http://www.diva-portal.org

Postprint

This is the accepted version of a paper presented at *19th International Conference on Transparent Optical Networks (ICTON)*.

Citation for the original published paper:

da Silva, C N., Yayimli, A., Wosinska, L., Furdek, M. (2017)
Assessing the effects of physical layer attacks on content accessibility and latency in optical CDNs.
In: Girona, Spain
https://doi.org/10.1109/ICTON.2017.8024993

N.B. When citing this work, cite the original published paper.

Permanent link to this version:
http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-214931

Assessing the Effects of Physical Layer Attacks on Content Accessibility and Latency in Optical CDNs

Carlos Natalino*, Marija Furdek*, Aysegul Yayimli, Lena Wosinska***

** Royal Institute of Technology (KTH), Stockholm, Sweden. E-mail: {carlosns, marifur, wosinska}@kth.se*

*** Istanbul Technical University (ITU), Istanbul, Turkey. E-mail: gencata@itu.edu.tr*

ABSTRACT

Content Delivery Networks (CDNs) are a major enabler of large-scale content distribution for Internet applications. Many of these applications require high bandwidth and low latency for a satisfactory user experience, e.g, cloud gaming, augmented reality, tactile Internet and vehicular communications [1]. Replication is one of the most prominent solutions to meet the requirements of latency-sensitive applications [1, 2]. However, infrastructure disruptions can greatly degrade the performance of such applications, or even cease their proper execution. The extent of degradation can be exacerbated by malicious attackers that target the critical elements of the CDN physical infrastructure to disconnect or severely degrade services.

In this work, we assess the effects of physical-layer attacks performed by cutting optical fiber links on content accessibility and latency in CDNs. We perform preliminary experiments on the Germany50 network with 50 nodes and 88 links, considering the scenario where attackers cut the links with highest importance, i.e., betweenness centrality (denoted by the number of shortest paths that traverse a link), in order to increase the effectiveness of the attack. By cutting links that are traversed by the majority of shortest paths, the attack forces a larger part of services to use longer paths, incurring higher latency. We consider the cases where each content has between 1 and 4 replicas placed at the network nodes with highest closeness centrality, i.e., nodes closest to all other nodes.

We evaluate the attack scenarios using the average content accessibility (ACA) [2] and the total replica latency. The ACA denotes the portion of nodes able to access any replica out of all nodes. To obtain the total replica latency, we summarize the propagation latency in the fiber and the switching latency at the network nodes along the paths connecting each node to the closest replica. The propagation latency is based on the average path length and light propagation speed in fiber equal to $2x10^8$ m/s. The switching latency is based on the number of hops, and for simplicity we assume that each node adds 20 µs delay, which is supported by commercially available switches.

Fig. 1a shows that the network maintains content accessibility for all nodes under attacks cutting up to 20% of the links. For more than 20% of the links cut, several nodes become unable to connect to a replica, even when 4 replicas are placed in the network. The latency results presented in Fig. 1b-1d focus on the case when up to 20% of links are cut and all nodes still have accessibility to content. In the single replica case, cutting 20% of the links doubles the average distance to replica (Fig. 1b), the average number of hops (Fig. 1c), and the total latency to replica (Fig. 1d). For 2 and 3 replicas, the latency degrades by 38% and 48%, respectively, while with 4 replicas the latency degradation amounts to 34%. Given that latency-sensitive applications can require latency in the order of 20 ms [1], propagation and switching latency account for 10% of the maximum latency under normal network working conditions. Under attacks, such latency can account for up to 20% of the maximum allowed value. Considering cloud processing and other factors that also contribute to the overall latency, current networks are operating at the limit of the latency requirements, and the increase caused by attacks can be the tipping factor causing the network to fail to meet the latency requirements.

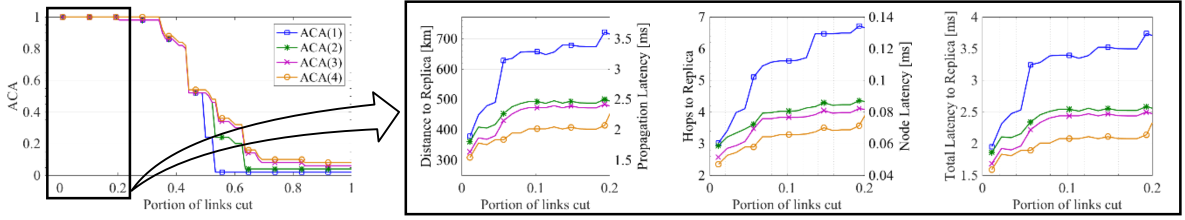

(a) (b) (c) (d)

Figure 1. Results for Germany50 network topology for (1) to (4) replicas: (a) the average content accessibility (ACA); (b) the shortest distance to replica and the corresponding propagation latency; (c) the number of hops to replica and the corresponding node latency; and (d) the total latency to replica.

Keywords: content delivery networks, content accessibility, targeted attacks, link cut attacks, latency.

ACKNOWLEDGEMENTS: This article is based upon work from COST Action CA15127 ("Resilient communication services protecting end-user applications from disaster-based failures – RECODIS") supported by COST (European Cooperation in Science and Technology).

REFERENCES

[1] S. Host, *et al.*: Network requirements for latency-critical services in a full cloud deployment, *in Proc. of SoftCOM 2016*, Split, Croatia, Sep. 2016, pp. 1-5.

[2] C. Natalino, *et al.*: Content accessibility in optical cloud networks under targeted link cuts, *in Proc. of ONDM 2017*, Budapest, Hungary, May 2017, pp. 1-6.