Stony Brook University
Honors College
The State University of New York

K-resolver: Towards Decentralizing Encrypted DNS Resolution

Ivan Lin, Nguyen Phong Hoang, Michalis Polychronakis
Stony Brook Computer Science

ABSTRACT

Recent protocols such as DNS-over-HTTPS (DoH) move towards privacy and security through encrypted domain resolution. However, these services will give providers full access to all user domain queries. We propose K-resolver, a DNS resolution mechanism that disperses DNS queries across multiple DoH resolvers, reducing the amount of information about a user's browsing activity exposed to each individual resolver. As a result, no resolver learns a user's entire web browsing history. We implement a prototype of our approach for Mozilla Firefox and used it to evaluate the performance of web page load time compared to the default centralized DoH approach.

INTRODUCTION

The domain name system (DNS) protocol was designed to resolve domain names to IP addresses. Its design did not consider security or privacy, permitting queries to be transmitted in cleartext. DoH introduced name resolution over secure encrypted channels, proxying resolution requests through third-party providers. However, resolver providers maintain full access to the queries.

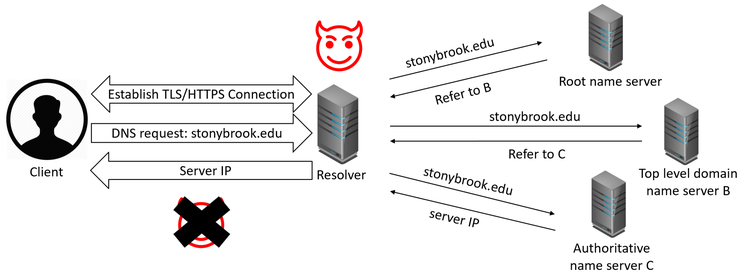

We consider several distribution strategies. Round robin and random distribution leak information to all resolvers over time. Distributing based on the queried domain could lead to browsing history being inferred (e.g. *yt.img* implies a visit to *youtube.com*). We distribute based on context domain.

IMPLEMENTATION

In order to resolve domain *d* given k available resolvers, use the modulo of *hash(d)* across the size of our resolver pool to select a resolver.
For a page's child domain resolutions, use the same resolver.

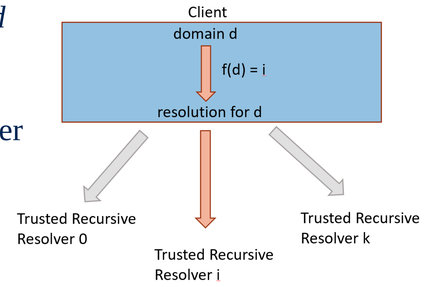

We propose K-resolver, which implements a bucketing method based on the primary page domain. Any requests originating from the context of *stonybrook.edu*, were sent to the resolver indexed by the hash of *stonybrook.edu*.

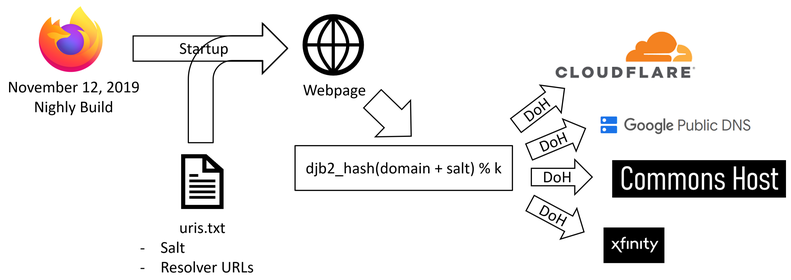

METHODS

Test profiles were developed as sets of 100 domains sourced from Tranco's top 1 million. 26 resolvers were manually selected. Using Browsertime, we visited each domain in the user profile and tracked measures for privacy and performance.

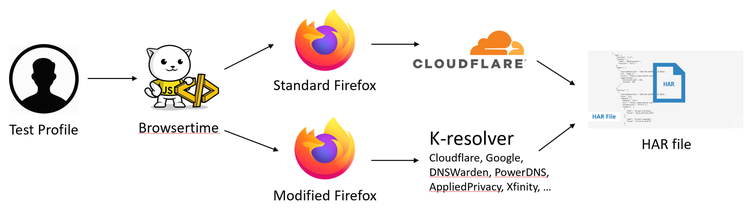

RESULTS

The hash function failed to uniformly distribute domains – some resolvers received as many as 10 requests per test while other received none.
Performance was also weaker compared to centralized DoH by a factor.

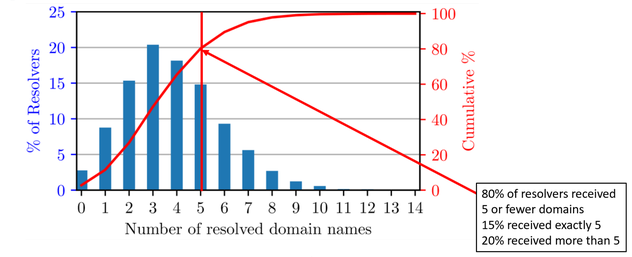

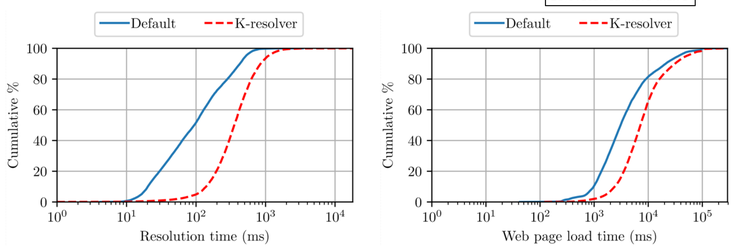

DISCUSSION & CONCLUSION

- Should domains with similar uses (e.g. health care, pornography) be sent to the same resolver?
- Is it possible to improve performance without sacrificing pribacy using alternative bucketing methods?
- Are client-endconvenience or performance features like DNS preloading and caching conducive to DNS-privacy?
- Performance will be variable – if a resolver is unavailable, should a secondary resolver be used?

KEY REFERENCES

Full paper: K-resolver: Towards Decentralizing Encrypted DNS Resolution, Nguyen Phong Hoang, Ivan Lin, Seyedhamed Ghavamnia, Michalis Polychronakis

ACKNOWLEDGEMENT

This was part of a work done with other authors. I would like to acknowledge my mentor Phong, the first author of the original full work, as well as the other authors and my advisor Michalis. For more information, see the full paper above: K-resolver: Towards Decentralizing Encrypted DNS Resolution