Measured Approaches to IPv6 Address Anonymization and Identity Association

ACM IMC PRIME Workshop – Madison, WI – 27 Oct 2025

David Plonka <plonka@wiscnet.net|research@wiscnet.net> & Arthur Berger

https://research.wiscnet.net/~plonka/pubs.html

“The agurify Tool: a kIP reference implementation” (IETF 105 Hackathon, 2019)
“kIP: a Measured Approach to IPv6 Address Anonymization” (2017)
https://arxiv.org/abs/1707.03900

WiscNet
Connecting People and Strategies

Premise: an Intersection of Privacy and Security

- **Privacy: Truncation and/or aggregation-based anonymization**
  i.e., for reporting traffic data, *e.g.,* correlating with network topology, routing, service providers, and geographic locations.

- **Security: in coordinated attack response**
  - Sharing IP address info while respecting victim or suspect attacker's privacy.
  - Mitigating abuse by blocking traffic associated with a victim and/or attacker.

Effectiveness depends on knowledge - or on assumptions - about globally-routed IP address prefixes – the *Identity Associations* (or IAs) – of the victims or attackers.

IP Address Anonymization and Identity Association

Consider these questions:

- *How can Internet measurements inform decisions about address anonymization and identity association?*

IP Address Anonymization and Identity Association

Consider these questions:

- How can passive and active Internet measurements inform decisions about address anonymization and identity association?

- *Is there any reason to believe that any one IP prefix length would perform satisfactorily for either? e.g., /24 in IPv4 or /48 in IPv6?*

IP Address Anonymization and Identity Association

Consider these questions:

- How can passive and active Internet measurements inform decisions about address anonymization and identity association?

- Is there reason to believe that any one IP prefix length would perform satisfactorily for either? e.g., /24 in IPv4 or /48 in IPv6?

IP Address Anonymization

- ***Truncation-based anonymization is ideal if, and only if, it can be guaranteed to improve privacy.***

We propose *k*IP anonymization, *i.e.*,
make an individual appear indistinguishable amongst a set of [*k*] individuals,
an *anonymity set*

[https://en.wikipedia.org/wiki/K-anonymity,
RFC 6973: “Privacy Considerations for Internet Protocols”]

IP Address Anonymization in Practice, Today

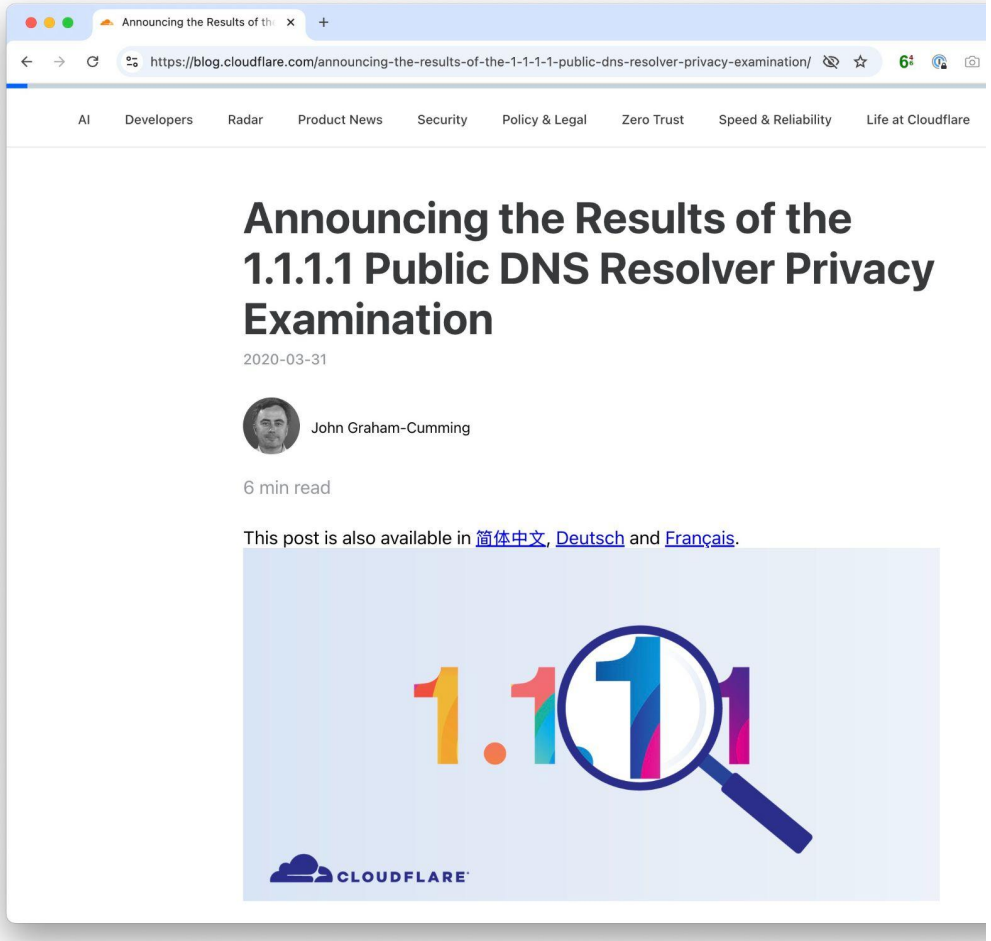

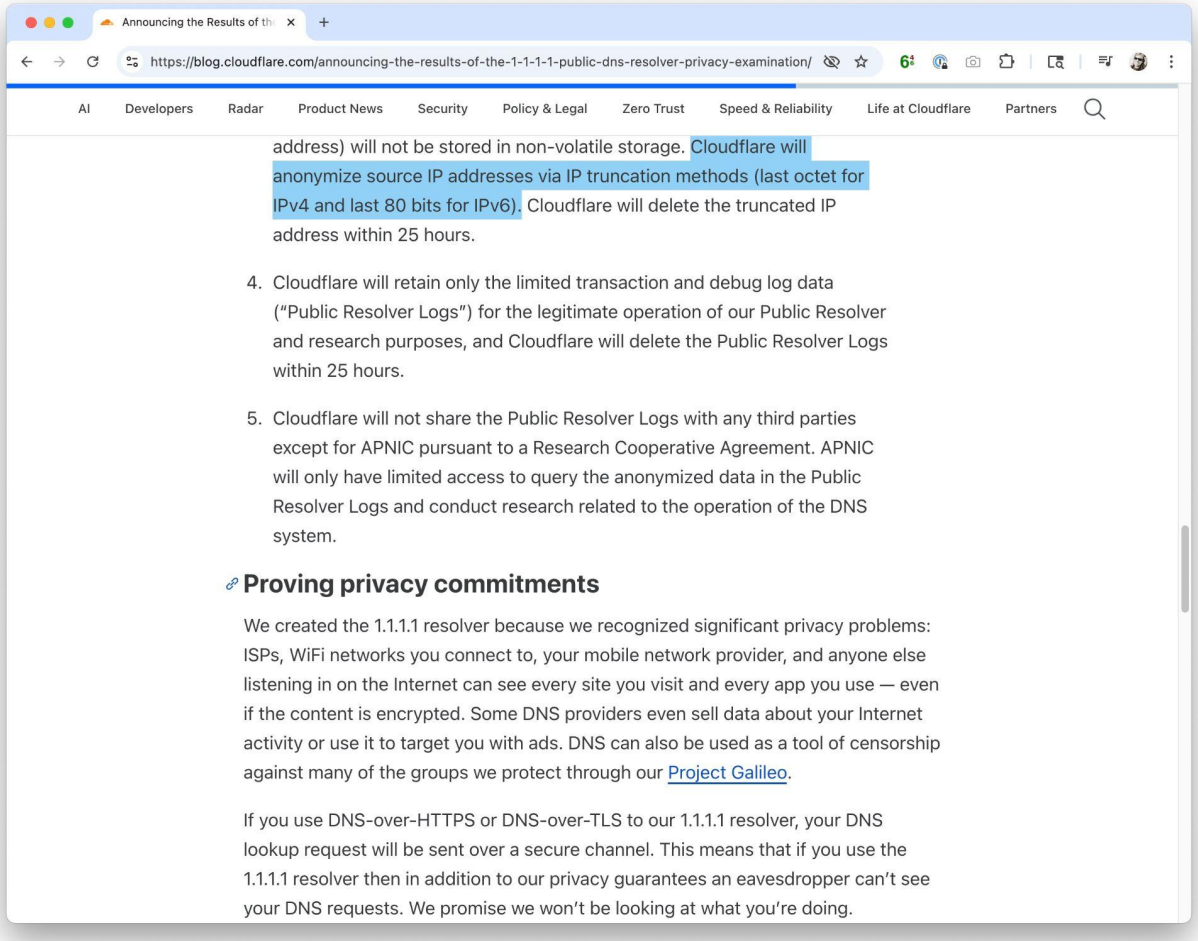

IP Address Anonymization in Practice, Today

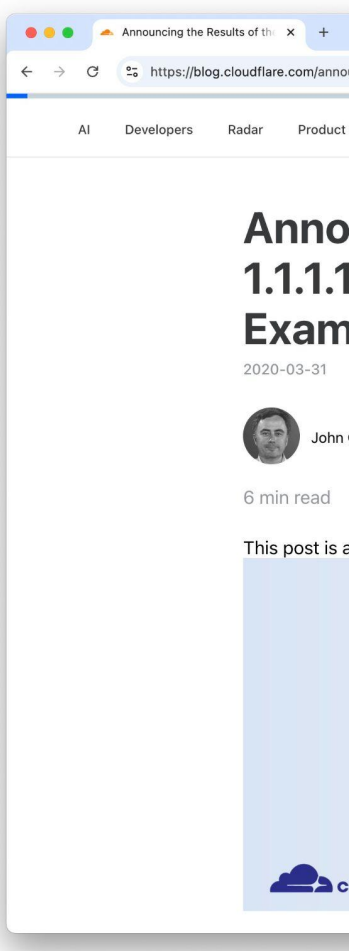

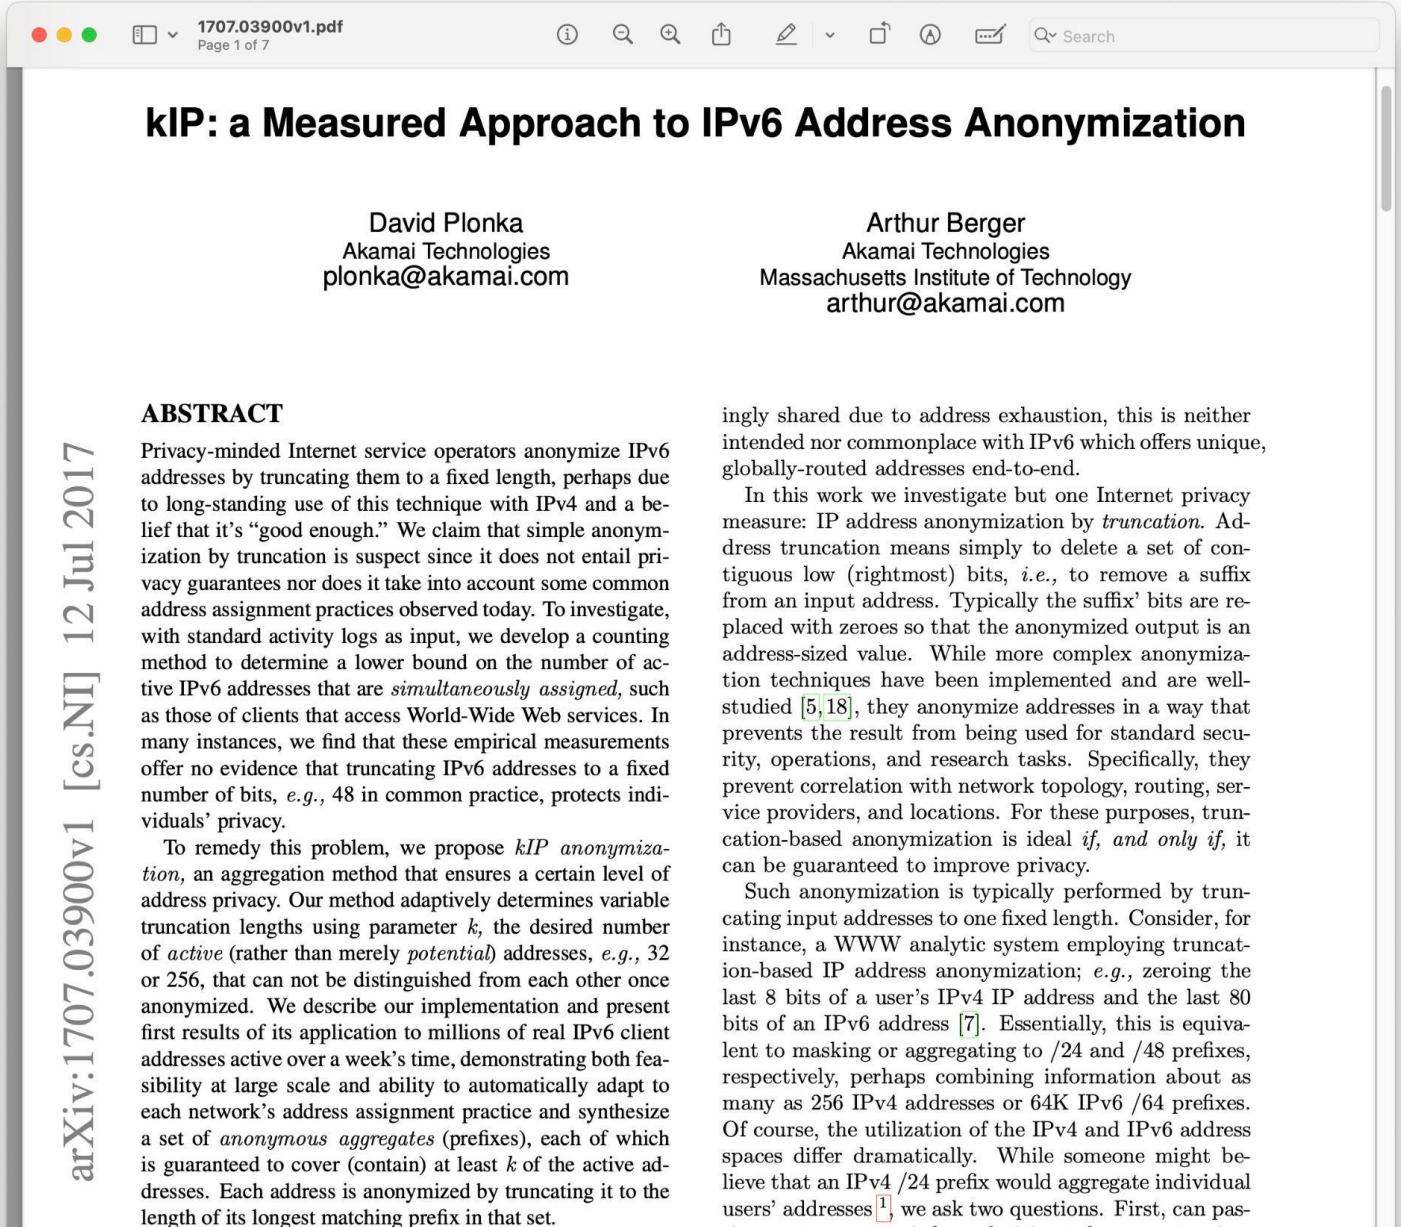

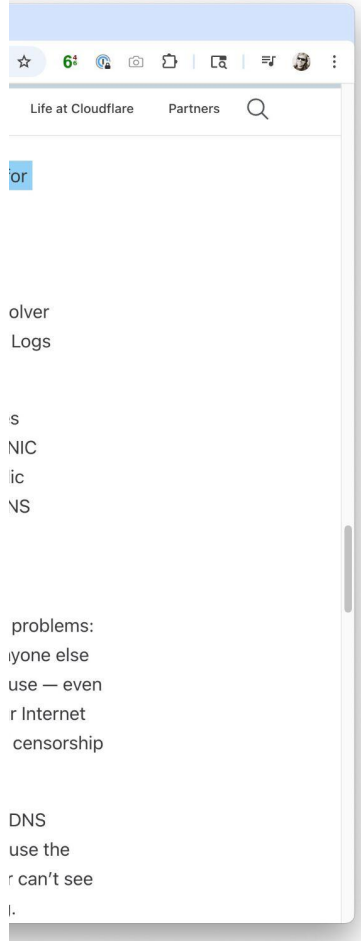

Histogram: k=32 anonymous aggregate prefix lengths (w=7d, i=1h)

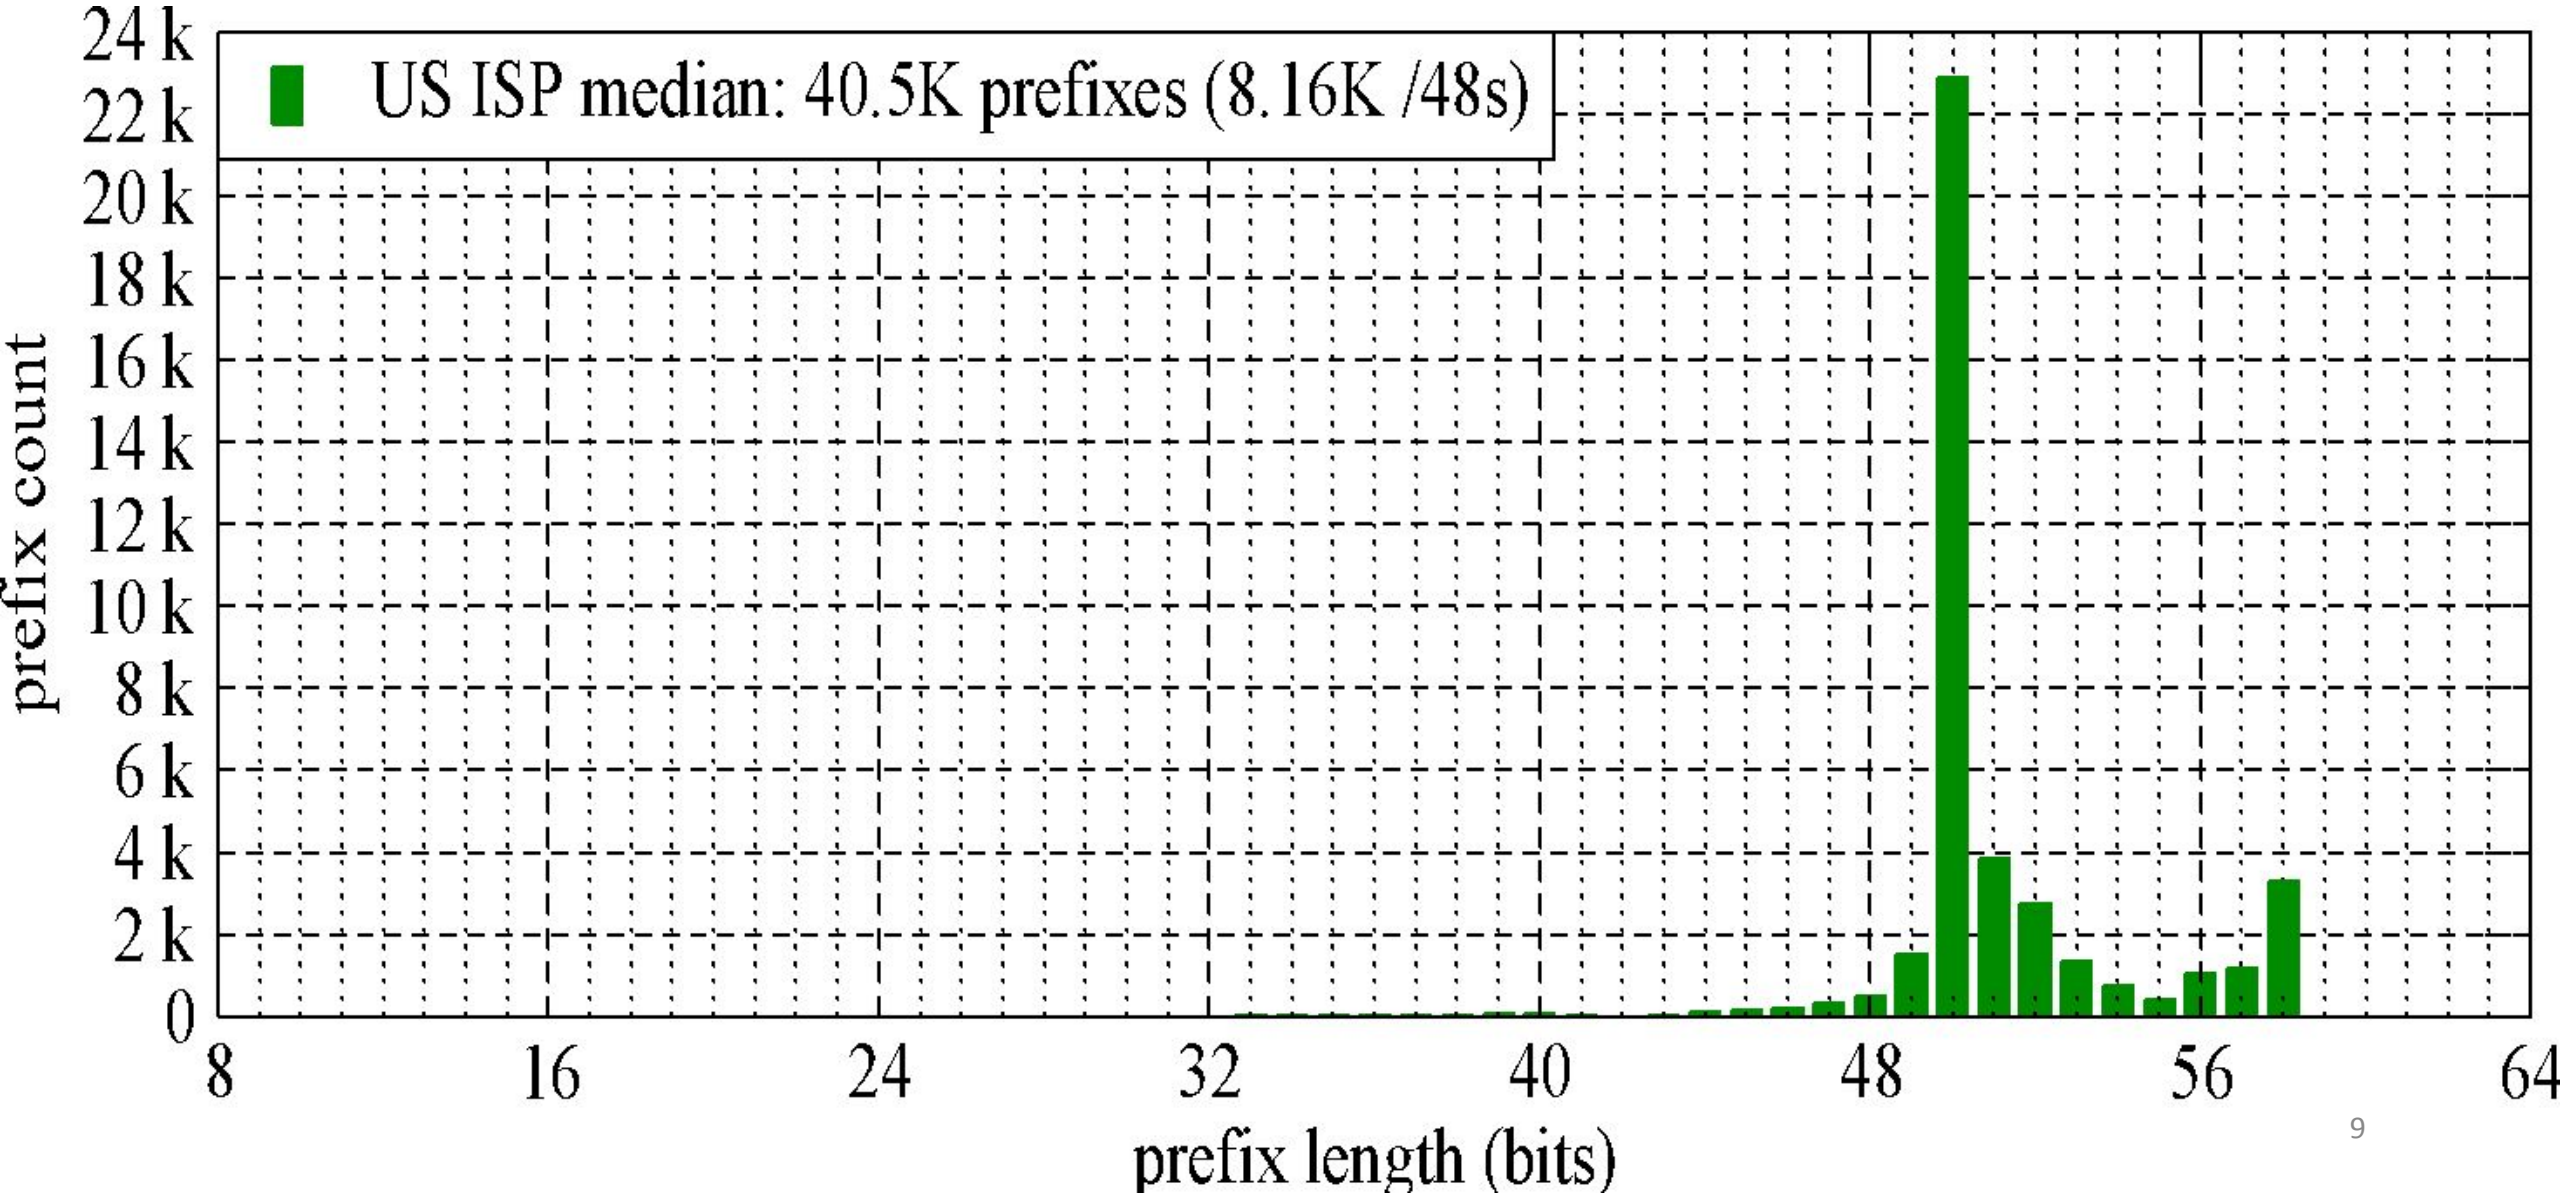

Histogram: k=32 anonymous aggregate prefix lengths (w=7d, i=1h)

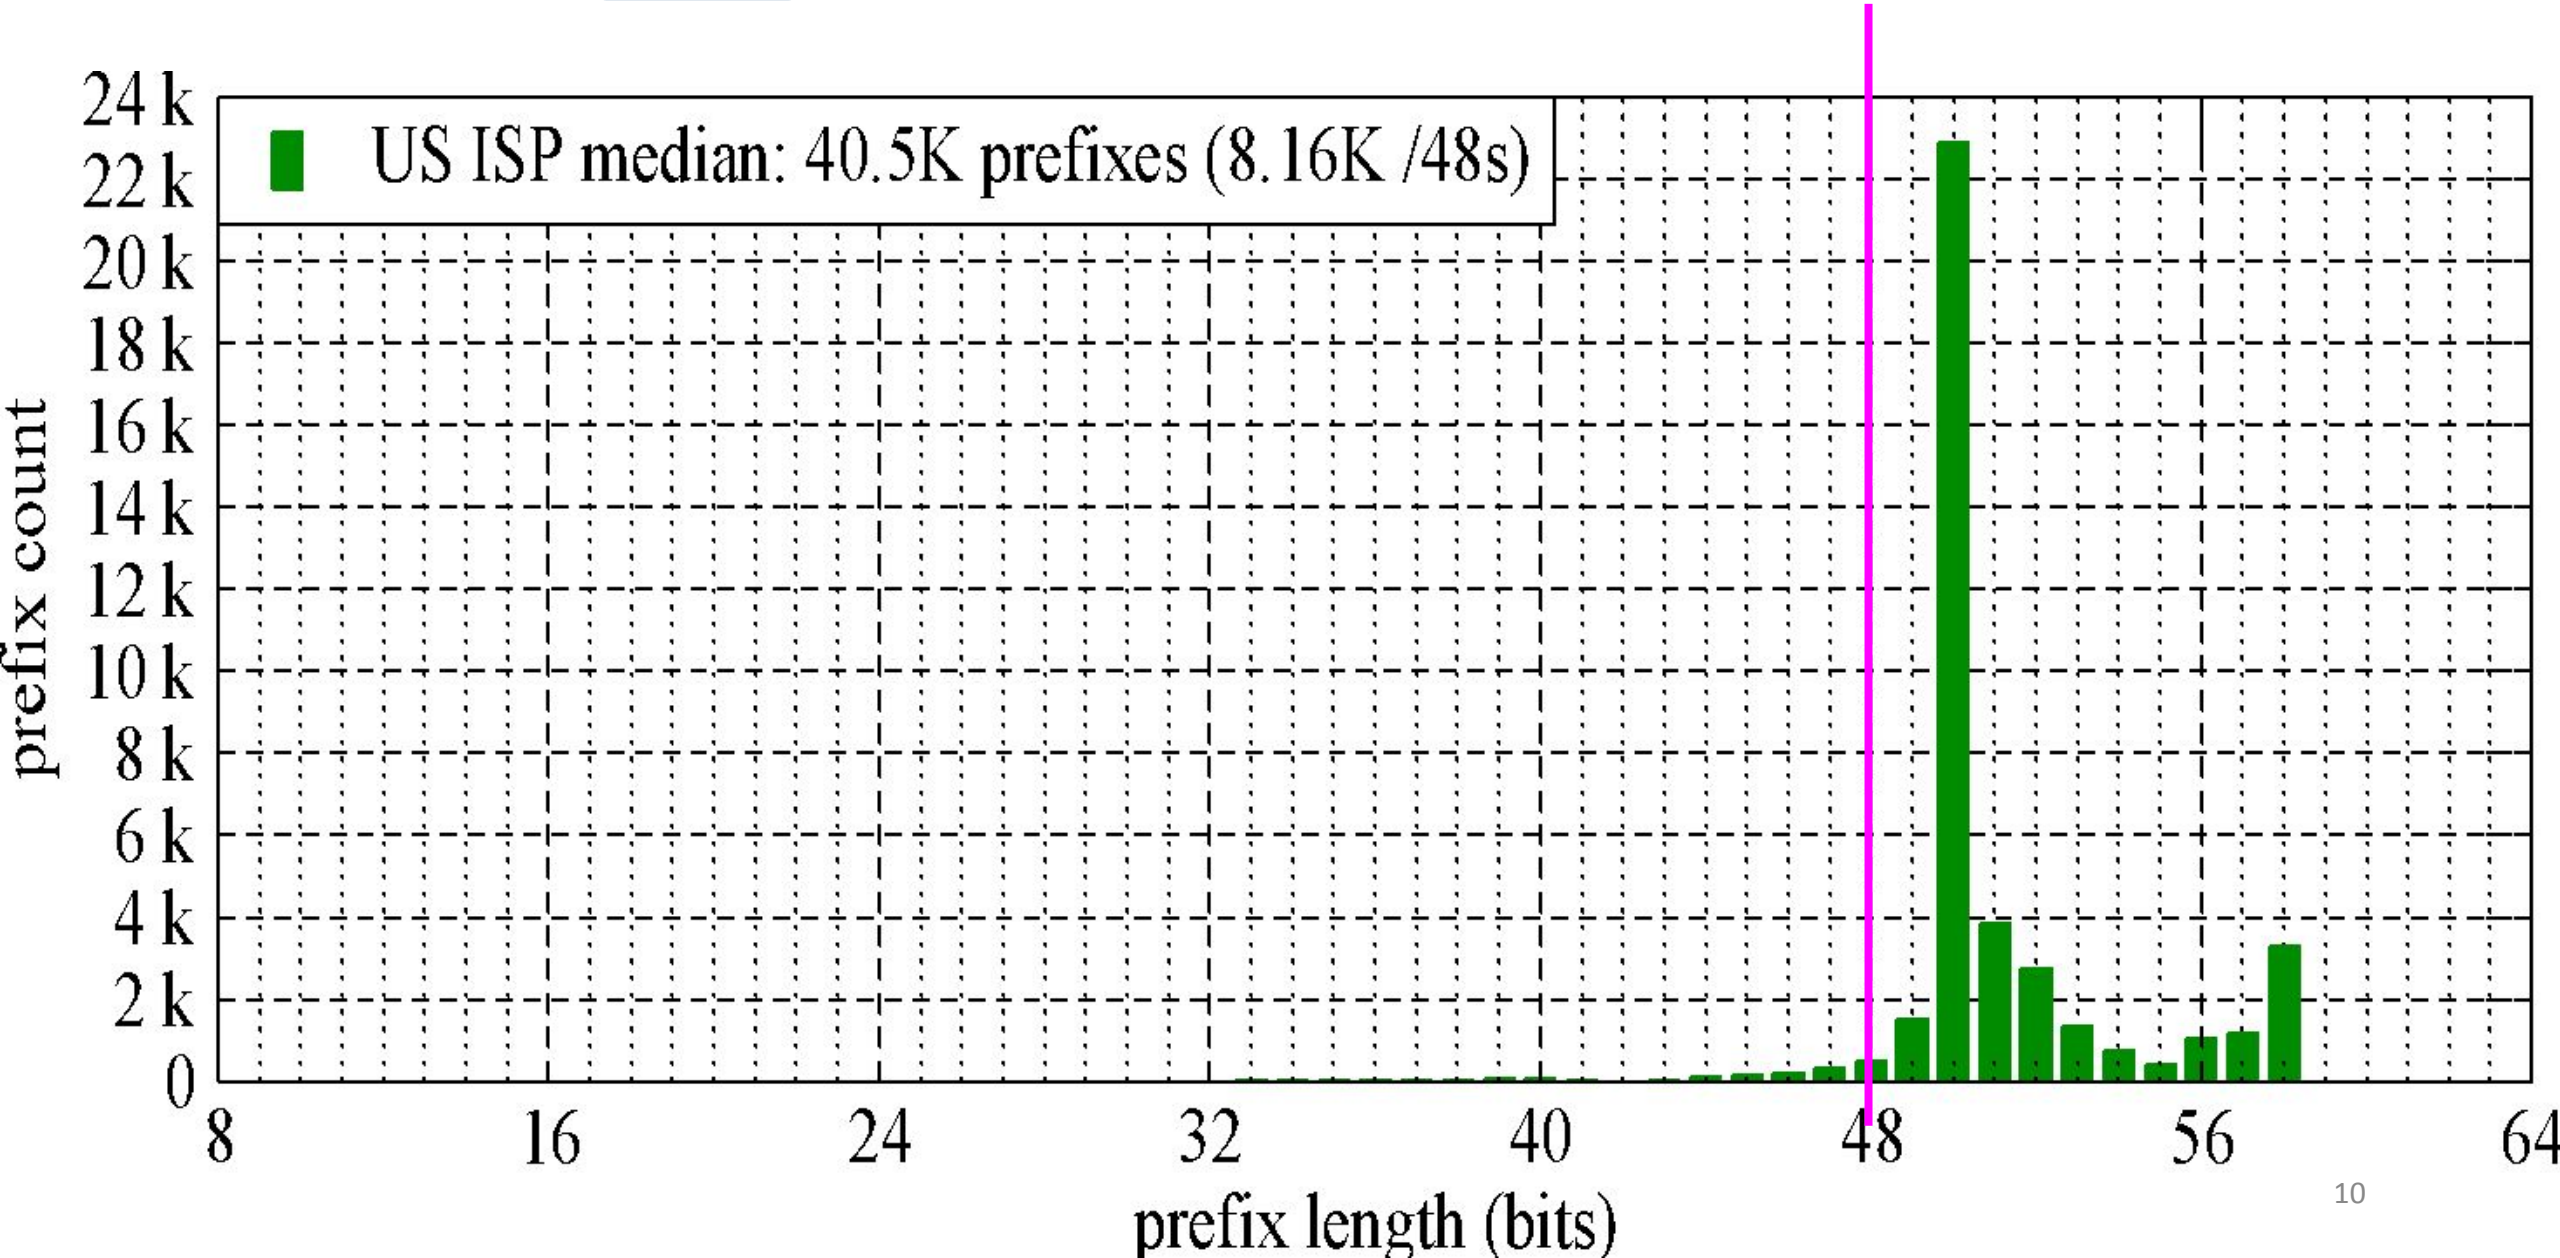

Histogram: k=32 anonymous aggregate prefix lengths (w=7d, i=1h)

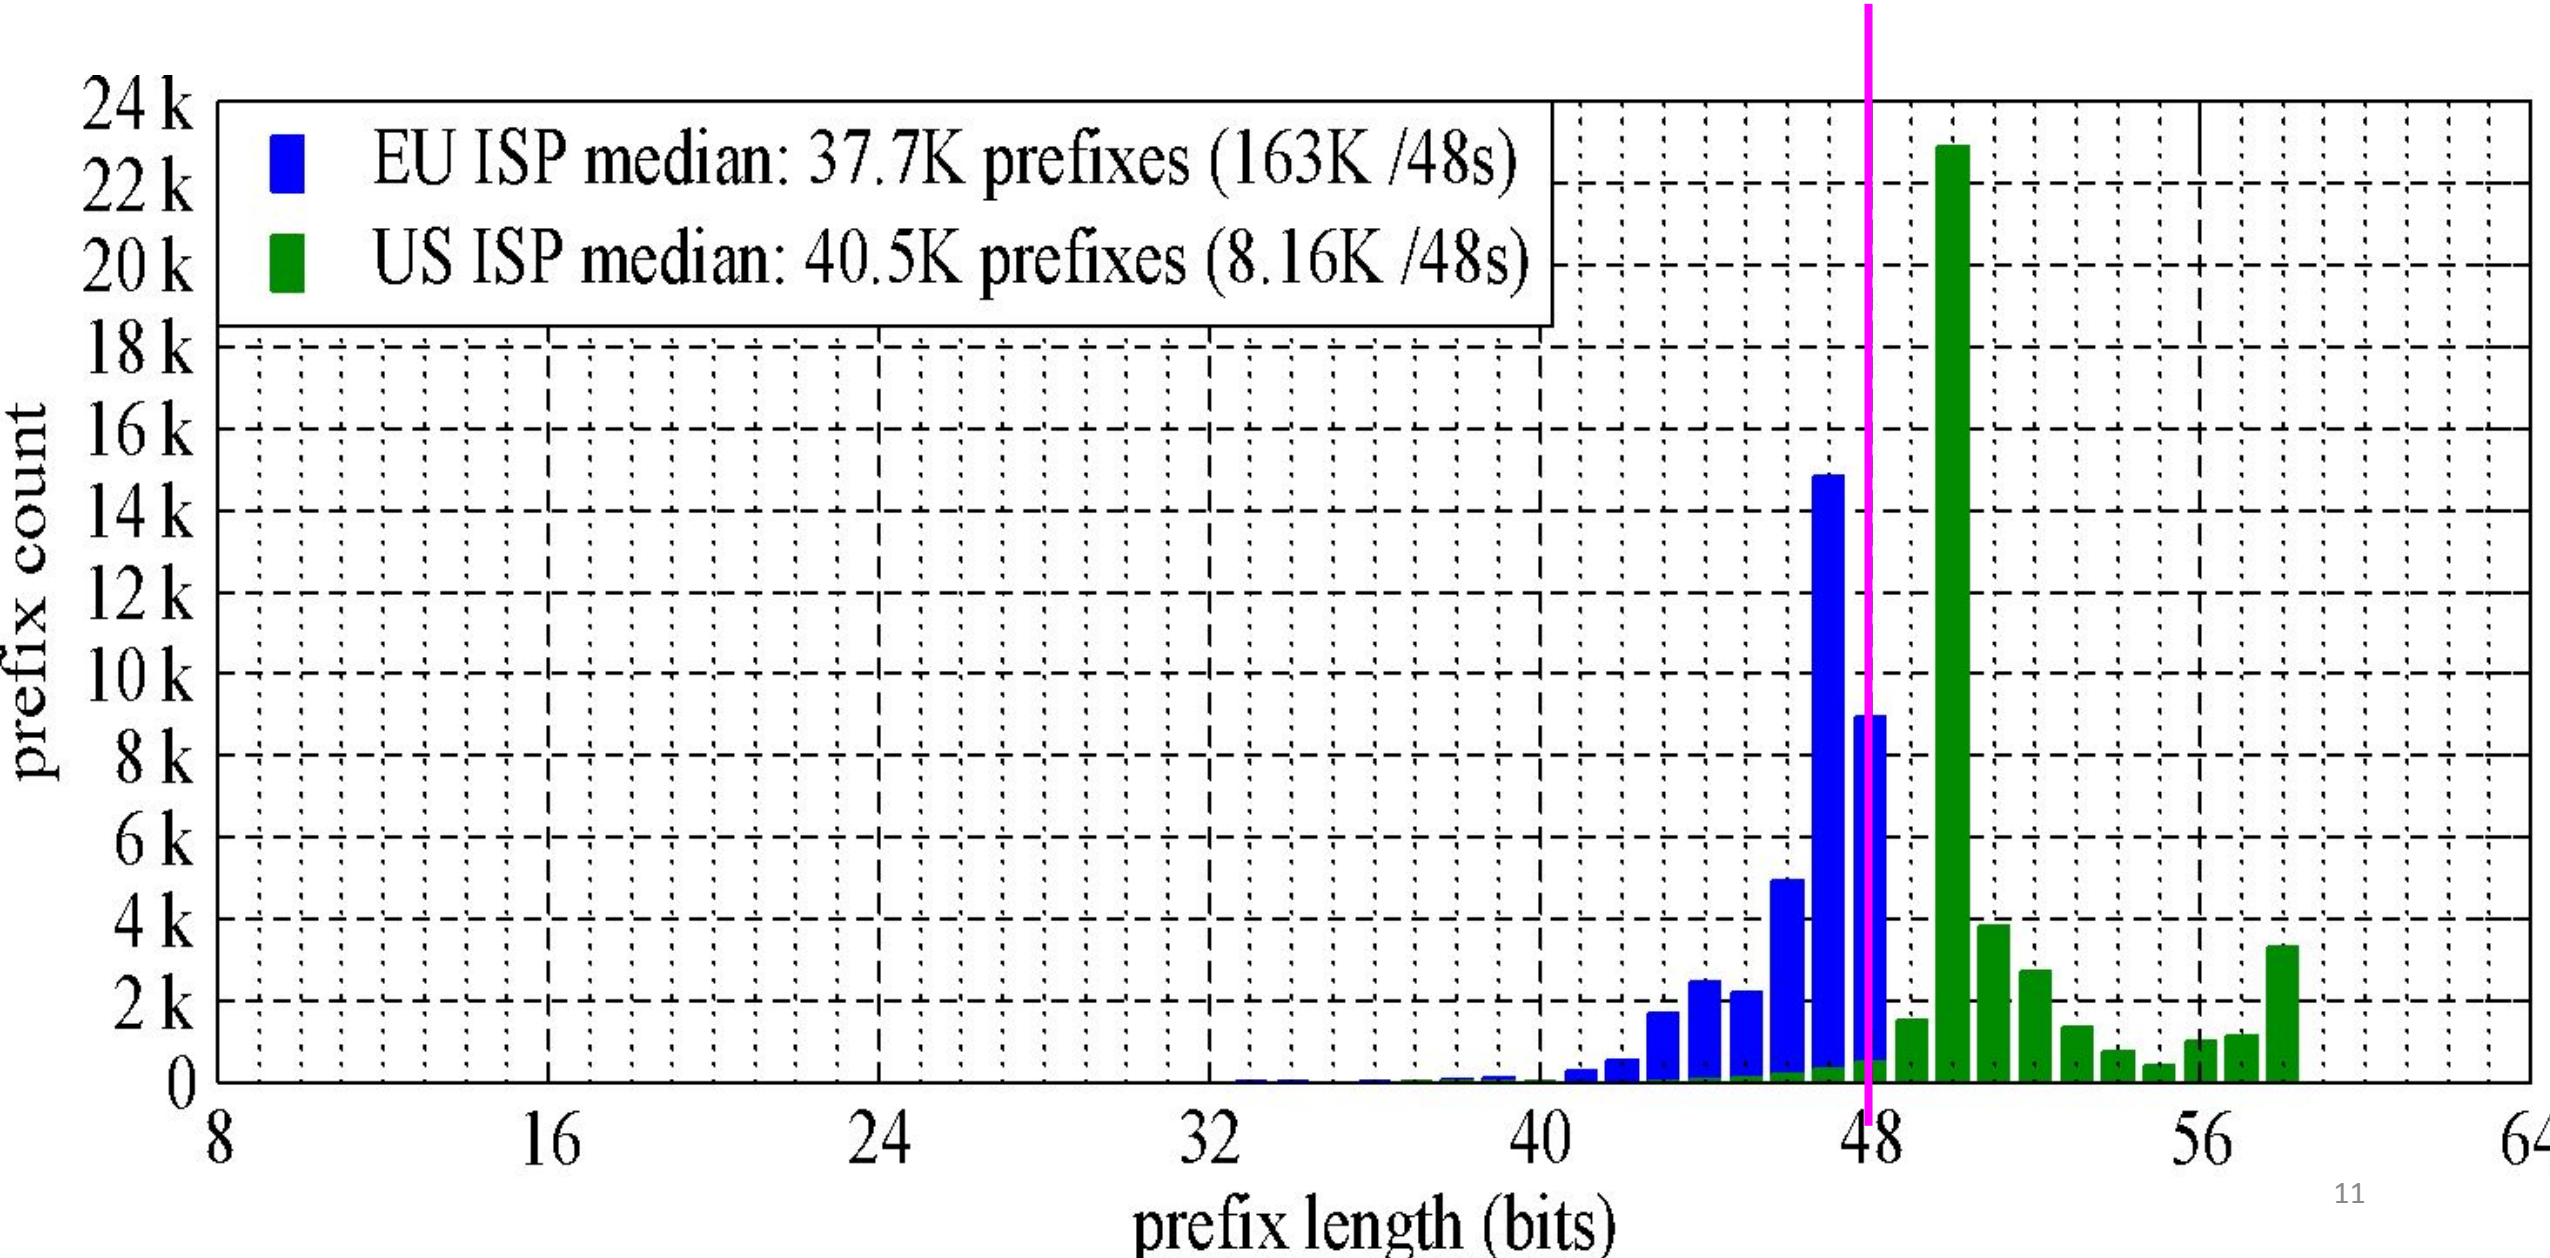

Histogram: k=32 anonymous aggregate prefix lengths (w=7d, i=1h)

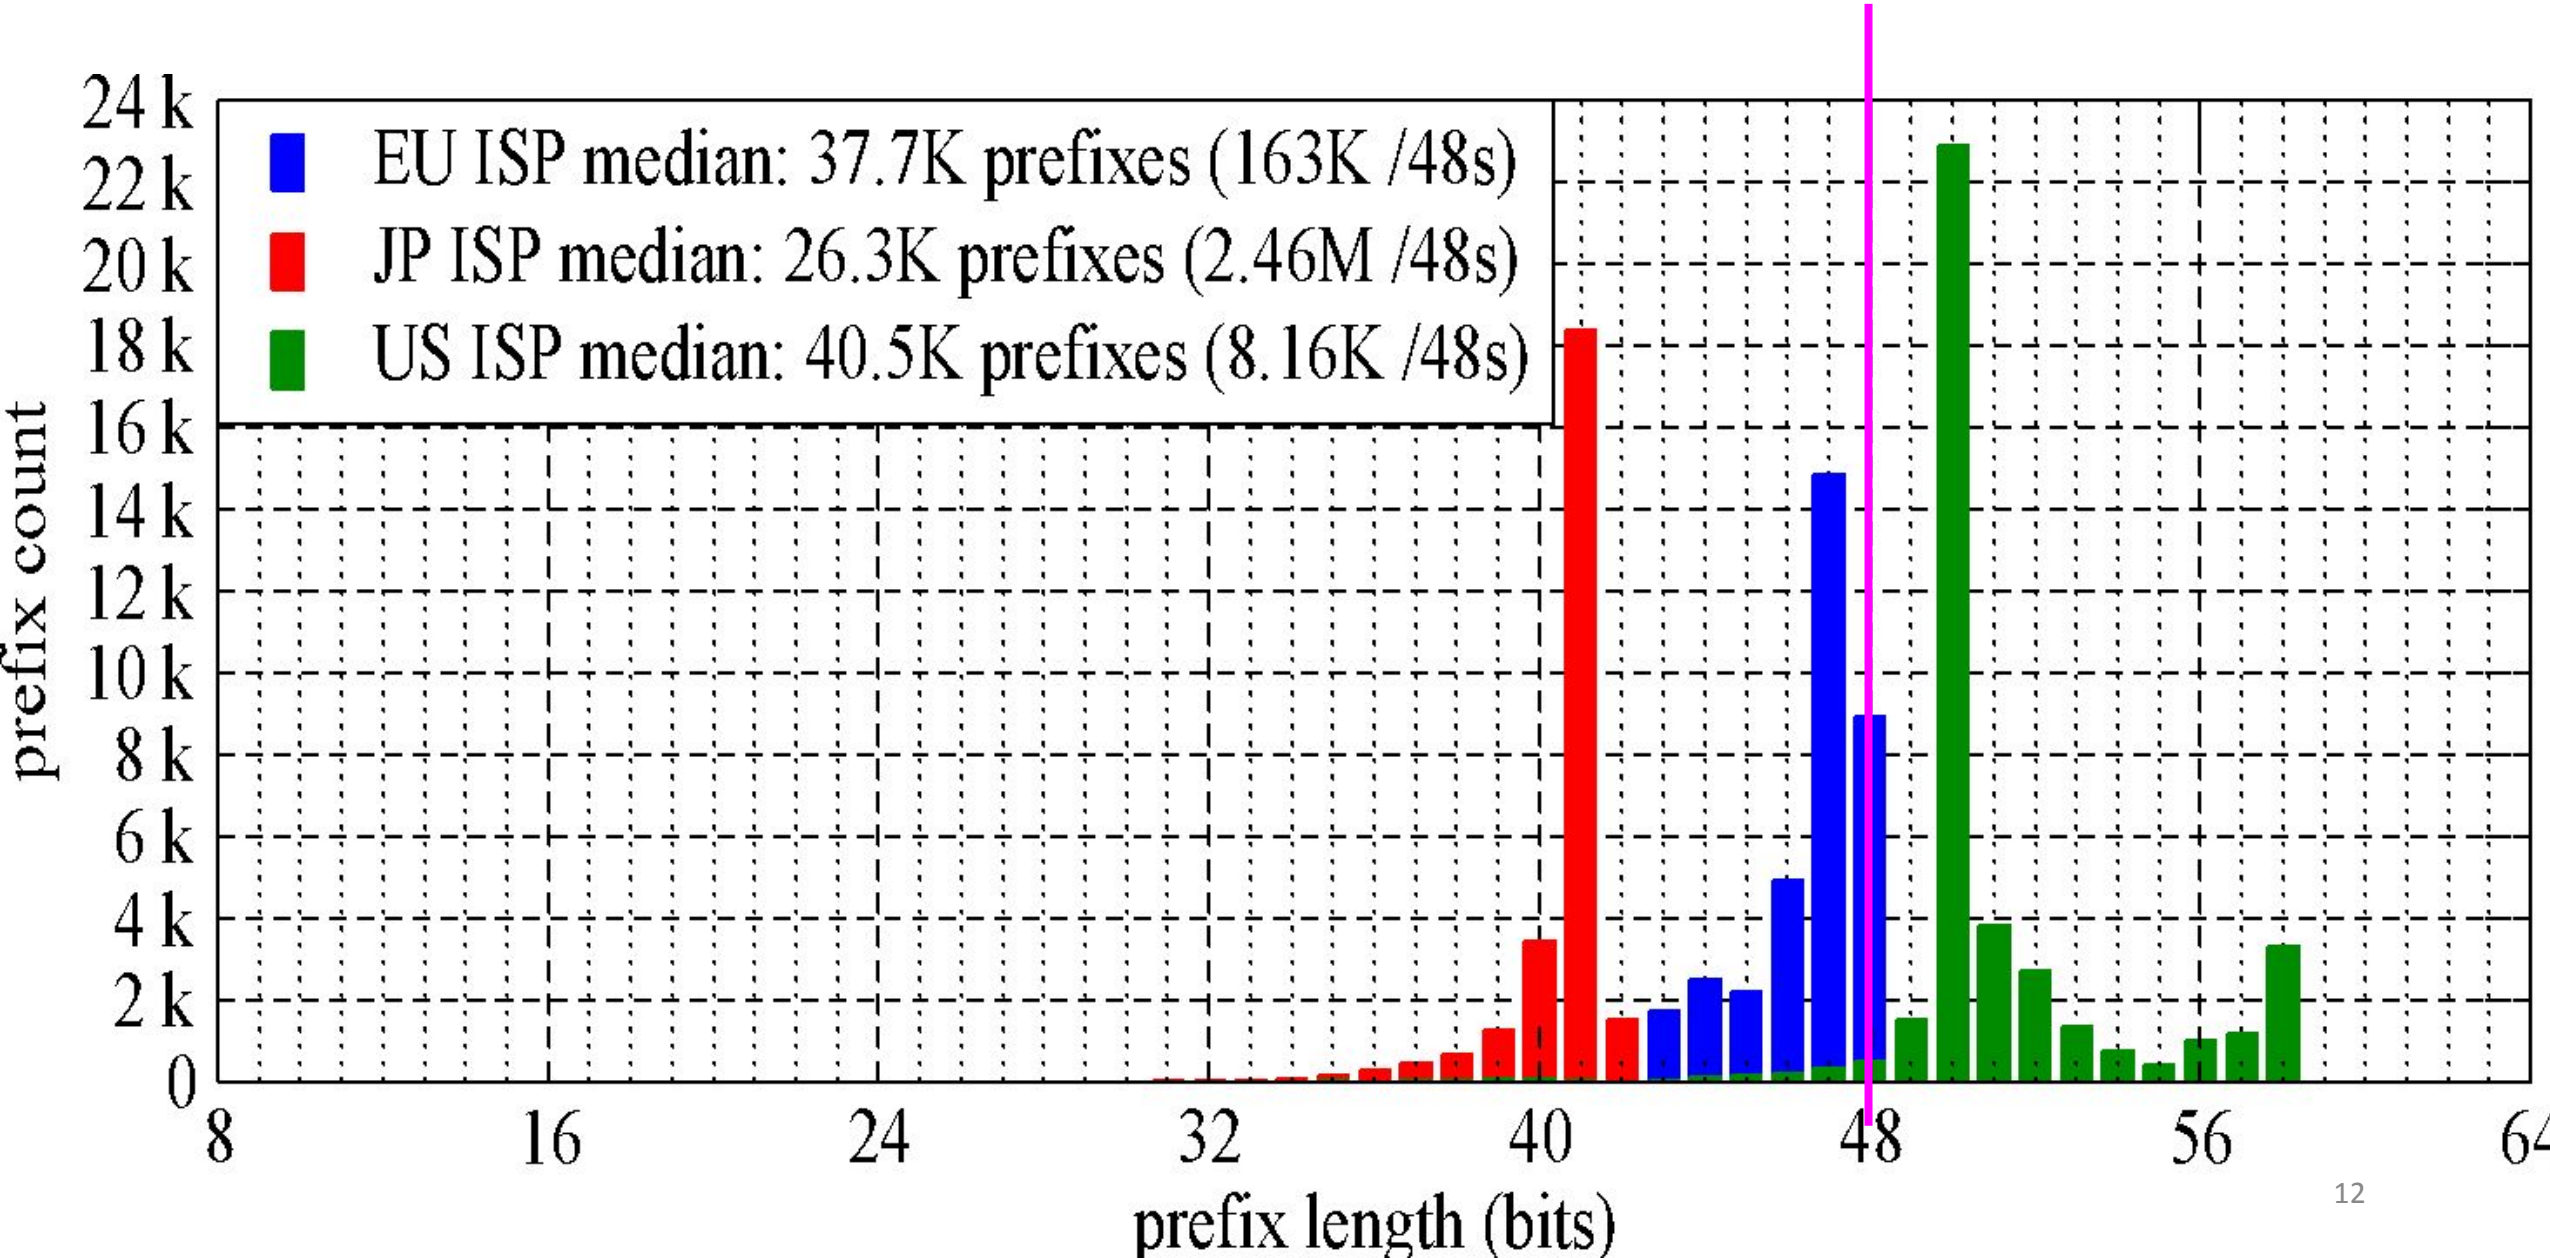

Some Prompts for Discussion

- Should there be a way to remotely **query for an IP address' preferred prefix length** to *(a)* **anonymization by truncation** or for its *(b)* **identity association**?

  If this were implemented, e.g., in the DNS similarly to `ip6.arpa`, the default value on lookup failure could be the length of the covering BGP prefix that must exist for globally-routable addresses.

  Is this a reasonable mechanism to also advertise different prefix lengths for different uses, e.g., anonymization within an organization (TLP:AMBER) versus outside (TLP:CLEAR)?

Some Prompts for Discussion

- Should there be a way to remotely **query for an IP address' preferred prefix length** to *(a)* **anonymization by truncation** or for its *(b)* **identity association**?

  If this were implemented, e.g., in the DNS similarly to `ip6.arpa`, the default value on lookup failure could be the length of the covering BGP prefix that must exist for globally-routable addresses.

  Is this a reasonable mechanism to also advertise different prefix lengths for different uses, e.g., anonymization within an organization (TLP:AMBER) versus outside (TLP:CLEAR)?

- Should administrative policy require **auditing** of the preferred prefix length to measure its effectiveness in **making an individual appear indistinguishable amongst a set of individuals**?

Measured Approaches to IPv6 Address Anonymization and Identity Association

ACM IMC PRIME Workshop – Madison, WI – 27 Oct 2025

David Plonka <plonka@wiscnet.net|research@wiscnet.net> & Arthur Berger

https://research.wiscnet.net/~plonka/pubs.html

“The agurify Tool: a kIP reference implementation” (IETF 105 Hackathon, 2019)
“kIP: a Measured Approach to IPv6 Address Anonymization” (2017)
https://arxiv.org/abs/1707.03900

WiscNet
Connecting People and Strategies

Histogram: k=256 anonymous aggregate prefix lengths (w=7d, i=1h)