End-To-End Privacy for Identity & Location with IP

*Saleem N. Bhatti, **Gregor Haywood**, Ryo Yanagida*

Service Routing and Addressing
IETF112 RTG/INT side meeting (online)
10 Nov 2021

This talk is based on a paper presented at:
NIPAA21 – 2nd Workshop on New Internetworking Protocols, Architecture and Algorithms,
29th IEEE International Conference on Network Protocols,
01 Nov 2021
Paper preprint: https://saleem.host.cs.st-andrews.ac.uk/publications/2021/nipaa21/nipaa21-bhy2021.pdf

University of
St Andrews

The Ties that un-Bind: Decoupling IP from web services and sockets for robust addressing agility at CDN-scale

Marwan Fayed[†], Lorenz Bauer[†], Vasileios Giotsas[†], Sami Kerola[†],
Marek Majkowski[†], Pavel Odinstov[†], Jakub Sitnicki[†], Taejoong Chung[⋆],
Dave Levin[‡], Alan Mislove[ø], Christopher A. Wood[†], Nick Sullivan[†]
[†] Cloudflare, Inc. [⋆] Virginia Tech [‡] University of Maryland [ø] Northeastern University

ABSTRACT

The couplings between IP addresses, names of content or services, and socket interfaces, are too tight. This impedes system manageability, growth, and overall provisioning. In turn, large-scale content providers are forced to use staggering numbers of addresses, ultimately leading to address exhaustion (IPv4) and inefficiency (IPv6).

In this paper, we revisit IP bindings, entirely. We attempt to evolve addressing conventions by decoupling IP in DNS and from network sockets. Alongside technologies such as SNI and ECMP, a new architecture emerges that "unbinds" IP from services and servers, thereby returning IP's role to merely that of reachability. The architecture is under evaluation at a major CDN in multiple datacenters. We show that addresses can be generated randomly *per-query*, for

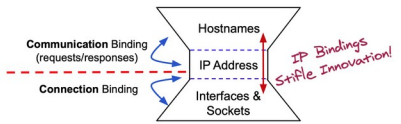

Figure 1: Conventional IP bindings to names, interfaces, and sockets, create transitive relationships between them that are difficult to track and reason about, which hinders changes to any binding without risking others.

SIGCOMM, Sep 2021 - https://dl.acm.org/doi/10.1145/3452296.3472922

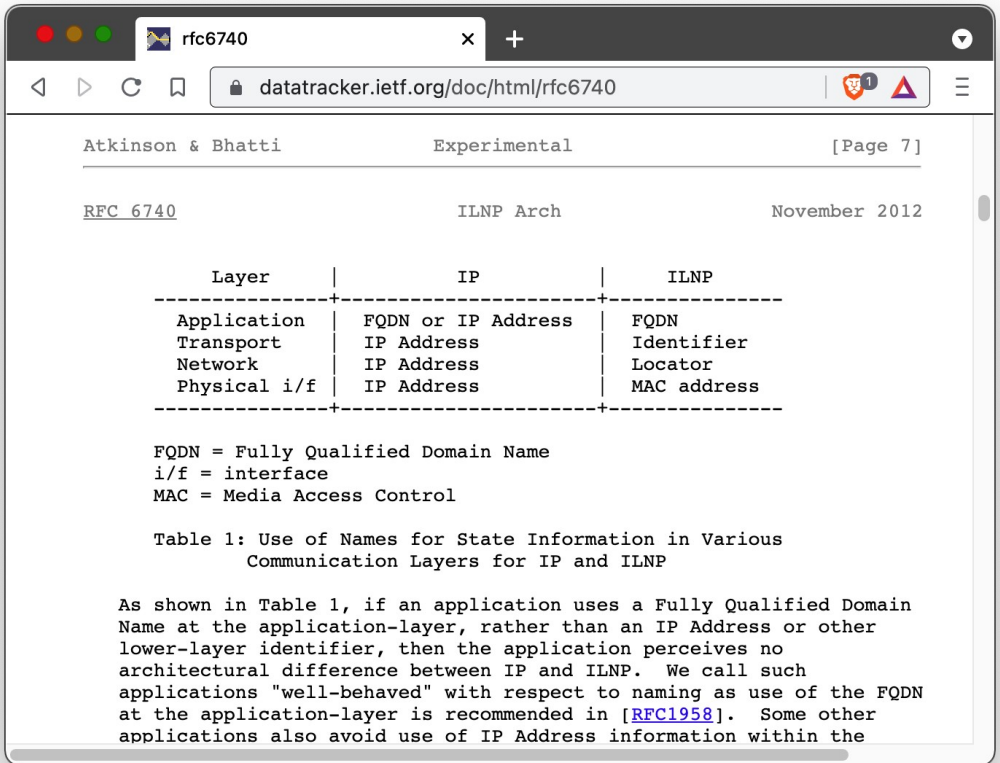

RFC6740(E), IRTF RRG, Nov 2012 - https://datatracker.ietf.org/doc/html/rfc6740

Identity and Location Privacy

University of St Andrews

- Modular network stack makes:
  - Design and implementation easy
  - Privacy hard
- Objectives:
  - Stop on-path attacks exploiting wire image
  - Avoid expanding trust boundary

Internet Location

- Upper 64 bits
- Used **globally** and managed **globally**
- Uniquely labels a **subnet**
- Determined by the ISP

Node Identity

- Lower 64 bits (IID)
- Used **globally** but generated **locally**
- Uniquely labels an **endpoint**
- Determined by node (e.g. SLAAC)

University of St Andrews

IPv6 address format (RFC4291 + RFC3587)

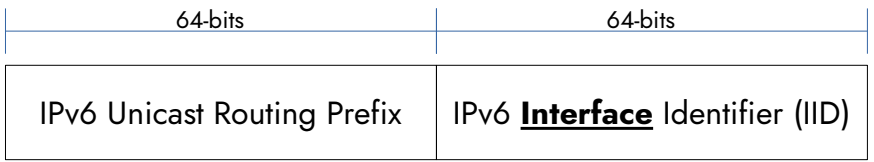

ILNP Identifier-Locator Vector (I-LV) (RFC6741)

Ephemeral Node Identifiers (NIDs)

University of St Andrews

- NIDs: transport-layer node identifiers
- Simultaneously use multiple
- Can be one-use

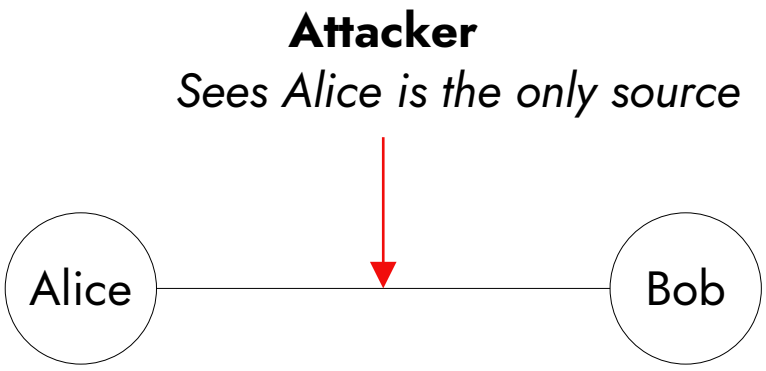

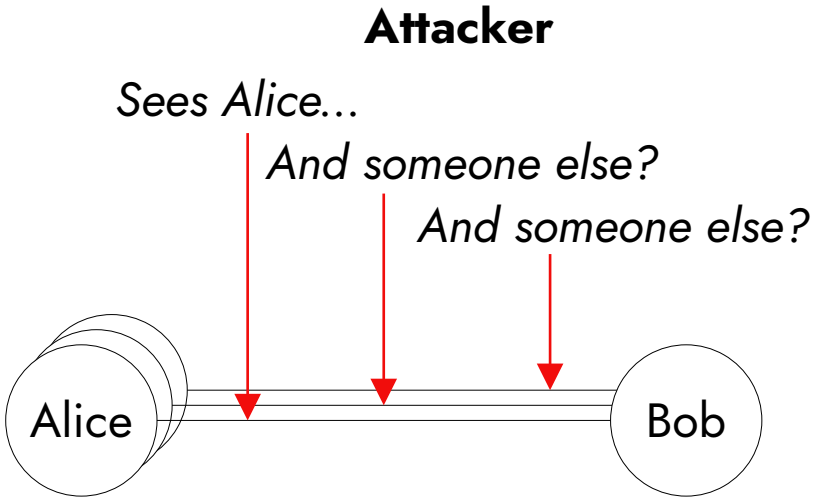

Location Privacy

University of St Andrews

- Routing information must be visible on path
- Solution: use multiple paths

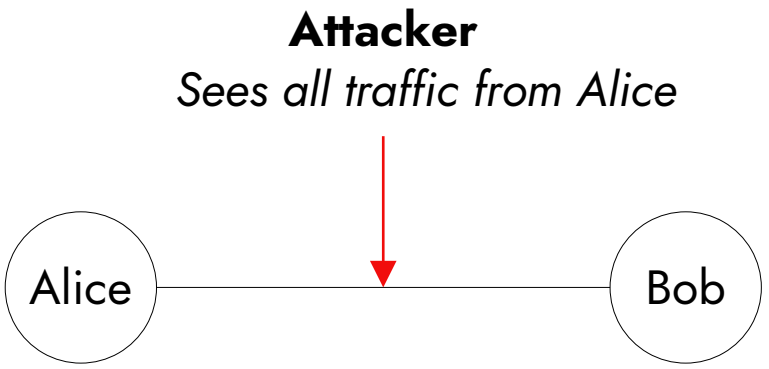

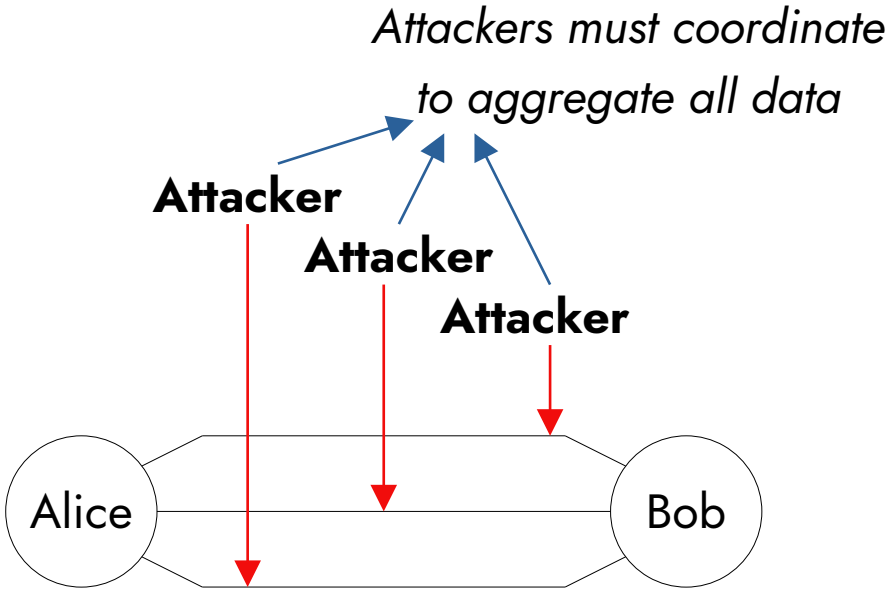

Location Privacy

University of St Andrews

- Location is still exposed unless using VPN/Tor
- Locator Rewriting Relays (LRRs) achieve this without tunneling
- Potentially easier for attacker to correlate
  - …but that may be inevitable either way

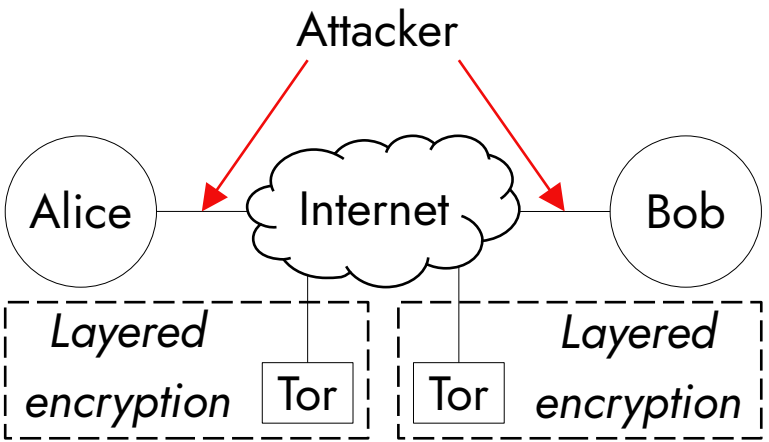

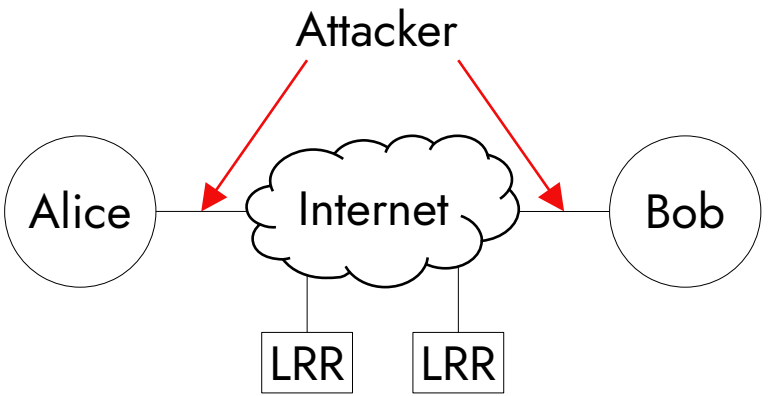

Emulation Overlay

University of St Andrews

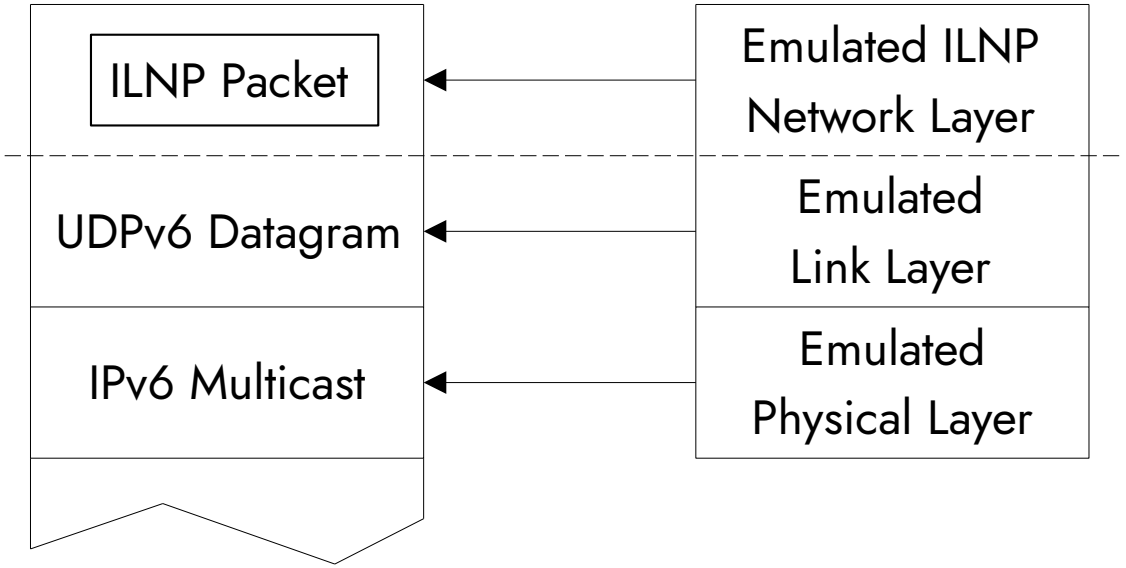

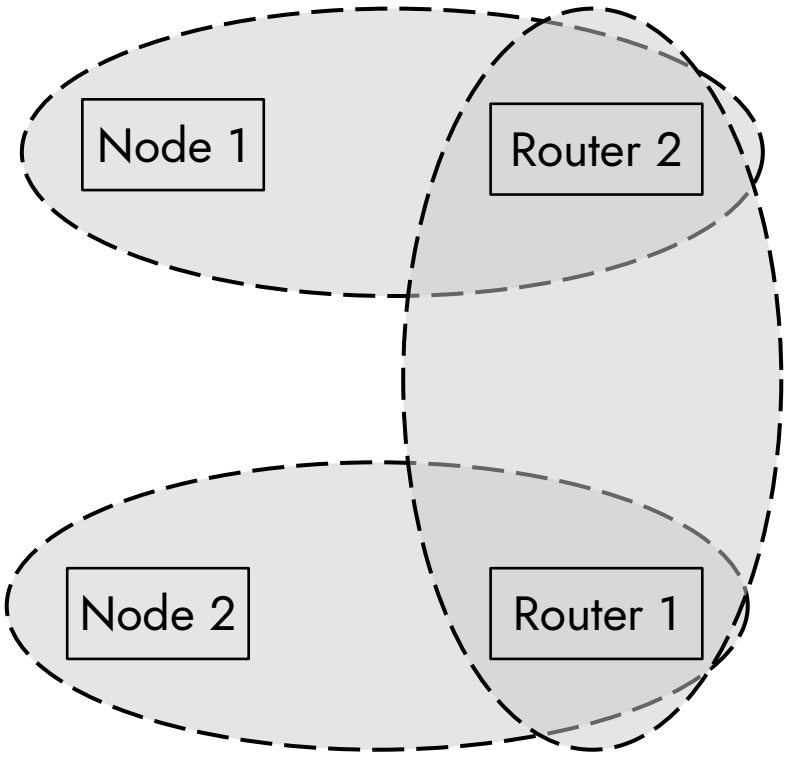

Results

University of St Andrews

No Defences

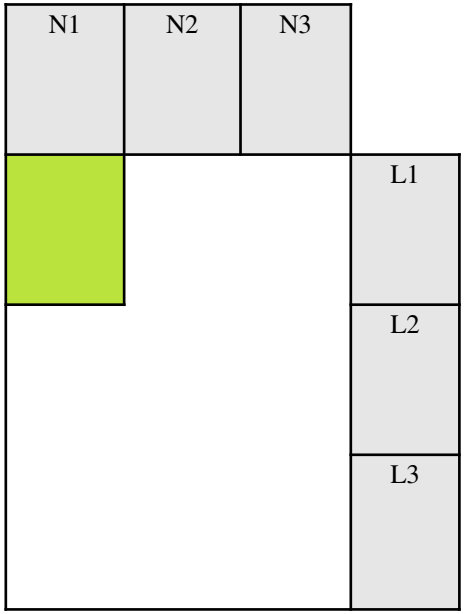

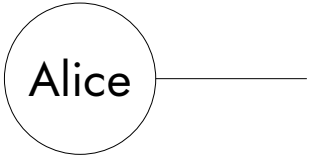

Ephemeral NIDs

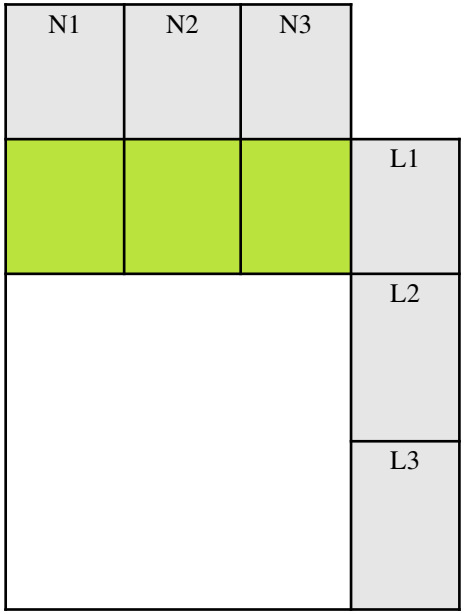

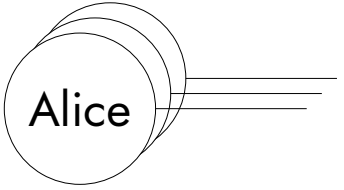

Ephemeral NIDs and Multihoming

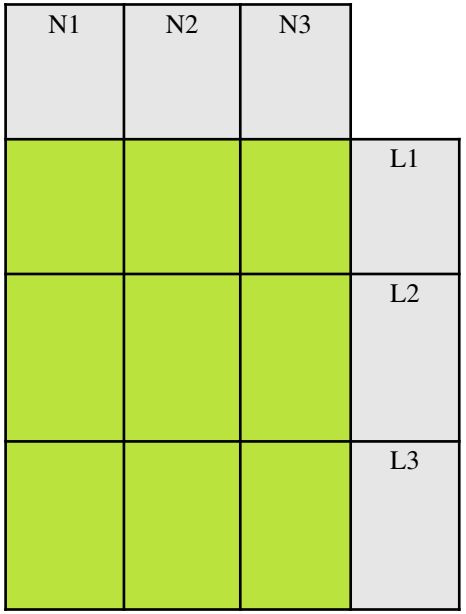

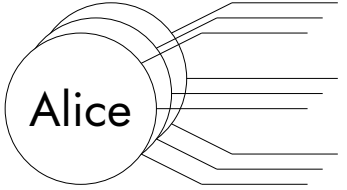

Concluding

University of St Andrews

- ILNP's architecture is useful for privacy
  - Isolate each flow with ephemeral NIDs
  - Multihoming makes attacker's job harder
  - LRRs provide low-cost location privacy
- Thank you!