Traffic Analysis Resistant Networks (TARN)

Richard R. Brooks, K.-C. Wang, Lu Yu,
Geddings Barrineau, Jon Oakley and Qing Wang

The Holcombe Department of
Electrical and Computer Engineering
Clemson University,
Clemson, SC, 29634, USA

Email:rrb@g.clemson.edu, kwang@clemson.edu,
lyu@g.clemson.edu, cbarrin@g.clemson.edu,
joakley@clemson,edu, qw@g.clemson.edu

March, 2017

Network vulnerabilities

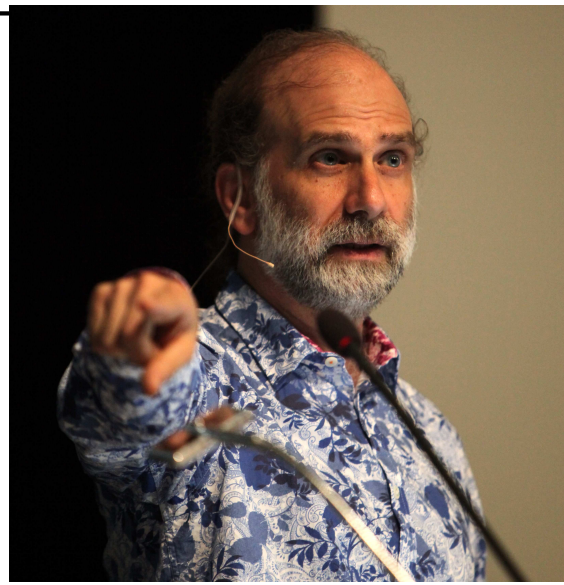

- Original Internet architecture did not consider security
  - Packets in clear text.
  - Meta-data in clear text.
- SSL/TLS and ssh suite encrypt payloads.
- Meta-data still in clear text.
- IP address meta-data leaks communications partners.
- Port numbers leak protocols.
- Packet sizes identify web content, even when encrypted.
- Timing side-channels leak information (including passwords), even when encrypted.

Why does this matter?

- Censorship
  - DNS filtering
  - IP filtering
- Surveillance
  - Watering hole attacks
  - Identify fans of Breitbart, Democracy Now!, Fox News, MSNBC, WSJ, NYTimes...
- Traffic hijacking
- Man-in-the-middle attacks
- DDoS

Existing solutions

- VPNs – Closed systems. Typically commercial. Easily detected and blocked. For example, Iran has simply blocked all encrypted connections.
- Single hop proxies – *Psiphon, Lantern, Clemson Internet Freedom in Africa...* Open systems. Subject to blocking of TLS/SSL. Some obfuscate their traffic. Must trust intermediates.
- Onion routing – *Tor* Increased latency and jitter. Access blocked by China and Iran. Bridge nodes get around blocking. Probing finds bridge nodes. Pluggable Transports.
- Decoy routing – *BBN, U Mich, UIUC* Router recognizes encrypted signals in headers and reroutes the traffic. Traffic has to pass through specific routers.Need to trust large multi-national telecom, US Federal Gov., or similar entity.
- Botnets – use DGA and double fast flux to hide their traffic and the location of their servers.

What else can be done?

- BGP has issues
  - Fat fingered route leaks
  - Pakistan removed Youtube from global Internet
  - Guadalajara – DC traffic routed through Belarus
  - UK nuclear arms traffic routed through Russia/Ukraine
- Few people monitor BGP route problems and detection takes days.
- IPv6 address space is enormous and almost empty.
- If you map all IPv4 addressed twice onto IPv6, the chance of a collision is essentially zero.
- Could we use Hedy Lamarr's insights on the IP space?

Option 1 – BGP abuse

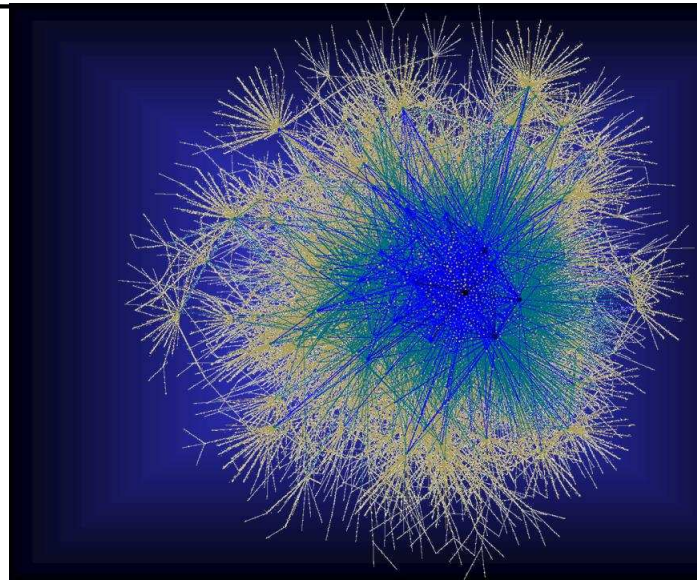

- We know how to foil DPI, mimic protocols, fool timing side-channels.
- Remaining problem – IP's in clear-text.
- Can we do IP redirection?
- Can we do IP randomization?
- Full randomization limited . BGP squatting violates civil contracts.
- We can test. BGP point of presence changes IP ranges (within allowable ranges) using a schedule shared with the client.
- Intermediate Floodlight Openflow controllers re-write IP destination addresses on the fly.
- Prototype tests concept of IP address hopping.
- Drawbacks: BGP injection overhead, requires synchronization, route update delays, limited number of updates.

Option 2 – SDX

- Next step.
- Instead of BGP, use SDN infrastructure.
- BGP route updates not needed. Software Defined Internet Exchanges (SDX) used.
- Re-routing potentially less transparent.
- Randomization not necessarily end-to-end. Potentially at multiple points.
- Re-routing not necessarily using IP addresses, we have used MAC addresses for similar things.
- Need to establish how to reduce overhead, make more robust, reduce information leakage.
- Lots of possibilities. More questions than answers for now.

PEERING Testbed

- Can an experimenter perform BGP interoperation experiments with today's Internet?
- Yes! The PEERING BGP testbed connects with real networks (via BGP), running experiments to exchange BGP routes and traffic directly with these networks.
- PEERING controls 8 ASNs (e.g. 47065, 61575, etc)
- PEERING controls 3 IPv4 prefixes (e.g. 184.164.224.0/19, 138.185.228.0/22, etc) and 1 IPv6 prefix (e.g. 2804:269c::/32)
- PEERING has points of presence (PoPs) at multiple internet exchange points and universities. Experimenters advertise and withdraw BGP prefixes at selected PoP.

How We Use PEERING

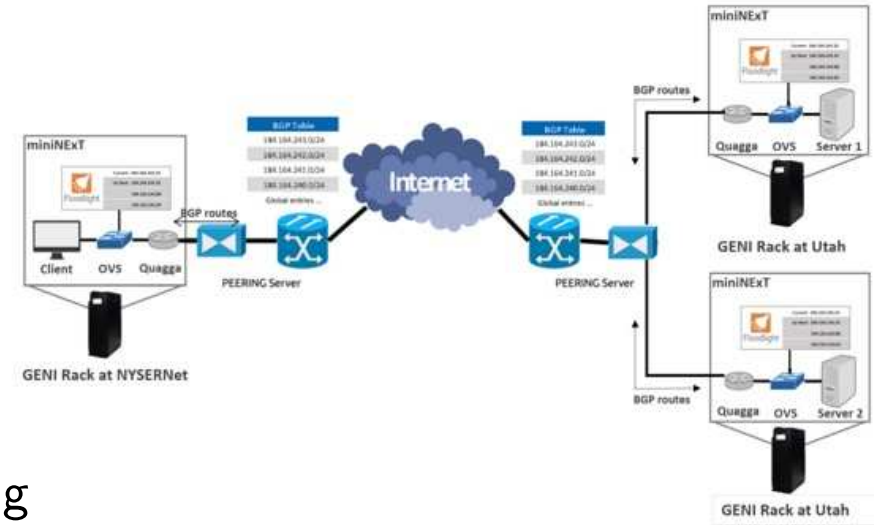

- Allotment of four prefixes: 184.164.240-243.0/24 Two PoPs: Amsterdam and Seattle
- miniNExT network on GENI node connects to PEERING using an OpenVPN tunnel
- BGP update limit of 18 update messages per day
- Advertise two type of BGP update messages
  - Announcement
  - Withdrawal
- Swap BGP prefix from one server to another:
  - Withdraw existing BGP prefix on one server
  - Insert that BGP prefix to another server

Maintaining Communication Sessions

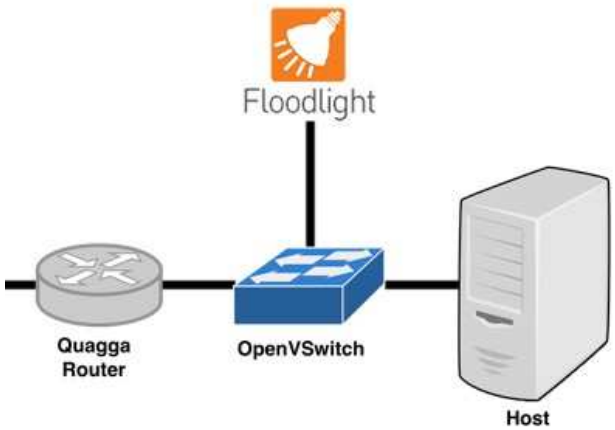

- Application on host communicates with a TARN server using server's identifier address (e.g. 10.0.0.1)
- Traffic from application travels through an OVS connected to the Floodlight OpenFlow controller
- Floodlight, recognizing the identifier address, inserts flows on the OVS to rewrite the identifier address to the current randomized IP address and vice versa
- New flows are inserted as the randomized IP address changes
- Because the rewrites are transparent to the application, the same socket can be used on both sides of the connection for the durationof the session
- The infrastructure can be located within a single host or on the edge of an