


Protocol Evolution

and its Impact on Network Operators

Ian Swett, IETF QUIC Editor

Data Flow



https://en.wikipedia.org/wiki/Internet_protocol_suite#/media/File:IP_stack_connections.svg

What Operators Want

1. Operate the Network

- **Allocate Resources** - link capacity, firewall capacity, services like proxy/cache, DNS...
- **Resolve Issues** - application faults, connectivity problems, excessive latency...
- **Assure Availability** - failover, redundancy...

2. Secure the Network

- **Identify** anomalous traffic / endpoints
- **Mitigate** threats
- **Scan** for virus / malware


3. Impose Policy

- Data Loss Prevention
- Content Filtering
- Cost Allocation / Charging
- “Quality of Service”
- Audit
- Access Control (e.g., Captive Portals)
- Child / Prisoner / Student / Employee / Citizen Monitoring

What's Changing

HTTP → HTTPS

August 2017



Globally, more than **half** of web browsing is HTTPS

<https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/felt>

HTTP/2

- Standard in 2015, now in all browsers, 45% of responses
- Major changes:
 - **Multiplexing**
 - **Header Compression**
 - **Server Push**
 - **Connection Coalescing**
 - (Practically) **Mandatory Encryption**

<https://http2.github.io>

HTTP/2 Operator Impact

- **New wire format** - if you intercept, don't assume 1.1
- **One connection/origin** - more fair, but loss more evident
- **More hosts than just SNI** - less fine grained
- **Forward Secrecy** - passive monitoring doesn't work

TLS 1.3

- Finishing touches on standard; support in Firefox Nightly and Chrome Canary. OpenSSL, et al coming.
- Major changes:
 - **1RT or 0RT Handshake**
 - **Pare down / modernise crypto**
- SNI still in the clear (for now)
- Operator impact:
 - **All PFS, all the time** - passive monitoring doesn't work

ORIGIN + Secondary Certs

- ORIGIN allows a server to specify which hosts a connection can be used for.
- Secondary Certificates allow a server to prove authority for new hosts.
- Use cases:
 - Advanced connection coalescing
 - Domain fronting
- Operator impact: **harder to identify/filter traffic**

QUIC

- Currently deployed by Google, others; in standardisation
- Major changes:
 - **UDP-based, multi-stream semantics**
 - **Avoids TCP HoL blocking**
 - **Collapses** transport/crypto/application protocol stack
 - **Allows mobility** - connection ID
 - **Encrypt all the things** - including transport metadata

<https://quicwg.github.io>

Privacy Concerns

Identifying HTTPS-Protected Netflix Videos in Real-Time

Andrew Reed, Michael Kranch

Dept. of Electrical Engineering and Computer Science

United States Military Academy at West Point

West Point, New York, USA

{andrew.reed, michael.kranch}@usma.edu

ABSTRACT

After more than a year of research and development, Netflix recently upgraded their infrastructure to provide HTTPS encryption of video streams in order to protect the privacy of their viewers. Despite this upgrade, we demonstrate that it is possible to accurately identify Netflix videos from passive traffic capture in real-time with very limited hardware requirements. Specifically, we developed a system that can report the Netflix video being delivered by a TCP connection using only the information provided by TCP/IP headers.

To support our analysis, we created a fingerprint database comprised of 42,027 Netflix videos. Given this collection of fingerprints, we show that our system can differentiate between

protected Netflix videos. We then improve upon the previous work by fully automating the fingerprint creation process, thereby enabling us to create an extensive collection of Netflix fingerprints which we then use to conduct a robust assessment of the attack. Finally, we developed a network appliance that can, in real-time, identify HTTPS-protected Netflix videos using IP and TCP headers obtained from passive capture of network traffic.

Our primary contributions are:

- A dataset that contains the fingerprints for 42,027 Netflix videos.
- An automated crawler that creates Netflix video fingerprints.

<https://dl.acm.org/citation.cfm?id=3029821>

QUIC Wire Format

5.2. Short Header

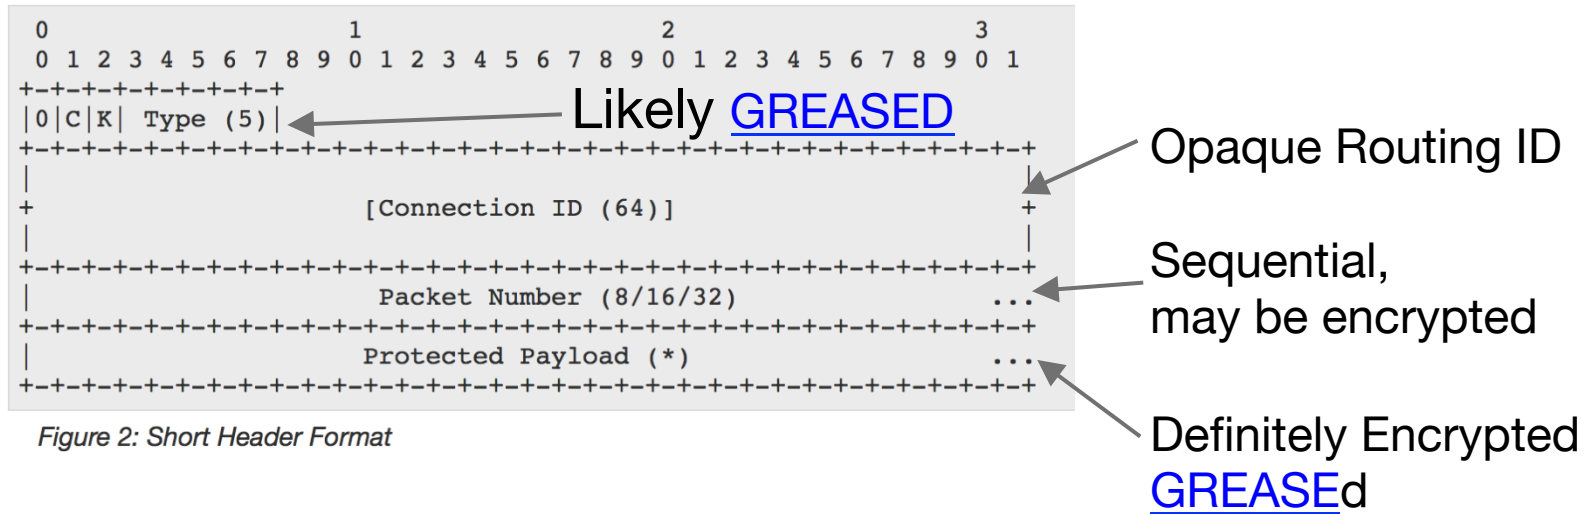


Figure 2: Short Header Format

QUIC Connection IDs

QUIC Connections are identified by Connection IDs

- Not identified by 5-tuple
 - 5-tuple may change (NAT rebinding)
- Connection IDs may change
 - During the handshake to a server chosen ID
 - During migration to a new network (WiFi to Cell)

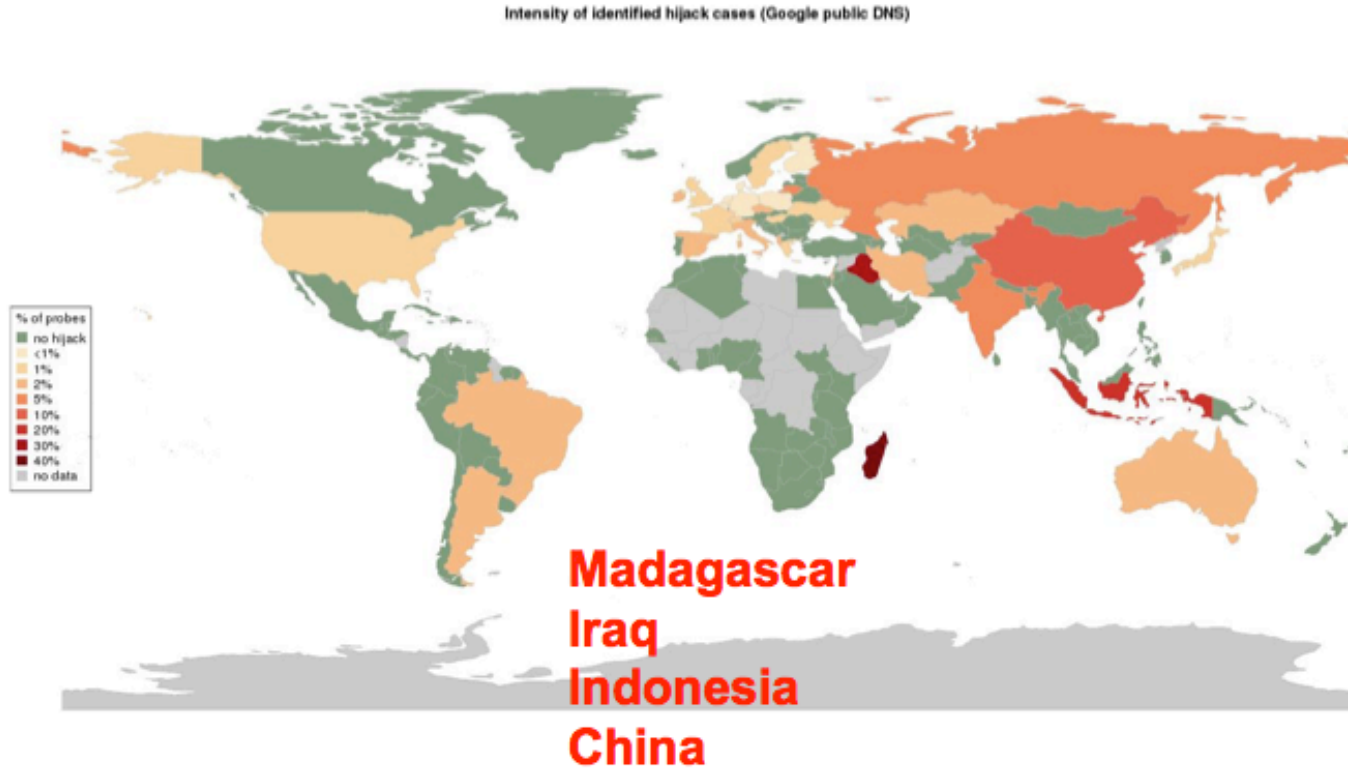
QUIC Operator Impact

- **New transport protocol** - tools, equipment support
- **Shift to UDP** - breaks assumptions
- **Encrypted metadata, incl ACKs, RST**
 - Passive estimation of latency / loss no longer feasible
 - Network can't just RST conns it doesn't like
- **Connections no longer identified by 5-tuple**
 - ... and connection-ID is optional

DOH!

- DNS-over-HTTPS
- Some ad hoc deployment (e.g., Google Public DNS)
- IETF WG formed in Sept 2017
- Use case?

Results: Google DNS hijacks (%)



DOH Operator Impact

- **Split DNS** - doesn't work (?)
- **DNS-based policy enforcement** - doesn't work
- **DNS-based data gathering** - doesn't work

Summary

- The Internet enables permissionless innovation by design; there's a lot of recent and ongoing activity
- Assumptions about availability of transport and application protocol information & control to networks are likely to be invalidated
- Focus on strong encryption, reduction of metadata
- Push towards applying policy / mitigations in endpoints
- If this causes issues in operability, **please get involved**
 - ... but be aware that there is a healthy amount of skepticism about unsupported claims!