Classifying Encrypted Traffic

PANRG Interim

June 3, 2020

John Border – Hughes Network Systems

John.Border@Hughes.com

Traffic Classification

- From Wikipedia

  Traffic classification is an automated process which categorises computer network traffic according to various parameters (for example, based on port number or protocol) into a number of traffic classes.[1] Each resulting traffic class can be treated differently in order to differentiate the service implied for the data generator or consumer.

  Reference [1] is RFC 2475 – An Architecture for Differentiated Services

- Classification is primarily based on transport header and application header information
  - TCP and UDP port numbers are the most basic classifiers
    - Port number overload (especially on Port 443) makes this difficult

The “First Mile”

- For network operators, traffic classification is important for meeting customer service expectations
- The last mile is really the first mile when it comes to traffic classification
  - The network for which classification is most important is the one to which the end user is directly connected
  - Besides prioritizing interactive over background with respect to available resources the network device may even have multiple paths with different characteristics available to it

Encryption

- Pervasive encryption is being used to hide metadata for privacy protection
  - **This is a good thing!**
- Encryption hides the transport and application information used for classification from the network operator
- Deep packet inspection, DNS correlation and SNI examination techniques can be used to try to classify traffic based on packet sizes, patterns, etc.
- Unfortunately, any technique which can be used by a network operator to classify traffic can also be used by other entities for other purposes

Signaling the Path

- Can the end user's device signal to the "first mile" network operator device the required classification information in such a way that it does not get propagated beyond that device?
  - DSCPs are one option but the information needs to be cleared before forwarding
    - Age old problem of can you trust the source to not be greedy?
  - Use a MASQUE-like technique to securely send metadata to the "first mile" device separate from the end to end connection?

Signaling the Path

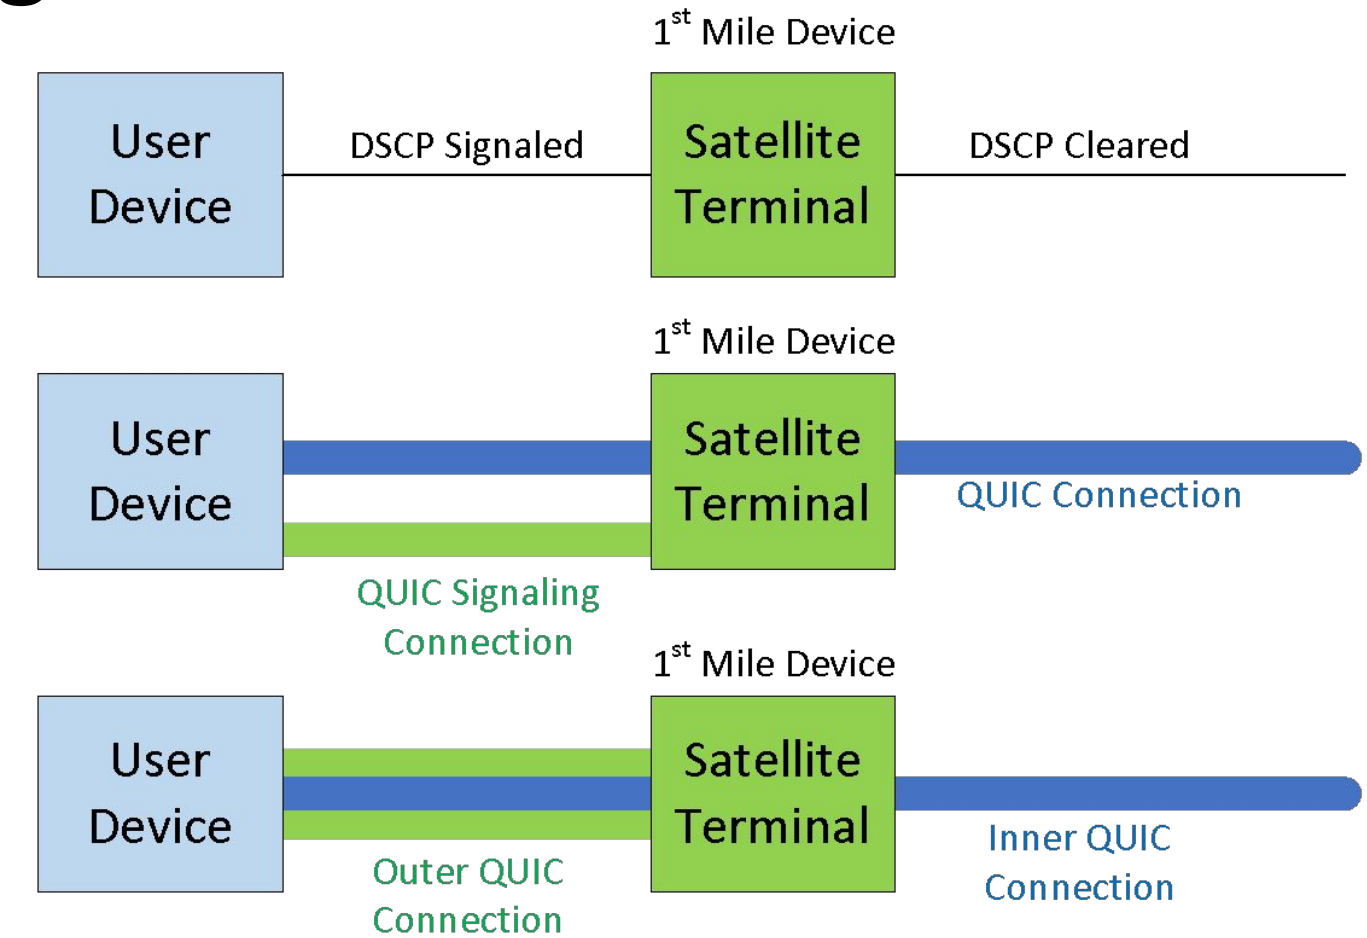

Discussion?

- How do we get end user (and, perhaps more importantly, application developer) buy-in?