Discovery of Network Rate-Limit Policies (NRLPs)

draft-brw-scone-rate-policy-discovery

IETF#121, Dublin

M. Boucadair (***Orange***)
D. Wing (***Cloud Software Group***)
T. Reddy.K (***Nokia***)
S. Rajagopalan (***Cloud Software Group***)
G. Mishra (***Verizon***)
M. Amend (***Deutsche Telekom***)
L. Contreras (***Telefonica***)

Networks Are Already Sharing Network Properties with Hosts

- Examples
  - DNR
  - Link Maximum Transmission Unit (MTU)
  - Prefixes of Network Address and Protocol Translation from IPv6 clients to IPv4 servers (NAT64)
  - …
- Additional signal advices can be enabled by networks if:
  - ***Appealing incentives to deploy***: Direct benefits a network can get
  - ***Tradeoffs*** (read ***costs***) ***are worth the benefits***
  - These are consistent with ***regulatory*** constraints (net neutrality, typically)
  - These do not ***distort the trust model***

Design Rationale: Focus on Deployability

- Requiring a new data plane feature to ***intercept, rewrite, etc. packets is too complex*** (let alone not justified) for the intended SCONE case

- Instead, we suggest to leverage existing/deployed protocols
  - ***Fate sharing***
  - ***Atomic configuration***
  - ***Updatability***: Change the policy at any time
  - ***Deployability***
    - Designed to ***ease integration*** with network management tools that are used to manage and expose policies
    - ***Consistent*** with rate-limit specifications that are deployed in the network (RFCs 2697/2698)

- NRLP is defined as a ***reusable blob*** that can be carried in DHCP, DHCPv6, Router Advertisement, PvD, etc.

NRLP as SCONE Protocol: Features

- ***Zero-trust security*** model, in which only a network element authorized to send RA/DHCP/PvD can provide the 'throughput' signal
- Application ***fairness*** to access the advice
- Applicable to ***any transport protocol***
- ***Feature parity*** upon transport fallback (e.g., QUIC to TCP)
- ***Cascaded environments*** are supported (e.g., LAN + Access Network)
- Ability to signal ***per-host*** and ***per-subscriber*** policy (tethering, etc.)
- Supports signaling policies ***bound to one or both traffic directions***
- ***Immune against*** changes of the UDP 4-tuple to bypass rate-limits
- ***Independent*** of access technology
- ***Clear accountability model*** for troubleshooting and service checks
- ***Extensible***

NRLP as SCONE Protocol DOES NOT…

- Require to ***reveal the identity of the server*** to consume the advice
- Allow for ***random on-path hosts*** on the Internet to originate the advice
- Require any ***inspection in the network of data packets*** to trigger or inject the advice
- Require any ***data plane change***
- Impact the ***forwarding performance*** of network nodes –no rewriting is needed
- Impact the ***MTU*** tweaking
- Impact the connection ***setup delay***
- Suffer from side effects of ***multi-layer encryption***
- Suffer from ***nested congestion control*** for tunneled proxy mode
- Incur the ***overhead of unauthenticated re-encryption*** of QUIC packets in the scramble transform of the forwarding mode
- Suffer from the complications of ***IP address sharing*** (RFC 6269)
- Require manipulating extra ***steering policies on the host*** to decide which flows can be forwarded over or outside a proxy connection

NRLP as SCONE Protocol DOES NOT…

Improvements vs. QUIC Tweaks

- Require to ***reveal the identity of the server*** to consume the advice
- Allow for ***random on-path hosts*** on the Internet to originate the advice
- Require any ***inspection in the network of data packets*** to trigger or inject the advice
- Require any ***data plane change***
- Impact the ***forwarding performance*** of network nodes –no rewriting is needed
- Impact the ***MTU*** tweaking
- Impact the connection ***setup delay***
- Suffer from side effects of ***multi-layer encryption***
- Suffer from ***nested congestion control*** for tunneled proxy mode
- Incur the ***overhead of unauthenticated re-encryption*** of QUIC packets in the scramble transform of the forwarding mode
- Suffer from the complications of ***IP address sharing*** (RFC 6269)
- Require manipulating extra ***steering policies on the host*** to decide which flows can be forwarded over or outside a proxy connection

NRLP as SCONE Protocol DOES NOT…

- Require to ***reveal the identity of the server*** to consume the advice
- Allow for ***random on-path hosts*** on the Internet to originate the advice
- Require any ***inspection in the network of data packets*** to trigger or inject the advice

- Require any ***data plane change***
- Impact the ***forwarding performance*** of network nodes –no rewriting is needed
- Impact the ***MTU*** tweaking
- Impact the connection ***setup delay***
- Suffer from side effects of ***multi-layer encryption***
- Suffer from ***nested congestion control*** for tunneled proxy mode
- Incur the ***overhead of unauthenticated re-encryption*** of QUIC packets in the scramble transform of the forwarding mode
- Suffer from the complications of ***IP address sharing*** (RFC 6269)
- Require manipulating extra ***steering policies on the host*** to decide which flows can be forwarded over or outside a proxy connection

Improvements vs. Masque proxy

Beyond Encoding bits: Ops Considerations Matter

How NRLP Fits in the WG Charter?

This WG aims to establish a mechanism for network elements capable of rate-limiting a UDP 4-tuple to communicate an upper bound on achievable bitrate, termed "throughput advice", to the sender of packets matching the UDP 4-tuple.

The WG is expected to:

- Develop a proposed standard **protocol to communicate an upper bound on achievable bitrate** – termed **"throughput advice"– from network elements to the endpoint**.

- Develop an Informational Applicability and **Manageability** specification.

Next Steps

- Welcome comments, suggestions, and contributions

Appendix

Example of the NRLP Blob in Nighbor Discovery

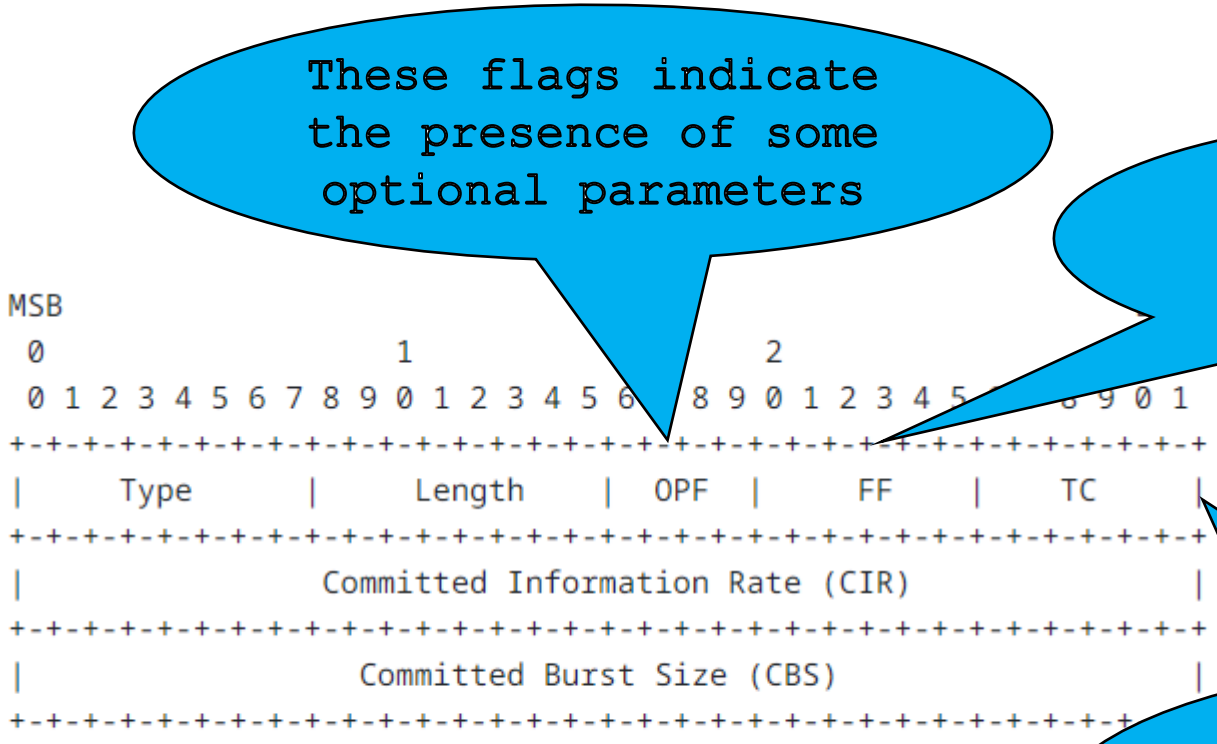

Figure 3: NRLP Option Format with Mandatory Fiel

These flags indicate the presence of some optional parameters

Flow Flags (FF):
- S (Scope)
- D (Direction)
- R (Reliablity)

TC (Traffic Category) specifies a traffic category to which a policy applies

Example of the NRLP Blob: PvD Example

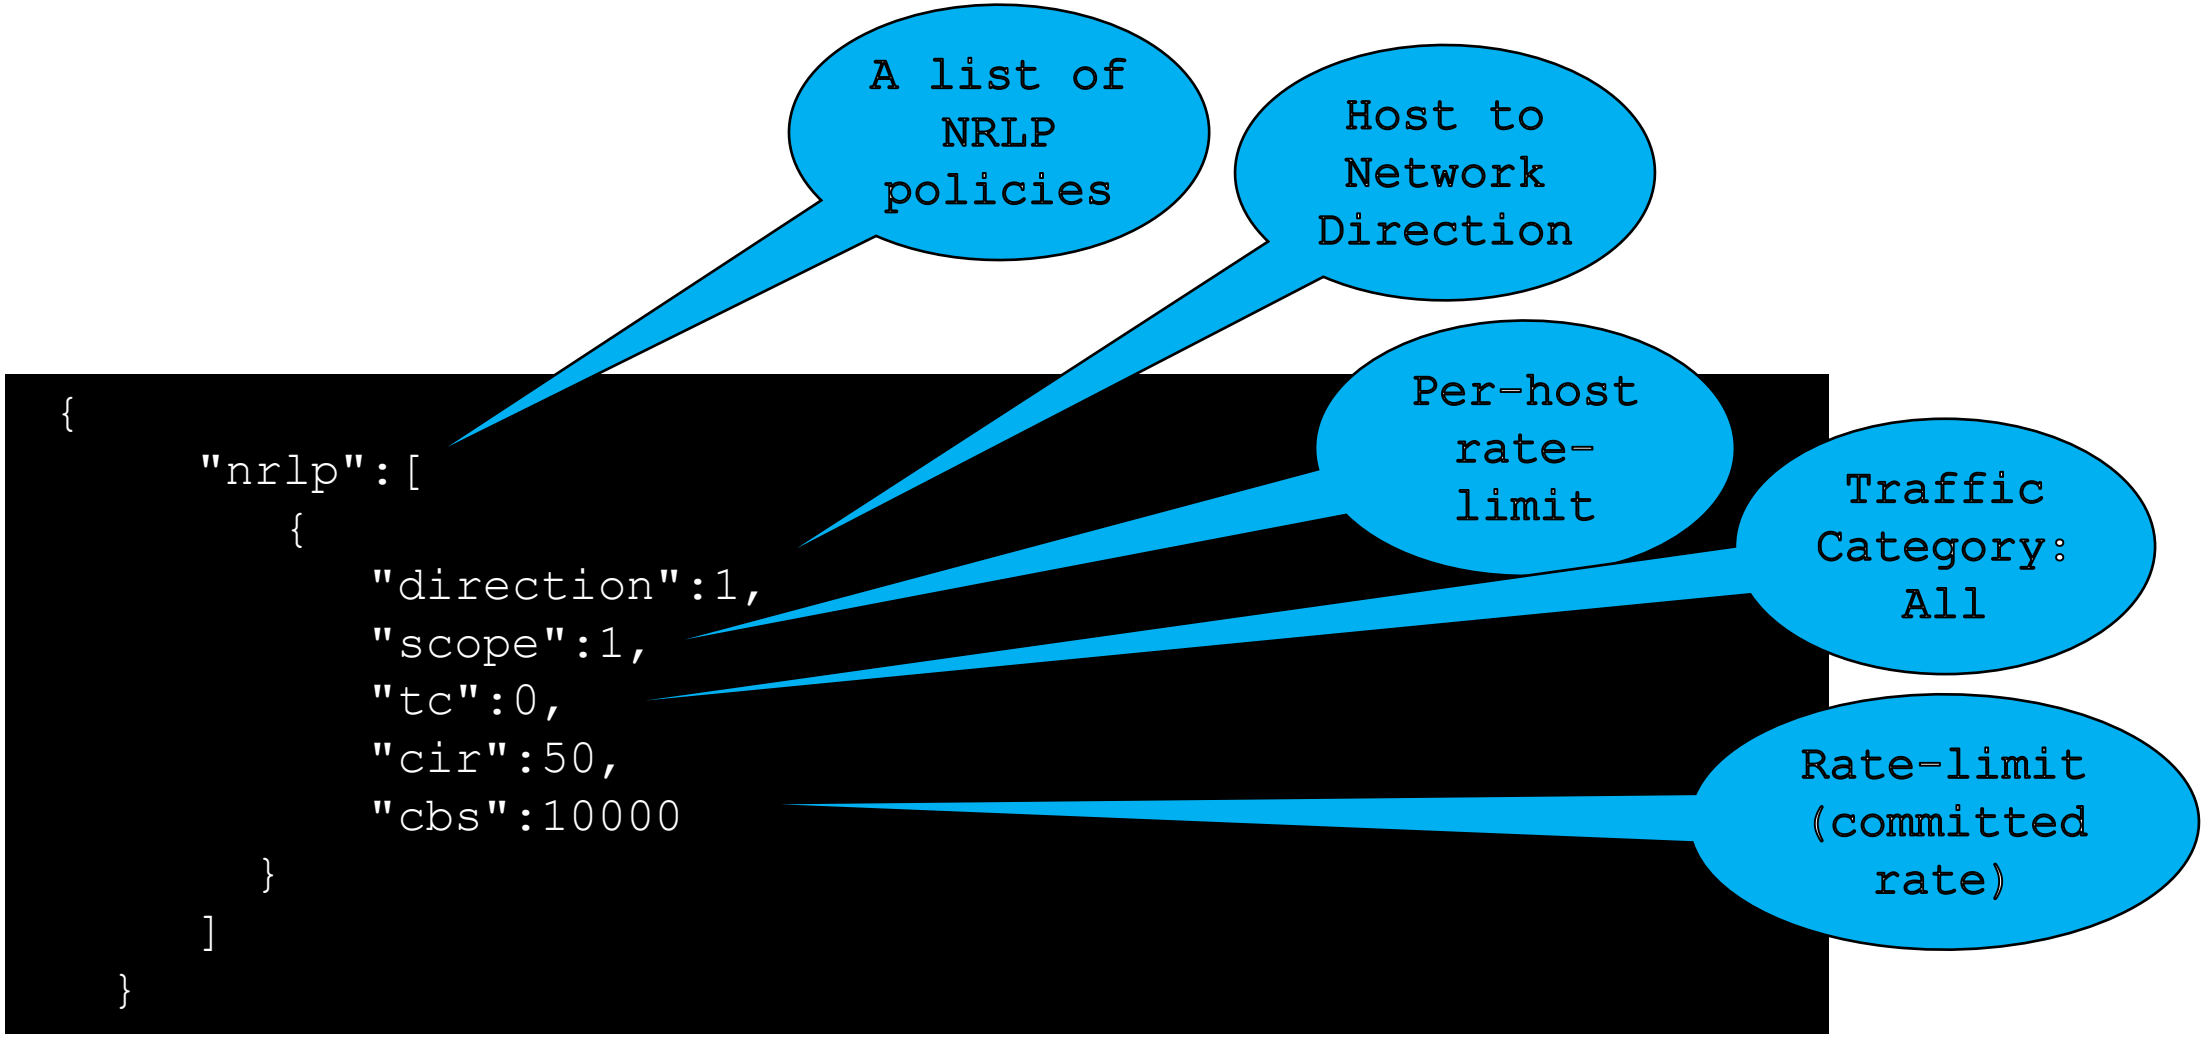

Sample Cascaded Deployment

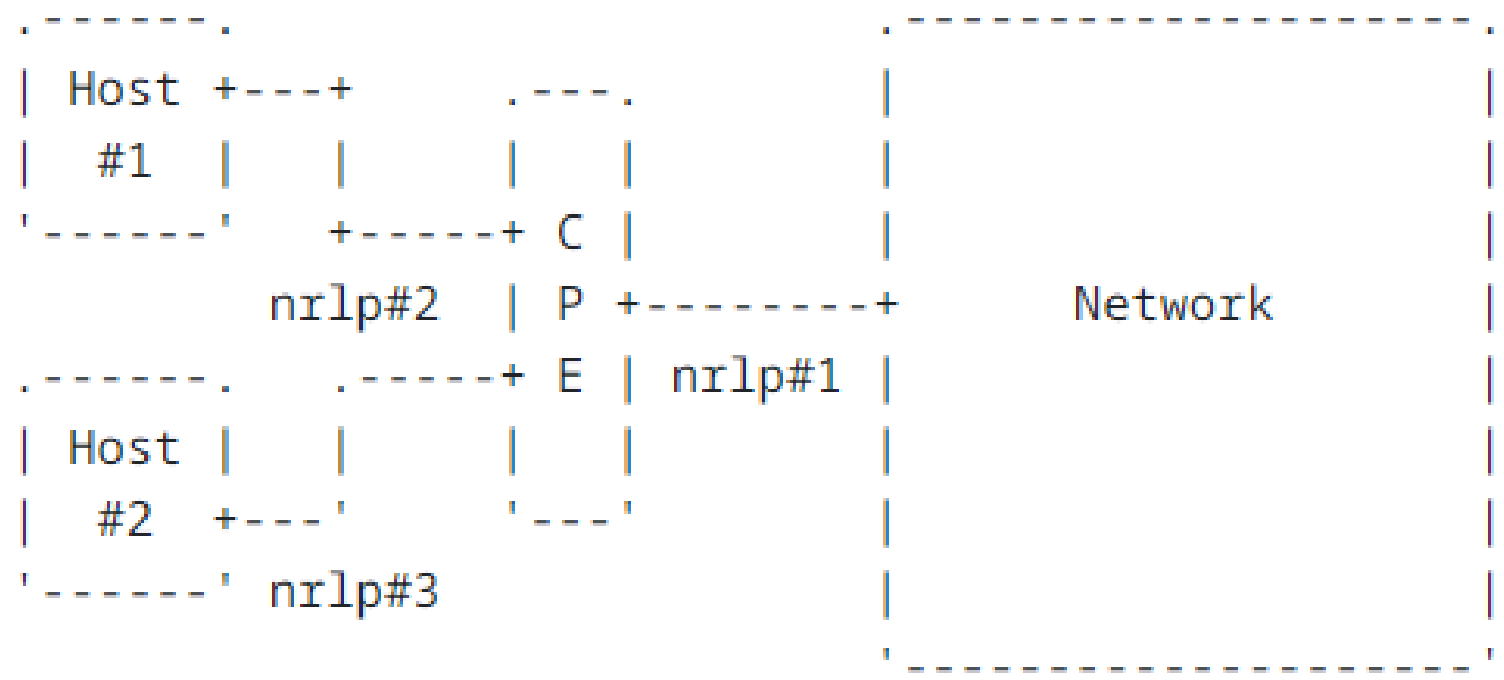

Leverage Existing AAA Architectures: An Example