scone
Internet-Draft
Intended status: Informational
Expires: 30 October 2025

D. Wing
Cloud Software Group
T. Reddy
Nokia
S. Rajagopalan
Cloud Software Group
L. Contreras
Telefonica
28 April 2025

SCONE Solution Analysis
draft-brw-scone-analysis-01

Abstract

This document provides an analysis of various SCONE solutions to share the throughput advice.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 30 October 2025.

Table of Contents

1. Introduction

The document provides an analysis of proposed SCONE solutions to share the throughput advice. The currently analyzed solutions (listed in alphabetic order) are as follows:

MASQUE: "MASQUE extension for signaling throughput advice" [I-D.ihlar-scone-masque-mediabitrate]

See Section 3.2.

NRLP: "Discovery of Network Rate-Limit Policies (NRLPs)" [I-D.brw-scone-rate-policy-discovery]

See Section 3.3.

TRONE: "Transparent Rate Optimization for Network Endpoints (TRONE) Protocol" [I-D.thoji-scone-trone-protocol]

See Section 3.4.

2. Criteria Classification

The following categories are used to classify the various criteria:

Security/Privacy (Sec): Indicates whether this impacts security/privacy. Some of the criteria that are classified as security-related may also have implications on the efficiency of sharing an advice (e.g., as that is likely to be ignored).

Some security/privacy criteria are as follows:

* Zero-trust security: Only authorized network elements must provide the throughput advice.

* Privacy: Indicates whether a solution does not reveal any details about the app or server identity.

* Mobility: Indicates whether a solution supports guards against a malicious app that keeps changing the 5-tuple to evade rate-limit enforcement by the network.

Deployability (Dep): Captures criteria that are important for unlocking the deployment of a solution at both network and host sides.

A deployability hurdle would be typically the misalignment of incentives between those receiving the benefit vs. those bearing the cost of providing the benefit (Section 3.3 of [I-D.narten-radir-problem-statement]). For example, the sender of the advice should see (immediate) benefits.

Some other deployability criteria are as follows:

* Fate sharing: reflects whether the mechanism used to advertise the throughput advice shares the fate of the rest of the network configuration on the host.

* Atomic configuration: Indicates whether the throughput advice can be learned using very few packets and whether changes to the policy require sharing the entire policy or just the relevant part.

Performance (Per): May impact the performance of the network device that enables the solution and/or the performance of the flow.

Service Interference (Int): Captures implications on other services (e.g., side effects).

For example, tweaking MTU may have an implication on all the flows that share the same network attachment, not only those that consumes an advice. Likewise, requiring address sharing has a plenty of issues that are discussed in [RFC6269]. Also, relying upon an explicit proxy would penalize the proxy which could serve both good and 'bad' clients (e.g., launching Layer 7 DDoS attacks).

Functional (Fun): Characterizes the functional capabilities offered by activating a solution.

Some examples of functional criteria are as follows:

* Updatability: indicates whether a solution allows to update hosts with policy changes at any time.

* Path coupled signaling/Path decoupled signaling: Indicates whether solution allows for the entity to share the advice be on-path or off-path. This criterion is also meant to assess the deployment flexibility offered by a solution.

* Support cascaded environments: Rate-limits may be enabled at several levels. For example, rate-limits may be enforced on the CPE in the home network for the endpoints attached to it and in the provider network to rate-limit the traffic from the subscriber. This criterion indicates whether such setups are supported.

A criterion may belong to one or more categories.

Criteria	Sec	Dep	Per	Int	Fun
Protocol ossification					X
Zero-trust security	X				

Privacy	X				
Guard against random advice injection by an on-path attacker	X				
Mobility (guard against changing 5-tuple)	X				X
Require guards against app abuse	X				X
Fate sharing		X			
Atomic configuration		X			
Updatability					X
Integration with network management tools		X			
Applicable to QUIC					X
Applicable to any application					X
Require an OS API		X			
Requires PvD		X			
Support cascaded environments					X
Path coupled signaling		X			X
Path decoupled signaling		X			X
Traffic direction (h2n, n2h, both)					X
Per-host policies					X
Per-subscriber policies					X
Extendable					X
Require data plane upgrade/ change		X			
Require transport payload inspection (network)		X			

Require transport payload inspection (host)		X			
Require flow inspection and tracking (network)	X				
Require steering policies on the host		X			
Depend on the server to consume the signal		X			
Impact the connection setup delay					X
Require the identity of the target server	X				X
Require MTU tweaking		X		X	
Incur multi-layer encryption		X	X		
Incur nested congestion control		X	X		
Incur multiple round-trips		X	X		
Forwarding peformance impact		X	X	X	
IP address sharing issues		X		X	
Penalizing the proxy		X		X	

Table 1: Criteria Classification

3. Detailed Analysis

3.1. Summary

Criteria	MASQUE	NRLP	TRONE	Else
Protocol ossification	TBC	N	TBC	TBC
Zero-trust security	TBC	Y	TBC	TBC
Privacy	TBC	Y	TBC	TBC

Guard against random advice injection by an on-path attacker	TBC	Y	TBC	TBC
Mobility (guard against changing 5-tuple)	TBC	Y	TBC	TBC
Require guards against app abuse	TBC	Y	TBC	TBC
Fate sharing	TBC	Y	TBC	TBC
Atomic configuration	TBC	Y	TBC	TBC
Updatability	TBC	Y	TBC	TBC
Integration with network management tools	TBC	Y	TBC	TBC
Applicable to QUIC	TBC	Y	TBC	TBC
Applicable to any application	TBC	Y	TBC	TBC
Require an OS API	TBC	Y/N(p)	TBC	TBC
Requires PvD	TBC	Y(p)/N	TBC	TBC
Support cascaded environments	TBC	Y	TBC	TBC
Path coupled signaling	TBC	Y	TBC	TBC
Path decoupled signaling	TBC	Y	TBC	TBC
Traffic direction (h2n, n2h, both)	TBC	Y	TBC	TBC
Per-host policies	TBC	Y	TBC	TBC
Per-subscriber policies	TBC	Y	TBC	TBC

Extendable	TBC	Y	TBC	TBC
Require data plane upgrade/change	TBC	N	TBC	TBC
Require transport payload inspection (network)	TBC	N	TBC	TBC
Require transport payload inspection (host)	TBC	N	TBC	TBC
Require flow inspection and tracking (network)	TBC	N	TBC	TBC
Require steering policies on the host	TBC	N	TBC	TBC
Depend on the server to consume the signal	TBC	N	TBC	TBC
Impact the connection setup delay	TBC	N	TBC	TBC
Require the identity of the target server	TBC	N	TBC	TBC
Require MTU tweaking	TBC	N	TBC	TBC
Incur multi-layer encryption	TBC	N	TBC	TBC
Incur nested congestion control	TBC	N	TBC	TBC
Incur multiple round-trips	TBC	N	TBC	TBC
Forwarding peformance impact	TBC	N	TBC	TBC
IP address sharing issues	TBC	N	TBC	TBC

Penalizing the proxy	TBC	N	TBC	TBC

Table 2: Analysis Summary

Notes: (p) indicates the assessment when PvD is used as NRLP mechanism.

3.2. MASQUE (to be completed by the authors of MASQUE)

3.2.1. Key Idea

3.2.2. Discussion

3.2.3. Main Expected Gains

3.2.4. Costs

3.3. NRLP

3.3.1. Key Idea

NRLP leverages existing discovery mechanisms (DHCP, RA, PvD) for networks to advertise throughout advices. The same generic blob is used independent of the signaling mechanism. NRLP operates within the existing network/host trust model.

Also, NRLP does not introduce additional dependency that would hinder having the benefits of enabling the NRLP feature.

3.3.2. Discussion

Only network elements that are entitled to send DHCP/RA/PvD configuration are allowed to share the throughput advices. As such, NRLP has built-in:

* zero-trust model

* Guard against random advice injection

Taking into account that NRLP advices are bound to a traffic category, NLRP relies upon the OS to enforce the received policies for applications falling under a traffic category (or all traffic). In doing so, NRLP adheres to the following:

* Mobility (guard against changing 5-tuple)

* Require guards against app abuse: The OS can allocate network resources more fairly among different processes, with NRLP signals, ensuring that no single process monopolizes the network.

NRLP meets the following criteria:

* Fate sharing: RA/DHCP are needed anyway so that connectivity is provided over a network attachment. NRLP ensures that throughput advices shares the fare of the other network configuration on the host.

* Atomic configuration: Only one packet (e.g., RA) is required to share the advice. Also, only a specific portion of the configuration can be provided.

* Updatability/Proactive signaling: It is possible to change the policy at any time and notify hosts (e.g., by sending a new RA).

Given that NRLP advices are shared during the establishment of a network attachment and then as part of the maintenance of the attachment, NRLP is therefore:

* Applicable to any transport protocol: This allows specifically to ensure a feature parity for applications that fallback to another transport protocol (e.g., QUIC to TCP).

* Applicable to QUIC

* Applicable to any application

To that aim:

* RA/DHCP NRLP requires an OS API to expose the signal to applications, and ensure application fairness.

* If PvD is used, an app only needs to learn the PvD ID from the OS (which is not specific to NRLP) and the PvD additional information can be retrieved by the app itself (without any dependency on the OS).

NRLP leverages existing mechanisms for the provisioning of network attachments, including supply of the various policies ([I-D.ietf-opsawg-ntw-attachment-circuit]). Also, NRLP leverages AAA mechanisms (e.g., [RFC9445]). Therefore, NRLP eases:

* Integration with network management tools

One of NRLP flavors:

* Requires PvD discovery. This is not required for DHCP/RA.

NRLP does not restrict the deployment options as providers can deploy distributed or centralized DHCP servers, use relays, enable NRLP RA in access routers, etc. Similar to other network configuration purposes, NRLP has the following capabilities:

* Support cascaded environments. The throughput advice can even be correlated with local conditions or policies as shown, e.g., in Figure 1.

* Path coupled signaling

* Path decoupled signaling

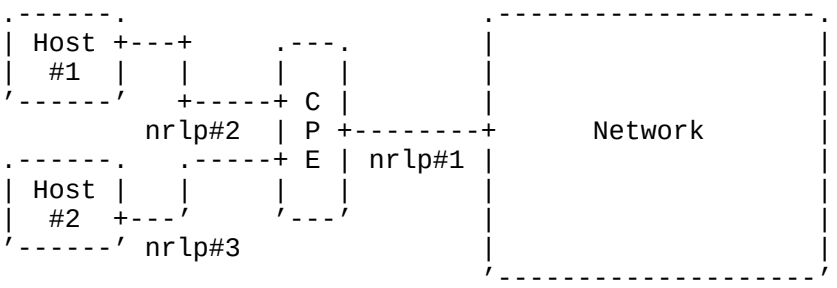

Figure 1: Example of Cascaded NRLPs

The same generic blob is used in NRLP independent of the signaling mechanism. The blob is designed with the following key characteristics:

* Traffic direction (h2n, n2h, both): policies for one or both directions can be supplied.

* Per-host policies: An explicit indication in inserted in the advice to tag per-host policies.

* Per-subscriber policies: An explicit indication in inserted in the advice to tag per-subscriber policies. This covers deployment scenarios such as tethering or CPE-based service offerings.

* Provide provisions for extensions: NLRP includes provisions for future attributes that are tracked in IANA registries.

Given that NRLP leverages existing control plane mechanisms, NRLP does not:

* Suffer from protocol ossification issues

* Require data plane upgrade/change

* Require transport payload inspection (network)

* Require transport payload inspection (host)

* Require flow inspection and tracking (network)

Also, given that NRLP signals are exchanged before connection establishment, NRLP does not:

* Depend on the server to consume the signal: NRLP advices are immediately consumable by applications and do not require involving a remote server.

* Require the identity of the target server to receive or consume the advices.

Moreover, NRLP does require any encapsulation or proxy function at the network. As such, NRLP does not:

* Require steering policies on the host to decide which flows are eligible to the proxy service.

* Impact the connection setup delay: NRLP signals are available on bootstrap of a host (and prior to any connection establishment).

* Require MTU tweaking

* Incur multi-layer encryption

* Incur nested congestion control

* Incur multiple round-trips: The signal is immediately available in one packet (RA NRLP, typically).

* Overhead of unauthenticated re-encryption

* Forwarding performance impact

* IP address sharing issues: NRLP does not require changing the source IP address used by a host.

* Penalize any network node (a proxy, typically) which could serve both good and bad clients (e.g., launching Layer 7 DDoS attacks).

3.3.3. Main Expected Gains

* Lower deployment barrier to experiment in large scale (no hardware or software change is needed in network components).

* Schedule network requests (independent of the transport protocol) more efficiently, preventing network congestion, and improving overall stability and network performance.

* Unlock new services in local networks and enhance the quality of experience at the LAN by providing a simple tool to communicate local policies to hosts.

* Provide a mechanism to assist networks managing the load at the source and, thus, contribute to better handle network overloads and optimize the use of resources under non nominal conditions.

3.3.4. Costs

* A simple configuration is required for IPv4: DHCP flavor can be provided by configuration of custom options. Refer to [NRLP-WIRE].

* A similar configuration approach can be followed for DHCPv6.

* A minor change to the network is required for NRLP RA: upgrade configuration of PE nodes with new Neighbor Discovery option. Note that all IPv6 hosts and networks are already required to support Neighbor Discovery [RFC4861].

* An API needs to be exposed on the host to share the advice with applications (e.g., scutil on MacOS). No additional API is needed if PvD is used.

3.4. TRONE (to be completed by the authors of TRONE)

3.4.1. Key Idea

3.4.2. Discussion

3.4.3. Main Expected Gains

3.4.4. Costs

4. Security Considerations

Security-related criteria are analyzed for each proposed solution.

5. IANA Considerations

This document does not make any IANA request.

6. References

6.1. Normative References

[I-D.brw-scone-rate-policy-discovery]
Boucadair, M., Wing, D., Reddy.K, T., Rajagopalan, S., Mishra, G. S., Amend, M., and L. M. Contreras, "Discovery of Network Rate-Limit Policies (NRLPs)", Work in Progress, Internet-Draft, draft-brw-scone-rate-policy-discovery-02, 16 December 2024, <https://datatracker.ietf.org/doc/html/draft-brw-scone-rate-policy-discovery-02>.

[I-D.ihlar-scone-masque-mediabitrate]
Ihlar, L. M. and M. K端hlewind, "MASQUE extension for signaling throughput advice", Work in Progress, Internet-Draft, draft-ihlar-scone-masque-mediabitrate-02, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-ihlar-scone-masque-mediabitrate-02>.

[I-D.thoji-scone-trone-protocol]
Thomson, M., Huitema, C., Oku, K., Joras, M., and L. M. Ihlar, "Transparent Rate Optimization for Network Endpoints (TRONE) Protocol", Work in Progress, Internet-Draft, draft-thoji-scone-trone-protocol-00, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-thoji-scone-trone-protocol-00>.

6.2. Informative References

[I-D.ietf-opsawg-ntw-attachment-circuit]
Boucadair, M., Roberts, R., de Dios, O. G., Barguil, S., and B. Wu, "A Network YANG Data Model for Attachment Circuits", Work in Progress, Internet-Draft, draft-ietf-opsawg-ntw-attachment-circuit-16, 23 January 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-ntw-attachment-circuit-16>.

[I-D.narten-radir-problem-statement]
Narten, T., "On the Scalability of Internet Routing", Work in Progress, Internet-Draft, draft-narten-radir-problem-statement-05, 17 February 2010, <https://datatracker.ietf.org/doc/html/draft-narten-radir-problem-statement-05>.

[NRLP-WIRE] "Examples of Wire Format Options", <https://github.com/boucadair/draft-xxx-ac-rate-policy-discovery/blob/main/example-nrlp-wire-format.md>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/rfc/rfc4861>.

[RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <https://www.rfc-editor.org/rfc/rfc6269>.

[RFC9445] Boucadair, M., Reddy.K, T., and A. DeKok, "RADIUS Extensions for DHCP-Configured Services", RFC 9445, DOI 10.17487/RFC9445, August 2023, <https://www.rfc-editor.org/rfc/rfc9445>.

Authors' Addresses

Dan Wing
Cloud Software Group Holdings, Inc.
United States of America
Email: danwing@gmail.com

Tirumaleswar Reddy
Nokia
India
Email: kondtir@gmail.com

Sridharan Rajagopalan
Cloud Software Group Holdings, Inc.
United States of America
Email: sridharan.girish@gmail.com

Luis M. Contreras
Telefonica
Spain
Email: luismiguel.contrerasmurillo@telefonica.com