Applicability & Manageablity SCONE WG Interim

April 30, 2026

Authors:
Sanjay Mishra, Verizon
Zahed Sarker, Nokia
Anoop Tomar, Meta
Khurram Abbas, Verizon

Post-01 Updates Addressing GitHub & IETF 125 Feedback

- Eliminated Protocol Redundancies: Pruned sections (including deleting the QoS section entirely) that merely restated core protocol mechanics, transitioning the document fully away from its old "requirements" roots into a concise guide for network operators.
- Clarified Stateless Operation: Rewrote the Flow Awareness section to explicitly state that the act of applying SCONE advice is inherently stateless, and that maintaining per-flow context is only necessary if operators choose to implement conformance monitoring.
- Congestion Control Interworking: Added explicit guidance separating SCONE's advisory throughput limits from transport-level congestion signals (like ECN/L4S), ensuring operators understand they are complementary but distinc
- Addressed no network element, single network element and multi network-element: Added new sections to address how multiple network elements operate independently (the most restrictive signal wins) and how stateless signaling handles access network handovers without requiring state transfers
- Added Operational Considerations section, more work to be done here
- Overhauled almost all sections. Need to publish -02 to ensure all outstanding comments are addressed and all redundant text is removed.

Issue #7 in Section 1: Forwarding Changes & Network Elements

Issues (summarized):

1. Operators will need to think through is that they do not have sole right to change the SCONE information, rather any network elememt along the path could reduce the value
2. A short paragraph explaining how the SCONE endpoint would operate with no SCONE network element on path and also with multiple network elements on the same path.
3. para on what happens when a forwarding path change results in a change of the SCONE Network Element that provides the advice. Section 6.1 of the protocol Spec suggests that missing the start of a flow (because a traffic starts on a different path, could result in the SCONE Network element on the new path not seeing the start of the flow

PR #32 Response:

- Introduction: Added text to recognize end-to-end data paths may contain zero, one, or multiple SCONE-capable elements.
- Added section: Presence of SCONE Network Elements
- Added section: Change of Network Element During an Active Flow

Issue: #8: Section 3.4: Please clarify words

Issue:

- Which element saves the mentioned CPU overhead? - I think this is the network element?

PR #28: Response:

1. Change Heading from SCONE Indication to the Network Element to Considerations of Processing Complexity (per Mirja)
2. Added Text:
    - SCONE-aware endpoints provide an indication to the SCONE Network Element, enabling it to identify the SCONE-capable flow without any need for compute-intensive flow classification. Additionally, SCONE-capable endpoints, through bit-rate self-adaptation, remove the need for complex rate-limiting functions in the network element. Support for SCONE indication and bit-rate self-adaptation reduces complexity and CPU processing load in the network element.

Issue #9: Section 3.5: How is retransmission realized?

Issues:

- (Gorry): Please say more about how a network element perform retransmission of the SCONE Packets? What specifically is retransmitted and how?
- On top of PR addressing above, additional review on PR were:
  - (Marcus): but the endpoints can not take element CPU into consideration. Endpoints can send as frequently as they see fit, but the elements make decisions on update frequency based on considerations of CPU-load etc
  - (Marcus): A network that has dynamic policies will likely care much more about timely updates than a network that has fixed policies (e.g., subscription-based policies that are not intended to change over SCONE-like timescales
  - (Mirja): Remove all text redundant to the protocol spec

PR #29 Response:

- Addresseed the core issue and reworked PR to address issues raised by Marcus & Mirja.

Issue 10#: In Section 3.6: Please clarify words around updates

Issue:

- The rate at which SCONE “updates are issued". By the sending endpoint I assume?

PR #30 Response:

- Noted, frequency of SCONE signaling is fundamentally driven by the application endpoint
- ABR video clients fetching short media segments, may choose to send SCONE packets frequently. non-ABR applications, for example bulk transfer applications such as background software updates, may function effectively by sending SCONE packets less frequently.
- While the core SCONE protocol defines the baseline timers to prevent advice from expiring, operators should expect the actual frequency of passing SCONE packets to vary significantly depending on the application type

Issue #11: Section 3.9, Please clarify words around compliance

Issue:

- (Gorry): Can this document explain what compliance means with respect to SCONE?
- (Mirja): should talk more about how to do compliance measurements, e.g. using a sliding window, and what to do if a violation is detected
- (Mirja): or not do anything at all. We should note that the scone signal is also valuable to provide without monitoring and compliance enforcement!!!

PR #33 Response:

- Added text per Gorry's initial comment followed by updates based on other comments from Mirja, For e.g.,
- Pure advisory nature of SCONE is highly valuable on its own, and it should not be implied that operators are expected to deploy strict rate limiters just because they are sending throughput advice. Updated the text to explicitly state that operators can choose to do nothing at all regarding enforcement, treating the SCONE signal as purely informational for cooperative endpoints

Issue #12: Section 3.1, Please clarify words about "stable”

Issue:

- Can you quantify relatively stable with an example, would an example be greater than multiple seconds? 10s seconds? Etc
- (Magnus): Wording on Transport is confusing in 1st Para of the PR #34
- (Mirja): More important to talk about application use cases here
- (Mirja): Why is operator enforcement out of policy?
- (Mirja): Questions use/contect of word “harmonizing”?

PR #34 Response:

- Updated PR to address above issue on top of core issue explaining CC vs SCONE. Example text
- For network operators considering co-deployment, SCONE throughput advice is strictly independent of the IP-layer ECN field. Because SCONE advice is carried within the QUIC payload, updating the advice does not interact with or modify ECN markings

Issue #14: Section 3.3: Please consider the interaction with other configured IP mechanisms

Issue:

- (Gorry): consideration of how a Scone signal interacts with diffserv and other methods?
- (Gorry): What is intended to happen when an application uses two (or more) DSCPs
- (Mirja): PR wording seems to contradict Section 5.3 in the scone protocol spec. Still not sure what this paragraph really tells me…really saying more or less nothing and I don't think needs to be further discussed here.

PR #27 Response:

- PR and the existing section on the chopping block. There is a new section “Determining Throughput Constraints" to cover how operators map policies to advice much more concretely.
- Given substantial make up to the document we can revisit document to see if Gorry’s original question is answered otherwise find the right section to add some text (authors’ tbd)

Issue #15: Section 3.2: Please clarify position for ECN

Issue:

- Some words explaining that the SCONE marking for a flow is independent of the ECN Field may be useful - since ECN appeared many times in the discussion of SCONE
- Other comments from Magnus and Mirja

PR #34 Response:

- Details (to get to the PR) same as on the previous slide

Issue #16: Add discussion about the relation of scone and L4S

Issue:

- Section 3.11 “abut (*sic*) “Talks about” congestion control and also mentioned L4S but it would be good to talk more about when it makes sense to use one of the other or both. the protocol spec already explains that scone does not replace congestion control. the applicability statement should focus more on when to use things and how.

PR #34 Response:

- Please see PR #34. Hopefully, it addresses the issue. Need more eyes!!

Issue #17: Discuss network deployment use cases with or without throttling fallback

Issue: Section 3.9 should talk more about how to do compliance measurements, e.g. using a sliding window, and what to do if a violation is detected

PR #34 Response:

- Conformance Monitoring now describes using a sliding window approach (evaluating flows against the highest limit advised over the preceding two monitoring periods) to account for application adaptation time. And also add explicit guidance on operator deployment choices when a violation is detected, for example, operators can either employ a throttling fallback (falling back to traditional rate-limiting mechanisms like dropping or delaying packets) or choose a purely advisory deployment without strict enforcement.

Issue #18: What's the difference between 3.1 (Flow session awareness) and 3.2 (Per-Flow Signaling)

Issue:

- Scone signalling is per flow; not sure what is meant with a session here. however, I also don't think the content of section 3.2 is fully correct. You can just put a certain rate into the scone field for every scone packet that you see; that would be fully stateless

PR #27 Response:

- Collapsed Sections 3.1 & 3.2 (Flow Session Awareness & Per-Flow Signaling).
- Created new section **Flow Awareness and Per-Flow Signaling**. Clarify
  - Throughput advice strictly applies to specific QUIC flows (UDP 4-tuple)
  - Stateless Signaling: Identifying a packet and applying advice is inherently stateless and does not require active flow context
  - Stateful Monitoring: Maintaining per-flow context is only necessary as an operational foundation if the operator chooses to execute monitoring, logging, and conformance evaluation

Issue #19: Section 3.4, QoS handling s not clear

Issue:

- (Mirja): Why is this needed? Scone provides one specific piece of information. Any other other QoS handling seems actual orthogonal to me. Can we maybe just remove this section? Or otherwise you would need to better explain what the scone specific considerations are.
- (Gorry): An operational concept could be how the network element ought to handle QoS classes, since Section 3.3 of the Scone Protocol states: "The signaled advice applies to the flow of packets on the same UDP address tuple for the duration of the current monitoring period" ... Likely the network element needs to provide advice consistent with the QoS class being supported - thinking the advice is provided for a longer time-scale?

Response:

- Same issue as Issue #19.
- Remove QoS Awareness section [PR to be submitted]. To do: Sanjay

Issue #20: Explain SCONE in MASQUE Proxies

Issue:

- (Mirja): A masque proxy could also serve as scone network element. Would be good to have a short section to describe that this is possible if scone advice is applied before/after en/de-capsulation. Also there might also be a case to use scone for tunnel traffic...? If that's a case maybe we need to discuss how to copy over the outer scone to the inner scone?

PR #37 Response:

- Mirja has proposed text. To be discussed.

Issue #21: Explain how throughput constraints are determined

Issue:

- do you measure actual bit rate per flow, do you observes its local radio and network conditions, do you subscription-based policy?
- Add Operational Consideration Section to align with RFC5706bis

PR #39 Response:

- Started a new section Operational Considerations and added subsection Determining Throughput Constraints, discussing:
  - Subscriber Policies and Data Plans
  - Application-Specific Policies
  - Dynamic Network Conditions
  - Capacity and Load Management

Issue #22: Explain Throttling Policy and impact of Throttling

Issue:

- (Qin): Explain Throttling Policy and impact of Throttling
- (Marcus): This text reads a bit like marketing, which isn't necessarily a bad thing. However, I think that text like this might be better suited in an introduction section

PR #36 Response:

- Trying to addres points raised at IETF 125, such as “should I implement and deploy this protocol or not?”
- Part of the text discusses, “By providing a standardized and scalable mechanism, SCONE allows network operators and QUIC endpoints to exchange bit-rate information without custom APIs or per-network integrations. SCONE improves user experience by enabling the network to provide bit-rate guidance directly to applications, allowing them to self-adapt instead of relying on network rate limiters such as policers or shapers. This avoids packet drops and throttling, resulting in better Quality of Experience (QoE)”
- Plan to revisit the text and perhaps distill down to a few key points and maybe move the text under Operations Considerations section (TBD)

Issue #23: SCONE at the application level or transport level?

Issue: CONE at the application level or transport level?

PR #34 Response:

- Reference back to PR #34 covering interworking with Other Congestion Control Mechanism. Need to revisit to make sure

Issue #24: Retransmission vs send upated advice more often every monitoring interval

Issue: Retransmission vs send upated advice more often every monitoring interval

PR #33 Response:

- When evaluating compliance, network operators will need to account for the time required for SCONE packets to be updated, received by endpoints, and acted upon by the application. Operators can accommodate this by utilizing a sliding window approach. Specifically, as established in {{Dynamic Updates}} section, operators should evaluate QUIC flows against the highest throughput limit advised over the preceding two monitoring periods (a span of 134 seconds). If a network element cannot update the throughput advice in every traversing SCONE packet, operators might configure a longer sliding window to account for the possibility of packet loss.
- To simplify the measurement function, reduce computational load, or offload this function to another node in the network, operators can select any value larger than the baseline 67-second window for their measurement and averaging period.

Issue #25: How do you know SCONE signaling fails to reach the intended endpoints

Issue:

- (Qin): How do you know SCONE signaling fails to reach the intended endpoints?
- (Gorry): The answer could fall into two parts: How does the sender detect? How can a network operator understand?

PR #29 Response:

- While endpoints send SCONE packets as frequently as they see fit to ensure reliable delivery, the Network Element makes independent decisions on how frequently to update those packets to mitigate packet loss. This decision relies on operational considerations such as CPU load and the nature of the network policies. A network enforcing dynamic policies will prioritize timely updates to minimize the delay in activating the advised bit-rate after a packet loss. Conversely, a network enforcing fixed, subscription-based policies that do not change over SCONE timescales can safely scale back the Network Element update frequency to conserve CPU resources, as timely recovery from packet loss is less critical.