Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 August 2026

W. Quan
W. Su
S. Pan
Beijing Jiaotong University
X. Liang
J. Liang
China Telecom
28 January 2026

IOAM network awareness for Low Latency, Low Loss, and Scalable Throughput (L4S)

draft-quan-l4s-ioam-01

Abstract

This specification defines a framework how to update L4S Dual-Queue Coupled AQM with In Situ Operations, Administration, and Maintenance (IOAM). These are designed for consistently very low queuing latency, very low congestion loss, and scaling of perflow throughput by using Explicit Congestion Notification (ECN) using the operational and telemetry information collected by IOAM. Since L4S lacks information about the use of network status and network nodes, which also affect network performance in practice, it is considered to use direct export mode for information collection of L4S-IOAM to strengthen the AQM regulation of L4S. It then gives the normative requirements that are necessary.

Copyright Notice

Table of Contents

1. Introduction

The L4S architecture [RFC9330] allows routers to use a different marking system that can provide early reaction to incipient congestion [RFC9332] and defines a reaction for this feedback when packets are marked with ECN. However, the single-bit ECN feedback in L4S provides limited visibility into the precise location and severity of transient congestion across the path. The application of IOAM technology in L4S framework can effectively solve this problem. In Situ Operations, Administration, and Maintenance (IOAM) is used for recording and collecting operational and telemetry information while the packet traverses a path between two points in the network. The IOAM data fields are further updated in [RFC9326] for direct export use cases. This document defines how to use the information collected by the front-end nodes to better update the L4S mechanism. IOAM can collect operational and telemetry information. L4S uses an Explicit Congestion Notification (ECN) scheme at the IP layer with the same set of codepoint transitions as the original (or 'Classic')

ECN.

The goal of L4S-IOAM is to solve the problem of how to get information awareness between the IOAM network and the L4S site. The basis to achieve this goal is network and computing. Therefore, Network Information Awareness (NIA) system architecture is proposed. As the control plane of the L4S-IOAM framework, NIA introduces the control center component on top of the L4S-IOAM framework to realize the management and comprehensive analysis of network information and encourage L4S site to take action to consider local network awareness.

This specification defines how the data collected by IOAM can be used to better update the Low Latency, Low Loss, and Scalable throughput (L4S). The terms "encapsulation" and "decapsulation" are used in this document in the same way as [RFC9197]. An IOAM encapsulating node incorporates one or more IOAM Option-Types into packets that IOAM is enabled for.

2. Terminology

L4S: Low Latency, Low Loss, and Scalable Throughput (L4S) as defined in [RFC9330].

L4S Dual-Queue Coupled AQM: A framework for coupling the Active Queue Management (AQM) algorithms in two queues intended for flows with different responses to congestion.

IOAM: In situ Operations, Administration, and Maintenance as defined in [RFC9197].

OAM: Operations, Administration, and Maintenance.

IOAM Transit Node (IOAM-T): An IOAM transit node updates one or more of the IOAM-Data-Fields. If both the Pre-allocated and the Incremental Trace Option-Types are present in the packet, each IOAM transit node will update at most one of these Option-Types.

IOAM Encapsulating Node (IOAM-E): An IOAM encapsulating node incorporates one or more IOAM Option-Types into packets that IOAM is enabled for. If IOAM is enabled for a selected subset of the traffic, the IOAM encapsulating node is responsible for applying the IOAM functionality to the selected subset.

IOAM Decapsulating Node (IOAM-DE): An IOAM decapsulating node removes any IOAM Option-Types from packetsand processes and/or exports the associated data.

IOAM NODE ID (T-ID): The combination of IOAM node_id and IOAM Namespace-ID should always be unique.

Direct Export: Direct Export is an IOAM mode of operation within which IOAM data are to be directly exported to a collector rather than be collected within the data packets. The IOAM Direct Export Option-Type consists of a fixed-size "IOAM direct export option header". Direct Export for IOAM is defined in [RFC9326].

3. IOAM network awareness in L4S framework

The following are system components for the L4S-IOAM.

```
                                  +--------------+       +------------+
                                 +-------------+ |       |    L4S     |
                                 | Monitoring/ | |======>|  QUAL-AQM  |
                                 | Analytics   | |       |            |
                                 | IOAM-C      |-+       +------------+
                                 +-------------+
                                        ^
                                        |  Processed/interpreted/
                                        |  aggregated IOAM data
                                        |
                                  +--------------+
                                 +-------------+ |
                                 | IOAM data   | |
                                 | processing  | |
                                 | system(s)   |-+
                                 +-------------+
                                        ^
                                        |  Raw export of
                                        |  IOAM data
                                        |
                 +--------------+-------+------+--------------+
                 |              |              |              |
                 |              |              |              |
    User     +---+----+     +---+----+     +---+----+     +---+-----+
    packets  | IOAM-E |     | IOAM-T |     | IOAM-T |     | IOAM-DE |
    -------->| Node   |====>| Node   |====>| Node   |====>| Node    |-->
             |        |     | A      |     | B      |     |         |
             +--------+     +--------+     +--------+     +---------+
```

Figure 1: L4S-IOAM Schematic

IOAM Control Center (IOAM-C): Store and manage network information and computing information, and make decisions through a comprehensive analysis of this information. The IOAM-C correlates the exported telemetry data (e.g., hop-by-hop delay, queue depth) to dynamically adjust the target and threshold parameters of the Coupled AQM at IOAM-capable nodes. This feedback loop enables the control plane to optimize L4S performance based on real-time network visibility.

IOAM Ingress Forwarder (IOAM-E): A network node with a similar SFC Classifier [RFC7665] forwarding function can classify, encapsulate (for example, add a packet header with a service path identifier using the NSH protocol [RFC8300]), and forward incoming traffic.

IOAM-E and IOAM-DE have a L4S Network Metric Agent (L-NMA), responsible for collecting network information. In L4S-IOAM, L-NMA reports the collected network information to IOAM-C through the IOAM-SBI Interface.

The following are system architecture for the L4S-IOAM.

```
                 (3)                  (2)                    (4)
               .-------^------..------------^------------------.
                                                            _________
  ,-(1)-----.                                 _____        | IOAM    |
 ; ________  :            L4S    -------.    |     | /_ _ _| Monitor |
 :|Scalable| :               _ __\     ||__\_|mark | \   | |_________|
 :| sender | :  __________  /    /     ||  / |_____| \   |  _________
 :|________|\; |          |/     -------'        ^    \  | |condit'nl|
  '---------'\_|  IP-ECN  |          Coupling    :     \_|_|priority |_\
   ________  / |Classifier|                      :     / | |scheduler| /
  |Classic |/  |__________|\       -------.    __:__  /  | |_________|
  | sender |                \_ __\ || | ||__\_|mark/|/   |
  |________|                     / || | ||  / |drop |/_ _|
                      Classic      -------'   |_____|\
```

 (1) Scalable sending host
 (2) Isolation in separate network queues
 (3) Packet identification protocol
 (4) Monitor in network queues

Figure 2: L4S-IOAM architecture

4. Information Details

IOAM for L4S is used to enhance diagnostics of L4S networks. It complements other mechanisms designed to enhance diagnostics of L4S networks, such as the "The Explicit Congestion Notification (ECN) Protocol" described in [RFC9331].

Figure 3 shows awareness information content examples for network aware which is used to provide L4S services.

Awareness information	Network information
Capability parameters	IOAM-F location; IOAM-F type; IOAM-F ID; Topology information.
Status parameters	Service request information; Traffic features; Communication information.

Figure 3: L4S-IOAM Information Details

"IOAM tracing data" is expected to be collected at every IOAM transit node that a packet traverses to ensure visibility into the entire path that a packet takes within an IOAM-Domain. In other words, in a typical deployment, all nodes in an IOAM-Domain would participate in IOAM and, thus, be IOAM transit nodes, IOAM encapsulating nodes, or IOAM decapsulating nodes. If not all nodes within a domain are IOAM capable, IOAM tracing information (i.e., node data, see below) will only be collected on those nodes that are IOAM capable.

IOAM tracing can, for example, collect the following types of information:

* Identification of the IOAM node. An IOAM node identifier can match to a device identifier or a particular control point or subsystem within a device.

* Identification of the interface that a packet was received on, i.e., ingress interface.

* Identification of the interface that a packet was sent out on, i.e., egress interface.

* Time of day when the packet was processed by the node as well as the transit delay. Different definitions of processing time are feasible and expected, though it is important that all devices of an IOAM-Domain follow the same definition.

* Generic data, which is format-free information, where the syntax and semantics of the information are defined by the operator in a specific deployment. For a specific IOAM-Namespace, all IOAM nodes should interpret the generic data the same way. Examples for generic IOAM data include geolocation information (location of the node at the time the packet was processed), buffer queue fill level or cache fill level at the time the packet was processed, or even a battery charge level.

* Information to detect whether IOAM trace data was added at every hop or whether certain hops in the domain weren’t IOAM transit nodes.

* Data that relates to how the packet traversed a node (transit delay, buffer occupancy in case the packet was buffered, and queue depth in case the packet was queued).

```
         Export of      Export of      Export of     Export of
         IOAM data      IOAM data      IOAM data     IOAM data
             ^              ^              ^             ^
             |              |              |             |
             |              |              |             |
User     +---+----+     +---+----+     +---+----+     +---+-----+
packets  | IOAM-E |     | IOAM-T |     | IOAM-T |     | IOAM-DE |
-------->| Node   |====>| Node   |====>| Node   |====>| Node    |-->
         |        |     | A      |     | B      |     |         |
         +--------+     +--------+     +--------+     +---------+
```

Figure 4: L4S-IOAM Direct Export Mode

Consider using Direct Export mode for L4S-IOAM information gathering. OAM information about each IOAM node a packet traverses is collected and immediately exported to a collector. Direct Export could be used in situations where per-hop tracing information is desired, but gathering the information within the packet -- as with per-hop tracing -- is not feasible. Rather than automatically correlating the per-hop tracing information, as done with per-hop tracing, Direct Export requires a collector to correlate the information from the individual nodes. In addition, all nodes enabled for Direct Export need to be capable of exporting the IOAM information to the collector.

Those content would allow L4S flows to achieve low latency, low loss and scalable throughput, but would sacrifice the more precise flow balance offered by.

5. UseCases

This section gives an example how the application of IOAM technology in L4S framework can effectively solve the problem that the forward node in the network is still congested before the L4S node, so the demand of L4S can also be met in L4S-IOAM, and it is conducive to reducing the overall delay of the network.

The following use cases for L4S are being considered by various interested parties:

* where the bottleneck is one of various types of access network, e.g., DSL, Passive Optical Networks (PONs), DOCSIS cable, mobile, satellite; or where it’s a Wi-Fi link

* private networks of heterogeneous data centres, where there is no single administrator that can arrange for all the simultaneous changes to senders, receivers, and networks needed to deploy DCTCP:

  - a set of private data centres interconnected over a wide area with separate administrations but within the same company

  - a set of data centres operated by separate companies interconnected by a community of interest network (e.g., for the finance sector)

  - multi-tenant (cloud) data centres where tenants choose their operating system stack (Infrastructure as a Service (IaaS))

* different types of transport (or application) congestion control:

  - elastic (TCP/SCTP);

  - real-time (RTP, RMCAT); and

  - query-response (DNS/LDAP).

* where low delay QoS is required but without inspecting or intervening above the IP layer :

  - Mobile and other networks have tended to inspect higher layers in order to guess application QoS requirements. However, with growing demand for support of privacy and encryption, L4S offers an alternative. There is no need to select which traffic to favour for queuing when L4S can give favourable queuing to all traffic.

* If queuing delay is minimized, applications with a fixed delay budget can communicate over longer distances or via more circuitous paths, e.g., longer chains of service functions [RFC7665] or of onion routers.

* If delay jitter is minimized, it is possible to reduce the dejitter buffers on the receiving end of video streaming, which should improve the interactive experience.

6. Contributors

Thanks to Xue Zhang, Ziheng Xu and Xuetong Hu for their contributions to this document.

7. IANA Considerations

None.

8. Security Considerations

For further study.

9. References

9.1. Normative References

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <https://www.rfc-editor.org/info/rfc8300>.

9.2. Informative References

[RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <https://www.rfc-editor.org/info/rfc7665>.

[RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <https://www.rfc-editor.org/info/rfc9197>.

[RFC9326] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", RFC 9326, DOI 10.17487/RFC9326, November 2022, <https://www.rfc-editor.org/info/rfc9326>.

[RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <https://www.rfc-editor.org/info/rfc9330>.

[RFC9331] De Schepper, K. and B. Briscoe, Ed., "The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9331, DOI 10.17487/RFC9331, January 2023, <https://www.rfc-editor.org/info/rfc9331>.

[RFC9332] De Schepper, K., Briscoe, B., Ed., and G. White, "Dual-Queue Coupled Active Queue Management (AQM) for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9332, DOI 10.17487/RFC9332, January 2023, <https://www.rfc-editor.org/info/rfc9332>.

Acknowledgments

The authors wish to thank Xue Zhang, Ziheng Xu and Xuetong Hu, for their invaluable comments.

Authors’ Addresses

Wei Quan
Beijing Jiaotong University
3 Shangyuan Cun, Haidian District
Beijing
P.R. China
Email: weiquan@bjtu.edu.cn

Wei Su
Beijing Jiaotong University
3 Shangyuan Cun, Haidian District
Beijing
P.R. China
Email: wsu@bjtu.edu.cn

Shuaihao Pan
Beijing Jiaotong University
3 Shangyuan Cun, Haidian District
Beijing
P.R. China
Email: 23111016@bjtu.edu.cn

Xiaobin Liang
China Telecom Research Institute
South District of Future Science and Technology, Changping District
Beijing
P.R. China
Email: liangxb2@chinatelecom.cn

Jie Liang
China Telecom Research Institute
South District of Future Science and Technology, Changping District
Beijing
P.R. China
Email: liangjie6@chinatelecom.cn