

Internet Engineering Task Force  
Internet-Draft  
Intended status: Informational  
Expires: January 4, 2015

P. Fan  
H. Deng  
China Mobile  
M. Boucadair  
France Telecom  
T. Reddy  
C. Eckel  
Cisco Systems, Inc.  
B. Williams  
Akamai, Inc.  
July 3, 2014

Application Enabled Collaborative Networking: Problem Statement  
draft-conet-aeon-problem-statement-01

## Abstract

Identification and treatment of application flows are important to many application providers and network operators. They often rely on these capabilities to deploy and/or support a wide range of applications. These applications generate flows that may have specific connectivity requirements that can be met if made known to the network. Historically, this functionality has been implemented to the extent possible using heuristics, which inspect and infer flow characteristics. Heuristics may be based on port ranges, network separation (e.g. subnets or VLANs, Deep Flow Inspection (DFI), or Deep Packet Inspection (DPI). But many application flows in current usages are dynamic, adaptive, time-bound, encrypted, peer-to-peer, asymmetric, used on multipurpose devices, and have different priorities depending on direction of flow, user preferences, and other factors. Any combination of these properties renders heuristic based techniques less effective and may result in compromises to application security or user privacy.

## 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 <http://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

Internet-Draft

AEON/CONET Problem Statement

July 2014

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 January 4, 2015.

## Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Definitions and Terminology . . . . .	<a href="#">3</a>
<a href="#">2.1.</a>	Types of Signaling . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Typical Workflows . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Limitations of Heuristic Based Solutions . . . . .	<a href="#">4</a>
<a href="#">5.</a>	Limitations of Existing Signaling Mechanisms . . . . .	<a href="#">5</a>
<a href="#">6.</a>	Efforts in Progress . . . . .	<a href="#">6</a>
<a href="#">7.</a>	Acknowledgements . . . . .	<a href="#">8</a>
<a href="#">8.</a>	Informative References . . . . .	<a href="#">8</a>
	Authors' Addresses . . . . .	<a href="#">9</a>

## [1.](#) Introduction

Networks today, whether public or private, are challenged with demands to support rapidly increasing amounts of traffic. New channels for creating and consuming rich media are deployed at a rapid pace. Pervasive video and access on demand are becoming second nature to consumers. Communication applications make extensive use of rich media, placing unprecedented quality of experience expectation on the underlying network. These trends present challenges for network forecast and planning operations.

Now more so than ever before, identification and treatment of application flows are critical for the successful deployment and operation of a growing number of business and household applications. These applications are based on a wide range of signaling protocols

and deployed by a diverse set of application providers that is not necessarily affiliated with the network providers across which the applications are used.

Historically, identification of application flows has been accomplished using heuristics, which infer flow characteristics based on port ranges, network separation, or inspection of the flow itself. Inspection techniques include DPI, which matches against characteristic signatures (e.g. key string, binary sequence, etc.) and DFI, which analyzes statistical characteristics and connection behavior of flows. Each of these techniques suffers from a set of limitations, particularly in the face of the network challenges outlined previously.

Heuristic-based approaches may not be efficient and require continuous updates of application signatures. Port based solutions suffer from port overloading and inconsistent port usage. Network separation techniques like IP subnetting are error prone and increase network management complexity. DPI and DFI are computationally expensive, prone to error, and become more challenging with greater adoption of encrypted signaling and secured media. An additional drawback of DPI and DFI is that any insights are not available, or need to be recomputed, at network nodes further down the application flow path.

As the IETF establishes default behaviors that thwart pervasive surveillance (e.g. [\[RFC7258\]](#)), it will be important to offer mechanisms that allow applications to request differential network treatment for their flows. The intent is to have applications protect the contents of their flows, yet have the ability to opt-in to information exchanges that provide a more precise allocation of network resources and thus a better user experience.

## [2.](#) Definitions and Terminology

### [2.1.](#) Types of Signaling

The following terms describe the relationship between signaling and the media to which it is associated.

- o off-path: signaling along a different network path than the media flow
- o on-path: signaling along the same network path as the media flow
  - \* in-band: signaling on the same port as the media flow
  - \* out-of-band: signaling on a different port than the media flow

### [3.](#) Typical Workflows

Various heuristic based approaches are used prevalently and successfully for the following workflows:

1. Provide network operators with visibility of traffic usage and patterns for troubleshooting, capacity planning, and other off network workflows. This is done by exporting observed traffic analysis via standard protocols such as IPFIX [[RFC7011](#)] and SNMP [[RFC3416](#)] as well as by proprietary protocols and methods.
2. Provide network operators with visibility of application and data usage for accounting and billing.
3. Provide differentiated network services for specific traffic classes according to network operator defined policies. Techniques to achieve this include traffic classification, policing and shaping (e.g. [[RFC2475](#)]), providing admission control (e.g. [[RFC6601](#)]), impacting routing, and permitting passage of specific traffic (e.g. firewall functions).

### [4.](#) Limitations of Heuristic Based Solutions

These workflows, visibility and differentiated network services, are critical in many networks. However, their reliance on inspection and observation limits their deployment. Reasons for this include the following:

- o Identification based on IP address lists is difficult to manage.

The addresses may be numerous and may change, they may be dynamic, private, or otherwise not meant to be exposed. With Content Delivery Network InterConnection (CDNI) [[RFC6770](#)], content could be served either from an upstream CDN (uCDN) or any of a number of downstream CDNs (dCDN), and it will not be possible to manually track the IP addresses of all the CDN surrogates. Even in cases where identification by IP addresses is possible, more granular identification of individual flows is not possible (e.g. audio vs. video vs. data).

- o Classification based on TCP/UDP port numbers often result in incorrect behaviour due to port overloading (i.e. ports used by applications other than those claiming the port with IANA).
- o More and more traffic is encrypted, rendering DPI and DFI impossible, inefficient, or much more complex, and sometimes at the expense of privacy or security (e.g. need to share encryption keys with intermediary proxy performing DPI/DFI).

- o Visibility generally requires inspecting the signaling traffic of applications. This traffic may flow through a different network path than the actual application data traffic. Impacting the traffic behavior is ineffective in those scenarios.
- o Extensions to signaling protocols and changes in the ways application use them can result in false negatives or false positives during inspection.
- o Inspection techniques are completely non-standard, so the ability and accuracy to identify traffic varies across vendors, and different implementation are likely to give different results for the same traffic.
- o Inspection techniques that require parsing the payload of packets (e.g. DPI) not only impact performance due to additional processing, but also impact memory due to the growing number and size of signatures to identify new protocols.
- o Network services leveraging heuristic based classification have a negative effect on the application behavior by impacting its traffic, while they do not provide explicit feedback to the

application. This results in a lost opportunity for the application to gain insight and adjust its operation accordingly.

## 5. Limitations of Existing Signaling Mechanisms

The IETF has standardized several mechanisms involving explicit signaling between applications and the network that may be used to support visibility and differentiated network services workflows. Unfortunately, none of these has experienced widespread deployment success, nor are they well suited for the applications usages described previously. Existing signaling options include the following:

- o RSVP [[RFC2205](#)] is the original on-path signaling protocol standardized by the IETF. It is transported out-of-band and could be used to signal information about any transport protocol traffic (it currently supports TCP and UDP). Its original goal was to provide admission control. Its requirement for explicit reservation of resources end to end proved too heavy for most network environments. Its success was further impacted by its reliance on router-alert, which often leads to RSVP packets being filtered by intervening networks, and by its requirement for access to a raw socket, something that is generally not available to applications running in user space. To date, more lightweight signaling workflows utilizing RSVP have not been standardized within the IETF.

- o NSIS (next Steps in Signaling) [[RFC5978](#)] is the next iteration of RSVP-like signaling defined by the IETF. It focused on the same fundamental workflow as RSVP admission control as its main driver, and because it did not provide significant enough use-case benefits over RSVP, it has seen even less adoption than RSVP.
- o DiffServ [[RFC4594](#)] and VAN Tagging [[IEEE-802.1Q](#)] style packet marking can help provide QoS in some environments, but such markings are often modified or removed at various points in the network or when crossing network boundaries. There are additional limitations when using DiffServ with real-time communications applications, and the DART working group has been chartered to write a document that explains the limitations that exist with DiffServ when used with RTP in general as well in the specific RTCWeb use cases [[I-D.ietf-rtcweb-use-cases-and-requirements](#)].

## 6. Efforts in Progress

Not surprisingly, there are several evolving proposals that aim to address the visibility and differentiated network services workflows where existing approaches are not sufficient. Protocol specific extensions are being defined, creating duplicate or inconsistent information models. This results operational complexity and a need to convert information between protocols to leverage the best protocol option for each specific use case. Examples of evolving signaling options include the following:

- o STUN [[RFC5389](#)] is an on-path, in-band signaling protocol that could be extended to provide signaling to on-path network devices. It provides an easily inspected packet signature, at least for transport protocols such as UDP. Through its extensions TURN [[RFC5766](#)] and ICE [[RFC5245](#)], it is becoming prevalent in application signaling driven by the initial use-case of providing NAT and firewall traversal capabilities and detecting local and remote candidates for peer-to-peer media sessions. The TRAM working group is chartered to update TURN and STUN to make them more suitable for WebRTC.
- o Port Control Protocol (PCP) [[RFC6887](#)] provides a mechanism to describe a flow to the network. The primary driver for PCP is creating port mappings on NAT and firewall devices. When doing this, PCP pushes flow information from the host into the network (specifically to the network's NAT or firewall device), and receives information back from the network (from the NAT or firewall device). It is not meant to be used end-to-end but rather independently on one "edge" of a flow. It is therefore an attractive alternative because it allows the introduction of

application to network signaling without relying on the remote peer. This is especially useful in multi-domain communications.

- o RESTCONF [[I-D.ietf-netconf-restconf](#)] is a REST-like protocol that provides a programmatic interface over HTTP for accessing data defined in YANG, using the datastores defined in NETCONF [[RFC6241](#)]. It is meant to provide a standard mechanism for web applications to access the configuration data, operational data,

data-model specific protocol operations, and notification events within a networking device, in a modular and extensible manner.

- o Interface to the Routing System (I2RS) is a working group chartered to provide interfaces for management applications, network controllers, and user applications to make specific demands on the network.
- o Abstraction and Control of Transport Networks (ACTN) is a non-working group mailing list intended to enable discussion of the architecture, use-cases, and requirements that provide abstraction and virtual control of transport networks to various applications/clients.
- o Prefix coloring has been proposed for use in HOMENET and 6MAN working groups to provide differentiated services to applications based on IP address.
- o RTP Media Congestion Avoidance Techniques (RMCAT) has been chartered to address the lack of generally accepted congestion control mechanisms for interactive real-time media, which is often carried via sets of flows using RTP over UDP. Explicit exchanges of flow characteristics and congestion information between applications and the network could play an important role in such mechanisms.
- o Transport Services (TAPS) is an effort to create a working group to define transport services that are exposed to internet applications. A TAP enabled application identifies its needs of the locally available transports services via an API. Some of the information provided is the same as what AEON proposes to have the application communicate to the network. Furthermore, the transport services of TAPS could benefit from this communication with the network.
- o Service Function Chaining (SFC) is a working group chartered to address issues associated with the deployment of service functions (e.g. firewalls, load balancers) in large-scale environments. Service function chaining is the definition and instantiation of an ordered set of instances of such service functions, and the

subsequent "steering" of traffic flows through those service



functions. Flow characteristics communicated via AEON could be used as input into an SFC classifier and it could be transported as SFC metadata.

## 7. Acknowledgements

The authors thank Toerless Eckert, Reinaldo Penno, Dan Wing, Amine Choukir, Paul Jones, and Bill VerSteeg for their contributions to this document.

## 8. Informative References

[I-D.ietf-netconf-restconf]

Bierman, A., Bjorklund, M., Watsen, K., and R. Fernando, "RESTCONF Protocol", [draft-ietf-netconf-restconf-00](#) (work in progress), March 2014.

[I-D.ietf-rtcweb-use-cases-and-requirements]

Holmberg, C., Hakansson, S., and G. Eriksson, "Web Real-Time Communication Use-cases and Requirements", [draft-ietf-rtcweb-use-cases-and-requirements-14](#) (work in progress), February 2014.

[IEEE-802.1Q]

"IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks", IEEE 802.1Q, 2005, <<http://standards.ieee.org/getieee802/download/802.1Q-2005.pdf>>.

[RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", [RFC 2475](#), December 1998.

[RFC3416] Presuhn, R., "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)", STD 62, [RFC 3416](#), December 2002.

[RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", [RFC 4594](#), August 2006.

- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", [RFC 5245](#), April 2010.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", [RFC 5389](#), October 2008.
- [RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", [RFC 5766](#), April 2010.
- [RFC5978] Manner, J., Bless, R., Loughney, J., and E. Davies, "Using and Extending the NSIS Protocol Family", [RFC 5978](#), October 2010.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", [RFC 6241](#), June 2011.
- [RFC6601] Ash, G. and D. McDysan, "Generic Connection Admission Control (GCAC) Algorithm Specification for IP/MPLS Networks", [RFC 6601](#), April 2012.
- [RFC6770] Bertrand, G., Stephan, E., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", [RFC 6770](#), November 2012.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [RFC 6887](#), April 2013.
- [RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, [RFC 7011](#), September 2013.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), May 2014.

Authors' Addresses

Internet-Draft

AEON/CONET Problem Statement

July 2014

Peng Fan  
China Mobile  
32 Xuanwumen West Street, Xicheng District  
Beijing 100053  
P.R. China

Email: fanpeng@chinamobile.com

Hui Deng  
China Mobile  
32 Xuanwumen West Street, Xicheng District  
Beijing 100053  
P.R. China

Email: denghui@chinamobile.com

Mohamed Boucadair  
France Telecom  
Rennes 35000  
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy  
Cisco Systems, Inc.  
Cessna Business Park, Varthur Hobli  
Sarjapur Marathalli Outer Ring Road  
Bangalore, Karnataka 560103  
India

Email: tiredddy@cisco.com

Charles Eckel  
Cisco Systems, Inc.  
170 West Tasman Drive

San Jose, CA 95134  
USA

Email: eckelcu@cisco.com

Fan, et al.

Expires January 4, 2015

[Page 10]

---

Internet-Draft

AEON/CONET Problem Statement

July 2014

Brandon Williams  
Akamai, Inc.  
8 Cambridge Center  
Cambridge, MA 02142  
USA

Email: brandon.williams@akamai.com

