

Transport Area Working Group
Internet-Draft
Intended status: Informational
Expires: August 2, 2008
Song
ETRI

J. Adams
Anagran, Inc
J. Joung
J.

June 24, 2008

**Progress and future development of Flow State Aware standards, and a
proposal for alerting nodes or end-systems on data related to a flow
draft-adams-tsvwg-flow-signalling-codepoint-00**

Status of this Memo

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST, AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#)."

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 2, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This document describes the work in progress on Flow State Aware standards activity in the ITU and proposes a new type of control packet to be identified that can alert downstream or upstream nodes on data related to an individual flow.

Table of Contents

1.	Introduction	5
2.	Background	5
3.	Issue	6
4.	Proposal	7
5.	Security considerations	8
6.	Informative References	8

1. Introduction

This contribution calls for a new type of control packet to be identifiable that will allow rate, security or other data related to an individual flow to be passed downstream and alerted to nodes and/ or the receiving end system, where such entities may want to manage flows based on this data, or passed upstream towards a source or Service Provider.

2. Background

Many ideas have converged under the general heading of Flow State Aware QoS control. The published papers on this topic date back to the late 90's, especially the early concepts of Flow Aware Networking from France Telecom [[1](#)] with no signalling and on the fly identification of flows in progress. It had admission control for both elastic and streamed flows, depending on the value of a parameter recording measured spare capacity of a network link. It was characterised by bufferless multiplexing and QoS routing, avoiding congested links and applied per flow.

The direction of the standards work in the ITU has, however, moved beyond these initial concepts. It is developing standards that extend the range of FSA-based QoS controls through explicit signalling with the inclusion of parameters indicating the treatment required. The first phase of standards work reached consent on:

- (a) A new transport service that provides, in effect, admission control in the form of the state assigned to a flow. Flows are allowed to proceed (usually) but the state determines the vulnerability to discard.
- (b) A basis for developing requirements for FSA.

In Q4/13 of the ITU, the second phase of work (on Recommendation Y.2121) includes in-band signalling protocol and procedure for establishing the fastest available rate that could be offered to a flow. Y.2121 also allowed for the option of an out-of-band signalling model as an optional alternative to the in-band approach, although with expected differences in the frequency of rate permissions that could be assigned to an elastic flow. Y.2121 also introduced requirements for FSA QoS controls applied to flow aggregates. The ITU sent the IETF a liaison attaching [Y.2121](#) in January 2008[2], thus Y.2121 is available to the IETF.

To answer a question about the expected benefit of all this activity:

- * One benefit is that Y.2121 provides the basis for developing an Intelligent Edge supporting end systems that do not have in-built FSA capability and merely launch new flows. The signalling path may be limited to Edge to Edge co-ordination to manage flows with statically-assigned preference priorities. Additionally, rate control can be applied per flow to manage any link with high delay or loss which makes TCP inefficient.
- * A specific example is to manage the limited use over a bad link like a satellite. The signalling is terminated at both ends of the bad link and is never seen by the network elsewhere. Current implementations of this have shown 50:1 improvement over satellites using FSA technology.

Network----Proxy--FSA Device---Satellite---FSA device--Proxy-----Network
(No FSA signalling) (No FSA signalling)

The current work in the ITU is in Q5/11. Whilst many of the concepts above can be applied end-to-end, the focus of the ITU work is only on service access scenarios.

The basic aim is the application of FSA QoS in the limited capacity pathway downstream from a service point of access towards an end system or upstream from an end system towards a service point of access. Such a pathway may cross through the core of an IP network, but any FSA signals would only be visible at FSA edge functions and FSA source/ receiver functions at either end (or both ends) of the access pathway.

FSA signals are deleted by FSA Edge functions or FSA source/ receiver functions so that they do not travel beyond the access pathway. This will require all FSA signal initiated at an FSA Edge (acting as the virtual signal source) to be marked so that response signals associated with such flows are deleted at the same Edge (even if this is separated into different sending and receiving physical units).

3. Issue

This contribution would argue that the current ITU work is concerned with scenarios and control protocols that do not affect IETF protocols, because of the restrictions placed on the scope of the work. However, it is recognised that there are significant advantages to extending the applicability of FSA controls to cases where (for example) a content source is anywhere on the internet and user space cannot convey information to or from the network as in encrypted IPv6.

When end systems are participating in FSA signalling, the full richness of FSA capabilities can be invoked. Available rate signalling could be used to determine an optimal file transfer rate. Preference priorities can be invoked dynamically. Feedback can be given to applications that allow them to adjust to make best use of the available capacity.

This contribution captures a more general requirement that follows from this (and is a general requirement in the sense that the utility is not restricted to FSA). It is a requirement for a new type of control packet to be identifiable that would facilitate:

- * The inclusion of new information obtained from (typically, but not necessarily) an edge function, providing rate or security or other data related to an individual flow.
- * The alerting of any of the following that new data is available:
 - o the receiving end system;
 - o downstream nodes with links that are affected by that flow;
 - o the source;
 - o an upstream function (for example a function that creates a copy of this packet and forwards it to a Service Provider function)

Whilst not providing an exhaustive list of the uses of such a new control packet, the following give some example uses, of which (a) and (b) are outside the scope of FSA signalling, but may be of interest to the IETF community.

- (a) A reporting function giving, for example, delivered instantaneous rate information relating to a selected flow, forwarding this data towards a source or an upstream Service Provider function, or both (allowing end-users to have some selection of flows they wish to monitor, as well as Service Providers). As stated above, this may be of use outside a purely FSA context.
- (b) A reporting function giving the instantaneous status that a flow is in competition with one or more flows that have an overriding priority. Again this information may be forwarded towards a source or Service Provider. Again, as stated above, this may be of use outside a purely FSA context.
- (c) FSA in-band signals that are forwarded from or towards any end-point on the Internet.

4. Proposal

There has been complete agreement at the ITU to take Y.2121 through as a consented Recommendation. This same cross-company support is continuing as part of the official and agreed workplan of Q 5/11.

It has been useful that new ideas have been allowed to develop and become captured in this process, for example from ETRI in Korea who were co-editors of Y.2121.

Nevertheless, this contribution invites comment on a proposal that would rapidly help the advancement of the work and where the help of the IETF would be invaluable. That the IETF Transport Area endorses the identification of new IPv6 control packets, supporting the requirement outlined in [section 3](#) of this contribution. Through the IETF working to assign a suitable codepoint for this purpose, our belief is that this would permit per-flow in-band signalling whilst ensuring interoperability with IETF protocols.

As a further clarification, encryption should not obscure the identification of such control packets at intermediate nodes in the network. Also, systems that do not recognize the codepoint should pass the packet forward ignoring the code point.

5. Security considerations

An eavesdropper, which can monitor the packets that correspond to the connection to be attacked could learn the IP addresses and port numbers in use (and also sequence numbers etc.) and easily attack the connection. In such situations, proper authentication mechanisms such as those described in [\[RFC4301\]](#) should be used.

Flow State Aware parameters must be protected against unauthorized modification while in transit. Furthermore, flow State Aware parameter requests must be protected against replay attacks, in conjunction with data integrity protection binding a set of Flow State Aware parameters to a specific flow.

FSA nodes within a domain may authenticate to peer FSA nodes within the domain. FSA nodes communicating as peers across a domain boundary should authenticate with each other.

6. Informative References

- [1] Jim Roberts, Flow Aware Networking for effective quality of service control, IMA Workshop on Scaling, October 1999.
- [2] ITU Liaison Statement to the IETF, Jan 2008, attachment Y.2121.

Authors' Addresses

Dr John Adams
Consultant,
Anagran, Inc.,
24, Keswick Close
Felixstowe,
Suffolk
IP11 9NZ

Phone: +44 1394 272713
email: j.l.adams1@btinternet.com

Professor Jinoo Joung
ETRI(Electronics and Telecommunications Research Institute)
Dept. of Computer Science
Sangmyung University
7 Hongji-dong, Jongro-gu,
Seoul, Korea

Phone: +82-2-2287-5452
email: jjoung@smu.ac.kr

Dr. Jongtae Song
Convergence Network Architecture Research Team
ETRI(Electronics and Telecommunications Research Institute)
161Gajeong-Dong,
Yuseong-Gu,
DAEJEON
305-350, Repulic of KOREA

Phone: +82-42-860-5789
Fax: +82-42-860-6342
email: jsong@etri.re.kr

Phone: +82-2-2287-5452
email: jjoung@smu.ac.kr