Transport Area Working Group J. Adams
Internet-Draft Razoom, Inc
Intended status: Informational
Expires: March 1st, 2010

September 15, 2009

Flow State Aware signalling standardisation, and a proposal for alerting nodes or end-systems on data related to a flow
draft-adams-tsvwg-flow-signaling-identification-00

Status of this Memo

Copyright Notice

Abstract

This document describes the motivation for Flow State Aware signalling and proposes a method of enabling Flow State Aware signalling packets to be identified.

Table of Contents

1. Introduction

This internet draft informs on the benefits of Flow State Aware (FSA) signalling, both as an ehancement of flow classification and, longer-term, as the engine to create new network capabilities.

This internet draft also informs on an existing proposal for Flow State Aware signal packet identification, and calls for further evaluation of such proposals and a recommendation of the best solution by the IETF.

2. Motivation and Early Deployment Scenarios for FSA QoS

Flow State Aware (FSA) technology operates at the per-flow level and provides functions to manage the QoS and monitoring actions that may be required on a specific flow. It distinguishes itself from Deep Packet Inspection and WAN Optimisation functions in that it supports FSA Transport Classes, i.e.

- the Maximim Rate (MR) class. Allowing a flow to start with no admission control and protecting other flows from any disturbance to their QoS experience through "remembering" and directing discard actions on this flow (until conditions change).
- the Guaranteed Rate (GR) class. A class of flows that are protected from discard actions under normal operating conditions.
- the Available Rate (AR) class. A class of flows that are supported by signalling tests of the current fastest end-to-end available rate
- the Variable Rate (VR) class. A class of flows with a minimum rate (in the same sense as an MR-level of guarantee) plus an additional available-rate "top-up".

FSA associates flow state with every flow (potentially) in any aggregate of flows and/or the aggregate itself. It uses this state information to perform the following:

2.1 QoS actions on any flow. In particular

- 2.1.1 Admission-based actions - allowing a flow to start but promoting it to QoS guarantees when resources are available.
- 2.1.2 Preference-based actions - controlling how soon a flow is promoted to QoS guarantees
- 2.1.3 Focused packet discards. "Remembering" flows whose packets are subject to some level of discard.
- 2.1.4 Guaranteed QoS. "Remembering" flows whose packets are not to be subjected to any discard under normal operating conditions.
- 2.1.5 Additional QoS actions if end-to-end Flow State Aware

signalling is operating. In particular:

2.1.5.1 Available Rate assignment of capacity. Based

on the signalling yielding the fastest end-to-end available rate.

2.2 Measurement actions on any flow or any flow aggregate. In particular:
2.2.1 Per-flow or per-aggregate flow policing, as required.
2.2.2 Per-flow or per-aggregate flow alerts, based on a measured parameter, and where the alert may trigger:
2.2.2.1 A change of state (per flow, per aggregate)
2.2.2.2 A change of QoS action
2.2.2.3 A change of monitoring action
2.2.3 Per-flow or per-aggregate flow alerts, based on a packet classifier parameter, and where the alert may trigger:
2.2.3.1 A change of state (per flow, per aggregate)
2.2.3.2 A change of QoS action
2.2.3.3 A change of monitoring action

This internet draft focuses on signalling development associated with flow state (through which the above action 2.1 and 2.2 are then managed). First it may be noted that not all of the actions above actually require any signalling. For example, 2.1.1 can allow a flow to start immediately, regardless of whether any signalling had been received related to that flow. But it would still require some actions related to flow classification so that the flow was properly classified as "needing QoS guarantees". Similarly, 2.1.2 could be done entirely with the aid of flow classification and no signalling.

Furthermore performing flow classification on high speed (10 Gbit/s) input links is possible (and is already being done) so no claim can be made that signalling is a necessary addition to make the classifiers operate at 10 Gbit/s or higher.

Essentially the argument for signalling is twofold:

A. Extending the information about what we know about a flow or what we want done to a flow.

B. Enabling path tests, such as the Available Rate path test.

In the next section of this internet draft, some future possibilities are listed that could extend what is being signalled and which add to the argument that an FSA signalling capability can do more than a classifier ever could do on its own.

It is clear that a richer level of information could be associated with a flow if flow-based signalling accompanies that flow. This would be true at any level of classifier processing power. By "accompanied" we could mean either in-band or path-coupled, but in-band seems to be the choice that causes the least processing hit, and opens up new possibilities like high-speed transfers based on the fastest available

rate.

An investigation into the use of GIST [1] as the signal transport for FSA yielded the following observations:

- FSA is a signal from one node to all downstream FSA nodes in the path of the flow. GIST is a signal between two nodes.
- GIST utilises UDP encapsulation. FSA uses the same type of encapsulation as is used for the data packets of the flow.

3. Architectural considerations

Reference [2] defines a Flow State Aware Signalling Edge Function. This is a function that provides the origin and/ or termination of the Flow State Aware end-to-end signalling path, and participates in requests and responses on behalf of a user application or management action. It may be located, for example, in a user end-system or at a network edge node where it serves as the signalling end-point of multiple users and associated applications. Alternatively, it may be located at an Aggregation End-point where it supports the signalling requirements of flow aggregates.

Figure 1 shows an overview of FSA functions

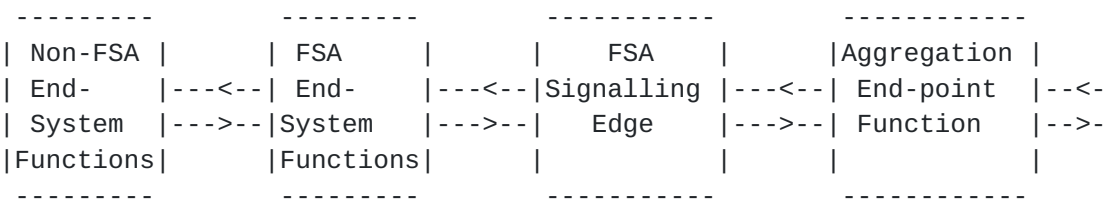

Figure 1: FSA Functional architecture

In Figure 1 an FSA End-System is shown with QoS and/ or monitoring capabilities but with no Signalling Edge Function. Of course, it is envisaged that End-Systems may evolve so that they may have:

- no FSA capabilities (as is the case today)
- limited FSA capabilities, for example at an End-System Hub
- Evolved FSA capabilities, including FSA signalling

The Aggregation End-point is a function which attaches or deletes the common Flow Aggregate Identifiers to ensure commencement of/ cessation of common routing and QoS treatment of packets. This end-point also initiates/ terminates in-band signalling to control Flow State information retained for treatment of the flow aggregate.

The implication of FSA signals terminating on a FSA Signalling Edge

Function (including Aggregation End-point functions) is that the End

Systems that are receiving or sending data do not receive extra packets (i.e. the signalling packets) that may otherwise cause data corruption or other problems relating to TCP or UDP transmission. Instead, the only packets that are forwarded further downstream from a FSA signalling Edge Function are the user-data packets.

At the far right-hand side of Figure 1 is shown both-way multiple packet flows going either downstream towards an End-System via a FSA Signalling Edge Function, or upstream towards a far-end Aggregation End Point function and far-end FSA Signalling Edge Function.

The FSA functions shown in Figure 1 may include layer 2 or layer 3 packet forwarding capability. Thus an example of a FSA function in Figure 1 could be:

- an Ethernet frame-forwarding function that is "IP aware"
- a router

4. Evolution considerations

More insight can be gained by looking at an early FSA deployment scenario and considering an evolution of this that includes FSA signalling and what is then added to the value.

(a) First stage: Stand-alone Flow State Aware functions managing the access bandwidth, QoS experience and monitoring actions associated with broadband service.

Without signalling, such a solution could focus on 2.1.1, 2.1.2, 2.1.3, 2.1.4, and 2.2.

(b) Second stage of evolution: With signalling on some flows only. Allows the support all of 2.1.1, 2.1.2, 2.1.3, 2.1.4, 2.2, plus 2.1.5 in the following limited sense:

- some flows are being controlled end-to-end if distributed content services utilise FSA managed access bandwidth and the content sources can transmit towards the users just using an access network.
- some flows are not associated with any signalling and are managed as in evolution stage (a)
- signalling (where used) increases the information available per flow or simplifies and extends the setting of preference priority or peak rate without a large number of classifier rules.

(c) Third stage: (Could also be a second stage without signalling) Access FSA functions are interconnected by fixed-bandwidth trunks and:

- Without signalling, such a solution could focus on 2.1.1, 2.1.2, 2.1.3, 2.1.4, and 2.2. (but with limited information, e.g. on

preference priority or peak rates).

- With signalling, extends the number of sources that can signal flow information towards an FSA QoS management function and utilise signal-supported services such as AR or VR.

(d) Forth stage of evolution: Signalling on most flows, where some sources are anywhere on the internet. Allows access to FSA signal-supported QoS and monitoring capabilities for content providers who have not built out content distribution to match scenario (b) or (c) above.

(e) Fifth stage of evolution: extensions to the signalling to incorporate further new capabilities (see next section).

Note that priority Assignment allows the sender to request the priority that the sender needs for each flow. For Emergency Services, Military use, and corporate use the ability for a user to assign different rate priorities to each flow is critical. If the flow may be available rate, the rate would then be assigned based on the priority. For maximum rate flows, the priority could allow critical flows the rate they need. This capability must be coupled with authorization security for the user to use this priority.

Authorization security for a sender to use a given priority can be verified with a AAA server but must be assured for each flow to insure security. Per packet assurance is considered infeasible and/or expensive, thus per flow security verification is the logical solution. This security is critical when priorities are used but also will greatly improve identifying who is sending malware or creating a DoS attack.

Note also that, being assigned an available rate by the network (through a path test via in-band signalling) ensures that TCP type flows from a FSA signalling-capable source end system can:
- stream at the maximum rate which is fair in the network
- flow at much higher rates than TCP can flow today
and for large distances greatly reduces the time to deliver a file.

5. FSA Signalling Standards Development

The direction of the standards work in the ITU is towards 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 work reached consent on:

(a) A new transport service that provides, in effect, admission control in the form of the state and packet deletion actions associated with a flow.

(b) Requirements for FSA in a Next Generation Network.

These requirements are captured in ITU-T Recommendation Y.2121 [2]. The ITU sent the IETF a liaison attaching Y.2121 in January 2008[3], thus **Y.2121 is available to the IETF.**

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

More explicitly, the scope of the current standards initiative in the ITU is on 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).

Thinking further about the longer-term, there could be extensions to what is conveyed in signalling packets to unlock new capabilities. So far, such extensions have not been addressed, but could include:

(i) in-band monitor on/ monitor off commands associated with monitoring a group of packets of a flow (or aggetegate) that trail behind the first such command and stop at the second such command.
(ii) in-band monitor-and-change-state-if commands associated with a group of packets that trail behind the first such command and change either a QoS state or a monitoring action if a given event is observed in the group of packets up until a second such command.

6. Signal Packet Identification

This internet draft is raising the solution to an issue already captured in [2]. It is a requirement for a new type of in-band control packet to be identifiable to any FSA node that would facilitate:

* The inclusion of new information obtained from (typically, but not necessarily) a source function, providing rate or or other data related to an individual flow or flow aggregate.
* The alerting of any of the following that new data is available:
    - the receiving end system;
    - downstream nodes with links that are affected by that flow;
    - the source.

So far, Q5/11 has considered a number of alternative solutions and has

proposed a candidate solution to signal packet identification. This

solution is that all signal packets should be marked with DiffServ=an EXP/LU DSCP set to "QoSSignal"

This is viewed, architecturally, in an analagous way to an "escape" key. The DSCP = "QoSSignal" creates an escape into a new set of QoS options that are conveyed in another part of the packet called the "QoS Structure".

Because packets marked with an EXP/LU DSCP set to "QoSSignal" originate and terminate on FSA Signalling Edge Functions, it is assumed that this chosen solution does not change any internet protocol as currently defined and operating on any network node or End-System node.

However, considering the evolution scenarios described in section 4 above, the proposed solution for FSA signal packet marking may need further study if the scenarios of exploiting the FSA signal-enhanced flow information and FSA QoS actions / monitoring actions extend towards more general use between any two points on the internet.

7. 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 try to attack the connection by generating signalling packets.

The protection of signalling packets with an authorization field is recommended, following the principles described in [2] Annex A.

Furthermore, FSA monitoring functions may be used to determine the frequency of signal packet arrivals and whether this frequency exceeds some threshold that would signify an alert and a possible DOS attack.

8. IANA considerations

This document has no actions for IANA with the existing proposal for signal packet identification. But this would not be the case for some alternative proposals.

9. Proposal

This internet draft indicates the background to a proposal to identify a new in-band signalling packet and indicates the current candidate solution that has been assumed in the ITU-T for the explicit limited scope described above in section 5. Members of the IETF are asked to consider this proposal and endorse the current solution for the scope currently under consideration and determine the best solution for extensions of this scope.

10. Informative References

[1] draft-ietf-nsis-ntlp-20 (2009), GIST: General Internet Signalling Transport
[2] ITU-T Recommendation Y.2121 (2008), Requirements for the support of flow state aware transport technology in an NGN
[3] ITU Liaison Statement to the IETF, Jan 2008, attachment Y.2121.

Authors' Address

Dr John Adams
CTO,
Razoom, Inc.,
24, Keswick Close
Felixstowe,
Suffolk
IP11 9NZ

Phone: +44 1394 272713
email: john@razoom.com