

Ingress Filtering Objectives - DRAFT

Soheil Samii (soheil dot samii at gm.com)

Yong Kim (ybkim at broadcom com)

Helge Zinner (Helge dot Zinner at de.bosch.com)

Recap Reminder – Reference Analysis

Ingress Policing from Markus Jochim @ GM

<http://www.ieee802.org/1/files/public/docs2013/new-tsn-jochim-ingress-policing-1113-v1.pdf>

	Per Stream <i>(= Potentially higher number of filters per port)</i>	Per Class <i>(= Small number of filters per port)</i>
Threshold Enforcing	<ul style="list-style-type: none"> A faulty stream sent by a faulty talker is not "silenced". Other streams from faulty / fault free talkers not affected. 	<ul style="list-style-type: none"> A faulty stream sent by a faulty talker is not "silenced". Non-faulty streams sent by faulty talkers can become faulty. A fault free stream sent by a fault free talker becomes faulty. (Fault propagation. Fault not contained)
Blocking	<ul style="list-style-type: none"> A faulty stream sent by faulty talker is "silenced". Non-faulty streams sent by faulty talker are not necessarily silenced. 	<ul style="list-style-type: none"> If a talker exceeds its configured bandwidth limit, the faulty talker is "silenced". In presence of a moderate babbler, a fault free stream sent by a fault free talker can become faulty. (Fault propagation. Fault not contained) Faulty streams sent by a faulty talker are not necessarily silenced.

Slide 24

	Per Stream <i>(= Potentially higher number of filters per port)</i>	Per Class <i>(= Small number of filters per port)</i>
Threshold Enforcing	<p>Not a preferred solution.</p> <p><u>Reason:</u> "Per Stream X Blocking" and "Per Stream X Thr. Enforcing" both require the same HW overhead associated with Per Stream IPFs, but "Per Stream X Blocking" exhibits preferable behavior.</p>	<p>Not a preferred solution.</p> <p><u>Reason:</u> Fault containment related problem. (Whether or not this is acceptable heavily depends on the specific application.)</p>
Blocking	<p>#1 Preferred solution !!!</p> <p>Very effectively addresses the Babbling Idiot problem.</p>	<p>#2 Potential alternative to "Per Stream X Blocking"</p> <p><u>Reason:</u> Effectively addresses Babbling Idiot problem in many cases with limited HW resources. Potential drawback: Depending on the application, the Moderate Babbler may need to be considered.</p>

Slide 26

Recap Reminder – Control & Mgmt

Slides 28-31 of Ingress Policing from Markus Jochim @ GM

- Ingress filtering function should be per-stream. Per-aggregated PCP acceptable but not preferred.
- Ingress rate measurement, by definition has latency. Minimize rate measurement latency.
 - Means of measurement interval configuration and latency TBD.
- Desired management actions include block ‘offending’ stream at ingress, filter-to-profile stream at ingress, or disable port.

What exists in Q

- The standard does not mandate Ingress Filtering (use of word *may*).
- The standard only mandates that decisions by the Ingress Filtering shall be port-based (per reception port) and shall not consider incoming or egressing frames of other bridge ports.
- The standard defines a set of flow classification rules that may be used by the flow metering mechanism (which itself may or may not exist in the bridge). There is no clear boundaries in the scope of Ingress Filtering (e.g., does not mandate per-stream or per-class metering) – it is essentially to be defined by bridge vendors.
- The drop_eligible parameter (DEI bit of the VLAN tag) of a frame can be changed by the ingress filter. Upon egress of a frame, the bridge may discard the frame. There is no requirement on certain action upon the value of the DEI bit.
- Referenced filtering algorithm in MEF 10.2 (Metro Ethernet Forum). Section 7.11.1 in MEF 10.2 defines the algorithm, but the presentation is not clear in the interpretation of the different introduced variables (e.g., B_c and B_e).

Refinements to Q

DETECTION

- If flow metering is supported by a bridge, then the following requirement shall be satisfied: Flow metering and monitoring shall be supported and frame classification shall be done based on a configurable set of parameters consistent with AVB/TSN streams, e.g. as noted in 802.1CB Annex C.
- Detection latency is to be minimized and less than <td> msec, and specify filtering and smoothing function, e.g. exponential rate gauge.
- Observation interval for the bandwidth measurement shall be in multiples of the class measurement interval
- The rate to measure by the ingress policing filter shall be configurable and the numerical granularity shall be in the same order as the one for stream reservation.

Addition to Q

ACTION

- Upon detection of excess bandwidth usage, the allowed action shall be configurable among the following alternatives:
- Block all future frames on the ingress port (i.e., block all incoming frames on the port where excess bandwidth was detected). This is sufficient in cases where the ingress port is connected to an end-station (one detected violation renders all data from the end-station unreliable).
- Enforce stream threshold (the MEF 10.2 algorithm essentially captures this property, but 802.1Q does not specify the usage of the algorithm on a per-stream basis).
- Block all future ingress frames of the stream. This is needed to properly handle errors (exceeding stream bandwidth) due to faults in a bridge, in cases where the ingress port is connected to a bridge