Flow ID Options

Caitlin.bestler@neterion.com

Neterion®

Three Different Approaches to Flow IDs

- **My perceptions on the different assumptions:**
  - Destination Differentiated
    - As in au-bestler-endstationrps-0708-05.pdf
    - Destination/Priority determine set.
    - Other L2-L4 headers hash within set for multi-pathing.
  - Non-Destination-Differentiated
    - Similar to above, but limited differentiation by Destination.
  - Opaque
    - Minimize standardized restrictions on End Station use of Flows so that the Flow ID can be used to facilitate identifying the true source of congestion in the upper layers.
- **Use of a salt to randomize the resulting Flow ID should work with all of these approaches.**

Neterion®

Destination Differentiated

- **Flow Queue Set is functionally derived from:**
  - FID of VLAN ID
  - Destination Address
  - Priority
- **Each Queue Set has *N* Flow Queues within it to allow for network multi-pathing.**
  - A default algorithm for selecting path from the Flow ID would be suggested, but not mandated. For link aggregation it may be elevated to a SHOULD.
- **Other L2-L4 headers determine hash within the Flow Queue Set**
  - Any given L4 flow should go to a single Flow ID.
  - Language from Link Aggregation can be referenced.
- **L4 Flow to Flow ID mapping is locked when Rate Limited.**
- **Optimized for CPs near the destination.**

Neterion®

Not Destination Differentiated

- **One weakness of the Destination Differentiation is the impact on CPs on Core Bridges.**
- **If there are many active destinations then a Core Bridge would be forced to send numerous CNMs before truly slowing traffic.**
- **Limiting the total number of Flow IDs that an End Station could use, regardless of Destination, would prevent this.**
- **This does result in undesirable Fate Sharing, but not having it could result in Priority-based Flow Control being invoked**
  - which is worse Fate Sharing.
- **The Global strategy might also lock Flow ID mapping totally.**

- **This is clearly the best algorithm for traffic patterns where most CNMs are generated by Core Bridges.**

Neterion®

Opaque

- **When congestion is being caused by relatively few flows then the method of choosing Flow ID is largely irrelevant to QCN performance.**
  - Statistically, the “Elephants” will be the one selected for CNM reduction, and they will be quickly reduced.
  - Therefore it would be enough to state that End Stations and/or applications should not artificially increase the number of Flow IDs by varying header parameters for what is a single application flow.
- **If unrestricted, End Stations would be better able to use the Flow ID to collaborate with upper layers (such as the network stack) on identifying the source of congestion.**
  - Working with the network stack can be more effective at eliminating congestion than L2 working alone.

Neterion®

Opaque Hybrid?

- A Flow ID could encode both a Rate Limiter ID (as in the first two options) and an Opaque l4-source-flow identification.
- But such a Flow ID probably would need to be larger than 16 bits.
- There may be implementation benefits to keeping the Cn-Tag structurally identical to a Q-Tag.
- Do the benefits of having both outweigh the cost of the larger size and extra format?

Neterion®

Metrics?

- **How many different flows contribute to congestion?**
  - When the network is congested, what % of the flows cause what % of the traffic.
  - The absence of scenarios where a large number of flows cause congestion would suggest not standardizing what a Flow ID was.
- **What is core/edge distribution of congestion?**
  - If edge congestion is more common then Destination differentiation will minimize penalizing innocent flows.
  - If core congestion is more common then Destination differentiation will result in Priority-based Flow Control, which penalizes even more innocent flows.

Neterion®