

PCN Status

Congestion and Pre-congestion Notification

IETF 70 - Vancouver
Monday 2007-12-03

Goals and Milestones (1)

- Nov 2007: Submit "Flow Admission and Termination Architecture within a Diffserv Domain" (Informational)
- Nov 2007: Submit "Survey of Encoding and Transport Choices of (Pre-)Congestion Information within a Diffserv Domain" (Informational)
- Mar 2008: Submit "(Pre-)Congestion Detection within a Diffserv Domain" (Proposed)
- Mar 2008: Submit "Requirements for Signaling of (Pre-)Congestion Information from Egress to Ingress in a Diffserv Domain" (Informational)
- Jul 2008: Submit "Encoding and Transport of (Pre-)Congestion Information from within a Diffserv Domain to the Egress" (Proposed)

Goals and Milestones (2)

- Nov 2008: Submit “Encoding and Transport of (Pre-) Congestion Information from the Domain Egress to the Ingress” (Proposed)
- Jul 2008: Submit “Suggested Flow Admission and Termination Boundary Mechanisms” (Informational)

Where are we?

- Making progress on our first milestone document, but not (yet) ready for WG last call. This milestone gates all of our other milestones.
- Need a WG document for the second milestone document; one candidate draft. This milestone gates our fifth milestone.
- Three proposals for the interior marking, egress metering, and admissions/termination algorithms. Selection/convergence amongst these would allow us to move forward on the rest of our milestones.

Why we are here

- To develop a simple, robust, and scalable MBAC mechanism to complement Diffserv. As such, PCN is designed for use in deployment scenarios where MBAC can be expected to work, and not in others.
- Not to invent a new BW reservation protocol (IETF already has a few of those).
- As such, minimizing admissions blocking/termination probability are not explicit goals.
- PCN should be easy to deploy, while solving a real problem for operators.
- PCN does not need to solve all corner cases.
- PCN is part of a toolkit, not a complete QoS solution. In different deployment scenarios PCN may be used along with other protocols. We don't need to develop overlapping functionality.

Moving forward (1)

- Mailing list discussion has been focused on corner cases. Need to weigh the value add of addressing them versus the cost in additional complexity.
- The PCN charter places constraints on the scope of our work. These constraints are deliberate. We can be mindful of the implications of our design decisions on possible future work, but we are not going to address topics that are out of scope of the charter.
- We need to resolve open issues with the architecture draft, and move the encoding draft forward to working group status and then to working group last call.
- When resolving open issues, err on the side of simplicity.

Moving forward (2)

- Remember: we are working with aggregate information. Even if we have per-flow state at the egress node in some deployment scenarios (to associate packets with ingress-egress aggregates, for instance), that shouldn't influence the behavior of the other PCN mechanisms.
- We need a single pre-congestion detection and marking mechanism in the interior routers, to ease deployment and interoperability.
- We want a single mechanism for metering pre-congestion information at the egress node. Innovation at the edge is easier than at the core, so we can imagine our mechanism here evolving in the future.

Moving forward (3)

- We need to be mindful of the Diffserv and ECN compatibility issues. The term “inelastic flows” can be and has been interpreted too narrowly.
- Need for speed: We are not reinventing email here, nor solving the routing scalability problem. We should be able to hit our milestone dates. It is important to make progress, because previous IETF QoS efforts have dragged on far longer than planned.

Operator Input

- The working group would like feedback from the operator community!
- Previous QoS efforts at IETF have suffered from the lack of early operator input. We want to avoid this problem.
- The chairs contacted a few operators already. Early feedback:
 - The architecture draft is moving in the right direction
 - PCN needs to be frugal in the use of IP header codepoints
 - PCN needs to follow the KISS principle
 - Hurry up with the current work so that we can look at the multi-domain problem.

Vendor Input

- The working group would like feedback from the router vendors!
- We would like input on how mechanisms (e.g., rate meters, DSCP remarking, ECN) already deployed in routers could be utilized by PCN.
- We would like feedback on how easy it would be to implement and deploy possibly new PCN mechanisms.