Draft-pcn-architecture-03

Philip Eardley

11 March 2008, IETF-71, Philadelphia

Since last ietf

- Revised, dealt with all comments
- No major changes
- Question posed on list of whether ready for WGLC

Lars's comments

- Draft sets out too many choices & options, want convergence on answers for:
- Encoding (coming next)
- Marking behaviours on PCN-interior-nodes
- Admission control & termination behaviours on PCN-boundary-nodes

Threshold Marking behaviour

- Aim: if PCN-traffic > configured PCN-lower-rate, then mark all pkts
- Pkt classify (see encoding)
- Meter behaviour MUST be (implementation unconstrained):
- a token bucket, which
- tokens are added at the PCN-lower-rate, to a maximum TB.size
- tokens are removed equal to the size of the PCN-packet, to a minimum TB.size=0
- if TB.fill < TB.threshold, then:
- Pkt mark (threshold-marking) (see encoding)

Excess rate Marking behaviour

- Aim: if rate of PCN-traffic > configured PCN-upper-rate, then mark pkts at rate {PCN-traffic-rate - PCN-upper-rate}
- Pkt classify (see encoding)
- Meter behaviour MUST be (implementation unconstrained):
- a token bucket, which
- tokens are added at the PCN-upper-rate, to a maximum TB.size
- tokens are removed equal to the size of the PCN-packet, to a minimum TB.size=0
- if TB.fill = 0, then
- Pkt mark (excess-rate-marking) (see encoding)
- Need to double check that Single Marking works ok

Admission ctrl, boundary-node

- Charter: Info doc or docs
- Aim: let's concentrate on getting one model done now, then do others
- Proposed model:
- egress makes admission decision & signals to the ingress "admit" or "block" on a per flow basis.
- egress bases this decision on the level of congestion on the ingress-egress-aggregate (eg EWMA)
  - Egress is configured with acceptable level of congestion (ie when switches from admit vs block decision).
- if no traffic from this ingress recently, then the "level of congestion" is based simply on whether this first pkt (which is the flow adm request msg) is marked or not (assumes signal approx follows data path)
- ingress abides by egress admission decision, unless eg management system tells it to over-ride.

Admission ctrl, boundary-node

- This excludes for now:
- Probing mechanisms that ensure signal follows data path even in ecmp environment
  - Because No agreement on a satisfactory way to solve this
- approach where: when egress sees a mark then signals to the ingress " block all requests”; when egress sees no marks for x secs the signals to egress “admit al flows”
  - Because we should do better understood approach first

Termination ctrl, boundary-node

- Charter: Info doc or docs
- Aim: let's concentrate on getting one model done now, then do others
- Proposed model:
- egress measures "sustainable rate" from specific ingress & reports to ingress
- Ingress measures rate sent towards this egress & calculates amount to terminate

Termination ctrl, boundary-node

- This excludes for now:
- (probabilistic) Terminating marked flows
  - This is simpler (no rate measurement) but other method (ie measuring rates) can be faster & less signalling & works if only “single marking” is done & works better if flows have different bitrates.
- Affected marking approach
  - Quantification of benefit is to be done
  - Not general (needs codepoint)