

draft-ietf-pcn-marking-
behaviour-01
(Standards track)

Philip Eardley

Minneapolis, IETF-73, 18 Nov 08

Since last ietf

- 2 new versions (WG -00, now at -01)
- Reflects consensus decisions at Philadelphia, re-confirmed at Dublin
- Traffic conditioning re-written
 - now a MAY;
 - ‘downgrading’ removed
 - “use flow termination”
- Competing-non-PCN-packets
 - now MAY be metered;
 - advice: don’t have any)
- I believe ready for WG last call

3 Discussion areas

- Topics raised recently on the mailing list
 - Not sure whether they're open issues

3 Discussion areas

1. Allow the (forthcoming) draft-satch-pcn-ST-marking-00
 - To be discussed later in meeting
2. For excess traffic marking, have a parameter N , such that only every N th pkt is marked (metering behaviour unchanged)
 - Relevant to edge behaviour where you terminate a flow with a marked pkt
 - Studies (eg draft-menth-pcn-emft) have shown that you can achieve the same overall behaviour by instead terminating $1/N$ flows that have a marked pkt
3. Make preferential dropping optional (instead of a **SHOULD**)
 - Draft says: preferentially drop pkts that are excess-traffic-marked
 - Have an option to allow “random dropping” (eg tail drop)

History of preferential dropping

- Philadelphia: choice between
 1. MUST preferentially drop pkts that are already excess-traffic-marked OR
 2. MUST drop independent of marking
 3. (preferentially drop pkts that are not excess-traffic-marked)
- Consensus for option 1, as a SHOULD
- Summary of discussion (impact on different edge approaches for doing flow termination):
 - Next slide, (I believe) this is still a fair summary
- New suggestion: have an Option, ie operator can choose whether to do 1 or 2 (or 3?)

Summary of Philadelphia discussion

Edge beh => ----- (below) drop pref	CL/SM	Marked flow termination
Prefer drop ExM pkts	Best for CL/SM	Works ok
'random' drop	breaks completely in some scenarios (terminates far too much)	works OK (slightly slower than above)
Prefer drop pkts not ExM	breaks completely in some scenarios (terminates all flows)	best for MFT (much faster than above)
	CL/SM: Can terminate fast ('1 shot')	MFT: Simpler (don't measure rates), but accept trade-off that won't terminate so fast

Configuration Option for dropping behaviour

- Michael's suggestion: have a configuration Option, ie operator can choose what sort of preferential dropping
- Pros:
 - All edge behaviours can choose their optimal preferential dropping behaviour
- Cons:
 - Extra complexity to standardise, implement & configure
- Question:
 - What would you say in the stds doc?
 - Eg what's the default behaviour? (tail drop?) (but breaks...)
- Suggestion:
 - Keep the SHOULD as now (prefer drop ExM)
 - Add a note about why you would do something different (eg tail drop is smallest implementation step from today's routers, but beware it has these trade-offs blah depending on your edge behaviour wibble)

WGLC

- I believe this doc is ready for WG last call
 - With the ‘Note’ suggested on previous slide; and a few nits
 - I’m not sure that 100% consensus will be possible, even with infinite discussion