

# draft-ietf-pcp-port-set-03

Qiong Sun <sunqiong@ctbri.com.cn>

Mohamed Boucadair <mohamed.boucadair@orange.com>

Senthil Sivakumar <ssenthil@cisco.com>

Cathy Zhou <cathy.zhou@huawei.com>

Tina Tsou <tina.tsou.zouting@huawei.com>

Simon Perreault <simon.perreault@viagenie.ca>

IETF 88, Berlin

PCP meeting

2013-11-08

# Status

- WGLC'ed
  - Comments received, issues raised

# Non idempotence (This is important. Pay attention.)

- Example: two requests:
  - a) Map internal ports 1-2
  - b) Map internal ports 2-3
- First request to be processed by the server “wins”.
- Second request treated as a refresh of the first one.
- However, the client is in control: it “knows” that it sent overlapping requests.
- Non-overlapping requests are idempotent.
- No security attack. (?)
- Proposal: describe this in the next draft revision.

# Simplicity

- “The claim that a port set makes code simpler needs an explanation to justify it. E.g., if an app really needs 16 ports, and it fails to get 16 back in its first request, it still has to handle that and say make multiple separate requests. So if it already has to have code to deal with multiple separate requests, why is it simpler to add port set functionality to and implement both? This seems counter-intuitive to a reader.”
- Answer:
  - Often, apps that ask for N ports don't really need N ports. Ideally they would like to get N ports, but they can live with  $<N$  ports. No need for single-port fallback in that case.
  - Asking multiple times for small allocations is a little bit evil, like opening multiple TCP connections in an attempts to trick an SFQ-based QoS into allocating more bandwidth.
- Proposal: document this in the next revision.

# Nonce reuse

- “Let's say internal port 100 and 110 both have the same mapping nonce, and internal port 105 has a different mapping nonce (e.g., from a different PCP client on the same host). If the first PCP client sends a refresh for port 100, Port Set Size = 11, what would the effect be at the PCP server?”
- Answer:
  - The mapping at port 100 would be refreshed.
  - The mapping at port 110 would not be refreshed, but its remaining lifetime would be returned by the server, as described in draft-cheshire-pcp-unsupp-family. The response's nonce would be copied from the request's nonce.
- Proposal: add this to the examples section.

# PORT\_SET in single-port mapping responses

- Example:
  - Client asks for port set size = 2.
  - Server does support PORT\_SET, but chooses to allocate port set size = 1.
  - Problem: the client believes that the server does not support the PORT\_SET attribute.
- Proposal: when the request contains PORT\_SET, always return PORT\_SET, even for port set size = 1.

# Normative references to non-RFCs

- Allow the draft to progress faster by cleaning up normative references to non-RFCs.
- Proposal:
  - Remove the reference to draft-boucadair-pcp-failure.
  - Remove the reference to draft-cheshire-pcp-unsupp-family. Move nonce text to draft-cheshire.
  - Leave section “5.2 Stateless Mapping Discovery” intact. It applies to generic stateless NAT, not draft-tsou-stateless-nat44.

# Next steps

- Publish new revision
- Ship to IESG