PING and TIMESTAMP for MASQUE

IETF 114, July 2022
Ben Schwartz
Slides v00

Overview

- Goals
  - Enable improved performance and reliability for MASQUE
  - Develop best practices for Capsule Protocol extensions
- Changes since -01
  - Redesigned to match final HTTP/3 Datagrams draft (mostly)
  - Added the TIMESTAMP extension
- Applicability
  - Works with any Capsule Protocol request that uses Context IDs
    - i.e. CONNECT-UDP and CONNECT-IP

PING

What is a PING Datagram?

- PING is a Capsule Protocol Extension
- Pings are sent between the HTTP client and origin
  - Can be sent in either direction
  - Opaque to intermediaries

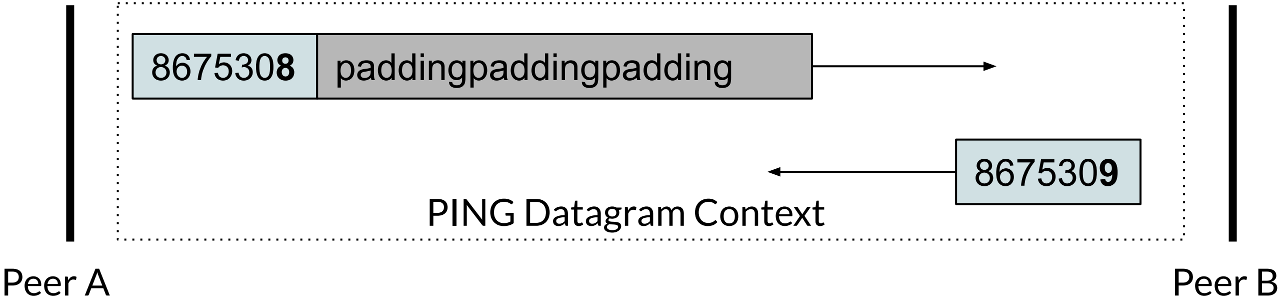

How do I enable PING?

- PING contexts are registered statically in the request/response.
- New HTTP header field: DG-Ping: <Context ID #>
  - There is only one PING context per request, and it is selected by the client.
  - The server echoes the header to confirm its support for PING.

Why use PING?

- Enables DPLPMTUD for the HTTP Datagram MTU
  - HTTP supports intermediaries, so endpoints can't determine the Datagram path characteristics from measurements at a lower layer
- Also good for measuring RTT and loss rates in any application (e.g. for performance debugging)

TIMESTAMP

What is a TIMESTAMP datagram?

- Attaches a transmission timestamp to any Datagram
  - e.g. UDP, IP, PING
- Uses NTP's 4-byte or 8-byte time encoding

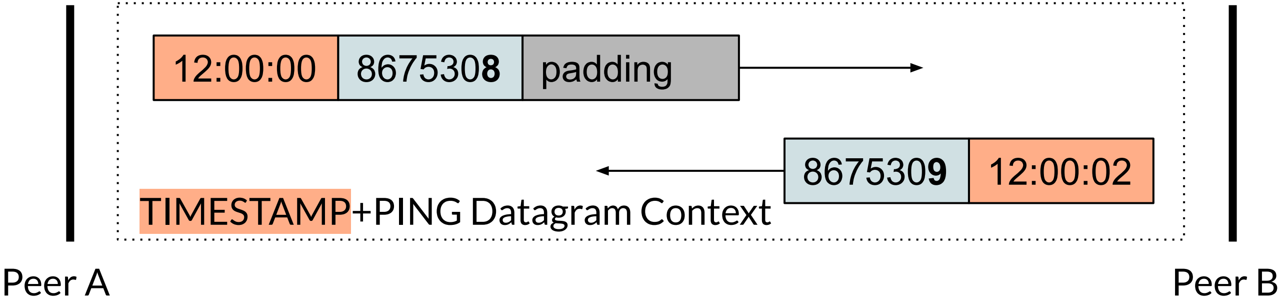

How do I enable TIMESTAMP?

- TIMESTAMP capability is negotiated by a header: DG-Timestamp: ?1
- TIMESTAMP contexts are registered dynamically using new Capsule types
- Each TIMESTAMP context corresponds to another context ID, and wraps its payload.

```
REGISTER_TIMESTAMP_CONTEXT Capsule
{
  Context ID (i)
  Inner Context ID (i)
  Short Format (1)
}
ACK_TIMESTAMP_CONTEXT Capsule {
  Context ID (i)
  Error Code (i)
}
CLOSE_TIMESTAMP_CONTEXT Capsule {
  Context ID (i)
}
```

Why use TIMESTAMP?

- Improved congestion control for proxying
  - Allows separation of congestion on the client-proxy and proxy-target legs.
  - Enables improved interaction between the client-proxy and end-to-end congestion controllers.
- Debugging latency issues
  - “Which queue is filling up?”
- Jitter reduction in highly interactive applications
  - e.g. gaming, robotics

Interesting questions

- Should we agree on a uniform prefix like “DG-” for header fields that negotiate datagram capabilities?
- If the other party allocates TIMESTAMP context ID X with Inner Context ID Y, should I always send on X instead of Y? Do we want a way to request timestamps on 1% of packets?
- Can I add timestamps to an Inner Context ID that was allocated by the other peer? What are the rules about closing contexts in complicated situations like this?

Status

- Seeking WG adoption in MASQUE
- May help to address outstanding issues related to MTU in CONNECT-IP