

MASQUE

HTTP Datagram Priorities

[draft-pardue-masque-dgram-priority](#)

IETF 114~~2~~² – Virtual ~~San Francisco~~Madrid –
2021-07~~10~~¹¹

Lucas Pardue – lucaspardue.24.7@gmail.com

Background

RFC 9000 says stream multiplexing can have a significant effect on application performance if resources allocated to streams are correctly prioritized.

QUIC does not provide a mechanism for exchanging prioritization information. Instead, it relies on receiving priority information from the application.

draft-ietf-quick-datagram (now **post-WGLC**) defines QUIC DATAGRAM frames. There is no multiplexing identifier but presumably DATAGRAM frames can be muxed with STREAM frames.

HTTP/3 doesn't define priority signals, punts to **draft-ietf-httpbis-priority** (now **post-WGLC**).

draft-ietf-masque-h3-datagram defines datagram multiplexing but doesn't say anything about priority.

Bouncing around

- HTTP Datagram Issue [#46](#) – The spec should discuss how h3-datagram works (or does not) with priority.
 - **Closed** with a PR that says:

Prioritization of HTTP/3 datagrams is not defined in this document. Future extensions MAY define how to prioritize datagrams, and MAY define signaling to allow endpoints to communicate their prioritization preferences.

- HTTP Priorities [#1550](#) – How are DATAGRAM frames prioritized?
 - **Closed** with a PR that says:

The priority scheme defined by this document considers only the prioritization of HTTP messages and tunnels ... Where HTTP extensions change stream behavior or define new data carriage mechanisms, they MAY also define how this priority scheme can be applied.



Problems for protocol and packet prioritization, processing and performance (the sinister 6 P's)

HTTP datagrams may have different priority needs from the related stream.

Datagrams are “atomic”, processed as they arrive.

Where QUIC connections coalesce application traffic there can be wins (congestion and fate sharing)

Where different application data flows exist, packetization and scheduling of data can put application performance at risk

Streams and datagrams could compete - starvation of flow type is a real problem.

draft-pardue-masque-dgram-priority

Extends HTTP Priorities

Captures input from issues and meetings.

Realizes that this is untrodden territory so starts as Experimental.



Workgroup: MASQUE
Internet-Draft: draft-pardue-masque-dgram-priority-latest
Published: 18 July 2021
Intended Status: Experimental
Expires: 19 January 2022
Author: L. Pardue
Cloudflare

HTTP Datagram Prioritization

Abstract

Application protocols using the QUIC transport protocol rely on streams, and optionally the DATAGRAM extension, to carry application data. Streams and datagrams can be multiplexed but QUIC provides no interoperable prioritization scheme or signaling mechanism itself. The HTTP Extensible Prioritization scheme describes how to prioritize streams in HTTP/2 and HTTP/3. This document adopts the scheme to support HTTP datagrams. ¶

Base HTTP Priorities recap

HTTP Extensible Prioritization for streams defines signals:

- *urgency* (“u”) - between 0 and 7. Smaller the value, higher the precedence
- *incremental* (“i”) - response can be processed incrementally (data as it arrives)

And some scheduling guidance:

- Expressing priority is only a suggestion.
- RECOMMENDED to respect urgency, serve in stream ID order.
- RECOMMENDED to respect incremental, fair bandwidth sharing between incremental at same urgency

The proposal's design philosophy and assumptions

Start simple.

HTTP datagrams **MAY** have different priority from the related stream.

But all datagrams related to a stream have the same priority. No per-context prioritization.

Flexibility with sensible defaults.

The proposal

Extend the priority scheme with a compatible parameter: *datagram-urgency* (“du”).

Identical to *urgency*, except that it applies to datagrams.

Omission of *datagram-urgency* is a signal to use the default. But there is no default value. Instead the default is to use the stream’s *urgency*.

Where stream and datagrams have the same urgency, default recommendation is to **alternate** between them when packetizing.

Examples

Explicit different. Datagrams lower precedence.	Priority: u=1, du=2
Explicit same.	Priority: u=1, du=1
Implicit same due to omission. datagram-urgency inherits the value of 1.	Priority: u=1
Implicit same via omission of both is legal. The stream urgency defaults to value of 3, datagram-urgency inherits the value of 3.	<intentionally blank>

~~Adopt this as a standalone draft?~~ Who cares?

Suspect datagram prioritization is not as easy as it seems.

Base QUIC and HTTP/3 drafts punt on priorities. ~~Let's not delay HTTP Priorities or HTTP Datagram by trying to perfect prioritization in them.~~ Too late now!

Need to consider extension of HTTP Datagram to HTTP/2.

Need to consider prioritization of Datagram Capsules? Is that even going to be possible...

~~If worth adopting, should it be in MASQUE or HTTPbis~~ Maybe its just too early to say and we can park this until people care? Should I renew the draft or leave it...

MASQUE

HTTP Priorities



draft-pardue-masque-dgram-priority

IETF 111 – Virtual Madrid – 2021-11

Lucas Pardue – lucaspardue.24.7@gmail.com