

IP Proxying Support for HTTP

[draft-ietf-masque-connect-ip-03](#)



Tommy Pauly, David Schinazi, **Alex Chernyakhovsky**,
Mirja Kühlewind, Magnus Westerlund

Interop progress

- Lots of successful interop!

| | | Server | | |
|--------|----------|--------|----------|-------|
| | | Google | Ericsson | Apple |
| Client | Google | ✓ | ✓ | ■ |
| | Ericsson | ✓ | ✓ | ■ |
| | Apple | ✓ | ✓ | ■ |

Changes in -03

All 7 issues closed!

- Clarified MTU responsibility, edge cases (#61, #62)
- Details on ICMP errors (#64, #71)
- Specify how to ask for "any" IP address (#66)
- Added request/response IDs to address capsules (#73)
- Discussion on error handling with dual-stack capsules (#72)

Clarified MTU responsibility, edge cases

IPv6 requires 1280 byte MTU. However, when intermediaries are involved, the Path MTU must be verified, not just the path from the client to the first server. The local link may not support inner MTU of at least 1280 bytes

Added text that endpoints **MUST** use the ICMPv6 method when intermediaries are in use, unless they have out-of-band information that the Path MTU is sufficient.

Added text that endpoints **MUST** abort the CONNECT-IP stream if the QUIC MTU does not allow sending 1280 bytes.

Details on ICMP errors

Previously, CONNECT-IP deferred all error signalling to recommendations to RFC 4443.

Details were added, citing RFC 792 and RFC 4443, as well as details on which errors to send when

- invalid source addresses are used
- unroutable destination addresses are encountered
- packets that cannot fit within the MTU

Added recommendation that endpoints request ICMP when settings IP PROTO, or if not using route advertisements, process the proxied ICMP packets to learn the error signals.

Specify how to ask for "any" IP address

Added a new feature for requesting "any" IP address, rather than an explicit one.

Endpoint requests an all-zeroes address in the desired address family, with a prefix-length of the desired subnet size.

Added request/response IDs to address capsules

"Ask for two, get one" problem. Previously CONNECT-IP didn't provide a guarantee that every request would get a response, or that requests would be responded to in-order.

Request/Response IDs were added to correlate the capsules, to avoid having to wait for a timeout to a capsule that will never arrive.

Discussion on error handling with dual-stack capsules

Proposal to split the capsules from holding both v4 and v6 addresses was considered, but ultimately declined as it had no clear benefit on error handling.

Next Steps

WGLC?

IP Proxying Support for HTTP

[draft-ietf-masque-connect-ip-03](#)