


TCP Fast Open for IPv6 prefixes

draft-hawari-tcpm-tfo-ipv6-prefixes-00

Introduction

- Research project about an IPv6-centric distributed storage system.
- Emulation of a hard drive by mapping each Logical Block Address to an IPv6 address.
- Send requests concerning a given block to the corresponding IPv6 address.



Motivation for the use of Fast Open

- A TCP connection is created for each request to the drive.
- Achievable throughput on the disk is very sensitive to RTT according to our measurement.
- Very transactional traffic : **Use of TCP Fast Open**

Current limitations of TCP Fast Open

- A TFO cookie is only valid for one IPv6 address
- It provides no real gain in the context of sparse connections to address within a same prefix.
- Idea : to extend the TFO cookie's validity domain to a whole IPv6 prefix.
- Applicability :
 - Assigning an IPv6 prefix to a server
 - Reusing a cookie for several servers in the same prefix

Proposed solution

- Let the server advertise the prefix length for which a TFO cookie is valid when the cookie is delivered to the client.
- TCP Fast Open Prefix Length Option proposed in draft-hawari-tcpm-tfo-ipv6-prefixes-00

TCP Option	Length	Prefix
Kind Number		Length
(1 BYTE)	(1 BYTE)	(1 BYTE)

- The client knows if it may reuse a cookie for an adjacent IPv6 address in a further connection.

Discussion

- This solution addresses the initial issue and is completely backwards-compatible with the previous behavior of TFO.
- **But it is far from ideal**
 - It requires a new TCP option and it is not necessarily worth it.
 - Perhaps we can let each application decide for which prefix length TFO cookies are valid. It would require that current implementations of TFO exposes a new API to the Application Layer.



Questions ?