Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 7, 2015

M. Boucadair
D. Binet
C. Jacquenet
France Telecom
L. Contreras
Telefonica I+D
Y. Lee
Comcast

March 6, 2015

On the Need for Transport Protocol Profiles & Investigating New Evolution Tracks
draft-boucadair-transport-protocols-01

Abstract

The world of Internet transport protocols is changing, after decades of TCP and UDP operation. Several proposals have been submitted for the past years (and counting) to introduce other transport protocols that aim at reducing the web latency of that of TCP or avoiding the burden of the various middle-boxes (NATs, firewalls, for one) encountered along the communication path. Such initiatives, although not new, are motivated by the complexity of some (non-transparent) networking functions.

This document advocates for the definition of transport profiles and the need to document recommendations for middleboxes, including Performance Enhancement Proxies (PEPs) behaviors. A collaboration among the involved players (service providers, vendors) is required to soften the current complications encountered in the Internet at large.

This Internet-Draft will expire on September 7, 2015.

Table of Contents

1. Introduction

The world of Internet transport protocols is changing, after decades of TCP and UDP operation. Several proposals have been submitted for the past years (and counting) to introduce other transport protocols or additional features to existing protocols that aim at reducing the web latency of that of TCP or avoiding the burden of the various middle-boxes (NATs, firewalls, for one) encountered along the communication path. Such initiatives, although not new, are motivated by the complexity of some (non-transparent) networking functions. Further collateral effects (including a thorough

identification of various network hindrances) are discussed in this document together with potential contributions from network operators to overcome some of the encountered issues.

Advanced service functions (e.g., Performance Enhancement Proxies ([RFC3135]), NATs, firewalls, etc.) are now required to achieve various objectives such as IP address sharing, firewalling, to avoid covert channels, to detect and protect against ever increasing DDoS attacks, etc. Removing those functions is not an option because they are used to address constraints that are often typical of the current yet protean Internet situation (global IPv4 address depletion comes to mind, but also the plethora of services with different QoS/security/robustness requirements, etc.), and this is even exacerbated by environment-specific designs (e.g., the nature and the number of service functions that need to be invoked at the Gi interface of a mobile infrastructure). Moreover, these sophisticated service functions are located in the network but also in service platforms, or intermediate entities (e.g., CDNs). This situation clearly complicates diagnostic procedures whenever service degradation is experienced, given That the responsibility is often shared among various players.

Also, there are performance issues that are specific to some wireless networks [I-D.manyfolks-gaia-community-networks].

An important effort was conducted by the IETF (e.g., BEHAVE, PCP, Performance Implications of Link Characteristics (pilc)), but we believe further work is still required to mitigate/soften some of the pending issues.

Note,

- "Middleboxes" or advanced Flow-Aware Service Functions are here to stay, whatever the progress of IPv6 adoption, in particular.
- Several experimental TCP extensions have been defined. These extensions (may) have merits when taken individually but further impact analysis is required when they have to co-exist in operational environments.
- HTTP/2 protocol ([I-D.ietf-httpbis-http2]) is being mostly implemented using TLS capabilities.
- More transport protocols encapsulated over TCP/UDP are being used by applications providers and vendors. Having a standard encapsulation scheme over TCP and UDP, including transport encapsulation recommendations, will help Network Providers fine tune their engineering rules and tweak of their networks.
- TCP proxies are widely present in operators architectures, specifically in mobile networks.

- Current evolution of transport and multiplexing services impact traffic patterns and optimization features set up to optimize resources and to improve Quality of Experience (QoS).
- Some application agents are not strictly following the "hard" limits of connections as indicated for instance in [RFC2068].
- Proposals to relax some of the TCP features (e.g., ordering) or to adopt an efficient byte stuffing schemes should be investigated.
- Transport over some media have specific requirements. An update of [RFC3150][RFC3481], for example, would be useful.
- Close collaboration and coordination between applications and networks can simplify if not improve network-inferred policy enforcement schemes. Applications may express their transport services requirements while transport protocols can expose, via advanced APIs, the functionalities they are offering and tweaking parameters that can be customized.
- Having a standard notification interface between the physical/link layer and the transport layer is likely to improve transport protocol performances in some networks.

Network Providers should be able to keep on delivering differentiated services as a competitive business advantage, while mastering the complexity of the applications, (continuously) evaluating the impacts on middleboxes, and enhancing customer's quality of experience. Because every (new) transport protocol will come with its own problems and perfectible features, leveraging skills and experience of TCP design and operation is a first major step for network providers.

This document advocates for the definition of transport profiles and the documentation of recommendations for middleboxes, including Performance Enhancement Proxy (PEPs) behaviors. A collaboration among the usual players is required to soften the current complications encountered in the Internet at large.

2. On Transport Services

Transport services refer to the set of features that are offered by protocols used to multiplex connections over IP. Examples of transport services include - but are not limited to- ordering delivery, reliable delivery, congestion control, or full or partial integrity protection.

A transport protocol can be abstracted as an implementation which exposes a set of transport services. For example, TCP (Transmission Control Protocol, [RFC0793]), which is the universally deployed and implemented transport protocol, offers reliable and ordered delivery, flow and congestion control, as well as primitives to manage a connection. Unlike TCP, UDP (User Datagram Protocol, [RFC0768]) is a

connectionless protocol that supports protection against data corruption using a checksum field.

3. Strategies to Enhance Transport services

Given the hurdles induced by advanced network-located service functions, "Make your own protocol" is not even an option.

"Encapsulate over your favorite existing protocol", if transported over TCP, has more chances to experience less session failures.

> This assumes that the remote server is also upgraded to support such transport scheme, while failures are likely to occur when the encapsulation is implemented over UDP.
>
> Examples of proposals that follow such mitigation strategy are [I-D.cheshire-tcp-over-udp], or [I-D.iyengar-minion-protocol].
>
> A fallback to TCP (or UDP) must be supported anyway, let alone the complications related to the discovery of the capabilities of the remote server.
>
> Even if protocols encapsulated over UDP can make use of NAT traversal techniques, these protocols are still suffering from issues related to the presence of NATs and firewalls. For example, there is no mechanism to notify endpoints that an entry is no more active in the NAT/Firewall. Immediate notification and state recovery can be solved by activating specific Port Control Protocol (PCP) feature: (PCP ANNOUNCE OPCODE, [RFC6887]).

The strategy that consists in "extending your favorite widely deployed transport protocol" is more viable from a deployment perspective.

> TCP can be extended [Options][ExtendTCP]. For example, extensions have been proposed to enhance user's quality of experience when TCP is in use such as: TCP Fast Open ([RFC7413]), Proportional Rate Reduction ([RFC6937]), increase the initial window ([RFC6928]), TCP Extensions for high performance ([RFC7323]), TCP-EDO [I-D.ietf-tcpm-tcp-edo], unordered TCP/TLS, etc.

4. Proliferation of Transport Protocols

Plethora of transport protocols have been proposed by the Internet community to accommodate requirements raised by emerging applications. Overall, these applications are either requiring more transport services than what is actually offered by TCP and UDP, or less transport services.

For example, SCTP (Stream Control Transmission Protocol, [RFC4960]) was specified to accommodate applications which need more transport services than what can be offered by TCP (e.g., preserve (application) data boundaries, support of out-of-order delivery, built-in support of multiple streams).

DCCP (Datagram Congestion Control Protocol, [RFC4340]) is another protocol that was promoted to accommodate requirements from applications which need more transport services than what is offered by UDP (e.g., congestion control), but without suffering from the constraints of a connection-oriented protocol like TCP (e.g., reliable delivery mechanisms).

UDP-lite ([RFC3828]) is a light version of UDP that was designed for applications that need less features than what is offered by UDP (e.g., partial data corruption detection), whereas DTLS (Datagram Transport Layer Security, [RFC6347]) and TLS ([RFC5246]) were specified for applications requiring encryption capabilities at the transport layer.

Other candidate transport protocols are currently investigated to reduce the delay required to invoke a resource located in the Internet. Typically, this consists in retrieving some contents by minimizing the delay induced by TCP or SCTP handshakes required for establishing a connection. Yet, such approaches can take advantage of the transport services provided by connection-oriented protocols.

It is worth mentioning that reducing the delay to access a requested resource is not only the responsibility of transport protocols, but also depends on various other services such as DNS and access service functions. The whole chain should be optimized! Reduce the delay when invoking a service objective should be moderated with other considerations such as policy enforcement at the server side (including rate-limit and actions taken to protect against DDoS attacks).

5. On TCP Hegemony

Despite the effort made by the Internet community to specify new transport protocols or propose improvements of existing ones (mainly TCP), the deployment reality is that TCP remains hegemonic. Even worse, only connections destined to some TCP port numbers are allowed in some networks.

Recent studies (e.g., [Traffic]) revealed that TCP accounts for 84.35% of the total amount of packets forwarded over the Internet and 92% of the bytes. DCCP and SCTP were not found in those studies.

The main reasons that explain the poor adoption of new transport-related features at the scale of the Internet are:

1. The presence of advanced network-located service functions (used to be called "middelboxes"), and

2. The lack of support by OSes.

Typical examples of service functions include: traditional NAT (Network Address Translation, [RFC3022]), CGN (Carrier Grade NAT; including IPv4-IPv4 CGN ([RFC6888]), DS-Lite AFTR ([RFC6333]) or NAT64 ([RFC6146])), firewall, application proxies, Performance Enhancement Proxies (PEP, [RFC3135]), traffic uniformizers, etc.

Transport-related solutions that assume that the remedy to the problem formulated above would be to withdraw all flow-aware service functions are not realistic. The presence of advanced service functions must be considered by solution designers as the rule rather than the exception.

Obviously, this does not mean that network providers should not question the pertinence to maintain some of these service functions active. Even if a rationalization effort is required in this area (still this is deployment-specific), solution designers should propose solutions that are robust in the presence of these functions.

6. Need for a Holistic View for TCP Variants

For example, variants have been proposed to enhance user's quality of experience when TCP is in use such as: TCP Fast Open ([RFC7413]), Proportional Rate Reduction ([RFC6937]), increase the initial window size ([RFC6928]), TCP Extensions for high performance ([RFC7323]) , unordered TCP/TLS, etc. More can be found in [RFC7414].

These variants may have merits when taken individually, but the question is whether those merits are still valid when co-existing with other features. In addition, these merits are a function of the deployment context (for example in fixed or mobile networks).

Implementing small changes at large is here to stay. Moreover, changing a transport protocol stack may is subject to the amplification principle (See Section 2.2.1 of [RFC3439]) since changes may not only have local impacts but may also impact the stability of a network (e.g., MPTCP hosts are more aggressive than TCP hosts). Assessing the impact of these variants on legacy hosts is critical.

7. Adoption Rate of TCP Extensions

According to [Traffic],

- 34% of TCP segments are data-less ACKs.
- 94% of SYN use SACK permitted option [RFC2018].
- MSS option in 96% TCP SYNs: A large number of TCP SYN messages advertise an MSS between 1000-1301 bytes (46% announces an MSS between 1300 bytes and 1460 bytes).
- Window Scale (WS) option and the TCP Timestamps (TS) [RFC7323] in 63.9% of TCP SYNs.
- 39.2% uses the TCP Timestamps (TS) [RFC7323] in TCP SYNs.
- Zero flows requesting ECN in TCP SYNs.

Risks of unordered delivery is often design-specific. Indeed, [Traffic] also showed that disordering is deployment-specific (because it was observed only in some networks); means that lead to such behavior should be disabled in those networks. This suggests reliable means to minimize such risks.

This data shows that several of TCP advances (e.g., WS) are not massively deployed or not deployed at all (e.g., ECN). A recent study about the support of ECN is available at [ECN].

More effort is required to evangelize recent TCP advances and their motivations.

8. A Network Provider's Perspective

8.1. Proposed Approach

Fortunately, there is still an opportunity for network providers to contribute to the improvement of transport services. A technical strategy that would focus on the root causes to properly derive associated recommendations should be encouraged.

Every (new) transport protocol will come with its own problems and perfectible features. Too many transport protocols are really painful for all actors, including for network operators (think about the configuration of class of services, fairness access and usage of network resources, and other traffic management services).

Leveraging skills and experience of TCP design as well as operation is a first major step for network providers. For example, in order to reduce latency for TCP-based applications, the following technical tracks can be investigated:

1. Deactivate ordering management;

2. Consider efficient byte stuffing schemes;

3. Get rid of the Three-Way Handshake; or

4. Consider persistent connections whenever suitable.

Network Providers should be able to keep on delivering differentiated services as a competitive business advantage, while mastering the complexity of the applications and enhancing customer's quality of experience. This can be achieved by exposing and communicating reachability information (i.e., routes to access desired contents) that will foster session establishment. This can be achieved using dedicated interfaces that can be used by applications.

Reduce complexity at the application level, strengthen the collaboration between the applications and the network layer via clear interfaces should also be encouraged, but this may be subject to agreements. Administrative-related considerations are out of scope of this document.

8.2. Some Risks:

From a network provider perspective, the following risks need to be taken into account when designing solution(s) that would enhance current transport services:

- Emergence of transport-specific proxies given that vendors promote their own transport protocols.

- In addition to the support of a fallback mode to TCP (or UDP), some of the proposals may lead to complex clients (application agents). This complexity should be avoided because this is likely to be a source of performance degradation, especially when other sophisticated features are required.

- Performance Enhancing Proxies are currently the rule to optimize TCP, especially in mobile networks. There is a need to agree on a TCP Profile, including required features to be supported by TCP acceleration engines (a.k.a., PEPs).

- Offloading some of the transport functions to the upper layers may be suitable for some cases (e.g., error detection) but this approach suffers from side effects such as buffering issues at the application level, potential misuse of the underlying transport service, complexity to diagnose degradation when it occurs, battery consumption for mobile devices because of frequent keepalives, etc.).

According to [Power], the consumption of a mobile device with a keep-alive interval of 20 seconds (that is the default value) is 29 mA (2G)/34 mA (3G). When no keep-alive is issued, the consumption would be 5.2 mA (2G)/6.1 mA (3G).

- Covert channels can be made possible if encapsulation schemes are allowed without any security features.

- The accuracy of the engineering and tuning of network devices for an optimized service delivery may be impacted by the variety of traffic profiles, and, especially the change of the transport behavior (e.g., aggressive vs. other flows, fairness to make use of network resources, etc.).

- Path diversity (e.g., be able to establish a TCP communication over different paths for the sake of optimized bandwidth usage) becomes a typical requirement given the current adoption rate of multi-interfaced devices.

  Networks can cooperate with applications to help selecting the best path(s) but diverse transport protocols can provoke service disruption when the device re-connects to another network (e.g., via a WLAN Hotspot, mobile, CPE, etc.), where network-assisted functions are hosted.

9. Previous IETF Works

Some recommendations to improve transport services have been documented for quite some time (e.g., [RFC4787], [RFC5382]).

Such recommendations are related to the design and the operation of services in the presence of flow-aware devices (in particular, NATs). A few examples: the use of endpoint-independent NAT mapping (EIM) and filtering (EIF) behaviors, IP address pooling behavior of "Paired" to not break protocols such as RTP/RTCP, the selection of long mapping lifetime values to avoid breaking some applications, the preservation of port parity for RTP/RTCP-based applications (like VoIP), the preservation of port contiguity for some applications, the use of port randomness to avoid session hijacking, the ability to discover the external IP address/port/lifetime ([RFC6887]) so that applications with referral behave with no degradation, the analysis of the use of the HOST_ID ([RFC6967]) to soften issues induced by address sharing at large ([RFC6269]), etc.

An effort to clarify some of the behave requirements is ongoing ([I-D.ietf-tsvwg-behave-requirements-update]).

Also, the Performance Implications of Link Characteristics (pilc) WG conducted an important effort which led to [RFC3135][RFC3150][RFC3155][RFC3366][RFC3449][RFC3481][RFC3819].

10. What's Next?

The following candidate actions are proposed (non-exhaustive list):

- Define a TCP profile for hosts. This profile can be an update of [RFC1122]. Or not.

- Update PEP recommendations. Edit a BCP document about TCP extensions to be supported by middleboxes vendors and activated by operators.

  E.g., Update the header compression features recommended in [RFC3150] to include [RFC4996].

- Specify a MPTCP Profile in network regions that are firewall- and NAT-free: One of the promising deployment scenario for MPTCP ([RFC6824]) is to aggregate the resources offered by a CPE that is connected to multiple networks (e.g., DSL, LTE, WLAN), see for example [I-D.deng-mptcp-proxy] or [I-D.lhwxz-hybrid-access-network-architecture].

  This deployment scenario requires a kind of "concentrator" at the network side to terminate the aggregated session before relaying it into a legacy TCP session. The concentrator is needed before the adoption rate of MPTCP at the server side is taking.

  Because the paths between the CPE and the concentrator are firewall- and NAT-free, the complexity of the MPTCP specification that was initially induced by handling the presence of firewalls and the routing asymmetry, is not justified anymore. Such context encourages the specification of a dedicated MPTCP profile that would in turn foster the adoption of MPTCP.

- Standardize encapsulation schemes over TCP and UDP.

11. IANA Considerations

This document makes no request of IANA.

12. Security Considerations

Add some text about privacy and security.

13. Acknowledgements

Many thanks to J. Touch for the comments.

14. Informative References

[ECN] B. Trammell, M. Kuehlewind, D. Boppart, I. Learmonth, G. Fairhurst, and R. Scheffenegger, "Enabling Internet-Wide Deployment of Explicit Congestion Notification", 2015, <http://ecn.ethz.ch/ecn-pam15.pdf>.

[ExtendTCP] Honda, M., Nishida, Y., Raiciu, C., Greenhalgh, A., Handley, M. and H. Tokuda,, "Is it still possible to extend TCP?", November 2011, <http://nrg.cs.ucl.ac.uk/mjh/tmp/mboxes.pdf>.

[I-D.cheshire-tcp-over-udp] Cheshire, S., Graessley, J., and R. McGuire, "Encapsulation of TCP and other Transport Protocols over UDP", draft-cheshire-tcp-over-udp-00 (work in progress), July 2013.

[I-D.deng-mptcp-proxy] Lingli, D., Liu, D., Sun, T., Boucadair, M., and G. Cauchie, "Use-cases and Requirements for MPTCP Proxy in ISP Networks", draft-deng-mptcp-proxy-01 (work in progress), October 2014.

[I-D.ietf-httpbis-http2] Belshe, M., Peon, R., and M. Thomson, "Hypertext Transfer Protocol version 2", draft-ietf-httpbis-http2-17 (work in progress), February 2015.

[I-D.ietf-tcpm-tcp-edo] Touch, J. and W. Eddy, "TCP Extended Data Offset Option", draft-ietf-tcpm-tcp-edo-01 (work in progress), October 2014.

[I-D.ietf-tsvwg-behave-requirements-update] Penno, R., Perreault, S., Kamiset, S., Boucadair, M., and K. Naito, "Network Address Translation (NAT) Behavioral Requirements Updates", draft-ietf-tsvwg-behave-requirements-update-01 (work in progress), February 2015.

[I-D.iyengar-minion-protocol]
Jana, J., Cheshire, S., and J. Graessley, "Minion - Wire Protocol", draft-iyengar-minion-protocol-02 (work in progress), October 2013.

[I-D.lhwxz-hybrid-access-network-architecture]
Leymann, N., Heidemann, C., Wasserman, M., Xue, L., and M. Zhang, "Hybrid Access Network Architecture", draft-lhwxz-hybrid-access-network-architecture-02 (work in progress), January 2015.

[I-D.manyfolks-gaia-community-networks]
Saldana, J., Arcia-Moret, A., Braem, B., Navarro, L., Pietrosemoli, E., Rey-Moreno, C., Sathiaseelan, A., and M. Zennaro, "Alternative Network Deployments. Taxonomy and characterization", draft-manyfolks-gaia-community-networks-02 (work in progress), January 2015.

[Options] Alberto Medina, Mark Allman, Sally Floyd, "Measuring Interactions Between Transport Protocols and Middleboxes", 2005, <http://conferences.sigcomm.org/imc/2004/papers/p336-medina.pdf>.

[Power] Haverinen, H., Siren, J., and P. Eronen, "Energy Consumption of Always-On Applications in WCDMA Networks", April 2007, <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4212635>.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.

[RFC2018] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgment Options", RFC 2018, October 1996.

[RFC2068] Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2068, January 1997.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

[RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, June 2001.

[RFC3150] Dawkins, S., Montenegro, G., Kojo, M., and V. Magret, "End-to-end Performance Implications of Slow Links", BCP 48, RFC 3150, July 2001.

[RFC3155] Dawkins, S., Montenegro, G., Kojo, M., Magret, V., and N. Vaidya, "End-to-end Performance Implications of Links with Errors", BCP 50, RFC 3155, August 2001.

[RFC3366] Fairhurst, G. and L. Wood, "Advice to link designers on link Automatic Repeat reQuest (ARQ)", BCP 62, RFC 3366, August 2002.

[RFC3439] Bush, R. and D. Meyer, "Some Internet Architectural Guidelines and Philosophy", RFC 3439, December 2002.

[RFC3449] Balakrishnan, H., Padmanabhan, V., Fairhurst, G., and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry", BCP 69, RFC 3449, December 2002.

[RFC3481] Inamura, H., Montenegro, G., Ludwig, R., Gurtov, A., and F. Khafizov, "TCP over Second (2.5G) and Third (3G) Generation Wireless Networks", BCP 71, RFC 3481, February 2003.

[RFC3819] Karn, P., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L. Wood, "Advice for Internet Subnetwork Designers", BCP 89, RFC 3819, July 2004.

[RFC3828] Larzon, L-A., Degermark, M., Pink, S., Jonsson, L-E., and G. Fairhurst, "The Lightweight User Datagram Protocol (UDP-Lite)", RFC 3828, July 2004.

[RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, March 2006.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

[RFC4996] Pelletier, G., Sandlund, K., Jonsson, L-E., and M. West, "RObust Header Compression (ROHC): A Profile for TCP/IP (ROHC-TCP)", RFC 4996, July 2007.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, October 2008.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.

[RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.

[RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

[RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, April 2013.

[RFC6928] Chu, J., Dukkipati, N., Cheng, Y., and M. Mathis, "Increasing TCP's Initial Window", RFC 6928, April 2013.

[RFC6937] Mathis, M., Dukkipati, N., and Y. Cheng, "Proportional Rate Reduction for TCP", RFC 6937, May 2013.

[RFC6967] Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Potential Solutions for Revealing a Host Identifier (HOST_ID) in Shared Address Deployments", RFC 6967, June 2013.

[RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, "TCP Extensions for High Performance", RFC 7323, September 2014.

[RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, December 2014.

[RFC7414] Duke, M., Braden, R., Eddy, W., Blanton, E., and A. Zimmermann, "A Roadmap for Transmission Control Protocol (TCP) Specification Documents", RFC 7414, February 2015.

[Traffic] David Murray, Terry Koziniec, "The State of Enterprise Network Traffic in 2012", 2012, <http://conferences.sigcomm.org/imc/2004/papers/p336-medina.pdf>.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

David Binet
France Telecom
Rennes 35000
France

Email: david.binet@orange.com

Christian Jacquenet
France Telecom
Rennes
France

Email: christian.jacquenet@orange.com

Luis M. Contreras
Telefonica I+D
Ronda de la Comunicacion, s/n
Madrid 28050
Spain

Email: lmcm@tid.es
URI: http://people.tid.es/LuisM.Contreras/

Yiu Lee
Comcast
US

Email: Yiu_Lee@Cable.Comcast.com