

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: September 14, 2017

A. Jain
A. Terzis
Google
H. Flinck
N. Sprecher
S. Arunachalam
Nokia Networks
K. Smith
Vodafone
V. Devarapalli
R. Bar Yanai
Vasona Networks
March 13, 2017

Mobile Throughput Guidance Inband Signaling Protocol
draft-flinck-mobile-throughput-guidance-04.txt

Abstract

The bandwidth available for end user devices in cellular networks can vary by an order of magnitude over a few seconds due to changes in the underlying radio channel conditions, as device mobility and changes in system load as other devices enter and leave the network. Furthermore, packets losses are not always signs of congestion. The Transmission Control Protocol (TCP) can have difficulties adapting to these rapidly varying conditions leading to inefficient use of a cellular network's resources and degraded application performance. Problem statement, requirements and the architecture for a solution is documented in [[Req_Arch_MTG_Exposure](#)].

This document proposes a mechanism and protocol elements that allow the cellular network to provide near real-time information on capacity available to the TCP server. This "Throughput Guidance" (TG) information would indicate the throughput estimated to be available at the radio downlink interface (between the Radio Access Network (RAN) and the mobile device (UE)). TCP server can use this TG information to ensure high network utilization and high service delivery performance. The document describes the applicability of the proposed mechanism for video delivery over cellular networks; it also presents test results from live operator's environment.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 14, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Contributing Authors	3
1.2.	Terminology	3
1.3.	Acronyms and Abbreviations	3
1.4.	Definitions	4
1.5.	Assumptions and Considerations for the Solution	4
2.	Protocol	6
2.1.	Message Format	8
2.2.	Authentication	10
3.	Applicability to Video Delivery Optimization	10
3.1.	Test Results	11
4.	Manageability considerations	12
5.	Security considerations	12
6.	IANA considerations	13
7.	Acknowledgements	13
8.	References	13
8.1.	Normative References	14
8.2.	Informative References	14
Appendix A.	15

Authors' Addresses	15
------------------------------	--------------------

[1.](#) Introduction

The problem statement related to the behavior of the TCP in cellular networks is provided in [[Req_Arch_MTG_Exposure](#)]. That same document specifies the requirements, reference architecture and proposed solution principles for a mobile throughput guidance exposure mechanism that can be used to assist TCP in cellular networks, ensuring high utilization and high service delivery performance.

This document presents a set of considerations and assumptions for the development of a solution. It specifies a protocol that addresses the requirements and the architecture stated in the [[Req_Arch_MTG_Exposure](#)]. This document describes also the applicability of the proposed mechanism to video delivery over cellular networks with test results from live production environment.

[1.1.](#) Contributing Authors

The editors gratefully acknowledge the following additional contributors: Peter Szilagyi/Nokia, Csaba Vulkan/Nokia, Ram Gopal/Nokia, Guenter Klas/Vodafone and Peter Cosimini/Vodafone.

[1.2.](#) Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

[1.3.](#) Acronyms and Abbreviations

This document uses the following acronyms:

ECGI	E-UTRAN Cell Global Identifier format
ECN	Explicit Congestion Notification
HMAC	Hash-based Message Authentication Code
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IP	Internet Protocol
IV	Initialization Vector
LTE	Long Term Evolution
MTG	Mobile Throughput Guidance
RAN	Radio Access Network
RCTP	RTP Control Protocol
RTT	Round Trip Time
SACK	Selective Acknowledgement
TCP	Transmission Control Protocol
TCP-EDO	TCP Extended Data option
TG	Throughput Guidance
UE	User Equipment

1.4. Definitions

Throughput Guidance Provider:

A functional element that has access to the radio network information and signals to the TCP server, information about the (near-real time) throughput estimated to be available to a UE at the radio downlink interface

1.5. Assumptions and Considerations for the Solution

This document specifies a solution protocol that is complies with the requirements and architecture specified in [[Req_Arch_MTG_Exposure](#)]. The protocol is used by the cellular network to provide throughput guidance information to the TCP server; this information indicates the throughput estimated to be available at the radio downlink interface for the TCP connection. The protocol allows the information to be provided in near real time in situations where the network conditions are changing frequently or the user is moving.

While the implementation details can vary according to the access technology, the resource allocation is abstracted as the capacity of the "radio link" between the RAN and the UE. For example, in the case of an LTE network, the number of physical resource blocks allocated to a UE, along with the modulation scheme and coding rate used, can be translated into radio link capacity in Megabits per second (Mbit/s). From the derived UE's total throughput and with the

UE's TCP flow information, Throughput guidance for the TCP connection can be computed.

The TCP server can use this explicit information to inform several congestion control decisions. For example: (1) selecting the initial congestion window size, (2) deciding the value of the congestion window during the congestion avoidance phase, and (3) adjusting the size of the congestion window when the conditions on the "radio link" change. In other words, with this additional information, TCP neither has to congest the network when probing for available resources (by increasing its congestion window), nor rely on heuristics to decide how much it should reduce its sending rate after a congestion episode.

The same explicit information can also be used to optimize application behavior given the available resources. For example, when video is encoded in multiple bitrates, the application server can select the highest encoding rate that the network can deliver.

This solution specified in this document also satisfies the following assumptions and considerations:

- o The end-to-end traffic is delivered via HTTP.
- o The end-to-end traffic is encrypted (through HTTPS), thus HTTP header enrichment cannot be used by intermediate elements between the client and the server.
- o TCP is used to deliver the HTTPS traffic.
- o The Real-time Transport Protocol (RTP) network protocol is not used for traffic delivery.

The protocol specified in this document assumes that a trustful relationship between the Throughput Guidance Provider and the TCP server has been formed using the means discussed in the Security considerations section.

The solution in this document satisfies the considerations and the assumptions presented above, and proposes an in-band exposure mechanism where the throughput guidance information is added to the TCP headers of the relevant upstream packets. HTTP and TCP are the most prevalent protocols in the Internet, used even by the most popular streaming application. Throughput guidance at TCP level can be shared among multiple applications; it is not limited to any particular application level optimization only but it offers a generic approach that works even if application level end-to-end encryption, such as HTTPS, is applied.

In particular, the Throughput Guidance Providers adds the throughput guidance information to the Options field of the TCP header (see [RFC 0793](#) [[RFC0793](#)]) of packets from the TCP client to the TCP server. An in-band mechanism is proposed because it does not require a separate interface, reference value, or correlation mechanism that would be needed with out of band approaches such as with RCTP that is limited to only certain types of applications. Furthermore, an in-band mechanism can keep up with the rapid changes in the underlying radio link throughput. Unlike existing mechanisms such as ECN, where an ECN-aware router sets a mark in the IP header in order to signal impending congestion (see [[RFC3168](#)]). The proposed scheme provides explicit information, (termed "Throughput Guidance") about the estimated throughput available for the TCP connection at the radio link between the RAN and the UE.

Note that once standardized and implemented, TCP Extended Data option (TCP-EDO) can be used to carry the throughput guidance information as specified in [[tcp-edo](#)] and simplify the use of the TCP Option fields by extending the space available for TCP options. Currently the TCP-EDO is still work in progress and not available in production. Therefore, the use of TCP-EDO to carry throughput guidance is left for the later drafts.

2. Protocol

This section describes the protocol mechanism and the message format that needs to be communicated from the RAN to the TCP remote endpoint. We describe the protocol mechanism and message format for throughput guidance. The protocol mechanism is defined in an extensible way to allow additional information to be specified and communicated. The protocol specification is based on the existing experiments and running code. It is recommended to insert the throughput guidance information to the TCP segments that flow from client to server (see reasoning in "Assumptions and Considerations" section). Most of the time, TCP segments are ACK packets from a client to the server and hence packets are unlikely to be fragmented. However, the described protocol solution can deal with fragmentation.

The Mobile Throughput Guidance Signaling message conveys information on the throughput estimated to be available at the down link path for a given TCP connection. The information is sent to the uplink endpoint of the connection (i.e, the TCP server). The TCP server MAY use this information to adapt TCP behavior and to adjust application-level behavior to the link conditions as defined in [[Req_Arch_MTG_Exposure](#)].

A good example is a content optimizer or a cache that can adapt the application-level coding to match the indicated downlink radio

conditions. As radio link conditions may change rapidly, this guidance information is best carried in-band using TCP options headers rather than through an out-of-band protocol.

Using the TCP options to carry throughput guidance associates the guidance information with an ongoing TCP connection and explicitly avoids separate session identification information. The proposed mechanism neither impacts the TCP state machine nor the congestion control algorithms of the TCP protocol.

The Options field enables information elements to be inserted into each packet with a 40-byte overall limit; this needs to be shared with the standardized and widely-used option elements, such as the TimeStamp and SACK. (Use of TCP-EDO will lift this constraint once available and deployed). The TCP Options field uses a Kind-Length-Value structure that enables TCP implementations to interpret or ignore information elements in the Options field based on the Kind.

In this draft, we define a message format for encoding information about the estimated capacity of a radio access link between the RAN and the UE which is traversed by a TCP connection. The intention is to define a generic container to convey in-band information within the limited TCP Option space with optional authentication capabilities. This document conveys throughput guidance information. Additional information can be specified in future.

The Throughput Guidance Provider functional element inserts Mobile Throughput Guidance TCP options only if there is enough space in the TCP header. The Throughput Guidance Provider has access to the radio network information and is typically co-located with the RAN functionality.

The Throughput Guidance information must be delivered in a secure way, such that an intermediate node cannot modify it. The information can be provided as plain text in a secure and closed network. In other cases, the information should be authenticated (between the Throughput Guidance Provider and the TCP server). An acceptable level of authentication (according to best common practices) may require more data than fits into a single TCP header (maximum of 40 bytes if no other options are present). As described below, fragmenting information across multiple packets will be used if such is the case.

Two transfer modes are defined to deal with security of exchanged throughput guidance information in this document; namely, plain-text mode and authenticated mode. A third mode, encryption with authentication mode, is equally feasible and may be described in a

future revision of this protocol. The flags field indicate which mode is used.

2.1. Message Format

Mobile Throughput Guidance Signaling uses the common TCP options structure as in [RFC0793] with experimental identifier as defined in [RFC6994]. The followign defines the message format:

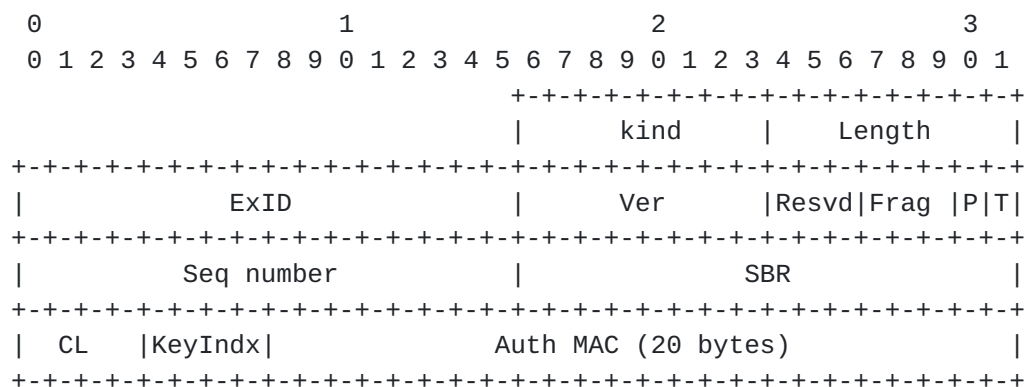


Figure 1

Kind:

Code point 253 for Experimental Option for 16-bit ExID [RFC6994]. The size of this field is 1 byte.

Length:

A 1-byte field, length of the option in bytes as defined in [RFC793].

ExID:

Two bytes Experimental Identifier according to [RFC6994]. Code point 0x6006.

Ver:

Version of the protocol, set to 1.

Flags:

One byte of MTG protocol flag field as defined below.


```

  0 1 2 3 4 5 6 7
+-+--+--+--+--+
|Resvd|Frag |P|T|
+-+--+--+--+--+

```

Flag field of common Kind-Lenght-Value header

Figure 2

Frag:

Three bits that provide information about how to reassemble information if fragmented into multiple packets. If no fragmentation across multiple TCP packet headers is needed, these bits are set to zero. Otherwise, Frag is a counter starting from 1 and incremented by 1 for each subsequent packet of the same type (see P- and T-bits below). For the last fragment, the Fragment is always 7 (binary 111) to indicate that the information is complete.

P and T bits:

These two bits encode the packet type: Plaintext (P=0, T= 0), Cipher text (P=0, T=1), Nonce (IV) (P=1, T=0) or Authentication (P=1, T=1). For Plaintext, the Fragment bits are always zero.

Seq Number:

16-bit sequence number to protect against replay attacks

SBR:

Suggested bit rate for the data session in Mbps. The 12 most significant bits are used for the integer value while the bottom 4 bits correspond to the decimal portion of the throughput value.

CL:

Cell Congestion Level (0, 1, 2, 3). A 4-bit field that indicates the current cell congestion level. "0" indicates no congestion and "3" indicates high congestion value.

Key Index:

A 4-bit field to identify the key used for integrity protection.

Auth MAC:

20 bytes of MAC that protects the TCP option

2.2. Authentication

Authentication covers the entire TCP option, excluding the Flags field and the Auth MAC field. The authentication uses HMAC codes (e.g. HMAC-SHA2-224), 128 bits (16 bytes) key size, 256 bits (32 bytes) digest size. Multiple keys (at most 256) for authentication with the same information receiver can be used. Truncation is possible but at least 160 bits (20 bytes) must be used from the digest to meet the typical security level of mobile networks.

Authentication takes a key, the input (arbitrary length) and produces a 32 byte long digest, which is truncated to 20 bytes (keeping the most significant bytes). The HMAC algorithm and truncation can be negotiated via key management (out of scope of this document).

The order in which the fields are included into the message authentication code is the same as the order in which the bytes appear in the message format.

In case the option packets used as input to the HMAC are fragmented into multiple TCP headers, they are processed so that headers with cipher text option are processed first, followed by IV/Nonce option packets.

The options containing the result of the HMAC are marked by setting both P- and T-bits of the flag-field to one. Key Index is set to point to the used authentication key, followed by the resulting authentication code. If the option doesn't fit into the free option space in the TCP header, it is fragmented across multiple TCP headers in the same way as the cipher text options.

3. Applicability to Video Delivery Optimization

The applicability of the protocol specified in this document to mobile video delivery optimization has been evaluated and tested in different network load scenarios.

In this use case, TCP traffic, for which throughput guidance information is required, passes through a Radio Analytics application which resides in a Mobile-edge Computing (MEC) server (see [\[MEC White Paper\]](#)). This Radio Analytics application acts as the Throughput Guidance Provider and sends throughput guidance information for a TCP connection using the Options field in the TCP header (according to the message specification provided in [section 2](#)). The TCP server MAY use this information to assist TCP congestion control decisions as described above. The information MAY also be

used to select the application level coding so that it matches the estimated capacity at the radio downlink for that TCP connection.

All of these improvements aim to enhance the quality of experience of the end user by reducing the time-to-start of the content as well as video stall occurrences.

3.1. Test Results

Nokia Networks and Google tested the video delivery optimization use case in a live production LTE network. Google server was placed close to the packet core network of LTE (SGi-interface of LTE). Different network load scenarios were taken into consideration. TCP CUBIC was used in these tests [[MTG ICCRG](#)].

Field trial performance results

Performance metric	Difference of Averages (%)	Diff of 99th percentiles
Time to play	-8.0%	-12%
Number of formats	+4.1%	+29.9%
Client bandwidth	+0.7%	+8.0%
Ave Video resolution	+6.2%	+5.6%
Re-buffer time	-19.7%	-5.1%

Table 1: Performance Data

These user experience improvements results into better video play and are likely to offer longer battery life.

Table 3 summarizes the results from a field trial in the 3G network of a Tier 1 mobile network operator in the Americas. As in the previous case, the Google servers were located close to the packet core network while the network elements from Vasona Networks generated the Throughput Guidance messages.

It is interesting to note that the improvements that Throughput Guidance provides are qualitatively different in LTE and 3G networks. LTE networks generally have capacities that cannot be fully utilized by CUBIC's initial sending rate. On the other hand, Throughput Guidance tends to be more conservative in 3G networks leading to higher time to play. Video re-buffering is smaller in both cases.

3G field trial performance results

Performance metric	Difference of Averages (%)	Diff of 99th percentiles
Time to play	+0.8%	+4.6%
Number of formats	-1.2%	+0.0%
Client bandwidth	+2.5%	-0.5%
Ave Video resolution	+0.5%	-1.9%
Re-buffer time	-13.9%	-10.9%

Table 2: Performance Data from 3G field trial

4. Manageability considerations

The application in the RAN SHOULD be configured with a list of destinations to which throughput guidance should be provided. The application in RAN will supply mobile throughput guidance information to more than one TCP server simultaneously based on the list of destinations.

In addition, it SHOULD be possible to configure the frequency (in milliseconds) at which throughput guidance needs to be signaled as well as the required security level and parameters for the encryption and the authentication if supported.

5. Security considerations

Throughput guidance SHOULD be provided in a secure way. The information can be provided as plain text in a secure and closed network (e.g. inside operator network). In other cases, the information should be authenticated (between the Throughput Guidance Provider and the TCP server).

[Section 2](#) described how the TCP Header information is protected. An out-of-band mechanism is currently used to agree upon the set of keys used to authenticate the messages exchanged between the endpoint and the network element that generates the throughput guidance headers. For example, service providers/OTTs can provide a portal that network providers can use to configure the keys they use to encrypt/sign the throughput guidance information. Then, a service behind the portal ensures that the keys are distributed to the servers that need them.

As stated in [\[Req_Arch_MTG_Exposure\]](#), the policy configuration of the Throughput Guidance Provider and the server endpoint, as well as the

key management are beyond the scope of this protocol definition. The protocol assumes that a trustful relationship has been formed between the Throughput Guidance Provider and the TCP server and that the required security level is already configured by the operator and agreed between the entities (i.e. authentication, encryption or both).

The identity of the Mobile Throughput Guidance provider that injects the throughput guidance header must be explicitly known to the endpoint receiving the information. Omitting such information would enable malicious third parties to inject erroneous information.

Fortunately, the issue of malicious disinformation can be easily addressed using well known techniques. First, the network entity responsible for injecting the throughput guidance header can include a cryptographically secure message authentication code. In this way the transport endpoint that receives the throughput guidance header can check that the information was sent by a legitimate entity and that the information has not been tampered with.

Furthermore, the throughput guidance information should be treated only as an estimate to the congestion control algorithm running at the transport endpoint. The endpoint that receives this information should not assume that it is always correct and accurate. Specifically, endpoints should check the validity of the information received and if they find it erroneous they should discard it and possibly take other corrective actions (e.g., discard all future throughput guidance information from a particular IP prefix). Endpoints MUST process throughput guidance information only from TCP segments that would otherwise be accepted as part of the standard TCP input process. For example, the receiver should ignore throughput guidance information included in TCP ACKs whose acknowledgement sequence numbers fall outside the range of valid sequence numbers.

6. IANA considerations

In the current version of the document and for field tests, the experimental value 253 is used for the "Throughput Guidance" TCP option kind. ExpID SHOULD be set to 0x6006 (16 bits)

7. Acknowledgements

8. References

8.1. Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<http://www.rfc-editor.org/info/rfc793>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6994] Touch, J., "Shared Use of Experimental TCP Options", [RFC 6994](#), DOI 10.17487/RFC6994, August 2013, <<http://www.rfc-editor.org/info/rfc6994>>.

8.2. Informative References

- [I-D.narten-iana-considerations-rfc2434bis]
Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [draft-narten-iana-considerations-rfc2434bis-09](#) (work in progress), March 2008.
- [MEC_White_Paper]
ETSI, "Mobile-Edge Computing - Introductory Technical White Paper", 2014.
- [MTG_ICCRG]
Szilagyi, P., and Terzis, A., "Mobile Content Delivery Optimization based on Throughput Guidance", Presentation at ICCRG meeting IETF93 (work in progress), July 2015.
- [Req_Arch_MTG_Exposure]
Jain, A., , Terzis, A., , Sprecher, N., , Arunachalam, S., , Smith, K., , and G. Klas, "Requirements and reference architecture for Mobile Throughput Guidance Exposure", [draft-sprecher-mobile-tg-exposure-req-arch-01.txt](#) (work in progress), February 2015.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), DOI 10.17487/RFC2629, June 1999, <<http://www.rfc-editor.org/info/rfc2629>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.

- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<http://www.rfc-editor.org/info/rfc3552>>.
- [RFC4413] West, M. and S. McCann, "TCP/IP Field Behavior", [RFC 4413](#), DOI 10.17487/RFC4413, March 2006, <<http://www.rfc-editor.org/info/rfc4413>>.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", [RFC 5925](#), DOI 10.17487/RFC5925, June 2010, <<http://www.rfc-editor.org/info/rfc5925>>.
- [tcp-ao-encrypt] Touch, J., , "A TCP Authentication Option Extension for Payload Encryption", [draft-touch-tcp-ao-encrypt-02.txt](#) (work in progress), November 2014.
- [tcp-edo] Touch, J., and Eddy, W., "TCP Extended Data Offset Option", [draft-ietf-tcpm-tcp-edo-01.txt](#) (work in progress), October 2013.

Appendix A.

Authors' Addresses

Ankur Jain
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Phone: +1-925-526-5879
Email: jankur@google.com

Andreas Terzis
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Phone: +1-650-214-5270
Email: aterzis@google.com

Hannu Flinck
Nokia Networks
Karaportti 13
Espoo
FI

Phone: +358504839522
Email: hannu.flinck@nokia-bell-labs.com

Nurit Sprecher
Nokia Networks
Hod HaSharon
IL

Phone: +97297751229
Email: nurit.sprecher@nokia.com

Swaminathan Arunachalam
Nokia Networks
Irving, TX
US

Phone: +19723303204
Email: swaminathan.arunachalam@nokia.com

Kevin Smith
Vodafone
One Kingdom Street, Paddington Central
London W2 6BY
UK

Phone: +19723303204
Email: kevin.smith@vodafone.com

Vijay Devarapalli
Vasona Networks

Email: vijay@vasonanetworks.com

Roni Bar Yanai
Vasona Networks

Email: rbaryanai@vasonanetworks.com

