

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 5, 2015

W. George
Time Warner Cable
P. Fan
China Mobile
S. Matsushima
SoftBank Telecom
T. Reddy
C. Eckel
Cisco Systems, Inc.
July 4, 2014

Application Enabled Collaborative Networking Use Cases
draft-conet-aeon-use-cases-01

Abstract

This document describes application enabled collaborative networking use cases. Application enabled collaborative networking has applications explicitly signal their flow characteristics to the network. This provides network nodes with visibility of the application flow characteristics, which enables them to apply the correct flow treatment and provide feedback to applications.

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 January 5, 2015.

Copyright Notice

Copyright (c) 2014 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
2.	Use Cases	4
2.1.	Firewall Traversal: Identification of new applications .	5
2.1.1.	Description of Problem	5
2.1.2.	Proposed Solution	5
2.1.3.	User/Application Benefit	6
2.1.4.	Operator Benefit	6
2.1.5.	Flow characteristics provided by application	6
2.1.6.	Action taken by firewall as result of receiving flow characteristics	6
2.1.7.	Feedback provided by firewall	6
2.1.8.	Security and Privacy Considerations	6
2.2.	Efficient Capacity Usage	6
2.2.1.	Description of Problem	6
2.2.2.	Proposed Solution	8
2.2.3.	User/Application Benefit	9
2.2.4.	Operator Benefit	9
2.2.5.	Flow characteristics provided by application	9
2.2.6.	Action taken by network as result of receiving flow characteristics	9
2.2.7.	Feedback provided by network	9
2.2.8.	Security and Privacy Considerations	10
2.3.	Video Adaptation	10
2.3.1.	Description of Problem	10
2.3.2.	Proposed Solution	11
2.3.3.	User/Application Benefit	11
2.3.4.	Operator Benefit	11
2.3.5.	Flow characteristics provided by application	11
2.3.6.	Action taken by network as result of receiving flow characteristics	12
2.3.7.	Feedback provided by network	12
2.3.8.	Security and Privacy Considerations	12
2.4.	Multi-interface selection: Use metadata to help interface selection or prioritization	12
2.4.1.	Description of Problem	12
2.4.2.	Proposed Solution	12
2.4.3.	User/Application Benefit	12
2.4.4.	Operator Benefit	13

2.4.5.	Flow characteristics provided by application	13
2.4.6.	Action taken by network as result of receiving flow characteristics	13
2.4.7.	Feedback provided by network	13
2.4.8.	Security and Privacy Considerations	13
2.5.	Session Identification: Identification of multiple media flows belonging to a common application session	13
2.5.1.	Description of Problem	13
2.5.2.	Proposed Solution	14
2.5.3.	User/Application Benefit	14
2.5.4.	Operator Benefit	14
2.5.5.	Flow characteristics provided by application	14
2.5.6.	Action taken by network as result of receiving flow characteristics	15
2.5.7.	Feedback provided by network	15
2.5.8.	Security and Privacy Considerations	15
2.6.	Content Based Charging	15
2.6.1.	Description of Problem	15
2.6.2.	Proposed Solution	15
2.6.3.	User/Application Benefit	16
2.6.4.	Operator Benefit	16
2.6.5.	Flow characteristics provided by application	16
2.6.6.	Action taken by network as result of receiving flow characteristics	16
2.6.7.	Feedback provided by network	16
2.6.8.	Security and Privacy Considerations	16
3.	Acknowledgements	17
4.	Informative References	17
	Authors' Addresses	17

1. Introduction

Identification and treatment of application flows are important to many application providers and network operators. They often rely on these capabilities to deploy and/or support a wide range of applications. These applications, and the packet flows they generate and consume, may require specific bandwidth, latency, etc., that can be better met if made known to the network. Historically, this functionality has been implemented to the extent possible using heuristics, which inspect and infer flow characteristics. Heuristics may be based on port ranges, network separation (e.g. subnets or VLANs, Deep Flow Inspection (DFI), or Deep Packet Inspection (DPI). But many application flows in current usages are dynamic, adaptive, time-bound, encrypted, peer-to-peer, asymmetric, used on multipurpose devices, and have different priorities depending on direction of flow, user preferences, and other factors. Any combination of these properties renders heuristic based techniques less effective and may

result in compromises to application security or user privacy, as described in detail in [[I-D.conet-aeon-problem-statement](#)].

Application enabled collaborative networking allows applications to explicitly signal their flow characteristics to the network. This provides network nodes with visibility of the application flow characteristics. These network nodes may take action based on this visibility and/or contribute to the flow description. The resulting flow description may be communicated as feedback from the network to applications. This proposes a way of building collaborative connections for network operators and application providers, benefiting both of them as well as users. Network provider is able to manage the traffic going through the network more effectively, and application provider utilizes action taken by network on its traffic to meet user requirement and expectation.

This document describes a set of use cases addressable by application enabled collaborative networking.

2. Use Cases

The following use cases have been identified.

1. Firewall Traversal: Identification of new applications
2. Efficient Capacity Usage
3. Video Adaptation
4. Multi-interface selection: Use metadata to help interface selection or prioritization.
5. Session Identification: Identification of multiple media flows belonging to a common application session.
6. Content Based Charging

In describing each use case, the following information is provided.

- o description of the problem
- o proposed solution
- o user/application benefit
- o operator benefit
- o flow characteristics provided by application

- o action taken by network as result of receiving flow characteristics
- o feedback provided by network
- o security and privacy considerations

2.1. Firewall Traversal: Identification of new applications

2.1.1. Description of Problem

Modern firewalls use application-layer gateways (ALGs) to perform policy enforcement. For example firewalls implement SIP-aware Application Layer Gateway function, which examines the SIP signaling and opens the appropriate pinholes for the RTP media. In particular firewall extracts media transport addresses, transport protocol and ports from session description and creates a dynamic mapping for media to flow through. This model will not work in the following cases:

1. Session signaling is end-to-end encrypted (say, using TLS).
2. Firewall does not understand the session signaling protocol, or extensions to the protocol, used by the endpoints (e.g. WebRTC signaling protocols).
3. Session signaling and media traverse different firewalls (e.g., signaling exits a network via one firewall whereas media exits a network via a different firewall).

Enterprise networks that use firewalls with restrictive policies block new applications like WebRTC and delay deployment of killer applications.

2.1.2. Proposed Solution

These problems can be addressed by the host providing authorization it received from an application server that is trusted by the network to authorize flows and associated actions (e.g., policies). PCP third party authorization ([\[I-D.wing-pcp-third-party-authz\]](#)) solves this problem by associating the media session with the signaling session. This is done by sending a cryptographic token in the signaling which authorizes the firewall mapping for the media session.

2.1.3. User/Application Benefit

Enterprise networks that use firewalls with restrictive policies can deploy new applications at a faster rate for user benefit.

2.1.4. Operator Benefit

Enterprise firewalls can enforce restrictive policies without the need to be enhanced to perform ALG on new applications. For example Enterprise firewall could have granular policies to permit peer-to-peer UDP media session only when the call is initiated using the selected WebRTC server (Dr. Good) it trusts and block others (Dr. Evil). PCP-aware firewalls can enforce such granular security policies without performing ALG on the session signaling protocols. This mechanism can be used by any other Application Function trusted by the network to permit time-bound, encrypted, peer-to-peer traffic.

2.1.5. Flow characteristics provided by application

The client requests dynamic mappings to permit flows required by the application. This request includes a cryptographic token and characteristics of the flow, such as the anticipated bandwidth needs as well as the tolerance to delay, loss, and jitter.

2.1.6. Action taken by firewall as result of receiving flow characteristics

The firewall uses the client request to permit and prioritize the traffic associated with those flows. The cryptographic token provides authorization for the flows and their prioritization.

2.1.7. Feedback provided by firewall

Firewall matches the authorization data with what is requested in the request sent by the client. If the authorization sets match, the firewall processes the request made by the client. If the token is invalid or the request exceeds what is authorized by the token then firewall rejects the request.

2.1.8. Security and Privacy Considerations

2.2. Efficient Capacity Usage

2.2.1. Description of Problem

Network traffic is bursty and often follows diurnal usage patterns such that there are times of day where traffic levels are at a peak, and other times of day where they are at a valley. Networks that are

properly capacity planned need to have enough capacity to service the traffic demands at peak. In a network with consistent demand and usage patterns, keeping up with demand is a matter of building capacity at a faster rate than the growth of the peak, in conjunction with any requirements for diversity and fault tolerance as well as any SLA for performance (latency, jitter, packet loss, etc) that may inform which traffic is passed, which is prioritized, and which may be dropped during periods of congestion. However, there are several problems to consider in this context:

1. Simply building enough capacity for peak usage is not always efficient and cost-effective, because not all traffic is the same in terms of its need to transit the network at the exact moment it is currently doing so, but few tools exist to provide applications with the information they need to make more intelligent decisions on demand, and thus they default to "as soon as possible." For example, those watching streaming video or doing real time communication or head to head gaming need immediate access, while less real-time activities such as data synchronization with the cloud for backups, or downloading software updates, or preloading content onto a CDN could potentially be deferred to times when more capacity is available, but today, all of that traffic competes for the same capacity at the same time.
2. QoS is not a substitute for capacity, and often a network designed for long periods of congested operation provides a poor user experience, since QoS ultimately is a method to identify which traffic should be dropped first.
3. When the network is not at peak usage, there is capacity sitting idle. Even in a well-used network capacity is built in increments that may not match up with growth rate i.e. if a network adds capacity in increments of 10G or 100G, but only needed a small fraction of that until growth catches up. This inefficiency is magnified when one considers the spare capacity designed into most networks to address the need to tolerate one or more failures in the network with minimal traffic impact. In many cases, the idle capacity even at peak may be up to 50%, and at off peak, it could be much higher.
4. Few networks have truly consistent demand and usage patterns. While the average usage may follow a rough pattern, this does not always provide for flash demands, where a large number of users are simultaneously downloading an OS update, or all watching the same event via streaming video, or more heavily using the network due to being stuck at home during a snowstorm, etc. The average usage patterns also do not take into account the effects of

outages at shifting large volumes of traffic around in the network, and so managing these exception events either requires further spare capacity, or acknowledgement that some traffic will be dropped due to congestion, with the attendant impact to end user experience.

2.2.2. Proposed Solution

Addressing this problem requires a multi-part solution:

- o Provide a mechanism for the network to communicate to applications when the network is busy and when it is not so that individual applications can manage their demand based on the nature of the application and its needs. This demand management helps to smooth the traffic at peak by redirecting some of the demand to off-peak, and has an analog in the power utility industry where demand based pricing or smart grid technology signals devices that use a large amount of power so that they can be intelligent about those demands and reduce the burden on the available capacity of the electrical grid.
- o Similarly, provide a means for applications to communicate their required performance envelope, as well as any data on how flexible the time of data transmission can be, i.e. "I need this transfer to complete by \$time on \$day" or "I need this transfer to complete within 12 hours" etc. This information can be used by the network to compute the best way to deliver the requested service, or to identify when it cannot provide the request and suggest an alternate.
- o Provide a means for "below best effort" or scavenger class data transmission so that traffic marked as scavenger will be carried in periods where no congestion is present, but may be discarded during periods of congestion due to either peak usage or outages.

This solution could also be used in conjunction with defined paths through the network (TE, Segment Routing, etc) to provide capacity for traffic that has specific performance requirements, or is not sensitive to using a sub-optimal path. i.e. capacity exists on this backup path that is much longer, so since this traffic does not care about a few 10s of milliseconds of additional latency, it should be marked to use the non-optimal path even if that path is not seen as best by the routing protocol.

2.2.3. User/Application Benefit

Key user benefits include:

- o Best service for real-time and other interactive applications (less interference from non real-time or non-interactive traffic)
- o More control over application bandwidth usage, potential for service guarantees for important applications

2.2.4. Operator Benefit

Reduced cost via better/more efficient management of capacity/growth while still providing first-class service to customers.

2.2.5. Flow characteristics provided by application

An application signals one or more of the following to the network:

- o level of service required (e.g. guaranteed service, best-effort, or below best effort)
- o minimum requirement for transmission rate/throughput
- o that it is tolerant/intolerant of high latency, high jitter, high packet loss
- o a request in the form "I need to deliver this data by X, when should I send, and how should I identify the flow?"

2.2.6. Action taken by network as result of receiving flow characteristics

Potential action taken by the network include:

- o Identify path through network that meets flow service requirements
- o Treat marked traffic according to identified service type (e.g. scavenger class carried in periods of low usage, and/or dropped during congestion)

2.2.7. Feedback provided by network

Feedback provided by the network includes:

- o Peak demand times, either proactively (e.g. this network peaks daily between the hours of X and Y) or reactively through

something like Explicit Congestion Notification (this network is at peak or is experiencing congestion right now)

- o ACK/NACK for requested level of service, throughput, etc.

2.2.8. Security and Privacy Considerations

This requires a trust model between application and network so that the information communicated about performance envelope requirements can be trusted. In the case where there are different costs, charging rates, tonnage limits by type of traffic, there is opportunity for abuse (maliciously marking all traffic such that it incurs additional cost, or such that it is dropped when it should not be, etc). Even simpler data such as IP Precedence is often remarked at the boundaries between networks as untrusted, so carrying this sort of metadata likely requires a method to ensure that it was set by a trusted entity and not manipulated in transit.

2.3. Video Adaptation

HTTP Adaptive Streaming (HAS) is an umbrella term for various HTTP-based streaming technologies that allow a client to adaptively switch between multiple bitrates, depending on current network conditions. HAS client first requests and receives a Manifest File, and then, after parsing the information in the Manifest File, proceeds with sequentially requesting the chunks listed in the Manifest File.

2.3.1. Description of Problem

The problems with HAS are:

- o HAS client selects the initial bitrate without knowing the current network conditions which could cause start-up delay and frame freezes while a lower bitrate chunk is being retrieved. HAS client does not have a mechanism to signal the flow characteristics and desired treatment to the network.
- o HAS server can mark the packets appropriately but setting DSCP has limitations. DSCP value may not be preserved or honored over the Internet and operating system may not allow to set DSCP values.
- o Content Providers may need a mechanism to convey the flow characteristics and desired treatment to the ISP. Existing mechanisms and the associated limitations are:
 1. ISP can be configured with the IP addresses of content providers to identify the traffic originating from those servers. The limitations with this approach are ISP has to

keep track of content providers IP addresses. With CDNI (Content Delivery Network InterConnection) content could be served either from uCDN (upstream CDN) or any of a number of dCDNs (downstream CDN) and it will not be possible to manually track the IP addresses of all the CDN surrogates. There is also no way to differentiate content which could be available in different bitrates.

2. If HAS client is behind NAT and content provider uses RESTful API (OneAPI) to install differentiated QoS then ISP will struggle to find the pre-NAT information. Content provider also needs to be aware of the ISP to which the client is attached and the IP address of the Policy Decision Point (PDP) in the ISP to which it needs to signal the flow characteristics.
- o ISP can use DPI to identify one-way video streaming content but is expensive and fails if the traffic is encrypted.

2.3.2. Proposed Solution

HAS client can use third party authorization to request network resources. At a high level, this authorization works by the client first obtaining a cryptographic token from the authorizing network element, then including that token in the request along with relevant flow characteristics. ISP validates the token and grants the request.

2.3.3. User/Application Benefit

This solution helps increase the average play quality, reduces the start-up delay and frame freezes by avoiding attempt to retrieve a too high-bit rate chunk etc thus improving the quality of experience for end user.

2.3.4. Operator Benefit

Network operators can better recognize and treat one-way video streaming content.

2.3.5. Flow characteristics provided by application

HAS client signals the flow characteristics such as the anticipated bandwidth needs as well as the tolerance to delay, loss, and jitter.

2.3.6. Action taken by network as result of receiving flow characteristics

Subject to local policies, a network node might perform bandwidth counting, or reconfigure the underlying network so that additional bandwidth is made available for this particular flow, or might perform other actions.

2.3.7. Feedback provided by network

The network responds that the client request can be fully or partially accommodated. It also notifies the client when conditions change.

2.3.8. Security and Privacy Considerations

2.4. Multi-interface selection: Use metadata to help interface selection or prioritization

2.4.1. Description of Problem

An increasing number of hosts are operating in multiple-interface environments and a host with multiple interfaces needs to choose the best interface for communication. Oftentimes, this decision is based on a static configuration and does not consider the link characteristics of that interface, which may affect the user experience. The network interfaces may have different link characteristics, but that will not be known without the awareness of the upstream and downstream characteristics of the access link.

2.4.2. Proposed Solution

The problem can be solved if a mechanism is provided for the applications to communicate required flow characteristics with the available interfaces, and know about network condition of each interface, or to what extent application requirement of flow characteristics can be met by each interface. Application can then prioritize the interfaces based on information gathered and select one or more interfaces that best meet its requirement.

2.4.3. User/Application Benefit

Applications can choose the interface that best meets their requirements for communication. User experience is improved because of the consistency between flow characteristics requested by application and network ability provided by the selected interface.

2.4.4. Operator Benefit

The network that can provide the requested flow characteristics will be selected by the application thus increasing the subscriber base of the operator.

2.4.5. Flow characteristics provided by application

Application signals flow characteristics over multiple interfaces and based on the response from its various interfaces sorts the source addresses according to the link capacity characteristics. Source addresses from the interface which best fulfills the desired flow characteristics are assigned the highest priority and would be tried first to communicate with the server or remote peer. For example [[I-D.reddy-mmusic-ice-best-interface-pcp](#)] explains the mechanism where Interactive Connectivity Establishment (ICE) agent on a host with multiple interfaces determines the link characteristics of the host's interfaces, which influences the ICE candidate priority. Similarly [[I-D.wing-mptcp-pcp](#)] explains how Multipath TCP (MPTCP) can select the best path when multiple paths are available.

2.4.6. Action taken by network as result of receiving flow characteristics

Network identifies flow characteristics requested by applications, and decides whether the request can be met or not.

2.4.7. Feedback provided by network

Link characteristics and ACK/NACK for flow requirement can be provided as feedback by network.

2.4.8. Security and Privacy Considerations

Users/applications are expected to consider security of interfaces, e.g. an untrusted public wifi access point will have lower priority than a trusted VPN tunnel, when prioritizing and selecting the interfaces.

2.5. Session Identification: Identification of multiple media flows belonging to a common application session

2.5.1. Description of Problem

Many end-to-end application sessions involve multiple application protocols, devices and administrative domains. These sessions involve multiple media flows (e.g. an audio flow and a video flow for a video call, media flows between different entities in a

supplementary service session consisting of multiple SIP dialogs or H.323 calls). Media flows may be added/removed from a application session during the lifetime of the session. From within the network, determining which media flows are associated with each application session is often difficult, making it hard to provide application level troubleshooting, traffic analysis, and QoS.

2.5.2. Proposed Solution

Including a session identifier (e.g. as defined in [[RFC7206](#)]) in the flow characteristics communicated by the application to the network would allow the network to identify media flows belonging to a common application session. This visibility would enable the following:

- o network troubleshooting and traffic analysis tools to correctly associate media flows with application sessions
- o media flows that are part of established application sessions to be identified (e.g. the triggered call in the case of a transfer) and given dedicated QoS. Preserving established sessions generally is higher priority than setting up new sessions (except when there is some form of multi-level preemption). Giving more bandwidth to additional flows on established sessions might cause some newer sessions to fail due to resource unavailability while established sessions continue without degradation, which is the preferred outcome in most cases.

2.5.3. User/Application Benefit

Users receive more predictable and reliable QoS for their application sessions.

2.5.4. Operator Benefit

Operators are able to perform traffic analysis and troubleshooting at the application level, and they are able to provide QoS at the application level rather than only at the media flow level.

2.5.5. Flow characteristics provided by application

The application provides a common session id as metadata for all its media flows throughout the lifetime of the session.

2.5.6. Action taken by network as result of receiving flow characteristics

The network identifies all media flows associated with a given session. This information may be used to provide application level QoS, preserving established sessions and/or giving more bandwidth to additional flows on established sessions.

2.5.7. Feedback provided by network

The network may provide feedback to the application indicating the amount of bandwidth it expects to be able to provide for its session. It may also be provide indications of the expected amount of delay, jitter, and loss the application should be prepared to tolerate.

2.5.8. Security and Privacy Considerations

2.6. Content Based Charging

2.6.1. Description of Problem

Commonly used billing method for internet subscribers, e.g. volume based charging, does not distinguish usage from the angle of applications. Under this billing model ISP cannot apply different pricing strategies to the applications it carries, users may hesitate to use certain types of applications, e.g. mobile apps consuming large volume, and application developers also have to strive for volume apart from user preference and usage time. Content based charging is an emerging billing method that takes content related information into account and enables smarter pricing strategy. Operators can place different prices for different types of traffic, and help content providers build tight relationship with their users, e.g. wholesaling the data volume of an application to its developer content provider so that users can use the application free of charge. Content based charging is required to precisely identify traffic that belongs to a certain pricing category in a way that is flexible and easy to manage, and granularity of traffic may range from types of applications to a detailed service function within an application. Those requirements reflex limitations in the current heuristics.

2.6.2. Proposed Solution

In order to address this problem a mechanism is needed to allow applications to notify its existence to the network by describing traffic flows needing a certain charging method. Network will then have direct visibility into the traffic and identify targeted traffic accordingly. This billing model usually involves collaboration

between network and content providers. The notification needs to go through an authentication function to guarantee the application is reliable and probably an identifier for network to identify the application that has service agreement with ISP. ISP will identify traffic based on characteristics notified by application and apply designated billing strategy.

2.6.3. User/Application Benefit

Users take advantage of the granular and customized charging model, and pay for different types of traffic at different rate. This charging model will reduce volume expense for users and stimulate internet usage. Content provider can benefit from providing users with exclusive payment function, e.g. pay for traffic volume or provide cheap volume package, and increase user enthusiasm and time to use its applications.

2.6.4. Operator Benefit

The solution will provide operators with a method to precisely identify and charge traffic based on content, and to agilely manage charging strategy of applications. Operators are able to cooperate with content providers to provide this new billing service to users and encourage encourage traffic consumption.

2.6.5. Flow characteristics provided by application

Application notifies network of its identifier and traffic description to enable network to recognize its traffic accordingly. Application may also signal its intended charging model as a request to the network.

2.6.6. Action taken by network as result of receiving flow characteristics

Network identifies traffic flows notified by the application and applies the designated billing model based on application request and business agreement with the content provider providing the application.

2.6.7. Feedback provided by network

2.6.8. Security and Privacy Considerations

There needs to be an authentication mechanism so as to ensure that traffic characteristics provider is right the authorized application the ISP has the charging agreement with.

3. Acknowledgements

The authors thank the attendees of the Bar BoF for contributing towards this set of use cases.

4. Informative References

[I-D.conet-aeon-problem-statement]

Fan, P., Deng, H., Boucadair, M., Reddy, T., and C. Eckel, "Application Enabled Collaborative Networking: Problem Statement and Requirements", [draft-conet-aeon-problem-statement-00](#) (work in progress), May 2014.

[I-D.reddy-mmusic-ice-best-interface-pcp]

Reddy, T., Wing, D., Steeg, B., Penno, R., and V. Varun, "Improving ICE Interface Selection Using Port Control Protocol (PCP) Flow Extension", [draft-reddy-mmusic-ice-best-interface-pcp-00](#) (work in progress), October 2013.

[I-D.wing-mptcp-pcp]

Wing, D., R, R., Reddy, T., Ford, A., and R. Penno, "Multipath TCP (MPTCP) Path Selection using PCP", [draft-wing-mptcp-pcp-00](#) (work in progress), October 2013.

[I-D.wing-pcp-third-party-authz]

Wing, D., Reddy, T., Patil, P., and R. Penno, "PCP Extension for Third Party Authorization", [draft-wing-pcp-third-party-authz-03](#) (work in progress), April 2014.

[RFC7206] Jones, P., Salgueiro, G., Polk, J., Liess, L., and H. Kaplan, "Requirements for an End-to-End Session Identification in IP-Based Multimedia Communication Networks", [RFC 7206](#), May 2014.

Authors' Addresses

Wesley George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
US

Email: wesley.george@twcable.com

Peng Fan
China Mobile
32 Xuanwumen West Street
Beijing 100053
China

Email: fanpeng@chinamobile.com

Satoru Matsushima
SoftBank Telecom
1-9-1 Higashi-Shinbashi, Munato-ku
Tokyo
Japan

Email: satoru.matsushima@g.softbank.co.jp

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tiredddy@cisco.com

Charles Eckel
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
US

Email: eckelcu@cisco.com

