

Internet Engineering Task Force (IETF)  
Request for Comments: 7778  
Category: Informational  
ISSN: 2070-1721

D. Kutscher  
F. Mir  
R. Winter  
NEC  
S. Krishnan  
Ericsson  
Y. Zhang  
Hewlett Packard Labs  
CJ. Bernardos  
UC3M  
March 2016

## Mobile Communication Congestion Exposure Scenario

### Abstract

This memo describes a mobile communications use case for congestion exposure (ConEx) with a particular focus on those mobile communication networks that are architecturally similar to the 3GPP Evolved Packet System (EPS). This memo provides a brief overview of the architecture of these networks (both access and core networks) and current QoS mechanisms and then discusses how congestion exposure concepts could be applied. Based on this discussion, this memo suggests a set of requirements for ConEx mechanisms that particularly apply to these mobile networks.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7778>.

## Copyright Notice

Copyright (c) 2016 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. Acronyms . . . . .	4
2. ConEx Use Cases in Mobile Communication Networks . . . . .	4
2.1. ConEx as a Basis for Traffic Management . . . . .	5
2.2. ConEx to Incentivize Scavenger Transports . . . . .	7
2.3. Accounting for Congestion Volume . . . . .	7
2.4. Partial vs. Full Deployment . . . . .	8
2.5. Summary . . . . .	9
3. ConEx in the EPS . . . . .	9
3.1. Possible Deployment Scenarios . . . . .	9
3.2. Implementing ConEx Functions in the EPS . . . . .	14
3.2.1. ConEx Protocol Mechanisms . . . . .	15
3.2.2. ConEx Functions in the Mobile Network . . . . .	15
4. Summary . . . . .	17
5. Security Considerations . . . . .	19
6. Informative References . . . . .	19
Appendix A. Overview of 3GPP's EPS . . . . .	22
Acknowledgements . . . . .	24
Authors' Addresses . . . . .	24

## 1. Introduction

Mobile data traffic continues to grow rapidly. The challenge wireless operators face is to support more subscribers with an increasing bandwidth demand. To meet these bandwidth requirements, there is a need for new technologies that assist the operators in efficiently utilizing the available network resources. Two specific areas where such new technologies could be deemed useful are resource allocation and flow management.

Analysis of data traffic in cellular networks has shown that most flows are short lived and low volume, but a comparatively small number of high-volume flows constitute a large fraction of the overall traffic volume [lte-sigcomm2013]. That means that potentially a small fraction of users is responsible for the majority of traffic in cellular networks. In view of such highly skewed user behavior and limited and expensive resources (e.g., the wireless spectrum), resource allocation and usage accountability are two important issues for operators to solve in order to achieve both a better network resource utilization and fair resource sharing. ConEx, as described in [RFC6789], is a technology that can be used to achieve these goals.

The ConEx mechanism is designed to be a general technology that could be applied as a key element of congestion management solutions for a variety of use cases. In particular, use cases that are of interest for initial deployment are those in which the end hosts and the network that contains the destination end hosts are ConEx-enabled but other networks need not be.

A specific example of such a use case can be a mobile communication network such as a 3GPP EPS networks where UEs (User Equipment) (i.e., mobile end hosts), servers and caches, the access network, and possibly an operator's core network can be ConEx-enabled; that is, hosts support the ConEx mechanisms, and the network provides policing/auditing functions at its edges.

This document provides a brief overview of the architecture of such networks (access and core networks) and current QoS mechanisms. It further discusses how such networks can benefit from congestion exposure concepts and how they should be applied. Using this use case as a basis, a set of requirements for ConEx mechanisms are described.

### 1.1. Acronyms

In this section, we expand some acronyms that are used throughout the text. Most are explained and put in a system context in Appendix A and the 3GPP, ECN, and ConEx specifications referenced there.

eNB

Evolved NodeB: LTE base station

HSS

Home Subscriber Server

S-GW

Serving Gateway: mobility anchor and tunnel endpoint

P-GW

Packet Data Network (PDN) Gateway: tunnel endpoint for user-plane and control-plane protocols -- typically the GW to the Internet or an operator's service network

UE

User Equipment: mobile terminals

GTP

GPRS Tunneling Protocol [TS29060]

GTP-U

GTP User Data Tunneling [TS29060]

GTP-C

GTP Control [TS29060]

## 2. ConEx Use Cases in Mobile Communication Networks

In general, quality of service and good network resource utilization are important requirements for mobile communication network operators. Radio access and backhaul capacity are considered scarce resources, and bandwidth (and radio resource) demand is difficult to predict precisely due to user mobility, radio propagation effects, etc. Hence, today's architectures and protocols go to significant lengths in order to provide network-controlled quality of service. These efforts often lead to complexity and cost. ConEx could be a simpler and more capable approach to efficient resource sharing in these networks.

In the following sections, we discuss ways that congestion exposure could be beneficial for supporting resource management in such mobile communication networks. [RFC6789] describes fundamental congestion exposure concepts and a set of use cases for applying congestion exposure mechanisms to realize different traffic management functions such as flow policy-based traffic management or traffic offloading. Readers that are not familiar with the 3GPP EPS should refer to Appendix A first.

## 2.1. ConEx as a Basis for Traffic Management

Traffic management is a very important function in mobile communication networks. Since wireless resources are considered scarce and since user mobility and shared bandwidth in the wireless access create certain dynamics with respect to available bandwidth, commercially operated mobile networks provide mechanisms for tight resource management (admission control for bearer establishment). However, sometimes these mechanisms are not easily applicable to IP- and HTTP-dominated traffic mixes; for example, most Internet traffic in today's mobile network is transmitted over the (best-effort) default bearer.

Given the above, and in the light of the significant increase of overall data volume in 3G networks, Deep Packet Inspection (DPI) is often considered a desirable function to have in the Evolved Packet Core (EPC) -- despite its cost and complexity. However, with the increase of encrypted data traffic, traffic management using DPI alone will become even more challenging.

Congestion exposure can be employed to address resource management requirements in different ways:

1. It can enable or enhance flow policy-based traffic management. At present, DPI-based resource management is often used to prioritize certain application classes with respect to others in overload situations, so that more users can be served effectively on the network. In overload situations, operators use DPI to identify dispensable flows and make them yield to other flows (of different application classes) through policing. Such traffic management is thus based on operator decisions, using partly static configuration and some estimation about the future per-flow bandwidth demand. With congestion exposure, it would be possible to assess the contribution to congestion of individual flows. This information can then be used as input to a policer that can optimize network utilization more accurately and dynamically. By using ConEx congestion contribution as a metric, such policers would not need to be aware of specific link loads (e.g., in wireless base stations) or flow application types.

2. It can reduce the need for complex DPI by allowing for a bulk packet traffic management system that does not have to consider either the application classes flows belong to or the individual sessions. Instead, traffic management would be based on the current cost (contribution to congestion) incurred by different flows and enable operators to apply policing/accounting depending on their preference. Such traffic management would be simpler and more robust (no real-time flow application type identification required, no static configuration of application classes); it would also perform better as decisions can be made based on real-time actual cost contribution. With ConEx, accurate downstream path information would be visible to ingress network operators, which can respond to incipient congestion in time. This can be equivalent to offering different levels of QoS, e.g., premium service with zero congestion response. For that, ConEx could be used in two different ways:
  - A. as additional information to assist network functions to impose different QoS for different application sessions; and
  - B. as a tool to let applications decide on their response to congestion notification while incentivizing them to react (in general) appropriately, e.g., by enforcing overall limits for congestion contribution or by accounting and charging for such congestion contribution. Note that this level of responsiveness would be on a different level than, say, application-layer responsiveness in protocols such as Dynamic Adaptive Streaming over HTTP (DASH) [dash]; however, it could interwork with such protocols, for example, by triggering earlier responses.
3. It can further be used to more effectively trigger the offload of selected traffic to a non-3GPP network. Nowadays, it is common that users are equipped with dual-mode mobile phones (e.g., integrating third/fourth generation cellular and Wi-Fi radio devices) capable of attaching to available networks either sequentially or simultaneously. With this scenario in mind, 3GPP is currently looking at mechanisms to seamlessly and selectively switch over a single IP flow (e.g., user application) to a different radio access while keeping all other ongoing connections untouched. The decision on when and which IP flows move is typically based on statically configured rules, whereas the use of ConEx mechanisms could also factor real-time congestion information into the decision.

In summary, it can be said that traffic management in the 3GPP EPS and other mobile communication architectures is very important. Currently, more static approaches based on admission control and

static QoS are in use, but recently, there has been a perceived need for more dynamic mechanisms based on DPI. Introducing ConEx could make these mechanisms more efficient or even remove the need for some of the DPI functions deployed today.

## 2.2. ConEx to Incentivize Scavenger Transports

3G and LTE networks are turning into universal access networks that are shared between mobile (smart) phone users, mobile users with laptop PCs, home users with LTE access, and others. Capacity sharing among different users and application flows becomes increasingly important in these mobile communication networks.

Most of this traffic is likely to be classified as best-effort traffic without differentiating, for example, periodic OS updates and application store downloads from web-based (i.e., browser-based) communication or other real-time communication. For many of the bulk data transfers, completion times are not important within certain bounds; therefore, if scavenger transports (or transports that are less than best effort) such as Low Extra Delay Background Transport (LEDBAT) [RFC6817] were used, it would improve the overall utility of the network. The use of these transports by the end user, however, needs to be incentivized. ConEx could be used to build an incentive scheme, e.g., by giving a larger bandwidth allowance to users that contribute less to congestion or lowering the next monthly subscription fee. In principle, this would be possible to implement with current specifications.

## 2.3. Accounting for Congestion Volume

3G and LTE networks provide extensive support for accounting and charging already, for example, see the Policy Charging Control (PCC) architecture [TS23203]. In fact, most operators today account transmitted data volume on a very fine granular basis and either correlate monthly charging to the exact number of packets/bytes transmitted or employ some form of flat rate (or flexible flat rate), often with a so-called fair-use policy. With such policies, users are typically limited to an administratively configured maximum bandwidth limit after they have used up their contractual data volume budget for the charging period.

Changing this data from volume-based accounting to congestion-based accounting would be possible in principle, especially since there already is an elaborate per-user accounting system available. Also, an operator-provided mobile communication network can be seen as a network domain that would allow for such congestion volume accounting. This would not require any support from the global Internet, especially since the typical scarce resources such as the

wireless access and the mobile backhaul are all within this domain. Traffic normally leaves/enters the operator's network via well-defined egress/ingress points that would be ideal candidates for policing functions. Moreover, in most commercially operated networks, accounting is performed for both received and sent data, which would facilitate congestion volume accounting as well.

With respect to the current Path Computation Client (PCC) framework, accounting for congestion volume could be added as another feature to the "Usage Monitoring Control" capability that is currently based on data volume. This would not require a new interface (reference points) at all.

#### 2.4. Partial vs. Full Deployment

In general, ConEx lends itself to partial deployment as the mechanism does not require all routers and hosts to support congestion exposure. Moreover, assuming a policing infrastructure has been put in place, it is not required to modify all hosts. Since ConEx is about senders exposing congestion contribution to the network, senders need to be made ConEx-aware (assuming a congestion notification mechanism such as Explicit Congestion Notification (ECN) is in place).

When moving towards full deployment in a specific operator's network, different ways for introducing ConEx support on UEs are feasible. Since mobile communication networks are multi-vendor networks, standardizing ConEx support on UEs (e.g., in 3GPP specifications) appears useful. Still, not all UEs would have to support ConEx, and operators would be free to choose their policing approach in such deployment scenarios. Leveraging existing PCC architectures, 3GPP network operators could, for example, decide policing/accounting approaches per UE -- i.e., apply fixed volume caps for non-ConEx UEs and more flexible schemes for ConEx-enabled UEs.

Moreover, it should be noted that network support for ConEx is a feature that some operators may choose to deploy if they wish, but it is not required that all operators (or all other networks) do so.

Depending on the extent of ConEx support, specific aspects such as roaming have to be taken into account, i.e., what happens when a user is roaming in a ConEx-enabled network but their UE is not ConEx-enabled and vice versa. Although these may not be fundamental problems, they need to be considered. For supporting mobility in general, it can be required to shift users' policing state during a handover. There is existing work on distributed rate limiting (see [raghavan2007]) and on specific optimizations (see [nec.euronf-2011]) for congestion exposure and policing in mobility scenarios.



Another aspect to consider is the addition of Selected IP Traffic Offload (SIPTO) and Local IP Access (LIPA) [TR23829]), i.e., the idea that some traffic such as high-volume Internet traffic is actually not passed through the EPC but is offloaded at a "break-out point" closer to (or in) the access network. On the other hand, ConEx can also enable more dynamic decisions on what traffic to actually offload by considering congestion exposure in bulk traffic aggregates, thus making traffic offload more effective.

## 2.5. Summary

In summary, the 3GPP EPS is a system architecture that can benefit from congestion exposure in multiple ways. Dynamic traffic and congestion management is an acknowledged and important requirement for the EPS; this is also illustrated by the current DPI-related work for EPS.

Moreover, networks such as an EPS mobile communication network would be quite amenable for deploying ConEx as a mechanism, since they represent clearly defined and well-separated operational domains in which local ConEx deployment would be possible. Aside from roaming (which needs to be considered for a specific solution), such a deployment is fully under the control of a single operator, which can enable operator-local enhancement without the need for major changes to the architecture.

In 3GPP EPS, interfaces between all elements of the architecture are subject to standardization, including UE interfaces and eNB interfaces, so that a more general approach, involving more than a single operator's network, can be feasible as well.

## 3. ConEx in the EPS

In this section, we discuss a few options for how such a mechanism (and possibly additional policing functions) could eventually be deployed in the 3GPP EPS. Note that this description of options is not intended to be a complete set of possible approaches; it merely discusses the most promising options.

### 3.1. Possible Deployment Scenarios

There are different possible ways for how ConEx functions on hosts and network elements can be used. For example, ConEx could be used for a limited part of the network only (e.g., for the access network), congestion exposure and sender adaptation could involve the mobile nodes or not, or, finally, the ConEx feedback loop could extend beyond a single operator's domain or not.

We present four different deployment scenarios for congestion exposure in the figures below:

1. In Figure 1, ConEx is supported by servers for sending data (web servers in the Internet and caches in an operator's network) but not by UEs (neither for receiving nor sending). An operator who chooses to run a policing function on the network ingress, e.g., on the P-GW, can still benefit from congestion exposure without requiring any change on UEs.
2. ConEx is universally employed between operators (as depicted in Figure 2) with an end-to-end ConEx feedback loop. Here, operators could still employ local policies, congestion accounting schemes, etc., and they could use information about congestion contribution for determining interconnection agreements. This deployment scenario would imply the willingness of operators to expose congestion to each other.
3. For Isolated ConEx domains as depicted in Figure 3, ConEx is solely applied locally in the operator network, and there is no end-to-end congestion exposure. This could be the case when ConEx is only implemented in a few networks or when operators decide to not expose ECN and account for congestion for inter-domain traffic. Independent of the actual scenario, it is likely that there will be border gateways (as in today's deployments) that are associated with policing and accounting functions.
4. [conex-lite] describes an approach called "ConEx Lite" for mobile networks that is intended for initial deployment of congestion exposure concepts in LTE, specifically in the backhaul and core network segments. As depicted in Figure 4, ConEx Lite allows a tunnel receiver to monitor the volume of bytes that has been lost, dropped, or ECN-CE (Congestion Experienced) marked between the tunnel sender and receiver. For that purpose, a new field called the Byte Sequence Marker (BSN) is introduced to the tunnel header to identify the byte in the flow of data from the tunnel sender to the tunnel receiver. A policer at the tunnel sender is expected to react according to the tunnel congestion volume (see [conex-lite] for details).




Figure 1: ConEx Support on Servers and Caches

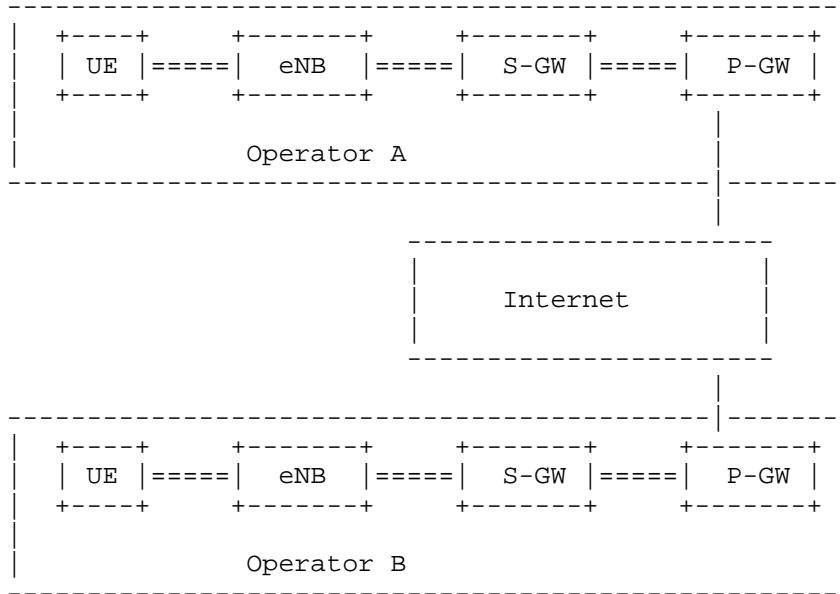


Figure 2: ConEx Deployment across Operator Domains

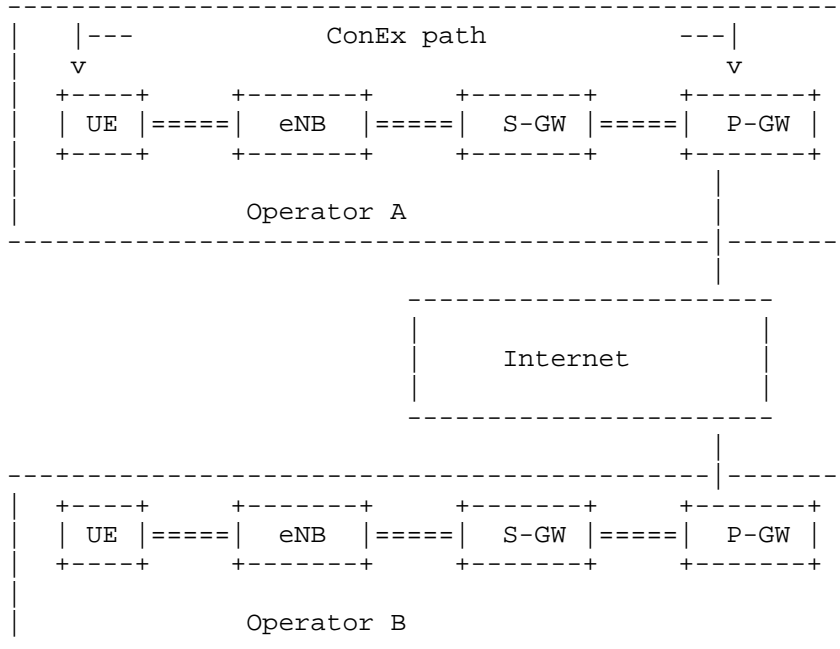


Figure 3: ConEx Deployment in a Single Operator Domain

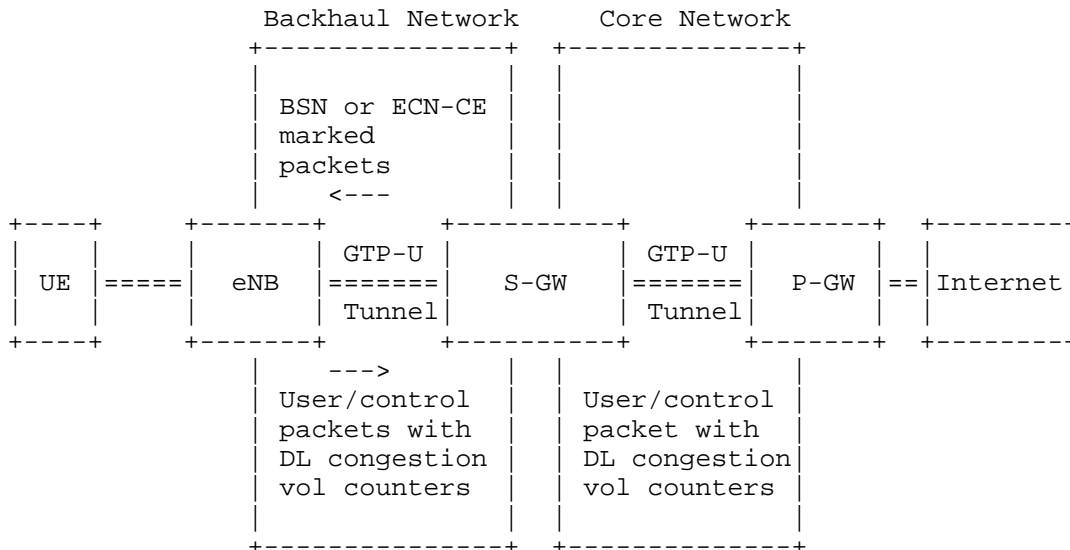


Figure 4: ConEx Lite Deployment

Note: DL stands for "downlink".

### 3.2. Implementing ConEx Functions in the EPS

We expect a ConEx solution to consist of different functions that should be considered when implementing congestion exposure in the 3GPP EPS. [RFC7713] describes the following congestion exposure components:

- o Modified senders that send congestion exposure information in response to congestion feedback.
- o Receivers that generate congestion feedback (leveraging existing behavior or requiring new functions).
- o Audit functions that audit ConEx signals against actual congestion, e.g., by monitoring flows or aggregate of flows.
- o Policy devices that monitor congestion exposure information and act on the flows according to the operator's policy.

Two aspects are important to consider: 1) how the ConEx protocol mechanisms would be implemented and what modifications to existing networks would be required, and 2) where ConEx functional entities would be placed best (to allow for a non-invasive addition). We discuss these two aspects in the following sections.

### 3.2.1. ConEx Protocol Mechanisms

The most important step in introducing ConEx (initially) is adding the congestion exposure functionality to senders. For an initial deployment, no further modification to senders and receivers would be required. Specifically, there is no fundamental dependency on ECN, i.e., ConEx can be introduced without requiring ECN to be implemented.

Congestion exposure information for IPv6 [CONEX-DESTOPT] is contained in a destination option header field, which requires minimal changes at senders and nodes that want to assess path congestion. The destination option header field does not affect non-ConEx nodes in a network.

In 3GPP networks, IP tunneling is used intensively, i.e., using either IP-in-GTP-U or Proxy Mobile IPv6 (PMIPv6) (i.e., IP-in-IP) tunnels. In general, the ConEx destination option of encapsulated packets should be made available for network nodes on the tunnel path, i.e., a tunnel ingress should copy the ConEx destination option field to the outer header.

For effective and efficient capacity sharing, we envisage the deployment of ECN in conjunction with ConEx so that ECN-enabled receivers and senders get more accurate and more timely information about the congestion contribution of their flows. ECN is already partially introduced into 3GPP networks: Section 11.6 in [TS36300] specifies the usage of ECN for congestion notification on the radio link (between eNB and UE), and [TS26114] specifies how this can be leveraged for voice codec adaptation. A complete, end-to-end support of ECN would require specification of tunneling behaviour, which should be based on [RFC6040] (for IP-in-IP tunnels). Specifically, a specification for tunneling ECN in GTP-U will be needed.

### 3.2.2. ConEx Functions in the Mobile Network

In this section, we discuss some possible placement strategies for ConEx functional entities (addressing both policing and auditing functions) in the EPS and for possible optimizations for both the uplink and the downlink.

In general, ConEx information (exposed congestion) is declared by a sender and remains unchanged on the path; hence, reading ConEx information (e.g., by policing functions) is placement-agnostic. Auditing ConEx normally requires assessing declared congestion contribution and current actual congestion. If the latter is, for example, done using ECN, such a function would best be placed at the end of the path.

In order to provide a comprehensive ConEx-based capacity management framework for the EPS, it would be advantageous to consider user contribution to congestion for both the radio access and the core network. For a non-invasive introduction of ConEx, it can be beneficial to combine ConEx functions with existing logical EPS entities. For example, potential places for ConEx policing and auditing functions would then be eNBs, S-GWs, or the P-GWs. Operator deployments may, of course, still provide additional intermediary ConEx-enabled IP network elements.

For a more specific discussion, it will be beneficial to distinguish downlink and uplink traffic directions (also see [nec.globecom2010] for a more detailed discussion). In today's networks and usage models, downlink traffic is dominating (also reflected by the asymmetric capacity provided by the LTE radio interface). That does not, however, imply that uplink congestion is not an issue, since the asymmetric maximum bandwidth configuration can create a smaller bottleneck for uplink traffic. There are, of course, backhaul links, gateways, etc., that could be overloaded as well.

For managing downlink traffic (e.g., in scenarios such as the one depicted in Figure 1), operators can have different requirements for policing traffic. Although policing is, in principle, location-agnostic, it is important to consider requirements related to the EPS architecture (Figure 5) such as tunneling between P-GWs and eNBs. Policing can require access to subscriber information (e.g., congestion contribution quota) or user-specific accounting, which suggests that the ConEx function could be co-located with the P-GW that already has an interface towards the Policy and Charging Rule Function (PCRF).

Still, policing can serve different purposes. For example, if the objective is to police bulk traffic induced by peer networks, additional monitoring functions can be placed directly at corresponding ingress points to monitor traffic and possibly drive out-of-band functions such as triggering border contract penalties.

The auditing function, which should be placed at the end of the path (at least after/at the last bottleneck), would likely be placed best on the eNB (wireless base station).

For the uplink direction, there are naturally different options for designing monitoring and policy enforcement functions. A likely approach can be to monitor congestion exposure on central gateway nodes (such as P-GWs) that provide the required interfaces to the PCRF but to perform policing actions in the access network (i.e., in eNBs). For example, the traffic is policed at the ingress before it reaches concentration points in the core network.



Such a setup would enable all the ConEx use cases described in Section 2 without requiring significant changes to the EPS architecture. It would also enable operators to re-use existing infrastructure, specifically wireless base stations, PCRF, and Home Subscriber Server (HSS) systems.

For ConEx functions on elements such as the S-GWs and P-GWs, it is important to consider mobility and tunneling protocol requirements. LTE provides two alternative approaches: PMIPv6 [TS23402] and the GPRS Tunneling Protocol (GTP). For the propagation of congestion information (responses), tunneling considerations are therefore very important.

In general, policing will be done based on per-user (per-subscriber) information such as congestion quota, current quota usage, etc., and network operator policies, e.g., specifying how to react to persistent congestion contribution. In the EPS, per-user information is normally part of the user profile (stored in the HSS) that would be accessed by PCC entities such as the PCRF for dynamic updates, enforcement, etc.

#### 4. Summary

We have shown how congestion exposure can be useful for efficient resource management in mobile communication networks. The premise for this discussion was the observation that data communication, specifically best-effort bulk data transmission, is becoming a commodity service, whereas resources are obviously still limited. This calls for efficient, scalable, and yet effective capacity sharing in such networks.

ConEx can be a mechanism that enables such capacity sharing while allowing operators to apply these mechanisms in different ways, e.g., for implementing different use cases as described in Section 2. It is important to note that ConEx is fundamentally a mechanism that can be applied in different ways to realize the policies of different operators.

ConEx may also be used to complement 3GPP-based mechanisms for congestion management that are currently under development, such as in the User Plane Congestion Management (UPCON) work item described in [TR23705].

We have described a few possibilities for adding ConEx as a mechanism to 3GPP LTE-based networks and have shown how this could be done incrementally (starting with partial deployment). It is quite feasible that such partial deployments be done on a per-operator-domain basis without requiring changes to standard 3GPP interfaces.

For network-wide deployment, e.g., with congestion exposure between operators, more considerations might be needed.

We have also identified a few implications/requirements that should be taken into consideration when enabling congestion exposure in such networks:

**Performance:** In mobile communication networks with more expensive resources and more stringent QoS requirements, the feasibility of applying ConEx as well as its performance and deployment scenarios need to be examined closer. For instance, a mobile communication network may encounter longer delay and higher loss rates, which can impose specific requirements on the timeliness and accuracy of congestion exposure information.

**Mobility:** One of the unique characteristics of cellular networks when compared to wired networks is the presence of user mobility. As the user location changes, the same device can be connected to the network via different base stations (eNBs) or even go through switching gateways. Thus, the ConEx scheme must be able to carry the latest congestion information per user/flow across multiple network nodes in real time.

**Multi-access:** In cellular networks, multiple access technologies can co-exist. In such cases, a user can use multiple access technologies for multiple applications or even a single application simultaneously. If the congestion policies are set based on each user, then ConEx should have the capability to enable information exchange across multiple access domains.

**Tunneling:** Both 3G and LTE networks make extensive usage of tunneling. The ConEx mechanism should be designed in a way to support usage with different tunneling protocols such as PMIPv6 and GTP. For ECN-based congestion notification, [RFC6040] specifies how the ECN field of the IP header should be constructed on entry and exit from IP-in-IP tunnels.

**Roaming:** Independent of the specific architecture, mobile communication networks typically differentiate between non-roaming and roaming scenarios. Roaming scenarios are typically more demanding regarding implementing operator policies, charging, etc. It can be expected that this would also hold for deploying ConEx. A more detailed analysis of this problem will be provided in a future revision of this document.

It is important to note that ConEx is intended to be used as a supplement and not a replacement to the existing QoS mechanisms in mobile networks. For example, ConEx deployed in 3GPP mobile networks

can provide useful input to the existing 3GPP PCC mechanisms by supplying more dynamic network information to supplement the fairly static information used by the PCC. This would enable the mobile network to make better policy control decisions than is possible with only static information.

## 5. Security Considerations

For any ConEx deployment, it is important to apply appropriate mechanisms to preclude applications and senders from misstating their congestion contribution. [RFC7713] discusses this problem in detail and introduces the ConEx auditing concept. ConEx auditing can be performed in different ways -- for example, flows can be constantly audited or only audited on demand when network operators decide to do so. Also, coarse-grained auditing may operate on flow aggregates for efficiency reasons, whereas fine-grained auditing would inspect individual flows. In mobile networks, there may be deployment strategies that favor efficiency over very exact auditing. It is important to understand the trade-offs and to apply ConEx auditing appropriately.

The ConEx protocol specifications [CONEX-DESTOPT] and [TCP-MOD] discuss additional security considerations that would also apply to mobile network deployments.

## 6. Informative References

[CONEX-DESTOPT]

Krishnan, S., Kuehlewind, M., Briscoe, B., and C. Ralli, "IPv6 Destination Option for Congestion Exposure (ConEx)", Work in Progress, draft-ietf-conex-destopt-12, January 2016.

[conex-lite]

Baillargeon, S. and I. Johansson, "ConEx Lite for Mobile Networks", In Proceedings of the 2014 ACM SIGCOMM Capacity Sharing Workshop, DOI 10.1145/2630088.2630091, August 2014.

[dash]

ISO/IEC, "Information Technology -- Dynamic Adaptive Streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats", ISO/IEC 23009-1:2014, May 2014.

- [lte-sigcomm2013]  
Huang, J., Qian, F., Guo, Y., Zhou, Y., Xu, Q., Mao, Z., Sen, S., and O. Spatscheck, "An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance", In Proceedings of the 2013 ACM SIGCOMM Conference, DOI 10.1145/2486001.2486006, August 2013.
- [nec.euronf-2011]  
Mir, F., Kutscher, D., and M. Brunner, "Congestion Exposure in Mobility Scenarios", In Proceedings of the 7th Euro-NF Conference on Next Generation Internet (NGI), DOI 10.1109/NGI.2011.5985948, June 2011.
- [nec.globecom2010]  
Kutscher, D., Lundqvist, H., and F. Mir, "Congestion Exposure in Mobile Wireless Communications", In Proceedings of 2010 IEEE Global Telecommunications Conference (GLOBECOM), DOI 10.1109/GLOCOM.2010.5684362, December 2010.
- [raghavan2007]  
Raghavan, B., Vishwanath, K., Ramabhadran, S., Yocum, K., and A. Snoeren, "Cloud Control with Distributed Rate Limiting", ACM SIGCOMM Computer Communication Review, DOI 10.1145/1282427.1282419, October 2007.
- [RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, DOI 10.17487/RFC6040, November 2010, <<http://www.rfc-editor.org/info/rfc6040>>.
- [RFC6789] Briscoe, B., Ed., Woundy, R., Ed., and A. Cooper, Ed., "Congestion Exposure (ConEx) Concepts and Use Cases", RFC 6789, DOI 10.17487/RFC6789, December 2012, <<http://www.rfc-editor.org/info/rfc6789>>.
- [RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, DOI 10.17487/RFC6817, December 2012, <<http://www.rfc-editor.org/info/rfc6817>>.
- [RFC7713] Mathis, M. and B. Briscoe, "Congestion Exposure (ConEx) Concepts, Abstract Mechanism, and Requirements", RFC 7713, DOI 10.17487/RFC7713, December 2015, <<http://www.rfc-editor.org/info/rfc7713>>.
- [TCP-MOD] Kuehlewind, M. and R. Scheffenegger, "TCP modifications for Congestion Exposure", Work in Progress, draft-ietf-conex-tcp-modifications-10, October 2015.

- [TR23705] 3GPP, "System Enhancements for User Plane Congestion Management", 3GPP TR 23.705 13.0.0, December 2015.
- [TR23829] 3GPP, "Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO)", 3GPP TR 23.829 10.0.1, October 2011.
- [TS23203] 3GPP, "Policy and charging control architecture", 3GPP TS 23.203 13.6.0, December 2015.
- [TS23401] 3GPP, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access", 3GPP TS 23.401 13.5.0, December 2015.
- [TS23402] 3GPP, "Architecture enhancements for non-3GPP accesses", 3GPP TS 23.402 13.4.0, December 2015.
- [TS26114] 3GPP, "IP Multimedia Subsystem (IMS); Multimedia telephony; Media handling and interaction", 3GPP TS 26.114 13.2.0, December 2015.
- [TS29060] 3GPP, "General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface", 3GPP TS 29.060 13.3.0, December 2015.
- [TS29274] 3GPP, "3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3", 3GPP TS 29.274 13.4.0, December 2015.
- [TS36300] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2", 3GPP TS 36.300 13.2.0, January 2016.

## Appendix A. Overview of 3GPP's EPS

This section provides an overview of the 3GPP "Evolved Packet System" (EPS [TS36300] [TS23401]) as a specific example of a mobile communication architecture. Of course, other architectures exist, but the EPS is used as one example to demonstrate the applicability of congestion exposure concepts and mechanisms.

The EPS architecture and some of its standardized interfaces are depicted in Figure 5. The EPS provides IP connectivity to UE (i.e., mobile nodes) and access to operator services, such as global Internet access and voice communications. The EPS comprises the radio access network called Evolved Universal Terrestrial Radio Access Network (E-UTRAN) and the core network called the Evolved Packet Core (EPC). QoS is supported through an EPS bearer concept, providing bindings to resource reservation within the network.




Figure 5: EPS Architecture Overview (Roaming Case)

Note:

- HPLMN - Home Public Land Mobile Network
- VPLMN - Visited Public Land Mobile Network
- AF - Application Function
- SGi - Service Gateway Interface
- LTE-Uu - LTE Radio Interface

The Evolved NodeB (eNB), the LTE base station, is part of the access network that provides radio resource management, header compression, security, and connectivity to the core network through the S1 interface. In an LTE network, the control-plane signaling traffic and the data traffic are handled separately. The eNBs transmit the control traffic and data traffic separately via two logically separate interfaces.

The Home Subscriber Server (HSS) is a database that contains user subscriptions and QoS profiles. The Mobility Management Entity (MME) is responsible for mobility management, user authentication, bearer establishment and modification, and maintenance of the UE context.

The Serving Gateway (S-GW) is the mobility anchor and manages the user-plane data tunnels during the inter-eNB handovers. It tunnels all user data packets and buffers downlink IP packets destined for UEs that happen to be in idle mode.

The PDN Gateway (P-GW) is responsible for IP address allocation to the UE and is a tunnel endpoint for user-plane and control-plane protocols. It is also responsible for charging, packet filtering, and policy-based control of flows. It interconnects the mobile network to external IP networks, e.g., the Internet.

In this architecture, data packets are not sent directly on an IP network between the eNB and the gateways. Instead, every packet is tunneled over a tunneling protocol -- the GPRS Tunneling Protocol (GTP) [TS29060] over UDP/IP. A GTP path is identified in each node with the IP address and a UDP port number on the eNB/gateways. The GTP protocol carries both the data traffic (GTP-U tunnels) and the control traffic (GTP-C tunnels [TS29274]). Alternatively, PMIPv6 is used on the S5 interface between S-GW and P-GW.

The above is very different from an end-to-end path on the Internet where the packet forwarding is performed at the IP level. Importantly, we observe that these tunneling protocols give the operator a large degree of flexibility to control the congestion mechanism incorporated with the GTP/PMIPv6 protocols.

#### Acknowledgements

We would like to thank Bob Briscoe and Ingemar Johansson for their support in shaping the overall idea and in improving the document by providing constructive comments. We would also like to thank Andreas Maeder and Dirk Staehle for reviewing the document and for providing helpful comments.

#### Authors' Addresses

Dirk Kutscher  
NEC  
Kurfuersten-Anlage 36  
Heidelberg  
Germany

Email: kutscher@neclab.eu



Faisal Ghias Mir  
NEC  
Kurfuersten-Anlage 36  
Heidelberg  
Germany

Email: [faisal.mir@gmail.com](mailto:faisal.mir@gmail.com)

Rolf Winter  
NEC  
Kurfuersten-Anlage 36  
Heidelberg  
Germany

Email: [rolf.winter@neclab.eu](mailto:rolf.winter@neclab.eu)

Suresh Krishnan  
Ericsson  
8400 Blvd Decarie  
Town of Mount Royal, Quebec  
Canada

Email: [suresh.krishnan@ericsson.com](mailto:suresh.krishnan@ericsson.com)

Ying Zhang  
Hewlett Packard Labs  
3000 Hannover Street  
Palo Alto, CA 94304  
United States

Email: [ying.zhang13@hp.com](mailto:ying.zhang13@hp.com)

Carlos J. Bernardos  
Universidad Carlos III de Madrid  
Av. Universidad, 30  
Leganes, Madrid 28911  
Spain

Phone: +34 91624 6236  
Email: [cjbc@it.uc3m.es](mailto:cjbc@it.uc3m.es)  
URI: <http://www.it.uc3m.es/cjbc/>