ConEx Working Group
Internet-Draft
Intended status: Experimental
Expires: April 17, 2015

S. Krishnan
Ericsson
M. Kuehlewind
ETH Zurich
C. Ralli
Telefonica

October 14, 2014

IPv6 Destination Option for ConEx
draft-ietf-conex-destopt-07

Abstract

ConEx is a mechanism by which senders inform the network about the congestion encountered by packets earlier in the same flow. This document specifies an IPv6 destination option that is capable of carrying ConEx markings in IPv6 datagrams.

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 April 17, 2015.

Table of Contents

1. Introduction

ConEx [I-D.ietf-ConEx-abstract-mech] is a mechanism by which senders inform the network about the congestion encountered by packets earlier in the same flow. This document specifies an IPv6 destination option [RFC2460] that can be used for performing ConEx markings in IPv6 datagrams.

This document specifies the ConEx wire protocol. The ConEx information can be used by any network element on the path to e.g. do traffic management or egress policing. Additionally this information will potentially be used by an audit function that checks the integrity of the sender's signaling. Further each transport protocol, that supports ConEx signaling, will need to specify precisely when the transport sets ConEx markings (e.g. the behavior for TCP is specified in [ID.conex-tcp-modifications]).

This specification is experimental to allow the IETF to assess whether the decision to implement the ConEx signal as a destination option fulfills the requirements stated in this document, as well as to evaluate the proposed encoding of the ConEx signals as described in [I-D.ietf-ConEx-abstract-mech].

The duration of this experiment is expected to be no less than two years from publication of this document as infrastructure is needed to be set up to determine the outcome of this experiment. Given ConEx is only chartered for IPv6, it might take longer to find a

suitable test scenario where only IPv6 traffic is managed using ConEx.

2. Conventions used in this document

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 [RFC2119].

3. Requirements for the coding of ConEx in IPv6

A set of requirement for an ideal concrete ConEx wire protocol is given in [I-D.ietf-ConEx-abstract-mech]. In the ConEx working group is was recognized that it will be difficult to find an encoding in IPv6 that satisfies all requirements. The choice in this document to implement the ConEx information in a destination option aims to satisfy those requirements that constrain the placement of ConEx information:

R-1: The marking mechanism needs to be visible to all ConEx-capable nodes on the path.

R-2: The mechanism needs to be able to traverse nodes that do not understand the markings. This is required to ensure that ConEx can be incrementally deployed over the Internet.

R-3: The presence of the marking mechanism should not significantly alter the processing of the packet. This is required to ensure that ConEx marked packets do not face any undue delays or drops due to a badly chosen mechanism.

R-4: The markings should be immutable once set by the sender. At the very least, any tampering should be detectable.

Based on these requirements four solutions to implement the ConEx information in the IPv6 header have been investigated: hop-by-hop options, destination options, using IPv6 header bits (from the flow label), and new extension headers. After evaluating the different solutions, the ConEx working group concluded that the use of a destination option would best address these requirements.

Choosing to use a destination option does not necessarily satisfy the requirement for on-path visibility, because it can be encapsulated by additional IP header(s). Therefore, ConEx-aware network devices, including policy or audit devices, might have to bury into inner IP headers to find ConEx information. This choice was a compromise between fast-path performance of Conex-aware network nodes and visibility, as discussed in Section Section 5.

4. ConEx Destination Option (CDO)

The ConEx Destination Option (CDO) is a destination option that can be included in IPv6 datagrams that are sent by ConEx-aware senders in order to inform ConEx-aware nodes on the path about the congestion encountered by packets earlier in the same flow or the expected risk of encountering congestion in the future. The CDO has an alignment requirement of (none).

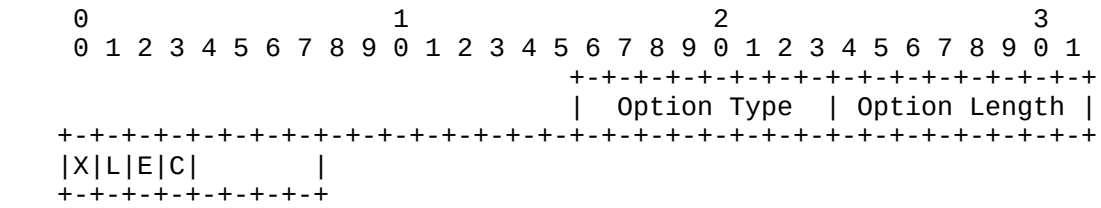

Figure 1: ConEx Destination Option Layout

Option Type

8-bit identifier of the type of option. The option identifier for the ConEx destination option will be allocated by the IANA.

Option Length

8-bit unsigned integer. The length of the option (excluding the Option Type and Option Length fields). This field MUST be set to the value 1.

X Bit

When this bit is set, the transport sender is using ConEx with this packet. If it is not set, the sender is not using ConEx with this packet.

L Bit

When this bit is set, the transport sender has experienced a loss.

E Bit

When this bit is set, the transport sender has experienced ECN-signaled congestion.

C Bit

When this bit is set, the transport sender is building up

congestion credit in the audit function.

Reserved

These bits are not used in the current specification. They are set to zero on the sender and are ignored on the receiver.

All packets sent over a ConEx-capable connection MUST carry the CDO. The CDO is immutable. Network devices with ConEx-aware functions read the flags, but all network devices MUST forward the CDO unaltered.

CDO MUST be placed as the first option in the destination option header before the AH and/or ESP (if present). IPsec Authentication Header (AH) MAY be used to verify that the CDO has not been modified.

If the X bit is zero all other three bits are undefined and thus should be ignored and forwarded unchanged by network nodes. The X bit set to zero means that the connection is ConEx-capable but this packet MUST NOT be counted when determining ConEx information in an audit function. This can be the case if no congestion feedback is (currently) available e.g. in TCP if one endpoint has been receiving data but sending nothing but pure ACKs (no user data) for some time. This is because pure ACKs do not advance the sequence number, so the TCP endpoint receiving them cannot reliably tell whether any have been lost due to congestion. Pure TCP ACKs cannot be ECN-marked either [RFC3168].

If the X bit is set, any of the other three bits (L, E, C) MAY be set. Whenever one of these bits is set, the number of bytes carried by this IP packet (including the IP header that directly encapsulates the CDO and everything that IP header encapsulates) SHOULD be counted to determine congestion or credit information. In IPv6 the number of bytes can easily be calculated by adding the number 40 (length of the IPv6 header in bytes) to the value present in the Payload Length field in the IPv6 header.

A transport sends credits prior to the occurrence of congestion (loss or ECN-CE marks) and the amount of credits should cover the congestion risk. Note, the maximum congestion risk is that all packets in flight get lost or ECN marked.

If the L or E bit is set, a congestion signal in the form of a loss or, respectively, an ECN mark was previously experienced by the same connection.

In principle all of these three bits (L, E, C) MAY be set in the same packet. In this case the packet size MUST be accounted more than once for each respective ConEx information counter.

If a network node extracts the ConEx information from a connection, it is expected to hold this information in bytes, e.g. comparing the total number of bytes sent with the number of bytes sent with ConEx congestion marks (L, E) to determine the current whole path congestion level. For ConEx-aware node processing, the CDO MUST use the Payload length field of the preceding IPv6 header for byte-based accounting. When a ratio is measured and equally sized packets can be assumed, counting the number of packets (instead of the number of bytes) should deliver the same result. But a network node must be aware that this estimation can be quite wrong, if e.g. different sized packed are sent and thus it is not reliable.

A ConEx sender SHOULD set the reserved bits in the CDO to zero. Other nodes MUST ignore these bits and ConEx-aware intermediate nodes MUST forward them unchanged, whatever their values. They MAY log the presence of a non-zero reserved field.

It might be possible to implement a proxy for a ConEx sender, as long as it is located where receiver feedback is always visible. A ConEx proxy MUST NOT introduce a CDO header into a packet already carrying one and it MUST NOT alter the information in any existing CDO header. However, it can add a CDO header to any packets without one, taking care not to disrupt any integrity or authentication mechanisms.

The CDO is only applicable on unicast or anycast packets (see [I-D.ietf-ConEx-abstract-mech] for reasoning). A ConEx sender MUST NOT send a packet with the CDO to a multicast address. ConEx-capable network nodes MUST treat a multicast packet with the X flag set the same as an equivalent packet without the CDO, but they SHOULD forward it unchanged.

There are no warning or error messages associated with the CDO.

5. Implementation in the fast path of ConEx-aware routers

The ConEx information is being encoded into a destination option so that it does not impact forwarding performance in the non-ConEx-aware nodes on the path. Since destination options are not usually processed by routers, the existence of the CDO does not affect the fast path processing of the datagram on non-ConEx-aware routers. i.e. They are not pushed into the slow path towards the control plane for exception processing.

The ConEx-aware nodes still need to process the CDO without severely affecting forwarding. For this to be possible, the ConEx-aware routers need to quickly ascertain the presence of the CDO and process the option if it is present. To efficiently perform this, the CDO needs to be placed in a fairly deterministic location. In order to facilitate forwarding on ConEx-aware routers, ConEx-aware senders that send IPv6 datagrams with the CDO MUST place the CDO as the first destination option in the destination options header.

6. Tunnel Processing

As with any destination option, an ingress tunnel endpoint will not natively copy the CDO when adding an encapsulating outer IP header. In general an ingress tunnel SHOULD NOT copy the CDO to the outer header as this would changed the number of bytes that would be counted. However, it MAY copy the CDO to the outer in order to facilitate visibility by subsequent on-path ConEx functions if the configuration of the tunnel ingress and the ConEx nodes is co-ordinated. This trades off the performance of ConEx functions against that of tunnel processing.

An egress tunnel endpoint SHOULD ignore any CDO on decapsulation of an outer IP header. The information in any inner CDO will always be considered correct, even if it differs from any outer CDO. Therefore, the decapsulator can strip the outer CDO without comparison to the inner. A decapsulator MAY compare the two, and MAY log any case where they differ. However, the packet MUST be forwarded irrespective of any such anomaly, given an outer CDO is only a performance optimization.

A network node that assesses ConEx information SHOULD search for encapsulated IP headers until a CDO is found. At any specific network location, the maximum necessary depth of search is likely to be the same for all packets.

7. Compatibility with use of IPsec

If the transport network cannot be trusted, IPsec Authentication should be used to ensure integrity of the ConEx information. If an attacker would be able to remove the ConEx marks, this could cause an audit device to penalize the respective connection, while the sender cannot easily detect that ConEx information is missing.

In IPv6 a Destination Option header can be placed in two possible position in the order of possible headers, either before the Routing header or after the Encapsulating Security Payload (ESP) header [RFC2460]. As the CDO is placed in the destination option header before the AH and/or ESP, it is not encrypted in transport mode

[RFC4301]. Otherwise, if the CDO were placed in the latter position and an ESP header were used, the CDO would also be encrypted and could not be interpreted by ConEx-aware devices.

The IPv6 protocol architecture currently does not provide a mechanism for new headers to be copied to the outer IP header. Therefore if IPsec encryption is used in tunnel mode, ConEx information cannot be accessed over the extent of the ESP tunnel.

8. Mitigating flooding attacks by using preferential drop

This section is aspirational, and not critical to the use of ConEx for more general traffic management. However, once CDO information is present, the CDO header could optionally also be used in the data plane of any IP-aware forwarding node to mitigate flooding attacks.

If a router queue experiences very high load so that it has to drop arriving packets, it MAY preferentially drop packets within the same Diffserv PHB using the preference order given in Table 1 (1 means drop first). Additionally, if a router implements preferential drop based on ConEx it SHOULD also support ECN-marking. Preferential dropping can be difficult to implement on some hardware, but if feasible it would discriminate against attack traffic if done as part of the overall policing framework as described in [I-D.ietf-ConEx-abstract-mech]. If nowhere else, routers at the egress of a network SHOULD implement preferential drop based on ConEx markings(stronger than the MAY above).

	Preference
Not-ConEx or no CDO	1 (drop first)
X (but not L,E or C)	2
X and L,E or C	3

Table 1: Drop preference for ConEx packets

A flooding attack is inherently about congestion of a resource. As load focuses on a victim, upstream queues grow, requiring honest sources to pre-load packets with a higher fraction of ConEx-marks.

If ECN marking is supported by downstream queues, preferential dropping provides the most benefits because, if the queue is so congested that it drops traffic, it will be CE-marking 100% of any forwarded traffic. Honest sources will therefore be sending 100% ConEx E-marked packets (and subject to rate-limiting at an ingress policer). Senders under malicious control can either do the same as

honest sources, and be rate-limited at ingress, or they can understate congestion and not set the E bit. If the preferential drop ranking is implemented on queues, these queues will preserve E/L-marked traffic until last. So, the traffic from malicious sources will all be automatically dropped first. Either way, malicious sources cannot send more than honest sources.

9. Acknowledgements

The authors would like to thank Marcelo Bagnulo, Bob Briscoe, Ingemar Johansson, Joel Halpern and John Leslie for the discussions that led to this document.

Special thanks to Bob Briscoe who contributed text and analysis work on preferential dropping.

10. Security Considerations

[I-D.ietf-ConEx-abstract-mech] describes the overall audit framework for assuring that ConEx markings truly reflect actual path congestion. This section focuses purely on the security of the encoding chosen for ConEx markings.

The chg bit in the CDO option type field is set to zero, meaning that the CDO option is immutable. If IPsec AH is used, a zero chg bit causes AH to cover the CDO option so that its end-to-end integrity can be verified, as explained in Section 4.

This document specifies that the Reserved field in the CDO must be ignored and forwarded unchanged even if it does not contain all zeroes. The Reserved field is also required to sit outside the encrypting security payload (ESP), at least in transport mode (see Section 7). This allows the sender to use the Reserved field as a 28-bit-per-packet covert channel to send information to an on-path node outside the control of IPsec. However, a covert channel is only a concern if it can circumvent IPsec in tunnel mode and, in the tunnel mode case, ESP would close the covert channel as outlined in Section 7.

11. IANA Considerations

This document defines a new IPv6 ConEx destination option for carrying ConEx markings. IANA is requested to assign a new destination option type in the Destination Options registry maintained at http://www.iana.org/assignments/ipv6-parameters <TBA1> ConEx Destination Option [RFCXXXX] The act bits for this option need to be 00. The destination IP stack will not usually process the CDO, therefore the sender can send a CDO without checking if the receiver

will understand it. The CDO MUST still be forwarded to the destination IP stack, because the destination might check the integrity of the whole packet, irrespective of whether it understands ConEx.

12. References

12.1. Normative References

[I-D.ietf-ConEx-abstract-mech] Mathis, M. and B. Briscoe, "Congestion Exposure (ConEx) Concepts and Abstract Mechanism", draft-ietf-ConEx-abstract-mech (work in progress), July 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.

[RFC6789] Briscoe, B., Woundy, R., and A. Cooper, "Congestion Exposure (ConEx) Concepts and Use Cases", RFC 6789, December 2012.

12.2. Informative References

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.

Authors' Addresses

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

Email: suresh.krishnan@ericsson.com

Mirja Kuehlewind
ETH Zurich

Email: mirja.kuehlewind@tik.ee.ethz.ch

Carlos Ralli Ucendo
Telefonica

Email: ralli@tid.es

CONEX WG
Internet-Draft
Intended status: Informational
Expires: March 21, 2015

D. Kutscher
F. Mir
R. Winter
NEC
S. Krishnan
Y. Zhang
Ericsson
CJ. Bernardos
UC3M

September 17, 2014

Mobile Communication Congestion Exposure Scenario
draft-ietf-conex-mobile-04

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). The draft provides a brief overview of the architecture of these networks (both access and core networks), current QoS mechanisms and then discusses how congestion exposure concepts could be applied. Based on this, this memo suggests a set of requirements for ConEx mechanisms that particularly apply to these mobile networks.

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 March 21, 2015.

Copyright Notice

Table of Contents

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 cellular network data traffic 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 congestion exposure 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. The IETF CONEX WG decided to initially start to work on a specific use case, where 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 Evolved Packet System (EPS) network, 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. I.e., 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.

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 extent in order to provide network-controlled quality of service. These efforts often lead to complexity and cost. ConEx could be simpler and more capable approach to efficient resource sharing in these networks.

In the following, we discuss ways how 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 Evolved Packet System (EPS) should refer to Appendix B 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 EPC -- despite its cost and complexity. With the increase of encrypted data traffic, traffic management using DPI alone however 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 effectively more users can be served 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 the application classes flows belong to and 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) and perform better as decisions can be taken 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:

    1. as additional information to assist network functions to impose different QoS for different application sessions; and

    2. 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 then, say, application-layer responsive in protocols such as 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 WiFi 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 in 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

As 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, application store downloads from web (browser)-based or other more real-time communication. For many of the bulk data transfers, completion times aren’t important within certain bounds and therefore if scavenger (or less-than best effort) transports like e.g. 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 allowing users that contribute less to congestion to give them a larger bandwidth allowance or e.g. to lower 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 cf. the Policy Charging Control (PCC) architecture. 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 volume-based accounting to a 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 within such congestion volume accounting would be possible, without requiring any support from the global Internet, in particular since the typical scare 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 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 any 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 mechanisms such as ECN is in place).

[I-D.briscoe-conex-initial-deploy] provides specific examples of how ConEx deployments can be initiated, focusing on unilaterial deployments by single networks, i.e., partial deployment.

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 hand-over. There is existing work in [raghavan2007] on distributed rate limiting and in [nec.euronf-2011] on specific optimizations for congestion exposure and policing in mobility scenarios.

Another aspect to consider is the addition of Selected IP Traffic Offload (SIPTO) and Local Breakout (LIPA), also see [3GPP.23.829], i.e., the idea that some traffic (e.g., 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, 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 eNodeB interfaces, so that a more general approach, involving more than one single operator's network, can be feasible as well.

3. CONEX in the EPS

At the time of writing, the CONEX mechanism is still work in progress in the IETF working group. Still, discussing a few options for how

such a mechanism (and possibly additional policing functions) could eventually be deployed in 3GPP's EPS is useful at this point. Note that this description of options is not intended as a complete set of possible approaches -- it is merely intended for discussing the most promising options.

3.1. Possible Deployment Scenarios

There are different possible ways 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 three different deployment scenarios for congestion exposure in the figures below:

1. In Figure 1 CONEX is supported by servers for sending data (here: 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. Isolated CONEX domains as depicted in Figure 3, where 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 or dropped (or ECN-CE marked) between the tunnel sender and

receiver. For that purpose, a new field is introduced to the tunnel header called the Byte Sequence Marker (BSN) that identifies the byte in the flow of data from the tunnel sender to the tunnel receiver. A policer at the tunnel sender is expected to re-act according to the tunnel congestion volume (see [conex-lite] for details.)

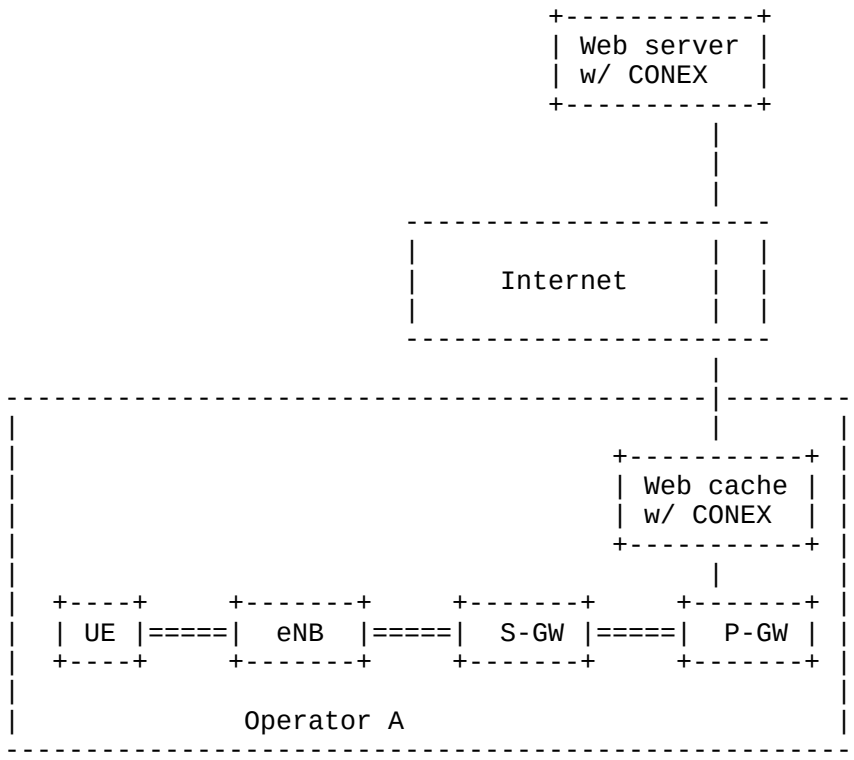

Figure 1: CONEX support on servers and caches

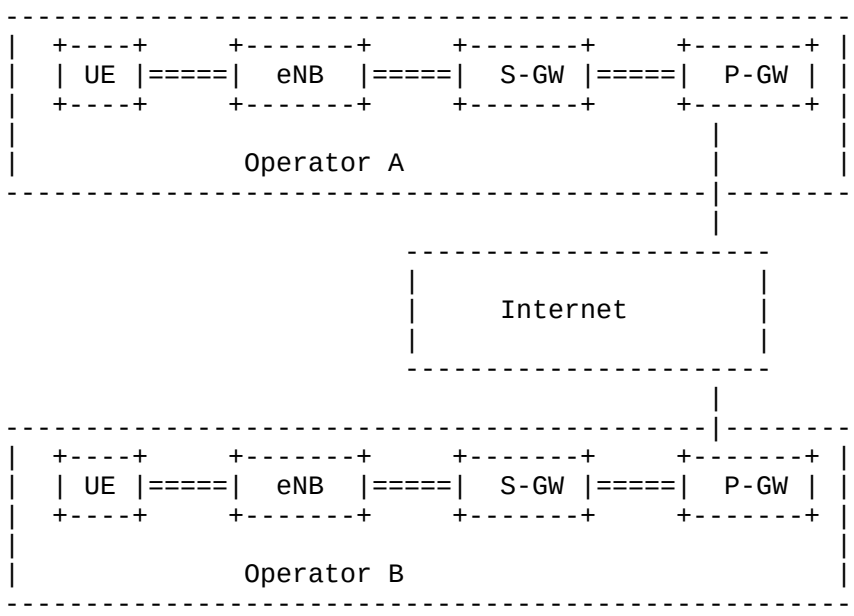

Figure 2: CONEX deployment across operator domains

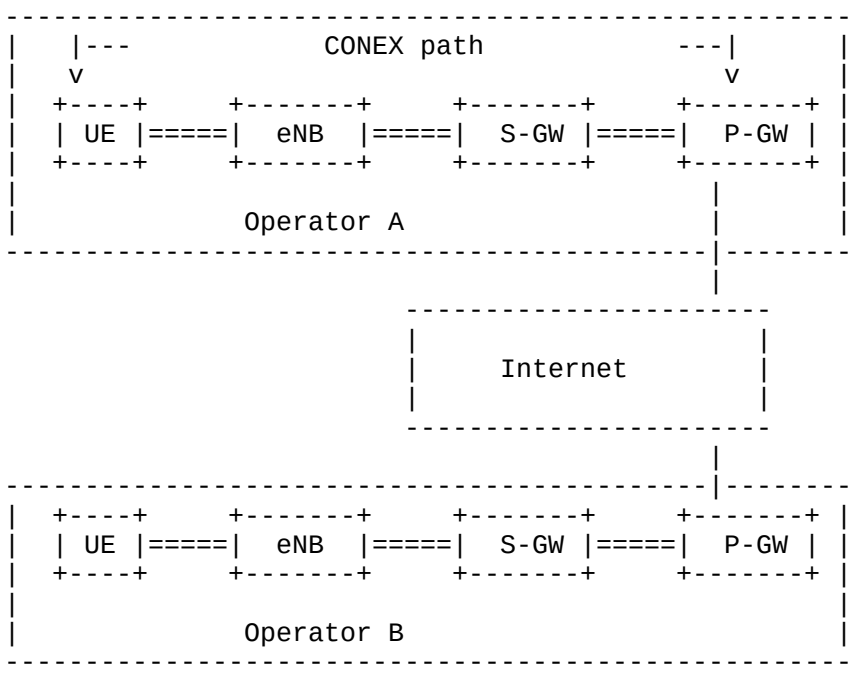

Figure 3: CONEX deployment in a single operator domain

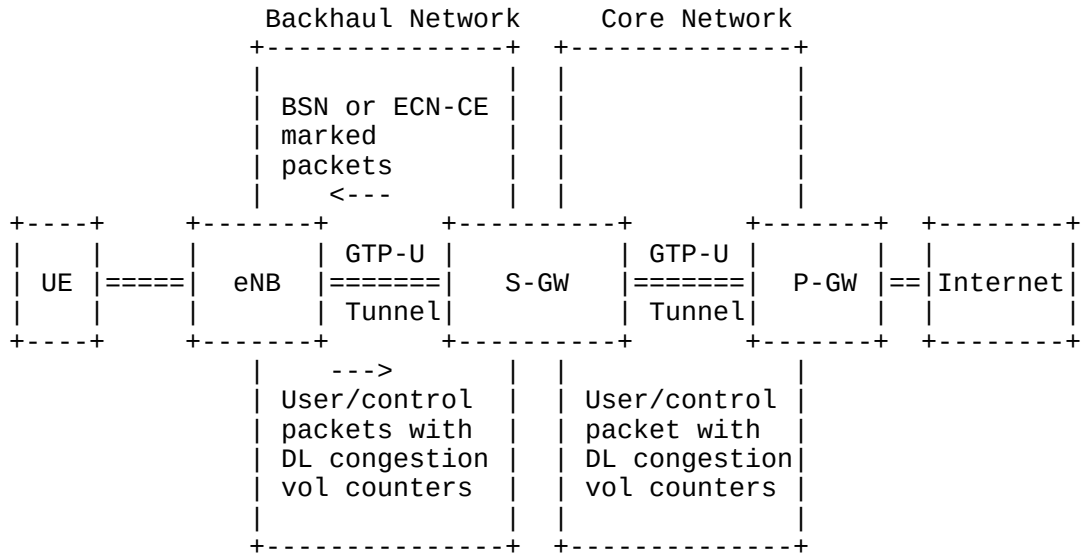

Figure 4: CONEX-lite deployment

We consider all three scenarios to be relevant and believe that all of them are within the scope of the CONEX WG charter. A more detailed description will be provided in a future version of this document.

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 3GPP's EPS. [I-D.ietf-conex-abstract-mech] is describing the following congestion exposure components:

- Modified senders that send congestion exposure information in response to congestion feedback.
- Receivers that generate congestion feedback (leveraging existing behavior or requiring new functions).
- Audit functions that audit CONEX signals against actual congestion, e.g., by monitoring flows or aggregate of flows.
- 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

As described in [I-D.briscoe-conex-initial-deploy], 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 [I-D.ietf-conex-destopt] is contained in a destination option header field, which requires minimal changes at senders and nodes that want to assess path congestion -- and that 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 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 an 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 their flows congestion contribution. ECN is already partially introduced into 3GPP networks: Section 11.6 in [3GPP.36.300] specifies the usage of ECN for congestion notification on the radio link (between eNB and UE), and [3GPP.26.114] 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) and on [I-D.briscoe-tsvwg-ecn-encap-guidelines]. Specifically, a specification for tunneling ECN in GTP-U will be needed.

3.2.2. CONEX Functions in the Mobile Network

In the following, 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 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 however not imply that uplink congestion is not an issue, since the

asymmetric maximum bandwidth configuration can create a smaller bottleneck for uplink traffic -- and 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 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 possible 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, e.g., to police traffic 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, while enabling operators to re-use existing infrastructure, specifically wireless base stations, PCRF and 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: Proxy-Mobile-IPv6 (PMIPv6, [3GPP.23.402]) and 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 -- which calls for efficient, scalable, 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 different operators policies.

CONEX may also be used to complement 3GPP-based mechanisms for congestion management which are currently under development, such as in the User Plane Congestion Management (UPCON) work item described in [3GPP.23.705].

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 a 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 in cellular network is the presence of user mobility compared to wired networks. As the user location changes, the same device can be connected to the network via different base stations (eNodeBs) or even go through switching gateways. Thus, the CONEX scheme must to be able to carry 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, and [I-D.briscoe-tsvwg-ecn-encap-guidelines] provides guidelines for adding congestion notification to protocols that encapsulate IP.

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. IANA Considerations

No IANA considerations.

6. Security Considerations

Security considerations for applying CONEX to EPS include, but are not limited to, the security considerations that apply to the CONEX protocols.

7. References

7.1. Normative References

[RFC6789] Briscoe, B., Woundy, R., and A. Cooper, "Congestion Exposure (ConEx) Concepts and Use Cases", RFC 6789, December 2012.

7.2. Informative References

[3GPP.23.401]
3GPP, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access", 3GPP TS 23.401 10.10.0, March 2013.

[3GPP.23.402]
3GPP, "Architecture enhancements for non-3GPP accesses", 3GPP TS 23.402 10.8.0, September 2012.

[3GPP.23.705]
3GPP, "System Enhancements for User Plane Congestion Management", 3GPP TR 23.705 0.8.0, October 2013.

[3GPP.23.829]
3GPP, "Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO)", 3GPP TR 23.829 10.0.1, October 2011.

[3GPP.26.114]
3GPP, "IP Multimedia Subsystem (IMS); Multimedia telephony; Media handling and interaction", 3GPP TS 26.114 10.7.0, June 2013.

[3GPP.29.060]
3GPP, "General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface", 3GPP TS 29.060 3.19.0, March 2004.

[3GPP.29.274]
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 10.11.0,

June 2013.

[3GPP.36.300]
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 10.11.0, September 2013.

[I-D.briscoe-conex-initial-deploy]
Briscoe, B., "Initial Congestion Exposure (ConEx) Deployment Examples", draft-briscoe-conex-initial-deploy-03 (work in progress), July 2012.

[I-D.briscoe-tsvwg-ecn-encap-guidelines]
Briscoe, B., Kaippallimalil, J., and P. Thaler, "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP", draft-briscoe-tsvwg-ecn-encap-guidelines-03 (work in progress), September 2013.

[I-D.ietf-conex-abstract-mech]
Mathis, M. and B. Briscoe, "Congestion Exposure (ConEx) Concepts and Abstract Mechanism", draft-ietf-conex-abstract-mech-08 (work in progress), October 2013.

[I-D.ietf-conex-destopt]
Krishnan, S., Kuehlewind, M., and C. Ucendo, "IPv6 Destination Option for ConEx", draft-ietf-conex-destopt-05 (work in progress), October 2013.

[RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, November 2010.

[RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, December 2012.

[conex-lite]
Baillargeon and Johansson, "ConEx Lite for Mobile Networks", in proceedings of ACM SIGCOMM-2014 Capacity Sharing Workshop (CSWS-2014), August 2014.

[dash] ISO/IEC, "ISO/IEC 23009-1: Information Technology -- Dynamic Adaptive Streaming over HTTP (DASH) --", April 2012.

[lte-sigcomm2013]
Huang, Qian, Guo, Zhou, Xu, Mao, Sen, and Spatscheck, "An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance", in proceedings of ACM SIGCOMM-2013, August 2013.

[nec.euronf-2011]
Mir, Kutscher, and Brunner, "Congestion Exposure in Mobility Scenarios", in proceedings of 7th EURO-NF CONFERENCE ON NEXT GENERATION INTERNET, June 2011.

[nec.globecom2010]
Kutscher, Lundqvist, and Mir, "Congestion Exposure in Mobile Wireless Communications", in proceedings of IEEE GLOBECOM 2010, December 2010.

[raghavan2007]
Raghavan, Vishwanath, Ramabhadran, Yocum, and Snoeren, "Cloud Control with Distributed Rate Limiting", in proceedings of ACM SIGCOMM 2007, 2007.

DOI: http://doi.acm.org/10.1145/1282427.1282419

Appendix A. Acknowledgments

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

Appendix B. Overview of 3GPP's Evolved Packet System (EPS)

This section provides an overview of 3GPP's "Evolved Packet System" (EPS [3GPP.36.300], [3GPP.23.401]) 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 1. The EPS provides IP connectivity to user equipment (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 UMTS Terrestrial Radio Access Network (E-UTRAN) and the core network called Evolved Packet Core (EPC). QoS is supported through an EPS bearer concept,

providing bindings to resource reservation within the network.

The evolved NodeB (eNB), the Long Term Evolution (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 Packet Data Network (PDN) Gateway, P-GW, is responsible for IP address allocation to the UE and is a tunnel endpoint for user 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 [3GPP.29.060]) 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 [3GPP.29.274]). Alternatively Proxy Mobile IP (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.

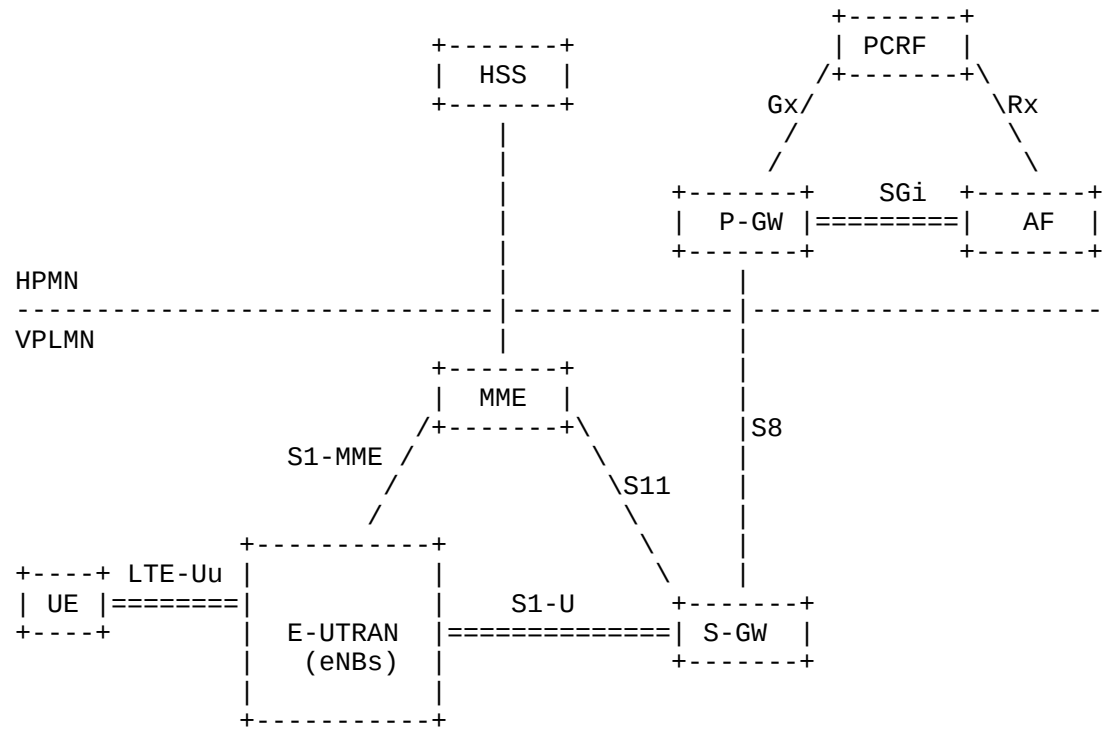

Figure 5: EPS architecture overview (Roaming Case)

Appendix C. ChangeLog

C.1. draft-ietf-conex-mobile-04

- added conex lite to deployment scenarios
- added reference to LTE study paper at SIGCOMM-2013"
- restructured section 2 and 3 (moved 3GPP specifics out of section2
- added normative references section and put RFC6789 there
- updated references

C.2. draft-ietf-conex-mobile-03

- implemented suggestions for improving 3GPP EPS descriptions by Andreas Maeder

- mentioned 3GPP UPCON and added reference
- updated references
- In section 3.1 (CONEX as a Basis for Traffic Management), changed the wording in the first abstract of the enumerated list to state that ConEx can enable/enhance flow policy-based traffic management -- not DPI (as we earlier said). DPI is not the objective -- it is the tool that is currently used...
- merged section 3.4 (CONEX as a Form of Differential QoS) into 3.1 (CONEX as a Basis for Traffic Management)
- moved section 2 (Overview of 3GPP’s Evolved Packet System (EPS)) to appendix.
- renamed section "CONEX Use Cases in the Mobile Communication Scenario" to "CONEX Use Cases in Mobile Communication Networks"
- updated TDF text in "CONEX as a Basis for Traffic Management"
- added reference to 3GPP UPCON to summary

C.3. Earlier

- changed title to "Mobile Communication Congestion Exposure Scenario" (was "use case")
- added new section 3 on "CONEX Uses Cases in mobile communication scenario"
- removed "Motivation" section in section 4
- removed "isolated connex deployment section in section 4"
- renamed "EPS integration" section in section 4 to "Additional EPS integration options"
- added a (still empty) summary section to section 4
- s/Re-ECN/CONEX/g
- added references
- added acknowledgments

Authors' Addresses

Dirk Kutscher
NEC
Kurfuersten-Anlage 36
Heidelberg,
Germany

Phone:
Email: kutscher@neclab.eu

Faisal Ghias Mir
NEC
Kurfuersten-Anlage 36
Heidelberg,
Germany

Phone:
Email: faisal.mir@neclab.eu

Rolf Winter
NEC
Kurfuersten-Anlage 36
Heidelberg,
Germany

Phone:
Email: rolf.winter@neclab.eu

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

Phone:
Email: suresh.krishnan@ericsson.com

Ying Zhang
Ericsson
200 Holger Way
San Jose, CA 95134
USA

Phone:
Email: ying.zhang@ericsson.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
URI: http://www.it.uc3m.es/cjbc/