

CONEX
Internet-Draft
Intended status: Informational
Expires: September 9, 2010

H. Tschofenig
Nokia Siemens Networks
A. Cooper
Center for Democracy &
Technology
March 8, 2010

Congestion Exposure Problem Statement
draft-tschofenig-conex-ps-02.txt

Abstract

The increasingly ubiquitous availability of broadband, together with flat-rate pricing, have made for increasing congestion problems on the network, which are often caused by a small number of users consuming a large amount of bandwidth. In some cases, building out more capacity to handle this new congestion may be infeasible or unwarranted. As a result, network operators have sought other ways to manage congestion both from their own users and from other networks. These different types of solutions have different strengths and weaknesses, but all of them are limited in a number of key ways.

This document discusses the problems created for operators by high-consuming users and describes the strengths and weaknesses of a number of techniques operators are currently using to cope with high bandwidth usage. The discussion of these solutions ultimately points to a need for a new kind of congestion accounting.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

Internet-Draft

Congestion Exposure PS

March 2010

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 9, 2010.

Copyright Notice

Copyright (c) 2010 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 BSD License.

Internet-Draft

Congestion Exposure PS

March 2010

Table of Contents

1.	Introduction	4
2.	Existing Approaches to Congestion Management	5
2.1.	Volume-Based Approaches	5
2.2.	Rate-Based Approaches	5
2.3.	Congestion-Based Approaches	5
2.4.	Applications-Based Approaches	6
3.	General Limitations of All Approaches	7
4.	New Activities	8
5.	Next Steps	9
6.	Security Considerations	10
7.	IANA Considerations	11
8.	Acknowledgments	12
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	13
Appendix A.	Example Policy Statement	15
A.1.	Fair Usage Policy	15
A.1.1.	What is the Fair Usage Policy?	15
A.1.2.	How do I know I'm a very heavy user?	15
A.1.3.	I have Contract Option 3, does the Fair Usage Policy apply to me?	15
A.1.4.	Peer to Peer (P2P)	15
	Authors' Addresses	17

1. Introduction

With the growth of "always on" broadband connections, network operators are increasingly facing congestion problems caused by a small number of network users occupying a large proportion of network capacity [[broadband-traffic-report](#)][traffic2]. This trend does not necessarily present a problem on its face, as increased traffic volumes do not automatically lead to congestion. However, in some cases where operators were not expecting these changes in growth rates and traffic consumption, their pricing models and congestion management architectures have proved inadequate. This is true of both congestion caused by an operator's own subscribers and congestion caused by an interconnecting network.

In some cases, building out more capacity to handle this new congestion may be infeasible or unwarranted. The cost of building new capacity may be prohibitive, especially for network operators that charge flat-rate fees to subscribers and are thus unable to charge heavier users more on the basis of their larger contribution to the congestion problem. For an operator facing congestion caused by other operators' networks, building out its own capacity is unlikely to solve the congestion problem. Operators are thus facing increased pressure to find effective solutions to dealing with high-consuming users, other than building out new capacity.

There are many factors to consider for each kind of solution, including how the solution performs, its cost, what its public relations impact might be, and what legal framework exists to support its use. The performance considerations must take into account the

balance between device performance and forwarding performance (since many of the solution mechanisms slow down forwarding performance), and this determination is intimately related to measuring a solution's overall cost.

The responses that operators have sought to manage congestion both from their own users and from other networks are generally based on one of four factors that the operator can observe for each subscriber or hand-off point on the network: volume of bandwidth used, rates of data transfer, congestion volume, and applications used. These different types of solutions have different strengths and weaknesses, but all of them are limited in a number of key ways. [Section 2](#) discusses the specific strengths and weaknesses of each approach. [Section 3](#), discusses the limitations that are common to all of the approaches.

[2.](#) Existing Approaches to Congestion Management

[2.1.](#) Volume-Based Approaches

The volume of traffic sent or received by a particular user or network is easy to measure. Operators have a number of different standardized protocols available to them to conduct volume accounting. Many information elements can be sent from an accounting client to an accounting server using standardized protocols, such as RADIUS (see [\[RFC2866\]](#) and [\[RFC2865\]](#)) and Diameter [\[RFC3588\]](#), to effectuate volume-based accounting. The existence of standardized protocols has allowed resource consumption measurement in roaming cases due to the interconnected AAA systems. These protocols are now used in almost every enterprise and operator network. The initial accounting mechanisms envisioned a rather non-real time nature in reporting resource consumption but with mechanisms like like Diameter Credit Control [\[RFC4006\]](#) allowed real-time credit control checks, allowing operators to make fine-grained congestion decisions based on volume.

If the collected accounting information leads to billable events then this typically leads to a quite effective countermeasure against

heavy usage. For example, data usage while roaming is often charged per volume (or per time) and heavy usage leads to huge costs. So, user's typically avoid heavy usage unless they are not aware of the fact that they are roaming, which may happen.

In any case, the drawback of all of these approaches is that they tend to be less user friendly nor do they reflect congestion in any way.

[2.2.](#) Rate-Based Approaches

Volume-based protocols are blind to protocols like LEDBAT [[ledbat](#)] that are specifically designed to be more sensitive to congestion than the underlying transport. One way that operators have used volume accounting is by including thresholds for baseline usage volume in end user contracts and reducing priority (or charging fees) when a user's consumption goes beyond the threshold. Under such a scheme, users would not be incentivized to use applications that support protocols like LEDBAT, because the volume accounting would not take the congestion benefits of using a LEDBAT-style protocol into consideration.

[2.3.](#) Congestion-Based Approaches

Initial attempts to capture congestion situations have simply focused on the peak hours and aimed at rate limiting heavy users during that

time. For example, users who have consumed a certain amount of bandwidth during the last 24 hours got elected as those who get their traffic shaped, in case the total amount of traffic reaches a congestion situation in certain nodes within the operators network.

More sophisticated schemes were developed to measure the congestion situation at different entities in the network and to take the overall consumption of the user's traffic into account when picking users for packet remarking that could lead, under congestion situation, to packet dropping. The Comcast FairShare solution [[I-D.livingood-woundy-congestion-mgmt](#)] belongs to the more sophisticated schemes.

[2.4.](#) Applications-Based Approaches

The use of deep packet inspection (DPI) allows operators to observe and analyze traffic at the application layer. This capability can be used to determine the applications and/or application-layer protocols that subscribers are using and respond on a per-application or per-protocol basis. An example of an ISP's fair usage policy describing how it manages specific protocols is included in [Appendix A](#).

If an operator experiences much congestion based on the use of easily identifiable applications protocols, or if the DPI capability can also be used for other purposes, operators may find this approach attractive. However, the applications-based approach has some drawbacks. The process of inspecting traffic, particularly in real time, can be highly performance-intensive. DPI equipment may also require continuous software updates to ensure that the detection engine recognizes the latest protocol variants, and its use can result in a cat-and-mouse game with applications developers constantly seeking ways to work around the packet inspection engine. Public policy concerns -- privacy, operator liability, user backlash, and others -- are another drawback.

[3.](#) General Limitations of All Approaches

All three of the approaches discussed above suffer from some general limitations. First, they introduce performance uncertainty. Flat-rate pricing plans are popular because operators know that users appreciate the certainty of having their monthly bill amount remain the same for each billing period, allowing users to plan their costs accordingly. But while flat-rate pricing avoids billing uncertainty,

it creates performance uncertainty: users cannot be sure that the performance of their connections is not being altered or degraded based on how the network operator manages congestion.

Second, none of the approaches is able to make use of what may be the most important factor in managing congestion: the amount that an endpoint contributes to congestion on the network. This information simply is not available to network nodes, and neither volume nor rate nor application usage is an adequate proxy for congestion volume, because none of these metrics measures a user or network's actual contribution to congestion on the network. [This point needs a little more discussion.]

Finally, none of these solutions accounts for inter-network congestion. The currently available mechanisms for identifying and mitigating congestion largely run wholly within an operator's network and without a lot of information exchange about congestion information to or from end hosts or other network operators. Exposing this information may allow end devices to make more informed decisions (although policy enforcement would still be required by the operator). [This point needs to be filled in more.]

Following the IETF Workshop on Peer-to-Peer (P2P) Infrastructure in 2008 (see [[RFC5594](#)]), two working groups and one research group were created that relate to the congestion issues created by peer-to-peer application usage: :

LEDBAT (Low Extra Delay Background Transport) [[ledbat](#)] is designed to keep the latency across a congested bottleneck low even as it is saturated. This allows applications that send large amounts of data, particularly upstream on home connections (such as peer-to-peer applications) to operate without destroying the user experience in interactive applications.

LEDBAT holds substantial promise should P2P clients adopt it widely. This solution has been focused on P2P applications, and its applicability to other applications, such as video using H.264, is unclear.

ALTO (Application-Layer Traffic Optimization) [[alto](#)] aims to design and specify mechanisms that will provide applications, typically P2P applications, with information to perform better-than-random initial peer selection to increase their performance and at the same time to avoid excessive cross-domain traffic that tends to be more expensive for the operator. ALTO services may take different approaches at balancing factors such as maximum bandwidth, minimum cross-domain traffic, or lowest cost to the user, but in all cases the goal is to expose information that can ameliorate the interactions between peer-to-peer usage and other usages of shared networks.

Peer to Peer Research Group [[p2prg](#)] aims to provide a discussion forum for researchers related to all sorts of challenges presented by P2P systems in general, such as P2P streaming, interconnecting distinct P2P application overlays, security and privacy. Current work on exposing myths about peer-to-peer filesharing [[I-D.irtf-p2prg-mythbustering](#)] provides a number of references to support some of the claimed benefits of ALTO solutions mechanisms, such as the expected decrease in cross- domain traffic.

5. Next Steps

Congestion is a reality. Operators that would like to counteract the impact of congestion on their networks have a fair number of tools at their disposal. These tools may allow operators to identify heavy users, collect performance and usage indications, and choose from a variety of mitigating steps depending on the operator's preferred business practices. Subscriber-specific information, including policies, resource consumption information, and details about the current network attachment point, may be available in accounting servers. Information about the network topology and the state of particular topology elements may be available in the network management infrastructure. Solution approaches similar to [\[I-D.livingood-woundy-congestion-mgmt\]](#) have demonstrated one way of taking congestion information into consideration.

The collection of congestion information poses the challenge of deciding where in the network to put the metering agents to ensure that enough information is collected at the right point in time. Distributed collection and the correlation of the information across different nodes is a complex task. An approach that collects this congestion information along the path of the data packet (via inband signaling) would simplify this task. Regardless of the technical solution utilized for collecting information, certain users will undoubtedly observe the effects of decisions that operators make about how to handle congestion. Allowing users to understand these decisions will be crucial and having a channel to send feedback to the end device and/or subscriber would be a helpful step towards increased transparency.

[6.](#) Security Considerations

This document highlights approaches for dealing with high-consuming network users and all of them raise security and privacy concerns. It does not introduce new mechanisms. The security considerations for the existing mechanisms mentioned apply.

[7.](#) IANA Considerations

This document does not require actions by IANA.

[8.](#) Acknowledgments

The authors would like to thank Alan DeKok, Jens-Peter Haack, Alexander Bachmutsky, Jonne Soininen, Joachim Charzinski, Hannu Flinck, Joachim Kross, Jouni Korhonen, Mayutan Arumaithurai, Richard Woundy, Daniel Correa Lobato, Luca Caviglione, Tommy Lindgren, Lars Eggert, and Jason Livingood for their time to discuss the topic. Additionally, we would like to thank Marcin Matuszewski for his help with the P2P infrastructure workshop paper (since it was used as a starting point for the work on this memo).

[9.](#) References

[9.1.](#) Normative References

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

[9.2.](#) Informative References

[I-D.irtf-p2prg-mythbustering]

Marocco, E., Fusco, A., Rimac, I., and V. Gurbani,
"Improving Peer Selection in Peer-to-peer Applications:
Myths vs. Reality", [draft-irtf-p2prg-mythbustering-01](#)
(work in progress), March 2010.

[I-D.livingood-woundy-congestion-mgmt]

Bastian, C., Klieber, T., Livingood, J., Mills, J., and R.
Woundy, "Comcast's Protocol-Agnostic Congestion Management

System", [draft-livingood-woundy-congestion-mgmt-03](#) (work in progress), February 2010.

- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", [RFC 2865](#), June 2000.
- [RFC2866] Rigney, C., "RADIUS Accounting", [RFC 2866](#), June 2000.
- [RFC2975] Aboba, B., Arkko, J., and D. Harrington, "Introduction to Accounting Management", [RFC 2975](#), October 2000.
- [RFC3576] Chiba, M., Dommety, G., Eklund, M., Mitton, D., and B. Aboba, "Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)", [RFC 3576](#), July 2003.
- [RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", [RFC 3588](#), September 2003.
- [RFC4006] Hakala, H., Mattila, L., Koskinen, J-P., Stura, M., and J. Loughney, "Diameter Credit-Control Application", [RFC 4006](#), August 2005.
- [RFC5594] Peterson, J. and A. Cooper, "Report from the IETF Workshop on Peer-to-Peer (P2P) Infrastructure, May 28, 2008", [RFC 5594](#), July 2009.
- [alto] "",
<<http://www.ietf.org/dyn/wg/charter/alto-charter.html>>.

Tschofenig & Cooper

Expires September 9, 2010

[Page 13]

Internet-Draft

Congestion Exposure PS

March 2010

[broadband-traffic-report]

Cho, K., "Broadband Traffic Report", Internet Infrastructure Review 4, 2009.

[ledbat]

"",
<<http://www.ietf.org/dyn/wg/charter/ledbat-charter.html>>.

[p2prg]

"", <<http://www.irtf.org/charter?gtype=rg&group=p2prg>>.

[traffic]

Cho, K., Fukuda, K., Kato, H., and A. Kato, "The impact and implications of the growth in residential user-to-user

traffic", SIGCOMM Comput. Commun. Rev. 36, 2006.

[traffic2]

Cho, K., Fukuda, K., Esaki, H., and A. Kato, "Observing slow crustal movement in residential user traffic, in International Conference On Emerging Networking Experiments And Technologies, Proceedings of the 2008 ACM CoNEXT Conference, Madrid, Spain, Article No. 12", , 2008.

[Appendix A](#). Example Policy Statement

[A.1](#). Fair Usage Policy

[A.1.1.](#) What is the Fair Usage Policy?

The Fair Usage Policy is designed to ensure that the service received by the vast majority of our customers is not negatively impacted because of extremely heavy usage by a very small minority of customers. This is why ISP X continuously monitors network performance and may restrict the speed available to very heavy users during peak time. This applies to customers on all Options. Note if you are a heavy user we will only restrict your speed, service will not be stopped so ability to upload and download remains. No restrictions will be imposed outside of the peak times. Only a very small minority of customers will ever be affected by this (less than 1 %).

[A.1.2.](#) How do I know I'm a very heavy user?

There is no hard and fast usage limit that determines if you are a heavy user as the parameters that determine heavy use vary with the demands placed on the network at that given time. If you have a query about fair usage related restrictions on your line please call us.

[A.1.3.](#) I have Contract Option 3, does the Fair Usage Policy apply to me?

Yes, the Fair Usage Policy applies to all customers on all Options, including Option 3. Option 3 allows unlimited downloads and uploads inclusive of the monthly rental price, so you will not be charged for over-use, however this does not preclude ISP X from restricting your speed at peak times if you are a heavy user. If you are an Option 3 heavy user this does not prevent you from continuing to use your service, nor does it cost you any more but it ensures that you do not negatively impact the majority of our customers who share the available bandwidth with you.

[A.1.4.](#) Peer to Peer (P2P)

[A.1.4.1.](#) I'm noticing slower P2P speeds at peak times even though I'm not a very heavy user, why is this?

P2P is the sharing and delivery of files amongst groups of people who are logged on to a file sharing network. P2P consumes a significant and highly disproportionate amount of bandwidth when in use even by small numbers of users.

This is why we have a peak time policy where we limit P2P speeds to manage the amount of bandwidth that is used by this application in particular.

Without these limits all our customers using their broadband service at peak times would suffer, regardless of whether they are using P2P or not. It's important to remember that P2P isn't a time-critical application so if you do need to download large files we advise you to do this at off-peak times when no restrictions are placed, not only will you be able to download faster but your usage will not negatively impact other users.

[A.1.4.2](#). Does this mean I can't use Peer-to-Peer (P2P) applications?

No, we are not stopping you from using any P2P service, P2P will just be slowed down at peak times. Again, P2P is not generally a time-sensitive application.

Internet-Draft

Congestion Exposure PS

March 2010

Authors' Addresses

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo FIN-02600
Finland

Phone: +358 (50) 4871445
Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Alissa Cooper
Center for Democracy & Technology
1634 I Street NW, Suite 1100
Washington, DC
USA

Email: acooper@cdt.org

Tschofenig & Cooper

Expires September 9, 2010

[Page 17]