CONEX
Internet-Draft
Intended status: Informational
Expires: April 29, 2010

H. Tschofenig
Nokia Siemens Networks
A. Cooper
Center for Democracy & Technology
October 26, 2009

Congestion Exposure Problem Statement
draft-tschofenig-conex-ps-01.txt

Status of this Memo

Copyright Notice

Abstract

The increasingly ubiquitous availability of broadband, together with flat-rate pricing, have made the use of new kinds of peer-to-peer applications increasingly common. From the perspective of the Internet's evolution and its contributions to end user value, this is very exciting. However, the uptick in peer-to-peer application usage has also contributed to the rise of "high-consuming users" who take their flat-rate contracts to the limit by continuously file-sharing to the maximum extent possible. Network operators have responded to this phenomenon in a number of different fashions.

This document discusses the problems created for operators by high-consuming users and illustrates a number of techniques operators are currently using to cope with high bandwidth usage.

Table of Contents

1. Introduction

In recent years, network operators around the world have begun to feel the affects of "high-consuming users" -- those who use the maximum amount of bandwidth possible, usually for the purpose of peer-to-peer file sharing. In 2006 K. Cho et al. [traffic] reported that for residential Japanese broadband connections, "a small number of users dictate the overall behavior; 4% of high-consuming users account for 75% of the inbound volume, and the fiber users account for 86% of the inbound volume." A more recent paper published December 2008, see [traffic2], confirms that the distribution has not changed much. User-to-user traffic comprised 63% of overall residential traffic volume. The authors noted not only a changing traffic distribution, but also a substantial increase in overall traffic growth (37% per year). Operators in other countries have experienced similar shifts.

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. In some countries, fierce competition among Internet access providers has yielded low profit margins. This has increased the pressure on operators to find effective ways to deal with high- consuming users who cost them more money than the bulk of their subscribers. Furthermore, some broadband networks (such as cable networks) may not have the ideal characteristics (routing topology, for example) to support high volumes of user-to-user traffic.

Congestion and cost are closely related. As stated in RFC 5594 [RFC5594], "... congestion can be viewed merely as a manifestation of cost. An ISP that invests in capacity could be considered to be paying to relieve congestion. Or, if subscribers are charged for congesting the network, then cost and congestion could be viewed as one and the same. The distinction between them may thus be

artificial.". The upshot for network operators is: those who produce a lot of traffic cost a lot.

Operators are now facing a range of options for addressing this problem. There are many factors to consider for each kind of solution, including how the solution performs, its cost, what the public relations impact of using a particular solution might be, and what legal framework exists to support the use of a particular solution. 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 intimiately related to measuring a

solution's overall cost.

In some cases, the popularity of flat-rate pricing plans exacerbate the congestion problem because an individual's bandwidth usage is not tied to his or her monthly bill, creating an incentive to use as much bandwidth as possible and leaving operators to cover the costs of all their users with essentially equal payments from each. Operators know that users appreciate the certainty of having the 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 degrated based on how the network operator manages congestion. Unfortunately, most of the solutions described below create some performance uncertainty, and thus users are unlikely to view them as ideal solutions, despite users' well known preference for flat-rate pricing.

2. State-of-the-Art Building Blocks

Two means of learning about the resource consumption and the traffic traveling through the network that are in use today are accounting and deep packet inspection.

2.1. Accounting

RFC 2975 [RFC2975] describes accounting as "The collection of resource consumption data for the purposes of capacity and trend analysis, cost allocation, auditing, and billing.".

Over the years the number of information elements that can be sent from an accounting client to an accounting server using standardized protocols, such as RADIUS (see [RFC2866] and [RFC2865]) and Diameter [RFC3588], has increased. The existence of standardized protocols has allowed different AAA networks to interconnect. 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.

Although they are popular, RADIUS and Diameter are not the only protocols that can be used to collect usage information and to trigger responses. Other approaches, as documented in [I-D.livingood-woundy-congestion-mgmt], lead to similar results.

2.2. Deep Packet Inspection

Deep packet inspection (DPI) refers to the observation and analysis of traffic that passes through operator networks up to the application layer. This allows operators to determine the applications and/or application-layer protocols that subscribers are using and respond on a per-application or per-protocol basis.

The process of inspecting traffic, particularly in real time, can be highly performance-intensive. DPI equipment may also require continous software updates to ensure that the detection engine recognizes the latest protocol variants.

There may be a number of other factors that contribute to a network operator's decision to use DPI, including potential user backlash, privacy impact, and legal concerns.

Depending on the configuration of the device doing the inspection, packet dropping/blocking and other usages might be applied. For example, content sharing p2p applications maintain many simultaneous

TCP connections with other nodes for the purpose of simultaneous downloads. An operator may, for example, limit . Certain end user contracts may also allow operators to ban servers from residential access.

3. Network Operator Responses to Congestion

Once they have collected congestion information using either of the techniques described above or others, network operators have a number of options for how to respond. For all of these options, it is up to the operator to decide the breadth and depth of its response: which users will be affected, the time frame in which congestion will be managed, whether specific applications or protocols will be targeted,

and so forth. Operators can chose from both technical and pricing/ contract-based options. Technical options include:

Wholistic traffic shaping:

End user contracts often provide users with a certain threshold for baseline usage volume (which is typically quite high). Subsequently, if consumption goes beyond the threshold, all of the user's traffic is given reduced priority vis a vis other users on the network. Some operators may only shape traffic during times of congestion or peak usage periods (even if a user has exceeded his or her baseline threshold).

Per-application or per-protocol shaping:

Network operators that can identify particular applications or protocols creating congestion may decide to throttle only those applications or protocols. They may also take indirect steps that result in the shaping of only certain applications, such as limiting the number of simultaneous TCP connection setups from a single subscriber (to handle peer-to-peer traffic), or preventing users from hosting servers on residential connections. An example of an ISP's fair usage policy describing how it manages specific protocols is included in Appendix A.

Class-Based Assignment:

In this technique users are classified into a set of classes depending on their past behavior. Subsequently, their traffic is treated according to their associated classes. This may prevent lightweight users from feeling the effects of sharing network capacity with heavy users. This mechanism requires some form of packet marking to be able to differentiate light users from heavy users.

Pricing/contract-based options include:

Charging for Excessive Traffic: Network operators may charge

users differently for traffic that exeeds a certain threshold compared to the traffic that falls below the threshold.

Suspending or Discontinuing Contracts: In some rare cases ISPs may decide to suspend or terminate the contracts of heavy users. In some cases this response may be associated with a security issue; when an operator recognizes a botnet-infected machine generating excessive traffic, it may hotline all the traffic of that particular machine to a separate network, and ultimately suspend or terminate the machine's connection. In some cases the same technique has been applied to users engaged in heavy P2P usage, either intentionally or due to a false alarm caused by a statistical traffic analysis.

4. New Activities

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 resarchers 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. Summary

High-consuming users are a reality. Operators that would like to counteract the impact of heavy users 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 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).

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, 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-00 (work in progress), August 2009.

[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-01 (work in progress), September 2009.

[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>.

[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.

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