Dealing with P2P Traffic in an Operator Network: State-of-the-Art

Hannes Tschofenig
Nokia Siemens Networks
hannes.tschofenig@nsn.com

Marcin Matuszewski
Nokia
marcin.matuszewski@nokia.com

ABSTRACT

The availability of broadband connections together with flat-rate pricing has made new types of peer-to-peer applications possible. From an Internet evolution and end user value point of view this is very exciting. As a consequence, an increase of user-to-user traffic was observable all around the world over the last few years. This, however, has lead to problems as well. Certain type of p2p applications upset the music & movie industry and again other aspects leave puzzled regulators behind. This paper focuses on small subset of aspects introduced by the usage of p2p systems, namely the observation that a certain group of users, called heavy hitters, decided to use their flat-rate contract excessively. This in turn seems to have caused network operators to take actions to deal with problems caused by this small group of people. This document illustrates a couple of techniques used by operators today to deal with excessive bandwidth usage. This should serve as a basis for the reader to judge whether state-of-the-art mechanisms are insufficient to deal with the perceived problem.

1. INTRODUCTION

In a recent publication by K. Cho et al. [2] about the growth of residential user-to-user traffic in Japan indicates that '... a small number of users dictate the overall behavior; 4% of heavy hitters account for 75% of the inbound volume, and the fiber users account for 86% of the inbound volume.'. The same paper also indicates a substantial increase in traffic growth, namely 37% per year according to [2], and not just a different distribution of traffic among the users. Furthermore, 63% of the residential traffic volume is contributed by user-to-user traffic.

These numbers itself do not represent a problem as such. However, some operators very likely had different expectations about the growth rates and traffic consumption of individual users and statistics (used for their pricing models) did not work out too well. The profit margins for Internet access are quite slim due to fierce competition. This puts a lot of pressure on operators to deal with customers, namely these heavy hitters, who cost a lot of money. Finally, some broadband networks may not have the ideal characteristics (such as the topology for routing traffic) for user-to-user traffic.

The existence of flat rate pricing contributes to some of the problems since the bandwidth usage in total needs to be covered by the money obtained from broadband customers but the usage of individual users is not reflected in the amount. As such, users that rarely utilize the network pay the same amount as someone who uses filesharing applications all day long. When increasing the capacity of a network, as described in Appendix A of [1], then the additional bandwidth is again largely eaten by heavy hitters. Thoughts like these have triggered the discussions around the meaning of 'fairness' (also in the context of TCP congestion control).

Operators in many countries have decided to take actions to deal with heavy hitters in order to ensure that congestion does not happen within their networks, users utilizing aggressive p2p applications do not push away traffic from other users, and that costs for leased lines from the ISP to the provider of the last mile is minimized. In any case, it appears that some operators have taken some actions to ensure that excessive traffic does not impact their regular network operation.

Similar problems are also expected in mobile networks. Some mobile operators that already launched flat rate data plans experienced similar usage patterns as in the fixed networks.

In the cases studied in this paper the actions focus more on the heavy hitters. As such, the aspects typically of concern in the context of network neutrality discussions, see for example [3], are not directly of relevance to this paper.

2. STATE-OF-THE-ART

Before starting new standardization work in the IETF, or in other standardization organizations, it is important to understand what is currently being done to deal with the perceived problems. The following list of solutions is provided without judgment whether they are better or worse than other solutions.

2.1 Volume Accounting

This is a classical approach that is still widely implemented in mobile networks today but pure volume based accounting is far less common in broadband access environments. It appears that customers in many countries prefer flat-rate pricing schemes because of their predictability.

2.2 Shaping

Often contracts offer a combination of 'flat-rate' scheme whereby a fixed price tariff is used up to a certain usage volume (typically quite high for regular usage). Subsequently, if the consumption goes beyond a certain threshold then further actions are taken. Contracts typically contain statements that support these actions. An example of such a fair use statement can be found in ẽfepolicy.

In many countries operators have to offer a clear description of the service they offer and since the term 'flat-rate' is already associated with a certain meaning the term 'Unlimited Data Rate' is often used for this type of service.

Note that traffic shaping is often applied only to heavy hitters (since they are known to the operators due to the internal accounting system) and often only applied during peak hours.

2.3 Deep Packet Inspection

This technique refers to inspecting traffic that passes through the operators networks and to determine the type of traffic up to the application layer. In case of P2P traffic, such as filesharing usage, rate limiting is applied when necessary. Typically, this approach has been quite expensive in terms of CAPEX and OPEX. Additionally, the attempt to selectively deal with applications (even though these applications might be the reason for the high traffic volume) has not been received well by the users and was highlighted in the media.

2.4 Class-Based Mechanism

This is a more sophisticated technique that does not seem to be used widely but interesting to mention. In this technique users are classified into a set of classes depending on their past behavior. Subsequently, the traffic is treated according to the associated class. It is ensured that the traffic of lightweight users, users that really rarely use their Internet connection, cannot be pushed away by heavy users.

2.5 Limiting Subscriber Flows

Typically content sharing p2p applications maintain many simultaneous connections with other nodes for p2p network maintenance and simultaneous download of content. A small number of operators limit the p2p traffic in their networks by limiting a number of connection setups from a single subscriber.

2.6 Banning Servers from Residential Access

Some operators are banning servers from residential access. There can be many actions that an operator can pursue by detecting a subscriber running a server. The actions are typically being pursued on heavy hitters. This technique is typically supported by the type of contracts customers signed.

2.7 Maintaining Super Peers or/and Content Caches

There are some attempts to optimize the p2p traffic by placing super peers close to the subscribers and cashing popular content. These attempts, however, have potential pitfalls related to copyright issues.

There are many actions that operators can pursue on heavy hitters to lower the costs minimize heavy-heaters impact on quality of service perceived by regular users.

2.8 Discontinuing Contracts

Some operators have decided to discontinue contracts with heavy hitters. From a legal point of view these actions are often supported by the type of contracts customers signed. However, cases in countries like Austria or Germany have caused a significant amount of bad press for the companies exercising these techniques.

2.9 Blocking Traffic

In some rare cases ISP also decided to block traffic if it exceeds a certain limit. Essentially, the connectivity is cut. In fact this might be justified in certain cases. For example, in case of new botnets malware distribution when the operator recognizes an infected machine and hotlines the entire traffic of that particular machine to a separate network and, like in WLAN hotspots, HTTP traffic is intercepted to display further information. In some cases the same technique has been applied with excessive usage of P2P traffic, either intentionally or due to a false alarm caused by a statistical traffic analysis technique.

2.10 Lowering Quality of Service

Using class-based mechanisms heavy hitters can be put into a separate class where they compete among themselves among the available resources. Operators may also apply deep packet inspection and in case of P2P traffic, such as file sharing usage, limiting data rates.

2.11 Charging for Excessive Traffic

As a possible action the user might get charged for excessive traffic or traffic is shaped to a lower transmission rate.

3. CONCLUSIONS

Almost all of the operators the authors have spoken with came to the conclusion that there is a problem with heavy hitters in their network. A few also indicate that their main solution is to throw more bandwidth into the network.

Operators already have a fair number of tools available to deal with heavy hitters. Even though they are technically quite simple they are applied widely today. Some with more success than others. Those who were asked by us interestingly argued that this is largely a business and marketing problem rather than a technical problem.

It is important to also understand that the business impacts are very different for the players in the individual markets. In some European countries there is quite tough competition when it comes to the last mile providers and hence operators might find more restrictions being placed on them. Additionally, operators with some degree of legacy infrastructure might find their networks not well suited for P2P traffic when it comes to the cost associated with user-to-user traffic.

In any case, it seems to be wise to study the currently deployed techniques in much greater detail to understand the business side (such as, OPEX and CAPEX costs), the economical environment and the techniques themselves. Then, one could make a much better decision whether there is (1) a need to go beyond the state of the art (and a need for standardization) and (2) whether there is interest by operators to look at additional tools that they would be able to deploy. From past experience in standardization work there is one important design aspect that has to be kept in mind, namely to offer a good incremental deployment story. This

is easier to accomplish when both the benefits and the costs are kept local rather than distributed throughout different stakeholders.

4. APPENDIX: EXAMPLE OF POLICY STATEMENTS

4.1 Fair Usage Policy

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%).

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.

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.

4.2 Peer to Peer (P2P)

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.

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.

Acknowledgments

The authors have had chats with 20+ operators to determine their current way of working. Most of the operators are from Europe but there are some from the US as well. The authors would like to thank them for their feedback.

Hannes Tschofenig additionally would like to thank Bob Briscoe, Louise Burness, and Toby Moncaster for their time to discuss reECN.

5. REFERENCES

[1] B. Briscoe, T. Moncaster, and A. Burness. Problem Statement: We Don't Have To Do Fairness Ourselves. Internet-Draft draft-briscoe-tsvwg-relax-fairness-00, Internet Engineering Task Force, Nov. 2007. Work in progress.

[2] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. *SIGCOMM Comput. Commun. Rev.*, 36(4):207–218, 2006.

[3] S. Marcus. Network neutrality: The roots of the debate in the united states. *Intereconomics: Review of European Economic Policy*, 43(1):30–37, January 2008.