Congestion-based Accounting with re-ECN

Mirja Kühlewind
Institute of Communication Networks and Computer Engineering
University of Stuttgart, Germany
mirja.kuehlewind@ikr.uni-stuttgart.de

Wolfram Lautenschläger and Michael Scharf
Alcatel-Lucent Bell Labs, Stuttgart, Germany
{wolfram.lautenschlaeger, michael.scharf}@alcatel-lucent.com

Making end-systems accountable for the congestion they cause will give an incentive to control one's congestion appropriate. re-ECN is a new Transmission Control Protocol (TCP) mechanism to expose the expected congestion on a network path. Based on such an announcement, the amount of congestion one end-system is allowed to introduce into the network can be limited. In the re-ECN framework a policer at network ingress is proposed which drops packets in congestion situations, if no so-called congestion credits are available. However, this will only allow to limit congestion of upstream traffic. If data is requested by a client, the server will not be able to decide about an appropriate data handling as the intention of the client is not known. To address this problem an architecture to transfer congestion credits from the client-side policer to the server-side policer could be used. This paper suggests and discusses solutions for such an accounting system.

1 Introduction

The amount of traffic in the Internet is more and more growing due to the increase in popularity of peer-to-peer file sharing, video streaming and other data-intensive services. Thereby only a few so called heavy users introduce the major part of the traffic. Internet Service Providers (ISPs) try to counteract by limiting the data rate or the traffic volume of such heavy users [1] or by distinguishing data-intensive services through Deep Packet Inspection. However, such approaches will block transmission requests even if enough network resources are available. Hence, policing would only be necessary if the available network resources are exhausted and congestion occurs.

Currently, there is an activity in the Internet Engineering Task Force (IETF) to achieve Congestion Exposure [2]. The object is to announce the congestion an end-system or a network component is expected to introduce into the network by sending or forwarding data. To expose this congestion information a mechanism called re-ECN [3] is proposed, which is based on the ECN [4] mechanism. A re-ECN sender will reinsert the ECN congestion information into the network as an estimation of the expected congestion. Given this, a sender can be made accountable for the congestion it is causing on a network path. Moreover, the introduction of a per-costumer congestion limit could provide an incentive for the end-system to not cause congestion unnecessarily. This could be achieved though appropriate congestion control in the end-system depending on the current congestion situation as well as certain application requirements and the user's intention.

However, if a client initiates a download from a webserver, the server-side will be accounted for the congestion caused. As the webserver is not the connection initiator, it does not know how to decide properly upon a congestion announcement. The reaction should depend on the intention of the client. The transfer of so called congestion credits from the client-side to the server-side would allow the client to assign a higher priority to a certain transmission, as those credits can be used from the server-side policer to conserve a certain data rate in a congestion situation. Of course, such a mechanism does not guarantee any data rate, but it will show the user-defined importance of a data transmission to the server. This could encourage the use of less aggressive congestion control mechanisms for e. g. background traffic as well as new high-speed congestion control mechanisms for time-critical transmissions.

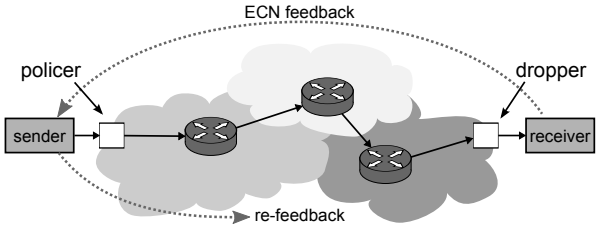

Figure 1: Re-feedback mechanism and re-ECN framework components

2 The re-ECN Protocol

re-ECN extends the ECN mechanism in such a way that the ECN feedback information is used to signal the expected congestion of a network path into the network. This basic re-feedback mechanism is shown in Fig. 1. The receiver will count the number of packets marked as congestion experienced (CE) by routers which use Active Queue Management (AQM) like Random Early Detection [5]. The current counter value is communicated in every TCP acknowledgement (ACK) packet back to the sender. For every counter increment the sender will as well mark an outgoing packet as re-echoed congestion (RE). The fraction of RE marked packets will expose the expected whole-path congestion for this data flow.
To ensure that a sender declares its expected congestion honestly it has been suggested to deploy a dropper at the network egress node [6]. At this point on the network path one should see as much RE marked packets (credit) as CE marked ones (debit). If the sender understates its congestion, the dropper will detect the absence of RE marks and start to punish the respective flow through packet drops.

The benefit of re-injecting the congestion information is that every intermediate node is informed about the congestion level along the path, i.e. not only the congestion that a packet suffered so far (upstream, would be possible with ECN alone) but also the anticipated congestion downstream. This information enables all stakeholders along the path (end system owner, access provider, transit network etc.) to build up technical mechanisms (policer, traffic shaper) but also economical mechanisms (charging, accounting) to deal with congestion. In this paper we investigate an end system centric application of the protocol.

3 Congestion based Traffic Engineering

A policer in the ingress is suggested which limits the amount of congestion marks one end-system can assign to its flows [7]. To handle the congestion credits in the policer a token bucket mechanism is proposed [8]. Such a mechanism has already been proposed with a former congestion charging approach [9]. Thereby, a certain token rate and a bucket size will be assigned to every end-system. The token rate determines when new credit points will be created while the bucket size gives a maximum number of credit points that can be stored if not used immediately. The policer will drop RE marked packets if no congestion credits are available anymore. Such a mechanism will give a strong incentive to save congestion credit points for later use and thus to make appropriate congestion control. A certain amount of congestion credits might also be necessary to balance congestion peaks. The token bucket mechanism thus provides a minimum access rate as well as continuous refill of the backup volume for high congestion situations.

Still there are many open research challenges. Regarding the token bucket approach, it still needs to be figured out how to set the token rate and the bucket size. If the token rate is too high or the bucket size too large, there is no incentive to react on congestion announcements at all as the bucket will never be completely empty. If the token rate is too low or the bucket size too small, the refill might be too slow to collect enough credit points to balance a congestion peak. Thus packets might get dropped even if the end-system reduced the sending data rate.
Furthermore, the parametrization might vary for costumers with different contracts or in different domains. The parameterization could be usage-dependent or simply be a part of a flatrate contract.

To achieve a more equal distributed resource allocation in the Internet, agreements on the caused congestion between the ISPs might be necessary as well. Such a decision might be influenced by the number of users in a domain, e. g. their traffic patterns and/or online time, as well as the congestion patterns at peering points or components close to the destination network. Further research is needed here.

4 Accounting Architecture

It seems reasonable to integrate the policer into the access router (AR), which is typically the first IP hop in an ISP network. To meet the various demands in a complex provider scenario, e. g. with dail-up or mobile user and advanced policer settings it is advisable to link the policer functionality to a respective user account, e. g. in a AAA (Authentication, Authorization, Accounting) server.

Having a download scenario the server-side normally does not know which transmissions are impor-

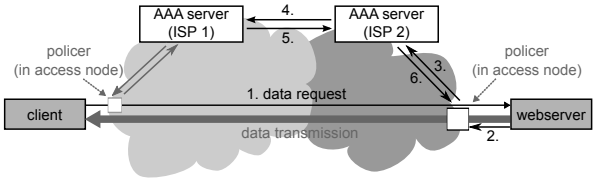

Figure 2: Congestion Accounting architecture and communication sequence

tant to a client and when congestion credits should be spent. Our solution to this proplem is to transfer a certain amount of congestion credits from a client's account to the server-side policer. Hence, the server application can spend these credits to implement a higher data rate for this respective data transmission in a congestion situation. Figure 2 shows one potential architecture. This approach involves communication between the end-system and its policer, between the policers in the AR and the AAA servers and among the AAA servers of different ISPs.

The architecture implies a trust relationship between different ISPs and new communication protocols between the end-system and the access router. The workshop contribution will discuss the implications of the proposed architecture on today's communication structure.

5 Conclusion and Outlook

In this paper we have shown that Congestion Accounting gives an incentive for appropriate congestion control and will provide a basis to enable new congestion control mechanisms for e. g. background traffic or high-speed transmissions. The already proposed framework based on the re-ECN mechanism for Congestion Exposure offers a solution to make the sender accountable for its congestion. In a client-server connection this is not sufficient as the server-side can not know about the client's intention. Therefore, we presented a reference accounting architecture for client-based congestion credit provisioning.

The final contribution to the workshop will include a comprehensive introduction of the re-ECN protocol architecture, an analysis how the protocol can be integrated into an overall accounting architecture, and an outlook to the potential impact on the Future Internet traffic management.

6 Achnowledgement

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n° 247674. Furthermore, this work is partially funded by the German Research Foundation (DFG) through the Center of Excellence "Nexus – Spatial World Models for Mobile Context-Aware Applications" .

References

[1] C. Bastian, T. Klieber, J. Livingood, J. Mills, and R. Woundy, "Comcast's Protocol-Agnostic Congestion Management System," internet draft, IETF, 2009.

[2] T. Moncaster, L. Krug, M. Menth, J. Araujo, and R. Woundy, "The need for Congestion Exposure in the Internet," internet draft, IETF, 2009.

[3] B. Briscoe, A. Jacquet, T. Moncaster, and A. Smith, "Re-ECN: Adding Accountability for Causing Congestion to TCP/IP," internet draft, IETF, Sept. 2009.

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

[5] S. Floyd and V. Jacobson, "Random Early Detection gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, pp. 397–413, Aug. 1993.

[6] B. Briscoe, A. Jacquet, C. Di Cairano-Gilfedder, A. Salvatori, A. Soppera, and M. Koyabe, "Policing Congestion Response in an Internetwork using Re-feedback," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 4, pp. 277–288, 2005.

[7] B. Briscoe, A. Jacquet, T. Moncaster, and A. Smith, "Re-ECN: The Motivation for Adding Congestion Accountability to TCP/IP," internet draft, IETF, Sept. 2009.

[8] A. Jacquet, B. Briscoe, and T. Moncaster, "Policing Freedom to Use the Internet Resource Pool," in *Proc Workshop on Re-Architecting the Internet (ReArch'08)*, ACM, Dec. 2008.

[9] D. D. Clark, "Internet cost allocation and pricing," in *Internet Economics* (L. W. McKnight and J. P. Bailey, eds.), ch. Usage Sensitive Pricing, pp. 215–252, MIT Press, Cambridge, MA, 1997.