Efficiency of PCN-Based Network Admission Control with Flow Termination

Michael Menth, University of Würzburg, Institute of Computer Science, Germany

***Abstract*— The pre-congestion notification (PCN) architecture marks packets based on the utilization of links and gives early warnings before congestion occurs. The IETF suggest to implement network admission control (NAC) and flow termination based on this feedback to enforce quality of service (QoS) under failure-free operation and the manageability of the network in failure scenarios. Admission control decision are taken only at the border routers. In case of link or node failures, traffic is just rerouted and no reservation states need to be recovered inside the network. Therefore, PCN-based NAC can be operated in a resilient mode for a set of protected failure scenarios if sufficient backup capacity is reserved to accommodate the redirected traffic. This paper presents ongoing work in the IETF, relates it to the theory of resilient NAC, and investigates the efficiency of the new approach.**

I. Introduction

Internet service providers (ISPs) recently offer increased access speeds, e.g., by digital subscriber lines (DSL), cable TV (CATV), and fiber to the home (FTTH). Their advent significantly increased the traffic volume in carrier networks and in 2005, the major traffic in Japan was already produced by residential users [1]. Popular video services like YouTube produce large traffic volumes, but are only weak precursors of high-quality IP-TV services. These present a challenge for ISPs which need to offer triple play, i.e. the integration of the transport of data, voice, and video. However, the resource management for triple play becomes more and more difficult due to the emerging interactive Web2.0 since residential users also become content providers. In particular, [2] has shown that normal users get accustomed with new services, change access technologies, and become "heavy hitters" such that the majority of the overall traffic is produced by a minority of residential users.

Today, ISPs rely on capacity overprovisioning (CO) to enforce quality of service (QoS) in terms of packet loss and delay. However, triple play requires guarantees that cannot be given by CO [3]. In [4] admission control (AC) was proposed for IP networks, but so far such techniques are only applied locally, they are rarely in use, and not deployed in core networks. If congestion occurs in core networks, this is mainly due to failures and redirected traffic, and only to a minor degree due to increased user activity [5]. Thus, both AC and CO require backup capacity that can be used under failure-free conditions to improve the transmission quality. Taking this into account, CO seems a viable alternative to AC in practice for networks with static traffic [6]. However, the above sketched dynamic behavior of users and services leads to an unpredictability of future demands such that QoS provisioning remains difficult. Therefore, ISPs see the need for AC to offer premium services over integrated IP networks in the future.

As a consequence, the Internet Engineering Task Force (IETF) has recently set up a working group called "Congestion and Pre-Congestion Notification" (PCN) [7] with the objective to standardize a measurement-based network admission control (NAC) for the Internet. It is compatible with a differentiated services architecture (DiffServ) [8], [9] in the sense that admitted flows receive a premium service. In contrast to the integrated services architecture (IntServ) [10] no per-flow signaling and reservation is performed by interior routers. Thus, core routers do not need to keep per-flow states, but measure and possibly mark the traffic per outgoing interface. The architecture is simple and only border routers need to take admission decisions. Therefore, if interior routers fail, traffic is rerouted, but no reservation states need to be recovered. Since PCN allows to allocate only a fraction of the link bandwidth, backup capacity can be reserved for redirected traffic during network failures, such that the PCN architecture can be configured as resilient NAC. As it is difficult for applications to indicate appropriate traffic descriptors for their generated flows, only a rough upper bound can be given in practice. Since the PCN-based admission control is based on measurements, the overestimated traffic descriptors are handled by implicit overbooking since admission decisions rely mainly on the feedback from the network. This opens the door for overadmission, i.e. for falsely admitted flows, especially in the presence of flash crowds which may be observed, e.g., before video transmission of mass events. To avoid severe service degradation in such cases and during severe failure scenarios, PCN-based NAC is accompanied by PCN-based flow termination which may take into account preemption priorities of flows.

The contribution of this paper is a simple description of the PCN architecture that is currently discussed in the IETF. As the PCN work is only in an early stage, the PCN architecture and its nomenclature are not stable yet. Therefore, we choose our own wording for the sake of simpler explanations. We relate the PCN-based NAC to the theory of resilient and non-resilient NAC and show that it is more efficient than existing solutions.

The paper is structured as follows. Section II presents the PCN architecture. Section III reviews related work which shows the historic roots of PCN and motivates the current

This work was funded by Deutsche Forschungsgemeinschaft (DFG) under grant TR257/18-2. The author alone is responsible for the content of the paper.

NAC approach. Section IV explains the theory of NAC and compares the efficiency of the PCN-based NAC with flow termination to other approaches. Section V summarizes this work.

II. Network Admission Control and Flow Termination Using Pre-Congestion Notification

In the following we explain NAC and flow termination based on PCN for deployment in DiffServ domains as suggested in [11]. The objective is to provide a controlled load (CL) service which offers the same QoS a flow would receive from lightly loaded network elements [12] and which is useful for inelastic flows, e.g., realtime media. Such a DiffServ domain is also called a CL region. The approach is in line with the "IntServ over DiffServ" framework [9]: the CL service is supported end-to-end by RSVP signalling [13] on a per hop basis, but the interior nodes of the DiffServ domain are bypassed and the DiffServ border routers act as neighboring hops for RSVP signalling.

A. Network Admission Control

NAC is required to control the load of an administrative domain and a domain applying PCN for NAC is called a PCN domain. QoS is signalled for each flow on an end-to-end basis by resource reservation protocols like RSVP [13]. Only the ingress and egress gateways of the CL region participate in this signalling, i.e., they admit or reject reservation requests of flows and hold per-flow reservation states. Interior routers of the PCN domain are unaware of individual microflows. Thus, only PCN ingress and egress gateways perform policing and traffic shaping if end-to-end signalling relies on traffic descriptors.

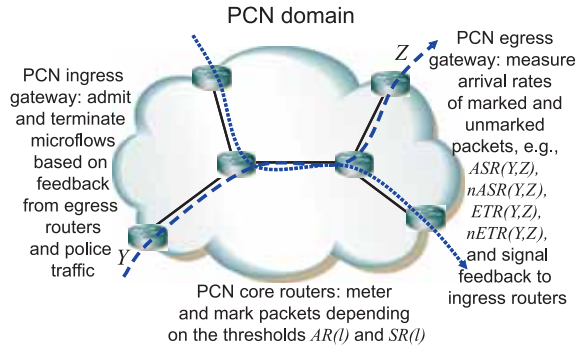

Fig. 1. PCN ingress gateways of a PCN domain take flow admission and termination decisions, core routers perform packet metering and marking, and egress gateways measure the rates of marked and unmarked packets.

The PCN architecture implements NAC for the CL region. PCN ingress gateways take AC decisions that are supported by PCN core routers metering and marking packets and by PCN egress gateways measuring marked and unmarked packets and providing feedback to PCN ingress gateways. Each link l of the PCN domain is associated with a configurable threshold for the admissible rate ($AR(l)$). When a link l carries more CL traffic (with rate $CL(l)$) than its admissible rate $AR(l)$, additional flows should not be admitted to the network. This is signalled as follows.

Each link l of the CL region is monitored and all packets are admission-stop marked in case that the CL traffic rate $CL(l)$ exceeds the admissible rate $AR(l)$. If the traffic is smooth, $CL(l)$ is either below or above $AR(l)$ and no or all packets are admission-stop marked, but bursty traffic may lead to oscillations in admission-stop marking. Each egress gateway Z estimates the bit rate of packets with and without admission-stop mark that are received from each ingress router Y. These measurements are performed based on a moving average and stored in the time-dependent variables $ASR(Y,Z)$ and $nASR(Y,Z)$. If the fraction $\frac{ASR(Y,Z)}{ASR(Y,Z)+nASR(Y,Z)}$ exceeds a certain threshold, egress gateway Z signals admission-stop to ingress gateway Y to avoid admission of additional flows from Y to Z. If the fraction drops below this threshold, the ingress gateway may continue to admit new flows. To avoid oscillations, a hysteresis should be applied.

B. Challenges of PCN-Based NAC

The PCN-based NAC is simple, but difficulties arise under some extreme conditions.

1) Impact of Equal-Cost Multipath Routing: In the presence of equal-cost multipath (ECMP) routing, the traffic is equally distributed to those outgoing interfaces that are part of a least-cost path. The split is done on a per flow basis to avoid packet reordering due to different transmission delays of the paths. The outgoing interface for a specific packet is determined by a hash function based on its header values that are invariant for all packets of the entire flow. This makes the traffic distribution stochastic [14], [15].

If ECMP routing is used in a PCN domain, it is possible that a core router of just one partial path marks packets with admission-stop while core routers of other partial paths do not mark them with admission-stop. Then, a certain fraction of the packets arrive with an admission-stop mark at the egress gateway such that this PCN egress probably signals admission-stop for the entire aggregate to the respective PCN ingress gateway. Thus, admission for a new flow is denied even if it will be carried on the lightly loaded path.

Probing can potentially avoid this problem. The PCN ingress issues a packet train with the same flow information (source and destination address and port) and the PCN egress router intercepts these packets and returns them to the PCN ingress. If these packets are admission-stop marked, the prospective path of the flow is already saturated and admission must be denied; otherwise, the new flow can be accepted.

2) Impact of Flash Crowds: As long as the PCN ingress router Y is not informed by egress router Z to stop the admission of additional flows from Y to Z, it can admit new flows. In the presence of a flash crowd a large burst of admission requests is accepted at once which can overload the links significantly. Several approaches can improve the handling of flash crowds: limiting the number of admitted flows per time, probing at the desired flow rate until the

application effectively sends data, but other approaches are also possible.

3) Impact of Delayed Transmission: We assume a realtime video transmission of a mass event. Many users log on the system before the event, but the application does not send any data. As a consequence, too many flows can be admitted such that the links are overloaded as soon as the transmission starts. This situation is similar to a flash crowd, but limiting the admission rate does not help.

4) Impact of Increased Transmission Rates: Flows in progress may increase their transmission rate. This can happen due to QoS attacks or due to realtime streaming of popular contents that in fact increases its transmission rates. If the rate increase of the admitted flows is very strong, it can lead to congestion.

5) Impact of Traffic Redirection due to Failures: If network failures occur in the PCN domain, traffic is rerouted. This redirected traffic can quickly lead to congestion on the backup paths.

C. Flow Termination

As outlined above, PCN-based NAC cannot avoid congestion under some extreme conditions. Therefore, a flow termination function is desirable to reenforce the CL service for non-terminated flows at the expense of some terminated flows, and the PCN architecture supports this function [11]. A configurable threshold for the supportable rate $SR(l)$ is associated with each link l which is usually significantly larger than its configurable threshold for the admissible rate $AR(l)$. If a link l carries more CL traffic than its supportable rate $SR(l)$, the ingress routers should take appropriate countermeasures to reduce the rate $CL(l)$ on that link. This is signalled as follows.

PCN core routers monitor the CL packet streams on each link and mark those packets exceeding the supportable rate $SR(l)$ with an excess-traffic flag. This may be done, e.g., based on a token bucket mechanism. The egress gateway Z determines for each ingress gateway Y the excess traffic rate $ETR(Y,Z)$ and the non excess traffic rate $nETR(Y,Z)$, i.e., the rate of packets received with and without excess-traffic flag. If packets of the aggregate from Y to Z are excess-traffic marked, ingress gateway Y should terminate some flows of its aggregate towards Z. The excess traffic rate $ETR(Y,Z)$ serves as an estimate for the overall rate of the flows that need to be terminated and is signalled from the egress gateway Z to the ingress gateway Y. The selection of the flows can be done by the ingress gateway Y based on preemption priorities. However, in the presence of multipath routing, it is crucial to terminate those flows that are effectively routed over the bottleneck link. After the termination of some flows, the rate of the bottleneck link l drops below its supportable rate $SR(l)$ and packets are no longer excess-traffic marked.

III. Related Work

In the following, we review related work regarding active queue management (AQM), explicit congestion notification (ECN), and marking-based admission control.

A. Active Queue Management

Active queue management (AQM) techniques differentiate packet loss among flows if the buffer space for an outgoing interface is about being exhausted. Different space priority schemes have been discussed and evaluated in [16], [17]. The drop tail approach is the simplest queue management scheme: packets are accepted as long as buffer space suffices, otherwise they are dropped; no loss differentiation is provided. As an alternative, buffer space can be allocated to different service classes: packets of class *A* exceeding the buffer space of class *A* are dropped to guarantee that packets of other classes do not suffer from the increased load of class *A*. Furthermore, a queue management similar to a Russian dolls model (RDM) [18] is possible.

The random early detection (RED) gateway was originally presented in [19], and in [20] it was recommended for deployment in the Internet. It was designed to detect incipient congestion by measuring the average buffer occupation in routers and to take appropriate countermeasures. That means, packets are dropped or marked to indicate congestion to TCP senders. RED was mostly combined with congestion state dependent random packet drops such that RED became a synonym for this buffer management strategy. RED calculates the average queue length using an exponentially weighted moving average (EWMA) to filter sudden increases due to traffic bursts. Hence, the average queue length avg is updated every time a packet arrives by $avg = (1-w_q) \cdot avg + w_q \cdot q$ with q being the current buffer occupation and w_q the weight parameter. The router uses the variable avg to determine the probability for random packet drops by the function given in Figure 2. The packet loss probability is typically zero for small values of $avg \leq r_{min}$ and increases then linearly up to a certain maximum value p_{max} for an average utilization of r_{max}. If the average queue length avg is larger than r_{max}, all packets are discarded. Several improvements have been suggested, e.g., to achieve fairness in the presence of non-adaptive connections and to introduce class priorities [21], [22]. Note that AQM techniques operate on the physical queue size which is unlike PCN-metering.

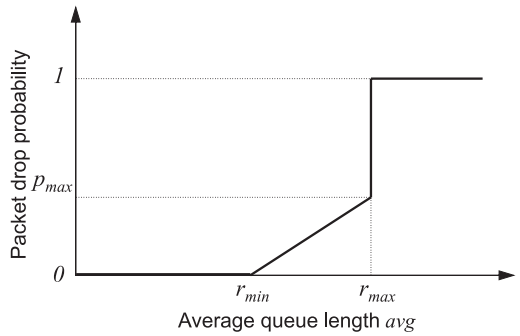

Fig. 2. RED discards packets early depending on the average buffer occupation.

B. Explicit Congestion Notification

Explicit congestion notification (ECN) is built on the above idea to signal incipient congestion to elastic TCP senders to reduce their sending window [23]. Packets of ECN-enabled TCP flows are labelled with an ECN-capable transport (ECT) flag such that RED gateways mark the packets with a congestion experienced (CE) value rather than discard them. This improves the TCP throughput since packet retransmission is no longer needed. Both the ECN marking and the behavior of senders and receivers after the reception of a marked packet is defined in [23]. It proposes to use the CU (currently unused) bits of the differentiated services codepoint (DSCP) in the IP header which is a redefinition of the type of service octet [24] for the encoding of the ECT, not-ECT and CE flags. These two bits may be reused for admission-stop and excess-traffic marking in the PCN architecture [11]. An appropriate DSCP indicates that they need to be interpreted as PCN marking instead of ECN marking.

C. Admission Control Based on Packet Marking

Admission control based on packet marking has been proposed in pervious work. Gibbens and Kelly [25] theoretically investigated AC based on the feedback of marked packets whereby packets were already marked by routers based on a virtual queue with configurable bandwidth. This enables early warning which is the core idea of pre-congestion notification. It also allows to limit the utilization of the link bandwidth by premium traffic to arbitrary values between 0 and 100%. Karsten and Schmitt [26], [27] integrated these ideas into the IntServ framework and implemented a prototype. They point out that the marking can also be based on the CPU usage of the routers instead of the link utilization if this turns out to be the limiting resource for packet forwarding.

IV. Efficiency of Resilient Network Admission Control Methods

We first give a short introduction to resilient NAC and then study the efficiency of different resilient and non-resilient NAC methods based on the performance evaluation method developed in [28].

A. Network Admission Control

In [28] admission control methods were subdivided into link and network admission control.

Link admission control (LAC) limits the traffic on a single link. This can be done, e.g., by bookkeeping the rates of admitted microflows. These flows require traffic descriptors to predict the expected packet loss and delay by queuing formulae. Such equations have been collected in [29] in the context of ATM systems for different traffic types. Based on such calculations, parameter-based LAC decides whether a new call can be accepted. Measurement-based LAC measures the actual traffic on the link to take the admission decision [30], [31].

In contrast, network admission control (NAC) limits the amount of high quality traffic for an entire network and avoids congestion on all links. To that end, the routing of the microflows is taken into account either implicitly or explicitly. NAC methods can be classified with respect to their flexibility regarding the use of the network resources. We briefly review two major NAC categories: border-to-border (b2b) budget (BBB) based NAC and link budget (LB) based NAC.

1) BBB NAC: Label switched paths (LSPs) may be set up using MPLS technology and the capacity of these tunnels can be prereserved. Then, an LSP looks like a fixed capacity pipe from the LSP ingress to the LSP egress, and the LSP ingress can perform admission control for microflows using LAC methods. This is depicted in Figure 3. If the LSP's capacity is fully allocated, no further microflows can be admitted. Abstracting from this example, we can talk of a b2b budget (BBB) $BBB(Y,Z)$ between ingress Y and egress Z. The capacity of $BBB(Y,Z)$ can be used only by microflows going from Y to Z, and if this capacity has been fully allocated, no further microflows from Y to Z can be admitted even if all links of the respective path have still unused capacity.

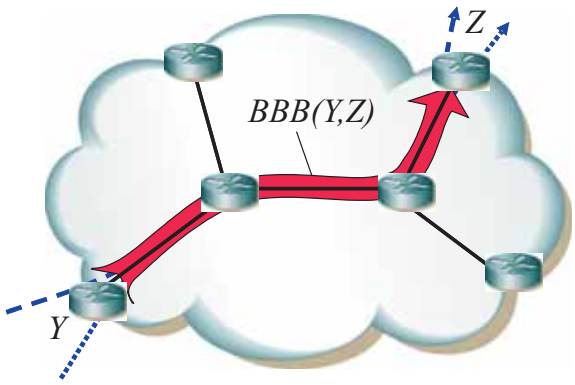

Fig. 3. The capacity of b2b budgets can be used only by the microflows of the corresponding b2b aggregate.

2) LB NAC: The most intuitive NAC is a link-by-link application of LAC methods. For instance, microflows use RSVP signalling [13] to install per-flow reservations in the routers for all links along their paths. Abstracting from this example, each link l is associated with a link budget (LB) $LB(l)$. The capacity of $LB(l)$ can be used for any microflow going over the link l. This is depicted in Figure 4. Thus, LBs can be used in a more flexible manner than BBBs, but no information about the ingress and egress routers of the admitted flows is a priori available. The presented PCN-based NAC does not constrain the use of the link bandwidth to any flows. Therefore, PCN-based NAC is categorized as LB NAC and the admissible rate $AR(l)$ of a link l corresponds to the link budget $LB(l)$.

3) Efficiency of Non-Resilient NAC Methods: Unlike BBB NAC, LB NAC has no constraints concerning the potentially admitted traffic. Therefore, networks using LB NAC encounter lower flow blocking probabilities than those using BBB NAC. In other words, networks using LB NAC need less bandwidth than those using BBB NAC to achieve the same flow blocking

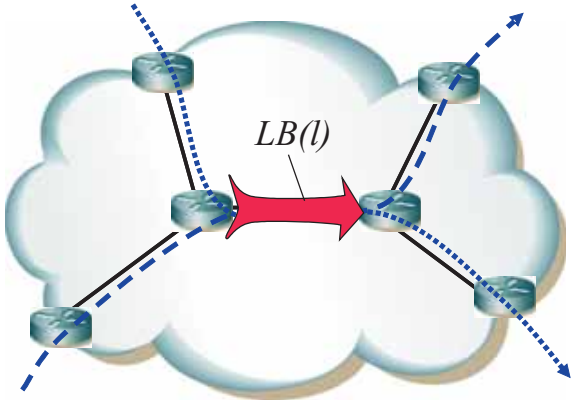

Fig. 4. The capacity of a link budget $LB(l)$ can be used by any flow being transported over link l.

probabilities. This consideration is based on a bandwidth dimensioning approach and we use it in the following to compare the efficiency of NAC methods.

We assume Poisson arrivals for microflows and control the traffic intensity by the average offered load a_{b2b} for each b2b pair. It is given in Erlang which is roughly the average number of simultaneous microflows. To achieve sufficiently low b2b blocking probabilities for microflows, the BBBs must be dimensioned large enough. Further details about flow rates and the applied equations are given in [28]. The capacities of the BBBs entail bandwidth requirements for the capacity tunnels along the respective paths and for the links. Finally, the fraction of the average transported traffic and the overall required capacity yields the average resource utilization which is our measure of efficiency. It is illustrated in Figure 5for the topology and traffic matrix of the Labnet03 network (cf. [28]) for a desired b2b flow blocking probability of 10^{-3}. The fact that the average resource utilization increases with increasing offered load a_{b2b} is called multiplexing gain.

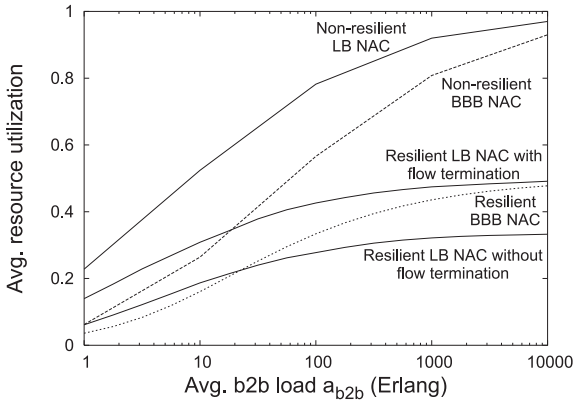

Fig. 5. NAC methods have different efficiency. Their average resource utilization depends on the offered load and resilience requirements.

In case of LB NAC, flows need to request admission to several links and, therefore, the encountered blocking probability at each LB needs to be smaller than the desired b2b blocking probability. Unlike BBB NAC, the capacity of an $LB(l)$ can be used by any flow. Therefore, its capacity is dimensioned based on the overall load expected from all b2b aggregates being carried over the link l. Figure 5 shows that non-resilient LB NAC leads to a higher resource utilization than non-resilient BBB NAC. This is due to a larger multiplexing gain which is caused by the fact that the dimensioning process for LBs is based on larger load quantities than the one for BBBs since BBBs support fewer flows than the LBs on the same paths.

4) Efficiency of Resilient NAC Methods without Flow Termination: A precondition for resilient NAC is that AC states are administered only at the network border such that no states need to be recovered inside the network when the routing changes due to failures if interior nodes fail. BBB NAC administers its budgets, e.g., at the ingress routers and LB NAC may be implemented, e.g., using the PCN framework which also has a stateless core. Furthermore, resilient NAC budgets need to be set small enough that redirected admitted traffic cannot cause overload. For the purpose of performance comparison, the links are dimensioned large enough that backup capacity is available for all admissible b2b traffic patterns in all protected failure scenarios.

In our study, we take all single link failures into account and the results are illustrated in Figure 5. When enough capacity is provided for all protected failure scenarios, resilient BBB NAC leads to larger average resource utilization than resilient LB NAC (without flow termination) for sufficiently large offered load. This is counterintuitive, but can be explained as follows. If a link l fails, the maximum b2b rate for any aggregate is known as it is constrained by the respective BBB. In case of LB NAC, the rates of the same b2b aggregates are limited by the LBs along the respective path which are usually significantly larger than the corresponding BBBs. Hence, the LB NAC can admit more and especially more extreme b2b traffic patterns. Thus, the budgets of the BBB NAC give tighter bounds on the potentially admitted traffic and, as a consequence, less backup capacity is required than for LB NAC. Therefore, resilient BBB NAC is more efficient than resilient LB NAC. More details can be found in [32].

5) Efficiency of Resilient LB NAC with Flow Termination: If flow termination is possible, we suggest to provide backup capacity only for likely traffic patterns and to terminate flows to restore the CL service when an unlikely admitted traffic pattern causes congestion in a failure case. To evaluate this approach, we dimension the backup paths only for the expected traffic given in the traffic matrix instead of providing backup capacity for all traffic patterns that are admissible in the failure-free scenario. Figure 5 shows that the average resource utilization for resilient LB NAC with flow termination is clearly above the one of resilient LB NAC without flow termination because significantly less backup capacity is required. The corresponding curve is also clearly above the one of resilient BBB NAC since resilient LB NAC with flow termination can realize more multiplexing gain especially for

little offered load.

Thus, resilient LB NAC with flow termination is very resource efficient. The probability of flow termination is low and can be neglected since flows are only terminated if a failure occurs and if the rerouting of the admitted traffic causes congestion. This happens only if an unlikely traffic pattern has been admitted whose structure deviates significantly from the one of the traffic matrix.

V. Conclusion

Network admission control (NAC) using pre-congestion notification (PCN) is a promising approach to achieve QoS in large networks. As it is based on PCN marking and measurement of current flow rates, it cannot guarantee QoS for extreme scenarios, e.g. flash crowds. A flow termination mechanism also based on PCN marking allows to restore a controlled load (CL) service at the expense of terminated flows. PCN-based NAC can be operated as resilient NAC without any modifications due to the following reasons. Firstly, interior nodes do not keep per flow reservations, therefore, no admission control related states need to be recovered in the presence of rerouting in case of network failures. Secondly, the admissible link rates $AR(l)$ can be set small enough such that redirected admitted traffic does not cause congestion in protected failure scenarios.

Link budget (LB) based NAC and border-to-border (b2b) budget (BBB) based NAC represent two fundamental NAC categories which are implemented by many AC protocols. Without resilience requirements, LB NAC is more efficient than BBB NAC, but BBB NAC is more efficient than LB NAC when resilience is required. The new PCN-based NAC can be categorized as resilient LB NAC. However, due to its flow termination capabilities, less backup capacity is required than for resilient LB NAC without flow termination. Thus it constitutes a new NAC category with is most efficient with and without resilience requirements.

Acknowledgements

The author would like to thank Stefan Kopf and Matthias Archut for their programming efforts as well as Andreas Binzenhöfer and Matthias Hartmann for their fruitful discussions.

References

[1] K. Fukuda, K. Cho, H. Esaki, and A. Kato, "The Impact of Residential Broadband Traffic on Japanese ISP Backbones," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 1, pp. 15–22, Jan. 2005.

[2] K. Cho, K. Fukuda, H. Esaki, and A. Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic," in *ACM SIGCOMM*, Pisa, Italy, Sept. 2006.

[3] D. M. Johnson, "QoS Control versus Generous Dimensioning," *British Telecom Technology Journal*, vol. 23, no. 2, pp. 81–96, Apr. 2005.

[4] S. Shenker, "Fundamental Design Issues for the Future Internet," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, pp. 1176–1188, Sept. 1995.

[5] S. Iyer, S. Bhattacharyya, N. Taft, and C. Diot, "An Approach to Alleviate Link Overload as Observed on an IP Backbone," in *IEEE Infocom*, San Francisco, CA, April 2003.

[6] M. Menth, R. Martin, and J. Charzinski, "Capacity Overprovisioning for Networks with Resilience Requirements," in *ACM SIGCOMM*, Pisa, Italy, Sept. 2006.

[7] IETF Working Group on Congestion and Pre-Congestion Notification (pcn), "Description of the Working Group," http://www.ietf.org/html.charters/pcn-charter.html, Feb. 2007.

[8] S. Blake, D. L. Black, M. A. Carlson, E. Davies, Z. Wang, and W. Weiss, "RFC2475: An Architecture for Differentiated Services," Dec. 1998.

[9] Y. Bernet, P. Ford, R. Yavatkar, F. Baker, L. Zhang, M. Speer, R. Braden, B. Davie, J. Wroclawski, and E. Felstaine, "RFC2998: A Framework for Integrated Services Operation over Diffserv Networks," Nov. 2000.

[10] B. Braden, D. Clark, and S. Shenker, "RFC1633: Integrated Services in the Internet Architecture: an Overview," June 1994.

[11] B. Briscoe, P. Eardley, D. Songhurst, F. L. Faucheur, A. Charny, J. Babiarz, K. Chan, S. Dudley, G. Karagiannis, A. Bader, and L. Westberg, "An Edge-to-Edge Deployment Model for Pre-Congestion Notification: Admission Control over a DiffServ Region," http://www.ietf.org/internet-drafts/draft-briscoe-tsvwg-cl-architecture-04.txt, Oct. 2006.

[12] J. Wroclawski, "RFC2211: Specification of the Controlled-Load Network Element Service," Sept. 1997.

[13] B. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "RFC2205: Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification," Sept. 1997.

[14] R. Martin, M. Menth, and M. Hemmkeppler, "Accuracy and Dynamics of Hash-Based Load Balancing Algorithms for Multipath Internet Routing," in *IEEE International Conference on Broadband Communication, Networks, and Systems (BROADNETS)*, San Jose, CA, USA, Oct. 2006.

[15] ——, "Accuracy and Dynamics of Multi-Stage Load Balancing for Multipath Internet Routing," in *IEEE International Conference on Communications (ICC)*, Glasgow, Scotland, UK, June 2007.

[16] C. Dovrolis and P. Ramanathan, "Proportional Differentiated Services, Part II: Loss Rate Differentiation and Packet Dropping," in *IEEE International Workshop on Quality of Service (IWQoS)*, Pittsburgh, PA, June 2000.

[17] M. Menth and R. Martin, "Service Differentiation with MEDF Scheduling in TCP/IP Networks," *Computer Communications*, vol. 29, no. 7, pp. 812–819, Apr. 2006.

[18] F. L. Faucheur, "RFC4127: Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering," June 2005.

[19] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.

[20] B. Braden et al., "RFC2309: Recommendations on Queue Management and Congestion Avoidance in the Internet," Apr. 1998.

[21] F. Anjum and L. Tassiulas, "Balanced-RED: An Algorithm to Achieve Fairness in the Internet," in *IEEE Infocom*, Mar. 1999.

[22] U. Bodin, O. Schelén, and S. Pink, "Load-Tolerant Differentiation with Active Queue Management," *ACM SIGCOMM Computer Communications Review*, July 2000.

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

[24] K. Nichols, S. Blake, F. Baker, and D. L. Black, "RFC2474: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," Dec. 1998.

[25] R. J. Gibbens and F. P. Kelly, "Distributed Connection Acceptance Control for a Connectionless Network," in 16^{th}*International Teletraffic Congress (ITC)*, Edinburg, 6 1999, pp. 941 – 952.

[26] M. Karsten and J. Schmitt, "Admission Control based on Packet Marking and Feedback Signalling – Mechanisms, Implementation and Experiments," Darmstadt University of Technology, Technical Report 03/2002, 2002.

[27] ——, "Packet Marking for Integrated Load Control," in *IFIP/IEEE Symposium on Integrated Management (IM)*, 2005.

[28] M. Menth, "Efficient Admission Control and Routing in Resilient Communication Networks," PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland, July 2004.

[29] J. Roberts, U. Mocci, and J. Virtamo, *Broadband Network Teletraffic - Final Report of Action COST 242*. Berlin, Heidelberg: Springer, 1996.

[30] M. Grossglauser and D. N. C. Tse, "A Framework for Robust Measurement-Based Admission Control," *IEEE/ACM Transactions on Networking*, vol. 7, no. 3, pp. 293–309, 1999.

[31] J. Qiu and E. W. Knightly, "Measurement-Based Admission Control with Aggregate Traffic Envelopes," *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, pp. 199–210, 2001.

[32] M. Menth, S. Kopf, and J. Charzinski, "Network Admission Control for Fault-Tolerant QoS Provisioning," in *IEEE High-Speed Networks for Multimedia Communication (HSNMC)*, Toulouse, France, June 2004, pp. 1 – 13.