

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: November 18, 2016

N. Kuhn, Ed.
CNES, Telecom Bretagne
P. Natarajan, Ed.
Cisco Systems
N. Khademi, Ed.
University of Oslo
D. Ros
Simula Research Laboratory AS
May 17, 2016

AQM Characterization Guidelines
draft-ietf-aqm-eval-guidelines-12

Abstract

Unmanaged large buffers in today's networks have given rise to a slew of performance issues. These performance issues can be addressed by some form of Active Queue Management (AQM) mechanism, optionally in combination with a packet scheduling scheme such as fair queuing. This document describes various criteria for performing characterizations of AQM schemes.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 18, 2016.

Copyright Notice

Copyright (c) 2016 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 Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Reducing the latency and maximizing the goodput	5
1.2. Goals of this document	5
1.3. Requirements Language	6
1.4. Glossary	6
2. End-to-end metrics	7
2.1. Flow completion time	7
2.2. Flow start up time	8
2.3. Packet loss	8
2.4. Packet loss synchronization	9
2.5. Goodput	9
2.6. Latency and jitter	10
2.7. Discussion on the trade-off between latency and goodput	10
3. Generic setup for evaluations	11
3.1. Topology and notations	11
3.2. Buffer size	13
3.3. Congestion controls	13
4. Methodology, Metrics, AQM Comparisons, Packet Sizes, Scheduling and ECN	14
4.1. Methodology	14
4.2. Comments on metrics measurement	14
4.3. Comparing AQM schemes	15
4.3.1. Performance comparison	15
4.3.2. Deployment comparison	16
4.4. Packet sizes and congestion notification	16
4.5. Interaction with ECN	16
4.6. Interaction with Scheduling	17
5. Transport Protocols	17
5.1. TCP-friendly sender	18
5.1.1. TCP-friendly sender with the same initial congestion window	18
5.1.2. TCP-friendly sender with different initial congestion windows	18
5.2. Aggressive transport sender	19
5.3. Unresponsive transport sender	19
5.4. Less-than Best Effort transport sender	20
6. Round Trip Time Fairness	20
6.1. Motivation	20
6.2. Recommended tests	21

6.3. Metrics to evaluate the RTT fairness	21
7. Burst Absorption	21
7.1. Motivation	21
7.2. Recommended tests	22
8. Stability	23
8.1. Motivation	23
8.2. Recommended tests	23
8.2.1. Definition of the congestion Level	24
8.2.2. Mild congestion	24
8.2.3. Medium congestion	24
8.2.4. Heavy congestion	25
8.2.5. Varying the congestion level	25
8.2.6. Varying available capacity	25
8.3. Parameter sensitivity and stability analysis	26
9. Various Traffic Profiles	27
9.1. Traffic mix	27
9.2. Bi-directional traffic	27
10. Multi-AQM Scenario	28
10.1. Motivation	28
10.2. Details on the evaluation scenario	28
11. Implementation cost	28
11.1. Motivation	28
11.2. Recommended discussion	29
12. Operator Control and Auto-tuning	29
12.1. Motivation	29
12.2. Recommended discussion	30
13. Conclusion	30
14. Acknowledgements	31
15. Contributors	31
16. IANA Considerations	32
17. Security Considerations	32
18. References	32
18.1. Normative References	32
18.2. Informative References	34
Authors' Addresses	35

1. Introduction

Active Queue Management (AQM) [RFC7567] addresses the concerns arising from using unnecessarily large and unmanaged buffers to improve network and application performance. Several AQM algorithms have been proposed in the past years, most notably Random Early Detection (RED), BLUE, and Proportional Integral controller (PI), and more recently CoDel [NICH2012] and PIE [PAN2013]. In general, these algorithms actively interact with the Transmission Control Protocol (TCP) and any other transport protocol that deploys a congestion control scheme to manage the amount of data they keep in the network. The available buffer space in the routers and switches should be

large enough to accommodate the short-term buffering requirements. AQM schemes aim at reducing buffer occupancy, and therefore the end-to-end delay. Some of these algorithms, notably RED, have also been widely implemented in some network devices. However, the potential benefits of the RED scheme have not been realized since RED is reported to be usually turned off.

A buffer is a physical volume of memory in which a queue or set of queues are stored. When speaking of a specific queue in this document, "buffer occupancy" refers to the amount of data (measured in bytes or packets) that are in the queue, and the "maximum buffer size" refers to the maximum buffer occupancy. In switches and routers, a global memory space is often shared between the available interfaces, and thus, the maximum buffer size for any given interface may vary over the time.

Bufferbloat [BB2011] is the consequence of deploying large unmanaged buffers on the Internet -- the buffering has often been measured to be ten times or hundred times larger than needed. Large buffer sizes in combination with TCP and/or unresponsive flows increases end-to-end delay. This results in poor performance for latency-sensitive applications such as real-time multimedia (e.g., voice, video, gaming, etc). The degree to which this affects modern networking equipment, especially consumer-grade equipment's, produces problems even with commonly used web services. Active queue management is thus essential to control queuing delay and decrease network latency.

The Active Queue Management and Packet Scheduling Working Group (AQM WG) was chartered to address the problems with large unmanaged buffers in the Internet. Specifically, the AQM WG is tasked with standardizing AQM schemes that not only address concerns with such buffers, but also are robust under a wide variety of operating conditions. This document provides characterization guidelines that can be used to assess the applicability, performance and deployability of an AQM, whether it is candidate for standardization at IETF or not.

AQM algorithm implemented in a router can be separated from the scheduling of packets sent out by the router [RFC7567]. The rest of this memo refers to the AQM as a dropping/marketing policy as a separate feature to any interface scheduling scheme. This document may be complemented with another one on guidelines for assessing combination of packet scheduling and AQM. We note that such a document will inherit all the guidelines from this document plus any additional scenarios relevant for packet scheduling such as flow starvation evaluation or impact of the number of hash buckets.

1.1. Reducing the latency and maximizing the goodput

The trade-off between reducing the latency and maximizing the goodput is intrinsically linked to each AQM scheme and is key to evaluating its performance. To ensure the safety deployment of an AQM, its behaviour should be assessed in a variety of scenarios. Whenever possible, solutions ought to aim at both maximizing goodput and minimizing latency.

1.2. Goals of this document

This document recommends a generic list of scenarios against which an AQM proposal should be evaluated, considering both potential performance gain and safety of deployment. The guidelines help to quantify performance of AQM schemes in terms of latency reduction, goodput maximization and the trade-off between these two. The document presents central aspects of an AQM algorithm that should be considered whatever the context, such as burst absorption capacity, RTT fairness or resilience to fluctuating network conditions. The guidelines also discuss methods to understand the various aspects associated with safely deploying and operating the AQM scheme. Thus, one of the key objectives behind formulating the guidelines is to help ascertain whether a specific AQM is not only better than drop-tail (i.e. without AQM and with a BDP-sized buffer) but also safe to deploy: the guidelines can be used to compare several AQM proposals with each other, but should be used to compare a proposal with drop-tail.

This memo details generic characterization scenarios against which any AQM proposal should be evaluated, irrespective of whether or not an AQM is standardized by the IETF. This documents recommends the relevant scenarios and metrics to be considered. The document presents central aspects of an AQM algorithm that should be considered whatever the context, such as burst absorption capacity, RTT fairness or resilience to fluctuating network conditions.

These guidelines do not define and are not bound to a particular deployment scenario or evaluation toolset. Instead the guidelines can be used to assert the potential gain of introducing an AQM for the particular environment, which is of interest to the testers. These guidelines do not cover every possible aspect of a particular algorithm. These guidelines do not present context-dependent scenarios (such as 802.11 WLANs, data-centers or rural broadband networks). To keep the guidelines generic, a number of potential router components and algorithms (such as DiffServ) are omitted.

The goals of this document can thus be summarized as follows:

- o The present characterization guidelines provide a non-exhaustive list of scenarios to help ascertain whether an AQM is not only better than drop-tail (with a BDP-sized buffer), but also safe to deploy; the guidelines can also be used to compare several AQM proposals with each other.
- o The present characterization guidelines (1) are not bound to a particular evaluation toolset and (2) can be used for various deployment contexts; testers are free to select a toolset that is best suited for the environment in which their proposal will be deployed.
- o The present characterization guidelines are intended to provide guidance for better selecting an AQM for a specific environment; it is not required that an AQM proposal is evaluated following these guidelines for its standardization.

1.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.4. Glossary

- o application-limited traffic: a type of traffic that does not have an unlimited amount of data to transmit.
- o AQM: the Active Queue Managment (AQM) algorithm implemented in a router can be separated from the scheduling of packets sent by the router. The rest of this memo refers to the AQM as a dropping/marking policy as a separate feature to any interface scheduling scheme [RFC7567].
- o BDP: Bandwidth Delay Product.
- o buffer: a physical volume of memory in which a queue or set of queues are stored.
- o buffer occupancy: amount of data that are stored in a buffer, measured in bytes or packets.
- o buffer size: maximum buffer occupancy, that is the maximum amount of data that may be stored in a buffer, measured in bytes or packets.
- o IW10: TCP initial congestion window set to 10 packets.

- o latency: one-way delay of packets across Internet paths. This definition suits transport layer definition of the latency, that shall not be confused with an application layer view of the latency.
- o goodput: goodput is defined as the number of bits per unit of time forwarded to the correct destination minus any bits lost or retransmitted [RFC2647].
- o SQRT: the square root function.
- o ROUND: the round function.

2. End-to-end metrics

End-to-end delay is the result of propagation delay, serialization delay, service delay in a switch, medium-access delay and queuing delay, summed over the network elements along the path. AQM schemes may reduce the queuing delay by providing signals to the sender on the emergence of congestion, but any impact on the goodput must be carefully considered. This section presents the metrics that could be used to better quantify (1) the reduction of latency, (2) maximization of goodput and (3) the trade-off between these two. This section provides normative requirements for metrics that can be used to assess the performance of an AQM scheme.

Some metrics listed in this section are not suited to every type of traffic detailed in the rest of this document. It is therefore not necessary to measure all of the following metrics: the chosen metric may not be relevant to the context of the evaluation scenario (e.g., latency vs. goodput trade-off in application-limited traffic scenarios). Guidance is provided for each metric.

2.1. Flow completion time

The flow completion time is an important performance metric for the end-user when the flow size is finite. The definition of the flow size may be source of contradictions, thus, this metric can consider a flow as a single file. Considering the fact that an AQM scheme may drop/mark packets, the flow completion time is directly linked to the dropping/mark policy of the AQM scheme. This metric helps to better assess the performance of an AQM depending on the flow size. The Flow Completion Time (FCT) is related to the flow size (F_s) and the goodput for the flow (G) as follows:

$$FCT [s] = F_s [Byte] / (G [Bit/s] / 8 [Bit/Byte])$$

If this metric is used to evaluate the performance of web transfers, it is suggested to rather consider the time needed to download all the objects that compose the web page, as this makes more sense in terms of user experience than assessing the time needed to download each object.

2.2. Flow start up time

The flow start up time is the time between the request has been sent from the client and the server starts to transmit data. The amount of packets dropped by an AQM may seriously affect the waiting period during which the data transfer has not started. This metric would specifically focus on the operations such as DNS lookups, TCP opens and SSL handshakes.

2.3. Packet loss

Packet loss can occur en-route, this can impact the end-to-end performance measured at receiver.

The tester SHOULD evaluate loss experienced at the receiver using one of the two metrics:

- o the packet loss ratio: this metric is to be frequently measured during the experiment. The long-term loss ratio is of interest for steady-state scenarios only;
- o the interval between consecutive losses: the time between two losses is to be measured.

The packet loss ratio can be assessed by simply evaluating the loss ratio as a function of the number of lost packets and the total number of packets sent. This might not be easily done in laboratory testing, for which these guidelines advice the tester:

- o to check that for every packet, a corresponding packet was received within a reasonable time [RFC2680].
- o to keep a count of all packets sent, and a count of the non-duplicate packets received [RFC2544].

The interval between consecutive losses, which is also called a gap, is a metric of interest for VoIP traffic [RFC3611].

2.4. Packet loss synchronization

One goal of an AQM algorithm is to help to avoid global synchronization of flows sharing a bottleneck buffer on which the AQM operates ([RFC2309],[RFC7567]). The "degree" of packet-loss synchronization between flows SHOULD be assessed, with and without the AQM under consideration.

Loss synchronization among flows may be quantified by several slightly different metrics that capture different aspects of the same issue [HASS2008]. However, in real-world measurements the choice of metric could be imposed by practical considerations -- e.g., whether fine-grained information on packet losses in the bottleneck available or not. For the purpose of AQM characterization, a good candidate metric is the global synchronization ratio, measuring the proportion of flows losing packets during a loss event. This metric can be used in real-world experiments to characterize synchronization along arbitrary Internet paths [JAY2006].

If an AQM scheme is evaluated using real-life network environments, it is worth pointing out that some network events, such as failed link restoration may cause synchronized losses between active flows and thus confuse the meaning of this metric.

2.5. Goodput

The goodput has been defined as the number of bits per unit of time forwarded to the correct destination interface, minus any bits lost or retransmitted [RFC2647]. This definition requires that the test setup needs to be qualified to assure that it is not generating losses on its own.

Measuring the end-to-end goodput provides an appreciation of how well an AQM scheme improves transport and application performance. The measured end-to-end goodput is linked to the dropping/marketing policy of the AQM scheme -- e.g., the fewer the number of packet drops, the fewer packets need retransmission, minimizing the impact of AQM on transport and application performance. Additionally, an AQM scheme may resort to Explicit Congestion Notification (ECN) marking as an initial means to control delay. Again, marking packets instead of dropping them reduces the number of packet retransmissions and increases goodput. End-to-end goodput values help to evaluate the AQM scheme's effectiveness of an AQM scheme in minimizing packet drops that impact application performance and to estimate how well the AQM scheme works with ECN.

The measurement of the goodput allows the tester evaluate to which extent an AQM is able to maintain a high bottleneck utilization.

This metric should be also obtained frequently during an experiment as the long-term goodput is relevant for steady-state scenarios only and may not necessarily reflect how the introduction of an AQM actually impacts the link utilization during at a certain period of time. Fluctuations in the values obtained from these measurements may depend on other factors than the introduction of an AQM, such as link layer losses due to external noise or corruption, fluctuating bandwidths (802.11 WLANs), heavy congestion levels or transport layer's rate reduction by congestion control mechanism.

2.6. Latency and jitter

The latency, or the one-way delay metric, is discussed in [RFC2679]. There is a consensus on an adequate metric for the jitter, that represents the one-way delay variations for packets from the same flow: the Packet Delay Variation (PDV) serves well all use cases [RFC5481].

The end-to-end latency includes components other than just the queuing delay, such as the signal processing delay, transmission delay and the processing delay. Moreover, the jitter is caused by variations in queuing and processing delay (e.g., scheduling effects). The introduction of an AQM scheme would impact end-to-end latency and jitter, and therefore these metrics should be considered in the end-to-end evaluation of performance.

2.7. Discussion on the trade-off between latency and goodput

The metrics presented in this section may be considered as explained in the rest of this document, in order to discuss and quantify the trade-off between latency and goodput.

With regards to the goodput, and in addition to the long-term stationary goodput value, it is RECOMMENDED to take measurements every multiple of the minimum RTT (minRTT) between A and B. It is suggested to take measurements at least every $K \times \text{minRTT}$ (to smooth out the fluctuations), with $K=10$. Higher values for K are encouraged whenever it is more appropriate for the presentation of the results. The value for K may depend on the network's path characteristics. The measurement period MUST be disclosed for each experiment and when results/values are compared across different AQM schemes, the comparisons SHOULD use exactly the same measurement periods. With regards to latency, it is RECOMMENDED to take the samples on per-packet basis whenever possible depending on the features provided by hardware/software and the impact of sampling itself on the hardware performance. It is generally RECOMMENDED to provide at least 10 samples per RTT.

From each of these sets of measurements, the cumulative density function (CDF) of the considered metrics SHOULD be computed. If the considered scenario introduces dynamically varying parameters, temporal evolution of the metrics could also be generated. For each scenario, the following graph may be generated: the x-axis shows queuing delay (that is the average per-packet delay in excess of minimum RTT), the y-axis the goodput. Ellipses are computed such as detailed in [WINS2014]: "We take each individual [...] run [...] as one point, and then compute the 1-epsilon elliptic contour of the maximum-likelihood 2D Gaussian distribution that explains the points. [...] we plot the median per-sender throughput and queueing delay as a circle. [...] The orientation of an ellipse represents the covariance between the throughput and delay measured for the protocol." This graph provides part of a better understanding of (1) the delay/goodput trade-off for a given congestion control mechanism Section 5, and (2) how the goodput and average queue delay vary as a function of the traffic load Section 8.2.

3. Generic setup for evaluations

This section presents the topology that can be used for each of the following scenarios, the corresponding notations and discusses various assumptions that have been made in the document.

3.1. Topology and notations

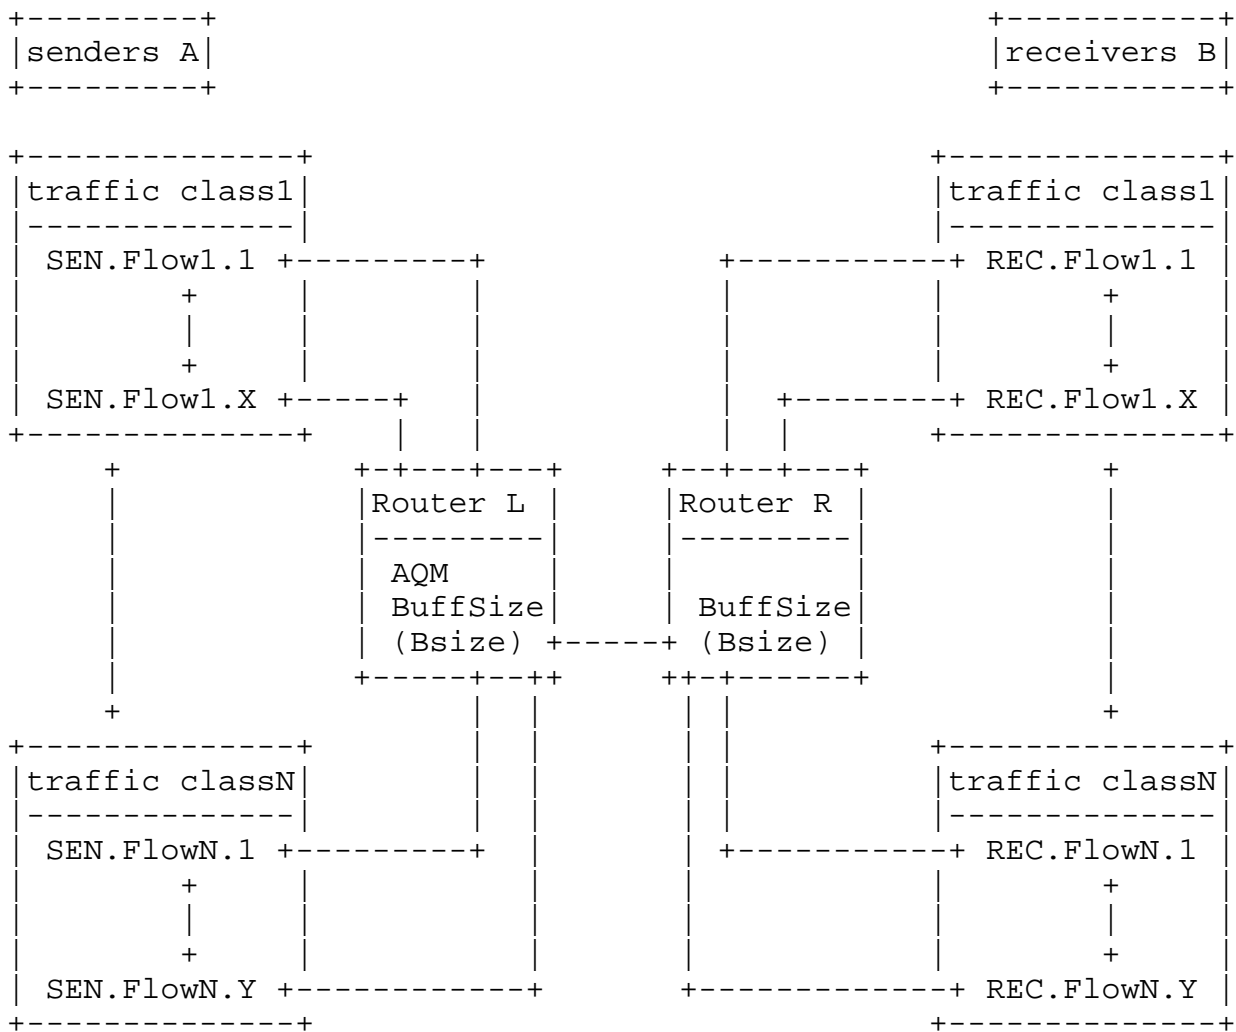


Figure 1: Topology and notations

Figure 1 is a generic topology where:

- o senders with different traffic characteristics (i.e., traffic profiles) can be introduced;
- o the timing of each flow could be different (i.e., when does each flow start and stop);
- o each traffic profile can comprise various number of flows;
- o each link is characterized by a couple (one-way delay, capacity);
- o flows are generated at A and sent to B, sharing a bottleneck (the link between routers L and R);

- o the tester SHOULD consider both scenarios of asymmetric and symmetric bottleneck links in terms of bandwidth. In case of asymmetric link, the capacity from senders to receivers is higher than the one from receivers to senders; the symmetric link scenario provides a basic understanding of the operation of the AQM mechanism whereas the asymmetric link scenario evaluates an AQM mechanism in a more realistic setup;
- o in asymmetric link scenarios, the tester SHOULD study the bi-directional traffic between A and B (downlink and uplink) with the AQM mechanism deployed on one direction only. The tester MAY additionally consider a scenario with AQM mechanism being deployed on both directions. In each scenario, the tester SHOULD investigate the impact of drop policy of the AQM on TCP ACK packets and its impact on the performance.

Although this topology may not perfectly reflect actual topologies, the simple topology is commonly used in the world of simulations and small testbeds. It can be considered as adequate to evaluate AQM proposals [I-D.irtf-iccrq-tcpeval]. Testers ought to pay attention to the topology that has been used to evaluate an AQM scheme when comparing this scheme with a newly proposed AQM scheme.

3.2. Buffer size

The size of the buffers should be carefully chosen, and MAY be set to the bandwidth-delay product; the bandwidth being the bottleneck capacity and the delay the largest RTT in the considered network. The size of the buffer can impact the AQM performance and is a dimensioning parameter that will be considered when comparing AQM proposals.

If a specific buffer size is required, the tester MUST justify and detail the way the maximum queue size is set. Indeed, the maximum size of the buffer may affect the AQM's performance and its choice SHOULD be elaborated for a fair comparison between AQM proposals. While comparing AQM schemes the buffer size SHOULD remain the same across the tests.

3.3. Congestion controls

This document considers running three different congestion control algorithms between A and B

- o Standard TCP congestion control: the base-line congestion control is TCP NewReno with SACK [RFC5681].

- o Aggressive congestion controls: a base-line congestion control for this category is TCP Cubic [I-D.ietf-tcpm-cubic].
- o Less-than Best Effort (LBE) congestion controls: an LBE congestion control 'results in smaller bandwidth and/or delay impact on standard TCP than standard TCP itself, when sharing a bottleneck with it.' [RFC6297]

Other transport congestion controls can OPTIONALLY be evaluated in addition. Recent transport layer protocols are not mentioned in the following sections, for the sake of simplicity.

4. Methodology, Metrics, AQM Comparisons, Packet Sizes, Scheduling and ECN

4.1. Methodology

A description of each test setup SHOULD be detailed to allow this test to be compared with other tests. This also allows others to replicate the tests if needed. This test setup SHOULD detail software and hardware versions. The tester could make its data available.

The proposals SHOULD be evaluated on real-life systems, or they MAY be evaluated with event-driven simulations (such as ns-2, ns-3, OMNET, etc). The proposed scenarios are not bound to a particular evaluation toolset.

The tester is encouraged to make the detailed test setup and the results publicly available.

4.2. Comments on metrics measurement

The document presents the end-to-end metrics that ought to be used to evaluate the trade-off between latency and goodput in Section 2. In addition to the end-to-end metrics, the queue-level metrics (normally collected at the device operating the AQM) provide a better understanding of the AQM behavior under study and the impact of its internal parameters. Whenever it is possible (e.g., depending on the features provided by the hardware/software), these guidelines advise to consider queue-level metrics, such as link utilization, queuing delay, queue size or packet drop/mark statistics in addition to the AQM-specific parameters. However, the evaluation MUST be primarily based on externally observed end-to-end metrics.

These guidelines do not aim to detail on the way these metrics can be measured, since the way these metrics are measured is expected to depend on the evaluation toolset.

4.3. Comparing AQM schemes

This document recognizes that these guidelines may be used for comparing AQM schemes.

AQM schemes need to be compared against both performance and deployment categories. In addition, this section details how best to achieve a fair comparison of AQM schemes by avoiding certain pitfalls.

4.3.1. Performance comparison

AQM schemes should be compared against the generic scenarios that are summarized in Section 13. AQM schemes MAY be compared for specific network environments such as data centers, home networks, etc. If an AQM scheme has parameter(s) that were externally tuned for optimization or other purposes, these values MUST be disclosed.

AQM schemes belong to different varieties such as queue-length based schemes (ex. RED) or queueing-delay based scheme (ex. CoDel, PIE). AQM schemes expose different control knobs associated with different semantics. For example, while both PIE and CoDel are queueing-delay based schemes and each expose a knob to control the queueing delay -- PIE's "queueing delay reference" vs. CoDel's "queueing delay target", the two tuning parameters of the two schemes have different semantics, resulting in different control points. Such differences in AQM schemes can be easily overlooked while making comparisons.

This document RECOMMENDS the following procedures for a fair performance comparison between the AQM schemes:

1. Similar control parameters and implications: Testers should be aware of the control parameters of the different schemes that control similar behavior. Testers should also be aware of the input value ranges and corresponding implications. For example, consider two different schemes - (A) queue-length based AQM scheme, and (B) queueing-delay based scheme. A and B are likely to have different kinds of control inputs to control the target delay - target queue length in A vs. target queuing delay in B, for example. Setting parameter values such as 100MB for A vs. 10ms for B will have different implications depending on evaluation context. Such context-dependent implications must be considered before drawing conclusions on performance comparisons. Also, it would be preferable if an AQM proposal listed such parameters and discussed how each relates to network characteristics such as capacity, average RTT etc.

2. compare over a range of input configurations: there could be situations when the set of control parameters that affect a specific behavior have different semantics between the two AQM schemes. As mentioned above, PIE has tuning parameters to control queue delay that has a different semantics from those used in CoDel. In such situations, these schemes need to be compared over a range of input configurations. For example, compare PIE vs. CoDel over the range of target delay input configurations.

4.3.2. Deployment comparison

AQM schemes MUST be compared against deployment criteria such as the parameter sensitivity (Section 8.3), auto-tuning (Section 12) or implementation cost (Section 11).

4.4. Packet sizes and congestion notification

An AQM scheme may be considering packet sizes while generating congestion signals [RFC7141]. For example, control packets such as DNS requests/responses, TCP SYNs/ACKs are small, but their loss can severely impact the application performance. An AQM scheme may therefore be biased towards small packets by dropping them with smaller probability compared to larger packets. However, such an AQM scheme is unfair to data senders generating larger packets. Data senders, malicious or otherwise, are motivated to take advantage of such AQM scheme by transmitting smaller packets, and could result in unsafe deployments and unhealthy transport and/or application designs.

An AQM scheme SHOULD adhere to the recommendations outlined in [RFC7141], and SHOULD NOT provide undue advantage to flows with smaller packets [RFC7567]. In order to evaluate if an AQM scheme is biased towards flows with smaller size packets, sender A in Figure 1 can be instantiated with two long standing TCP flows with different packet sizes; such as 500 bytes vs. 1500 bytes respectively, and metrics can be compared for these two flows.

4.5. Interaction with ECN

Deployed AQM algorithms SHOULD implement Explicit Congestion Notification (ECN) as well as loss to signal congestion to endpoints [RFC7567]. ECN [RFC3168] is an alternative that allows AQM schemes to signal receivers about network congestion that does not use packet drop. There are benefits of providing ECN support for an AQM scheme [WELZ2015]. AQM schemes SHOULD NOT drop or remark packets solely because the ECT(0) or ECT(1) codepoints are used, and when ECN-

capable SHOULD set a CE-mark on ECN-capable packets in the presence of incipient congestion.

If the tested AQM scheme can support ECN [RFC7567], the testers MUST discuss and describe the support of ECN. Since these guidelines can be used to evaluate the performance of the tested AQM with and without ECN markings, they could also be used to quantify the interest of enabling ECN.

4.6. Interaction with Scheduling

A network device may use per-flow or per-class queuing with a scheduling algorithm to either prioritize certain applications or classes of traffic, limit the rate of transmission, or to provide isolation between different traffic flows within a common class [RFC7567].

The scheduling and the AQM conjointly impact on the end-to-end performance. Therefore, the AQM proposal MUST discuss the feasibility to add scheduling combined with the AQM algorithm. This discussion as an instance, MAY explain whether the dropping policy is applied when packets are being enqueued or dequeued.

These guidelines do not propose guidelines to assess the performance of scheduling algorithms. Indeed, as opposed to characterizing AQM schemes that is related to their capacity to control the queuing delay in a queue, characterizing scheduling schemes is related to the scheduling itself and its interaction with the AQM scheme. As one example, the scheduler may create sub-queues and the AQM scheme may be applied on each of the sub-queues, and/or the AQM could be applied on the whole queue. Also, schedulers might, such as FQ-CoDel [HOEI2015] or FavorQueue [ANEL2014], introduce flow prioritization. In these cases, specific scenarios should be proposed to ascertain that these scheduler schemes not only helps in tackling the bufferbloat, but also are robust under a wide variety of operating conditions. This is out of the scope of this document that focus on dropping and/or marking AQM schemes.

5. Transport Protocols

Network and end-devices need to be configured with a reasonable amount of buffer space to absorb transient bursts. In some situations, network providers tend to configure devices with large buffers to avoid packet drops triggered by a full buffer and to maximize the link utilization for standard loss-based TCP traffic.

AQM algorithms are often evaluated by considering Transmission Control Protocol (TCP) [RFC0793] with a limited number of

applications. TCP is a widely deployed transport. It fills up available buffers until a sender transferring a bulk flow with TCP receives a signal (packet drop) that reduces the sending rate. The larger the buffer, the higher the buffer occupancy, and therefore the queuing delay. An efficient AQM scheme sends out early congestion signals to TCP to bring the queuing delay under control.

Not all endpoints (or applications) using TCP use the same flavor of TCP. Variety of senders generate different classes of traffic which may not react to congestion signals (aka non-responsive flows [RFC7567]) or may not reduce their sending rate as expected (aka Transport Flows that are less responsive than TCP[RFC7567], also called "aggressive flows"). In these cases, AQM schemes seek to control the queuing delay.

This section provides guidelines to assess the performance of an AQM proposal for various traffic profiles -- different types of senders (with different TCP congestion control variants, unresponsive, aggressive).

5.1. TCP-friendly sender

5.1.1. TCP-friendly sender with the same initial congestion window

This scenario helps to evaluate how an AQM scheme reacts to a TCP-friendly transport sender. A single long-lived, non application-limited, TCP NewReno flow, with an Initial congestion Window (IW) set to 3 packets, transfers data between sender A and receiver B. Other TCP friendly congestion control schemes such as TCP-friendly rate control [RFC5348] etc MAY also be considered.

For each TCP-friendly transport considered, the graph described in Section 2.7 could be generated.

5.1.2. TCP-friendly sender with different initial congestion windows

This scenario can be used to evaluate how an AQM scheme adapts to a traffic mix consisting of TCP flows with different values of the IW.

For this scenario, two types of flows MUST be generated between sender A and receiver B:

- o A single long-lived non application-limited TCP NewReno flow;
- o A single application-limited TCP NewReno flow, with an IW set to 3 or 10 packets. The size of the data transferred must be strictly higher than 10 packets and should be lower than 100 packets.

The transmission of the non application-limited flow must start first and the transmission of the application-limited flow starts after the non application-limited flow has reached steady state.

For each of these scenarios, the graph described in Section 2.7 could be generated for each class of traffic (application-limited and non application-limited). The completion time of the application-limited TCP flow could be measured.

5.2. Aggressive transport sender

This scenario helps testers to evaluate how an AQM scheme reacts to a transport sender that is more aggressive than a single TCP-friendly sender. We define 'aggressiveness' as a higher increase factor than standard upon a successful transmission and/or a lower than standard decrease factor upon a unsuccessful transmission (e.g., in case of congestion controls with Additive-Increase Multiplicative-Decrease (AIMD) principle, a larger AI and/or MD factors). A single long-lived, non application-limited, TCP Cubic flow transfers data between sender A and receiver B. Other aggressive congestion control schemes MAY also be considered.

For each flavor of aggressive transports, the graph described in Section 2.7 could be generated.

5.3. Unresponsive transport sender

This scenario helps testers to evaluate how an AQM scheme reacts to a transport sender that is less responsive than TCP. Note that faulty transport implementations on an end host and/or faulty network elements en-route that "hide" congestion signals in packet headers [RFC7567] may also lead to a similar situation, such that the AQM scheme needs to adapt to unresponsive traffic. To this end, these guidelines propose the two following scenarios.

The first scenario can be used to evaluate queue build up. It considers unresponsive flow(s) whose sending rate is greater than the bottleneck link capacity between routers L and R. This scenario consists of a long-lived non application limited UDP flow transmits data between sender A and receiver B. Graphs described in Section 2.7 could be generated.

The second scenario can be used to evaluate if the AQM scheme is able to keep the responsive fraction under control. This scenario considers a mixture of TCP-friendly and unresponsive traffics. It consists of a long-lived UDP flow from unresponsive application and a single long-lived, non application-limited (unlimited data available to the transport sender from application layer), TCP New Reno flow

that transmit data between sender A and receiver B. As opposed to the first scenario, the rate of the UDP traffic should not be greater than the bottleneck capacity, and should be higher than half of the bottleneck capacity. For each type of traffic, the graph described in Section 2.7 could be generated.

5.4. Less-than Best Effort transport sender

This scenario helps to evaluate how an AQM scheme reacts to LBE congestion controls that 'results in smaller bandwidth and/or delay impact on standard TCP than standard TCP itself, when sharing a bottleneck with it.' [RFC6297]. There are potential fateful interactions when AQM and LBE techniques are combined [GONG2014]; this scenario helps to evaluate whether the coexistence of the proposed AQM and LBE techniques may be possible.

A single long-lived non application-limited TCP NewReno flow transfers data between sender A and receiver B. Other TCP-friendly congestion control schemes MAY also be considered. Single long-lived non application-limited LEDBAT [RFC6817] flows transfer data between sender A and receiver B. We recommend to set the target delay and gain values of LEDBAT respectively to 5 ms and 10 [TRAN2014]. Other LBE congestion control schemes [RFC6297], MAY also be considered.

For each of the TCP-friendly and LBE transports, the graph described in Section 2.7 could be generated.

6. Round Trip Time Fairness

6.1. Motivation

An AQM scheme's congestion signals (via drops or ECN marks) must reach the transport sender so that a responsive sender can initiate its congestion control mechanism and adjust the sending rate. This procedure is thus dependent on the end-to-end path RTT. When the RTT varies, the onset of congestion control is impacted, and in turn impacts the ability of an AQM scheme to control the queue. It is therefore important to assess the AQM schemes for a set of RTTs between A and B (e.g., from 5 ms to 200 ms).

The asymmetry in terms of difference in intrinsic RTT between various paths sharing the same bottleneck SHOULD be considered so that the fairness between the flows can be discussed since in this scenario, a flow traversing on shorter RTT path may react faster to congestion and recover faster from it compared to another flow on a longer RTT path. The introduction of AQM schemes may potentially improve the RTT fairness.

Introducing an AQM scheme may cause the unfairness between the flows, even if the RTTs are identical. This potential unfairness SHOULD be investigated as well.

6.2. Recommended tests

The RECOMMENDED topology is detailed in Figure 1.

To evaluate the RTT fairness, for each run, two flows are divided into two categories. Category I whose RTT between sender A and receiver B SHOULD be 100ms. Category II which RTT between sender A and receiver B should be in the range [5ms;560ms] inclusive. The maximum value for the RTT represents the RTT of a satellite link [RFC2488].

A set of evaluated flows MUST use the same congestion control algorithm: all the generated flows could be single long-lived non application-limited TCP NewReno flows.

6.3. Metrics to evaluate the RTT fairness

The outputs that MUST be measured are: (1) the cumulative average goodput of the flow from Category I, `goodput_Cat_I` (Section 2.5); (2) the cumulative average goodput of the flow from Category II, `goodput_Cat_II` (Section 2.5); (3) the ratio `goodput_Cat_II/goodput_Cat_I`; (4) the average packet drop rate for each category (Section 2.3).

7. Burst Absorption

"AQM mechanisms need to control the overall queue sizes, to ensure that arriving bursts can be accommodated without dropping packets" [RFC7567].

7.1. Motivation

An AQM scheme can face bursts of packet arrivals due to various reasons. Dropping one or more packets from a burst can result in performance penalties for the corresponding flows, since dropped packets have to be retransmitted. Performance penalties can result in failing to meet SLAs and be a disincentive to AQM adoption.

The ability to accommodate bursts translates to larger queue length and hence more queuing delay. On the one hand, it is important that an AQM scheme quickly brings bursty traffic under control. On the other hand, a peak in the packet drop rates to bring a packet burst quickly under control could result in multiple drops per flow and severely impact transport and application performance. Therefore, an

AQM scheme ought to bring bursts under control by balancing both aspects -- (1) queuing delay spikes are minimized and (2) performance penalties for ongoing flows in terms of packet drops are minimized.

An AQM scheme that maintains short queues allows some remaining space in the buffer for bursts of arriving packets. The tolerance to bursts of packets depends upon the number of packets in the queue, which is directly linked to the AQM algorithm. Moreover, an AQM scheme may implement a feature controlling the maximum size of accepted bursts, that can depend on the buffer occupancy or the currently estimated queuing delay. The impact of the buffer size on the burst allowance may be evaluated.

7.2. Recommended tests

For this scenario, tester MUST evaluate how the AQM performs with the following traffic generated from sender A to receiver B:

- o Web traffic with IW10;
- o Bursty video frames;
- o Constant Bit Rate (CBR) UDP traffic.
- o A single non application-limited bulk TCP flow as background traffic.

Figure 2 presents the various cases for the traffic that MUST be generated between sender A and receiver B.

+-----+-----+-----+-----+-----+				
Case	Traffic Type			
	Video	Web (IW 10)	CBR	Bulk TCP Traffic
I	0	1	1	0
II	0	1	1	1
III	1	1	1	0
IV	1	1	1	1
+-----+-----+-----+-----+-----+				

Figure 2: Bursty traffic scenarios

A new web page download could start after the previous web page download is finished. Each web page could be composed by at least 50

objects and the size of each object should be at least 1kB. 6 TCP parallel connections SHOULD be generated to download the objects, each parallel connections having an initial congestion window set to 10 packets.

For each of these scenarios, the graph described in Section 2.7 could be generated for each application. Metrics such as end-to-end latency, jitter, flow completion time MAY be generated. For the cases of frame generation of bursty video traffic as well as the choice of web traffic pattern, these details and their presentation are left to the testers.

8. Stability

8.1. Motivation

The safety of an AQM scheme is directly related to its stability under varying operating conditions such as varying traffic profiles and fluctuating network conditions. Since operating conditions can vary often the AQM needs to remain stable under these conditions without the need for additional external tuning.

Network devices can experience varying operating conditions depending on factors such as time of the day, deployment scenario, etc. For example:

- o Traffic and congestion levels are higher during peak hours than off-peak hours.
- o In the presence of a scheduler, the draining rate of a queue can vary depending on the occupancy of other queues: a low load on a high priority queue implies a higher draining rate for the lower priority queues.
- o The capacity available can vary over time (e.g., a lossy channel, a link supporting traffic in a higher diffserv class).

Whether the target context is a not stable environment, the ability of an AQM scheme to maintain its control over the queuing delay and buffer occupancy can be challenged. This document proposes guidelines to assess the behavior of AQM schemes under varying congestion levels and varying draining rates.

8.2. Recommended tests

Note that the traffic profiles explained below comprises non application-limited TCP flows. For each of the below scenarios, the graphs described in Section 2.7 SHOULD be generated, and the goodput

of the various flows should be cumulated. For Section 8.2.5 and Section 8.2.6 they SHOULD incorporate the results in per-phase basis as well.

Wherever the notion of time has explicitly mentioned in this subsection, time 0 starts from the moment all TCP flows have already reached their congestion avoidance phase.

8.2.1. Definition of the congestion Level

In these guidelines, the congestion levels are represented by the projected packet drop rate, had a drop-tail queue was chosen instead of an AQM scheme. When the bottleneck is shared among non application-limited TCP flows. l_r , the loss rate projection can be expressed as a function of N , the number of bulk TCP flows, and S , the sum of the bandwidth-delay product and the maximum buffer size, both expressed in packets, based on Eq. 3 of [MORR2000]:

$$l_r = 0.76 * N^2 / S^2$$

$$N = S * \text{SQRT}(1/0.76) * \text{SQRT}(l_r)$$

These guidelines use the loss rate to define the different congestion levels, but they do not stipulate that in other circumstances, measuring the congestion level gives you an accurate estimation of the loss rate or vice-versa.

8.2.2. Mild congestion

This scenario can be used to evaluate how an AQM scheme reacts to a light load of incoming traffic resulting in mild congestion -- packet drop rates around 0.1%. The number of bulk flows required to achieve this congestion level, N_{mild} , is then:

$$N_{\text{mild}} = \text{ROUND}(0.036 * S)$$

8.2.3. Medium congestion

This scenario can be used to evaluate how an AQM scheme reacts to incoming traffic resulting in medium congestion -- packet drop rates around 0.5%. The number of bulk flows required to achieve this congestion level, N_{med} , is then:

$$N_{\text{med}} = \text{ROUND}(0.081 * S)$$

8.2.4. Heavy congestion

This scenario can be used to evaluate how an AQM scheme reacts to incoming traffic resulting in heavy congestion -- packet drop rates around 1%. The number of bulk flows required to achieve this congestion level, N_{heavy} , is then:

$$N_{\text{heavy}} = \text{ROUND} (0.114 * S)$$

8.2.5. Varying the congestion level

This scenario can be used to evaluate how an AQM scheme reacts to incoming traffic resulting in various levels of congestion during the experiment. In this scenario, the congestion level varies within a large time-scale. The following phases may be considered: phase I - mild congestion during 0-20s; phase II - medium congestion during 20-40s; phase III - heavy congestion during 40-60s; phase I again, and so on.

8.2.6. Varying available capacity

This scenario can be used to help characterize how the AQM behaves and adapts to bandwidth changes. The experiments are not meant to reflect the exact conditions of Wi-Fi environments since it is hard to design repetitive experiments or accurate simulations for such scenarios.

To emulate varying draining rates, the bottleneck capacity between nodes 'Router L' and 'Router R' varies over the course of the experiment as follows:

- o Experiment 1: the capacity varies between two values within a large time-scale. As an example, the following phases may be considered: phase I - 100Mbps during 0-20s; phase II - 10Mbps during 20-40s; phase I again, and so on.
- o Experiment 2: the capacity varies between two values within a short time-scale. As an example, the following phases may be considered: phase I - 100Mbps during 0-100ms; phase II - 10Mbps during 100-200ms; phase I again, and so on.

The tester MAY choose a phase time-interval value different than what is stated above, if the network's path conditions (such as bandwidth-delay product) necessitate. In this case the choice of such time-interval value SHOULD be stated and elaborated.

The tester MAY additionally evaluate the two mentioned scenarios (short-term and long-term capacity variations), during and/or including TCP slow-start phase.

More realistic fluctuating capacity patterns MAY be considered. The tester MAY choose to incorporate realistic scenarios with regards to common fluctuation of bandwidth in state-of-the-art technologies.

The scenario consists of TCP NewReno flows between sender A and receiver B. To better assess the impact of draining rates on the AQM behavior, the tester MUST compare its performance with those of drop-tail and SHOULD provide a reference document for their proposal discussing performance and deployment compared to those of drop-tail. Burst traffic, such as presented in Section 7.2, could also be considered to assess the impact of varying available capacity on the burst absorption of the AQM.

8.3. Parameter sensitivity and stability analysis

The control law used by an AQM is the primary means by which the queuing delay is controlled. Hence understanding the control law is critical to understanding the behavior of the AQM scheme. The control law could include several input parameters whose values affect the AQM scheme's output behavior and its stability. Additionally, AQM schemes may auto-tune parameter values in order to maintain stability under different network conditions (such as different congestion levels, draining rates or network environments). The stability of these auto-tuning techniques is also important to understand.

Transports operating under the control of AQM experience the effect of multiple control loops that react over different timescales. It is therefore important that proposed AQM schemes are seen to be stable when they are deployed at multiple points of potential congestion along an Internet path. The pattern of congestion signals (loss or ECN-marking) arising from AQM methods also need to not adversely interact with the dynamics of the transport protocols that they control.

AQM proposals SHOULD provide background material showing control theoretic analysis of the AQM control law and the input parameter space within which the control law operates as expected; or could use another way to discuss the stability of the control law. For parameters that are auto-tuned, the material SHOULD include stability analysis of the auto-tuning mechanism(s) as well. Such analysis helps to understand an AQM control law better and the network conditions/deployments under which the AQM is stable.

9. Various Traffic Profiles

This section provides guidelines to assess the performance of an AQM proposal for various traffic profiles such as traffic with different applications or bi-directional traffic.

9.1. Traffic mix

This scenario can be used to evaluate how an AQM scheme reacts to a traffic mix consisting of different applications such as:

- o Bulk TCP transfer
- o Web traffic
- o VoIP
- o Constant Bit Rate (CBR) UDP traffic
- o Adaptive video streaming

Various traffic mixes can be considered. These guidelines RECOMMEND to examine at least the following example: 1 bi-directional VoIP; 6 Web pages download (such as detailed in Section 7.2); 1 CBR; 1 Adaptive Video; 5 bulk TCP. Any other combinations could be considered and should be carefully documented.

For each scenario, the graph described in Section 2.7 could be generated for each class of traffic. Metrics such as end-to-end latency, jitter and flow completion time MAY be reported.

9.2. Bi-directional traffic

Control packets such as DNS requests/responses, TCP SYNs/ACKs are small, but their loss can severely impact the application performance. The scenario proposed in this section will help in assessing whether the introduction of an AQM scheme increases the loss probability of these important packets.

For this scenario, traffic MUST be generated in both downlink and uplink, such as defined in Section 3.1. These guidelines RECOMMEND to consider a mild congestion level and the traffic presented in Section 8.2.2 in both directions. In this case, the metrics reported MUST be the same as in Section 8.2 for each direction.

The traffic mix presented in Section 9.1 MAY also be generated in both directions.

AQM proposals SHOULD provide pseudo-code for the complete AQM scheme, highlighting generic implementation-specific aspects of the scheme such as "drop-tail" vs. "drop-head", inputs (e.g., current queuing delay, queue length), computations involved, need for timers, etc. This helps to identify costs associated with implementing the AQM scheme on a particular hardware or software device. This also facilitates discussions around which kind of devices can easily support the AQM and which cannot.

11.2. Recommended discussion

AQM proposals SHOULD highlight parts of their AQM logic that are device dependent and discuss if and how AQM behavior could be impacted by the device. For example, a queueing-delay based AQM scheme requires current queuing delay as input from the device. If the device already maintains this value, then it can be trivial to implement the their AQM logic on the device. If the device provides indirect means to estimate the queuing delay (for example: timestamps, dequeuing rate), then the AQM behavior is sensitive to the precision of the queuing delay estimations are for that device. Highlighting the sensitivity of an AQM scheme to queuing delay estimations helps implementers to identify appropriate means of implementing the mechanism on a device.

12. Operator Control and Auto-tuning

12.1. Motivation

One of the biggest hurdles of RED deployment was/is its parameter sensitivity to operating conditions -- how difficult it is to tune RED parameters for a deployment to achieve acceptable benefit from using RED. Fluctuating congestion levels and network conditions add to the complexity. Incorrect parameter values lead to poor performance.

Any AQM scheme is likely to have parameters whose values affect the control law and behaviour of an AQM. Exposing all these parameters as control parameters to a network operator (or user) can easily result in a unsafe AQM deployment. Unexpected AQM behavior ensues when parameter values are set improperly. A minimal number of control parameters minimizes the number of ways a user can break a system where an AQM scheme is deployed at. Fewer control parameters make the AQM scheme more user-friendly and easier to deploy and debug.

"AQM algorithms SHOULD NOT require tuning of initial or configuration parameters in common use cases." [RFC7567]. A scheme ought to

expose only those parameters that control the macroscopic AQM behavior such as queue delay threshold, queue length threshold, etc.

Additionally, the safety of an AQM scheme is directly related to its stability under varying operating conditions such as varying traffic profiles and fluctuating network conditions, as described in Section 8. Operating conditions vary often and hence the AQM needs to remain stable under these conditions without the need for additional external tuning. If AQM parameters require tuning under these conditions, then the AQM must self-adapt necessary parameter values by employing auto-tuning techniques.

12.2. Recommended discussion

In order to understand an AQM's deployment considerations and performance under a specific environment, AQM proposals SHOULD describe the parameters that control the macroscopic AQM behavior, and identify any parameters that require tuning to operational conditions. It could be interesting to also discuss that even if an AQM scheme may not adequately auto-tune its parameters, the resulting performance may not be optimal, but close to something reasonable.

If there are any fixed parameters within the AQM, their setting SHOULD be discussed and justified, to help understand whether a fixed parameter value is applicable for a particular environment.

If an AQM scheme is evaluated with parameter(s) that were externally tuned for optimization or other purposes, these values MUST be disclosed.

13. Conclusion

Figure 4 lists the scenarios and their requirements.

Scenario	Sec.	Requirement
Interaction with ECN	4.5	MUST be discussed if supported
Interaction with Scheduling	4.6	Feasibility MUST be discussed
Transport Protocols	5.	
TCP-friendly sender	5.1	Scenario MUST be considered
Aggressive sender	5.2	Scenario MUST be considered
Unresponsive sender	5.3	Scenario MUST be considered
LBE sender	5.4	Scenario MAY be considered
Round Trip Time Fairness	6.2	Scenario MUST be considered
Burst Absorption	7.2	Scenario MUST be considered
Stability	8.	
Varying congestion levels	8.2.5	Scenario MUST be considered
Varying available capacity	8.2.6	Scenario MUST be considered
Parameters and stability	8.3	This SHOULD be discussed
Various Traffic Profiles	9.	
Traffic mix	9.1	Scenario is RECOMMENDED
Bi-directional traffic	9.2	Scenario MAY be considered
Multi-AQM	10.2	Scenario MAY be considered
Implementation Cost	11.2	Pseudo-code SHOULD be provided
Operator Control	12.2	Tuning SHOULD NOT be required

Figure 4: Summary of the scenarios and their requirements

14. Acknowledgements

This work has been partially supported by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700).

15. Contributors

Many thanks to S. Akhtar, A.B. Bagayoko, F. Baker, R. Bless, D. Collier-Brown, G. Fairhurst, J. Gettys, P. Goltsman, T. Hoiland-Jorgensen, K. Kilkki, C. Kulatunga, W. Lautenschlager, A.C.

Morton, R. Pan, G. Skinner, D. Taht and M. Welzl for detailed and wise feedback on this document.

16. IANA Considerations

This memo includes no request to IANA.

17. Security Considerations

Some security considerations for AQM are identified in [RFC7567]. This document, by itself, presents no new privacy nor security issues.

18. References

18.1. Normative References

[I-D.ietf-tcpm-cubic]

Rhee, I., Xu, L., Ha, S., Zimmermann, A., Eggert, L., and R. Scheffenegger, "CUBIC for Fast Long-Distance Networks", draft-ietf-tcpm-cubic-01 (work in progress), January 2016.

[I-D.irtf-iccrgr-tcpeval]

Hayes, D., Ros, D., Andrew, L., and S. Floyd, "Common TCP Evaluation Suite", draft-irtf-iccrgr-tcpeval-01 (work in progress), July 2014.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<http://www.rfc-editor.org/info/rfc793>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, 1997.

[RFC2488] Allman, M., Glover, D., and L. Sanchez, "Enhancing TCP Over Satellite Channels using Standard Mechanisms", BCP 28, RFC 2488, DOI 10.17487/RFC2488, January 1999, <<http://www.rfc-editor.org/info/rfc2488>>.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <<http://www.rfc-editor.org/info/rfc2544>>.

[RFC2647] Newman, D., "Benchmarking Terminology for Firewall Performance", RFC 2647, DOI 10.17487/RFC2647, August 1999, <<http://www.rfc-editor.org/info/rfc2647>>.

- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, DOI 10.17487/RFC2679, September 1999, <<http://www.rfc-editor.org/info/rfc2679>>.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, DOI 10.17487/RFC2680, September 1999, <<http://www.rfc-editor.org/info/rfc2680>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC3611] Friedman, T., Ed., Caceres, R., Ed., and A. Clark, Ed., "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, DOI 10.17487/RFC3611, November 2003, <<http://www.rfc-editor.org/info/rfc3611>>.
- [RFC5348] Floyd, S., Handley, M., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 5348, DOI 10.17487/RFC5348, September 2008, <<http://www.rfc-editor.org/info/rfc5348>>.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, DOI 10.17487/RFC5481, March 2009, <<http://www.rfc-editor.org/info/rfc5481>>.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<http://www.rfc-editor.org/info/rfc5681>>.
- [RFC6297] Welzl, M. and D. Ros, "A Survey of Lower-than-Best-Effort Transport Protocols", RFC 6297, DOI 10.17487/RFC6297, June 2011, <<http://www.rfc-editor.org/info/rfc6297>>.
- [RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, DOI 10.17487/RFC6817, December 2012, <<http://www.rfc-editor.org/info/rfc6817>>.
- [RFC7141] Briscoe, B. and J. Manner, "Byte and Packet Congestion Notification", RFC 7141, 2014.
- [RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <<http://www.rfc-editor.org/info/rfc7567>>.

18.2. Informative References

- [ANEL2014] Anelli, P., Diana, R., and E. Lochin, "FavorQueue: a Parameterless Active Queue Management to Improve TCP Traffic Performance", *Computer Networks* vol. 60, 2014.
- [BB2011] "BufferBloat: what's wrong with the internet?", *ACM Queue* vol. 9, 2011.
- [GONG2014] Gong, Y., Rossi, D., Testa, C., Valenti, S., and D. Taht, "Fighting the bufferbloat: on the coexistence of AQM and low priority congestion control", *Computer Networks*, Elsevier, 2014, 60, pp.115 - 128 , 2014.
- [HASS2008] Hassayoun, S. and D. Ros, "Loss Synchronization and Router Buffer Sizing with High-Speed Versions of TCP", *IEEE INFOCOM Workshops* , 2008.
- [HOEI2015] Hoeiland-Joergensen, T., McKenney, P., Taht, D., Gettys, J., and E. Dumazet, "FlowQueue-Codel", *IETF (Work-in-Progress)* , January 2015.
- [JAY2006] Jay, P., Fu, Q., and G. Armitage, "A preliminary analysis of loss synchronisation between concurrent TCP flows", *Australian Telecommunication Networks and Application Conference (ATNAC)* , 2006.
- [MORR2000] Morris, R., "Scalable TCP congestion control", *IEEE INFOCOM* , 2000.
- [NICH2012] Nichols, K. and V. Jacobson, "Controlling Queue Delay", *ACM Queue* , 2012.
- [PAN2013] Pan, R., Natarajan, P., Piglione, C., Prabhu, MS., Subramanian, V., Baker, F., and B. VerSteeg, "PIE: A lightweight control scheme to address the bufferbloat problem", *IEEE HPSR* , 2013.

- [RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, April 1998.
- [TRAN2014] Trang, S., Kuhn, N., Lochin, E., Baudoin, C., Dubois, E., and P. Gelard, "On The Existence Of Optimal LEDBAT Parameters", IEEE ICC 2014 - Communication QoS, Reliability and Modeling Symposium , 2014.
- [WELZ2015] Welzl, M. and G. Fairhurst, "The Benefits to Applications of using Explicit Congestion Notification (ECN)", IETF (Work-in-Progress) , June 2015.
- [WINS2014] Winstein, K., "Transport Architectures for an Evolving Internet", PhD thesis, Massachusetts Institute of Technology , 2014.

Authors' Addresses

Nicolas Kuhn (editor)
CNES, Telecom Bretagne
18 avenue Edouard Belin
Toulouse 31400
France

Phone: +33 5 61 27 32 13
Email: nicolas.kuhn@cnes.fr

Preethi Natarajan (editor)
Cisco Systems
510 McCarthy Blvd
Milpitas, California
United States

Email: prenatar@cisco.com

Naeem Khademi (editor)
University of Oslo
Department of Informatics, PO Box 1080 Blindern
N-0316 Oslo
Norway

Phone: +47 2285 24 93
Email: naeemk@ifi.uio.no

David Ros
Simula Research Laboratory AS
P.O. Box 134
Lysaker, 1325
Norway

Phone: +33 299 25 21 21
Email: dros@simula.no