

Network Working Group
Request for Comments: 4883
Category: Informational

G. Feher
K. Nemeth
A. Korn
BUTE
I. Cselenyi
TeliaSonera
July 2007

Benchmarking Terminology for Resource Reservation Capable Routers

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

The primary purpose of this document is to define terminology specific to the benchmarking of resource reservation signaling of Integrated Services (IntServ) IP routers. These terms can be used in additional documents that define benchmarking methodologies for routers that support resource reservation or reporting formats for the benchmarking measurements.

Table of Contents

1. Introduction	2
2. Existing Definitions	3
3. Definition of Terms	4
3.1. Traffic Flow Types	4
3.1.1. Data Flow	4
3.1.2. Distinguished Data Flow	4
3.1.3. Best-Effort Data Flow	5
3.2. Resource Reservation Protocol Basics	5
3.2.1. QoS Session	5
3.2.2. Resource Reservation Protocol	6
3.2.3. Resource Reservation Capable Router	7
3.2.4. Reservation State	7
3.2.5. Resource Reservation Protocol Orientation	8
3.3. Router Load Factors	9
3.3.1. Best-Effort Traffic Load Factor	9
3.3.2. Distinguished Traffic Load Factor	10
3.3.3. Session Load Factor	11
3.3.4. Signaling Intensity Load Factor	11
3.3.5. Signaling Burst Load Factor	12
3.4. Performance Metrics	13
3.4.1. Signaling Message Handling Time	13
3.4.2. Distinguished Traffic Delay	14
3.4.3. Best-effort Traffic Delay	15
3.4.4. Signaling Message Deficit	15
3.4.5. Session Maintenance Capacity	16
3.5. Router Load Conditions and Scalability Limit	17
3.5.1. Loss-Free Condition	17
3.5.2. Lossy Condition	18
3.5.3. QoS Compliant Condition	19
3.5.4. Not QoS Compliant Condition	20
3.5.5. Scalability Limit	20
4. Security Considerations	21
5. Acknowledgements	21
6. References	21
6.1. Normative References	21
6.2. Informative References	21

1. Introduction

Signaling-based resource reservation using the IntServ paradigm [4] is an important part of the different Quality of Service (QoS) provisioning approaches. Therefore, network operators who are planning to deploy signaling-based resource reservation may want to examine the scalability limitations of reservation capable routers and the impact of signaling on their data forwarding performance.

An objective way of quantifying the scalability constraints of QoS signaling is to perform measurements on routers that are capable of IntServ-based resource reservation. This document defines terminology for a specific set of tests that vendors or network operators can carry out to measure and report the signaling performance characteristics of router devices that support resource reservation protocols. The results of these tests provide comparable data for different products, and thus support the decision-making process before purchase. Moreover, these measurements provide input characteristics for the dimensioning of a network in which resources are provisioned dynamically by signaling. Finally, the tests are applicable for characterizing the impact of the resource reservation signaling on the forwarding performance of the routers.

This benchmarking terminology document is based on the knowledge gained by examination of (and experimentation with) different resource reservation protocols: the IETF standard Resource ReSerVation Protocol (RSVP) [5], Next Steps in Signaling (NSIS) [6][7][8][9], and several experimental ones, such as YESSIR (Yet Another Sender Session Internet Reservation) [10], ST2+ [11], Session Description Protocol (SDP) [12], Boomerang [13], and Ticket [14]. Some of these protocols were also analyzed by the IETF NSIS working group [15]. Although at the moment the authors are only aware of resource reservation capable router products that interpret RSVP, this document defines terms that are valid in general and not restricted to any of the protocols listed above.

In order to avoid any confusion, we would like to emphasize that this terminology considers only signaling protocols that provide IntServ resource reservation; for example, techniques in the DiffServ toolbox are predominantly beyond our scope.

2. Existing Definitions

RFC 1242 "Benchmarking Terminology for Network Interconnection Devices" [1] and RFC 2285 "Benchmarking Terminology for LAN Switching Devices" [3] contain discussions and definitions for a number of terms relevant to the benchmarking of signaling performance of reservation-capable routers and should be consulted before attempting to make use of this document.

Additionally, this document defines terminology in a way that is consistent with the terms used by the Next Steps in Signaling working group laid out in [6][7][8].

For the sake of clarity and continuity, this document adopts the template for definitions set out in Section 2 of RFC 1242.

Definitions are indexed and grouped together into different sections for ease of reference.

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 [2].

3. Definition of Terms

3.1. Traffic Flow Types

This group of definitions describes traffic flow types forwarded by resource reservation capable routers.

3.1.1. Data Flow

Definition:

A data flow is a stream of data packets from one sender to one or more receivers, where each packet has a flow identifier unique to the flow.

Discussion:

The flow identifier can be an arbitrary subset of the packet header fields that uniquely distinguishes the flow from others. For example, the 5-tuple "source address; source port; destination address; destination port; protocol number" is commonly used for this purpose (where port numbers are applicable). It is also possible to take advantage of the Flow Label field of IPv6 packets. For more comments on flow identification, refer to [6].

3.1.2. Distinguished Data Flow

Definition:

Distinguished data flows are flows that resource reservation capable routers intentionally treat better or worse than best-effort data flows, according to a QoS agreement defined for the distinguished flow.

Discussion:

Routers classify the packets of distinguished data flows and identify the data flow to which they belong.

The most common usage of the distinguished data flow is to get higher-priority treatment than that of best-effort data flows (see the next definition). In these cases, a distinguished data flow is sometimes referred to as a "premium data flow". Nevertheless, theoretically it is possible to require worse treatment than that of best-effort flows.

3.1.3. Best-Effort Data Flow

Definition:

Best-effort data flows are flows that are not treated in any special manner by resource reservation capable routers; thus, their packets are served (forwarded) in some default way.

Discussion:

"Best-effort" means that the router makes its best effort to forward the data packet quickly and safely, but does not guarantee anything (e.g., delay or loss probability). This type of traffic is the most common in today's Internet.

Packets that belong to best-effort data flows need not be classified by the routers; that is, the routers don't need to find a related reservation session in order to find out to which treatment the packet is entitled.

3.2. Resource Reservation Protocol Basics

This group of definitions applies to signaling-based resource reservation protocols implemented by IP router devices.

3.2.1. QoS Session

Definition:

A QoS session is an application layer concept, shared between a set of network nodes, that pertains to a specific set of data flows. The information associated with the session includes the data required to identify the set of data flows in addition to a specification of the QoS treatment they require.

Discussion:

A QoS session is an end-to-end relationship. Whenever end-nodes decide to obtain special QoS treatment for their data communication, they set up a QoS session. As part of the process, they or their proxies make a QoS agreement with the network, specifying their data flows and the QoS treatment that the flows require.

It is possible for the same QoS session to span multiple network domains that have different resource provisioning architectures. In this document, however, we only deal with the case where the QoS session is realized over an IntServ architecture. It is assumed that sessions will be established using signaling messages of a resource reservation protocol.

QoS sessions must have unique identifiers; it must be possible to determine to which QoS session a given signaling message pertains. Therefore, each signaling message should include the identifier of its corresponding session. As an example, in the case of RSVP, the "session specification" identifies the QoS session plus refers to the data flow; the "flowspec" specifies the desired QoS treatment and the "filter spec" defines the subset of data packets in the data flow that receive the QoS defined by the flowspec.

QoS sessions can be unicast or multicast depending on the number of participants. In a multicast group, there can be several data traffic sources and destinations. Here the QoS agreement does not have to be the same for each branch of the multicast tree forwarding the data flow of the group. Instead, a dedicated network resource in a router can be shared among many traffic sources from the same multicast group (cf. multicast reservation styles in the case of RSVP).

Issues:

Even though QoS sessions are considered to be unique, resource reservation capable routers might aggregate them and allocate network resources to these aggregated sessions at once. The aggregation can be based on similar data flow attributes (e.g., similar destination addresses) or it can combine arbitrary sessions as well. While reservation aggregation significantly lightens the signaling processing task of a resource reservation capable router, it also requires the administration of the aggregated QoS sessions and might also lead to the violation of the quality guaranties referring to individual data flows within an aggregation [16].

3.2.2. Resource Reservation Protocol

Definition:

Resource reservation protocols define signaling messages and message processing rules used to control resource allocation in IntServ architectures.

Discussion:

It is the signaling messages of a resource reservation protocol that carry the information related to QoS sessions. This information includes a session identifier, the actual QoS parameters, and possibly flow descriptors.

The message processing rules of the signaling protocols ensure that signaling messages reach all network nodes concerned. Some resource reservation protocols (e.g., RSVP, NSIS QoS NSLP [8]) are only concerned with this, i.e., carrying the QoS-related

information to all the appropriate network nodes, without being aware of its content. This latter approach allows changing the way the QoS parameters are described, and different kinds of provisioning can be realized without the need to change the protocol itself.

3.2.3. Resource Reservation Capable Router

Definition:

A router is resource reservation capable (it supports resource reservation) if it is able to interpret signaling messages of a resource reservation protocol, and based on these messages is able to adjust the management of its flow classifiers and network resources so as to conform to the content of the signaling messages.

Discussion:

Routers capture signaling messages and manipulate reservation states and/or reserved network resources according to the content of the messages. This ensures that the flows are treated as their specified QoS requirements indicate.

3.2.4. Reservation State

Definition:

A reservation state is the set of entries in the router's memory that contain all relevant information about a given QoS session registered with the router.

Discussion:

States are needed because IntServ-related resource reservation protocols require the routers to keep track of QoS session and data-flow-related metadata. The reservation state includes the parameters of the QoS treatment, the description of how and where to forward the incoming signaling messages, refresh timing information, etc.

Based on how reservation states are stored in a reservation capable router, the routers can be categorized into two classes:

Hard-state resource reservation protocols (e.g., ST2 [11]) require routers to store the reservation states permanently, established by a setup signaling primitive, until the router is explicitly informed that the QoS session is canceled.

There are also soft-state resource reservation capable routers, where there are no permanent reservation states, and each state has to be regularly refreshed by appropriate refresh signaling

messages. If no refresh signaling message arrives during a certain period, then the router stops the maintenance of the QoS session assuming that the end-points do not intend to keep the session up any longer or the communication lines are broken somewhere along the data path. This feature makes soft-state resource reservation capable routers more robust than hard-state routers, since no failures can cause resources to stay permanently stuck in the routers. (Note that it is still possible to have an explicit teardown message in soft-state protocols for quicker resource release.)

Issues:

Based on the initiating point of the refresh messages, soft-state resource reservation protocols can be divided into two groups. First, there are protocols where it is the responsibility of the end-points or their proxies to initiate refresh messages. These messages are forwarded along the path of the data flow refreshing the corresponding reservation states in each router affected by the flow. Second, there are other protocols, where routers and end-points have their own schedule for the reservation state refreshes and they signal these refreshes to the neighboring routers.

3.2.5. Resource Reservation Protocol Orientation

Definition:

The orientation of a resource reservation protocol tells which end of the protocol communication initiates the allocation of the network resources. Thus, the protocol can be sender- or receiver-oriented, depending on the location of the data flow source (sender) and destination (receiver) compared to the reservation initiator.

Discussion:

In the case of sender-oriented protocols (in some sources referred to as sender-initiated protocols), the resource reservation propagates in the same direction(s) as of the data flow(s). Consequently, in the case of receiver-oriented protocols, the signaling messages reserving resources are forwarded backward on the path of the data flow. Due to the asymmetric routing nature of the Internet, in this latter case, the path of the desired data flow should be known before the reservation initiator would be able to send the resource allocation messages. For example, in the case of RSVP, the RSVP PATH message, traveling from the data flow sources towards the destinations, first marks the path of the data flow on which the resource allocation messages will travel backward.

This definition considers only protocols that reserve resources for just one data flow between the end-nodes. The reservation orientation of protocols that reserve more than one data flow is not defined here.

Issues:

The location of the reservation initiator affects the basics of the resource reservation protocols and therefore is an important aspect of characterization. Most importantly, in the case of multicast QoS sessions, the sender-oriented protocols require the traffic sources to maintain a list of receivers and send their allocation messages considering the different requirements of the receivers. Using multicast QoS sessions, the receiver-oriented protocols enable the receivers to manage their own resource allocation requests and thus ease the task of the sources.

3.3. Router Load Factors

When a router is under "load", it means that there are tasks its CPU(s) must attend to, and/or that its memory contains data it must keep track of, and/or that its interface buffers are utilized to some extent, etc. Unfortunately, we cannot assume that the full internal state of a router can be monitored during a benchmark; rather, we must consider the router to be a black box.

We need to look at router "load" in a way that makes this "load" measurable and controllable. Instead of focusing on the internal processes of a router, we will consider the external, and therefore observable, measurable and controllable processes that result in "load".

In this section we introduce several ways of creating "load" on a router; we will refer to these as "load factors" henceforth. These load factors are defined so that they each impact the performance of the router in a different way (or by different means), by utilizing different components of a resource reservation capable router as separately as possible.

During a benchmark, the performance of the device under test will have to be measured under different controlled load conditions, that is, with different values of these load factors.

3.3.1. Best-Effort Traffic Load Factor

Definition:

The best-effort traffic load factor is defined as the number and length of equal-sized best-effort data packets that traverse the router in a second.

Discussion:

Forwarding the best-effort data packets, which requires obtaining the routing information and transferring the data packet between network interfaces, requires processing power. This load factor creates load on the CPU(s) and buffers of the router.

For the purpose of benchmarking, we define a traffic flow as a stream of equal-sized packets with even interpacket delay. It is possible to specify traffic with varying packet sizes as a superposition of multiple best-effort traffic flows as they are defined here.

Issues:

The same amount of data segmented into differently sized packets causes different amounts of load on the router, which has to be considered during benchmarking measurements. The measurement unit of this load factor reflects this as well.

Measurement unit:

This load factor has a composite unit of [packets per second (pps); bytes]. For example, [5 pps; 100 bytes] means five pieces of one-hundred-byte packets per second.

3.3.2. Distinguished Traffic Load Factor**Definition:**

The distinguished traffic load factor is defined as the number and length of the distinguished data packets that traverse the router in a second.

Discussion:

Similarly to the best-effort data, forwarding the distinguished data packets requires obtaining the routing information and transferring the data packet between network interfaces. However, in this case packets have to be classified as well, which requires additional processing capacity.

For the purpose of benchmarking, we define a traffic flow as a stream of equal-sized packets with even interpacket delay. It is possible to specify traffic with varying packet sizes as a superposition of multiple distinguished traffic flows as they are defined here.

Issues:

Just as in the best-effort case, the same amount of data segmented into differently sized packets causes different amounts of load on the router, which has to be considered during the benchmarking

measurements. The measurement unit of this load factor reflects this as well.

Measurement unit:

This load factor has a composite unit of [packets per second (pps); bytes]. For example, [5 pps; 100 bytes] means five pieces of one-hundred-byte packets per second.

3.3.3. Session Load Factor

Definition:

The session load factor is the number of QoS sessions the router is keeping track of.

Discussion:

Resource reservation capable routers maintain reservation states to keep track of QoS sessions. Obviously, the more reservation states are registered with the router, the more complex the traffic classification becomes, and the more time it takes to look up the corresponding resource reservation state. Moreover, not only the traffic flows, but also the signaling messages that control the reservation states have to be identified first, before taking any other action, and this kind of classification also means extra work for the router.

In the case of soft-state resource reservation protocols, the session load also affects reservation state maintenance. For example, the supervision of timers that watchdog the reservation state refreshes may cause further load on the router.

This load factor utilizes the CPU(s), the main memory, and the session management logic (e.g., content addressable memory), if any, of the resource reservation capable router.

Measurement unit:

This load component is measured by the number of QoS sessions that impact the router.

3.3.4. Signaling Intensity Load Factor

Definition:

The signaling intensity load factor is the number of signaling messages that are presented at the input interfaces of the router during one second.

Discussion:

The processing of signaling messages requires processor power that raises the load on the control plane of the router.

In routers where the control plane and the data plane are not totally independent (e.g., certain parts of the tasks are served by the same processor; or the architecture has common memory buffers, transfer buses or any other resources) the signaling load can have an impact on the router's packet forwarding performance as well.

Naturally, just as everywhere else in this document, the term "signaling messages" refer only to the resource reservation protocol related primitives.

Issues:

Most resource reservation protocols have several protocol primitives realized by different signaling message types. Each of these message types may require a different amount of processing power from the router. This fact has to be considered during the benchmarking measurements.

Measurement unit:

The unit of this factor is signaling messages/second.

3.3.5. Signaling Burst Load Factor

Definition:

The signaling burst load factor is defined as the number of signaling messages that arrive to one input port of the router back-to-back ([1]), causing persistent load on the signaling message handler.

Discussion:

The definition focuses on one input port only and does not consider the traffic arriving at the other input ports. As a consequence, a set of messages arriving at different ports, but with such a timing that would be a burst if the messages arrived at the same port, is not considered to be a burst. The reason for this is that it is not guaranteed in a black-box test that this would have the same effect on the router as a burst (incoming at the same interface) has.

This definition conforms to the burst definition given in [3].

Issues:

Most of the resource reservation protocols have several protocol primitives realized by different signaling message types. Bursts built up of different messages may have a different effect on the router. Consequently, during measurements the content of the burst has to be considered as well.

Likewise, the first one of multiple idempotent signaling messages that each accomplish exactly the same end will probably not take the same amount of time to be processed as subsequent ones. Benchmarking methodology will have to consider the intended effect of the signaling messages, as well as the state of the router at the time of their arrival.

Measurement unit:

This load factor is characterized by the number of messages in the burst.

3.4. Performance Metrics

This group of definitions is a collection of measurable quantities that describe the performance impact the different load components have on the router.

During a benchmark, the values of these metrics will have to be measured under different load conditions.

3.4.1. Signaling Message Handling Time

Definition:

The signaling message handling time (or, in short, signal handling time) is the latency ([1], for store-and-forward devices) of a signaling message passing through the router.

Discussion:

The router interprets the signaling messages, acts based on their content and usually forwards them in an unmodified or modified form. Thus the message handling time is usually longer than the forwarding time of data packets of the same size.

There might be signaling message primitives, however, that are drained or generated by the router, like certain refresh messages. In this case, the signal handling time is not necessarily measureable, therefore it is not defined for such messages.

In the case of signaling messages that carry information pertaining to multicast flows, the router might issue multiple signaling messages after processing them. In this case, by definition, the signal handling time is the latency between the incoming signaling message and the last outgoing signaling message related to the received one.

The signal handling time is an important characteristic as it directly affects the setup time of a QoS session.

Issues:

The signal handling time may be dependent on the type of the signaling message. For example, it usually takes a shorter time for the router to remove a reservation state than to set it up. This fact has to be considered during the benchmarking process.

As noted above, the first one of multiple idempotent signaling messages that each accomplish exactly the same end will probably not take the same amount of time to be processed as subsequent ones. Benchmarking methodology will have to consider the intended effect of the signaling messages, as well as the state of the router at the time of their arrival.

Measurement unit:

The dimension of the signaling message handling time is the second, reported with a resolution sufficient to distinguish between different events/DUTs (e.g., milliseconds). Reported results MUST clearly indicate the time unit used.

3.4.2. Distinguished Traffic Delay**Definition:**

Distinguished traffic delay is the latency ([1], for store-and-forward devices) of a distinguished data packet passing through the tested router device.

Discussion:

Distinguished traffic packets must be classified first in order to assign the network resources dedicated to the flow. The time of the classification is added to the usual forwarding time (including the queuing) that a router would spend on the packet without any resource reservation capability. This classification procedure might be quite time consuming in routers with vast amounts of reservation states.

There are routers where the processing power is shared between the control plane and the data plane. This means that the processing of signaling messages may have an impact on the data forwarding performance of the router. In this case, the distinguished traffic delay metric also indicates the influence the two planes have on each other.

Issues:

Queuing of the incoming data packets in routers can bias this metric, so the measurement procedures have to consider this effect.

Measurement unit:

The dimension of the distinguished traffic delay time is the second, reported with resolution sufficient to distinguish between different events/DUTs (e.g., millisecond units). Reported results MUST clearly indicate the time unit used.

3.4.3. Best-effort Traffic Delay

Definition:

Best-effort traffic delay is the latency of a best-effort data packet traversing the tested router device.

Discussion:

If the processing power of the router is shared between the control and data plane, then the processing of signaling messages may have an impact on the data forwarding performance of the router. In this case, the best-effort traffic delay metric is an indicator of the influence the two planes have on each other.

Issues:

Queuing of the incoming data packets in routers can bias this metric as well, so measurement procedures have to consider this effect.

Measurement unit:

The dimension of the best-effort traffic delay is the second, reported with resolution sufficient to distinguish between different events/DUTs (e.g., millisecond units). Reported results MUST clearly indicate the time unit used.

3.4.4. Signaling Message Deficit

Definition:

Signaling message deficit is one minus the ratio of the actual and the expected number of signaling messages leaving a resource reservation capable router.

Discussion:

This definition gives the same value as the ratio of the lost (that is, not forwarded or not generated) and the expected messages. The above calculation must be used because the number of lost messages cannot be measured directly.

There are certain types of signaling messages that reservation capable routers are required to forward as soon as their processing is finished. However, due to lack of resources or other reasons, the forwarding or even the processing of these signaling messages might not take place.

Certain other kinds of signaling messages must be generated by the router in the absence of any corresponding incoming message. It is possible that an overloaded router does not have the resources necessary to generate such a message.

To characterize these situations we introduce the signaling message deficit metric that expresses the ratio of the signaling messages that have actually left the router and those ones that were expected to leave the router. We subtract this ratio from one in order to obtain a loss-type metric instead of a "message survival metric".

Since the most frequent reason for signaling message deficit is high router load, this metric is suitable for sounding out the scalability limits of resource reservation capable routers.

During the measurements one must be able to determine whether a signaling message is still in the queues of the router or if it has already been dropped. For this reason we define a signaling message as lost if no forwarded signaling message is emitted within a reasonably long time period. This period is defined along with the benchmarking methodology.

Measurement unit:

This measure has no unit; it is expressed as a real number, which is between zero and one, including the limits.

3.4.5. Session Maintenance Capacity

Definition:

The session maintenance capacity metric is used in the case of soft-state resource reservation protocols only. It is defined as the ratio of the number of QoS sessions actually being maintained and the number of QoS sessions that should have been maintained.

Discussion:

For soft-state protocols maintaining a QoS session means refreshing the reservation states associated with it.

When a soft-state resource reservation capable router is overloaded, it may happen that the router is not able to refresh all the registered reservation states, because it does not have the time to run the state refresh task. In this case, sooner or later some QoS sessions will be lost even if the endpoints still require their maintenance.

The session maintenance capacity sounds out the maximal number of QoS sessions that the router is capable of maintaining.

Issues:

The actual process of session maintenance is protocol and implementation dependent, thus so is the method to examine whether a session is maintained or not.

In the case of soft-state resource reservation protocols, where the network nodes are responsible for generating the refresh messages, a router that fails to maintain a QoS session may not emit refresh signaling messages either. This has direct consequences on the signaling message deficit metric.

Measurement unit:

This measure has no unit; it is expressed as a real number, which is between zero and one (including the limits).

3.5. Router Load Conditions and Scalability Limit

Depending mainly, but not exclusively, on the overall load of a router, it can be in exactly one of the following four conditions at a time: loss-free and QoS compliant; lossy and QoS compliant; loss-free but not QoS compliant; and neither loss-free nor QoS compliant. These conditions are defined below, along with the scalability limit.

3.5.1. Loss-Free Condition

Definition:

A router is in loss-free condition, or loss-free state, if and only if it is able to perform its tasks correctly and in a timely fashion.

Discussion:

All existing routers have finite buffer memory and finite processing power. If a router is in loss-free state, the buffers of the router still contain enough free space to accommodate the next incoming packet when it arrives. Also, the router has enough processing power to cope with all its tasks, thus all required operations are carried out within the time the protocol specification allows; or, if this time is not specified by the protocol, then in "reasonable time" (which is then defined in the benchmarks). Similar considerations can be applied to other resources a router may have, if any; in loss-free states, the utilization of these resources still allows the router to carry out its tasks in accordance with applicable protocol specifications and in "reasonable time".

Note that loss-free states as defined above are not related to the reservation states of resource reservation protocols. The word "state" is used to mean "condition".

Also note that it is irrelevant what internal reason causes a router to fail to perform in accordance with protocol specifications or in "reasonable time"; if it is not high load but -- for example -- an implementation error that causes the device to perform inadequately, it still cannot be said to be in a loss-free state. The same applies to the random early dropping of packets in order to prevent congestion. In a black-box measurement it is impossible to determine whether a packet was dropped as part of a congestion control mechanism or because the router was unable to forward it; therefore, if packet loss is observed except as noted below, the router is by definition in lossy state (lossy condition).

If a distinguished data flow exceeds its allotted bandwidth, it is acceptable for routers to drop excess packets. Thus, a router that is QoS Compliant (see below) is also loss-free provided that it only drops packets from distinguished data flows.

If a device is not in a loss-free state, it is in a lossy condition/state.

Related definitions:

- Lossy Condition
- QoS Compliant Condition
- Not QoS Compliant Condition
- Scalability Limit

3.5.2. Lossy Condition

Definition:

A router is in a lossy condition, or lossy state, if it cannot perform its duties adequately for some reason; that is, if it does not meet protocol specifications (except QoS guarantees, which are treated separately), or -- if time-related specifications are missing -- doesn't complete some operations in "reasonable time" (which is then defined in the benchmarks).

Discussion:

A router may be in a lossy state for several reasons, including but not necessarily limited to the following:

- a) Buffer memory has run out, so either an incoming or a buffered packet has to be dropped.

- b) The router doesn't have enough processing power to cope with all its duties. Some required operations are skipped, aborted or suffer unacceptable delays.
- c) Some other finite internal resource is exhausted.
- d) The router runs a defective (non-conforming) protocol implementation.
- e) Hardware malfunction.
- f) A congestion control mechanism is active.

Loss can mean the loss of data packets as well as signaling message deficit.

A router that does not lose data packets and does not experience signaling message deficit but fails to meet required QoS parameters is in the loss-free, but not in the QoS compliant state.

If a device is not in a lossy state, it is in a loss-free condition/state.

Related definitions:

- Loss-Free Condition (especially the discussion of congestion control mechanisms that cause packet loss)
- Scalability Limit
- Signaling Message Deficit
- QoS Compliant Condition
- Not QoS Compliant Condition

3.5.3. QoS Compliant Condition

Definition:

A router is in the QoS compliant state if and only if all distinguished data flows receive the QoS treatment they are entitled to.

Discussion:

Defining what specific QoS guarantees must be upheld is beyond the scope of this document because every reservation model may specify a different set of such parameters.

Loss, delay, jitter etc. of best-effort data flows are irrelevant when considering whether a router is in the QoS compliant state.

Related definitions:

- Loss-Free Condition
- Lossy Condition
- Not QoS Compliant Condition
- Scalability Limit

3.5.4. Not QoS Compliant Condition

Definition:

A router is in the not QoS compliant state if and only if it is not in the QoS compliant condition.

Related definitions:

- Loss-Free Condition
- Lossy Condition
- QoS Compliant Condition
- Scalability Limit

3.5.5. Scalability Limit

Definition:

The scalability limits of a router are the boundary load conditions where the router is still in the loss-free and QoS compliant state, but the smallest amount of additional load would drive it to a state that is either QoS compliant but not loss-free, or not QoS compliant but loss-free, or neither loss-free nor QoS compliant.

Discussion:

An unloaded router that operates correctly is in a loss-free and QoS compliant state. As load increases, the resources of the router are becoming more and more utilized. At a certain point, the router enters a state that is either not QoS compliant, or not loss-free, or neither QoS compliant nor loss-free. Note that such a point may be impossible to reach in some cases (for example if the bandwidth of the physical medium prevents increasing the traffic load any further).

A particular load condition can be identified by the corresponding values of the load factors (as defined in 3.3 Router Load Factors) impacting the router. These values can be represented as a 7-tuple of numbers (there are only five load factors, but the traffic load factors have composite units and thus require two numbers each to express). We can think of these tuples as vectors that correspond to a state that is either both loss free and QoS compliant, or not loss-free (but QoS compliant), or not QoS compliant (but loss-free), or neither loss-free nor QoS compliant. The scalability limit of the router is, then, the boundary between

the sets of vectors corresponding to the loss-free and QoS compliant states and all other states. Finding these boundary points is one of the objectives of benchmarking.

Benchmarks may try to separately identify the boundaries of the loss-free and of the QoS compliant conditions in the (seven-dimensional) space defined by the load-vectors.

Related definitions:

- Lossy Condition
- Loss-Free Condition
- QoS Compliant Condition
- Non QoS Compliant Condition

4. Security Considerations

As this document only provides terminology and does not describe a protocol, an implementation, or a procedure, there are no security considerations associated with it.

5. Acknowledgements

We would like to thank Telia Research AB, Sweden and the High Speed Networks Laboratory at the Department of Telecommunication and Media Informatics of the Budapest University of Technology and Economics, Hungary for their support in the research and development work, which contributed to the creation of this document.

6. References

6.1. Normative References

- [1] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.

6.2. Informative References

- [4] Braden, R., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.

- [5] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [6] Hancock, R., Karagiannis, G., Loughney, J., and S. Van den Bosch, "Next Steps in Signaling (NSIS): Framework", RFC 4080, June 2005.
- [7] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signaling Transport", Work in Progress, April 2007.
- [8] Manner, J., Ed., Karagiannis, G., and A. McDonald, "NSLP for Quality-of-Service Signaling", Work in Progress, June 2007.
- [9] Ash, J., Bader, A., Kappler, C., and D. Oran, "QoS NSLP QSPEC Template", Work in Progress, March 2007.
- [10] P. Pan, H. Schulzrinne, "YESSIR: A Simple Reservation Mechanism for the Internet", Computer Communication Review, on-line version, volume 29, number 2, April 1999
- [11] Delgrossi, L. and L. Berger, "Internet Stream Protocol Version 2 (ST2) Protocol Specification - Version ST2+", RFC 1819, August 1995.
- [12] P. White, J. Crowcroft, "A Case for Dynamic Sender-Initiated Reservation in the Internet", Journal on High Speed Networks, Special Issue on QoS Routing and Signaling, Vol. 7 No. 2, 1998
- [13] J. Bergkvist, D. Ahlhard, T. Engborg, K. Nemeth, G. Feher, I. Cselenyi, M. Maliosz, "Boomerang : A Simple Protocol for Resource Reservation in IP Networks", Vancouver, IEEE Real-Time Technology and Applications Symposium, June 1999
- [14] A. Eriksson, C. Gehrman, "Robust and Secure Light-weight Resource Reservation for Unicast IP Traffic", International WS on QoS'98, IWQoS'98, May 18-20, 1998
- [15] Manner, J. and X. Fu, "Analysis of Existing Quality-of-Service Signaling Protocols", RFC 4094, May 2005.
- [16] Baker, F., Iturralde, C., Le Faucheur, F., and B. Davie, "Aggregation of RSVP for IPv4 and IPv6 Reservations", RFC 3175, September 2001.

Authors' Addresses

Gabor Feher
Budapest University of Technology and Economics
Department of Telecommunications and Media Informatics
Magyar Tudosok krt. 2, H-1117, Budapest, Hungary

Phone: +36 1 463-1538
EMail: Gabor.Feher@tmit.bme.hu

Krisztian Nemeth
Budapest University of Technology and Economics
Department of Telecommunications and Media Informatics
Magyar Tudosok krt. 2, H-1117, Budapest, Hungary

Phone: +36 1 463-1565
EMail: Krisztian.Nemeth@tmit.bme.hu

Andras Korn
Budapest University of Technology and Economics
Department of Telecommunication and Media Informatics
Magyar Tudosok krt. 2, H-1117, Budapest, Hungary

Phone: +36 1 463-2664
EMail: Andras.Korn@tmit.bme.hu

Istvan Cselenyi
Teliasonera International Carrier
Vaci ut 22-24, H-1132 Budapest, Hungary

Phone: +36 1 412-2705
EMail: Istvan.Cselenyi@teliasonera.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

