Network Working Group
INTERNET-DRAFT
Expiration Date: January 2002

Gabor Feher, BUTE
Istvan Cselenyi, TRAB
Andras Korn, BUTE

July 2001

Benchmarking Methodology for Routers Supporting Resource Reservation

<draft-ietf-bmwg-benchres-method-00.txt>

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft shadow directories can be accessed at http://www.ietf.org/shadow.html

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

2. Table of contents

3. Abstract

The purpose of this document is to define benchmarking methodology measuring performance metrics related to IP routers supporting resource reservation signaling. Apart from the definition and discussion of these tests, this document also specifies formats for reporting the benchmarking results.

4. Introduction

The IntServ over DiffServ framework [1] outlines a heterogeneous Quality of Service (QoS) architecture for multi domain Internet services. Signaling based resource reservation (e.g. via RSVP [2]) is an integral part of that model. While this significantly lightens the load on most of the core routers, the performance of border routers that handle the QoS signaling is still crucial. Therefore network operators, who are planning to deploy this model, shall scrutinize the scalability limitations in reservation capable routers and the impact of signaling on the forwarding performance of the routers.

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

This benchmarking methodology document is based on the knowledge gained by examination of (and experimentation with) several very different resource reservation protocols: RSVP [2], Boomerang [3], YESSIR [4], ST2+ [5], SDP [6], Ticket [7] and Load Control [8]. Nevertheless, this document aspires to compose terms that are valid in general and not restricted to these protocols.

5. Existing definitions

A previous document, "Benchmarking Terminology for Routers Supporting

Resource Reservation" [9] defines performance metrics and other terms that are used in this document. To understand the test methodologies defined here, that terminology document must be consulted first.

6. Methodology

6.1 Evaluating the Results

RFC2544 [10] describes considerations regarding the implementation and evaluation of benchmarking tests, which are certainly valid for this test suite also. Namely, the authors intended to create a system from commercially available measurement instruments and devices for the sake of easy implementation of the described tests. Simple test scripts and benchmarking utilities for Linux are publicly available from the Boomerang homepage [11].

During the benchmarking tests, care should be taken for selecting the proper set of tests for a specific router device, since not all of the tests applicable to a particular Devices Under Test (DUT).

Finally, the selection of the relevant measurement results and their evaluation requires experience and it must be done with an understanding of generally accepted testing practices regarding repeatability, variance and statistical significance of small numbers of trials.

6.2 Test Setup

The ideal way to perform the measurements is to connect a passive tester device (or, in short, passive tester) to all network interfaces of the DUT, enabling the tester to capture all signaling and data traffic that enters into or leaves from the DUT. Based on the captured data packets and signaling messages along with the proper time stamps the investigated performance metrics can be computed. In addition to the passive tester there are signaling and data traffic end-points that are responsible to generate and terminate the required signaling and data flows going through the DUT. These flows are used to generate router load in the DUT and the measurements are also performed using them. This scenario is illustrated in Figure 1.

Probably, the best solution is to connect the tester via network traffic repeater devices (e.g. hubs) to the network interfaces of the DUT. These repeaters cause very small delay in the ongoing packets, and therefore their effect is insignificant in the measurements.

```
                             +------------+
                             |            |
                        +--->|  Passive   |<---+
                        |    |   tester   |    |
                        |    +------------+    |
                        |                      |
    +---------------+   |    +------------+    |    +---------------+
    | Signaling and |   |    |            |    |    | Signaling and |
    |  data traffic |----+--->|    DUT     |----+--->|  data traffic |
    |   end-point   |        |            |         |   end-point   |
    +---------------+        +------------+         +---------------+
```

Figure 1

Moreover, tester devices should not have to be passive during the measurement, rather they can generate the signaling and data flows as well. This way the signaling and data traffic end-point and the traffic capturing device can be combined into a single tester device, called active tester. In this case the signaling and traffic flow, the initiator tester device is the driver of the input network interfaces of the DUT, while the second one, the signaling and traffic terminator tester device is connected to the output network interfaces of the tested device and captures signaling messages and data packets leaving the DUT. Figure 2 shows this scenario.

```
        +---------------+      +------------+      +---------------+
        |               |      |            |      |               |
        | Active tester |----->|    DUT     |----->| Active tester |
        |               |      |            |      |               |
        +---------------+      +------------+      +---------------+
```

Figure 2

In this scenario, the performance metrics are calculated from the log of initiated packets and their initiation time in the first active tester device and the log of captured packets and their capture time in the second active tester. Obviously, the measurements do worth nothing if the two testers are not clock-synchronized, since the difference of the packet initiation times and packet capture times is biased by the clock skew of the testers. For this reason, the clock of the testers must be synchronized before the measurements are performed. Nevertheless, scalability tests do not depend on the clock synchronization and therefore they can be performed without any preparation on the testers.

It is also possible to use only one active tester, which is the signaling and traffic flow initiator and terminator device in the same time. Although, this way the clock synchronization problem can be avoided, but the tester should be powerful enough to generate and capture all the test flows required by the measurements.

During the benchmarking tests, if the clocks are properly synchronized when it is necessary, each test configuration is suitable for the measurements. For this reason, we have not defined different test methodologies for each test scenarios. Instead, we use terms "initiator tester" and "terminator tester", which have their equivalent appliances in each test configuration.

Initiator tester is the device that generates the signaling and data flows, while terminator tester is the device that terminates the signaling and data flow. In addition, the performance metrics measurement is also performed by the tester(s). Evidently, in the case of the configuration, where there is only one active tester, the initiator tester and the terminator tester is the same appliance.

6.2.1 Testing Unicast Resource Reservation Sessions

Testing unicast resource reservation sessions requires that the initial tester is connected to one of the network interfaces of the DUT and the terminator tester is connected to a different network interface of the tested device.

During the benchmarking tests, the initiator tester must use unicast addresses for data traffic flows and the resource reservation requests must refer to unicast resource reservation sessions. In order to be able to compute the performance metrics, all data packets and signaling messages transmitted by the DUT must be perceivable for the tester.

6.2.2 Testing Multicast Resource Reservation Sessions

Testing multicast resource reservation sessions requires the initial tester to be connected to more than one network interfaces of the DUT, while the terminator tester is connected to more than one network interfaces of the tested device whose interfaces are different from the previous ones.

Furthermore, during the measurements, the data traffic flows originated from the initiator tester must be sent to multicast addresses and the reservation sessions must refer to one or more of the multicast flows. Of course, just like in the case of unicast resource reservation sessions, all data packets and signaling messages transmitted by the DUT must be perceivable for the tester.

Since there are protocols supporting more than one resource reservation schemes for multicast reservations (e.g. RSVP SE/FF/WF); and in a view of the fact that the number of incoming and outgoing network interface combinations of the DUT might be almost countless; the benchmarking tests, described here, do not require measuring all imaginable setup situation. Still, routers supporting multicast resource reservations must be tested against the performance metrics and scalability limits on at least one multicast scenario. Moreover, there is a suggested multicast test configuration that consists of a multicast group with four signaling end-points including one traffic originator and three traffic destinations residing on different network interfaces of the DUT.

The benchmarking test reports taken on DUTs supporting multicast resource reservation sessions always have to contain the proper multicast scenario description.

6.2.3 Signaling Flow

This document often refers to signaling flows. A signaling flow is sequence of signaling messages.

In the case of the measurements defined in this document there are two types of signaling flows: First, there is a signaling flow that is constructed from signaling primitives of the same type. Second, there is a signaling flow that is constructed from signaling primitive pairs. Signaling primitive pairs are needed in situations where one of the signaling primitive alters the states of the DUT, but the test demand constant DUT conditions during the test. In this case, to avoid the effect of the state modification, the second signaling primitive should restore the states modification in the DUT. A typical example for the second type of signaling flow is a flow of alternating reservation set-up and tear-down messages.

Moreover, the signaling messages should be equally spaced on the time scale when they are forming a signaling flow. This is mandatory in order to obtain measurements that can be repeated later. Since modern resource reservation protocols are designed to avoid message synchronization, thus, equally spaced signaling messages are not unrealistic in the real life.

The signaling flow is characterized with the type of the signaling primitive or the pair of signaling primitives along with the period time of the signaling messages.

6.2.4 Signaling Message Verification

Although, the conformance testing of the resource reservation is beyond the scope of this document, defective signaling message processing can be expected in an overloaded router. Therefore, during the benchmarking tests, when signaling messages are processed in the DUT, the terminator device must validate the messages whether they are fully conform to the message format of the resource reservation protocol specification and whether they are the expected signaling messages at the given situation. If any of the messages are against the protocol specification then the benchmarking test report must indicate the situation of the failure.

Verifying data traffic packets are not required, since the signaling performance benchmarking of reservation capable routers should not deal with data traffic. For this purpose there are other benchmarking methodologies that verify data traffic during the measurements, like the one described in RFC 2544.

6.3 Scalability Tests

Scalability tests are defined to explore the scalability limits of a reservation capable router. This investigation focuses on the scalability limits related only to signaling message handling and therefore examination of the data forwarding engine is out of the scope of this document.

6.3.1 Maximum Signaling Message Burst Size

Objective:
Determine the maximum signaling burst size, which is the number of the signaling messages in a signaling burst that the DUT is able to handle without signaling loss.

Procedure:
1. Select a signaling primitive or a signaling primitive pair and construct a signaling flow. The signaling messages should follow each other back-to-back in the flow and after "n" number of messages the flow should be terminated. In the first test sequence the number "n" should be set to one.

Additionally, all the signaling messages in the signaling flow must conform to the resource reservation protocol definition and must be parameterized in a way to avoid signaling message processing errors in the DUT.

2. Send the signaling flow to the DUT and count the signaling messages received by the terminator tester.

3. When the number of sent signaling messages ("n") equals to the number of received messages, then the number of messages forming the signaling flow ("n") should be increased by one; and the test sequence has to be repeated. However, if the receiver receives less signaling messages than the number of sent messages, it indicates that the DUT is beyond its scalability limit. The measured scalability limit for the maximum signaling message burst size is the length of the signaling flow in the previous test sequence ("n"-1).

In order to avoid transient test failures, the whole test must be repeated at least 30 times and the report should indicate the median of the measured maximum signaling message burst size values as the result of the test. Among the test runs, the DUT should be reset to its initial state.

There are signaling primitives, such as signaling messages indicating errors, which are not suitable for this kind of scalability tests. However, each signaling primitive suitable for the test should be investigated.

Reporting format:
The report should indicate the type of the signaling primitive or

signaling primitive pair and the determined maximum signaling message burst size.

Note:
In the case of routers supporting multicast resource reservation sessions, the signaling burst can be also constructed by sending signaling messages to multiple network interfaces of the DUT at the same time.

6.3.2 Maximum Signaling Load

Objective:
Determine the maximum signaling load, which is the maximum number of signaling messages within a time unit that the DUT is able to handle without signaling loss.

Procedure:
1. Select a signaling primitive or a signaling primitive pair and construct a signaling flow. The period of the signaling flow should be adjusted in a way that exactly "s" signaling messages arrive within one second. In the first test sequence the number "s" should be set to one (i.e. 1 message per second).

Additionally, all the signaling messages in the signaling flow must conform to the resource reservation protocol definition and must be parameterized in a way to avoid signaling message processing errors in the DUT.

2. Send the signaling flow to the DUT for at least one minute, and count the signaling messages received by the terminator tester.

3. When the number of sent signaling messages ("s" times the duration of the signaling flow) equals to the number of received messages, the signaling flow period should be decreased in a way that one more signaling message fits into a one second interval of the signaling flow ("s" should be increased by one). But, if the receiver receives less signaling messages than the number of sent messages, it indicates that the DUT is beyond its scalability limit. The measured scalability limit for the maximum signaling load is the number of signaling messages fitting into one second of the signaling flow in the previous test sequence ("s"-1).

In order to avoid transient test failures, the whole test must be repeated at least 30 times and the report should indicate the median of the measured maximum signaling load values as the result of the test. Among the test runs, the DUT should be reset to its initial state.

In the case of this test, there are also signaling primitives which are not suitable for this kind of scalability tests. However, each signaling primitive that is suitable for the test should be investigated just like in the case of the maximum signaling burst size test.

Reporting format:

The report should indicate the type of the signaling primitive or signaling primitive pair and the determined maximum signaling load value.

6.3.3 Maximum Session Load

Objective:
Determine the maximum session load, which is the maximum number of resource reservation sessions that can be maintained simultaneously in a reservation capable router. The maximum number of session relies on two architectural components of the DUT. First, the DUT should have enough memory space to store the attributes of the different resource reservation sessions. Second, the DUT has to be powerful enough to maintain all the reservation sessions if they require actions during the lifetime of the sessions.

In the case of hard-state protocols we cannot speak of reservation session maintenance, therefore in this situation the available memory space is the only limit for the session number. Moreover, there are also resource reservation protocols that handle only the aggregates of reservation sessions (e.g. Load Control [8]) and do not distinguish the separate traffic flows referring to reserved resources. Of course, in this situation there is no session maintenance either, since there are no reservation sessions, plus the memory allocation for the aggregates is limited. In this latter case, the maximum session load is defined to be unlimited and the test can be skipped.

According to the dual limits of the measurement, the benchmarking procedure is separated into two tests. The first test investigates the session number limit due to the memory space, while the second test explores the reservation session maintenance capability of the DUT.

The first test is applied to every resource reservation protocol, which stores reservation sessions separately and not only an aggregate of them. Resource reservation protocols that are capable for session aggregation, but still have the capability to handle separate sessions (e.g. Boomerang [3]) are still subject of this test.

Procedure:
1. Set up a reservation session in the reservation capable router by sending the appropriate signaling messages to the DUT.

2. Establish one more reservation session in the DUT using the appropriate signaling messages. In the case of soft-state protocols, all the reservation sessions existing in the DUT must be maintained using refresh messages.

3. Repeat step 2 until the router signs that there is not enough memory space to establish the new reservation session. In this case,

the test is finished and the maximum memory capacity available to store the sessions is reached.

Note:
Not all the resource reservation protocols support to signal the overrun of the maximum memory capacity limit directly. However, certain behavior of the router may also indicate the memory overrun.

The second test is applied to those resource reservation capable routers only that run reservation session maintenance mechanisms to refresh internal states belonging to reservation sessions. Here, we investigate the DUT whether it is able to cope with the refresh signaling message handling that shows also the capability to refresh the internally stored reservation sessions.

Procedure:
1. Set up "n" number of reservation session in the reservation capable router by sending the appropriate signaling messages to the DUT. In the first test sequence the number "n" should be set to one. Beside the reservation session generation, the initiator tester must also take care of the reservation session refreshes.

2. Capture the refresh signaling messages leaving the DUT for a specified amount of time ("T") while still maintaining the established reservations with refresh signaling messages. Time "T" must be at least as long as the protocol specifies as reservation time out.

3. Check whether each reservation session is refreshed during the refresh period that was examined in step 2. The proof of the session refresh is a leaving refresh signaling message referring to the corresponding reservation session. If all sessions that were set up in step 1 are refreshed during step 2, then repeat the test sequence by increasing the number of reservations by one ("n"+1). However, when any of the reservations was dropped by the DUT, then the test sequence should be cancelled and the determined maximum session load is the number of resource reservation sessions maintained successfully in the previous test sequence ("n"-1).

In order to avoid transient test failures, the whole test must be repeated at least 30 times and the report should indicate the median of the measured maximum signaling load values as the result of the test. Among the test runs, the DUT should be reset to its initial state.

Reporting format:
The report should indicate determined maximum session load value, which is the lowest value between the two test results.

Note:
When the number of reserved sessions grows over a number that counts to a very high value in the given technology conditions, then the test can be canceled and the report can state that the resource

reservation protocol implementation performs the maximum number of reservation sessions over that limit (e.g. "Over 100.000 sessions").

Also note, that testing the DUT in the case of multicast and unicast scenario, it may result different maximum session load values.

6.4 Benchmarking Tests

Benchmarking tests are defined to measure the QoS signaling related performance metrics on the resource reservation capable router device.

Since the objective of the benchmarking is to characterize routers performing resource reservation in real-life situations, therefore during the tests the DUT must not bump into its scalability limits determined by the previous test.

Each performance metric is measured when the DUT is under different router load conditions. The router load is generated and characterized using combinations of independent load types:

a. Signaling load
b. Session load
c. Premium traffic load
d. Best-effort traffic load

The initiator tester device generates the signaling load on the DUT by sending a signaling flow to the terminator tester. This signaling flow is constructed from a specific signaling primitive or a signaling primitive pair and has the appropriate period parameter.

The session load is generated by the signaling end-points setting up resource reservation sessions in the DUT via signaling. In the case of soft-state protocols, the initiator tester device must also maintain the reservation sessions with refresh signaling messages periodically.

The initiator tester device generates the premium traffic load by sending a data traffic flow to the terminator tester across the DUT. This traffic flow should have dedicated resourced in the DUT, set up previously using signaling messages. The traffic must consist of equally spaced and equally sized data packets. Although any transfer protocol is suitable for traffic generation, it is highly recommended to use UDP packets, since this data flow is totally controllable, unlike TCP that uses congestion avoidance mechanism. The premium traffic must be characterized by its traffic parameters: data packet size in octets, the calculated bandwidth of the stream in kbps unit and the transfer protocol type. The data packet size should include both the payload and the header of the IP packet.

The initiator tester device generates the best-effort traffic load by sending a data traffic flow (that refers to no resource reservation sessions) to the terminator tester across the DUT. Any other attributes of the traffic flow must meet the conditions described previously in the case of premium traffic load.

Note, that these four load types have influence on each other from their nature that may spoil the measurements. Therefore, in order to have accurate results these cross-effects must be minimized during the benchmarking tests. The signaling load can cause interference with the session load, when certain signaling messages alter the number of reservation session in the DUT. To cancel this influence the signaling flow should contain signaling message pairs, where the message pairs has opposite effect restoring the changes caused in the DUT. On the other hand, in the case of soft-state protocols, sessions must be refreshed by periodically sent signaling messages. Although refresh messages are used to maintain the reservation sessions, still they are counted as signaling messages. Furthermore, signaling messages are realized as data packets. Such way signaling messages must be taken into account in the traffic flow calculation as well.

6.4.1 Performing the Benchmarking Measurements

Objective:
The goal is to take measurements on the DUT running a resource reservation protocol implementation in the case of different load conditions. The load on the DUT is always the combination of the four load components described before.

Procedure:
The procedure is to load the router with each load component at a desired level and measure the investigated performance metrics. The load condition on the DUT should not change during the test. Once, the measurement is complete, repeat the test with different load distributions.

During the test sequences, in order to avoid transient flow behavior influencing the measurements, the measurements should begin after a delay of at least "T" time and after the setup of the common load on the DUT. The value of "T" depends on the parameters of the load components and the resource reservation protocol implementations, but, as a rule of thumb, it should be enough for at least 10 packets from the traffic flows and 10 signaling messages from the signaling flow to pass through the DUT and at least one refresh period to expire in the case of soft-state protocols.

During the measurement of the performance metrics in a practical load setup, not just one, but 100 measurement samples should be collected. Normally, the empirical distribution function of the tests is similar to the curve of a Gaussian distribution, and therefore the modus and the median are in the same location. Such case, the result of the test sequence is the median of the samples. In the case of different shaped empirical distribution functions, the curve must be further analyzed and the result should describe the curve well enough.

In order to avoid transient test run failures that may cause invalid results for the entire test, the whole test must be repeated at least 10 times and the report should indicate the median of the measured values filtering out the extreme results. Moreover, after each test run the DUT should be reset to its initial state.

In order to perform a complete benchmarking test, every performance metrics must be measured using signaling flows made of every applicable signaling primitives or primitive pairs.

Since the test methodology is the same for all the different performance metric benchmarking procedure, it is also recommended to perform the measurements for all performance metrics at the same time in one test cycle.

At first sight, this procedure may look easy to carry out, but in fact there are lots of difficulties to overcome. The following guidelines may help in reducing the complexity of creating a conforming measurement setup.

1. It is reasonable to define different amounts for each load component (load levels) before benchmarking and then measure the performance metrics with all possible combinations of these individual load levels.

2. The number of different load combinations depends on the number of different load levels defined for a load component. Working with too much number of load levels is very time-consuming and therefore not suggested. Instead, there are proposed levels and parameters for each load component.

The data traffic parameters for the traffic load components have to be selected from generally used traffic parameters. It is recommended to choose a packet size of: 54, 64, 128, 256, 1024, 1518, 2048 and 4472 bytes (these are the same values that are used in RFC 2544 that introduces methodology for benchmarking network interconnect devices). Additionally, the size of the packets should always remain below the MTU of the network segment. The packet rate is recommended to be one of 0, 10, 500, 1000 or 5000 packets/s. Since the number of combinations for these traffic parameters is still large, the highly recommended values are 64, 128 and 1024 bytes for the packet size and 10 and 1000 packets/s packet rate. These values adequately represent a wide range of traffic types common in today's Internet.

The number of session load levels should be at least 4 and it is recommended to share them equally between 0 and the maximum session load value.

The number of signaling load levels should be at least 4 as well, and the actual value of the signaling load is also recommended to be equally distributed between 0 and the maximum signaling load value.

Zero load level means that the actual load component is not involved in the router load.

Reporting format:
As the whole report description requires a four-dimension table (four load components plus the results), which is hard to visualize for a human being, therefore the results are extracted into ordinary two-dimensional tables. Each table has two fixed load component quantities and the other two load component levels are the row and column for the table. Such way, one set of such tables describe the benchmarking results for one certain type of signaling flow used in the generation of the signaling load. Naturally, each different signaling flow requires separate tables.

Note:
Of course in the case of multicast resource reservation sessions, the combination number of the different multicast scenarios multiplies the number benchmarking tests also.

7. Acknowledgement

The authors would like to thank the following individuals for their help in forming this document: Norbert Vegh and Anders Bergsten from Telia Research AB, Sweden, Krisztian Nemeth, Peter Vary, Balazs Szabo and Gabor Kovacs from High Speed Networks Laboratory of Budapest University of Technology and Economics.

8. References

[1] Y. Bernet, et. al., "A Framework For Integrated Services Operation Over Diffserv Networks", Internet Draft, work in progress, May 2000, <draft-ietf-issll-diffserv-rsvp-05.txt>

[2] B. Braden, Ed., et. al., "Resource Reservation Protocol (RSVP) - Version 1 Functional Specification", RFC 2205, September 1997.

[3] J. Bergkvist, I. Cselenyi, D. Ahlard, "Boomerang - A Simple Resource Reservation Framework for IP", Internet Draft, work in progress, November 2000, <draft-bergkvist-boomerang-framework-00.txt>

[4] P. Pan, H. Schulzrinne, "YESSIR: A Simple Reservation Mechanism for the Internet", Computer Communication Review, on-line version, volume 29, number 2, April 1999

[5] L. Delgrossi, L. Berger, "Internet Stream Protocol Version 2 (ST2) Protocol Specification - Version ST2+", RFC 1819, August 1995

[6] 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

[7] A. Eriksson, C. Gehrmann, "Robust and Secure Light-weight Resource Reservation for Unicast IP Traffic", International WS on QoS'98, IWQoS'98, May 18-20, 1998

[8] L. Westberg, Z. R. Turanyi, D. Partain, Load Control of Real-Time Traffic, A Two-bit Resource Allocation Scheme, Internet Draft, work in progress, April 2000, <draft-westberg-loadcntr-03.txt>

[9] G. Feher, I. Cselenyi, A. Korn, "Benchmarking Terminology for Routers Supporting Resource Reservation", Internet Draft, work in progress, July 2001, <draft-ietf-bmwg-benchres-term-00.txt>

[10] S. Bradner, J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999

[11] Boomerang Team, "Boomerang homepage - Benchmarking Tools", http://boomerang.ttt.bme.hu

9. Authors' Addresses:

Gabor Feher
Budapest University of Technology and Economics (BUTE)
Department of Telecommunications and Telematics
Pazmany Peter Setany 1/D, H-1117, Budapest,
Phone: +36 1 463-3110
Email: feher@ttt-atm.ttt.bme.hu

Istvan Cselenyi
Telia Research AB
Vitsandsgatan 9B
SE 12386, Farsta
SWEDEN,
Phone: +46 8 713-8173
Email: istvan.i.cselenyi@telia.se

Andras Korn
Budapest University of Technology and Economics (BUTE)
Institute of Mathematics, Department of Analysis
Egry Jozsef u. 2, H-1111 Budapest, Hungary
Phone: +36 1 463-2475
Email: korn@math.bme.hu