SNMP-driven Active Measurements in DiffServ Networks

G. Xilouris, G. Gardikis, K. Sarsembagieva, A. Kourtis
Institute of Informatics and Telecommunications
National Centre for Scientific Research "Demokritos"
Athens 153 10, Greece
{xilouris; gardikis; katias; kourtis}@iit.demokritos.gr

Abstract—**Active network monitoring techniques have been traditionally used in network management frameworks for OAM operations. They can also provide vital end-to-end network status information which can be exploited for application optimization (Network-Aware Applications). The paper proposes a novel, standards-compliant management framework for conducting active measurements on DiffServ networks, solely based on SNMP as control protocol. A specialized SNMP agent (SNMP for Active Measurements – SAM) is introduced for this purpose, accompanied with a custom MIB, which utilizes the standardized OWAMP (One-Way Active Measurement Protocol) for establishing active measurement sessions. A proof-of-concept implementation deployed on a laboratory testbed shows that the proposed framework exhibits satisfactory accuracy and scalability, while at the same time being compatible with off-the-shelf SNMP-based network managers. The SAM framework has been released as open-source for Linux-based routers.**

Keywords- active measurements, SNMP, OWAMP, DiffServ

I. INTRODUCTION

As Internet applications become more and more a part of our everyday life and also scale rapidly in terms of complexity and traffic volume, proper monitoring of the network infrastructure, access and core, wired and wireless, becomes imperative. From the network operators' point of view, network monitoring facilitates FCAPS (fault, configuration, accounting, performance, security) tasks, optimizes the end-users' experience, as perceived at application level, and secures customer fidelity.

There is also an additional aspect which has gained significant attention in the research community during the last years; the capability of applications to self-optimise in real time exploiting network monitoring data – what is referred to as Network-Aware Applications (NAA). Traditionally, network status has been implicitly inferred by measuring the end-to-end performance at application level (such as the TCP congestion control mechanism). While this seems enough for one-to-one, client-server communication, more contemporary communication paradigms, such as peer-to-peer transfers, distributed caching mechanisms, in-network clouds and information-centric networking, exhibit significant potential for self-optimisation if explicit network information status is provided. Such interplay is promoted by several research efforts on application/network coupling and also by initiatives in standardization bodies such as the Application-Layer Traffic Optimisation (ALTO) architecture [1], [2]. The advent of such paradigms increase the value of network monitoring data, which can now be exploited not only for network management purposes, but also for application optimization.

In this context, in addition to passive network metrics, the evaluation of end-to-end performance using probe traffic (active measurements) can be of added value. While several proprietary protocols and mechanisms exist for the conduction of in-network active measurements, it would be useful to have a common standardized architecture for the control of monitoring procedures –both passive and active- and the communication of the results.

Addressing this lack, in this paper we propose and implement a novel integrated SNMP-based network monitoring framework for DiffServ networks, able to conduct active measurements in a standardized manner exploiting enhanced agents (SNMP for Active Measurements – "SAM" agents). The proposed framework is totally compatible with existing SNMP-based network managers and thus can be integrated into the network without requiring additional components at the network management subsystem.

The rest of the paper is structured as follows: Section II provides background information and identifies the need for the proposed monitoring framework. Section III describes the functionality and architecture of the "SNMP for Active Measurements" (SAM) agent for DiffServ networks and Section IV presents the experimental testbed which was set-up for proof-of-concept validation, followed by evaluation results.Conclusions are discussed in Section V.

II. BACKGROUND

Most network monitoring tasks normally rely on passive monitoring, i.e. the observation of information and statistics which are reported by network elements and refer to user traffic. On the other hand, active monitoring involves the insertion of artificial (probe) traffic into the network, generated with specific pre-configured characteristics, such as bit rate, packet size, inter-packet time etc. Probe traffic is observed simultaneously at the sender and the receiver and reflects the end-to-end status of the traversed network path.

Active measurements can yield very useful results and are applicable even on core operational backbones [3], complementing passive measurements. They can be used to assess one- or two-way metrics for end-to-end paths, which cannot be derived via passive monitoring. Commonly collected end-to-end metrics are packet loss, one-way delay (OWD), jitter/delay variation and reordering. Especially, accurate measurement of one-way delay (OWD), as discussed in [4], requires tight synchronization between the sender and the receiver parts, which can be achieved –to a certain degree- via NTP queries to a low-stratum time server.

The active measurement task needs to follow a specific protocol with regard to i) the probe traffic format, ii) the establishment and control of the test session and iii) the communication of the results. Usually, each commercial network testing system follows its own proprietary protocol. Some of these mechanisms are discussed in the literature [5][6], also for MPLS networks [7] while others have been made publicly available [8]. Efforts to standardize this procedure within the IETF resulted in the One-Way Active Measurement Protocol (RFC 4656) [9]. OWAMP, which is also used by the SAM agent described in this paper, defines the establishment of an active measurement session between two peer parties based on time-stamped UDP traffic. It also specifies the eventual exchange of the measured quantities, such as one-way delay, loss and jitter. The main OWAMP session parameters are: packet size, number of packets in the probe traffic burst, inter-packet time and statistical distribution. OWAMP also allows setting the DSCP field in the IP header of the probe traffic packets, thus making it suitable for measuring DiffServ networks, a feature which is exploited in this paper.

While OWAMP defines the establishment of the session between the communicating peers and the exchange of the results between them, it does not specify the communication with a third management entity (e.g. a network manager). This task could be well undertaken by SNMP, which is already being used for network management purposes. Controlling in-network active measurements via SNMP, using an appropriate Management Information Base (MIB), would allow to integrate the OWAMP mechanism into an existing network management framework, without modifying the network manager itself. The approach we propose realizes the SNMP/OWAMP integration via an enhanced "SNMP for Active Measurements" (SAM) agent, accompanied with a tailored MIB.

The SAM agent can be installed either in layer-3 network elements (e.g. routers) or in dedicated measurement servers at the edge of the network. The advantage stemming from this approach is that both active and passive measurements can be integrated under the SNMP "umbrella", making it possible for a standard SNMP-compliant network manager to configure and collect both types of metrics.

The proposed framework is demonstrated and evaluated in a Diffserv network scenario, although it can be well applied to non-DiffServ networks also.

III. An Agent for SNMP-driven Active Measurements

In order to instruct over SNMP the establishment of an OWAMP session and also to retrieve the results, a dedicated Management Information Base (MIB) needs to be designed and implemented. There exist several MIB structures for DiffServ-enabled elements such as the one proposed by IETF [10], or several proprietary variants. However, they are all restricted to accommodating passive metrics only and no one supports the execution of active measurements. The proposed MIB ("SAM-MIB") adds this functionality, and may coexist with any of the aforementioned existing DiffServ-MIBs. The SAM-MIB, whose structure and objects shown in Table I, contains:

- Configuration parameters (characteristics of the probe traffic) such as peer IP address, packet size, DSCP field etc. An optional "VPath ID" corresponds to the network path which is measured in the case of a switched network. For instance VPath ID can correspond to a specific Label Switched Path (LSP) in the case of an MPLS setup.
- Measurement results, such as hops traversed, packet loss, jitter and OWD. The results are provided in two sets (Inbound and Outbound), corresponding to two unidirectional probe streams which are dispatched in opposite directions.

Since multiple OWAMP sessions can be run in parallel, the information contained in the MIB is organized in a table *(samActMsmtTable).* Each entry of the table corresponds to a set of configuration parameters and results for a single OWAMP session.

TABLE I. OBJECTS OF THE SAM-MIB

Object name	**Description**
samDaemonRunning	Shows whether the daemon should be active or not. (a value of 0 suspends the active measurement procedure)
samDaemonInterval	Time Interval between two consecutive measurements (msec)
samActMsmtTable	Active Measurements and Configurations Table
└samActMsmtEntry	Row of the samActMsmtTable
└samActMsmtConfVPathId	(Configuration) Unique index corresponding to the network path measuredin the case of switched networks (e.g. LSP in the case of MPLS)
└samActMsmtConfPeerIP	(Configuration) The IP Address (or Host Name) of the peer node with which the OWAMP session will be established
└samActMsmtConfDSCP	(Configuration) Differentiated Serviced Code Point to be set to probe packets for assessing DiffServ networks.
└samActMsmtConfNoPkts	(Configuration) Number of probe packets to be sent during the OWAMP session
└samActMsmtConfPktSize	(Configuration) Size of the packets to be sent during the OWAMP session
└samActMsmtConfInterPktTime	(Configuration) Time interval between two consecutive packets (msec)
└samActMsmtConfLossTimeout	(Configuration) Time threshold which, if exceeded, a packet will be considered lost (msec)

Object name	**Description**
└samActMsmtOutboundNoHops	(Result) Number of Hops measured in the Outbound direction(from the local node to the peer node)
└samActMsmtOutboundDuplicateCount	(Result) Number of Duplicate Packets measured in the Outbound direction
└samActMsmtOutboundPktJitter	(Result) Packet Jitter value obtained in the Outbound direction (msec)
└samActMsmtOutboundPktsLost	(Result) Number of packets lost in the Outbound direction
└samActMsmtOutboundMinDelay	(Result) Minimum one-way delay in the Outbound direction (msec)
└samActMsmtOutboundMaxDelay	(Result) Maximum delay value in the Outbound direction (msec)
└samActMsmtOutboundAvrgDelay	(Result) Average delay value in the Outbound direction (msec)
└samActMsmtOutboundStDevDelay	(Result) Standard Deviation delay value in the Outbound direction (msec)
└samActMsmtInboundNoHops	(Result) Same as OutboundNoHops, in the Inbound direction (i.e. from the peer node to the local node)
└samActMsmtInboundDuplicateCount	(Result) Same as OutboundDuplicateCount, in the Inbound direction
└samActMsmtInboundPktJitter	(Result) Same as OutboundPktJitter, in the Inbound direction
└samActMsmtInboundPktsLost	(Result) Same as OutboundPktsLost, in the Inbound direction
└samActMsmtInboundMinDelay	(Result) Same as OutboundMinDelay, in the Inbound direction
└samActMsmtInboundMaxDelay	(Result) Same as OutboundMaxDelay, in the Inbound direction
└samActMsmtInboundAvrgDelay	(Result) Same as OutboundAvrgDelay, in the Inbound direction
└samActMsmtInboundStDevDelay	(Result) Same as OutboundStDevDelay, in the Inbound direction
└samActMsmtLastUpdate	The timestamp (Date and Time) of the more recent measurement results available

The configuration objects of the MIB (*samActMsmtConf** objects) are set by the Network Manager via SNMP SET commands. After an OWAMP session has been configured in the MIB, the enhanced SNMP agent (SAM Agent) automatically schedules and executes OWAMP-based measurements to the remote node, using the configured parameters. These measurements are carried out continuously and periodically in the background. The time interval between two consecutive OWAMP measurements to the same peer node is configured via the *samDaemonInterval* object and is common for all sessions. At any time, the Network Manager may retrieve the most recent measurement results by executing SNMP GET queries to:

- the *samAcsMsmtOutbound** objects for the results of the outbound session (probe traffic dispatched to the peer node)
- the *samAcsMsmtInbound** objects for the results of the inbound session (probe traffic received from the peer node)

The SAM Agent is a composite module, developed in C on Linux Debian, also working on other Linux flavors. It consists of the following submodules:

a) a standard SNMP agent (linux net-snmp library), so as to respond to SNMP queries.

b) an OWAMP client, to initiate sessions with a remote peer. Code from the open-source RFC implementation [11] was reused.

c) an OWAMP Server, to listen and respond to active measurement requests initiated by remote peers.

d) a Monitoring Daemon, especially developed from scratch, which coordinates the aforementioned three sub-modules and ensures the communication among them. The Monitoring Daemon also utilises an NTP client so as to synchronise to an external NTP server, since time synchronisation is essential for one-way delay measurements and the proper operation of the OWAMP protocol in general.

The overall high-level sequence diagram showing the interactions among the SAM agent submodules is shown in Fig.1. The agent functionality can be briefly summarized in the following steps:

- First, the Monitoring Daemon uses the local SNMP agent to connect to the MIB and retrieve the general configuration parameters (measurement interval, running on/off flag).
- Following, the Monitoring Daemon retrieves from the MIB the configuration of each of the scheduled OWAMP sessions i.e. the entries of the *samActMsmtTable*, including the address of the peer node.
- For each entry, the Monitoring Daemon spawns a child thread, which invokes an instance of the OWAMP Client to connect to the OWAMP Server of the remote peer and to initiate a pair of active measurement sessions (Inbound and Outbound) with the configured parameters.
- After the session is over, the results are stored in the MIB via the SNMP agent. The procedure is repeated periodically and infinitely, until the configuration entry is deleted from the MIB table by the Network Manager.

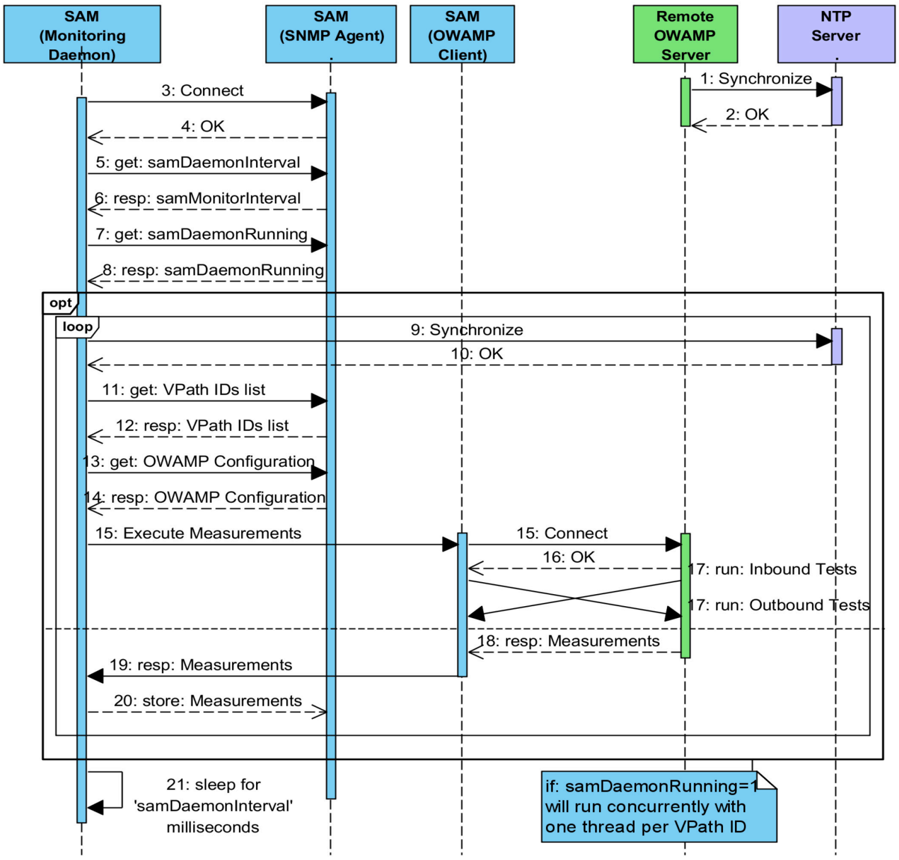

Figure 1. Sequence diagram of the active measurement procedure

The SAM agent and the SAM MIB have been released on open-source basis [12], with the corresponding documentation, for further experimentation on Linux-based network testbeds.

IV. VALIDATION TESTBED

The testbed which we set up for the validation and assessment of the SAM framework is depicted on Fig.2 and consists of:

- six Linux-based routers (R1-R6), based on Debian 6.0 (kernel version 3.2, iproute2 version 2.36). The routers form a DiffServ domain and are connected with 100Mbps Ethernet links. Four PHBs are defined: EF (using pfifo queuing discipline), AF11/AF12 (using GRED qdisc) and BE (using HTB/RED qdisc). The SAM agent was installed at the two edge elements (R1 and R6)
- a Network Manager, running the open source Multi Router Traffic Grapher (MRTG) SNMP manager [13], to which the SAM-MIB has been loaded. Alternative, any network manager can be used, provided it can load custom MIBs.
- an NTP Server, for providing time synchronisation, relying on a Stratum 1 GPS time source. It allows R1 and R6 to be in-sync with 0.01-msec-order accuracy.
- a commercial traffic generator platform (Traffic Source and Traffic Sink) to load the network, attached to its edges.

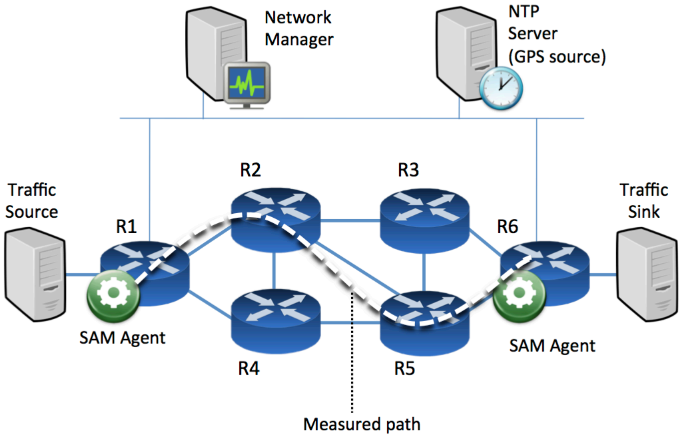

Figure 2. Sequence diagram of the active measurement procedure

R1 is configured, via SNMP queries, to execute three simultaneous OWAMP sessions to R6. Configuration parameters for the three sessions, as set in the SAM-MIB, are almost identical: 100 packets per session, 1000 bytes packet size, 2 seconds of loss timeout. Each session lasts approximately 13 seconds, using the default OWAMP inter-packet timing, and the inter-session interval is 10 sec, yielding an average bitrate of 3x35 kbps for the probe traffic. The only different element is the DSCP field (*samActMsmtConfDSCP*) of the probe packets; the first session is marked as EF, the second as AF11 and the third as BE, so as to assess the network status, as reflected to three different traffic classes.

We used the traffic generator to load the network with UDP CBR traffic, marked as EF, linearly sweeping from 0 to 100 Mbps. The MRTG monitor at the network manager periodically issues SNMP queries to the R1's SAM agent to retrieve the active measurement results, which are depicted in Fig.3, where the MRTG graph screenshots have been grouped in a single figure. (Please note the differences in y axis scales, which were automatically set by MRTG). A normal DiffServ end-to-end behaviour is depicted, and especially:

- for the BE session, the end-to-end OWD increases from almost 0 to 90 msec , while packet losses climb from 0 to 40%
- for the AF11 session, the end-to-end OWD increases from almost 0.1 to 0.6 msec, while packet losses climb from 0 to 5%
- for the EF session, losses and delay stay at the minimum (0.2 msec and 0%), apart from the moment where the 100 Mbps links are saturated, where losses raise to 1% and OWD to 0.4 msec.

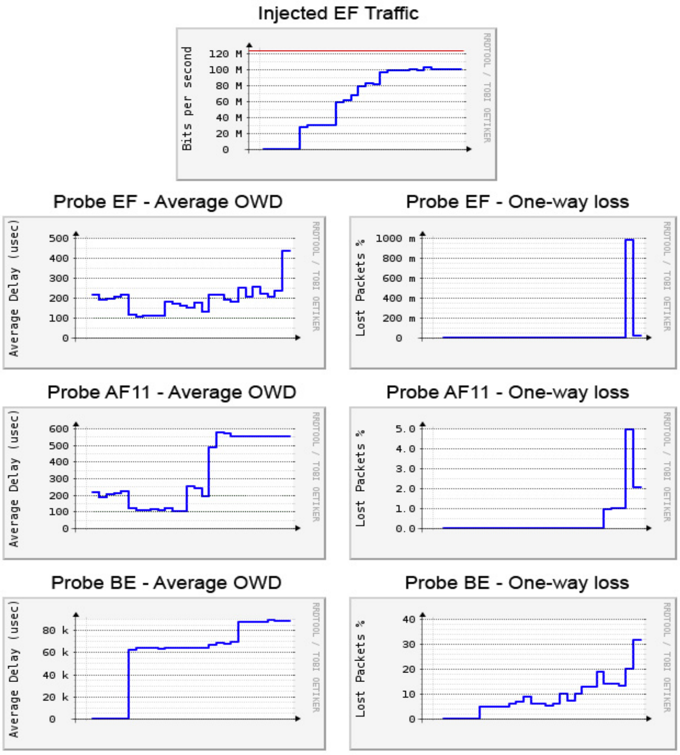

Figure 3. Screenshots of the MRTG grapher, depicting the results reported by the SAM Agent over SNMP

It is thus validated that the SAM framework can assess with satisfactory accuracy and granularity the end-to-end per-traffic-class performance of the DiffServ domain, using only 100 kbps of probe traffic, therefore introducing only 0.1% overhead to the overall network traffic.

V. CONCLUSIONS

In this paper, we described a novel framework for SNMP-driven active monitoring in DiffServ (and also non-DiffServ) domains, using open standards for the control and execution of the measurements and also for the communication of the results. The SAM framework can be integrated with any existing SNMP Network Manager to provide end-to-end, one-way performance metrics, which cannot be derived using passive measurement mechanisms.

Further research, as planned, involve the development of a “Network Cost” proxy for communication of the measured metrics to applications towards their self-optimisation (Network-Aware Applications). Also, a simulation-based study on a large-scale configuration will assess the scalability of the proposed method and determine the trade-off between measurement accuracy and overhead to the network.

ACKNOWLEDGMENT

The work described in this paper is partially supported by the research project FP7/ALICANTE, funded by the European Commission under grant agreement n° 248652.

REFERENCES

[1] J. Seedorf, E. Burger, “Application-Layer Traffic Optimization (ALTO) Problem Statement”, RFC 5693, September 2009, http://tools.ietf.org/html/rfc5693

[2] R. Alimi, R. Penno, Y. Yang, “ALTO Protocol”, Work in progress, March 2012, http://datatracker.ietf.org/doc/draft-ietf-alto-protocol

[3] L. Ciavattone, A. Morton and G. Ramachandran, “Standardized Active Measurements on a Tier 1 IP Backbone”, IEEE Communications Magazine, June 2003, pp. 90-96

[4] L. de Vito, S. Rapuano, L. Tomaciello, “One-Way Delay Measurement: State of the Art”, IEEE Trans. On Instrumentation and Measurement, 57(12), December 2008, pp. 2742-2750

[5] U. Hofmann, T. Pfeiffenberger, and B. Hechenleitner, “One-way-delay measurements with CM toolset,” in Proc. Int. Performance Comput.Commun. Conf., Phoenix, AZ, 2000, pp. 41–47.

[6] J. Jeong, S. Lee, Y. Kim, and Y. Choi, “Design and Implementation of One-Way IP Performance Measurement Tool”, vol. 2343. London, U.K.:Springer-Verlag, Jun. 2002, pp. 673–686. no. 2.

[7] R. Mishra, V. Sharma, "QoS routing in MPLS networks using active measurements", in Proc. TENCON 2003. Conference on Convergent Technologies for Asia-Pacific Region , vol.1, no., pp. 323- 327 Vol.1, 15-17 Oct. 2003, doi: 10.1109/TENCON.2003.1273338

[8] RIPE Test Traffic Measurement Service, http://www.ripe.net/data-tools/stats/ttm/test-traffic-measurement-service

[9] S. Shalunov et al, “A One-way Active Measurement Protocol”, RFC 4656, September 2006, http://tools.ietf.org/html/rfc4656

[10] F. Baker, K. Chan, A. Smith, “Management Information Base for the Differentiated Services Architecture”, RFC 3289, http://tools.ietf.org/html/rfc3289

[11] Internet2, OWAMP version 3.3. http://www.internet2.edu/performance/owamp/

[12] SNMP for Active Measurements (SAM) Framework, http://www.medianetlab.gr/opensource/

[13] Tobi Oetiker's MRTG - The Multi Router Traffic Grapher, http://oss.oetiker.ch/mrtg/