Internet Engineering Task Force
INTERNET-DRAFT
draft-svdberg-temon-00

February 2001
expires: August, 2001

Steven Van den Berghe
Pim Vanheuven
Piet Demeester
Ghent University/IMEC

Hamid Asgari
Thales Research Ltd

Some Issues for Desiging a Measurement Architecture for Traffic Engineered IP Networks

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

To view the list Internet-Draft Shadow Directories, see http://www.ietf.org/shadow.html.

Copyright Notice

1 Abstract

This memo describes some requirements that need to be taken into account when developing a measurement architecture for use in IP-traffic engineering. It looks at the methods for collecting measurement information and possible problems that might be encountered while doing so. It links the measurement requirements to the efforts made in other IETF-working groups and tries to state some guidelines for defining a formal framework to describe how measurements can be organised for traffic engineering purposes.

2 **Conventions** used in this document

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.

3 **Introduction**

This memo describes a set of requirements for constructing measurement architecture in a traffic engineered IP-network. It elaborates on the description of the measurement subsections of the framework document [4] as being written by the traffic engineering working group, where traffic monitoring is defined as: "the process of observing traffic characterisitics at a given point in a network and collecting the traffic information for analysis and further action". The measurement architecture we will be talking about, only contains the functionalities of organizing the low-level measurements and disseminating their results. The goal of this memo is to split this general definitition into more specific functional requirements. The scope we will be looking at is limited to intra-domain measurements. Furthermore, we will assume no input from link layer information. This memo doesn't describe which measurements should be performed. Furthermore the guidelines on how the measurement should be performed on the network is also outside the scope of this document. Within the functional description, a relation will be made with requirements, frameworks and definitions defined by other IETF working groups.

3.1 **Background:** Traffic Engineering Aware Measurements

To clarify the scope of the measurements we will first elaborate on the reasons why current measurement methodologies aren't sufficient in Traffic Engineered (TE) IP-networks. The technologies used within IP traffic engineering can have a wide range influences on the network behaviour. They can have an impact on the level of the which paths are used to forward packets through the network (e.g. MPLS-based TE), the way packets are distributed among those paths (multipath load balancing), the way packets are treated within the nodes along a path (e.g. DiffServ[2]), or any combination of these elements.
As a result traffic forwarded in the network might encounter a differentiation into several classes of service (CoS). As traffic belonging to each QoS has certain requirements and exhibit a certain behaviour, there is no longer a single metric result adequate for all packets belonging to a different CoS. Also, the measurement methodology

must be aware of these classes of service. This is taken into account by the IPPM-working group by parameterizing the metrics to be analysed according to so-called "type-P-packets", where type-P implies the CoS[3].

3.2 Granularity and Complexity of Measurements

The measurement architecture must not only be aware of the classes present in the network, but it can also be affected by this service differentiation. Not every CoS needs to be monitored in the same way. Different CoSs can have different QoS indicators that will require processing of different types of information, a different timing for the measurements, etc. Both the granularity and the size of the test intervals can be related to the CoS (e.g. a high priority class might need to be monitored at a higher frequency). For the granularity, an architecture could for instance provide three incremental levels: service availability, throughput, or full-metric level. The last one would then mean that a certain subset of delay, delay variation and loss are also being measured. Monitoring different CoS related traffic using different levels of granularity makes the measurement architecture far more complex, especially when it is performed in a traffic engineered environment with dynamic network configurations.

4 Terminology

The terminology used in this document conforms to the definitions in the "Framework for Internet Traffic Engineering" draft. Specific for their use in this document some additional definitions, describing two complementary methods of performing measurements, are used:

- Active measurements: this method uses the injection of test-traffic to perform its measurements. The behaviour of the packets in this teststream should conform to the behaviour of the CoS under test.

- Passive measurements: this method is to observe the traffic passing through a network element and performs basic measurements, based on counters built into those elements.

In both cases, the measurement results of these methods need to be retrieved for analysing the behaviour of the CoS under test.

5 Functional Description

In this section a possible functional description of the measurement role within traffic engineering is made based on an overview of some specific properties that are needed to be achieved.

5.1 Diagnostic versus Operational Measurements

When measurement functions are deployed on current networks, they mostly have a diagnostic role. They evaluate the current status of the network, or analyse the network behaviour during a certain time period, and report their findings to a management layer. When adding traffic engineering to the network, the algorithms used will also need an

overview of the network status (e.g. to perform constraint based routing). The measurement functionality that delivers this status can be designated as operational measurements. The operational aspects of measurements could also be used in other areas, e.g. restoration.

5.2 Basic versus Aggregated Results

When looking at the output needed from a measurement architecture, two levels of results can be identified. The first level gives a basic description of the network status in terms of the metrics defined by the IPPM working group. Within this memo, these metrics will be referred to as basic metrics. The second level uses this input during post-processing to perform an analysis of the status of the paths/network. This will result in actions like trend analysis, threshold checking and related alarm triggering. This second level of results will thus aggregate the basic results collected in the network in order to provide a more abstract meaning to other traffic engineering functions interested in measurement results, e.g. for triggering a modification of link provisioning.

5.3 Scope of the Measurements: End-to-End vs Hop-by-Hop

A very delicate point in the definition of a measurement architecture is the scope of measurements. Several different approaches can be taken here. The measurements can be either done on an "end-to-end" basis, thus performing measurements from ingress to egress node(s) in the domain or a part of the domain. On the other hand, the measurements can be made on a "per hop" basis, meaning that every node in the network performs measurements to determine the status on the links (and associated queues, schedulers, meters etc.) to its neigbours. The former method has as its major drawback the fact that in a multipath environment (when the traffic is dynamically assigned to multiple paths), the result of a measurement may not characterize the behaviour of all the traffic going from an ingress to an egress node, since a single measurement may not be able describe the behaviour of all paths between these two end-nodes nodes. Another issue introduced by this method might be the scaleability, as the ingress node needs to inject traffic (in the case of active measurements) conforming to the different CoSs under test, to all associated egress nodes. The latter one can overcome these difficulties by taking the multipath situation into account while concatenating the hop-by-hop results into an end-to-end status of a path in the network (and for instance take the worst case of the individual multipaths). When using this method, the status of every individual link is known, but a larger error will be introduced due to the concatenating (e.g. measurement faults, errors due to the unknown internal behaviour of a network element, etc.).

6 Monitoring Methodology

While the previous section has been focusing on some general functionalities needed, the next part will describe the more practical issues when developing a measurement architecture.

6.1 Active Measurements

When performing active measurements in order to determine the value of some metric, the IPPM-framework document prescribes that the points in time at which the measurements are performed, should be determined following some random distribution. The reason for this, lies in the ability to avoid synchronisation effects in the network. Now, in the context of traffic engineering, this methodology guideline could prove to introduce a new problem. When wanting to compare the measurement results, or perform calculation with them, constant measurement intervals might be needed. To be able to combine this with the prescription of the IPPM-framework, a discretisation of the measurement samples into fixed intervals could be performed. This action could be done by taking all the "random" measurements in such a fixed interval and aggregating them into a new result (e.g. the average of all measurements, or a min/max/avg triplet), and use these "fixed timing" resuts for future analysis.

6.2 Passive Measurements

The description of the metrics that are being determined using passive measurements at a network node, must be given as references to information bases as they are developed by other working groups within the IETF. This means that the measurement entity performing the passive measurements will read the counters determined by a reference to an entry of either a Management Information Base or a Policy Information Base, e.g. referencing to the teTunnelPackets entry in the TE MIB. Since the passive results are available at any given time, a measurement architecture should therefore try to perform the probing at the end of the constant interval introduced in the previous paragraph.

6.3 Abstract Information Tree

The following tree gives an overview on how the information stored in a measurement architecture can be organized.

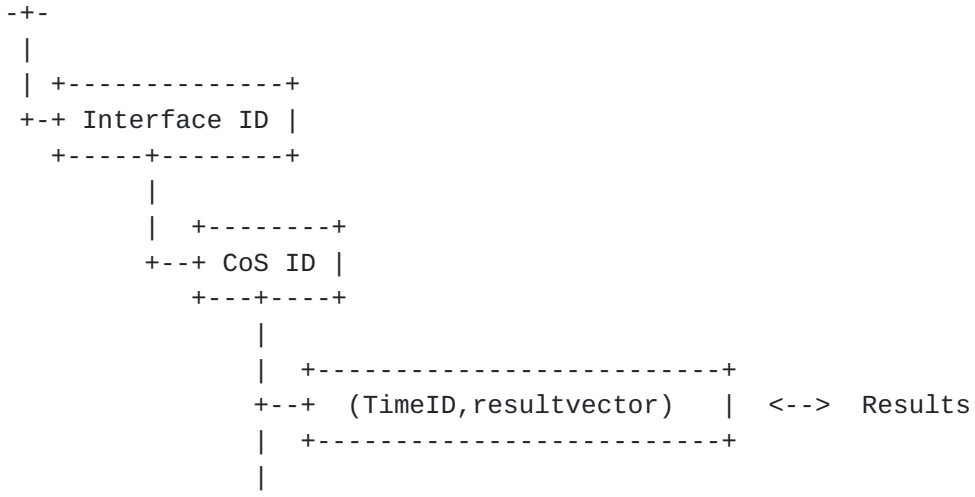

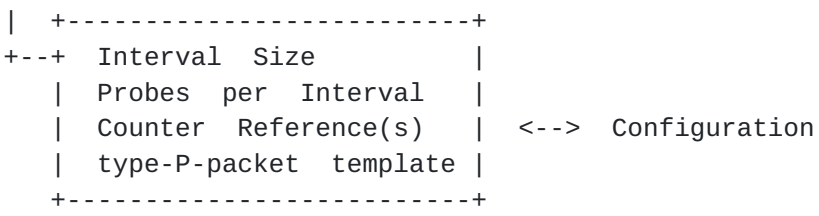
Interval Size
Probes per Interval
Counter Reference(s)
type-P-packet template
<--> Configuration

In this tree a per interface, per CoS approach is taken. Every measurement unit will then need two substructures: one to keep the results, and one to provide the configuration. The results should be stored together with a key identifying the time interval the vector of results is coming from. The configuration will need to determine (amongst other elements) the size of that interval, the average number of probes per interval, a reference to the counters that need to be read (e.g. MIB-references) and a template identifying how a type-P-packet should look like (e.g. in a diffserv environment this is only identified by the DiffServ Code Point).

6.4 Aggregating the Measurements

While describing the basic measurements that can be done, and can be performed by using the definitions proposed by the IPPM Working Group, in quite a formal way, doing the same for aggregated measurements might pose some more problems. Trend analysis, threshold alarms, etc. might need to be described formally, enabling a measurement architecture to co-operate with other IP-traffic engineering functionalities. This is also the issue for the correlation between operational and diagnostic monitoring, since the information delivered by the first function might need to be transformed into information needed by the second.

7 Conclusion

An measurement architecture and a formal description of its functions will need to be developed in order to provide network status information to other traffic engineering functionalities. While the major outlines of such an architecture might seem obvious, this document tries to point at some possible problems and to give some basic guidelines to be taken into account when developing such a measurement architecture. Several specific aspects remain however to be filled in.

8 Secuity Considerations

The solution developed to adress these requirements will also need to adress security and authentification issues, in order to ensure correctness and reliability of the measurements.

9 Acknowledgements

Part of this work has been funded under the European Commission 5th framework IST programme.
The authores would like to acknowledge all their colleugueas in the TEQUILA project for their input and reflection on this work.

10 References

[1] IST-Tequila project http://www.ist-tequila.org/
[2] RFC 2475, "An Architecture for Differentiated Services", S. Blake, D. Black, M.Carlson,E.Davies,Z.Wang,W.Weiss,
[3] RFC 2330, "Framework for IP Performance Metrics", V. Paxson, G. Almes, J. Mahdavi, M. Mathis
[4] draft-ietf-tewg-framework, "A Framework for Internet Traffic Engineering", D. Awduche, A. Chiu, A. Elwaldid, I. Widjaja, X. Xiao, July 2000

11 Authors' Adresses

Steven Van den Berghe,
Pim Vanheuven,
Piet Demeester,
St. Pietersnieuwsstraat 41
B-9000 Ghent, Belgium
Phone. ++32 9 267 35 86
E-mail:steven.vandenberghe@intec.rug.ac.be,
pim.vanheuven@intec.rug.ac.be,
piet.demeester@intec.rug.ac.be

Hamid Asgari
Thales Research Limited
Worton Drive, Worton Grange Industrial Est.
Reading RG2 0SB, UK
Phone:
Email: Hamid.Asgari@rrl.co.uk