Internet Draft
Expiration Date: December 2001

R.G. Cole
AT&T
R. Dietz
Hifn, Inc.
C. Kalbfleisch
Verio, Inc.
D. Romascanu
Avaya Inc.

A Framework for Synthetic Sources for Performance Monitoring

<draft-cole-sspm-03.txt>

Status of this Memo

1. Copyright Notice

2. Abstract

This memo discusses the use of synthetic sources (or 'active' probes) within the context of remote performance monitoring. It discusses the importance of developing an 'active' probe monitoring capability within the Internet. It develops a framework for synthetic sources in performance monitoring against the backdrop of previous and current, related work within the IETF. Specifically, the

relationship of this work to current activities in the RMON and IPPM working groups is discussed. It further reports on the broad agreements reached in the rperfman BOF held in Adelaide in March 2000 on furthering work in this area within the IETF. It is expected that this work will become part of the RMON WG Charter soon.

Distribution of this memo is unlimited.

3. Objectives and Motivation

3.1 Introduction

There is much utility in fully defining a performance monitoring capability within the IETF. As the Internet architecture becomes more complex, as enhanced QOS capabilities are defined and deployed, performance monitoring capabilities must be developed to account for this richer transport and service infrastructure. ISP's will be offering enhanced transport services, content hosting services will offer differentiated hosting services, and customers will demand methods to monitor the quality of the services to which they subscribe.

This memo defines a framework for the development of a synthetic source (or 'active' probe) capability for the purpose of enhancing remote performance monitoring capabilities within IP networks and services. By an 'active' probe, we mean a device or embedded software which generates a data packet (or packets) and injects it (them) into the network to a corresponding probe or existing server for the primary purpose of measuring some aspect of the performance of the end-to-end path or service. By performance monitoring we mean the act of collecting a specific set of measurements, either actively or passively, for the purpose of evaluating the quality of the path or the service. Much work within the IETF exists related to performance monitoring. One interesting aspect of this body of work is that it does not explicitly define an 'active' probe capability. An active probe capability is complimentary to existing capabilities, and should be developed by building as much as possible on this existing work.

3.2 History of This Document

This document was first published as an Internet draft to help motivate the rperfman BOF at the IETF meeting held in Adelaide in March 2000. At that time it was issued as <draft-cole-appm-00.txt>, dated March 2000 and was titled "A Framework for Active Probes for Performance Monitoring". Following the BOF in Adelaide, this second draft was issued under a new name "A Framework for Synthetic Sources for Performance Monitoring" to better reflect the nature of the

capability being proposed and to avoid confusion with other documents currently under development within the RMONMIB WG.

The major updates to this document include:

+ The second draft updates the first draft in several areas, including the results of the rperfman BOF, new developments within the RMON WG and an improved understanding of the capabilities and work items being suggested by this draft.

+ The third draft updates new work developing within the IPPM working group to define a protocol for one-way measurements [1]. During the development of the early drafts of the sspmmib [2], it became apparent that a one-way measurement protocol was required. This draft also incorporates a discussion of the SSPM MIB work documented in [2].

+ The fourth draft includes a discussion, found in Section 4.2, of an overall performance management architecture for application and transport level monitoring and traffic generation. This architecture intends to address both application level traffic generation and monitoring as well as the work within the IPPM WG on the development of a One Way Delay Protocol (OWDP).

3.3 Terms

This section defines the terms used throughout this memo.

+ 'Performance monitoring' is the act of monitoring traffic for the purpose of evaluating a statistic of a metric related to the performance of the system. A performance monitoring system is comprised of a) traffic generators, b) measurement, c) data reduction, and d) reporting. The traffic generators may be natural sources, synthetic sources or intrusive sources.

+ A 'probe' is a device or embedded software program that is placed in the data flow path or on a client or server to provide a performance monitoring function.

+ A 'synthetic source' is a device or an embedded software program which generates a data packet (or packets) and injects it (them) onto the path to a corresponding probe or existing server solely in support of a performance monitoring function. A synthetic source may talk intrusively to existing application servers.

+ 'Natural sources' are those that generate traffic to accomplish some unit of work and are measured passively by a measurement device or probe.

+ An 'intrusive source' is one that modifies an existing traffic flow in some manner.

+ An 'active probe' is a device or embedded software program that combines both synthetic source and probe functionality.

+ A 'passive probe' is a probe, which non-intrusively listens to packets flowing across the 'wire' or monitors request/responses on a client or server, and provides a performance monitoring function based upon its observations. Within the context of this discussion, it is synonymous with the term 'probe'.

+ A 'path' is a set of network transport components that provide a transport service between a given source and destination address pair. In its simplest form the network components are a series of routers interconnected by links. In more complex scenarios, a path has a more complex topology due to asymmetric routes, alternate paths, load balancing and redirection.

+ A 'service' is a collection of network components and servers designed to deliver a capability to an end user. The service could be a transport capability, a processing capability, etc.

+ 'Instrumentation' is the machinery required for the low-level programming of the probe's protocol interactions.

+ 'Instrumentation control' is the high level supervision of the probe's instrumentation, e.g., probe on/off, probe lifetime, etc.

+ A 'metric' is a carefully specified quantity related to some aspect of an Internet service [4].

+ A 'singleton metric' is a single measurement of a given metric.

+ A 'sample metric' is a set of measurements related by a common metric, traffic source(s) and measurement parameters, e.g., sample points, end points, path, etc.

+ A 'statistical metric' is a value derived by computing a specified statistical quantity on the sample metric. This accomplishes a reduction of the overall data.

+ 'Data reduction control' is the high-level supervision of the probe's (or distributed set of measurement points') statistical data reduction, e.g., the selection of a given statistic from a pre-specified set to perform data reduction.

3.4 Motivation

The bulk of the current development within the IETF is in the area of defining 'passive' monitoring, either self-monitoring as counters of local metrics or external-monitoring as defined within the RMON working group [5]. In contrast to passive monitoring is, what we refer to as, active monitoring. Active monitoring relies upon the injection of probe data or 'request' packets into the transport network or into a service. The active monitoring probe (or cooperating probes) then performs some type of measurement based upon the specific packet(s) it injects.

There are distinct advantages and disadvantages of both passive and active performance monitoring. These two approaches are very complimentary in nature. Passive probes are, by their very nature, non-intrusive; they add no additional load on the network or service. Passive monitors can provide a more extensive measurement capability (not only in the type of measurements but also in the amount of samples collected). Passive monitors do not, however, control the generation of data for the measurement samples. In contrast, active monitors are intrusive; they add load to the network or service. Because they control the generation of the packets, they also control the volume of traffic they introduce. In general, it is not expected that the objectives for generating active probes would necessitate high volumes of traffic.

Combined, these attributes limit the volume of measurements collected from active monitoring probes. However, this will allow for a richer set of historical data to be maintained in the probe due to the relatively low volume of measurement data (as compared, say, to an RMON probe sitting on a highly utilized fast ethernet LAN segment).

There are a number of reasons to develop an active probe capability for performance monitoring within the Internet. However, they all fundamentally boil down to the single issue of control. As discussed at length in the IPPM framework document [4], if you do not control the nature of the traffic generation, then you do not control the sampling and hence you do not control the quality of the respective statistics. It is important to control the timing of the packet generation to ensure the quality of the statistic (i.e., the random nature of the underlying sample). It is important to control the path of the test packets (at least the source and destination) to ensure that enough measurements are taken over the path in order to accurately identify the quality of the path. It is important to control the 'size' of the transactions to ensure that the measurements are relevant to the metric, e.g., throughput statistics should be based upon measurements with large files.

The utility of active probe capabilities will be found in:

+ troubleshooting paths - a pingMIB [6] identifies that connectivity exists but additional capabilities are required to determine the quality of the connectivity,

+ circuit pre-test and turn-up - prior to turning up a capability or customer, there is much value in monitoring the quality of their path or service prior to putting the customer on-line (without the capability of generating probe traffic this can be problematic),

+ fault management - allows determination of whether the application is operating or not,

+ base lining enhancements - active probes could be used to base-line BEFORE and measure AFTER a certain set of QoS or routing policies are applied. This would try to provide an answer to the question 'how effective is my proposed policy strategy?'.

+ capacity management - typical capacity management programs monitor local, utilization statistics to drive a capacity management decision, e.g., upgrade a facility, a CPU, etc. An active probe could be used to monitor complimentary aspects of network performance, more akin to an end-to-end metric, whose results could drive capacity management decisions as well. (This can be correlated to component level measures and can trigger specific capacity upgrades.),

+ Service Level Agreement (SLA) monitoring - because the nature of the probe packets used to measure a metric are tightly specified, the corresponding statistics will have significance within the context of an SLA.

In the next section we discuss issues of an architectural nature. We follow this with a section on related work, both previous and current, within various working groups at the IETF. Then, we present thoughts on Configuration Issues and Implementation Issues. Finally, the rperfman BOF [7] was held at the Adelaide IETF meeting in March 2000, which addressed the merits of the IETF specifying synthetic traffic sources for performance monitoring. The recommendations of that BOF are summarized at the end of this document along which proposed work items to follow up on this development.

4. Performance Management Architectural Considerations

In this section we first present some general considerations for the development of a synthetic source within the existing Internet

architecture. We then follow this up with a more specific proposal for the role and inter-relationship of various working drafts covering the overall performance management architecture.

4.1. General Architectural Considerations

There are several capabilities required which comprise a performance monitoring system. These include traffic generation, monitoring or measurement, data reduction and their respective control, as well as the various performance monitoring applications. Further, and as discussed throughout this document, there are various synchronization control functions necessary, e.g., clock synchronization between synthetic traffic source and sink or between synthetic traffic source and the metric monitoring functions. These are identified in Figure 1, along with an indication of their interrelationship.

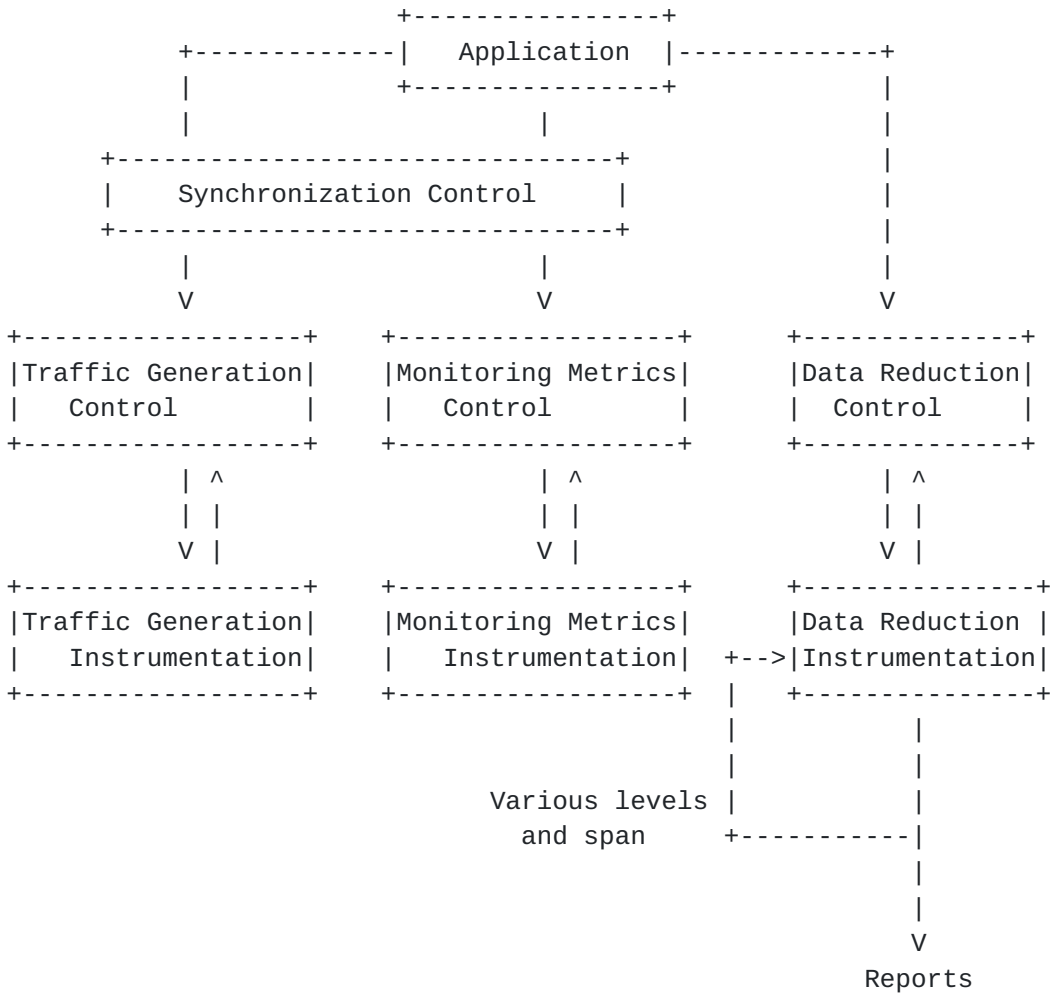

Figure 1: A performance monitoring system

Related to each defined transport or application service, we introduce the concept of a monitoring service, characterized by type of service, passive traffic generation method (if relevant), active traffic generation method (if relevant), metrics, monitoring and data reduction methods. In this context, a passive probe is an implementation of a passive monitoring method. An active probe is an implementation of an active traffic generation along with a passive monitoring method. Such an approach is currently being discussed within the context of a passive monitoring capability in the RMON working group. See, for example, [8] and [9].

One can expand upon this notion beyond performance monitoring. In fact, there are very few pieces of information that one might extract from a resource that are only useful for just one purpose, e.g., fault, policy or performance monitoring. For most of the attributes available today, the differences are in the use to which the information is put, not the data itself. It is only after we have defined higher-level objects (computed from existing ones) that we really have "performance data" or "fault data" or "policy data". Thus it should be possible to report basic fault information as well as gather performance statistics and policy baselines (see the discussion of base lining policies in Section 3.4 above). For instance, at a minimum the detected operational state should be reportable with a notification to indicate the transitions.

Given a monitoring service, a framework can be built that looks something like that shown in Figure 2.

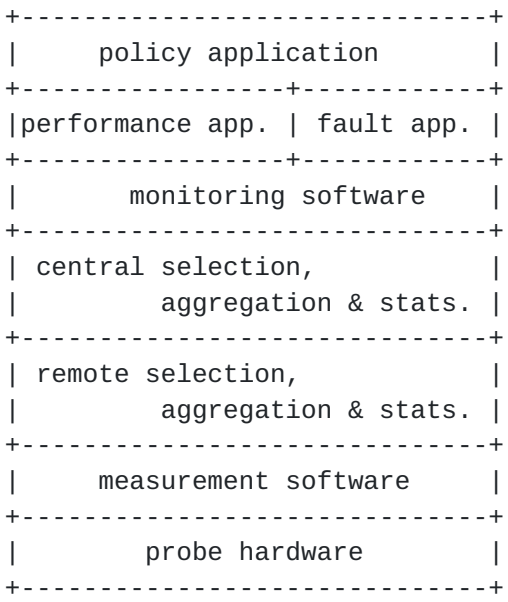

Figure 2: A framework for a monitoring service

In the context of performance, fault can be viewed as not performing at all and policy data can be obtained by comparing performance data from measurements of different networking scenarios. They should all be monitored with the same probes to reduce network traffic.

Much work within the IETF has addressed various of these capabilities (see the discussion in the section below on 'Related Work'). However, very little work within the IETF addresses the traffic generation capabilities for a monitoring service. In this section we focus on the traffic generation capabilities required for an overall performance monitoring system. Further, we discuss various architectural issues relating to the generation of 'synthetic traffic' for performance monitoring purposes.

There are various architectural considerations when discussing 'synthetic traffic sources' (or active probes) within the context of the Internet and it's standards. These include:

+ the target of the monitoring process, e.g., network transport versus server or process,

+ the 'layer' at which the probe functions, e.g., connectivity probes versus synthetic applications,

+ configuration - how to setup the behavior of the probe through R/W MIB objects for configuring the probe,

+ communication channels to remote probes,

+ the deployment architecture and its relationship to other monitoring methods, e.g., passive monitoring devices, and

+ security - related to probe control and generation.

We consider each of these issues in this section.

It is envisioned that specific probes/monitoring capabilities are to be developed specific to the service being monitored. When the target of the monitoring process is a transport service, then one naturally thinks of delay probes, loss probes, throughput and jitter probes, etc. When one thinks of database access services, one naturally thinks of various types of application request probes. We will talk of 'network' or 'connection' probes when monitoring transport services. We will speak of 'process-level', 'application-level or 'synthetic-application' probes when speaking of monitoring applications or a combination of transport and application services depending upon the location of the probes. It may even make sense to define an intermediate probe type, e.g., a 'session' probe, for the

purpose of monitoring some common aspects of the service and transport services.

Examples of 'connection' probes are delay, loss, delay variation, jitter, and throughput probes. Examples of 'synthetic-application' probes would be Oracle or SAP transaction probe or HTTP_get request probes, etc. Examples of 'session' probes might be DNS or DHCP probes, SIP probes for monitoring aspects of call setup delays, etc.

The configuration of an active probe ranges from full probe programming to a simpler 'control' of a synthetic traffic source. Full programming is viewed as providing too much flexibility to a remote application and hence is deemed a general security risk. The definition of a capability such as this was deemed dangerous and will not be addressed. Thus, we are left with the 'control' of an active traffic source from a remote application.

The active probes could be developed along the lines of the DISMAN's pingMIB [6], i.e., it is defined within the context of a MIB, directly accessible through SNMP and resident on a remote device. It could, instead be developed within the framework of the DISMAN's scriptMIB [10], where the active probe is an application which is distributed to the remote monitoring device and run on that remote device. Within this latter architecture, access to the probe's configuration, etc., may be through means other than SNMP and a MIB.

Depending upon the nature of the probes, some form of communication and control may be necessary between the communicating probes themselves (in addition to the probe packets). This is probably best addressed through SNMP communication to read/write MIB objects controlling the actions of the traffic source. The traffic stream generated by the synthetic source could be sent to a standard or well known destination port. In this case, the read/write MIB objects are required only to control the operation of the traffic source. However, for certain measurements or metrics, e.g., jitter metric, one way delay metric, etc., it may be necessary to invoke certain capabilities on the destination as well. This would require read/write MIB objects for the synthetic traffic generation destination as well as the source. This later case is the approach taken in the development of the sspmMib [2].

For metrics requiring multiple measurement points, e.g., a one-way delay metric requiring cooperation between a transmitter and a receiver (as discussed in the previous paragraph), a problem of time synchronization between the multiple measurement points exists. There are several possible solutions for this problem, some of them may be at the level of the application, others may result in requirements imposed on devices like support for a network time

protocol [11] or other clock synchronization methods.

Various deployment scenarios are feasible, depending upon the functionality desired and the allocation of that functionality across components. Clearly, active and passive probes can be implemented as either stand-alone devices that sit on the wire, or they can be implemented as embedded software within specific network elements or clients or server applications. An architecture can be envisioned which combines synthetic sources and passive probes, where the synthetic source is designed for the sole purpose of generating traffic at particular time points and the sample collection and statistical computations occur in already defined passive probes, e.g., RMON probes. This later case is the approach assumed in the current RMONMIB Working Group's drafts on performance monitoring, see [2], [8] and [9].

With respect to security considerations, past discussions related to active monitoring encountered a certain degree of pessimism, as did many other SNMP applications that involved configuration operations. However, the recent development of the SNMPv3 [12-16] security model, improved this situation, and we are witnessing the increased acceptance of SNMP as a 'trusted' and 'secure' protocol. This framework will analyze the issue of security and propose if necessary extra measures for ensuring a safe and secure utilization of the active monitoring capabilities.

Several security issues exists, including:

+ the security of the communication between a management application and the remote, synthetic traffic source - At a minimum, SNMPv3 authentication mechanisms should be considered for this aspect of configuration control. In some scenarios, it may be desirable to invoke the encryption capabilities within SNMPv3 as well. One specific concern wrt the ability of SNMPv3 to prevent replay attacks has been raised [3]. This issue should be addressed within the sspmMib work [2].

+ when using application level probes, we need to discuss the security of those applications - For instance, we may need to use secure protocols within the synthetic traffic streams. This raises the issue that an active probe should actually support the security protocols at the highest level of the devices in the network, and maybe share the secrets specific to the application. Active and passive probes may need to share secrets. This adds another dimension to the already complex problem of monitoring secure protocols. This is an example where SNMPv3 encryption is necessary to prevent snooping of control data containing shared, application-level, secrets.

+ there is the potential that the probes for monitoring will be perceived as security violations - e.g., port scans.

+ the nature of the communications between the active probes themselves - In the event that the control of both the synthetic source and destination is required, there are several ways to accomplish this level of coordination. The coordination could be left within the jurisdiction of the management application, in which case SNMP v3 security mechanisms may be invoked. Alternately, this level of coordination may be left to the source/destination probes themselves, in which case some secure communications protocol is required. As an example of this later situation, the OWDP work ongoing in the IPPM WG is developing a OWDP-control protocol with associated security capabilities built into the control protocol [1].

+ spoofing results - potentially disrupting communications, and

+ using the active probes in denial of service attacks. For example, using replay attacks to configure multiple probes, as previously mentioned.

4.2. A Proposed Performance Management Architecture

Here we present some thoughts on a proposed Performance Management Architecture for the IETF. The proposal builds upon current ongoing work in various existing working groups within the IETF; most notably the RMONMIB, the IPPM and the DISMAN Working Groups. The proposal references several existing drafts in various states of maturity within the above working groups. The current drafts we reference are:

+ The Application Performance Monitoring MIB (APM MIB) [8], which defines a method for identifying and reporting application level performance metrics. This is being defined within the RMONMIB WG.

+ The One Way Delay Protocol (OWDP) [1], which defines a method for controling and measuring various one-way metrics. This is being defined within the IPPM WG.

+ The Synthetic Source for Performance Monitoring MIB (SSPM MIB) [2], which defines a method to control the remote generation of measurement traffic for performance monitoring purposes. This work is to be defined within the RMON MIB WG.

+ The Transport Performance Monitoring MIB (TPM MIB) [9], which defines a method for identifying, measuring and reporting transport level metrics. This work is currently being defined

within the RMONMIB WG, but it's immediate future is uncertain.

+ Various documents from the IPPM WG which define transport metrics, e.g., [17-24].

Using these drafts as a foundation we propose the following Performance Management Architecture. Noter, there exists holes in this architecture if one strictly reads the drafts and attributes their current state of development to the below architecture. We list the gaps at the end of this section. The proposed architecture makes the following assumptions:

+ All application-level metrics are 'transactional' in nature and hence can be monitored at a single point within the traffic stream.

+ Transport level metrics are either transactional and one-way and hence the architecture must incorporate both types.

+ Monitors can (and often will) be replicated along the measurement path in order to attempt isolation of the end-to-end performance down to sub-section specific measurements.

+ It is highly desirable to rely on existing network management standards for the control and collection of data within the Performance Management Architecture. I.e., there is no need to re-invent secure management protocols.

We begin with the presentation of the Performance Management Architecture for One-Way Measurements. In a sense, this is the more complicated of the situations to consider. Figure 3 diagrams a situation where a network management application is setting up a one-way measurement test and monitoring. The network management application sits at the top of the diagram and controls the traffic generation through the SSPM MIB and the traffic monitoring through the TPM MIB (or its reincarnation, see the RMONMIB WG meeting minutes at the 50th IETF [37]). The OWDP-test function generates the traffic and runs the test protocol between the source on the left and the sink on the right.

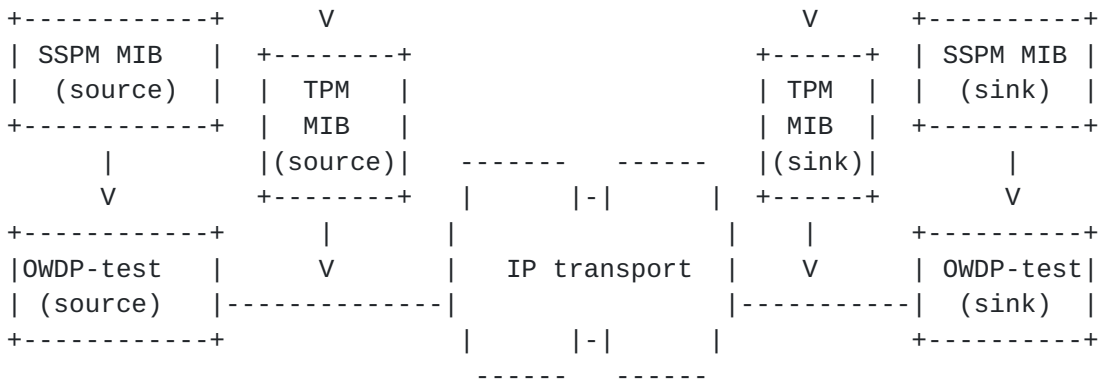

Figure 3: An Architecture of One-Way Performance Monitoring

The following functions are suggested by this architecture.

The SSPM MIB functions include:

+ Source control - test scheduling, end-point configuration.

+ Sink control - test scheduling, end-point configuration.

+ Interface to network management application through SNMP.

The OWDP functions include:

+ Handshake - the OWDP-test handles the initial handshake, i.e., the "Start Sessions"/"Control ACK" message exchange to start the actual test traffic flow.

+ Packet Generation - the OWDP-test would run the test measurement protocol, e.g., packet creation (sequence numbering, time stamping, etc) and packet injection handling.

+ Protocol exchange termination - the OWDP-test would terminate the protocol excahnge at the completion of the test measurements, i.e., the "Stop Sessions"/"Control ACK" message exchange to terminate the test traffic flow.

The TPM MIB functions include:

+ Measurement collection - the collection and storage of the raw measurement results, e.g., a History Table.

+ Statistical data aggregation - the temporal aggregation of local data, e.g., a Reports Table, with aggregation according to IPPM

referenced documents.

+ Metric definition - the TPM MIB would provide references to clearly defined metric reference to ensure unambiguous interpretation of results.

The associated control required to setup a test within this architecture is divided up into "Traffic Generation Control" and "Monitoring and Reports Control". Specifically, we envision the following steps to establish a test and data collection measurement:

TRAFFIC GENERATION CONTROL

+ Network management application builds the SSPM source and sink Control Table entries on the traffic source and the traffic sink Then the OWDP requires:

- Source and destination IP addresses

- UDP source and destination port numbers,

- Packet rate and pattern information,

- Total packets to be sent,

- TOS field values.

+ SSPM schedules OWDP-test:

- OWDP-test sends the OWDP Session-Start handshake,

- OWDP-test sends measurement packets,

- OWDP-test receiver collects packets,

- OWDP-test terminates test with OWDP Stop-Session handshake.

+ SSPM ages out Control Tables.

MONITORING AND REPORTS CONTROL:

+ The network management application builds the TPM Report Control table entries on two monitoring points, which may or may not be coincident with the traffic source and sink.

- TPM Control now specifies, e.g., "IPPM-one-way-delay" metric and associated "IPPM-one-way-delay" statistics

- Report presents the statistics, and time stamp accuracy information.

+ Network management application may build a TPM History Control Table entry.

- History Table contains the raw measurement data,

- OWDP specifies the following information be collected and stored: sequence numbers, send time (or presumed time if lost), received time (or zero if lost).

+ Network management application collects the statistical report from the Reports Table and/or raw measurement data from History Table.

We now cover the Performance Management Architecture for Round-Trip Measurements. In a sense, this is the simplier of the situations to consider. Figure 4 diagrams a situation where a network management application is setting up a round-trip measurement test and monitoring. The network management application sits at the top of the diagram and controls the traffic generation through the SSPM MIB and the traffic monitoring through the TPM MIB (or its reincarnation, see the RMONMIB WG meeting minutes at the 50th IETF [37]) and the APM MIB. The SSPM MIB controls the generation of traffic, running the application between the source on the left, e.g., the client, and the server on the right. By way of an example, we use a Web-based client/server application and indicate this in Figure 4 showing an HTTP client on the left and an HTTP server on the right.

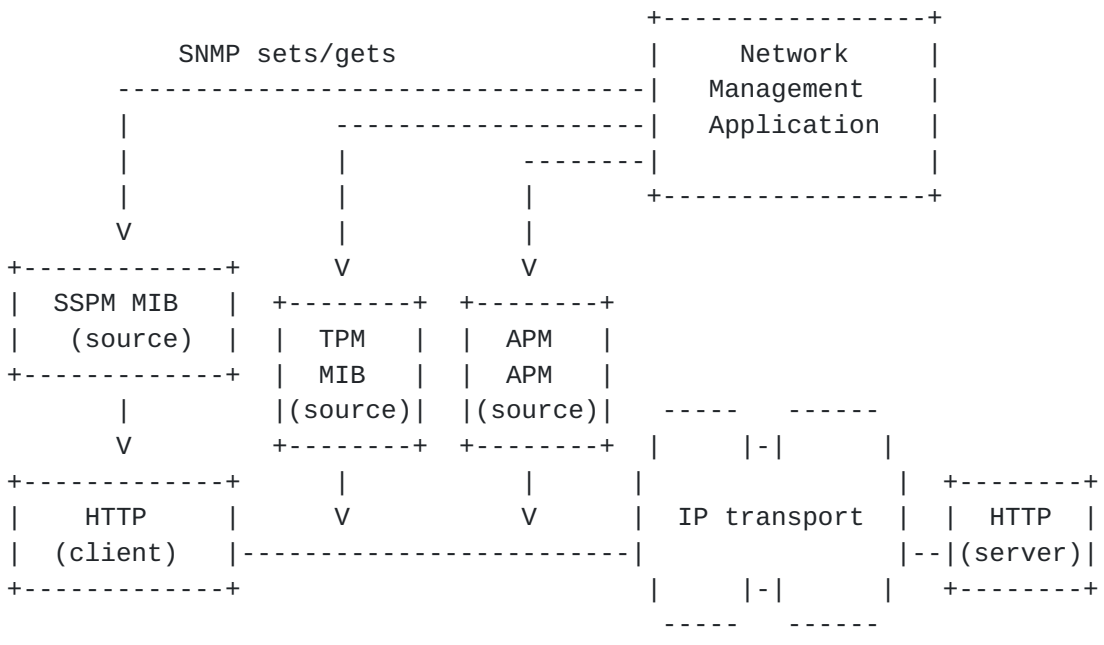

Figure 4: An Architecture of Round Trip Performance Monitoring

The following functions are suggested by this architecture for the round trip measurements.

The SSPM MIB functions include:

+ Source/Sink control - common platform, test scheduling, end-point configuration.

+ Configuration - may include the source/destination IP addresses, HTTP header information, TOS bit settings, timeouts, etc.

+ Single "interface" to network management application through SNMP.

The HTTP Client functions include:

+ Builds the DNS request to resolve the hostname to an IP address

+ Establishes a TCP connection to the IP address on the specified port

+ Build HTTP Get request packets

+ Issue the request

+ Parse HTML response for embedded objects

+ (potentially) establishes more TCP connections

+ Issue requests for unique embbed objects, etc.

The TPM MIB functions include:

+ Measurement collection - the collection and storage of the raw measurement results, e.g., a History Table.

+ Statistical data aggregation - the temporal aggregation of local data, e.g., a Reports Table, with aggregation according to IPPM referenced documents, e.g., pointers to IPPM standards and associated statistics such as the ippm-round trip-delay average, distribution, variance, etc.

+ Sub-transaction level data - the collection and reporting on data on the individual sub-transactions that comprise the total

application-level transaction, e.g., DNS, TCP and HTTP sub-transaction level information within a Web browser application.

+ Metric definition - the TPM MIB would provide references to clearly defined metric reference to ensure unambiguous interpretation of results, e.g., pointers to IPPM standards and associated statistics such as the ippm-round trip-delay metric.

The APM functions include:

+ Availability and responsiveness reporting - the end-user experience is captured within the context of an availability and a responsiness metric as discussed within the APM MIB draft, and

+ Aggregation of reporting information - the APM MIB provides various types of statistical data aggregation and sample statistics.

The associated control required to setup a test within this architecture is similar in spirit to that discussed above for the one way delay measurements. Hence we will not discuss these again.

There are several issues associated with this high level architectural discussion. They are:

+ The current plan for the development of the OWDP protocol includes it's own, unique controlarchitecture. Further, it is not clear that the appropriate separation between the control part and the test part of the protocol exists. For example, see the discussion of the one way delay measurement control flow above and compare this to the functional allocation of the OWPD into a test portion and a control portion.

+ The current TPM MIB development work is going away. This would leave a hole in the overall architecture. For example, refer above to the discussion of the TPM functions in the one way and round trip measurements.

+ There needs to be more clarity in the role of the APM MIB and the initially proposed TPM MIB functionality. We suspect that the TPM should include access to raw measurements and a breakdown of the APM aggregated data into subtransaction level data and error code information, e.g., timeouts, codes, etc.

5. Relationship to Other Work

Much work has already occurred within the IETF which has a direct bearing on the development of active performance probe definitions. This body of work is addressed in various working groups over the

years. In this section we focus our attention to the work of a) the IPPM working group, b) the DISMAN working group, c) the RMON working group, d) the ApplMIB working group, and e) the RTFM working group.

5.1 IPPM

The IPPM working group has defined in detail a set of performance metrics, sampling techniques and associated statistics for transport-level, or connectivity-level, measurements. The IPPM framework document [4] discusses numerous issues around sampling techniques, clock accuracy, resolution and skew, wire time versus host time, error analysis, etc. Much of these are considerations for Configuration and Implementation Issues discussed below. The IPPM working group has defined several metrics and their associated statistics, including

+ a connectivity metric [17]

+ one-way delay metric [18]

+ one-way loss metric [19]

+ round trip delay and loss metrics [20]

+ delay variation metric [21]

+ a streaming media metric [22]

+ a throughput metric [23] and [24], and

+ others are under development.

These (or a subset) could form the basis for a set of active, connectivity-level, probe types designed for the purpose of monitoring the quality of transport services. A consideration of some of these metrics may form a set of work activities and a set of early deliverables out of a group developing an active probe capability.

During the early development of the sspmmib drafts [2], it became apparent that a one-way measurement protocol was required in order for the ssmpMib to control. This helped led to the current work withi the IPPM WG on the development of the One-Way Measurement Protocol (OWDP) [1]. This protocol work includes both the measurement protocol itself, as well as the development of a seperate control protocol. This later control protocol is rendundant with the current work on the ssmpMib, so it appears that the IPPM WG will seperate their protocol into two seperate drafts, one for the

measurement protocol and one for the control protocol. But this remains to be finally agreed to in the working group.

5.2 DISMAN

The DISMAN working group is defining a set of 'active' tools for remote management. Of relevance to this draft are:

+ the pingMIB [6],

+ the DNS Lookup MIB [6],

+ the traceroute MIB [6],

+ the scriptsMIB [10], and

+ the expressionMIB [25].

The pingMIB and traceroute MIB define an active probe capability, primarily for the remote determination of path and path connectivity. There are some performance related metrics collected from the pingMIB and one could conceivably use these measurements for the evaluation of a limited set of performance statistics. But there is a fundamental difference in determining connectivity versus determining the quality of that connectivity. However, in the context of performance monitoring, a fault can be viewed as not performing at all. Therefore, they should both be monitored with the same probes to reduce network traffic. This was discussed further in the Architecture section above.

The DNS Lookup MIB also includes some probe-like capabilities and performance time measurements for the DNS lookup. This could be used to suggest details of a related session-level, active probe.

Also mentioned in the Architecture section above, the scriptsMIB allows a network management application to distribute and manage scripts to remote devices. Conceivably, these scripts could be designed to run a set of active probe monitors on remote devices.

5.3 RMON

The RMON working group has developed a extensive, passive monitoring capability defined in [5], [26], ... Initially, the monitors collected statistics at the MAC layer, but has now been extended to high-layer statistics. Higher-layer statistics are identified through the definition of a Protocol Directory [5]. The working group is recently re-chartered and is now concentrating on, among other items, monitoring at the application level.

The minutes of the Boston interim meeting in January 2000 are a good source for information about these ongoing activities in the RMON WG [27]. A number of individual drafts exist which discuss a number of interesting areas such as:

+ application typing and relevant metrics [8] and [28]

+ transaction level statistics collection and reporting [9] and [28]

Within this context (and discussed within the Architecture Section above), the development of an active traffic source for performance monitoring fits well within the overall performance monitoring architecture being defined within the RMON WG.

Indeed, based upon the agreements from the rperfman BOF, it appears that the development of the ssmpMib will occur within the RMONMIB WG (see the discussion of the rperfman BOF below).

5.4 ApplMIB

The ApplMIB working group defined a series of MIBs which monitor various aspects of applications, processes and services.

The System Application MIB [29] describes a basic set of managed objects for fault, configuration and performance management of applications from a systems perspective. More specifically, the managed objects it defines are restricted to information that can be determined from the system itself and which does not require special instrumentation within the applications to make the information available.

The Application MIB [30] complements the System Application MIB, providing for the management of applications' common attributes which could not typically be observed without the cooperation of the software being managed. There are attributes which provide information on application and communication performance.

The WWW MIB [31] describes a set of objects for managing networked services in the Internet Community, particularly World Wide Web (WWW) services. Performance attributes are available for the information about each WWW service, each type of request, each type of response and top accessed documents.

In the development of synthetic application-level probes, consideration should be given to the relationship of the application MIBs to the measurements being performed through a synthetic application-level probe. Similar, cross-indexing issues arise within

the context of the RMON monitoring and synthetic application-level active probes.

5.5 SNMPCONF

The snmpconf working group will create a Best Current Practices document [32] which outlines the most effective methods for using the SNMP Framework to accomplish configuration management. The scope of the work will include recommendations for device specific as well as network-wide (Policy) configuration. The group is also chartered to write any MIB modules necessary to facilitate configuration management, specifically they will write a MIB module which describes a network entities capabilities and capacities which can be used by management entities making policy decisions at a network level or device specific level.

Currently the snmpconf working group is focused on the SNMP Configuration MIB for policy [33]. For synthetic probes there is need to have configuration of a) a single probe, b) several probes, c) source and destination probes and d) intermediate probes. In addition, it may be necessary to configure any or all of these combinations simultaneously. It is hoped that the work of snmpconf will suffice. The scripting language defined by the SNMP Configuration MIB could allow for active monitoring to be activated and configured from a policy management script. Further, the results of active monitoring could become arguments in further policy decisions. This notion is reflected in the decision flow outlined in Figure 5 below.

5.6 RTFM

The Realtime Traffic Flow Measurement (RTFM) working group is concerned with issues relating to traffic flow measurements, usage reporting for network traffic and Internet accounting. Various documents exist which describe requirements [34], traffic flow measurement architectures [35], and a traffic flow MIB [36]. The work in this group is focused on passive measurements of user traffic. As such, its work is related to the monitoring work within the RMON WG. Fundamentally, their attention has not been concerned with methods of active traffic generation.

5.7 Relationship to Other Work: Summary

In summary, the development of an active traffic generation capability primarily for the purpose of performance monitoring should draw upon various activities, both past and present within the IETF. Redrawing Figure 1 in Figure 5, but now with annotations to the various work activities briefly touched upon in this section, is a

means to position the development of a traffic generation capability within the larger context of a performance monitoring system.

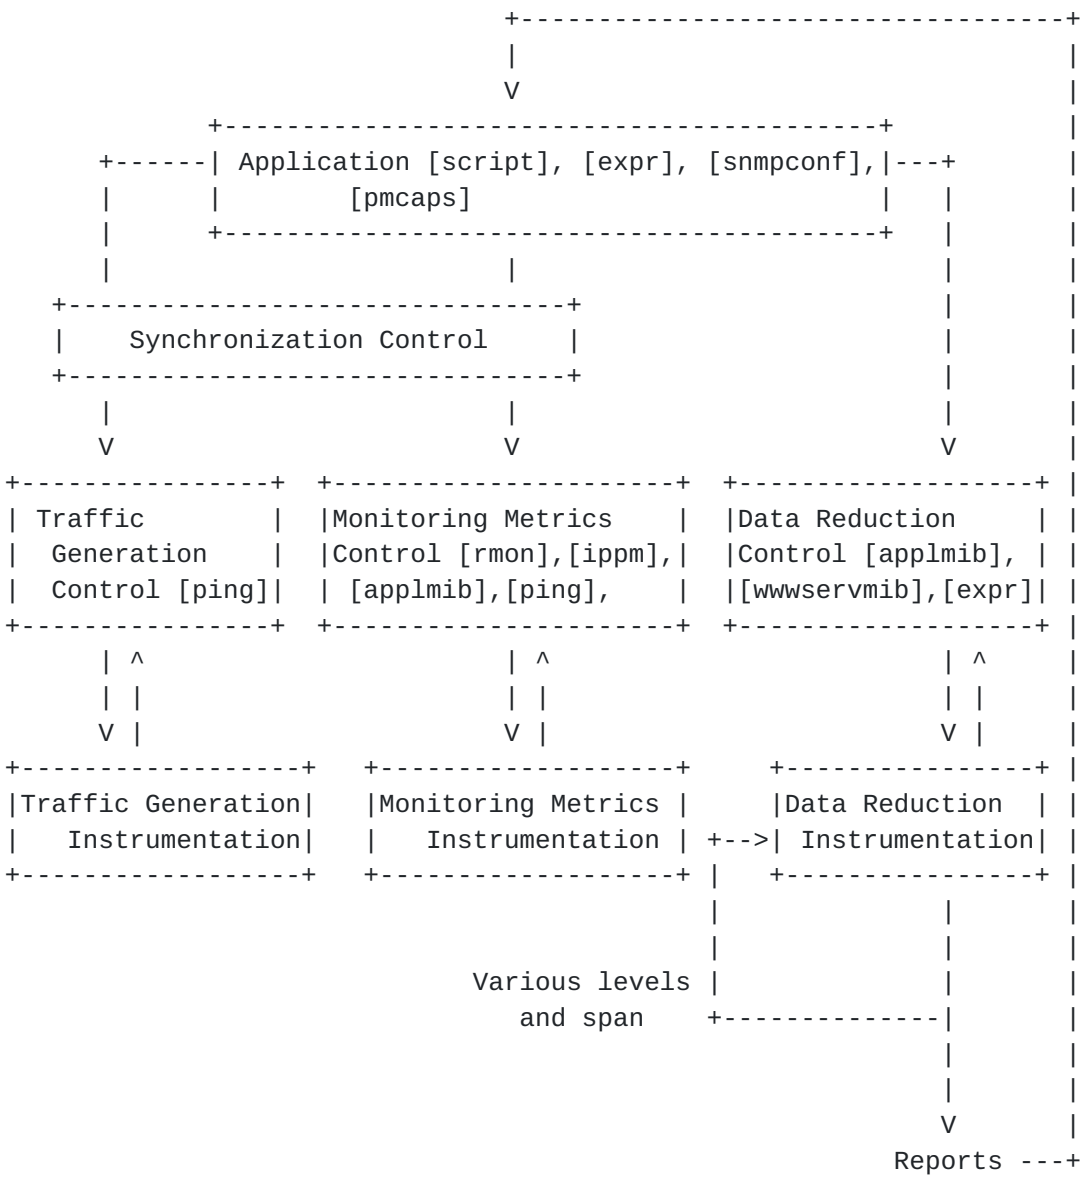

Figure 5: Coverage for an overall performance monitoring system

6. Configuration Issues

It is primarily assumed within this memo that the configuration of the probes is accessible through a MIB and communications to the remote probe is via SNMP. Other options, exist; one such option was briefly discussed above in the Architecture section.

The remainder of this section focuses on various configuration issues surrounding the definition and development of an active traffic generation capability. Here we discuss a) sampling methodologies, b) useful probe configuration options, c) statistics, reporting and historical data, and d) correlation of results to other measurements.

6.1 Sampling

Controlling the generation of traffic has numerous advantages as discussed above in the Motivation section. However, in the context of performance monitoring, a key advantage is being able to control the sampling. As discussed within the various IPPM documents, especially within the IPPM Framework document [4], it is critical to the quality of the statistical metric to be able to control the sampling. In particular, a performance monitoring application should be able to control the beginning and end of a sampling period, as well as the frequency and nature of the sampling within that period. The lifetime of the test may be finite or infinite, i.e., the test has an on/off switch settable by a management application. The frequency range should be carefully considered. The frequency may be tied to the type of the test probe, e.g., it may be fine for ping to have a 1 second retry, but for higher level applications we may not want to allow 1 second retries. Desirable sampling methods would include, at a minimum, both deterministic, i.e., generating probe traffic at fixed intervals, and Poisson, i.e., generating probes with exponentially distributed inter-arrival times.

6.2 Probe Configurations

The configuration of the specific probes can be quite extensive, given all of the potential options. The options would cover areas such as:

+ static, read-only information related to the implementation of the active probes and their capabilities,

+ timing and frequency of the probe packets (see Sampling section above),

+ data configuration (protocol selection, payload size, data fill, etc),

+ protocol options (could include multiple layers of protocol processing),

+ source and sink probe configuration in the case that the active probes are for the purpose of activating one-way measurements,

+ path configuration options (source and destination addresses, TOS field settings, do not fragment settings, ifNumber, TTL, source route, etc.), and

+ link level, quality of service type parameter settings, e.g., priority bit settings, loss priority bit settings, etc.

6.3 Statistics, Data Reduction, Reports and Historical Data

This section covers the statistics computed locally, the nature of the reports generated, and the storage of historical data. Reference [9] has a good discussion of a general set of statistics to maintain in probes, the complexities involved and the utility of the various statistics. Also, the work of the IPPM working group and their specific documents discusses or recommends statistics related to the metrics they define.

As discussed in the Architecture section above, traffic generation and performance measurements are separate functions within an overall performance monitoring service. Further, other work is in progress which addresses the measurement, data reduction and reporting of performance monitoring results, specifically [8] and [9]. Therefore, we concern ourselves here with those aspects of measurement and data reduction which may, in some sense, be unique to an overall performance monitoring service which is relying upon active traffic generation. Specifically, because we are controlling the nature and rate of the sampling, it is reasonable to expect that the measurement system will be capable of maintaining (maybe in an exception condition) the full historical data from active probe test periods. In general, measurement systems will perform some level of data reduction to minimize the data storage burden. However, this burden can be tightly controlled within a performance management service relying on active traffic generation.

6.4 Indexing to Other Measurements

There will potentially be a great deal of performance related information collected across numerous MIBs. The definition of a set of active probes only adds to this data. Methods are available within subsets of this data to cross-correlate results through standard indexing tables. Various MIBs from the Appl working group, i.e., [29], [30], and [21], are related through a service instance identifier. To quote [31],

"The WWW Service MIB interfaces to the Application MIB [30] by using the service instance identifier {applSrvIndex} for wwwServiceIndex if an applicable instance of applSrvIndex is

available."

The discussion and early drafts from the RMON working group, i.e., [8] and [9], discuss the relationship between the metrics of application-level and transport-level measurements and their cross-indexing. To quote [9], ...

"This document is intended to create a general framework for the collection and reporting of performance related metrics on traffic flows in a network. The MIB in this document is directly linked to the current RMON-2 MIB and uses the Protocol Directory as a key component in reporting the layering involved in the traffic flows."

The definition of active probes and their related statistics should be defined in such a way that useful cross-correlation of results is possible.

This type of correlation is currently possible for certain definitions of "service" in [30]. For instance in Section 6.1 of [30] indicates that for long lived services like http and smtp there would be instances in the service-level tables. For finger there may not be an entry. From here we can determine the reference points back to system application MIB and determine all of the information about the application.

Clearly, it would be desirable to be able to correlate, e.g., the results of a synthetic application probe running on a remote device into an application server with the measurements found within the applMIB for that same application running on that server. To take this example further, then to correlate the applications-level probe's measurements to transport-level measurements and even to the individual component level. This would require the ability to relate the path of the probes to the specific components, which may be complicated due to asymmetries in routing, load balancing across paths and servers, etc.

7. Implementation Issues

Implementation of active probes and their corresponding measurements is a tricky business, as discussed in detail in the body of the IPPM WG documents, in particular references [4] and [18]. In this section we reinforce some of the discussion in these references in the area of measurement accuracy, etc. Specifically, we discuss a) requirements on implementations, b) error analysis statements, and c) compliance tests.

7.1 Requirements on Implementations

There are a number of areas where implementation capabilities can affect the quality of the statistical metrics. These include, but are not limited to, items such as clock resolution, and skew, types of packet injection process supported, upper and lower bounds on packet generation rates, etc. Although not obvious at the time of this writing, it may be desirable to define a set of requirements on implementations of synthetic traffic generation devices. We suspect, however, that a better approach is to have an statement from the vendors of the various components of an overall performance monitoring service presenting an error analysis of their products and their respective output. This is discussed in the following section.

7.2 Error Analysis Statements

Performance measurements, whether they are based on active or passive monitoring, are error prone. It may make sense to define an error analysis statement/methodology so that implementations can clearly define their source of errors and hence the accuracy of their results. There is a fair amount of discussion within the IPPM framework document [4] surrounding this issue, which should be drawn upon extensively.

7.3 Compliance Tests and Statements

Implementations often surprise their implementers. For this reason it may be useful to define a compliance test covering the nature of the traffic generation, as well as the measurement system within an overall performance monitoring service. This would most likely be an activity separate from the definition of a traffic generation MIB and related monitoring MIBs.

Further, a statement of the types of synthetic probes supported is necessary.

8. Next Steps

There are several steps to move this work forward. A BOF was held in Adelaide to discuss this area of work as a potential basis for a working group at the IETF. The discussions during this BOF are documented in a set of meeting notes [7]. The broad agreements reached during the BOF were succinctly stated by Randy Presuhn in a mail message to the disman mailing list on 30 March 2000:

> "The rperfman BOF met for one session in Adelaide on Thursday, March 30, 2000. We covered all the items on the agenda and reached broad agreement that the following disposition of the work would make sense:

1) work on the control of active probes appears to belong in the rmonmib working group. It may be helpful to limit the scope of such work to the high-level control/supervision of such probes, rather than getting involved in the low-level programming of their protocol interactions. The rmonmib WG chair will give this topic due consideration in planning future activities.

2) While probe-level data summarization belongs in rmonmib, the control of the summarization of information from multiple systems is better pursued in disman. The reporting of the summarized information should be consistent with the techniques being developed in rmon where practical. The disman WG chair will raise this issue in the disman WG as a topic of possible future work.

3) It is believed that snmpconf work will provide adequate means to support the coordination of probe and data summarization function configuration. Those working on this topic will provide feedback into the snmpconf work.

With all of the topic areas either handled by existing WG activities or by the above proposed disposition, we agreed that there is no need for a new working group nor for a follow-up BOF at this time."

Within this context, we believe that the following work is appropriate:

+ Further develop this framework/architecture document defining the architecture of an active performance monitoring capability, its tradeoffs relative to other potential architectures, and its relationship to other, already defined monitoring capabilities. Roughly, the idea is that the synthetic sources' capabilities are listed in ssmpMIB [2]. Then this MIB would expand those entries with N entries for instances of the actual synthetic sources. Reporting is proposed through the apmMIB [8] and the tpmMIB [9].

+ (possibly) Develop a separate security document,

+ Develop a MIB for active probes and another for a usage of that MIB for some specific network or application layer synthetic sources. This work has begun and is documented in the ssmpMIB draft [2].

9. Intellectual Property

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication 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 Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

10. Security Considerations

This needs a very close examination, probably more than usual. Some security issues are briefly mentioned in the Architecture section above, but the issue of security was one of the reasons for this work being deferred in the past. It may be necessary to create a special document that deals specifically with the security issues related to the development of active, traffic generation MIBs.

11. List of Outstanding Issues

In writing this document several issues are uncovered that are yet to be resolved. This section summarizes this list of outstanding issues in one place. The intent is to resolve all of these issues prior to finalizing this document.

The list of outstanding issues is as follows:

+ So far the MIB object discussion has focused around the source of traffic generation. There is also a need to configure a destination/reflector. Probably we should have separate R/W objects in the MIB for source and destination configuration. We would then need to be able to co-ordinate the configuration of these two devices. Need some more discussion about how this might work. One item to consider is what attributes are needed for one way delay and jitter measurements?

+ One proposal is to rely solely on the reporting capabilities

within the apmMIB [8] and the tpmMIB [9]. However, it may not be prudent to limit performance monitoring to only the data in the apmMIB and tpmMIB. For example, there is a need to consider how reporting of say 100 one-way delay measurements would happen. This type of historical data is not currently available through the apmMIB or the tpmMIB. This area of performance monitoring is still up for discussion.

+ How to coordinate with lower level protocol parameters, e.g., link level QOS parameters such as the 802.1 priority levels? Could we consider a way to specify link layer information generically rather than through specific attributes? Or should we develop specific tables for specific link layers?

+ It is currently planned that the capabilities of the synthetic sources are listed through the sspmMIB [2]. Then, the apmMIB [8] and tmpMIB [9] would monitor the traffic for performance monitoring purposes. Within this context, there is a need to consider indexing to handle the situation where multiple managers are configuring the synthetic sources.

+ The current OWDP work within the IPPM WG needs to allow for better integration with the current work in the RMONMIB WG.

12. Acknowledgements

The authors gratefully acknowledge the contributions and discussions they have had with Randy Presuhn of BMC Software, Inc.

13. References:

[1] Shalunov, S., Teitelbaum, B. and M. Zekauskas, "A One-Way Delay Protocol for IP Performance Measurements", <draft-ietf-ippm-owdp-02.txt>, December 2000.

[2] Kalbfleisch, C., Cole, R.G. and D. Romascanu, "A Synthetic Source for Performance Monitoring MIB", <draft-kalbfleisch-sspmmib-02.txt>, June 2001.

[3] Private communications with S. Bellovin, December 2000.

[4] Paxson, V., Almes, G., Mahdavi, J. and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

[5] Waldbusser, S., "Remote Network Monitoring Management Information Base Version 2 using SMIv2", RFC 2021, January 1997.

[6] White, K., "Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations", RFC 2925, September 2000.

[7] Bierman, A., "Minutes of rperfman BOF in Adelaide", in an email message to the disman and rmon WGs' mailing list on 11 April 2000 from R. Presuhn.

[8] Waldbusser, S., "Application performance measurement MIB", <draft-ietf-rmonmib-apm-mib-00.txt>, May 2000.

[9] Dietz, R. "Application Performance Measurement Framework Transport Performance Metrics MIB", Internet Draft, <draft-ietf-rmonmib-tpm-mib-00.txt>, May 2000.

[10] Levi, D. and J. Schoenwaelder, "Definitions of Managed Objects for the Delegation of Management Scripts", RFC 2592, May 1999.

[11] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation and Analysis", RFC 1305, March 1992.

[12] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing SNMP Management Frameworks", RFC 2271, Cabletron Systems, Inc., BMC Software, Inc., IBM T. J. Watson Research, January 1998

[13] Case, J., Harrington D., Presuhn R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", RFC 2272, SNMP Research, Inc., Cabletron Systems, Inc., BMC Software, Inc., IBM T. J. Watson Research, January 1998.

[14] Blumenthal, U., and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", RFC 2274, IBM T. J. Watson Research, January 1998.

[15] Levi, D., Meyer, P., and B. Stewart, "SNMPv3 Applications", RFC 2273, SNMP Research, Inc., Secure Computing Corporation, Cisco Systems, January 1998

[16] Wijnen, B., Presuhn, R., and K. McCloghrie, "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)", RFC 2275, IBM T. J. Watson Research, BMC Software, Inc., Cisco Systems, Inc., January 1998

[17] Mahdavi, J. and V. Paxson, "IPPM metrics for Measuring Connectivity", RFC 2678, September 1999.

[18] Almes, G., Kalidindi, S. and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[19] Almes, G., Kalidindi, S. and M. Zekauskas, "A One-Way Packet Loss Metric for IPPM", Internet Draft, <draft-ietf-ippm-loss-07.txt>, May 1999.

[20] Almes, G., Kalidindi, S. and M. Zekauskas, "A Round-Trip Delay Metric for IPPM", RFC 2681, September 1999.

[21] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IPPM", Internet Draft, <draft-ietf-ippm-ipdv-06.txt., October 1999.

[22] Raisanen, V. and G. Grotefeld, "Network Performance Measurement for Periodic Streams", Internet Draft, <draft-ietf-ippm-npmps-00.txt>, March 2000.

[23] Mathis, M. and M. Allman, "Empirical Bulk Transfer Capacity", Internet Draft, <draft-ietf-ippm-btc-framework-02.txt>, Octobet 1999.

[24] Mathis, M., "TReno Bulk transfer Capacity", Internet Draft, <draft-ietf-ippm-treno-btc-03.txt>, February 1999.

[25] Stewart, B. and R. Kavasseri, "Distributed Management Expression MIB", RFC 2982, October 2000.

[26] Waldbusser, S., "Remote Network Monitoring Management Information Base", RFC 1757, February 1995.

[27] Meeting minutes from the interim meeting of the RMON working group on January 11 and 12, 2000 in Boston, MA.

[28] Warth, A. and J. McQuaid, "Application Response Time (ART) MIB", Internet Draft, <draft-warth-art-mib-01.txt>, October 1999.

[29] Krupczak, C. and J. Saperia, "Definitions of System-Level Managed Objects for Applications", RFC 2287, February 1998.

[30] Kalbfleisch, C., Krupczak, C., Presuhn, R. and J. Saperia, "Application Management MIB", RFC 2564, May 1999.

[31] Hazewinkel, H., Kalbfleisch, C., and J. Schoenwaelder, "Definitions of Managed Objects for WWW Services", RFC 2594, May 1999.

[32] MacFadden, M., and J. Saperia, "Configuring Networks and Devices with SNMP", Internet Draft, ,draft-ietf-snmpconf-bcp-01.txt., May 2000.

[33] Waldbusser, S., Saperia, J., and T. Hongal, "Policy Based

Management MIB", Internet Draft, <draft-ietf-snmpconf-pm-01.txt>, May 2000.

[34] Mills, C., Hirsch, G., and Ruth, G. "Internet Accounting Background", RFC 1272, November 1991.

[35] Browlee, N., Mills, C. and Ruth, G. "Traffic Flow Measurement: Architecture", RFC 2063, January 1997.

[36] Brownlee, N. "Traffic Flow Measurement: Meter MIB", RFC 2064, January 1997.

[37] Bierman, A. "Minutes of the RMONMIB WG at the 50th IETF meeting in Minneapolis", www.ietf.org, March 2001.

14. Author Information

Robert G. Cole
AT&T Laboratories
Network Design and Performance Analysis Department
330 Saint John Street, 2nd Floor
Havre de Grace, MD 21078

Phone: +1 410-939-8732
Fax: +1 410-939-8732
Email: rgcole@att.com

Russell Dietz
Hifn, Inc.
750 University Ave
Los Gatos, CA 95032-7695

Phone: + 1 408-399-3623
Fax: + 1 408-399-3501
Email: rsdietz@hifn.com

Carl W. Kalbfleisch
Verio, Inc.
1950 Stemmons Freeway
Suite 2026
Dallas, TX 75207

Phone: + 1 214-853-7339
Fax: +1 214-744-0742
Email: cwk@verio.net

Dan Romascanu
Avaya Inc.
Atidim Technology Park, bldg. #3
Tel Aviv, 61131
Israel

Phone: +972-3-645-8414
Fax: +972-3-648-7146
Email: dromasca@avaya.com

A. Full Copyright Statement