

Network Traffic Measurement: a closer look

Hiroshi Tsunoda[†] Glenn Mansfield Keeni[‡]

1. Introduction

Though network traffic measurement is widely deployed, the quality and nature of the statistics obtained have not been closely scrutinized. In this paper, we take a closer look at the measurement practices widely deployed and discuss the inaccuracies. Latency and its variation degrade the accuracy of network measurements. We point out that the inaccuracy essentially has its origins in the timestamp attribute of a measured value. Timestamp is a significant attribute but in current practices the definition and/or usage is inherently imprecise.

2. Network traffic measurement

Network traffic measurement is conducted for understanding overall volume and peak of the traffic. It has significant implications in accounting, operations, security and quality of service management of networks.

In this work, we discuss the inaccuracy of traffic measurement in large distributed network environments illustrated in Figure 1. End users' networks are connected to a provider's access router via customer premises equipment (CPEs). For network management purposes, a manager is located on a network management station in the provider network, and agents are deployed on the CPEs. An agent on a CPE maintains counters of the various facets of network traffic such as the number of packets, bytes, and/or errors etc. These counters are in general cumulative.




Figure 1 Assumed environment

A manager polls the agent periodically using some management protocol. It sends a request for the value of a counter. The agent samples the counter and sends back its value to the manager in response to the request. Then the manager computes the delta of the two samples and thereby computes the bandwidth utilization for the interval between the two samples.

Figure 2 illustrates how the bandwidth utilization is computed. v_t denotes the value of the cumulative traffic counter at the agent at time t . From two samples $v_{t_{i-1}}$ and v_{t_i} , the bandwidth utilization (Bw_i) between t_{i-1} and t_i is calculated as

$$Bw_i = \frac{\Delta v_i}{\Delta t_i} = \frac{v_{t_i} - v_{t_{i-1}}}{t_i - t_{i-1}}. \quad (1)$$

Note that, in periodic polling, a manager tries to keep Δt_i constant and equal to Δt , the specified polling interval.




Figure 2 How to calculate bandwidth utilization

In the request-response mode of management, there are latencies in the manager, intermediate networks, and the agent. Latency in the manager and the agent include the delay for processing the contents of request and response packets. Latency in intermediate networks includes transmission and propagation delays. In this work we focus on the cases where the latency and its variation are significant.

3. Inaccuracies in network traffic measurement

For properly understanding network dynamics, Δv_i and the corresponding Δt_i in Eq. (1) must be accurately measured with appropriate granularity. In real world situations this turns out to be difficult.

3.1 Inaccuracy of data time-stamp

In the ideal case, v_{t_i} is the value of the counter at time t_i . However, in reality, in the absence of explicit time-tags, the manager cannot know the exact value of t_i and uses t_i^A which is an approximation of t_i .

Figure 3 illustrates two successive polls, $(i-1)$ -th and i -th. According to this figure, v_{t_i} should be the value at time t_i^A , the time at which the value is actually sampled. Thus, t_i^A is a data time-stamp and t_i should be t_i^A .

However, a manager, in general, does not have the means of knowing the exact value of t_i^A unless data itself has explicit time-tags. Instead of t_i^A , a manager will generally use t_i^{SRq} (t_i^{RRS}), the time when the manager sent the request to (received the response from) the agent, as an approximation.

However, the request/response latency, d_i^{Rq} and d_i^{Rs} , between the manager and the agent, and, the request processing time at the agent (d_i^{PrC}) are all variables depending on the network conditions and processing load at the agent. Hence, even if the manager adjusts t_i^{SRq} and t_{i-1}^{SRq} so that the polling interval $\Delta t_i' = t_i^{SRq} - t_{i-1}^{SRq}$ becomes constant, data interval $\Delta t_i = t_i^A - t_{i-1}^A$ will vary.

[†] Tohoku Institute of Technology, [‡] Cyber Solutions Inc.






Figure 3 Sequence diagram of a polling process

Figure 4 and Figure 5 show the variations of polling interval and estimated data interval, respectively. These data were obtained from an experiment using simple network management protocol (SNMP) [1] as a management protocol in a real network. Since smaller Δt provides finer-grained measurement and fine-grained measurement, Δt was set to 2 seconds. The manager and the agent were connected the same intranet.

Figure 4 Variation of polling interval $\Delta t'_i$ Figure 5 Variation of actual data interval Δt_i

We estimate t_i^A by using the Managed Object *sysUpTime* which gives the time (in hundredths of a second) since the agent was last re-initialized. In the absence of explicit time-tags, this object fetched from the agent along with the traffic counter values is a good, not exact, estimate of t_i^A .

As shown in Figure 4, $\Delta t'_i$ is almost constant at about 2.2 seconds. On the other hand, as shown in Figure 5, Δt_i varies from 1.98 seconds to 2.4 seconds. The standard deviation is 0.065 seconds. Since the manager and the agent are connected to the same intranet, d_i^{Rq} is stable and almost negligible. Hence, variation of Δt_i must be due to the variation of d_i^{Prc} .

3.2 Non-realtimeness of the sampled value

Even if the data time-stamp is obtained, the manager can sample only an approximation of the value v_{t_i} of the traffic counter at time t_i , depending on agent implementations.

The agent, when queried, refers to a traffic counter which is generally some kernel variable and provides the corresponding value to the manager. This reference in general involves multiple lookups of kernel tables. To optimize the load due to such lookups, the looked up value is cached and reused for a small

cache-lifetime. So, the counter value, v_{t_i}' , sampled and returned by the agent is not updated in real-time, but is updated discretely.

The real example of this issue is shown using a *net-snmp* agent on a Linux device. For a constant traffic rate, counter value v_t increases linearly, but the counter value returned by the agent, v_t' , increases in steps. As shown in the figure, the value of the *ifInOctets* counter is updated every 15 seconds.




Figure 6 Example of discrete update

4. Discussion

As a result, the bandwidth utilization is computed based on v_{t_i}' and t_i' as shown in Eq. (2) and includes inaccuracies.

$$Bw_i' = \frac{\Delta v_i'}{\Delta t_i'} = \frac{v_{t_i}' - v_{t_{i-1}}'}{t_i' - t_{i-1}}. \quad (2)$$

These inaccuracies may be corrected if the variation of data interval Δt_i and the counter update interval can be predicted or known in advance. However, the former depends on measurement environment and is difficult to be predicted. The later is implementation dependent and its setting is not widely known.

For accurate measurement, an agent must provide explicit time-tags for the data. The High Resolution Traffic Measurement MIB [2], and Managed Object Aggregation MIB in RFC 4498 [3] use such time-tags. In these MIBs an agent adds time-tags to the data and stores the tagged data in a time sequenced manner at regular intervals. However, to the best of authors' knowledge, such time-tagging technique is not widely deployed. The network administrator must know the existence of inaccuracies for strict quality and security management.

5. Conclusion

In this work, we have examined the issues of inaccuracies in traffic measurement widely deployed in network management. The time at which the data was actually measured, is generally treated loosely. This causes inaccuracies in traffic measurements. For accurate measurements, explicit time-tagged data is essential. In the absence of time-tagged data, one needs to be aware of the limited accuracy of the results.

Acknowledgement

A part of this R&D has been promoted under the support of the Green ICT Innovation Promotion (PREDICT #115102001) of Ministry of Internal Affairs and Communications of Japan.

References

- [1] J. Case, R. Mundy, D. Partain and B. Stewart, Introduction and Applicability Statements for Internet-Standard Management Framework, RFC3410, 2002.
- [2] G. Mansfield, S. Karakala, T. Saitoh and N. Shiratori, "High Resolution Traffic Measurement," in *A workshop on Passive and Active Measurements on the Internet (PAM2001)*, 2001.
- [3] G. Keeni, The Managed Object Aggregation MIB, RFC4498, 2006.