On the 95-Percentile Billing Method

Xenofontas Dimitropoulos[1], Paul Hurley[2], Andreas Kind[2], and Marc Ph. Stoecklin[2]

[1] ETH Zürich
fontas@tik.ee.ethz.ch
[2] IBM Research Zürich
{pah,ank,mtc}@zurich.ibm.com

Abstract. The 95-percentile method is used widely for billing ISPs and websites. In this work, we characterize important aspects of the 95-percentile method using a large set of traffic traces. We first study how the 95-percentile depends on the aggregation window size. We observe that the computed value often follows a noisy decreasing trend along a convex curve as the window size increases. We provide theoretical justification for this dependence using the self-similar model for Internet traffic and discuss observed more complex dependencies in which the 95-percentile increases with the window size. Secondly, we quantify how variations on the window size affect the computed 95-percentile. In our experiments, we find that reasonable differences in the window size can account for an increase between 4.1% and 42.5% in the monthly bill of medium and low-volume sites. In contrast, for sites with average traffic rates above 10Mbps the fluctuation of the 95-percentile is bellow 2.9%. Next, we focus on the use of flow data in hosting environments for billing individual sites. We describe the *byte-shifting effect* introduced by flow aggregation and quantify how it can affect the computed 95-percentile. We find that in our traces it can both decrease and increase the computed 95-percentile with the largest change being a decrease of 9.3%.

1 Introduction

Transit ISPs and hosting providers monitor the traffic usage of their customers and typically charge them using the 95-percentile method. A period of a month is split into fixed size time intervals and each interval yields a sample that denotes the transferred bytes during the interval. An automated tool polls the SNMP counters of the appropriate router(s) or switch(es) and finds the transferred bytes. Then, the 95-percentile of the distribution of samples is used for billing. Often the 95-percentile is computed both on the inbound and the outbound directions and the smaller value is ignored.

In a billing cycle of 30 days, the 95-percentile method filters out 36 hours of spikes, which may include Denial of Service (DoS) attacks, flash crowds, and back-up traffic. The method essentially realizes a compromise between two objectives. The first objective is billing a customer based on its absolute traffic usage, whereas the second objective is billing based on the capacity of the provisioned links and the peak rates. If we consider the traffic rate as a continuous signal, then the first objective suggests using the average of the signal for billing, whereas the second objective suggests using the maximum. The 95-percentile is typically between the average and the maximum balancing, in this way, the two objectives.

S.B. Moon et al. (Eds.): PAM 2009, LNCS 5448, pp. 207–216, 2009.

Nevertheless, the 95-percentile is not the result of a sophisticated optimization. Certain large ISPs started using it many years ago and, over time, it became more widely-used and established. The properties and limitations (also due to operational constrains) of the method have not been systematically studied and are not well understood. For example, the window size used for computing the 95-percentile can vary between different providers. Most commonly a 5-minute interval is used, however used values range to as low as 30 seconds [1]. The effect of such variations on the 95-percentile has not been analyzed. In addition, the 95-percentile method is often applied on traffic flow data for billing, e.g., individual websites. The relationship between flow aggregation and the 95-percentile, as we discuss further in our paper, is a challenging research problem.

In this work, we characterize important aspects of the 95-percentile billing method using several traffic traces. We first analyze how the 95-percentile depends on the window size and find that it typically exhibits a noisy convex decreasing trend, although more complex inter-dependencies are also possible. Then, we make the assumption that network traffic is self-similar to provide a mathematical explanation of the observed decreasing dependence. We quantify fluctuations on the 95-percentile due to different window sizes and find that fluctuations are 1) significant, i.e., between 4.1% and 42.5% for low volume sites with average rates bellow 10 Mbps, and 2) negligible, i.e., bellow 2.9%, for high volume sites with average rates above 10 Mbps. Next, we describe how the byte-shifting effect of flow aggregation can affect the computed 95-percentile. We characterize the extent of byte-shifting and find that in our traces up to 35.3% of the total number of bytes can be shifted between neighboring windows causing a decrease on the 95-percentile by 9.3%.

The remainder of this paper is structured as follows: in the next section we describe the traffic traces we used for our experiments. In Section 3, we characterize the dependence of the 95-percentile on the size of the aggregation window. Then, in Section 4 we analyze the effect of flow aggregation on the 95-percentile billing method and provide supporting measurement results. Finally, we conclude this paper in Section 5.

2 Data Traces and Preprocessing

We used traffic traces collected with *tcpdump* [2] and *NetFlow* [3] on two distinct networks. The first network provided web hosting services to 46 websites of varying sizes. We collected unsampled NetFlow version 9 packets from the border router of the network that transferred more than 6 TBytes a day with average sending and receiving rates of 550Mbps and 100Mbps, respectively. The NetFlow trace spanned 27 days during April 2008. In addition, we used tcpdump to collect packet headers destined to or originating from an individual medium-volume website. The tcpdump trace spanned 30 days starting on the 17th of July 2007 and the average rate was 615 Kbps.

The second network was a medium-size enterprise campus network that receives transit services from a commercial and an academic ISP. In particular, we collected data from the IBM Zurich Research Laboratory campus that hosts approximately 300 employees and at least as many networked computers. For our experiments, we used a tcpdump trace collected over a period of approximately 63 continuous days, from the 2nd of March until the 5th of May 2008. The trace includes all the outgoing and incoming packet headers and the overall average traffic rate was 7.536 Mbps.

We processed the tcpdump and NetFlow traces and created a set of traffic volume time series for our experiments. We first parsed the tcpdump data collected from the campus network and computed the total (both inbound and outbound) number of bytes observed in each consecutive second. In this way, we derived a sequence of Byte counts, which we split into two time series that spanned approximately one month each.

In addition, we derived a time series for each individual website in the hosting environment. We used the NetFlow trace and associated flows with sites based on the known IP addresses of the latter. Then, we distributed the size of a flow uniformly over its time-span and derived how much each flow contributed to each one-second window. By aggregating the contributions of the flows, we constructed a baseline time series that indicates the total bytes sent and received from a site during each consecutive second. In this way, we derived a time series for each individual website. Out of the 46 sites, we ignored the time series that corresponded to the 12 lowest-volume sites that had on average a rate smaller than 1 Kbps. These sites appeared virtually unused and, therefore, we used the remaining 34 sites for our experiments. In addition, we derived one last time series from the tcpdump trace for the individual website using the procedure we outlined above.

Overall, we used 37 time series with an one-second time resolution. We call these time series baselines. To measure the effect of the aggregation window size, we aggregated the baselines using windows of varying size and computed the 95-percentile of the aggregated series that we report in our experiments.

3 95-Percentile versus Window Size

3.1 Measurements

Using the different traces, we first examine how the 95-percentile depends on the window size. In Figure 1(a), we illustrate the relationship between the 95-percentile and the window size for the traffic of the enterprise campus network. The computed 95-percentile reflects what a transit provider would charge the network. As the size of the aggregation window decreases, we observe that the 95-percentile increases. The 95-percentiles corresponding to a window size of 30 seconds increases by 5% with respect to the 95-percentile of a 300-second window.

Secondly, we examine how the 95-percentile depends on the window size in the case of a web-hosting provider charging a high-volume website. In Figure 1(b), we illustrate the corresponding plot for the website in our traces that had the highest mean traffic volume. We observe again that the 95-percentile exhibits a noisy decreasing trend as the window size increases. In this case the relative fluctuations of the 95-percentile are smaller than in the campus network. The 95-percentile increases by only 0.7% between a 300 and a 30-second window. In the set of 95-percentile values that correspond to window sizes between 30 and 400 seconds, the maximum 95-percentile increase we observed was 1.3%[1].

[1] Note that in Figure 1(b) we also illustrate the behavior of the 95-percentile for window sizes between 2 and 30 seconds, however, we do not use this range of values to find the maximum 95-percentile increase, as in practice such low window sizes are unlikely.

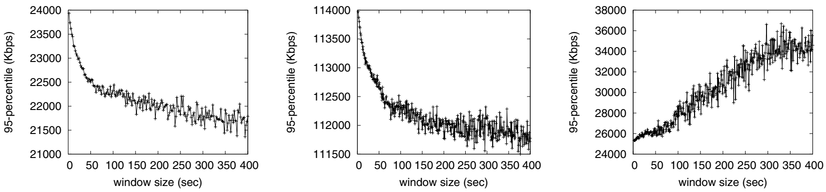

(a) medium-volume campus network (b) high-volume website (c) atypical website

Fig. 1. 95-percentile versus aggregation window size

We plotted and examined the dependence of the 95-percentile on the window size for each individual site. For 23 out of the 34 sites the dependence on the window size had two characteristics in common with Figures 1(a) and 1(b): 1) the 95-percentile gradually decreases following approximately a convex curve; and 2) the 95-percentile curve is noisy often exhibiting significant fluctuations between nearby points. Among the remaining 11 sites that did not follow the identified trend, 5 exhibited a very low monthly traffic volume that on average remained bellow 50 Kbps. The last 6 medium-volume sites had irregularities in their traffic patterns, like lack of time-of-the-day or day-of-the-week effects, occasionally very high bit-rates, or long down-time periods. In Figure 1(c), we illustrate the dependence for one of the sites we identified irregularities. In this case the 95-percentile increases with the window size, which phenomenon we traced back to very high periodical traffic spikes. The spikes were close to a maximum two orders of magnitude larger than the mean traffic rate.

We further study how the fluctuations on the 95-percentile relate to the mean traffic rates of the sites. We define the *maximum fluctuation* of the 95-percentile to be the increase of the largest over the smallest value in a set of 95-percentile values. For each baseline time series, we computed aggregate time series using window sizes between 30 to 400 seconds. Then, for the aggregate time series we computed the corresponding 95-percentiles and found their maximum fluctuation. Figure 2 plots the maximum fluctuation of the sites versus their mean traffic rate. We observe that as the mean traffic rate of a site increases, the fluctuation on the computed 95-percentile decreases. For high-volume sites with mean traffic rate above 10 Mbps, the maximum fluctuation is bellow 2.9%. On the other hand, for medium and low-volume sites with mean traffic rate lower than 10 Mbps the maximum fluctuation is larger reaching up to 915% in one extreme case, but mainly varying between 4.1% and 42.5%. This observation suggests that changes in the window size can introduce notable variations in the computed 95-percentile value only on sites and networks with small and medium traffic rates. On the other hand, high-volume sites or large networks and ISPs are not significantly affected by varying the window size.

3.2 Analysis

We can model the effect of the window size as an aggregation process that takes the mean of each consecutive m samples of a traffic volume series a_s to construct an aggregate time series $a_s^{(m)}$. Taking the mean tends to decrease the volume of large samples

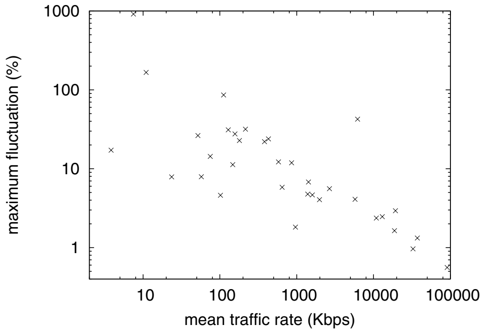

Fig. 2. Maximum fluctuation of 95-percentile versus mean traffic rate of sites

and makes the 95-percentile smaller. In theory, however, increasing the window size may also result in an increase on the 95-percentile. We illustrate this with a simple example. Assume that the baseline time series a_s corresponds to one-second time intervals and that the maximum number of aggregated samples m_u corresponds to one month. In this extreme case, the aggregated time series $a_s^{(m_u)}$ has a unique sample that is equal to the overall mean of a_s. The 95-percentile of $a_s^{(m_u)}$ is also equal to the overall mean traffic volume $\bar{a_s}$. In addition, consider that the 95-percentile of the baseline a_s can be smaller than $\bar{a_s}$, which indicates that the aggregation process will tend to increase the 95-percentile. In practice, however, real-world traffic signals typically have a mean that is smaller than their 95-percentile. The mean and 95-percentile of the baseline time series corresponding to the 1st-month trace of the campus network are 7.3 Mbps and 24.1 Mbps, respectively. As a result, increasing the window size tends to decrease the 95-percentile.

Internet traffic is known [4,5] to exhibit scaling effects. Let $X(s)$ and $X^{(m)}(s)$ denote the processes that generate a_s and $a_s^{(m)}$, respectively. If we assume that the two processes exhibit *exact self-similarity* [4] with Hurst parameter H, then their distributions are related:

$$X^{(m)}(s) \stackrel{d}{=} m^{H-1} X(s).$$

In addition, the $(1-\gamma)$-quantiles $X_{1-\gamma}(s)$ and $X_{1-\gamma}^{(m)}(s)$ of the distributions are related:

$$X_{1-\gamma}^{(m)}(s) = m^{H-1} X_{1-\gamma}(s). \quad (1)$$

For Internet traffic, the Hurst parameter H takes values between 0.5 and 1. Fixing $X_{1-\gamma}(s)$ in the last equation and setting γ to 0.05, we get that the 95-percentile of the aggregated time series decreases polynomially as the aggregation window m increases. This behavior is consistent with our observations in Figure 1 and with the remaining figures for the 23 other sites.

Besides, Figure 2 indicates that fluctuations on the 95-percentile are larger for low and medium-volume sites and smaller for high-volume sites. Higher volume sites are associated with a higher degree of statistical multiplexing. As a result, they exhibit a lower traffic burstiness than low and medium-volume sites. We speculate that this lower burstiness results in smaller fluctuations on the 95-percentile. We can use the *relative standard deviation (RSD)* of a distribution, i.e., the standard deviation divided by the mean, to quantify the burstiness of a traffic signal. The RSD is computed on the baseline time series before aggregation. In Figure 3 we illustrate the rank correlation between and RSD and the maximum 95-percentile fluctuation of the websites. We see that the points are aligned mostly along the line on the 45-degree angle, which indicates a strong correlation between the RSD and the maximum 95-percentile fluctuation. The Spearman correlation coefficient, in particular, is 0.84. This high correlation suggests that a high (low) degree of traffic burstiness results in more (fewer) 95-percentile fluctuations. Traffic burstiness can easily be quantified (using RSD) and, therefore, it can serve as indicator on how susceptible a site and a network is to 95-percentile fluctuations.

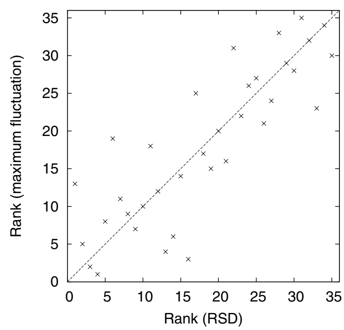

Fig. 3. Rank correlation between the traffic burstiness, i.e., relative standard deviation (RSD) of traffic rate distribution, and the maximum 95-percentile fluctuation of the different sites

In summary, our characterization and analysis of the dependence of 95-percentile on the aggregation window size yields the following important observations:

– For sites and networks of high and medium traffic rates, the 95-percentile follows a noisy decreasing trend, as the aggregation window size increases. This trend can be modeled as polynomial decrease.
– Fluctuations on the computed 95-percentile due to the size of the aggregation window are higher in low and medium-volume traffic mixes and negligible in high-volume mixes. Traffic burstiness is indicative of how susceptible the 95-percentile is to fluctuations.

4 95-Percentile from Flow Data

4.1 Measurements

NetFlow data are typically used to bill individual sites in web-hosting environments, where the traffic crossing router interfaces and incrementing SNMP counters might be destined to many different customers. NetFlow aggregates the packets of a flow and reports its duration, size, and timestamps among other attributes. Then, the 95-percentile is computed from the flow records by uniformly distributing the size of a flow over its lifespan and by counting the overall contribution of the flows in each aggregation window. The volume of a flow may exhibit variations, which are smoothed by aggregating the size of a flow across its duration. This “horizontal” aggregation is illustrated in Figure 4. It effectively shifts bytes between neighboring windows, which affects the estimated traffic volume for a window and may, therefore, skew the computed 95-percentile.

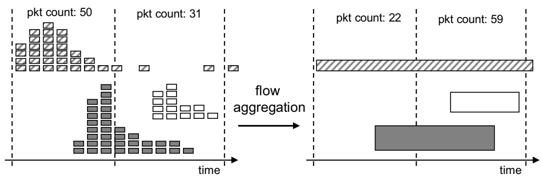

Fig. 4. Effect of flow aggregation on traffic volume observed during a time interval

To investigate the effect of flow aggregation on the 95-percentile we used the tcpdump traces. We first computed the 95-percentiles of time series constructed from packet-level data. These time series attribute each packet to the correct window and do not suffer from the problem discussed above. Then, we grouped packets in the tcpdump traces into flows, using the standard five-tuple flow definition, and computed the amount that each flow contributed to each window. In this case, the resulting time series were skewed due to flow aggregation.

In Figure 5, we plot the 95-percentile versus the window size for the packet and flow-based data using the tcpdump traces of the individual website and of the enterprise campus network. The plots correspond to three distinct behaviors. In Figure 5(a), we observe that the curves corresponding to using flow and packet-level data to compute the 95-percentile are almost indistinguishable, indicating negligible artifacts introduced from flow aggregation. The maximum increase of the 95-percentile in the range of window sizes between 30 and 400 seconds in only 0.42%. In Figures 5(b) and 5(c), we show the corresponding plots for the first and second month of the campus trace, respectively. Figure 5(b) exhibits small differences between the the packet and flow-based curves. For a window size of 300 seconds, the 95-percentile increases by 1%, whereas the maximum increase in the range above 30 seconds is 2.89%. Figure 5(c) demonstrates a significant decrease on the flow-based 95-percentile. This decrease is consistent through out the range of window values and has a maximum value of 9.3% at a window size of 200 seconds. At the commonly-used 300-second window size the decrease is 5.8%.

In Table 1 we illustrate the total traffic volume that was shifted due to the effect of Figure 4 between windows. We observe as expected that the volume of shifted bytes decreases as the window size becomes larger, since fewer flows cross window boundaries. In agreement with Figures 5, the amount of shifted bytes is smaller for the website, larger for the 1st month of the campus trace and even larger for the 2nd month. In addition, in Table 1 we mark the fraction of the total traffic that was shifted to a neighboring window. The traffic fraction is as large as 35.3% indicating that the effect of Figure 4 can be prevalent leading to significant distortion of a flow-based traffic signal.

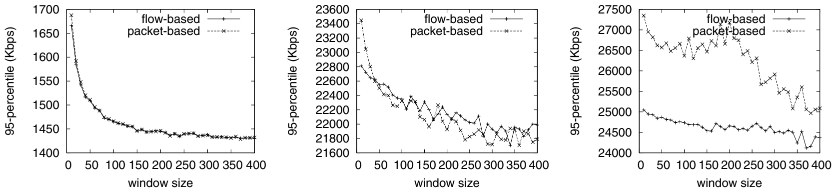

(a) Website (b) Campus (1st month) (c) Campus (2nd month)

Fig. 5. 95-percentile versus window size computed from packet-level and flow-level data

Table 1. Shifted traffic between neighboring windows in our experiments

	Website	Campus 1	Campus 2
Window size (sec)	Shifted Mbytes / Shifted traffic fraction		
30	1,270 / 2.4%	201,431 / 30.2%	248,892 / 35.3%
100	134 / 0.8%	48,705 / 24.4%	63,748 / 30.1%
200	38 / 0.5%	20,112 / 20.1%	27,583 / 26.1%
300	19 / 0.3%	11,820 / 17.7%	16,714 / 23.7%

4.2 Analysis

We further investigated the traces to understand the reasons leading to the three distinct behaviors illustrated in the above figures. Figures 5(b) and 5(c) correspond to two consecutive months in the same network and the significant difference in the second month worthed further examination. Our investigation revealed during a period of a week in the second month hourly large traffic spikes that persisted even during low-volume periods, like nights and weekends. Figure 6 compares the traffic patterns in the week with the spikes with another week in the first month of the trace. The periodic (hourly) spikes

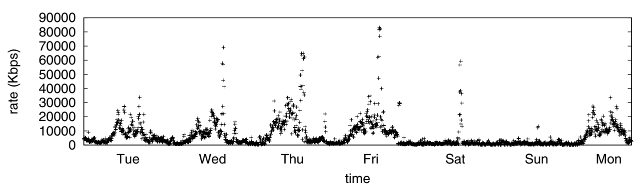

(a) first month

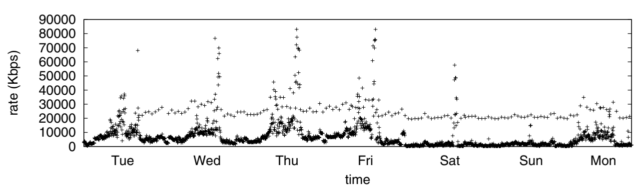

(b) second month

Fig. 6. Weekly traffic rate variations in the first and second month of the enterprise campus network

shown in Figure 6(b) are ranked in the top 5% of the monthly samples. Without flow aggregation they increase the 95-percentile, however, with flow aggregation the spikes are smoothed and therefore do not significantly affect the 95-percentile.

On the other hand, the high agreement between the packet and flow-based curves in Figure 5(a) results from the distinct properties of the trace. In particular, the website trace includes mainly short-lived http flows, which are less susceptible to the effect of Figure 4 than the more diverse set of traffic flows encountered in the campus network. For this reason, we only observe very few bytes shifted between windows in Table 1.

In summary, our analysis illustrates that ‘horizontal” aggregation due to flow accounting skews the computed 95-percentile. The (lack of) decrease/increase on the 95-percentile depends significantly on the individual characteristics of the examined traffic traces and ranges from a 2.89% increase to a 9.3% decrease.

5 Discussion and Conclusions

In this work, we used a large set of data to study the widely-used 95-percentile billing method. We make a number of observations: 1) for medium and high-volume traffic rates the 95-percentile typically decreases as the aggregation window size increases; 2) more complex relations between the window size and the 95-percentile are possible; 3) the observed 95-percentile fluctuations were significant only for medium and small-volume traffic rates and rather negligible for high-volume sites; 4) flow aggregation can skew the computed 95-percentile value causing in our data a decrease up to 9.3%. Besides, we used certain properties of Internet traffic to justify our observations and provide a mathematical basis.

A natural question to ask is how to compute the 95-percentile correctly? One could make the assumption that the traffic rate is a continuous signal and could try to find the 95-percentile of the continuous signal. In this case, the 95-percentile would be well-defined. In our early work, we experimented with trying to find an “ideal” 95-percentile of an assumed underlying continuous signal. However, it turned out that the Fourier spectrum of network traffic has many high frequency components, which by the well-known standard Shannon sampling theorem would require sampling network traffic with a very high frequency or equivalently aggregating network traffic using a very small window. Such a “correct” 95-percentile, therefore, would be impossible to compute in practice due to the high measurement and instrumentation overhead it would require.

One take-away of our work is that providers should *all* use a fixed, ideally standardized, window size to charge their customers, in order to enable a fare comparison between different billing rates. This might be already happening to a certain extent, as the 5-minute window size is popular. However, not everybody uses the same window size and more importantly the over-charging consequences of varying the window size were not known, to the best of our knowledge, before our work.

A second take-away of our work is that “horizontal” aggregation introduced by flow accounting can skew the number of bytes during a window interval, which, in turn, can bias the computation of the 95-percentile. This observation is significant, as flow technologies are widespread for billing. A possible simple solution to this problem is to have synchronized routers/switches that at fixed timestamps, e.g., at 16:00, 16:05, ..., expire flows. This aligns flow durations within the aggregation window intervals and, therefore, the described byte-shifting effect is avoided. This approach, if not implemented intelligently, however, could lead to flow export synchronization problems.

Summing up, in the future we would like to understand better the properties of network traffic that affect the computed 95-percentile from flow data. An important challenge is collecting long, i.e., ideally at least one month long, tcpdump traces from different networks. In addition, a model of the byte-shifting process described in Section 4 could have several applications, like in predicting 95-percentile changes or in reconstructing accurate traffic time series from flow data.

References

1. Webhostingtalk Forum: 95th percentile billing polling interval (2008) (last accessed: 09/23/2008), `http://www.webhostingtalk.com/showthread.php?t=579063`
2. The Tcpdump team: tcpdump, `http://www.tcpdump.org/`
3. Cisco: IOS NetFlow, `http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html`
4. Leland, W.E., Taqqu, M.S., Willinger, W., Wilson, D.V.: On the self-similar nature of ethernet traffic. SIGCOMM Comput. Commun. Rev. 23(4), 183–193 (1993)
5. Paxson, V., Floyd, S.: Wide-area traffic: the failure of poisson modeling. In: SIGCOMM 1994: Proceedings of the conference on Communications architectures, protocols and applications, pp. 257–268. ACM, New York (1994)