

MRTG 95th Percentile

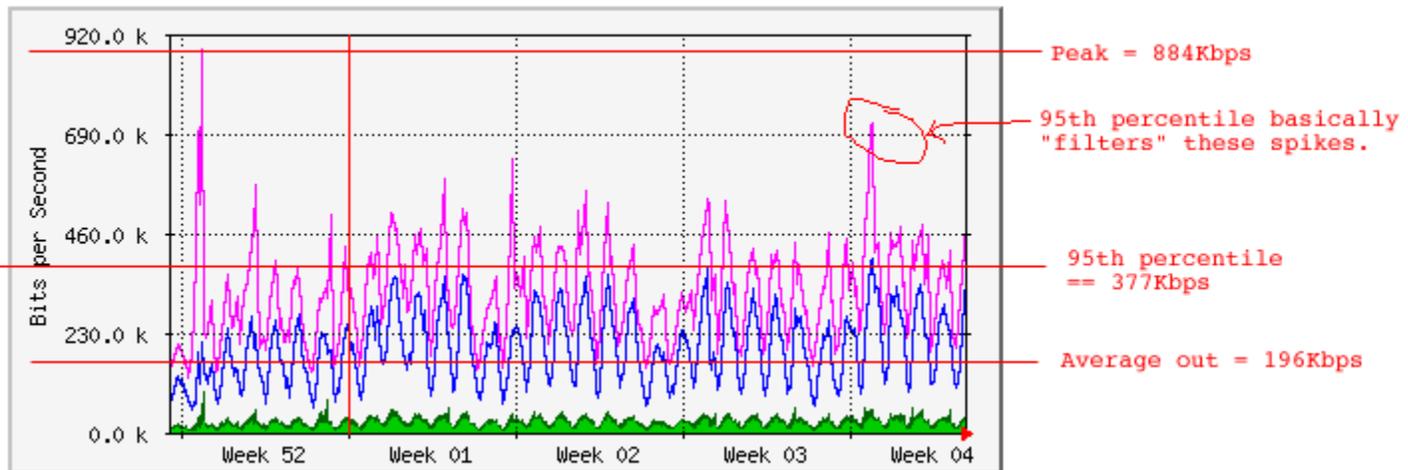
This is Mirrored Information from SeanAdams.com Website

MRTG 95th Percentile

95th Percentile Billing is the Standard way Bandwidth is billed by ISPs and the standard way that ISPs are themselves billed. Typically it means that samples are taken every 5 minutes for a month and the top 5% are discarded-- this gets rid of the so called "spikes". The client is then billed based on 95th percent of their bandwidth usage. ISPs have to reserve capacity on their fiber for you and they pay each month - just like you do.. If you are using a certain capacity even for one day then that means that no one else can use it and therefore that capacity has to be reserved for you even if you only use it for one day! This is why ISPs don't usually like "spikey" traffic -- clients with "spikey" traffic don't like to pay for the whole month for their traffic if they only need it one day a week. And unless an ISP has purchased more bandwidth then they are using, they typically will want to bill you the same way that they themselves are billed-- based on 95th percentile - the fair capacity bandwidth measuring tool. Typically you will see either MRTG, RTG, or CACTI graphs but they are all basically the same -- graphing samples of your traffic.

MRTG is an SNMP monitoring/graphing program by Tobias Oetiker and Dave Rand.

Example: busy web server



What is the 95th percentile, and why is it useful in measuring bandwidth?

The 95th percentile is the smallest number that is greater than 95% of the numbers in a given set. The reason this statistic is so useful in measuring data throughput is that it gives a very accurate picture of the cost of the bandwidth. Here's an example. Suppose an ISP sells you a T1 line, but you're only using it to access the web. Even though you might frequently download very large files (filling the pipe) your cost to the ISP is negligible, because your usage is intermittent. A single T3 connection to the backbone could easily support hundreds of such downstream customers, and never become saturated. As another example, suppose you are hosting a very busy web site that half-way fills your T1 for several hours every day. This type of bandwidth is more expensive, because your ISP can't oversell their connection to the backbone as effectively. The important thing to realize is that it doesn't cost your ISP anything to sell you a pipe of any particular size - it is the sustained rate of data transfer that costs them money. The sum of the 95th percentile usage of all of an ISP's customers predicts the peak amount of backbone traffic that the ISP will incur (in a given direction).

Here are some examples. ISPs must charge for bandwidth by one of three means:

1. Sell a flat rate, possibly bandwidth limited connection, and try to sell to customers whose usage patterns are not so intense. Nearly all DSL providers do this. The customers like it because they don't have to worry about how much bandwidth they use, and ISPs like it because it simplifies billing, and they make more money as long as they have plenty of low-usage customers. The problem, particularly if the ISP is selling very fast connections, is that the ISP can become overwhelmed by even a small number of high-usage customers. Even residential customers can be such high-usage clients, thanks to recently popular services such as peer-to-peer file sharing.
2. Sell a fast connection (eg 100Mbit Ethernet, which is inexpensive) and charge for the volume of data transfer - eg number of Gigabytes per month. This model works great for web sites, which almost always generate traffic in a predictable bell curve. However, it severely penalizes customers who use bandwidth intermittently. For example, suppose a customer runs an automated off-site backup every night. This brief usage spurt costs the ISP almost nothing. Although the recurring sustained data rate is low, the customer gets charged for a huge amount of bandwidth.
3. Sell a fast connection and bill by 95th percentile. By now this should make sense - it's a fair system where everybody pays for what they get. The advantage to the customer is that they get the performance of a high-speed connection, while paying only for their actual usage. ISPs like it because they don't have to worry about high-usage customers upsetting their overselling ratios.

Irrespective of billing concerns, the 95th percentile is a very interesting and useful figure. Bottom line is it tells you how much of your connection you're really using (and really need).