Jonathan Foulkes
CEO – evenroute.com

Metrics helpful in assessing Internet Quality

This paper presents some Internet quality metrics that have proven to work well in real-world deployments and proposes additional metrics to further refine connectivity quality identification.

The existing metrics and results presented below have evolved from developing and deploying a fleet of smart routers (the IQrouter) that deliver an optimized and improved Internet user experience, even on challenging lines. The data set spans five+ years and has data points primarily from North America.

Focusing on Internet Quality improvement in that product has shown that that the following properties contribute to a good user experience for common Internet usage (In descending order of priority):

- Low latencies under working loads
- Low and consistent jitter
- Traffic Prioritized per type and DSCP (e.g., Cake multi-tier)
- Sufficient capacity for the load (often a lot less than people imagine is needed)
- Stable lines with good uptimes

The metrics necessary to capture and report in support of the above properties are discussed, as is their rationale. A proposal for additional metrics to understand on-premises vs. external quality and performance issues is presented.

Internet Quality Metrics

Line capacity and latency under working loads

Why low latency under working load is critical to Internet quality is answered by the extensive work and documentation published at bufferbloat.net, so it will not be rehashed here. The deployments and gathered metrics show that both the prevalence of the problem and the value of remediation are significant, with reductions in latency under load of up to 100x.

Capturing the latency under load as well as actual line capacity is best performed on the router controlling the traffic at the border between the local network and the ISP service, as it can measure the capacity of all flows and not just the sample from the synthetic load test.

Latency measurements using both active and passive methods are captured to assess loaded (working) vs. unloaded performance, with and without traffic management, along with varying traffic management settings.

It is critical that the traffic manager be correctly configured with respect to multiple properties, such as the line type, and it must be correctly set for the current line capacity. In addition, the data shows that line capacity varies (discussed further below) over time, which requires a dynamic adaptation system to keep traffic management effective. Therefore, measurement and adaptation are necessary on an ongoing basis.

Determining the actual line capacity is an activity that must be performed repetitively, as our data shows that 36% of lines have capacity metrics that vary by more than 10% throughout the day and across multiple days.

Accurately measuring line capacity is critical to the points above, and it requires accounting for all network usage at the CPE edge router.

End-user device-based speed tests are notoriously inconsistent due to many factors, including local connectivity (WiFi vs. Ethernet), competing LAN traffic, device load/configuration, time of day (Local loop/backhaul capacity).

So even if the ISP has provisioned 200Mbps down on the cable modem (often with a decent margin), end-user device metrics might not reliably reflect that due to the above factors.

Only a capacity test run from the router (assuming enough router resources) can account for the totality of local traffic and perform an independent capacity test. And even then, various factors might limit the accuracy of said test. And as noted previously, multiple samples are required to increase accuracy in finding the high-water mark.

Documenting the actual capacity of the line with the router-based test helps users improve their experiences by eliminating sources of capacity cuts (e.g., poor WIFI) and latency increases (e.g., PC 'utilities' or A/V).

This report is a useful tool for ISPs as it reports the actual provisioned capacity (pre-traffic manager) and helps their support staff direct customers to identify local issues.

Captured metrics

- Line capacity (in Mbps) in both directions with no traffic controls
  - Latency metrics with the line fully loaded, including drops
- Line capacity (in Mbps) in both directions with traffic management enabled
  - Latency metrics with the line fully loaded, including drops
- Continuous latency metrics accumulated and averaged by hour by day
- Ongoing capacity metrics collected and averaged by hour by day
- Traffic manager capacity targets dynamically changed (date/time, target values)

These metrics are used two ways:

1. To control the traffic manager to always maintain low latencies
2. To inform the user of line capacity usage and latency performance

Reporting current traffic manager capacity targets and their variance over time helps users understand the quality of the service they receive and correlate it with observed experiences.

End-users tend to focus on line speed (capacity) as the most critical line metric yet lack information on how much capacity they actually use regularly. Reporting average and maximum capacity usage over time helps them achieve an ideal balance between cost and user experience. Users are universally surprised at their actual capacity usage metrics.

Following are illustrations of the points and metrics mentioned so far.

First, an example of a line whose capacity varies throughout the day.
This cable user with a 180Mbps line requires multiple adjustments to traffic management targets to maintain low latencies.

The orange line is download, and the green line is upload capacity targets. Brown lines indicate dynamic lowering, bright green indicates dynamic increases to targets.

Our data set shows an average of 72 dynamic changes since deployment per unit on cable lines and an average of 144 dynamic changes since deployment for DSL lines. This is because limited capacity lines, such as DSL, are more likely to reach max line capacity and therefore encounter sag-induced bufferbloat requiring dynamic adjustments.

We also observe that once the industry started provisioning service at speed points >100Mbps, the incidence of line capacity variance dramatically increased.

Our data shows that 24% of cable lines <100Mbps have >10% capacity variance, but 52% of lines >100 show greater than 10% variance.

The next example is about capacity reporting from a 300Mbps cable line with two WFH tech workers and 60+ network-connected devices, the green elements in the bars represent average capacity usage, and it only shows up in the evening when streaming 4K video.

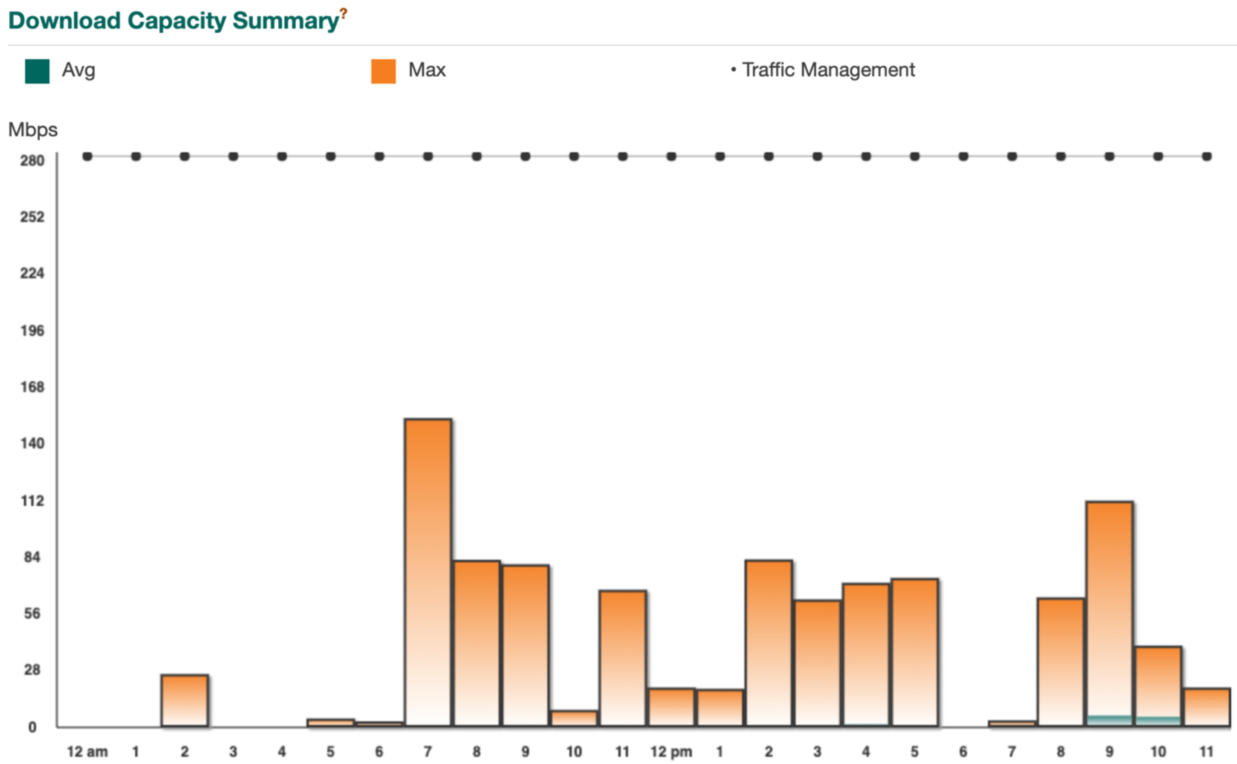

This is an excellent example of someone who could dial back to 200 or even 100Mbps service without noticing the difference.

In contrast, the following example could use substantially more.

This is from a 50 Megabit cable line serving a dorm with 44 college students and >150 devices. The average ping hardly varies despite the onslaught of traffic. The download capacity averages for a couple of the hours are practically the full provisioned capacity.

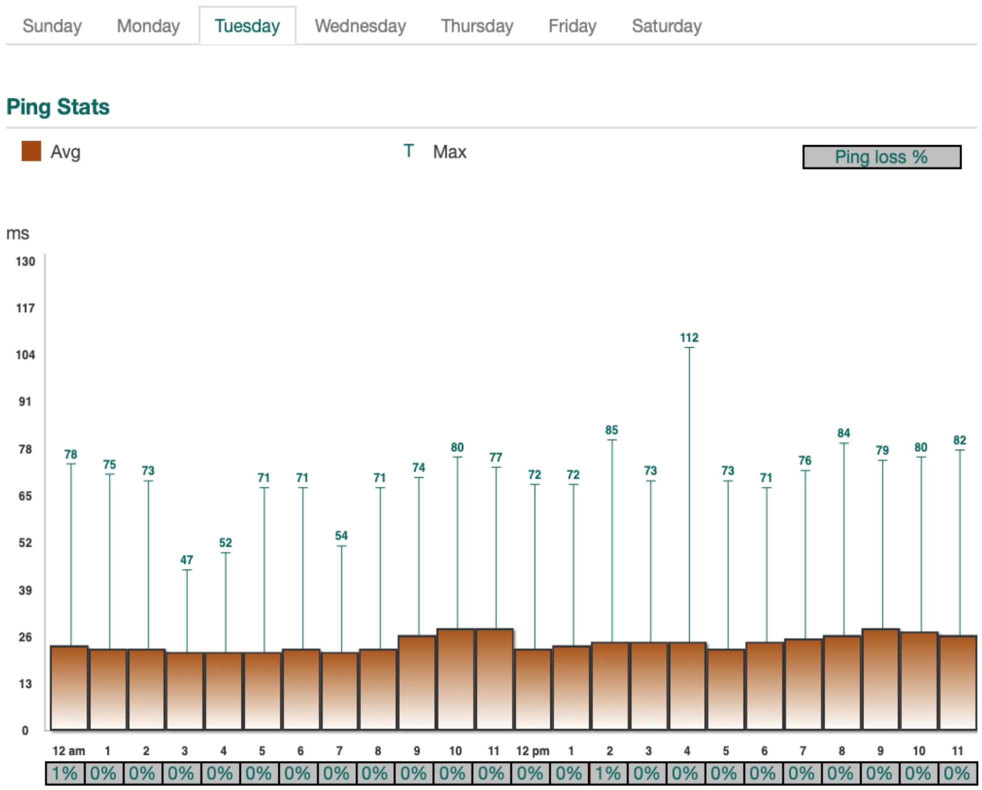

That is a great example of a deployment where a LOT more capacity would help. But thanks to effective traffic management in the automatically tuned IQrouter, they are happy enough.

By reviewing these graphs, end users can decide if they have too much or too little capacity and thus make an informed decision about which service plan would suit them best.

Router-based ping statistics also help users determine whether latencies are from local causes (PC 'utilities' interfering) or are intrinsic to the line.

These metrics are also valuable to ISPs and others looking at capacity planning.

Line stability metrics

These metrics are about two line-related properties:

- Was the line down, and for how long
- Was the line unstable, and when

Stable lines are critical to good user experiences, yet it is a metric rarely even mentioned.

We define an unstable line condition as follows:

A line with multiple instances of higher-than-expected latencies recorded when simultaneously seeing very little actual traffic on the line.

Or more simply: an increase in latencies with little to no load.

Bad connections or failing modems typically cause this. However, it also can be caused by backhaul congestion driving up bloat on DSLAMs or CMTS with limited backhaul.

There is not much an end-user can do about this other than raise the incidents to their ISP to demand a resolution. However, it is a handy metric to get to the bottom of line-related issues and pressure the ISP with facts about the ongoing condition of the line or their supplied CPE hardware.

A more common issue is line drops or total loss of connectivity to the Internet. Obviously, losing Internet access is a bad user experience, but few people know precisely when or how frequently outages happen or how long they last. Logging and reporting on these help end-users work with the ISPs to resolve localized line problems. These are often associated with lines exhibiting unstable line metrics as well.

Captured metrics (per unit)

Unstable event start, event count, actual latency, actual traffic, current traffic manager limits

Unstable event end, event count, actual latency, actual traffic, current traffic manager limits

Line down start date/time

Line up date/time

Following are illustrations of the line stability metrics.

This is an example of the unstable line section in an IQrouter Customer Dashboard report for a 100Mbps cable line near the start of the Pandemic in 2020.

Unstable Line Changes ?

stable line	baseline latency	actual latency	actual upload	actual download	qosup	qosdown	unstable server date	unstable local date	unstable cnt
NO	14	2116	18	62	5796	95000	2020-03-29	2020-03-29	1
YES	14	18	13	21	5796	95000	2020-03-29	2020-03-29	1
NO	14	2232	5	4	5796	95000	2020-03-30	2020-03-30	4
YES	14	18	3	3	5796	95000	2020-03-30	2020-03-30	4
NO	14	1223	43	330	5796	95000	2020-03-31	2020-03-31	8
YES	14	16	5	5	5796	95000	2020-03-31	2020-03-31	8
NO	14	1669	7	5	5796	95000	2020-04-01	2020-04-01	8
YES	14	18	4	3	5796	95000	2020-04-01	2020-04-01	8
NO	14	1407	4	3	5796	95000	2020-04-02	2020-04-02	5
YES	14	17	8	12	5796	95000	2020-04-02	2020-04-02	4
YES	14	16	9	14	5796	95000	2020-04-03	2020-04-03	1
NO	14	260	22	78	5796	95000	2020-04-18	2020-04-18	1
YES	14	16	2	2	5796	95000	2020-04-18	2020-04-18	1
NO	14	4938	41	60	5796	95000	2020-08-02	2020-08-02	1
YES	14	16	68	1939	5796	95000	2020-08-02	2020-08-02	1
NO	14	4309	22	57	5796	95000	2020-08-03	2020-08-03	5
YES	14	18	21	264	5796	95000	2020-08-03	2020-08-03	4
NO	14	1960	5	66	4694	68638	2021-02-04	2021-02-04	4
YES	14	18	1	8	5796	95000	2021-02-04	2021-02-04	3
YES	14	22	3	36	5796	95000	2021-02-05	2021-02-05	9
NO	14	5727	9	129	5796	95000	2021-02-05	2021-02-05	8
NO	14	1272	8	94	5216	80750	2021-02-06	2021-02-06	3

Here is an article that helps interpret the above report.

Our data shows that 27% of all lines experience one or more unstable line events and that 16% experience over 50 unstable events.

By connectivity, cable/fiber has 19% with >50 events vs DSL/VDSL with 14% > 50 events.

Next is an example of the Line Drops section, where this DSL line was having a terrible day with constant drops due to modem re-syncs:

Line Drops ?

line state	start loss	end loss
Lost Connectivity	2017-06-07 21:45:02	2017-06-07 21:46:21
Lost Connectivity	2017-06-07 22:47:16	2017-06-07 22:49:05
Lost Connectivity	2017-06-08 00:19:16	2017-06-08 00:20:18
Lost Connectivity	2017-06-08 03:27:15	2017-06-08 03:28:21
Lost Connectivity	2017-06-08 05:53:16	2017-06-08 05:54:20
Lost Connectivity	2017-06-08 08:42:16	2017-06-08 08:44:18
Lost Connectivity	2017-06-08 10:16:17	2017-06-08 10:17:20
Lost Connectivity	2017-06-08 11:31:17	2017-06-08 11:32:19
Lost Connectivity	2017-06-08 11:48:17	2017-06-08 11:51:20
Lost Connectivity	2017-06-08 12:16:59	2017-06-08 12:18:20
Lost Connectivity	2017-06-08 13:35:16	2017-06-08 13:36:20
Lost Connectivity	2017-06-08 13:51:15	2017-06-08 13:52:20
Lost Connectivity	2017-06-08 13:58:16	2017-06-08 13:59:21
Lost Connectivity	2017-06-08 14:13:16	2017-06-08 14:14:18
Lost Connectivity	2017-06-08 14:30:16	2017-06-08 14:31:20
Lost Connectivity	2017-06-08 15:21:15	2017-06-08 15:22:20
Lost Connectivity	2017-06-08 17:45:15	2017-06-08 17:47:22
Lost Connectivity	2017-06-08 19:36:16	2017-06-08 19:37:20
Lost Connectivity	2017-06-08 19:56:16	2017-06-08 19:57:20
Lost Connectivity	2017-06-08 20:19:16	2017-06-08 20:20:18
Lost Connectivity	2017-06-08 20:36:16	2017-06-08 20:37:23
Lost Connectivity	2017-06-08 21:10:17	2017-06-08 21:11:23

Pretty much all lines will eventually log a line drop. But 22%, like the DSL line example above, will log greater than 100.

Local network metrics

What is not yet commonly captured are metrics related to local network performance and latencies. Many user-experience issues are due to things like congested WiFi, poor WiFi range, and throughput. Yet, other than contrasting with the metrics above, it is hard to determine if poor performance is local or due to problems from the router outward.

We see a need for a test between the end-user devices and the CPE router that would report the capacity and latency, or alternatively, a single value such as returned from the Apple RPM test of the local link during the test. Poor results on that link would help users track down WiFi, powerline links, or cabling issues. The challenge is that there is no industry standard set of tests and metrics (ideally simple) that network equipment and network client vendors can support. We hope that an outcome of this workshop is further industry discussion and work to define such a test.

As a vendor who boldly proclaims to help improve a user's Internet Quality, we deal intimately with the challenges involved in identifying where there might be a breakdown in that quality.

We find it helpful to segment the problem as follows:

Segmentation and localization of issues. Where is the lack of quality?

- In-home – non-ethernet links subtracting quality, or even a flaky Ethernet cable
- Last-mile, CPE to ISP core
    - Idle and under load
- Peering lag – my ISP to ISP hosting target server
    - In gaming, much lag comes from this area
- Service lag – the responsiveness of the actual service might be low
- For each of the above, is it just delay, or are there variances in loss and jitter?

We believe there is room for more nuanced testing that helps identify which of the segments is subtracting the most quality from the device to target link.

Developing standards for the client device to CPE router test is a good first step, as many user-controllable mitigations can be applied and likely the highest payback in terms of improving user experience.

Conclusions and Recommendations

In summary, reporting on:

- Actual line capacity, with minimum and maximums, measured
- Hourly/daily latency metrics
- Actual line capacity usage per hour per day
- Line related issues with ‘unstable’ lines and line drops
- Traffic Management targets over time

These metrics all help vendors, providers, and end-users get a clear picture of the elements that contribute to their Internet Quality, or lack thereof, and are very useful in communicating and resolving issues as they come up.

Using these metrics to tune and control traffic management is required to mitigate bufferbloat under varying conditions.

Adding local network metrics and insights into elements between the CPE and the wider Internet would allow more targeted recommendations and further dynamic mitigations.

In conclusion, these metrics, reports, and tuned traffic management help improve Internet Quality for end-users.