# Stop Guessing, Start Measuring: Collaborative Data Collection and Internet Fitness

***Leslie Daigle**, Thinking Cat Enterprises LLC*[1]

December 9, 2016.

How recently have you asked yourself a question about the state of the Internet: Is it secure? Is it progressing? Is it progressing "well enough"? Can it handle <X>? New technologies and operational practices are developed by individuals and entities that perceive shortcomings: how to know if they are they needed? No (network or infrastructure) operator wants to make changes based on fear, uncertainty and doubt, and yet those are often the clearest driving forces for change.

A better approach would make use of concrete measurements and collected data about the Internet, the networks that compose it, and infrastructure services upon which it depends to provide a baseline understanding of its actual state.

One of the consistent features of the Internet is that it is relatively accessible for measurement by independent observers. There are many existing measurements activities that provide important information about different aspects of the Internet, and each has strengths for particular purposes[2]. A different approach is to coordinate measurements and observations taken by first-hand observers.

With a baseline in hand, a next question is "how does the current state compare to the baseline?" Rather than thinking in terms of "health", a focus "fitness" is appropriate – of the Internet and its component parts, against selected ranges of expectation, and for different specific purposes.

Nevertheless, it is very hard to make broad observations or measurements to provide answers to such questions. The Internet achieved its global reach by facilitating the inclusion of independent networks, and each network is designed and operated based on local needs. There is no single ideal Internet; there isn't even a single ideal for the individual network that makes up part of the Internet. Eric Jardine's paper on the normalized perspective of cybersecurity risks paints a pretty clear picture of how wide a variance there is in published data when

---

[1] *This paper was written as part of the "Network Operator Measurement Activity" of Thinking Cat Enterprise LLC's TechArk project (http://www.techark.org/noma), which is partially funded by the Internet Society.*

[2] *http://www.techark.org/wp-content/uploads/2016/12/20161208-NOMA-SurveyPaper.pdf*

estimates (about size of the Internet, number of security attacks) are compiled from external sources.[3]

To build effective understandings of the Internet's fitness (to various purposes), it's time to stop guessing and start measuring the Internet infrastructure in a distributed yet collaborative fashion.

## Observe directly, Share collaboratively

The most accurate representations of any aspect of the Internet come from the operators of its networks and infrastructure services. They see the network traffic, and are in a position to observe, measure, count and otherwise reflect features of the traffic, the state of the network (links and connections), and the servers that support the infrastructure services. In many cases, the observations that would most usefully contribute to building an Internet fitness report are ones that provide important operational and business insights for the operators themselves. That gives them a business reason to make the observations on a regular and ongoing basis.

In order to create a coherent picture of the whole elephant, some reflection of these observations and measurements must be shared externally to the operator's organization. Of course, the detailed information that can be so informative internally to an organization is likely also business-sensitive and may intrude on users' privacy.

Therefore, to be able to create a global fitness metric, individual organizations' observation and measurement contributions must be filtered to address basic concerns of business and personal privacy, before being shared publicly.

In summary, the important points in this observation/sharing framework are:

1. (Network and infrastructure) operators are in control of the observations in their own systems
    - And the selected observations must be ones that provide useful business and operational insights in order to motivate their instrumentation
2. Any shared data must be filtered before public presentation
    - And the resulting fitness measures must be useful to the operators making the observations in order to motivate their participation (e.g., by building a credible benchmark of performance, they can avoid being subjected to imposed, imperfect "guestimates" from external sources)

[3] *Eric Jardine, "Global Cyberspace Is Safer than You Think: Real Trends in Cybercrime" , GCIG Paper Series #16, July 2015. https://www.cigionline.org/sites/default/files/no16_web_1.pdf*

## A Bigger Picture

The section above outlines an approach and requirements in order to be able to make direct observations available publicly through collaboration. However, individual measurement types don't tell a story in and of themselves. Individual observations can be used as "stand ins" (metrics) for some aspect of fitness that should be tracked, shared, and observed over time. This is similar to human (physical) fitness measurements: tracking resting heart rate is one useful metric, but it doesn't necessarily give the whole picture of a person's readiness to undertake running a marathon, for example.

In the context of the Internet, in the long run, the framework above can be used to develop composite metrics establishing baselines for, and determine progress on, such things as:

- Function/performance of the Internet
  - Speed, access times to popular websites
- Growth of the Internet/its use
  - Number of users, amount of traffic
- Attacks on infrastructure
  - Size, scale and impact
  - Malicious interventions, normalized to the size of the Internet
- Stability
  - Of the network itself, in terms of actual usage patterns

## Conclusions

The strategy to achieve the objective above must acknowledge that operators will only do what they get some kind of value from. Tactically, the approach of "self-instrumentation" can give network and infrastructure operators insight into what they need to know to improve their networks' performance and efficiency. Additional steps to achieve the objective include facilitating self-instrumentation and developing "support" groups. Then, benchmarking (e.g., using other external measurements activities) and lesser metrics may help get operators to consider sharing their actual views.

## Appendix: An Example – IPv6 Usage

A very timely question is "How much IPv6 usage is there?" The first thing anyone involved with measurements will say is: that's not well-formed. What do you mean by "usage"? And in what scope – globally, per country, per capita? And so begins the first stage of approximation – not asking the question that is actually on your mind, but rather asking a question for which you might be able to obtain an answer.

To be able to talk about the state of the Internet, and its progress, a first step is to determine some form of *metric* that reflects something about reality.

One metric is to look at the amount of traffic that the world's top (IPv6-capable) content providers see over IPv6. In fact, Google does collect that information and provides a per-country analysis[4]. While quite useful, that's pretty large scale, and it doesn't reflect how much IPv6 capability there is in networks (for doing things other than web searches).

Another metric would be to look at the number of IPv6 packets within access networks. That data is available to network operators. It can be aggregated to geographic units much smaller than countries. For individual operators, this information provides concrete data about the viability and importance of IPv6 in their networks. If collected and shared, a global picture of IPv6 usage over time could be built.

[4] *https://www.google.com/intl/en/ipv6/statistics.html*

## Appendix: Network Operator Measurements Activity (NOMA)

The example above is pretty simplistic, but it gives a sense of what can be built from individual operators' efforts to self-instrument, with some data shared collaboratively.

The NOMA activity is working on a pilot project to demonstrate the feasibility and value of this approach, with a view to expanding to define and undertake collaborative collection for other metrics as outlined above.

In an ideal world, the 5 year objective for NOMA can be articulated as follows:

- All major operators will be self-instrumented
- There will be a culture of shared Internet health data, and practices for instrumentation
- Up-and-coming network operators will have some (real world) benchmarks to shoot for in terms of their performance expectations
- There will be a multi-year repository of data about specific shared metrics that should show something about how the Internet has evolved and performance has improved over that time
- We'll have figured out how to apply the same pattern (collaborative observation and/or instrumentation and sharing) to address other questions of interest in Internet health. Take, for example, DNSSEC uptake. Or, numbers relating to unwanted traffic and mitigation.