Comments on the Identifier Technology Health Indicators (ITHI) Initiative

Steve Crocker, 23 January 2017

I write in my personal capacity, not in my role as chair of the ICANN Board nor as a member of SSAC. Normally, I would not enter the public dialog in a personal capacity, but on this particular topic I have spent a lot of time in the past, partly as prior chair of SSAC and partly in my professional life outside of ICANN, and I was urged to share my views and to do so as part of the community.

None of what I'm writing here should be taken as representing the view of the ICANN Board nor as an indicator of potential Board action. This project has not been on the Board's agenda, nor do I expect it to become so.

I very much like the idea of listing several issues related to the domain name system and tracking them over time. I think the notion of "disease" may be a bit strained and that not all of the issues fit comfortably into this metaphor. Attempts to put metrics on the health stretches the metaphor even further, and doesn't work very well for some of the issues. On the other hand, the template of definition, symptoms, causes, risk factors, complications, impact and potential treatment is useful.

Further, my own list of issues numbers a bit more than five, which I offer for consideration. Aside from whether my own list of issues matches anyone else's, the main point of my comment is the use of metrics to indicate health will work in some cases but not in others, and I think it's more important to treat each of the issues on its own terms. Force fitting metrics onto them may not lead in a useful direction.

Issues related to the Root

The Root includes the process for adding or modifying entries in the root zone, and the process for responding to queries to the root servers. The first is carried out by ICANN/PTI and Verisign; the latter is carried out by the root server operators. The entire operation has been extraordinarily reliable over the lifetime of this operation, so by any measure one has to say its health is good. That said, it's prudent to examine the root service in detail to anticipate potential problems.

- Confusion of top level name spaces

  Names that look like domain names but are not part of the global Internet domain name system often leak onto the public net. New names and new name systems come into existence periodically. There isn't any cooperative, effective agreement or control on these. As a result, there are both

security hazards and political controversies. The IETF has some partial processes related to this topic, but nothing is definitive so far.

Metrics? I don't know of how to assign a health index to this.

- Controversy over confusability of top level names

  Even within the ICANN administered space of top level names for the root, there are controversies as to when a name is or might be confusable, visually or otherwise, with other names.

  Metrics? It would be wonderful to have a meaningful way to measure confusability. I don't think we have it yet.

- Accuracy and timeliness of root zone updates

  We have a lot of data on root zone updates. The process has been nearly perfect over close to two decades. The process has also been reasonably prompt, though details on how long it takes to update a root zone entry have been partial.

  Metrics? Yes, timeliness and accuracy are easily measured.

- Trust in the root zone update process

Despite the essentially perfect record in maintaining, updating and publishing the root zone, some ccTLD operators or some governments, remain concerned that adverse changes might be made in the future. This leads to the question of trust.

Metrics? Is it possible to quantify trust in this situation? Probably so. Not by measuring anything about the actual operation, but by polling the ccTLD operators or their governments.

- Root Server operations

  The root zone is published through thirteen distinct constellations of root servers. Is this service responsible, accurate and reliable? Yes, and these attributes are measurable.

  A separate question is how strong is the root server system when faced with denial of service attacks. There is a fair amount of data regarding past attacks, and the root server system has performed well. What is its capacity to perform well in the face of future attacks, and how massive are future attacks likely to be.

  Metrics? All of the questions above can be quantified and measured.

Issues related to the ccTLDs and gTLDs

- DNSSEC deployment across the ccTLDs and gTLDs

  The root was signed in 2010 and a large number of of the TLDs are now signed and accepting signed registrants. The Internet Society, as part of its Deploy360 programme, publishes weekly the current status of DNSSEC deployment across the ccTLDs.

  Metrics? Yes, the deployment status of DNSSEC across the ccTLDs is easily measurable. (The same is true for the gTLDs. All gTLDs that are part of the current round of gTLDs must be signed, so there are only a few of the legacy gTLDs that are not yet signed.)

- Financial health of the ccTLDs and gTLDs

  Each TLD operator is a separate business operation. As might be expected in any set of businesses, some are in better shape financially than others. Service levels, reliability or even continued existence might be at risk in the financially weaker TLDs.

  Metrics? Financial health is typically measurable via credit ratings and similar measures.

- Security of the TLD operation

  Each TLD operation holds the data from its registrants and publishes that data through its name servers. Instances have been reported of a TLD operation being penetrated, resulting in either unauthorized changes to some of the entries or loss of functionality of the entire TLD operation.

  Metrics? There are two sets of metrics applicable here. One is the instances of penetration or disruption of service. The other is the potential for measuring the strength of the operation in comparison to industry best practices.

- Complaints in the TLD marketplace

  Some registrants have bad experiences dealing with registrars or registries. Sometimes registrants have lost track of the registration details or failed to renew their registrations. In other cases the registrar or registry has not protected the registrant. Both of these result in complaints. Measuring the number of complaints each year is relatively straightforward, but interpreting the meaning of those complaints is more difficult.

  Metrics? Yes, but assigning a meaning to the metrics requires some care.

- Trust in the TLD marketplace

  The above addresses the actual experience of the registrants. "Trust" is presumably related but reflects that the registrants and future registrants expect when they enter the marketplace.

  Metrics? Trust is measurable but has to be done through polling or other social science processes.

- Utility of the registration data

  Much has been said about the accuracy of whois data. What's implied but not said explicitly is that whois data is used by others, i.e. not the registrants, registrants or registries, to find out who is responsible for a particular registration and, usually, to contact that party. The contacting parties might law enforcement agencies, intellectual property plaintiffs, potential buyers, stalkers or others.

  We do not yet have a strong model for what constitutes healthy registration data and healthy use of that registration data.

  Metrics? This is an area where metrics associated with utility as opposed to accuracy or availability might be helpful.

System-wide issues

- DNS filtering

  Countries and other parties seem to be imposing filters and presenting non-uniform views of the domain name system. How widespread is this, and is it really growing? What impact will it have?

  Metrics? Is it possible to measure the degree of fragmentation or other effect of DNS filtering? What would such a measure mean?

- DDoS attacks

  DDoS attacks on DNS service providers seems to be a growing threat.

  Metrics? Is there a way to measure the threat? The degree of robustness or resilience in the face of an attack? The likely damage due to an attack?

- DNS surveillance

  DNS queries contain a lot of information about the interests and attention of the querying parties. To what extent is this information being collected and used? By whom and for what?

Metrics? I have not seen any metrics associated with surveillance. It may be hard to come to agreement on whether this is measurable, whether it should be measured, and how to interpret such measurements.

- DNS software reliability

  There are relatively few sources of DNS software and much of it is provided for free. The groups that produce the software are often underfunded. Much of their attention is on raising funds through donations or service contracts, and their donors often apply pressure for new features as opposed to increased reliability.

  There is a potential for a broad scale failure due to a flaw in one or more of the widely used DNS software packages.

  Metrics? Is there a way to quantify the risk involved and to relate that risk to the resources applied to the software development and distribution process?

The above is a personal catalog of DNS issues, and I cannot claim it is complete. My main point, as I

mentioned above, is concern that attempting to fit all of these into a single template, and more particularly an attempt to assign meaningful metrics to all of them may not work out once the details are examined.