

Net.info

A proposal for making network
service information easily available

Steven Bauer

Slides from 2010

MIT

Problem

- No easy way to identify network service information
 - Initial motivation is to make very basic network information easily available (e.g. contracted upload and download speeds)
 - Lots of potential information, but lets start very simple

Why is this important?

- Increasing number of studies and tests need this information
 - Comscore
 - Samknow's Ofcom study
- Attempts to infer this information from measurement data is problematic
- Customers don't know it and often have a hard time finding it when they do want to know it
- Details of service tiers change over time
- Revealing testing panel to service providers to have them manually do the tier identification is methodologically worrisome
- **We are wasting time repeatedly solving this problem in sub-optimal ways**

Now is the time to develop a convention for making network service information easily accessible both to human users and software at the edge

Net.info

- An ICANN reserved DNS domain so it may be available for use as a coordination point
- “Info” stands for *information* in about 37 languages, and is a neutral name

Net.info use case

1. A test of a broadband connection is run
2. Upon test completion, http get of net.info is issued by client
3. Net.info server redirects request to provider (e.g. <http://net.info.csail.mit.edu>) based upon client IP address
 - Actual URL conventions subject to community consensus
4. Information returned is specific to client IP address (e.g just like http://youtube.com/my_speed)
5. Integrate and record provider response data with test data
 - Vast amounts of raw test data is being collected... trivial to collect just a bit more from the providers as well

Provider net.info data

- Upload and download speeds (bps) at time of request
- Powerboost upload and download speeds if applicable
- Volume cap information if applicable
- Distance to DSLAM if applicable
- **No pricing information, just simple technical data about the network connection**

This can be more than just a boring burden for providers...

Providers should see this as a way to inject correct, novel, and informative data into the edge based testing datasets.

Provider net.info data

2. Network “traffic alerts” (entirely optional, but maybe really interesting data)
 - “FAP is temporarily in-effect.”
 - “You are currently over your volume limit so your traffic is de-prioritized.”
 - “Pardon our performance, we are currently experiencing abnormal conditions and are working to resolve the problem. Thank you.”
 - “Your network is scheduled for an upgrade, you should see improved performance soon.”
 - “Upstream provider is experiencing traffic delays.”
 - **Let providers expose information *if they choose and in the manner they choose***

Traffic alerts allow providers should to inject “their side of the story” when they still have information easily available that could contribute factually relevant information to understanding *why* a measurement result occurred.

Net.info benefits

- Minimizes time-of-test to time-of-service-identification errors
- Multiple tools/studies can leverage information
- Making net.info generally available eliminates one method providers could use to identify specific test panels (e.g. Samknows)

Net.info benefits

- Establishes a communication channel between provider and client
 - Other information could potentially be returned e.g. see <http://tools.ietf.org/html/draft-livingood-web-notification-08>
 - **“You went to the doctor to talk about one thing but while there learned of a completely unrelated more serious problem.”**

Provider incentives on traffic alerts

- Some ISPs already provide such information to some customers and/or such information is sometimes available to telephone support personnel.
- Incentive to only post traffic alert when a provider suspects or knows of an issue with their network.
- Providers would presumably not want to have the aggregate data show they are *always* claiming to be “temporarily” experiencing problems.

How unstructured traffic alert data could be useful to large scale measurement studies...

- Analyze performance under different subsets of data
 - No traffic alert data
 - Traffic alert group data
 - Traffic alert group data subsets (formed by key word searches or manual inspection).
 - Maybe conventions develop over time or providers contribute their own preferred groupings.
 - This is something that has to evolve in the community.
- **The point is this is useful even if it doesn't have a complex ontology.**

Privacy issues

- By default, IP address specific net.info information will be available to clients coming from an IP address without any restrictions (just like http://youtube.com/my_speed)
- Potential privacy problem: devices (e.g. Apple or Google) may survey customers net.info data using their phones or other devices
- Customers can prevent this by opting out or restricting access (perhaps with passwords) by simply following links on their provider's net.info site
- Re-activation of net.info information without restrictions would require solving a captcha or logging into an account (any action that would require a human)

Implementation notes

- Net.info is an ICANN reserved name so it is (potentially) available
- Restful API required
- Format data in multiple formats (e.g. HTML, RDF, XML, text) for easy consumption in different usage scenarios by programs or a human user

Net.info potential next steps

- Talk to W3C folks about knowledge representation
- Talk to privacy experts
- Solicit feedback on idea from:
 - broadband providers
 - Samknows/FCC/M-labs
- ***Identify holes, gauge how challenging this really is***
- Talk to ICANN, understand reserved DNS names issue
- Build a demonstration prototype of service