The Road to an Open Internet is Paved With Pragmatic Disclosure & Transparency Policies

Bill Lehr
MIT
wlehr@mit.edu

Erin Kenneally
CAIDA/UCSD
erin@caida.org
(corresponding author)

Steve Bauer
MIT
bauer@mit.edu

Table of Contents

1. Introduction

Ensuring a healthy ecosystem for broadband services is critical to securing the future of a vigorous and open Internet. From the perspective of social welfare maximization, this means collective management of the decision-making regarding how we design, operate, provide access to, use, and pay for our broadband access networks.[1] Realizing this collective goal requires balancing the interests of multiple market participants[2] that are often in conflict and evolve in light of changing technical, business, and policy conditions.

The efficiency of markets and regulatory interventions depends on whether decision-makers at all market levels are appropriately informed. This requires the selective sharing of information. Consumers need information about their broadband access options in order to make informed decisions about which (if any) broadband services to subscribe to, how to use those services, and what investments to make in complementary assets (devices, content, applications). Providers of content, applications, and other complementary goods and services need to know about broadband access options to appropriately position their offerings in the market. ISPs need information about usage to appropriately design and provision their networks. Moreover, since end-to-end performance depends on the decisions of multiple stakeholders, market participants need information about how content and application edge providers are provisioning their services and the impact that those decisions have on networks. And, regulators need information about broadband access options and edge provider practices to design and enforce policies that will promote competition and ensure appropriate market choices exist.

All of these stakeholders need information about broadband service availability, pricing, performance, and to the extent discernible, about trends and plans that will shape future choices. Notably, the providers of broadband access services ("ISPs" hereafter) either already are in possession of a significant portion of the information or are in a better position to acquire the information needed by market participants to render effective decisions. However, the information sharing challenges are far from simple. Stakeholders' information needs vary, information is costly to collect and share, and information that matters for decisions has strategic value. For example, better informed consumers might be more inclined to switch providers, thereby intensifying price competition; while better informed regulators may be better able to limit supra-

[1] An earlier version of the paper referred to "reasonable network management" in the title. We have renamed the paper to signal that what we mean by management is much broader, engaging all market participants, rather than what is meant by the narrower use of the term *Reasonable*

[2] The market participants, or equivalently, stakeholders that we will focus on here include the ISPs, edge providers, consumers, and regulators.

competitive profit opportunities.[3] Additionally, sharing too much information to certain parties about the performance of specific broadband connections might threaten subscriber privacy or render broadband networks more vulnerable to attack.

Disclosure and Transparency ("D&T") is an umbrella term used here to encompass a toolset of often interrelated policies, rules, processes, practices and mechanisms (e.g., performance testing platforms) that are used by market participants to help structure and manage the flow of information that is needed for informed decision-making.[4] D&T policies comprise a significant component of the regulatory provisions in the FCC's 2015 OIO, which sets forth the FCC's approach for regulating providers of broadband access services.[5] These OIO D&T provisions are neither without interpretation challenges nor are they self-executing but they do induce other *intervention* tools-- market and regulatory practices and mechanisms-- that shape how broadband management relevant information is discovered, shared, and interpreted. This paper's focus, therefore, is on providing a framework within which to compare the relative effectiveness of the D&T interventions in translating the OIO's D&T policies within the larger market context and relative to the OIO's objectives.

As we shall explain, having a rich portfolio of D&T tools is desirable in order to address the diverse and complex questions that require information sharing. Moreover, understanding how these D&T tools interact and complement (and/or substitute) for each other is helpful if these they are to be appropriately applied and appreciated. Application of these tools should be nuanced and evolvable to incentivize cooperation and voluntary disclosure by the ISPs while also safeguarding the interests of end users and intermediaries in the broadband Internet ecosystem.

In Section 2, we review the specific D&T provisions in the FCC's 2015 OIO and situate these within the larger D&T policy framework. We introduce a meta-tool, the D&T Coordinator, to assist in better understanding the landscape of potential interventions and with which to contrast the relative merits of different interventions in different contexts.

---

[3] While the focus of much of the debate and current regulations is on limiting the potential market power of broadband access providers, this is not the only possible source of market power that the FCC may need to address. For example, edge providers or providers of other complementary services might also be found to have market power (e.g., providers of scarce programming content, search services, or end-user devices). Moreover, abuses of potential market power are not the only market imperfection that might justify regulatory interventions. Public goods problems and coordination problems stemming from asymmetric information are also relevant challenges that might justify regulatory interventions that would address information sharing incentives and practices among stakeholders.

[4] The focus here will be on the information that needs to be shared in the context of network management (as broadly construed herein); however, the D&T tools are also relevant and used for other regulatory policy goals and same sort of analysis explicated here applies more generally.

[5] For full reference, see Note 1 *supra*.

In Section 3, we apply our framework to two prototypical examples of the sorts of questions that confront the challenge of how to best manage broadband networks. At one end, we have what appears to be the narrow and specific question of crafting an appropriate set of D&T mechanisms to ensure adequate reporting of packet loss by ISPs. At the other extreme, we consider open-ended questions that relate to society's aspirations or goals for what the Internet and broadband services should be. We argue that an assortment of D&T tools are needed for the array of questions confronting broadband stakeholders, but with different emphasis, because the contexts within which they arise engage both specific and general, closed and open-ended details to be addressed appropriately.

Section 4 offers our concluding summary and directions for future work.

2. Framework for Analyzing D&T Policies

The FCC's 2010[6] and 2015 OIOs identified five core objectives: (a) maintaining an open Internet for the purposes of ensuring that consumers have informed choices regarding broadband Internet access services; (b) that application and content providers can engage productively in the marketplace; (c) that ISPs will continue to promote a healthy Internet; and (d) that regulators can assess and intervene to protect and promote these goals.

The Disclosure and Transparency (D&T) policies that are the focus of this paper play a key role in helping ensure these goals may be realized. In the following sub-sections, we provide a general overview of what we mean by D&T polices and review how these are addressed in the OIOs and some of the problems with interpretation that arise. We also introduce an analytic framework that is useful for understanding the capabilities, limitations, and interaction effects of different D&T policy interventions.

2.1. What are Disclosure and Transparency Policies

As we explain more fully in following sub-sections, D&T policies include a diverse set of policy tools or mechanisms that specify how information is shared among industry stakeholders. Hereafter, we follow the FCC, and focus on four classes of relevant stakeholders: ISPs,[7] edge providers, consumers, and regulators.[8]

---

[6] See FCC (2010), *Report and Order*, In the Matter of Preserving the Open Internet (GN Docket No. 09-191) and Broadband Industry Practices (WC Docket No. 07-52), adopted December 21, 2010, available at https://apps.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf (hereafter, "FCC 2010 OIO").

[7] For consistency and simplicity purposes we use "ISP" herein to refer to "broadband Internet providers." More generally, the term ISP may refer to providers or transit or peering services, and are interconnected to and upstream of the access ISP that provides the broadband Internet service. For our discussion here, we would class these non-access ISP provider services along with Edge Providers (see next note).

The D&T policies affect how information is shared along five interrelated dimensions:

- *Why* – to what problems is data disclosure a solution?
- *What*- what classes and/or types of data address those problems?
- *Who* - what entity has the data and should disclose it?
- *How Used* – for what purposes may the disclosed data be used?
- *How Disclosed* – what are the processes and mechanisms to make the data available (reporting, retention, accessibility, evidentiary proceedings, etc.)?

D&T policies provide a toolset with which to address asymmetric or imperfect information problems that might otherwise result in market imperfections such as abuses of market power, coordination failures, and the like. This toolset includes regulatory codes that establish disclosure obligations, but also performance measurement and reporting platforms and practices, and processes for enabling multi-stakeholder participation in decision-making.[9] D&T policies may be explicit as when D&T provisions are specified in regulatory codes, as well as implicit. The latter may occur in the case where other aspects of the regulatory framework (e.g., other portions of code or institutional aspects) necessarily interact with and impact D&T interpretations or applications. Finally, it is worth noting that all of the stakeholders mentioned above may play a role in the design and operation of D&T policies. For example, industry self-regulation and end-user-based measurements are examples of bottom-up, market-based actions that may complement or substitute for government-initiated D&T policies.

The full range of D&T tools includes FCC orders and consent decrees, transparency reports, market research reports, and consumer complaints.[10] Moreover, new disclosure capabilities may arise from edge measurement tools like *net.info*,[11] or be adapted from models such as the Key Facts Indicator,[12] ISP Censorship Transparency Reports,[13] and FCC NORS[14] reporting.

---

[8] The FCC uses Edge Provider "to refer to content, application, service, and device providers, because they generally operate at the edge rather than the core of the network." Consumers are included in the class of "End-Users," which also includes "any individual or entity that uses a broadband Internet access service" (See 2010 OIO, footnote 2).

[9] This list is illustrative and not intended to be comprehensive, as we make clear below.

[10] See https://consumercomplaints.fcc.gov/hc/en-us

[11] *net.info* will be discussed further below. It refers to a tool that we are investigating that is intended to facilitate better information sharing between ISPs and individual broadband subscribers.

[12] In 2011, a number of the larger ISPs in the UK, organized as the Broadband Stakeholder Group (BSG, http://www.broadbanduk.org/), agreed to publish a common Key Facts Indicators (KFI) that describe their traffic management practices (see http://media.ofcom.org.uk/news/2011/improving-traffic-management-transparency/; or Figure 1 in http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/statement/statement.pdf). For links to the KFI disclosures of the participating ISPs, see http://www.broadbanduk.org/bsg-openinternettrafficmanagement/trafficmanagementkfis/.

2.2. FCC Open Internet Order and Disclosure and Transparency Policy

D&T provisions were a core component of the FCC's 2010 OIO. *Transparency* was one of the "three basic rules" that the FCC asserted were "grounded in broadly accepted Internet norms, as well as" the FCC's "prior decisions."[15] In the 2015 OIO, the FCC reaffirmed and expanded the explicit provisions establishing its D&T policies under the Order (see paragraphs 23-24, 109, 154-175). In addition to these, other portions of the 2015 OIO are also relevant, including:

- Disclosure process *Safe Harbor* (in paragraphs 176-181), which asserts that ISPs that voluntarily disclose information that complies with the consumer disclosure format (to be proposed by October 2015) by the Consumer Advisory Committee (CAC) will be deemed to be in compliance with the transparency rules;
- *Bright Line* rules (in paragraphs 14-19, 104-107, 110-132) that prohibit ISPs from traffic blocking, throttling, or paid prioritization of lawful content, applications, services, or devices;
- *General Conduct Standard* (in paragraphs 20-22, 108, 133-145) that prohibits engaging in practices that would unreasonably interfere with or disadvantage Consumers or Edge Providers, which imposes a public interest test standard; and,
- *Reasonable Network Management (RNM) exception* (in paragraphs 32-34, 69, 85, 214-224) that establishes the basis for exceptions to the bright line rules on reasonable network management that otherwise prohibits traffic throttling or blocking.[16]

---

[13] For example, for Google's self-disclosure of its transparency policies, see http://www.google.com/transparencyreport/removals/government/. There are also third-party initiatives to publish information on transparency policies (https://transparency.automattic.com/) and to stimulate more such activity (https://www.eff.org/deeplinks/2013/01/its-time-transparency-reports-become-new-normal). Many on-line service providers, including ISPs disclose many of their policies as part of their posted Acceptable Use Policies (AUPs). As examples, see AT&T (http://www.corp.att.com/aup/); Time Warner Cable (http://help.twcable.com/twc_misp_aup.html); Amazon Web Services (https://aws.amazon.com/aup/); Level 3 (http://www.level3.com/en/security-law-enforcement-and-acceptable-use-policy/acceptable-use-policy/); Verizon (https://www.verizon.com/about/terms/); or Comcast (http://www.xfinity.com/corporate/customers/policies/highspeedinternetaup.html). Although there are many commonalities in these various AUPs (e.g., statements prohibiting use of broadband service to infringe copyright), the reports are not standardized.

[14] See https://www.fcc.gov/nors/outage/StartUp.cfm.

[15] The FCC identified "Transparency" as the first of its three rules, specifying that this mean that "Fixed and mobile broadband providers must disclose the network management practices, performance characteristics, and terms and conditions of their broadband services" (see 2010 OIO, paragraph 1). The other two rules were "No Blocking" and "No Unreasonable Discrimination." It is worth noting that while the second two rules were successfully challenged in Court, the Transparency rule was affirmed by the Court.

[16] FCC 2015 OIO paragraph 21, 32-34, 69, 85, 214-224.

As we will explain further below, how these sections are interpreted and applied will directly and indirectly impact the application and interpretation of the D&T policies. Each of these elements have associated gray areas that resist a simple, one-size-fits all interpretation or application. For example, the OIO provides guidance but remains ambiguous with respect to determining which business practices might be determined unreasonable under the conduct standard. And the *Bright Line* rules and *RNM exception* raise interpretation challenges that require a context-based and nuanced application of D&T tools.

While the OIO provides the raw materials for D&T policies, it does not provide a blueprint for how these should be applied to the myriad real world scenarios that demand multi-dimensional considerations. To this end we introduce a conceptual tool - the D&T Coordinator (Figure 1) - to facilitate comparing D&T policies in alternate decision-making contexts. The D&T Coordinator is intended to assist in understanding and ensuring that appropriate D&T policies are in place to ensure transparency for consumers and intermediaries regarding network management practices, performance, and terms; to help ISPs avail themselves of the safe harbor and reasonable network management provisions; and to arm regulators with empirical data to enforce D&T policies.

2.3. Confronting Issues with the Order's Rules and Principles

The FCC 2015 OIO is a hybrid regulatory framework comprised of both rule and principal-based provisions. The former specify proscriptions and prescriptions, while the latter express more general goals. The allure of principle-based provisions -- flexibility, universality, and discretion— comes at the risk of ambiguity in how the regulations will be applied in practice, uncertainty regarding regulatory outcomes, and an increased need for ex-post remediation.[17] Similarly, the advantages of its rule-based components – ex ante compliance specificity and certainty – may render regulatory decision-making rigid with reduced capacity to adapt as markets and technologies evolve. We contend that for both types of provisions related to D&T, there are decision, application and evaluation gray zones that warrant intervention tools to address ambiguous and emergent interpretations.

The so-called *bright line rules* are manifest in provisions such as the Transparency Rule requirements that ISPs disclose network practices,[18] characteristics,[19] and commercial

[17] Burgemeestre, Brigitte, Joris Hulstijn, and Yao-Hua Tan. "Rule-based versus Principle-based Regulatory Compliance." *JURIX*. 2009 (Available at http://homepage.tudelft.nl/w98h5/Articles/jurix.pdf).

[18] FCC 2015 OIO at ¶215: "A network management practice is a practice that has a primarily technical network management justification, but does not include other business practices. A network management practice is reasonable if it is primarily used for and tailored to achieving a legitimate network management purpose, taking into account the particular network architecture and technology of the broadband Internet access service."

[19] FCC 2015 OIO at ¶165.

terms;[20] and the prohibitions on blocking, throttling and paid prioritization for broadband access.[21] For example, the OIO requires specific and detailed performance disclosures for users and edge providers, such as actual speed, latency, and packet loss.[22] For network practices, they mandate disclosure of congestion management, application-specific behavior, device attachment rules, security, and practices that are applied to traffic associated with a particular user or user group (including any application-agnostic degradation of service to a particular end user).[23]

Yet, metrics for these performance characteristics are far from standardized or settled. For example, these *bright line rules* appear less bright when one recognizes that the OIO does not specify how loss should be measured or acknowledge that different measurement methods for loss yield different answers.[24] How long may a packet be delayed before it should be included in packet loss?[25] Or, how a consumer might use a report of ISP loss rates associated with the ISP's router on the customer-side of an interconnection link to inform the customer about the effect of upstream congestion on the consumer's experience?[26]

In addition to providing guidance on the content ("what") that must be disclosed, the OIO enhances the prescribed format or means ("how") of disclosure to include "a requirement that broadband providers notify end users directly if their individual use of a network will trigger a network practice, based on their demand prior to a period of congestion, that is likely to have a significant impact on the use of the service."[27] Here again, the precondition to this disclosure introduces a value judgment that renders ambiguous the application of the bright line rule.

A plausible scenario giving rise to bright line rule conflict is if a subscriber's edge behavior causes an ISP to throttle/block traffic in congestion circumstances that are, in part, the result of an Interconnection Agreement (IA) between the broadband service provider and an upstream ISP or Edge Provider.[28] Does the OIO requirement for

---

[20] FCC 2015 OIO at ¶164.

[21] FCC 2015 OIO at ¶110; 8.5, Order at ¶8.7, 8.9

[22] FCC 2015 OIO at ¶168.

[23] FCC 2015 OIO at ¶169.

[24] Infra §3.1.

[25] Timeouts: Beware Surprisingly High Delay:
http://www.caida.org/workshops/aims/1503/slides/aims1503_nspring.pdf.

[26] Of equal importance, current measurements don't capture loss on the other side of interconnection links by upstream ISPs or edge providers. Discussed infra §3.1.1.2, in routers queues generally build on the outgoing interface.

[27] 2015 OIO paragraph 156. See more generally paragraphs 154-175, and 169.

[28] The IA is likely conditioned on aggregate user traffic, rather than the traffic of a particular subscriber. Users with Gbps access could impact aggregate traffic loads.

notifying customers of "network practices" therefore require the broadband service provider to disclose the details of the IA?

Generally, Internet interconnection agreements are considered commercial arrangements that the FCC refrained from subjecting to the Order.[29] These IAs are typically the result of private bilateral negotiations that are regarded as confidential by the parties. The FCC might conclude that the IA is a business practice, rather than a (technical) network practice, and hence exempt disclosure of IAs from the notification requirement. However, such an interpretation might allow ISPs to circumvent the intent of the disclosure notification requirement by inserting practices they would like to avoid disclosing in IAs.[30] Figuring out what constitutes a business versus a technical practice will not be easy and so partitioning what portion of IAs might have to be disclosed and what might remain confidential will be difficult, and likely, may require case-by-case consideration.

The argument could then ping-pong to the claim that the OIO disclosure trigger—"significant impact on use"[31]-- only applies to a network practice that is "caused by" a user's individual action at the edge, not to an action that is triggered by terms of an IA. The IA behaviors are determined on the basis of aggregate traffic that flows across the

---

[29] The 2015 OIO asserts the FCC's authority to regulate Interconnection, but does not seek to actively regulate Interconnection under the OIO today. For example at ¶28, the FCC states that the sale of broadband access service to retail customers with the representation "that they will be able to reach 'all or substantially all Internet endpoints' necessarily includes the promise to make the interconnection arrangements necessary to allow that access," and note that in classifying broadband access as a telecommunications service under Title II, "commercial arrangements for the exchange of traffic with a broadband Internet access provider are within the scope of Title II" (¶29). The FCC explains its "watch and see" approach regarding interconnection regulations as follows: "Three factors are critical in informing this approach to interconnection. First, the nature of Internet traffic, driven by massive consumption of video, has challenged traditional arrangements—placing more emphasis on the use of CDNs or even direct connections between content providers (like Netflix or Google) and last-mile broadband providers. Second, it is clear that consumers have been subject to degradation resulting from commercial disagreements, perhaps most notably in a series of disputes between Netflix and large last-mile broadband providers. But, third, the causes of past disruption and—just as importantly—the potential for future degradation through interconnection disputes—are reflected in very different narratives in the record.... Thus, we find that the best approach is to watch, learn, and act as required, but not intervene now, especially not with prescriptive rules. This Order—for the first time—provides authority to consider claims involving interconnection, a process that is sure to bring greater understanding to the Commission."

[30] It is also possible that an edge provider might also use the confidentiality of IAs as a way to avoid disclosing their content or application provisioning practices, making it difficult for other market participants to identify the source of end-to-end performance problems that may arise.

[31] OIO at paragraph 171: "[we] ... enhance the rule to require a mechanism for directly notifying end users if their individual use of a network will trigger a network practice, based on their demand prior to a period of congestion, that is likely to have a significant impact on the end user's use of the service."

interconnection links. If it is aggregate traffic behavior that triggers the ISPs congestion management response under an IA, what obligation does the ISP then have regarding its disclosure requirements to the specific user? How is the ISP to separately identify the contribution of the individual's traffic to the aggregate traffic? The takeaway is that this illustrates how arguments can spiral in the face of information asymmetries. Disclosure practices that facilitate a more proactive expectation of transparency around network management behaviors can pre-empt costly adjudication of duties and obligations.

As for the principle-based provisions, the so-called *light touch* strategy shrouds the "reasonable network management" exception to the no -blocking and –throttling rules,[32] the no-unreasonable interference/disadvantage standard,[33] and the disclosure format and content safe harbor.[34] While on its face the OIO explicitly forbears application to Internet traffic exchange,[35] its assertion of authority to govern interconnection via the prohibition on unjust and unreasonable practices standard places interconnection within the OIO's discretionary authority.[36] For all of these provisions, the FCC plans to take a case-by-case enforcement strategy.

2.4. D&T Coordinator

Network management occurs amidst a backdrop of changing technologies, business relations, consumer demand, regulatory and market conditions, and other factors that impact choice, innovation and oversight. All of these may impact stakeholder decision-making as it relates to the OIO objectives cited earlier. Nuanced and flexible D&T policies are needed to respond to these challenges.

The D&T Coordinator is inspired by the risk assessment framework proposed in Coull & Kenneally (2012).[37] The D&T Coordinator is intended to aid in understanding how

---

[32] OIO 214-224

[33] OIO 133-149; 8.11

[34] OIO 176-181

[35] OIO at ¶194-206. Order at ¶195; see Order at ¶. 206 ("To be clear, we are not applying the open Internet rules we adopt today to Internet traffic exchange.").

[36] Order at ¶ 513: "The Commission retains authority under sections 201, 202 and the open Internet rules to address interconnection issues should they arise, including through evaluating whether broadband providers' conduct is just and reasonable on a case-by-case basis." Order at ¶. 205. See FN 525: "We observe that should a complaint arise regarding BIAS provider Internet traffic exchange practices, practices by edge providers (and their intermediaries) would be considered as part of the Commission's evaluation as to whether BIAS provider practices were "just and reasonable" under the Act.

[37] This model was inspired by and adapted from the Disclosure Control Framework, see, Coull, Scott E. and Kenneally, Erin, A Qualitative Risk Assessment Framework for Sharing Computer Network Data (March 31, 2012). 2012 TRPC. Available at SSRN: http://ssrn.com/abstract=2032315; and, Coull, Scott E. and Kenneally, Erin E., "Toward a Comprehensive Disclosure Control Framework for Shared Data", IEEE International Conference

different D&T interventions might relate to each other and fit with different decision-making contexts along three dimensions: (1) the nature of the disclosure target recipient (to whom is the data being disclosed), (2) the specificity of the disclosure content (what is being disclosed and potentially for what purpose or why), and (3) the temporal nature of the disclosure (when is the data disclosed). The D&T Coordinator provides a conceptual model for visualizing the comparative space of possible D&T tools or interventions, the diversity of questions to be addressed, and how multiple tools might address those questions. The D&T Coordinator may help in guiding the development and selection of tactical and strategic disclosure and transparency practices by providers; informing accountability and enforcement oversight by regulators; and steering and galvanizing public research that can inform evidence-based open Internet policy and intervention strategies. The exercise of mapping existing tools into the space defined by the D&T Coordinator may assist in understanding where additional tools may be needed and how tools may or should evolve over time as experience accumulates and markets evolve.

Although the current landscape is populated with multiple sources of information relevant to the management of broadband services, assessing the value of existing sources and how best to improve these or fill gaps with new sources and how to integrate and share the information requires multiple D&T tools. An underlying assumption is that disclosure (via bottom-up measurement or top-down reporting) should provide empirical, objective data that should be embedded into practice and translated into pragmatic decision-making in interpreting and applying the OIO's provisions. We currently face gaps in our knowledge about how networks are being managed and what the effects of those decisions are on stakeholders, how such understanding (or lack thereof) can or should inform stakeholder behavior, and what intervention strategies (techniques) can effectively protect Internet openness consistent with promoting innovation and investment.

Partially owing to the relative immaturity of the OIO, there is little consensus even within the various stakeholder communities regarding the details or general principals that might characterize best-practice standards for network management of broadband services, and the related D&T policies. For example, are traffic level metrics available to individual subscribers more or less effective than collective learning disclosure strategies? Can we distinguish the relative effectiveness for consumer protection between disclosure to regulators and disclosures targeting the public or third parties? What is a successful strategy for measuring congestion? What characteristics are common to successful strategies such as public measurement tools (net.info) and transparency reports for ensuring competitive markets? To what degree are strategies more effective in ascertaining what are appropriate broadband management practices when they capture data in real time at the event-level or longitudinal at the traffic-level? How might these

---

on Technologies for Homeland Security (November 2013, Boston, MA), Available at SSRN: http://ssrn.com/abstract=2326264. Its visual depiction was also informed conceptually by evidence-based policing research, see, Lum, Cynthia, Christopher S. Koper, and Cody W. Telep. "The evidence-based policing matrix." *Journal of Experimental Criminology* 7.1 (2011): 3-26.

insights guide the development and/or selection of disclosure strategies across different problems and contexts?

As a concrete example, consider the question of D&T related to the measurement of broadband speeds. The Measuring Broadband America ("MBA") program to be discussed further below relies on a purpose-built measurement infrastructure that relies on hardware boxes that are situated so as to enable isolation and flexible measurement of a representative set of broadband subscriber lines in place at a select group of the largest ISPs. The measurement data is aggregated, summarized, and published in a set of periodic reports issued once a year. However, the periodic reports are unlikely to provide much that would be of use to an individual end-users seeking to diagnose real-time broadband service issues. For such purposes, the end-users may choose to make use of speed measurements using web-based platforms, which offer a different set of cost/benefit trade-offs.[38] On the one hand, the web-based measurements allow subscribers to do real-time measurements and measurements for ISPs that are not part of the MBA program. On the other hand, the lack of standardization of web-based platforms and potential issues associated with attributing what portion of the results is due to the broadband access service makes interpretation of the web-based speed measurements difficult. Having both sets of tools offers multiple perspectives and those multiple perspectives can enhance the usefulness of both. Both the MBA and web-based speed measurement data provide longitudinal and cross-sectional data, but whereas the MBA data is designed to facilitate "apples-to-apples" comparisons across time and across ISPs, the web-based data provides a less systematic set of crowd-sourced data that can fill in gaps. The two samples provide (partially) independent tools for mutual cross-validation.

[38] For discussion of complexity and issues associated with measuring broadband speeds, see Bauer, Steven and Clark, David D. and Lehr, William, Understanding Broadband Speed Measurements (August 15, 2010). TPRC 2010. Available at SSRN: http://ssrn.com/abstract=1988332.

Figure 1: Disclosure & Transparency Coordinator

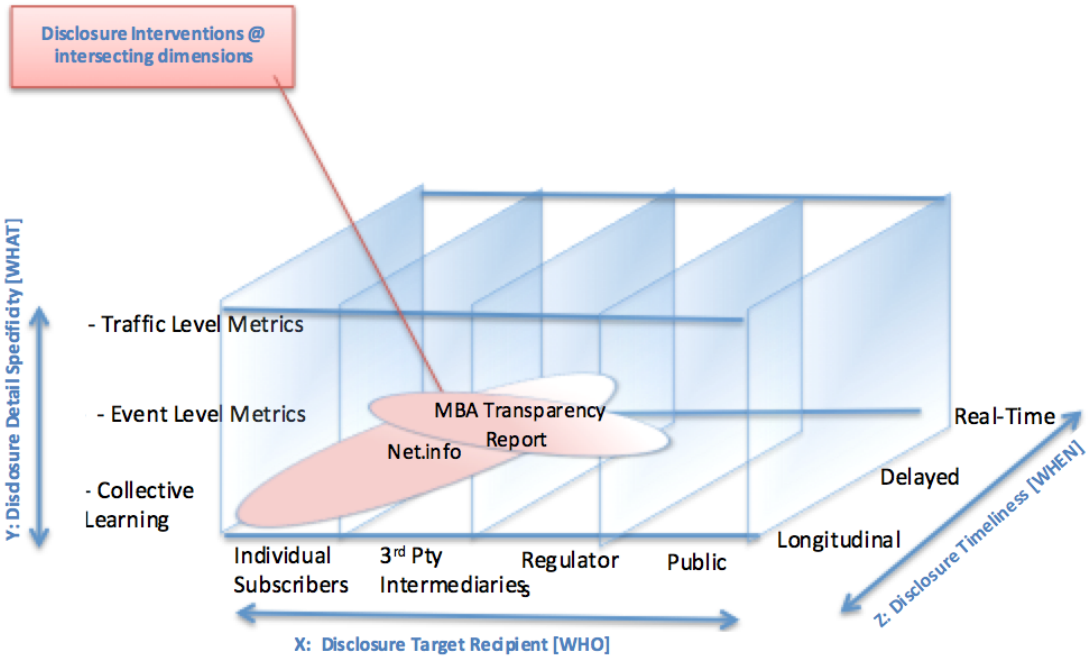

The X-axis focuses on the "WHO" dimension of the disclosure. The presumption here is that the ISP is the entity disclosing information, and the X-dimension characterizes to type of entity that is the principal intended beneficiary of the disclosure. The WHO might be individual end-users as might be the case if the ISP provided access to a real-time dashboard that enabled the subscriber to track his or her individual usage in real-time; to third-parties such as market research firms who might aggregate the data to produce summaries of focused performance studies; to a regulator (in this case, the FCC) which might use the data to inform and enforce policymaking; or the general public that might rely on the reports to inform its understanding about the overall broadband market.

The Y-axis focuses on the specificity of the data. This includes both "WHAT" is to be disclosed, and implicitly, also "WHY" it is to be disclosed. The conflation of WHAT/WHY questions highlights the fact that in order to assess whether information would be informative for decision-making, one has to understand what the decision is. This WHAT-axis is directly related to the disclosure purpose specificity such that open-ended disclosures are most often mapped to collective learning type data disclosures, and narrow decision questions such as whether congestion is causing performance degradation are more apt to require specific traffic-level data. Intermediate between the two may be event-level disclosures that aggregate multiple types of specific traffic-level metrics. For example, disclosures of information about broadband speed or packet loss (traffic-level disclosures) might be used to inform disclosures of measures of the quality of experience or service reliability (event-level disclosures) that might feed into our collective assessment of what constitute appropriate expectations for broadband services (collective learning).

The Z-axis captures the temporal dimension of the timeliness of the data disclosed. This Z-axis ranges from longitudinal to delayed to real-time. Real-time data may be useful for diagnosing real-time operational problems, such as consumer QoS demands, but may be more error prone and less valuable for comparing across service providers. The disclosure of delayed data may be less expensive to manage and offer comparatively greater descriptive and evaluative certainty. Longitudinal data, such as for macro level assessment, provides a basis for evaluating trends and may prove useful for forecasting.

2.5. D&T Interventions

While the regulatory foundations for the OIO give rise to ambiguities regarding D&T rights and obligations, they also assert the authority and provide the impetus for regulators and market participants to develop, evaluate and enforce D&T intervention mechanisms.

We generally concur with the light touch approach that FCC has taken in leaving the application of D&T to the discretion of the ISPs (self-regulation) and market forces; and with the FCC's nuanced use of multiple D&T tools to facilitate expanded information sharing among market participants. First, the FCC has asserted its authority to compel certain disclosures from ISPs to the FCC and potentially to the general public (even if the FCC's specific mandates today are limited).[39] Second, the FCC has mandated that ISPs are required to issue public "transparency reports" on broadband performance sufficient to enable consumers to make informed subscription/usage decisions. While the FCC is currently leaving it to individual ISPs to determine the details of such reports, the FCC has specified a safe harbor for compliance with this requirement. Third, the FCC imposes an obligation that all information disclosed by ISPs be accurate, including the duty to maintain the accuracy of disclosures when there is a material change to an ISP's terms, practices or performance. This last provision is akin to the accounting standard for financial disclosures. That is, an ISP might be liable for ex post penalties if it fails to disclose something that could reasonably have been anticipated to be material had it been disclosed.[40]

---

[39] Of course, market participants are likely to differ with respect to what constitutes a reasonably "limited" set of disclosure mandates. In the 2015 OIO at ¶ 166, the FCC identifies some quite specific disclosure mandates.

[40] OIO at ¶161 ("For these purposes a "material" change is any change that a reasonable consumer or edge provider would consider important to their decisions on their choice of provider, service, or application.") For financial disclosures, a materiality standard is applied to what information management is obligated to disclose. This "limits the information required to those matters to which there is a substantial likelihood that a reasonable investor would attach importance in determining whether to purchase the security registered." Finally, as with any 'reasonableness-based' standard, lawyers may disagree on the appropriate scope of its interpretation (e.g., whether it only applies to those matters expressly identified in the OIO or may be taken to apply more generally).

There is no one-size-fits-all set of D&T policies that would work or be appropriate. Sticks that impose D&T mandates, threaten more burdensome regulatory oversight, or limit opportunities to manage networks flexibly need to be balanced with carrots that induce voluntary compliance with D&T goals. Ultimately, the issue is whether stakeholders have the information they need, not precisely how they obtained such information. However, the HOW matters since it is related to the quality of the information and the stakeholders incentives to provide the information or trust its provenance. For example, we might expect end-user measurements that are understood to be subject to significant measurement error to be less accurate than truthfully revealed ISP performance reports; however, how do we convince stakeholders that the ISP reports are, indeed, truthful. The opportunity to have multiple perspectives and sources of information can contribute to the trustworthiness and verifiability of all of the sources. In light of this observation, the key to good D&T policy is to have a multiplicity of D&T tools. A mix of measurement tools and appropriate hard disclosure mandates for ISPs can complement and partially substitute for each other.

2.6. Transparency Reports: Measuring Broadband America

As for existing mechanisms, the measurement infrastructure and associated reporting that is rooted in the MBA exemplifies a type of transparency report that is helpful for translating and applying the ambiguous and emergent interpretations of specific D&T provisions in the OIO. With respect to the D&T Coordinator, MBA may be viewed as a disclosure intervention that is focused on traffic level data and event level measurements (WHAT-axis) that is primarily targeted for the FCC and third party intermediaries (WHO-axis), and that results in periodic (delayed) reporting of results (WHEN-axis).

Juxtaposed with top-down compulsion, transparency reports can facilitate disclosure that advances regulatory goals without overburdening ISPs, thus achieving a balance between competing stakeholder interests. ISPs that participate in the MBA are eligible for the safe harbor compliance with the OIO's *bright line* disclosure requirements.

The OIO's safe harbor provision with respect to MBA serves multiple purposes. First, it offers a response to those who argue that the OIO lacks adequate specificity and is too vague. Providing a safe harbor reduces regulatory uncertainty that might otherwise threaten ISPs. Second, participation in the MBA creates network effects, the results of which may include higher quality data and increased likelihood that the MBA can provide a nexus for consensus about appropriate industry metrics and reporting standards. Indeed, the fact that the MBA has been developed in close consultation with ISPs and other industry stakeholders, independent researchers, and the FCC lends it a degree of credibility and rigor that is lacking in some of the other measurement platforms.[41] Third,

[41] The MBA's design allows it to isolate performance measurements to the broadband access service provided by the ISPs since the hardware boxes are upstream of the end-user's premises network with measurement to servers on well-managed connections; and the MBA supports a rich and flexible set of performance measurement options. On the other hand, the MBA boxes are relatively sparsely deployed across the subscriber base and do not provide performance data for

it provides transparent contours for tests within which various measurement techniques can be compared.

Without the benefit of this safe harbor, and reliant solely on ISP initiated/designed reporting, there would be a greater risk that ISP reporting would be inconsistent, rendering cross-provider comparisons more difficult. There would also be an increased risk (real and imagined) that ISP reports would not be trustworthy.

Also, although periodic reporting is not as valuable as real-time data for some questions, requiring real-time continuous reporting for all ISP links and routers would be overly burdensome in terms of the sheer volume of data and overhead associated with collecting and managing the data. Without a focused rationale for how the data would be used (what important questions are to be answered?), such a "full disclosure" mandate would be inefficient.

From 2011 through July 2015, eleven of the largest wireline ISPs have been participating in this program. Performance disclosure for wireless ISPs is yet to be decided but may well end up mirroring the wireline MBA model that has been collecting data for a number of years but has not yet released a report or any of the raw data collected so far.

2.7. Edge Measurement Tool: Net.info

We propose "net.info" as a possible example of an emergent D&T intervention. It has several intended purposes, one of which includes providing a technical channel for broadband providers to communicate customer-specific information that may enhance the capabilities of edge-based measurement strategies.

Net.info is a project underway at MIT with the goal of developing an architecture and implementation for a new communication channel between ISPs and their downstream users. There does not currently exist a trusted, secure and easy to use channel for communicating information between access networks and their connected end users. It is currently surprisingly difficult for access providers to communicate notifications, alerts and other types of relevant network information with their attached users. The choice of the name reflects the desire to make this broadly and easily adoptable.[42]

Some of the kinds of information that should be able to flow from network access providers to end users over a secure, authenticated and trusted channel might include:

---

the many ISPs that do not participate in the program, or with respect to end-to-end measurements that may be of greater interest in certain contexts. The MBA does a good job at measuring what it was designed to do, but having additional tools and perspectives is also important for comprehensiveness and accuracy.

[42] Most obviously, *net.info* concatenates "network" and "information" in a short domain name that is globally accessible. Info is a shortening of information in 39 languages and net is equally recognizable. We believe users could become readily accustomed to typing net.info into web browsers or using applications that displayed net.info information and notifications.

1. Network performance characteristics (e.g., a user's subscribed upload and download speeds and data caps);
2. Network practices (e.g. network management information, notification of security threats such as botnet and malware infections); and/or,
3. Commercial terms (e.g., changes in privacy policies, Copyright Alert System notifications[43])

As a disclosure intervention in the context of the D&T Coordinator, net.info addresses event- and collective learning- level metrics (WHAT-axis), across the spectrum of stakeholders, although primarily intended initially for individual consumers and 3rd party researchers (WHO-axis), and in real time (although that data can be aggregated to produce longitudinal data as well) (WHEN-axis).

The *net.info* mechanism acts as a coordination point where such information would be readily accessible and capable of being easily shared (similar to how 9-1-1 is well known in the US telephone network). The information could be made available both in human readable forms (unstructured or semi-structured) accessible with any web browser and in structured formats (e.g. JSON, CBOR, etc) for use by software systems (e.g. menu bar notification apps, phone apps, and edge based measurement tools such as FCC's Measuring Broadband America program).

Today, access networks attempt to push notifications to their users via one of the following communication methods: email, phone, text messages, mail courier, and walled gardens.[44] In addition to being expensive, these channels are vulnerable to attack. Notably, malicious hackers have exploited the first four of these using social engineering tactics to worsen, not improve, end user security.[45]

If successful, *net.info* should offer an additional and improved channel for ISPs to communicate with end-users, but today it remains an early-stage work-in-progress. As such it highlights the way in which the D&T toolset may be expected to evolve, with the addition of new D&T capabilities altering the effectiveness of existing tools.

---

[43] See https://en.wikipedia.org/wiki/Copyright_Alert_System

[44] A walled garden is a network provider operated mechanism that restricts or modifies a user's web browsing experience for the purpose of ensuring that a user sees an piece of information, such as a notification of a malware infection.

[45] For example, we are personally aware of two incidents where end-users were called and partially talked through the installation of malware -- both with the justification that their computer was infected and needed to be "cleaned up."

3. D&T Application is Context Dependent

As described in the previous section, the D&T policies in the OIO are the source of both interpretation and application issues and disclosure intervention authority. Existing and forthcoming D&T intervention mechanisms enable translation of and compliance with those policies, and vary according to how well/to what extent they address those interpretation issues raised by the OIO's construction. In this section we illustrate how the D&T Coordinator tool helps compare the relative effectiveness of the D&T interventions in translating the D&T policies according to several use contexts.[46] We engage this comparative tool to illuminate how a specific edge measurement mechanism is more effective than a type of transparency report in disclosing certain required performance metrics and practices (loss, security) because of the level of detail, timeliness and targeted recipient. Further, we show how the same two disclosure interventions perform differently when the transparency issue concerns aspiration-level policies with different time and stakeholder needs.

At this point, it should be clear that the right D&T approach depends on the context of the policy question that is being addressed. In the following sub-sections, we examine two (seemingly) polar D&T contexts. These contexts vary from questions or issues that are narrowly focused and seemingly easy to define or specify (loss metrics, or "how should ISPs report packet loss?") to questions or issues that are broadly focused and amorphous (Internet aspirations, or "what is the right Internet for society?"). What we will show is that while these contexts vary along a continuum from the specific to the general in ways that match the initial intuition highlighted above, each also is more complex than might be initially presumed and each has need for a portfolio of D&T tools, but the appropriate mix and use of tools varies with the context.

Narrower questions like "loss metrics" are amenable to instantiation in specific metrics, reporting protocols, and standardization. Service Level Agreements (SLAs),[47] regulatory standards, and market research (e.g., consumer reports) already reference specific traffic metrics like upstream/downstream data rates, latency, jitter, and potentially, packet loss that are amenable to relatively clear and narrowly defined specification. But even with such data, there is a need for a multiplicity of definitions, concerns about information

---

[46] We assume that there is a positive correlation between the D&T policies themselves and the overarching objectives of the OIO detailed earlier. It is beyond the scope of the D&T Coordinator, and this paper, to address the extent to which the various D&T policies actually effectuate those objectives. Such focus should be the subject of future policy-centric research & analysis.

[47] SLAs are short-hand for contracts that may exist in many forms both retail service agreements between end-users and ISPs, and wholesale interprovider agreements. While consumer broadband contracts do not presently (true?) include specification of loss metrics, there is implicit specification that losses are within bounds to make service usable. For example (trivially), packet loss has to be less than 100% for service to exist for any service; but what constitutes acceptable packet loss may vary by application (real-time vs. delay-tolerant or bufferable). Other types of agreements (e.g., enterprise-level service agreements often do specify a host of detailed traffic metrics that may include packet loss statistics).

sharing, and a need for flexibility to accommodate diversity in definitions and interpretation, and flexibility to evolve the definitions as the Internet's technologies and markets change (mix of services shift, industry structure changes, and users grow in sophistication).[48]

At the other extreme, questions that relate to the sort of social values we would like our Internet infrastructure to embody are much more amorphous, and seemingly infeasible to instantiate in specific metrics, reporting protocols, and standardization. For example, market participants might agree that we would all like to have "trusted" Internet services, while failing to agree on what we mean by trust. One aspect of ensuring trust is associated with protecting Internet services from a plethora of threat vectors such as malware, phishing attacks, denial of service attacks, and the like. Security mechanisms like encryption tools, authentication mechanisms, and network monitoring comprise some of the apparatus with which stakeholders seek to protect services from such attacks. However, these tools, by their very nature need to evolve as threats, technologies, and markets evolve. Moreover, aspects of the tools, to be effective, must remain confidential. D&T policies can play an important role in structuring how security "best practices" and threat notifications might be shared and supported to help protect trust. And, this "security from attacks" is just one element of what might constitute the challenge of ensuring a "trusted Internet" that confronts market participants.

For a range of issues that we might consider as social aspirations, we would not expect to have a narrow or static set of metrics and will need to contend with ambiguous and ill-defined problems. Moreover, we can anticipate that for many important problems, stakeholders will not agree and may actively be in conflict with respect to the goals. For example, even when folks agree about the need for certain trust-supporting mechanisms, they may be expected to disagree when it comes time to apportion obligations for paying for those mechanisms.

3.1. Loss metrics

As noted earlier, one of the key goals of the FCC 2015 OIO is to enable consumers to make informed decisions about broadband subscription services and to determine whether providers are abiding by network information requirements. Obviously, having access to appropriate network performance information is important for this goal. The FCC 2015 OIO notes:

> "The existing transparency rule requires disclosure of actual network performance. In adopting that requirement, the Commission mentioned speed and latency as

[48] For example, in a world where telephone is only service, the choice of metrics was relatively simple; but in the Internet cloud, we need many more and more complex metrics and summary statistics to describer the diverse services (see Lehr, W., D. Clark, and S. Bauer (2013), "Measuring Performance when Broadband is the New PSTN," Journal of Information Policy, Vol. 3 (2013), pg. 411-441, available at: http://jip.vmhost.psu.edu/ojs/index.php/jip/article/view/94).

> two key Measures. Today we include packet loss as a necessary part of the network performance disclosure."[49]

Thus, the FCC 2015 OIO is reaffirming but also expanding the scope of the network performance disclosures that were originally mandated in the FCC 2010 OIO. Information about packet losses as a performance-relevant metric is interesting because such losses may be indicative of network congestion or other problems that may adversely impact the consumer's Internet experience. While the details of this packet loss performance disclosure have not yet been worked out, it is worthwhile exploring in some detail the challenges of framing appropriate loss metric disclosures. In this way we can understand the likely impacts on the various stakeholders' and decide how to improve disclosures in accordance with their respective rights and interests (as per the objective of the OIO).

The FCC offered limited guidance on the nature of the network performance disclosure in the FCC 2015 OIO leaving the details to the FCC CTO to work out. They noted that disclosures "should be reasonably related to the performance the consumer would likely experience in the geographic area in which the consumer is purchasing service" and that "network performance will be measured in terms of average performance over a reasonable period of time and during times of peak usage." Implicit in this guidance is that performance measurements are spatially and temporally relative to other similarly situated consumers, thus arguing for disclosure of manifold measurements. In this section we describe the co-evolving relationship between loss metrics and several disclosure interventions in terms of interpretation issues and comparative disclosure technique shortcomings.

3.1.1. D&T with Measuring Broadband America

As discussed above, MBA provides a valuable D&T tool, but with respect to loss-metrics, it has a number of deficiencies.

[49] 2015 OIO, paragraph 166.

Test	Primary measure(s)
Download speed	Throughput in Megabits per second (Mbps) utilizing three concurrent TCP connections
Upload speed	Throughput in Mbps utilizing three concurrent TCP connections
Web browsing	Total time to fetch a page and all of its resources from a popular website
UDP latency	Average round trip time of a series of randomly transmitted UDP packets distributed over a long timeframe
UDP packet loss	Fraction of UDP packets lost from UDP latency test
Video streaming	Initial time to buffer, number of buffer under-runs and total time for buffer delays[23]
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
DNS resolution	Time taken for the ISP's recursive DNS resolver to return an A record[24] for a popular website domain name
DNS failures	Percentage of DNS requests performed in the DNS resolution test that failed
ICMP latency	Round trip time of five regularly spaced ICMP packets
ICMP packet loss	Percentage of packets lost in the ICMP latency test
Latency under load	Average round trip time for a series of regularly spaced UDP packets sent during downstream/upstream sustained tests
Availability[25]	Total time the connection was deemed unavailable for any purpose, which could include a network fault or unavailability of a measurement point
Consumption[26]	A simple record of the total bytes downloaded and

Figure 2: Measuring Broadband America (MBA) Test Suite

First, different measurement methods for loss may give very different answers such that we cannot rely on the MBA, by itself, to resolve the ambiguity inherent in the OIO. For example, the existing set of MBA tests explicitly mention "loss" in multiple metrics (see *Table 1*). These include: UDP packet loss, ICMP packet loss, Voice over IP. Loss metrics can also be derived from a variety of the other tests, including download speed, upload speed, and latency under load. These different ways of measuring loss provide quite different estimates

Perhaps as significant as what the tests measure is what they do not. The MBA tests are not end-to-end tests. These tests measure the path between the customer and a test server that is located over an interconnection point that is designed to be uncongested. The focus of the MBA program has been to measure performance in the access network itself. Indeed if the interconnection point between the access network and the network hosting the test server is congested, as has occasionally happened, tests from the congested period are not included in the annual FCC report.

Second, in addition to loss metrics derived from active measurements like in MBA, loss can also be reported from router interfaces. Some of the loss measurements that are of greatest potential interest to regulators and the general public would be the losses that may occur at the interconnection links to large content and service providers, and hence are not included in the current MBA measurements. If one looks at the research literature

to understand how loss measurements from active probes compare to loss measurements derived from router interfaces there is evidence to suggest there might be some significant differences. See Figure 3 Comparison Loss Rates on Two Comparable Routes below which is copied from a research report in 2004.[50] Exactly why these loss measurements are so different is unclear, and these results from 2004 might not be replicated today.

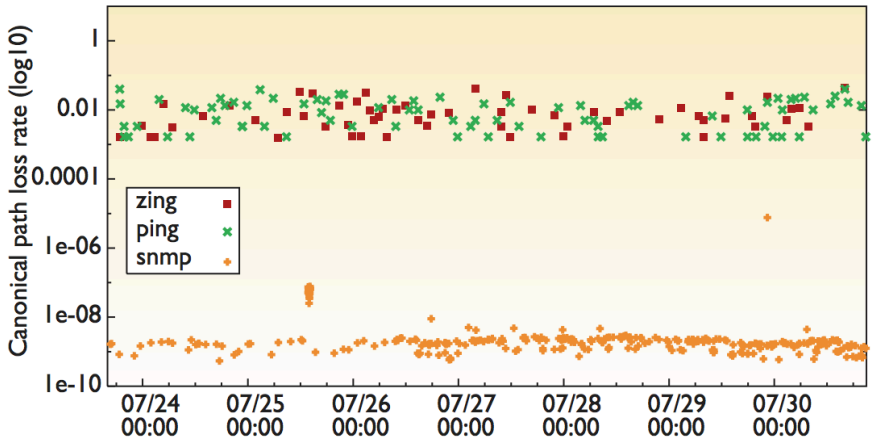

Figure 3 Comparison Loss Rates on Two Comparable Routes

When interconnection links are congested, there is a question as to which provider (the ISP or the upstream content/application provider) should report the congestion.[51] In routers queues generally build on the outgoing interface. Thus if the bulk of traffic is headed toward a broadband access provider and an interconnection link to the access provider is congested, the sender (e.g., Netflix or YouTube), not the broadband access provider will see the losses occurring on their router. Monitoring the interface on the broadband access provider's router on the customer-side of the interconnect will not show any congested related losses. While the existing MBA test suite does not cover these links, the suite could be evolved to provide a better picture of performance on those links. Indeed, currently being trialed is a video performance test to actual video providers like YouTube and Netflix.

Third, regulators in a number of other countries have deployed measurement platforms similar to the MBA program and have issued reports. *Table 2* summarizes a survey of

---

[50] The chart shows the comparison of loss rates on the canonical path for traffic from Indianapolis to Loss Angeles, using a 20-Hz probe rate (see, http://pages.cs.wisc.edu/~pb/intcomp_final.pdf.

[51] It is worth noting that earlier, the FCC had considered but ultimately rejected requiring that broadband providers disclose "meaningful information regarding the source, location, timing, speed, packet loss, and duration of network congestion." In explaining its decision, the Commission noted that "congestion may originate beyond the broadband provider's network" and that there are "limitations" on "a broadband provider's knowledge of some of these performance characteristics." (2015 OIO, ¶168).

these reports, indicating that loss metrics have not been a focus of any of the reports to date. To the extent loss metrics are mentioned, the focus has been on the UDP test. Most regulators have reported average loss rates over different time periods, monthly or by average loss per hour or peak period. Brazil had a slightly different metric and reported on the percentage of time loss was below 2%. They noted a target of 90% of the time loss should be below 2%.

Regulator	Report date	Loss metric	Notes
United States FCC	July 2014, Feb 2013, July 2012, Aug 2011		Never reported on loss
European Commission	October 2013, March 2012	Averages derived from UDP loss tests	Compare loss EU averages to loss averages in US
Singapore	Ongoing	Averages derived from UDP loss tests	Calculate loss to targets in US
UK Ofcom	12 reports from 2008-2014 (every 6 months)		Only reported on loss in first report
Brazil	7 reports from 2013 – 2015 (every six months)	Threshold metric drived from UDP loss tests	Report metric is % of time loss is below 2%

Table 2: Summery of how regulators using SamKnows measurement boxes have reported on loss.

Finally, it is worth noting that loss rates are not independent of other network performance relevant metrics. Changes in loss rates are one of the ways the Internet controls congestion. Achieving zero loss rates would not be an appropriate goal. Indeed, the appropriate standard for packet loss rates depends on how TCP behavior is managed, and ISPs do not control TCP behavior. This raises the question of how should we measure loss rates and what loss-relevant disclosures might it be useful to obtain from access ISPs. Further, measurement when the network is idle (as is the current MBA methodology) will most certainly differ from the exact same measurement when the network is under load, as well as that of an ongoing transfer.

3.1.2. D&T with net.info

We next consider some motivating examples of how *net.info* can assuage some measurement interpretation discrepancies via supplanting or augmenting other D&T techniques (such as the MBA) via efficient and secure disclosure of information.

It is currently exceedingly difficult for end users to identify basic information such as the headline upload and download speeds of their broadband access service.[52] While this information is readily available when signing up for a broadband service it changes over time as service is gradually upgraded. At least for the authors of this report, learning the current service levels of their residential broadband service involves a complicated procedure of logging into their broadband account and comparing the service name on their billing information with a lookup on a zip code specific mapping of service names to speeds on a different part of the provider's website.

Especially at a time when over 80% of households already have broadband service in the United States, real time and easy access to this basic service information is arguably more important than broadband labeling efforts aimed at point of purchase time periods. It is unfortunate that when tens of millions of broadband speedtests are run each year users are not capable of easily comparing their actual performance with the advertised speeds of their service in real-time.

One of the reasons that the FCC's MBA project only releases its data and reports once a year is that the procedure for validating panelist service information is run by the FCC and ISPs only once a year and only covers the one month of the report period. This limits the analysis that is conducted by the FCC to a single, hopefully representative, month as well. Singapore on the other hand releases data from their very similar, albeit smaller, broadband measurement effort every three months with data from every month included.

The adoption of *net.info* would enable testing tools such as the FCC's measurement devices and the popular speedtest.com to acquire accurate service detail information automatically. This could enhance both regulatory and end user monitoring of broadband performance.

A bit of new D&T mechanism like *net.info* may also contribute to enhancing Internet security. The prevention, detection and mitigation of cyber threats are increasingly being offered by and/or expected of service providers, including ISPs, which may be expected to have relative advantages in terms of mitigating threats to the Internet. Botnets are one such prominent malicious threat.[53] In September 2014, the Messaging Anti-Abuse Working Group (MAAWG) reported the percentage of subscribers that participating network operators had identified as being infected with a botnet. Some the key results are included in Figure 4: M3AAWG Bot Metrics Report, below.[54]

---

[52] Other notions like "actual speeds," "expected speeds," or "maximum speeds" are even more difficult since there are important issues associated with the design of appropriate metrics for these concepts.

[53] Ramneek, Puri (2003-08-08). "Bots &; Botnet: An Overview". SANS Institute. Retrieved 9 August 2015.

[54] For the full report, see https://www.m3aawg.org/for-the-industry/bot-metrics-report.

2012	Q1 2012	Q2 2012	Q3 2012	Q4 2012
Subscribers Represented	37,707,435	37,358,206	36,991,516	37,383,662
Subscribers Deemed Infected	317,064	402,585	249,492	440,746
% Infected	0.84%	1.08%	0.67%	1.18%
Infected Subscribers Notified	314,295	400,439	245,522	437,253
% Notified	99.13%	99.47%	98.41%	99.21%

2013	Q1 2013	Q2 2013	Q3 2013	Q4 2013
Subscribers Represented	37,270,265	37,735,195	37,639,022	43,550,674
Subscribers Deemed Infected	388,152	435,921	493,572	346,615
% Infected	1.04%	1.16%	1.31%	0.80%
Infected Subscribers Notified	387,221	435,149	492,382	325,787
% Notified	99.76%	99.82%	99.76%	93.99%

Figure 4: M3AAWG Bot Metrics Report

Around 1% of subscribers are detected as being infected with malware, and attempts were made to notify almost all those infected. The study does not report on the costs or effectiveness of the mitigation efforts, but it is noteworthy that the incidence of infections was not substantially reduced over the eight quarters. A persistent supply of up to 500k bots poses a significant threat to Internet security.

We posit that *net.info* has the potential to drive down by one or two orders of magnitude the number of subscribers infected with a botnet or malware that can be detected by their ISPs. There are two key elements: 1) building an authenticated, secure, and trusted channel (that is not vulnerable to social engineering attacks that worsen instead of improve the problem); and 2) significantly reducing the response time between when an infection is identified and when it may be addressed. By design, net.info focuses on the communication of information about the infection, while leaving the response to that infection to be determined by other elements of the complete security control system. This partial approach is intentional, and we argue, consistent with enhancing the likelihood of incentive-compatible adoption in today's Internet environment.

There is a wide-range of possible responses to infections that may include human or software helpers (by end-user, the ISP, third-party actors, or regulators) that may be used in conjunction with *net.info* to affect a full response to the Botnet threat. Separating the communication of the threat from the choice of how the threat is addressed partitions the incentive problem so as to allow greater scope for local, context-based decision-making, and hence, is consistent with the goal of enabling an open, decentralized control plane.

3.2. Managing Disclosure Policy and Internet Aspirations

For our final use context, we address what we refer to as aspirational goals for the Internet. As with loss metrics above, we explore the relative merits of the same disclosure interventions in terms of interpretation and application issues with so-called *Internet aspirations*. Clark & Claffy (2015) identify a range of goals that have been articulated as societal aspirations for what the Internet *should* be. To paraphrase, some of the

aspirations they identify include that the Internet should be secure (safe, trusted);[55] be universally available (ubiquitous, affordable, accessible);[56] support personal autonomy (voluntary, choice);[57] and an be an efficient[58] platform for economic activity (commerce, competition, innovation, growth).[59] This list is neither comprehensive nor disjoint, but is sufficient for our purposes here.

Each of these goals – as articulated – would appear to be universally appealing. While it is hard to imagine aversion to these attributes, it is also hard to imagine subscribers universally agreeing on what any of these vague aspirations mean. A critical step to steward more shared understanding is translating and mapping technical network layer terms and meanings to their regulatory human layer counterparts. For example, a subscriber's subjective assessment of his/her quality of experience (QoE) using the network (e.g., application, reliability, speed, etc.) needs to be translated into the equivalent measurements that are meaningful to network engineers (e.g., throughput, delay and jitter, and packet loss).

Moreover, as one refines these notions further (and as Clark & Claffy note) many of these refinements and the larger notions are likely to conflict. For example, personal privacy and system security are often in conflict; or, one individual's exercise of free choice may

---

[55] We are using "Security" here broadly, as in a secure system is one that users can trust and feel safe when using; and that promotes/supports these feelings of security, trust, and safety more generally in the wider economy and society. (Users includes everyone who participates in the Internet ecosystem so consumers, but also businesses that use the Internet or contribute to supporting the Internet.)

[56] We are using "availability" broadly to encompass not just ubiquitous accessibility of infrastructure (everywhere/always connected), but also that "users" have the complementary tools (devices, software) and skills (education) needed to make effective use of the Internet, and this is affordable. This broader interpretation is required to make sure we are not just building highways that no cars are able to actually drive on (because there are no cars, drivers, or driving is too expensive).

[57] A fundamental human right is that individuals have (as much) "free choice" over their actions (as is compatible with society). (Note, different societies and at different times, folks have very different views about what should constitute free choice and how societies should be organized to manage it. With respect to the Internet, users should have choices in whether and how they use the Internet. While a goal may be to promote universal adoption/usage (potentially a refinement of availability), there is generally a presumption that the decision to join the Internet be voluntary.

[58] Economic efficiency implies that the Internet infrastructure satisfies productive (resource/cost minimizing), allocative (scarce resources directed to highest value uses), and dynamic efficiency (investment is efficient over time). In real world with uncertain/stochastic demand and supply shocks, tolerating congestion (which requires allocation of scarce resources) is likely to be efficient.

[59] All of these goals underpin what it means to regard the Internet as basic economic infrastructure akin to our road, electricity, or water systems. These are infrastructures that are used by virtually everyone in the normal course of their social and economic lives.

interfere with the choices of others (e.g., pollution, congestion); or, competing firms may favor architectures that give one or the other a strategic advantage. The fact that social aspirations often result in conflicts among stakeholders and/or among goals is hardly surprising since how these aspirations are defined and implemented will impact the allocation of costs and benefits. For example, meeting universal service goals is likely to require cross-subsidies. Further, even when stakeholders may reach agreements on how to define and implement collective Internet aspirations, it is reasonable to expect that choices will change as markets, technologies, and relationships among stakeholders evolve.

In spite of the vagueness of such Internet social aspirations, these do represent an important class of questions that need to be addressed. How stakeholders frame, prioritize, and resolve such questions will impact Internet architecture, investment, and regulatory policy, and thus are of substantive importance. It may seem as if these sorts of questions are categorically different from the loss metrics context case discussed earlier. This certainly bears some truth, but the differences are more ones of degree than kind. With social aspirations, we might expect greater difficulty in even framing the problem to be addressed (e.g., what to disclose about packet losses versus how to measure universal availability), let alone reaching agreement on a well-defined set of D&T metrics and reporting obligations.

A core conceptual bridge between this *Internet aspirations* use context and a disclosure lever is the OIO's *No Unreasonable Interference/Disadvantage Standard,* a D&T policy intended to protect the open nature of the Internet by granting the FCC the discretion to prohibit practices that unreasonably interfere with or unreasonably disadvantage the ability of consumers to reach the Internet content, services, and applications of their choosing or of edge providers to access consumers using the Internet.[60] Disclosures that prevent and detect harmful traffic discrimination against citizen-consumers by ISPs (for financial advantage or other unreasonable justification) is a feature of D&T interventions that will help translate this D&T policy.

In terms of the D&T Coordinator, questions related to aspirations are likely to require collective learning (WHAT), and engage all stakeholders (i.e., all types of disclosure recipients, WHO), and at all time-scales (i.e. real to longitudinal in the WHEN). Informing these sorts of questions will also depend on the full panoply of D&T interventions. For example, the detailed data included in MBA transparency reports will be supplemented with other third-party edge-based measurement data and reports from marketing research firms, public interest groups, academics, and industry representatives that will collectively comprise the public discourse. Process rules will help ensure

[60] OIO at ¶136. "Any person engaged in the provision of broadband Internet access service, insofar as such person is so engaged, shall not unreasonably interfere with or unreasonably disadvantage (i) end users' ability to select, access, and use broadband Internet access service or the lawful Internet content, applications, services, or devices of their choice, or (ii) edge providers' ability to make lawful content, applications, services, or devices available to end users. Reasonable network management shall not be considered a violation of this rule."

adequate participation in the collective decision-making discourse, while empirical data will help provide the evidence-based foundation for informed collective decision-making. New bits of mechanism like *net.info* that can enhance the efficiency of existing platforms (like MBA or web-based speed measurements) or create new ways to exchange data (like information about malware infections) can help build trust and provide a way to discover and share new classes of information.

As with loss metrics, the D&T Coordinator can illuminate the relative merits of engaging MBA transparency reports and net.info tool to address the vague and dichotomous nature of Internet aspirations described above. According to their previously detailed properties, MBA transparency reports focus on longitudinal, mid- to fine-grained data for regulators or third party intermediaries could be assistive in strengthening trust among the collective subscribers viz-a-viz third party research reporting. Those contours however would not make for progress in stewarding network-to-human-layer concepts like QoE, ameliorating conflicts between stakeholders, or impacting social level autonomy or economic efficiency. Similarly, a tool like net.info that can in more real time and directly enable communication with individual subscribers and third party analysts may also engender consumer trust by allowing corroboration or refutation of ISP discretionary disclosures. Because of the independent nature of it's design, it also portends a stronger likelihood of being embraced as a disclosure mechanism by those same stakeholders, which is a necessary prerequisite for the high level collective learning that underpins the *aspirations* challenges.

3.3. Loss Metrics v. Aspirations: Context matters

Whereas one might characterize the earlier contexts as relating to questions amenable to objective specification and within scope for well-defined metrics and D&T standards, the discussion of questions like social aspirations for the Internet are inherently more open-ended.

D&T interventions that encourage and support voluntary information sharing to support active and better-informed interactive discussions about social aspirations should facilitate collective decision-making. These mechanisms need to be open to raising new questions and explorations as aspirational concerns arise and evolve.[61]

With respect to open-ended questions, the role of D&T is not simply to resolve the issue but to provide structure and mechanisms for on-going discussion. The open-endedness of social aspiration-like questions may make it more difficult to predict or limit the use of information made available as part of D&T processes. This need to evolve and respond also applied to contexts like loss metrics, but to a sufficiently lesser degree that coming up with concrete metrics and clear specification of reporting standards seems more tractable. However, this is only partially true since debates over social aspirations will be

[61] For example, the definition of what constitutes acceptable broadband service may evolve as needs for asymmetric or symmetric upstream/downstream bandwidth evolve with Internet services and markets.

informed by the evidence and information associated with other contexts. For example, data on loss metrics or reliability will be part of the context and will help frame and inform discussions over aspirations for universal availability and network performance (e.g., how close is today's Internet's performance to what society wants?).

4. Conclusions

In preceding sections, we have explained that although it seems obvious that for broadband markets to work properly, participants (ISPs, edge providers, consumers, and regulators) need information about network performance and practices to make informed decisions (What services to develop/sell/buy? How to use those services? How to manage networks? Etc.), it is far from trivial to figure out what information needs to be disclosed by who and to whom, for what purpose, and when to actually accomplish the seemingly simple goal of ensuring appropriate information transparency and disclosure.

To better understand Disclosure and Transparency (D&T) policies, we have presented a framework that translates the way in which the decision context differs depending on the type of question to be answered and how these are interrelated. Our analysis leads us to several high-level conclusions.

First, we conclude that in order to effectively address the complex challenges that D&T policies are directed at solving (i.e., promoting informed collective decision-making), a multiplicity of D&T intervention tools are needed. Second, to understand the efficacy of these tools, they need to be comparatively assessed with respect to the features that distinguish the context of specific decisions. By articulating the question to be answered according the dimensions of the D&T Coordinator -- *what* information is needed, by *whom*, and *when*—we can more consistently and comprehensively address the information deficiency at play. Context matters, and nuanced application of D&T tools is important. While the Coordinator does not necessarily prescribe D&T requirements, by facilitating a more standardized comparison of the various tools, it provides a much needed ability to intelligently distinguish contexts and inform the associated demands for more specific statistics and conditions.

Second, we conclude that the FCC's current approach toward D&T as directed in the current OIO and on-going efforts to expand disclosure are generally appropriate. However, we note that because the OIO is directed at ISPs and does not apply to the conduct of edge providers of content, applications, or devices, whose service provisioning and traffic management decisions are also important determinants of end-to-end performance, the OIO's D&T framework is incomplete.

A richer understanding of the D&T challenge, enabled by the framework explicated here, should give caution to both those who argue that the FCC should mandate much more complete disclosure by ISPs through orders that embody greater specificity that leaves less room for industry self-regulation, and those who argue that the FCC should scale back disclosure requirements and leave it to markets to reveal whatever information is needed. Our own research on Internet metrics and focused examination of the D&T

challenges helps us better appreciate the complexity of the policy challenge confronting policymakers generally and the FCC more narrowly, and highlights the gaps in D&T policies that still need to be addressed.

Third, we believe that the current D&T environment will benefit from the growth and maturation of edge-based, independent third-party measurement infrastructure. This is important not only as a complement to ISP self-disclosures and the MBA, but also to provide insight into portions of the end-to-end service that the current OIO does not adequately address (i.e., performance issues arising in end user or edge provider-controlled portions of the end-to-end path).

Contributing to the development of such infrastructure has been the focus of a significant portion of our research efforts in recent years (CAIDA-MIT, MITAS, etc.). We explain why better edge-based measurement capabilities are necessary and complement our first conclusion (i.e., in the absence of such capabilities, we would be inclined to favor more aggressive mandatory disclosure policies by the FCC). We also believe that the efficacy of edge-based measurements would be enhanced by some additional Internet "D&T infrastructure" that would facilitate controlled information sharing (disclosure) between ISPs and users at the edge. We call this proposal "net.info" and briefly outline the concept and our plans for developing this idea further.

5. References

Lehr, W., D. Clark, and S. Bauer (2013), "Measuring Performance when Broadband is the New PSTN," Journal of Information Policy, Vol. 3 (2013), pg. 411-441 (available at: http://jip.vmhost.psu.edu/ojs/index.php/jip/article/view/94)

Lehr, W. (2012), "Measuring the Internet: the data challenge," Organization for Economic Cooperation and Development (OECD) Digital Economy Working Paper 184, ISSN 2071-6826, April 2012, available at: http://www.oecd-ilibrary.org/science-and-technology/measuring-the-internet_5k9bhk5fzvzx-en.

Lehr, W., S. Bauer, M. Heikkinen, and D. Clark (2011) "Assessing broadband reliability: Measurement and policy challenges," 39th Research Conference on Communications, Information and Internet Policy (www.tprcweb.com), Alexandria, VA, September 2011. (pdf=http://people.csail.mit.edu/wlehr/Lehr-Papers_files/Lehr%20et%20al%20TPRC2011%20Assessing%20Broadband%20Reliability.pdf) (slides=http://people.csail.mit.edu/wlehr/Lehr-Papers_files/TPRC-2011%20Reliability%20Lehr.pdf

Bauer, Stephen, David Clark, and William Lehr (2009), "Broadband Micro Foundations: the Need for Traffic Data," invited paper prepared for Beyond Broadband Access conference, NewAmerica Foundation, Washington DC, September 22-24, 2009. (pdf= http://people.csail.mit.edu/wlehr/Lehr-Papers_files/Bauer_Clark_Lehr_Traffic.pdf). Published in Richard Taylor and Amit Schejter (2013), Beyond Broadband Access: Developing Data-based Information, Fordham University Press, 2013.

Lehr, W. (2015), "Reliability and the Internet Cloud," C.Yoo and J-F. Blanchette (eds), in Regulating the Cloud: Policy for Computing Infrastructure, MIT Press: Cambridge, MA, 2015.