Reliability and the Internet Cloud

William Lehr[1]
Massachusetts Institute of Technology

June 2012

1. Introduction

The Internet is becoming the new PSTN at the same time as it is evolving into the Cloud (a public utility for networked computing resources). These concurrent transitions will simultaneously increase the saliency of and complexity of ensuring reliability. Understanding the nature of this challenge requires bridging divergent views of reliability as it has been considered in the Internet, what it has meant in the telephony-centric PSTN, and what it will mean in the evolving Internet Cloud of the future. This will have cross-layer implications for the entire Internet cloud computing ecosystem, where the "layers" refer not just to the protocol layers in an IT-architecture sense, but the industry/market structure, business processes, and regulatory environment in which the Internet Cloud will exist. This paper will review how the challenge of ensuring reliability will evolve and what this will mean for policymakers and industry stakeholders. The challenges of insuring high-levels of reliability for critical infrastructure is not unique to the Internet, and much can be learned from other domains, although the legacy PSTN provides an obvious touchstone. This paper will help frame the discussion of ensuring reliability in an Internet cloud ecosystem, and will interpret some of these lessons in light of current directions in future Internet architectures. Of specific interest, this paper will discuss some of the challenges and opportunities presented by the design of a highly reliable core router architecture that will be analogous to the "carrier-grade" switching fabric of the legacy PSTN. In addition, this paper will comment on the need for and challenges for developing public metrics for assessing cloud reliability performance.

In Section 2, I trace the evolution of the Internet and PSTN, and explain what this implies in terms of technology, industry structure, and policy for the Internet ecosystem. In Section 3, I focus on the challenge of ensuring reliability in the new Internet cloud. In Section 4, I take up the special problem posed by the transition to a hyper-reliable core routing architecture. Section 5 concludes.

2. Changing Internet Ecosystem

[1] Email: wlehr@mit.edu. In completing this work, Dr. Lehr would like to acknowledge support from NSF Awards 1040020, 1040023, and the MIT Communications Futures Program. All opinions expressed herein are those of the authors alone

In this section, I describe how the Internet has evolved from a best-effort network into a platform for cloud computing services. In realizing this evolution, the Internet has become the new PSTN, basic essential infrastructure for our information economy. This transition has important technical, market structure, and regulatory implications.

2.1. From Telephone Network Application to the Internet Cloud Utility

Since its origins in the 1960s, the Internet has evolved from an application supported on top of the Public Switched Telephone Network (PSTN) into *the* platform for all global electronic communications.[2] As a consequence of this evolution, the Internet has experienced exponential growth in its capacity, capabilities, and the volume of traffic and diversity of applications it now supports.

In the 1990s, the Internet emerged as the first successful mass-market platform for data communications, adding the third crucial element needed to realize the world of pervasive *cloud* computing that we are still in the midst of transitioning towards. The other two legs were the concurrent PC revolution that delivered mass-market computing resources to end-users desktops and the growth of mobile telephony that brought us personalized mass-market communication services. The prototypical Internet services were delay-tolerant email, Web access, chat, and file-sharing.

Mass-market access to networked computing resources proved sufficiently compelling to spur exponential growth in eCommerce and investments in telecommunications infrastructure and complementary goods and services all across the ICT value chain. New ventures with novel business models like Amazon.com, Google.com, and eBay.com – and a host of others – proliferated to take advantage of the market opportunities that became available with the growth of the Internet. Unfortunately, realization of the Internet's potential was hampered by the slow speeds of dial-up access connections, the lack of mobility support, and the limited capabilities of user devices, applications, and the Internet in those days. These limitations contributed to the Dot.com bust of 2000 when ambitious hopes for growth collided with real-world challenges.

With the migration to broadband and now mobile broadband, with enhanced interactivity through technologies like Web2.0 and new user devices/interfaces (like tablets, eBook readers, connected TVs, and smartphones), and with the big expansion in the range of applications and content (such as social networking, interactive multimedia, and video conferencing), the Internet is increasingly pervasive in society and the economy. With Moore's-Law-driven advances in computing, storage, and communications technology, we now are able to foresee a future of pervasive computing where we are always/everywhere connected, and where all manner of activities may be computing-assisted. The assistance may or may not require human interaction or awareness, may be passive or active, and might be machine-to-machine.

[2] Much of the discussion in this section parallels our discussion in Lehr, Bauer, and Clark (2012).

This future is sometimes referred to as the *Internet of Things* (IoT), and its fullest realization, would merge the real and virtual worlds.[3] Such a future will require embedding computer intelligence in all sorts of devices and network elements, rendering *smart* our end-to-end computing/communications systems. Such systems underlie visions of smart grids,[4] smart infrastructures (highways, buildings, transport grids),[5] smart supply-chains,[6] smart healthcare,[7] and so on. These *SmartX* systems are at the core of national strategies for economic growth and environmental sustainability.[8]

The Internet is a central element in this vision of pervasive computing/communications resources. Whether the Internet's role should be principally to provide the telecommunications services to connect intelligent devices at the edges (CPUs and storage in data centers); or whether such services and resources should be embedded in the Internet is a question of active debate among network researchers, industry participants, end-users, and policymakers.[9] Ignoring for the moment where the *smart* functionality should be located (and who should control or own the assets that support it), it is clear that there are many things that today's Internet does not handle well that might be better addressed if the Internet's functionality were expanded. This includes things like better support for trust (security, privacy), better support for context-differentiated services (for quality-of-service, location awareness, or other "context"-related

---

[3] For example, in a report prepared for the European Commission, Botterman (2009) describes IoT as "'a world-wide network of interconnected objects uniquely addressable, based on standard communication protocols,' or, more widely: 'Things having identities and virtual personalities operating in smart spaces using intelligent interfaces to connect and communicate within social, environmental, and user contexts'." Instead of just imaging a world with computing/communications "*anytime, any place* connectivity for *anyone*," we will have world where such connectivity is extended to *"anything"* (see ITU, 2005).

[4] See Department of Energy at http://energy.gov/oe/technology-development/smart-grid.

[5] See "Smart roads, smart bridges, smart grids," Wall Street Journal, February 17, 2009 (available at: http://online.wsj.com/article/SB123447510631779255.html).

[6] See "The Smarter Supply Chain of the Future: Global Chief Supply Chain Officer Study," IBM, 2009 (available at: http://www-148.ibm.com/tela/servlet/Asset/297861/CSCO Study 10_21_09.PDF).

[7] See "Connected Health," Cisco Healthcare Solutions, 2012, (available at: http://www.cisco.com/web/strategy/docs/healthcare/cisco_connectedhealthcare_overview.pdf).

[8] See "Strategy for American Innovation," White House of the United States, February 2011 (available at: http://www.whitehouse.gov/innovation/strategy). One of the key building blocks is identified as "develop an advanced information technology ecosystem…a 'virtual infrastructure'" that encompasses the "critical information, computing and networking platforms that increasingly support our national economy." Or, see, "Connection technologies to play critical role in building sustainable future – UN" (7 February 2012, available at: http://www.un.org/en/development/desa/news/sustainable/connection-technologies-to-play-critical-role-building-sustainable-future.html)

[9] For example, see Thierer (2006), Odlyzko (1998), Lucky (1997) or Isenberg (1997).

differentiators in service characteristics[10]), better support for network management (to allow better dynamic resource allocation), as well as support for cloud-based computing and storage resources. Much of this functionality is already supported via a hodge-podge of Internet add-ons and fixes provided as value-added services by participants in the Internet value chain. Meanwhile, as part of the NSF's Future Internet Architecture (FIA) program, several teams of network researchers are redesigning the Internet to expand the range of intelligent functionality to support finer-grained context-dependent resource assignment, including to shared computing and storage resources.[11]

Enabling on-demand access to computing and storage resources via the Internet is a motivating characteristic of "cloud computing." A common taxonomy for cloud services identifies three tiers of access:[12]

- (i) Software-as-a-service (SaaS), which provides access to cloud-hosted applications, enabling thin-client users to access software applications via the Internet. Examples include web-based email, office software like Microsoft Office[13], or Google Apps[14].
- (ii) Platform-as-a-service (PaaS), which provides a platform for hosting applications in the Internet, with tools for accessing and managing the underlying computing, storage, and networking resources. Examples of this include Microsoft's Azure[15] and Google App Engine[16].
- (iii) Infrastructure-as-a-service (IaaS), which provides access to the underlying core computing, storage, and network resources. These can be used to contruct on-demand, virtual enterprise computer networks. Examples of IaaS providers include Amazon's Elastic Cloud (EC2)[17], Rackspace[18], and IBM Computing on Demand[19].

---

[10] Traditionally, much of the discussion over Quality of Service (QoS) differentiated services has focused on the need to address differential requirements for latency or other technical service attributes. For example, delay-tolerant applications like email may be better supported than delay-intolerant applications like telephony over a best-effort Internet service in the face of congestion, inducing some to advocate using technologies like MPLS, DiffServ or other techniques to support more fine-grained (service-specific) service provisioning. However, "context" may be thought of more broadly as a characteristic of the type of application (telephony v. email), the identity of the parties communicating, the time/location of the communication, or anything else that might make it appropriate to manage the resources used to support the activity more effectively.

[11] See Jianli, Paul and Jain (2011) for a survey of Internet architecture research. The author is a participant in two of the projects mentioned, MobilityFirst and Nebula.

[12] See Zhu (2010), Armbrust et al. (2009), or Rimal, Choi, and Lumb (2010).

[13] See "Office365" at http://www.microsoft.com/en-us/office365/online-software.aspx -fbid=LpfpXGYHiYs.

[14] See "Google Apps for Business" at http://www.google.com/enterprise/apps/business/.

[15] See http://www.windowsazure.com/en-us/.

[16] See https://developers.google.com/appengine/.

[17] See http://aws.amazon.com/ec2/.

[18] See http://www.rackspace.com/.

[19] See http://www.ibm.com/cloud-computing/us/en/.

Some of the essential attributes that characterize the "cloud computing" vision include:[20]

- On-demand access to resources (storage, computing, network)
- Dynamic scaling of capacity (up and down)
- Broad network access (flexible "anywhere" access support)
- Resource pooling (shared resources via virtualization)
- Measured services (pay-as-you-go support for on-demand resources)

The economic benefits of enabling such functionality are several. On-demand access to resources allows better dynamic matching between resource needs and capacity. From a business/economic perspective, this can translate capital (fixed) costs for excess capacity to operating (variable) expenses, with the user paying only for the capacity the user needs and uses. These savings can also translate into savings in power and other shared operating costs. Whether opting for cloud-based services to meet a user's need for computing/communication services is a good decision obviously depends on how the cloud services are provided and priced relative to available alternatives. In principal at least, there may be significant scale and scope economies realizable from relying on shared resources to meet heterogeneous (and uncorrelated in time/location) demands.[21]

In addition to cost or resource utilization benefits, cloud-based resources may offer benefits in supporting flexible access with support for thin-clients, mobile,[22] or ad hoc usage[23] and enhancing service reliability. When services are distributed in a highly connected cloud there are many more routes to support robustness in the face of one or multiple link failures.

Indeed, the reliance on such resource sharing was fundamental to the economic design of the PSTN as a general telephone "utility." The PSTN relied on shared transport and switching to allow anyone-to-anyone telephone calling. The "cloud utility" model generalizes that model to include computing and storage resources. As we discuss further below, this generalization implies increased complexity.

2.2. From Service to Basic Infrastructure

[20] See Mell and Grance (2009).

[21] Note, if demands are strongly correlated in time/location (everyone wants the same computing resources at the same time), then sharing will not efficiently address the peak capacity challenge.

[22] Mobile clients may also be thin clients because of the device power, portability, and other inherent design constraints. Thin clients may make it easier to port applications to new devices, especially as we expand the range of connected entities in an Internet-of-Things world.

[23] Ad hoc here refers to unplanned or disruption-prone applications. If you cannot predict where/when you will need resources, your only option may be to provision on the fly. Demand and supply (capacity) shocks may be the case of such uncertainty. A natural disaster is an example of just such a shock.

The expansion in Internet capacity and capabilities described above has been driven by, and in turn, helps drive the virtuous cycle of service demand and supply growth. The Internet has grown in scale (globalization, adoption saturation, exponential traffic growth) on an aggregate and per-subscriber basis. That is, more users are using the Internet for a wider range of applications that engage an ever increasing range of activities in our social and economic lives. In light of this transformation, the Internet is appropriately regarded as essential basic infrastructure.

Like with electric power, roads, water, *and* the telecommunication services supported by the PSTN, policymakers recognize that ensuring universal access to reliable Internet service is essential for the health of our economy and society. This means that there is an enduring public regulatory interest in ensuring the health of the Internet ecosystem, and its broad availability (universal access) for all citizens and businesses. This responsibility and its relevance for the overall economy is explicitly articulated in the US National Broadband Plan.[24]

Recognizing that there is an enduring public (regulatory) interest, however, does *not* mean that the appropriate model for regulating the Internet cloud is legacy PSTN regulation. Even in the absence of the growth of the Internet, we would be continuing with our decades long project to dismantle and overhaul traditional PSTN regulation, increasingly transitioning from a command-and-control Public Utility model for regulation to one that relies ever more on market-forces. Earlier examples of this trend include the successive opening of customer premises equipment, long distance telephone service, and local telephone service markets to competition; the transition from rate of return to price cap; and the de-tariffing and de-regulation of a growing range of services, including broadband services.

This transition was motivated by the recognition that competition was viable in a wider-range of PSTN elements and services (making reliance on market-forces a more reasonable alternative), while the burdens of enforcing legacy PSTN rules became increasingly intractable. In addition to the deadweight costs of regulatory bureaucracy and the attenuated incentives for efficient resource allocation that the lack of a profit

[24] See FCC (2010). President Obama has affirmed this position. Speaking for his administration, Susan Crawford commented in a speech on May 14, 2009 that "Broadband is the new essential infrastructure" (see http://www.broadcastingcable.com/article/232506-President_Obama_Focused_On_Broadband.php). Similar positions have been adopted in Europe, where the European Commission has concluded that "widespread and affordable broadband access is essential to realize the potential of the Information Society" (see http://ec.europa.eu/information_society/eeurope/2005/all_about/broadband/index_en.htm); in Australia, where a government report concludes that "ubiquitous, multi-megabit broadband will underpin Australia's future economic and social prosperity" (see http://www.dcita.gov.au/communications_for_consumers/internet/broadband_blueprint/broadband_blueprint_html_version/chapter_one_broadband_as_critical_infrastructure); in Japan, where the Japanese have joined with regional partners to "enable all people in Asia to gain access to broadband platforms" by 2010 (see http://www.dosite.jp/asia-bb/en/pdf/abp005.pdf).

motive implies for government operations,[25] there is a fundamental information asymmetry: the regulated firms generally know much more about market and technical trends and conditions and are much more agile at adapting (unless constrained by cumbersome regulations) than regulators. As the environment gets more complex and information asymmetries amplify, the case for allowing regulated firms greater discretion increases. This allows firms more scope to optimize in the face of a dynamic environment. It also allows firms greater scope to potentially behave in ways that are adverse to the public interest (e.g., abuse any market power they may have). If competition is sufficiently viable and robust, then market forces can constrain these abuses and the delegation of regulatory authority to markets instead of via direct regulation is a win-win proposition: overall efficiency is enhanced while regulatory costs are reduced.

In contrast to the PSTN, the Internet was largely unregulated. It began as an application that existed on top of the PSTN, implemented in equipment and software owned-and-operated by end-users. The Internet was designed as a peer-to-peer packet data transport network that required only very limited intelligence in the network to support end-to-end connectivity at the network layer. (However, there was a lot of network intelligence supporting the switched telephone network that underlay the Internet). Most of the incremental investment to create the Internet was in end-user equipment and applications at the edge, but most of the total investment (when one includes the PSTN) was still associated with the telecommunications infrastructure of the PSTN. Nevertheless, the Internet, like the markets for computer equipment, software, and services, remained largely unregulated. There were thousands of access ISPs, and although there were only a large handful of Tier 1 ISPs, most analysts regarded the Internet as robustly competitive and pointed to the Internet's record for rapid growth and graceful scaling to meet new challenges in the absence of regulation as strong justifications for preserving its unregulated character.

In transitioning from the voice telephony PSTN to the Internet as the new PSTN, we have replaced the basic circuit-switched paradigm with packet-switching. Voice telephony is now just another application on the Internet (VoIP). We have replaced the central office switches with routers,[26] the copper wires with fiber and wireless,[27] and the centralized control of Signaling System 7 (SS7) with distributed/decentralized Internet routing and network management. However, we have also seen the traditional Internet enhanced by adding new access network infrastructures below the narrow waist of the Internet protocols (IP on fiber rather than SONET, mobile wireless, ad hoc networks) and

[25] Government bureaucrats lack profit incentives and market discipline that can give rise to X-inefficiency (Leibenstein, 1966).

[26] Although these switches were essentially special-purpose computers, and today, the functionality of legacy central office switches may be emulated in soft switches hosted on Internet servers.

[27] There is still a lot of copper wire in use, and much of the outside plant investment is in conduit and other assets that remain important even today. Moreover, the transition to fiber or other very-high-capacity last-mile technologies like cable (with DOCSIS 3.0) reignites questions about last-mile bottlenecks and market power.

overlays above (routing, security, content-delivery).[28] Edge-boxes and software applications (Browsers, client applications) have become more capable, adding new functionality like support for end-to-end encryption, modified congestion control, support for caching, and other functionality that is intended to enhance the quality of the user-experience even in the face of the variable performance of the best-effort Internet.[29]

These trends have blurred the traditional boundaries between peer and network-based functionality. Relative to the voice-only PSTN, figuring out what technical functions belong where and how to regulate them poses a much more complex problem for regulators. Ceteris paribus, this increased technical and marketplace complexity strengthens the preference for relying on market-forces relative to direct regulatory oversight.

At the same time, and with the relaxation of regulatory restrictions, we have seen the rise of intermodal facilities-based competition between telephone, cable, satellite, and mobile service providers as the scope of services that can be supported on each platform has converged (so each can offer a mix of voice, video, and data services). We have seen the emergence of new types of all-IP providers like Global Crossing and Level 3. Deregulation also eliminated line-of-business restrictions that limited the ability of local telephone companies to compete in markets for Internet services. As these providers expanded their Internet offerings, the boundary between Internet and telecom assets blurred. Today, the legacy access providers, who were also the dominant facilities-based providers of PSTN infrastructure, are among the largest Internet Service Providers (ISPs). Meanwhile, new types of service providers have emerged that rely on and interconnect with the ISPs like Akamai, Google, Facebook, Netflix, Twitter, and Amazon that provide Internet functionality but are not typically regarded as ISPs.[30]

In this new environment, we confront a quandary. On the one hand, we recognize that the Internet is no longer just an application on the PSTN but is the *new PSTN*, and that means that there is a heightened public interest in regulating the Internet.[31] On the other hand, the Internet ecosystem is fundamentally more complex than the telephony-centric world of the old PSTN, the inefficiencies of legacy PSTN regulation are well-understood, and the prospects for the viability of competition across the Internet ecosystem remain uncertain. The largest access providers, content providers, and overlay network functionality providers have increased their market shares, but many performance indicators suggest competition remains robust. At this point, it seems reasonable to

[28] See Lehr et al. (2006).

[29] See Bauer, Clark, and Lehr (2011) for a discussion of how faster-than-realtime broadband service may be used by streaming media applications to compensate for variable performance over time.

[30] See Labovitz et al (2009) "Arbor Networks Traffic Study" which documents the rise in recent years of the "hyper-giants" as the Internet ecosystem has expanded.

[31] See Lehr, Bauer, and Clark (2012) for further discussion of some of the regulatory issues that are on the FCC's current and prospective regulatory agenda.

expect that the dominant model for regulating the Internet will remain reliance on market forces.[32]

Although primary reliance on markets to govern the Internet as the new PSTN may be inevitable or even desirable, it is important to remember that markets failure may arise in multiple ways. Market power may be excessive (competition is not vigorous enough) *or* competition may fail to be efficient for a number of other reasons. It is important to remember that an excess of market power is not the only market failure that regulation may be called upon to address. For example, there may be fundamental non-convexities that preclude existence of a sustainable pricing equilibrium under competition (e.g., marginal cost pricing fails to recover long run incremental costs[33]); or information imperfections may preclude equilibria supporting efficiency-enhancing quality differentials (e.g. a 'Lemons' problem[34]); or incomplete contracts (e.g., a lack of enforcement mechanisms for service level agreements) may prevent coordination even when it is in everyone's best interests (e.g., a potential free-rider or Prisoner's Dilemma problem).[35] Ceteris paribus, increased complexity would suggest an increase both in the desire for increased reliance on market-forces, but also an increased potential (perhaps) for market power and (more likely) for non-market-power-related market failures.

This is not meant to imply that the increased complexity of the Internet (relative to the Internet of old and relative to the telephony PSTN) warrants more direct regulatory intervention, but only that our decision to rely on market forces comes with a challenge. Markets are not unregulated, they are regulated differently (relative to legacy PSTN public utility regulation).

In the next section, I focus on the policy challenge of ensuring reliability in the new environment of the Internet cloud.

3. Reliability and the Policy Challenge

Reliability means different things in different contexts.[36] At the highest level, reliability implies that systems behave as we expect them to, consistently. Generally, we also assume that a reliable system is one that performs well. A common metric for reliability

---

[32] I say "will" to avoid offering an opinion here as to the desirability of more direct regulatory intervention. I believe a strong case might be made for limited forms of regulatory interventions under certain conditions, but precisely what these might be would take the discussion too far afield.

[33] Fixed, sunk, or shared costs may be a sufficient share of total costs as to preclude any sustainable pricing equilibrium under competition.

[34] See Akerlof (1970).

[35] A free-rider or Prisoner's Dilemma problem may be solved if the players could contract over their actions.

[36] See Lehr, Bauer, Heikenen & Clark (2011) for a discussion of broadband reliability.

is *availability* which is the amount of time that a system is expected to be in-service. It is often expressed as a statistical time measure (e.g., Mean Time to Failure) or the percent of time over some period that the system is available for service (free of failures).

3.1. Reliability in the legacy telephone networks and the Internet

What constitutes a failure depends on what the system is supposed to do. The legacy fixed line PSTN was designed to support anytime/anyone-to-anyone voice telephony. Because telephone service was regarded as critical infrastructure (for business, for public safety, for daily life), it was expected that out-of-service events would be infrequent events for individuals, and for large groups of individuals (or for big chunks of the PSTN) would be very rare. It was expected that the telephone calling experience would be relatively homogeneous and of "good audio quality" across calls (at different times and between any two parties with fixed-line phones).[37] Absent excessive line noise or the occasional fast-busy signal resulting from switch congestion, fixed line telephony provided over the PSTN achieved a very high standard of reliability. It was not uncommon for fixed line telephony to continue to work even when storms had disrupted electric service.[38]

This high availability standard for the PSTN was consistent with the view that the telephone network was essential infrastructure. It was generally accepted that businesses could not function and lives depended on continuously available telephone service (e.g., the ability to call an ambulance in an emergency). Achieving this goal motivated the end-to-end design of legacy telephone networks.

The design of the PSTN was optimized to support voice grade end-to-end circuits with tight technical performance characteristics and low blocking probabilities (i.e., fast busy signals should be infrequent). Interface standards imposed tight latency bounds to ensure that end-to-end latency did not exceed two hundred milliseconds, the threshold for real-time voice telephony to be viable. Core components of the PSTN like the telephone switches were designed for *five nines* (99.999%) reliability, or less than 6 minutes of out-of-service time per year. This required full (1+1) redundancy for core switch and other critical network components. That is, there were full capacity hot spares ready to assume the load if the active unit failed. If the probability of one failing is p, and the failures are independent, then the probability of both failing at the same time is p^2. Adding redundancy provides significant gains in terms of enhanced availability, but comes at a significant cost in terms of increased capital intensity. That expenditure is warranted when a prolonged outage of a single central office switch would pose significant harm on a large number of telephone subscribers.

[37] Like other basic infrastructure, most users do not think about the quality unless there is something wrong, and then they notice in a hurry (pot holes in the road, power surges that burn out electric appliances, water that tastes bad, or dropped telephone calls). This ability to "take the infrastructure for granted" is a design goal.

[38] The copper telephone lines were powered.

Meeting the rigorous technical requirements of supporting voice telephony with the desired high standard for reliability, given the state of technology at the time, required significant centralized, hierarchical control. The out-of-band Signaling System 7 (SS7) network was put in place to support such control, allowing network-wide resource allocation. Additionally, the desire to ensure ubiquitous coverage and connectivity required significant on-going capital investment as the national, and soon global, PSTN was being first established. To manage the technical and economic challenges of supporting the PSTN, it was long believed that monopoly provisioning was desirable. It was only over time, with the advance of technology and market growth, that the technical design requirements and economics of the PSTN were rendered compatible with increased competition and the decentralization and distribution of control that that implied.

With mobile telephony, users tolerated much more variable performance. Calls could not be placed everywhere (coverage was limited), congestion problems were not uncommon, and call quality could be quite variable (dropped calls were a common occurrence). But mobile telephones allowed calling where no fixed telephones were available, and allowed users greater personal control over their calling.[39] Also, competition was built into mobile telephony from the start, with two operators licensed initially in each market in the US. Although the technical architectures for each provider were hierarchical and centrally managed, control of core assets was inherently more distributed and decentralized (across service provider networks). A cost of this was that roaming across provider networks introduced additional quality degradation and might incur additional end-user charges.[40] On the other hand, competition allowed end-users service choices and helped drive down prices, both of which might be regarded as important quality improvements. Viewed in this light, we see that mobile telephony was not so much *less* reliable as *differently* reliable than fixed-line telephony. As we shall see shortly, a similar interpretation was applicable to the Internet.

The purpose of the Internet was to support asynchronous data communications rather than voice telephony. The Internet was not expected to meet the same sort of availability standards as the PSTN. In its original incarnation, it was a research network designed to support data communications between mainframe computers – while this was important, it was not viewed as essential basic infrastructure with the accompanying public interest mandate that implies. The Internet was not supporting business operations or other mission-critical functionality. The best-effort packet-delivery model provided a graceful way for asymmetrically sized *and* delay-tolerant datagrams to share transport capacity. The simple, lightweight Internet Protocols (the "narrow waist") allowed interoperable data connections between heterogeneous peers over variable capacity transmission links, without requiring much in the way of intelligent support from the network. Compared to the complex switches at the core of the PSTN, the routers that switch packets are simple

[39] When away from a user's home fixed-line telephone, payphones, credit cards, or "borrowing" another person's fixed line telephone were cumbersome alternatives to mobile telephony.

[40] For example, 2G digital handsets initially roamed using analog AMPS, an older and lower-quality 1G technology.

packet-forwarding devices. Routers were not typically designed with 1+1 redundancy, and with there greater simplicity, were far less expensive than PSTN switching equipment.

The Internet was not designed to meet tight latency bounds, but to ensure data connectivity across variable quality data links. The packets would get from source to destination, but they may take a while, and follow different routes along the way. In achieving this goal, the Internet did not require significant intelligence from the routers that forwarded packets. They just needed to know where next to send arriving packets. Control and network intelligence were highly decentralized and distributed. There was no provision for centralized information sharing about the overall state of the Internet. When the network was congested, routers buffered packets and when buffers over-flowed, packets were dropped – and end-to-end latency increased. When the network was not congested, sending hosts were permitted to increase their data rate until either they completed sending the desired data or the dropping of packets indicated that congestion was occurring somewhere downstream and sending hosts should slow down and resend packets. This variable-bit-rate capability allowed applications like VoIP or streaming video (e.g., YouTube) to take advantage of higher bit rate opportunities to send improved quality audio/video or faster-than-needed delivery to support buffering to smooth performance when slower-than-needed data rates were available.

By continuously expanding the capacity of links throughout the Internet and moving to bigger and faster routers, the best effort Internet was able to scale to meet exponential traffic growth without realizing debilitating end-to-end latency problems. When congestion threatened, it generally proved more efficient to simply expand capacity than to introduce significant network intelligence to support quality-of-service differentiation. VoIP services like Skype using better codecs are able to offer higher than legacy telephony audio quality, and can be easily extended to introduce interactive multimedia like video-conferencing or text/file sharing. The potential to expand functionality was an original driver for computer-based telephony, but in the 1990s when mass-market VoIP services took off, there was the added attraction of "free" telephone calling.[41]

Over time, and as noted above, the architectures of legacy electronic communication networks and the Internet have converged. Historically, silo-based service provider networks have moved towards a common architecture with the broadband Internet as the common platform. While single-service, best-effort transport is still the dominant mode for exchanging Internet traffic between ISPs, the Internet ecosystem has grown substantially more complex both from a technical and business/industry structure perspective. New technical functionality and service capabilities are being supported as intelligence and cloud-based services grow.

[41] Legacy long distance telephone calls were typically prices on a per minute of use basis that included significant additional regulatory-mandated "access" charges. For international calls, these so-called "settlement" charges could be quite large. With flat-rate (volume insensitive) dial-up Internet subscriptions used over flat-rate telephone lines which became the norm in the US in the 1990s, using VoIP either end-to-end or in the network provided an arbitrage opportunity to bypass those charges/

3.2. Reliability in the Internet Cloud

Whereas the legacy Internet was a general purpose *utility* packet transport network, the emerging cloud is that plus a platform for utility computing and storage. This adds complexity. More users with different goals in using the Internet (to make telephone calls, to watch movies, or to access emergency services) in different ways (real-time or delay-tolerant) and with different tolerances for performance-based prices may legitimately have very different perspectives on what constitutes an appropriate level of reliability.

In Lehr, Bauer et al. (2011), we discussed what this means for broadband reliability, suggesting at least three ways in which a consumer might regard their broadband service as being reliable: (1) performance metrics (e.g., probability bitrates are in some expected range, potentially exceeding some minimum threshold that would identify a service failure event); (2) connectivity metrics (e.g., the ability to connect to Internet servers); and (3) core service availability metrics (e.g., availability of core services like email or DNS). As was common with traditional telephony service monitoring, data could be tracked across a large sample of subscribers on competing service provider networks and benchmarked against appropriate standards. With millions of subscribers and service-events occurring all the time, even a very high standard of service reliability will yield statistically significant samples of failure events that might be used to track service quality.

From a policy-perspective, the challenge of ensuring adequate broadband reliability amounts to a customer-protection activity, akin to ensuring truth in advertising, product safety, and a well-functioning market of quality-differentiated services.[42] An extensive framework of standards, regulations, and reporting requirements were established over time to provide such customer-protection oversight for legacy telephone service. The need for such oversight was motivated in no small part by the almost $8 billion dollar per year subsidy program designed to promote and secure affordable universal service for basic telephony service.

With the transition to broadband and a new technology, encompassing an expanded range of services, and involving a large number of new players that have been previously exempt from legacy telecom regulation, the challenges of designing and implementing such a consumer protection policy framework for the Internet cloud is daunting. Moreover, because management and the ownership of assets is more decentralized and distributed in the Internet cloud (and hopefully will remain so if viable competition is to be sustained), there will be multiple competing and complementary providers of key components for the end-to-end system. Diagnosing performance problems, assigning responsibility, and implementing remediation is more challenging in such an environment. As our earlier work on speed measurement pointed out, end-to-end performance may suffer because of problems in the end-users home network (e.g., a misconfigured PC or poor WiFi access connection), the last-mile access network, or problems further

[42] Broadband services are offered in quality/usage tiers, with higher-priced,

upstream (e.g., a congested content server).[43] And Internet speed measurement is much easier than evaluating Internet reliability.

As discussed in Lehr, Baur, and Clark (2012), the added complexity will perforce drive us to rely on market forces rather than legacy PSTN regulatory models. Accomplishing that goal will depend critically on the markets' abilities (potentially aided by policymakers) to aggregate and disseminate appropriate information to end-users and service providers across the value chain if the markets are to work efficiently; and should market failures require regulatory intervention, permit regulators to intervene effectively.

With respect to access to cloud computing or storage resources that may be made available in data centers distributed across the Internet, like raisins in a muffin or intrinsic to the fabric of a Future Internet Architecture,[44] appropriate standards for reliability will need to focus not just on availability but also consistency and accuracy of access, with heightened requirements for data security and protection. Legacy data centers were built from meshes of multiple computers which provided "N+M" redundancy. Data was replicated across multiple computers and drives, with smaller N and larger M providing greater protection and security but at a higher cost in committed resources. With large N and small M, data was stored over many computers, any of which could fail with a fairly high-frequency (like low-end routers in the Internet) without threatening a loss of data. Over time, data centers have evolved into rack-mounted systems where multiple CPUs share more reliable and efficient power supplies. While this offers performance improvements, it requires back-up rack power supplies (i.e., 1+1 redundancy for certain key data center components) since the failure of these new more powerful power supplies would no longer mean the loss of a single CPU, but all of the CPUs supported by that rack. The data centers begin to look like large computers themselves.

Reliable access to data is further achieved by distributing copies of the data across multiple data centers. A fire that destroys a data center in Seattle would leave the one in New York intact. Moreover, a user in New York could generally access the data from the data center in New York more quickly than if the data had to be pulled from Seattle. Providing such redundancy presents fundamental challenges that are embodied by the *CAP Theorem* that says you can guarantee, at most, two out of the three desirable properties of distributed databases: consistency, availability, and partition tolerance. When someone makes a change to the data in New York, that change needs to be reflected in the database in Seattle. Ensuring that the two copies of data are consistent, while allowing users to access the same data (partition tolerance) at the same time (high

[43] See Bauer, Clark, & Lehr (2010).

[44] For example, the FIA Nebula Project (Smith et al., 2011) is an architecture for a hyper-reliable cloud that would offer a sufficiently reliable and secure service to support remote monitoring and control of a diabetic's insulin pump via the cloud. And, MobilityFirst (http://mobilityfirst.winlab.rutgers.edu/) would provide support for disruption tolerant networking and access to in-network storage and computing resources spread across the Internet.

availability) is not possible.[45] Many database service providers sacrifice consistency (serve the latest available version of the database) to enable high availability. As we move to cloud-based resources, we will need to integrate metrics that properly capture such trade-offs.

While we expect to principally manage reliability (in the customer protection sense) via market forces, there are many ways in which markets might fail, and the range of potential failures is exacerbated by the increased distribution and decentralization of control and functionality across enterprises and end-user networks and equipment. It is far from clear what intelligence or functionality should be placed where, what may need to be regulated, and what should be insulated from the market distortions regulation imposes. As noted above, today much of the focus for broadband regulation is on last-mile access providers who may be deemed to have market power. In the future, it is not unreasonable to hypothesize that important functionality might be provided by data center service providers or the entities that provide identity management services for authenticating users and services or some other critical functionality associated with accessing cloud resources. It is far from clear what (if any) regulatory entities have or should have jurisdiction over which firms. While Google (cloud resources), Akamai (content delivery network), and Verizon (last-mile access and transport) all provide complementary services that may each be critical for the reliable operation of the Internet cloud, only Verizon as a telecommunications service provider is subject to significant regulatory oversight from the FCC. These providers do not have explicit agreements with each other that help ensure that the reliability of the overall Internet is ensured. When service failures occur, figuring out where the problem arose and who is responsible is quite complicated. Neither the alternative of regulating everyone nor the alternative of deregulating everyone seem desirable, but retaining asymmetric regulation distorts market incentives and potentially adversely impacts innovation and investment.

The problem of designing an appropriate customer protection framework to ensure appropriate reliability in the Internet cloud is not a problem with a single solution. It will require collaboration among multiple regulatory authorities (e.g., competition authorities like the Department of Justice, communications regulators like the FCC, and commerce regulators like the FTC) as well as key stakeholders. The key stakeholders will include the service providers as well as end-users and other government interests like the security/public safety communities and international trade and standardization communities. It will remain a work-in-process.

In the next section, I consider some of the special problems associated with ensuring reliability in the Internet core.

4. Reliability in a Hyper-reliable Core

[45] Strictly speaking, partition tolerance refers to the ability of the system to tolerate partitions without failure. That is the distributed database will continue to work even if the distinct database servers cannot communicate with each other (i.e., the distributed database system is partitioned).

Today's data centers composite multiple computers that are each much faster and with denser storage than those that were in service only a few years ago. The 1GigE links that were put in place between servers in the data centers after 2000 are being upgraded to 10GigE connections. The increase in capacity at the edges drives a concurrent need for even higher capacity connections in the core. This is driving the demand for 40GigE links and beyond, and the ability to switch packets across links at line speeds.

Today, among the largest core routers currently deployed are the Cisco CRS-1s. These "carrier-grade" routers are deployed by large service providers and enterprise customers to support traffic switching in the core of the Internet. The CRS-1 is expandable to support up to 72 40-Gbps line cards for a total switching capacity of 92Tbps.[46] When fully loaded, these boxes consume 1000s of watts of power, imposing significant air conditioning loads and requiring strong structures to support the routers and all of the associated paraphernalia. These boxes are much closer to the telecom switches of old in complexity, cost, and infrastructure requirements than to the low-cost simple packet-forwarding routers of old. If one of these boxes is out of service then the volume of traffic and potential loss to businesses and end-users could be extremely large. Like the switches of old, these mega-routers need 1+1 redundancy capabilities to ensure adequate reliability.

To continue to meet the needs for increased traffic growth and keep ahead of ever-expanding capacity in edge-networks (data centers and customer premises),[47] router functionality needs to be spread over multiple boxes since the power density for a single box is getting too high. This requires developing a distributed software architecture that is sufficiently reliable and robust to look to the other network elements that rely on these core routers as if it were a single router. Allowing capabilities for deploying software updates, routine router maintenance, and changes or additions to router feature sets – all tasks that are critically important in today's dynamic cloud services markets – poses significant technical challenges. The goal is to design such systems that offer even better than five nine's reliability -- in effect, that approach 100% availability in the limit.[48]

From a policy perspective, meeting such a policy goal of hyper-reliability presents a number of interesting problems. When systems get that reliable, we never observe failures. Potential failures are so rare as to be "Black Swans."[49] This poses a number of challenges for efficient markets. First, estimating the probability of such rare events is very difficult. Reasonable analysts might have widely different estimates. Moreover, estimating the potential harm or costs from moderating either the harm (in the unlikely

[46] http://www.cisco.com/en/US/products/ps5842/index.html.

[47] Today's core router is tomorrow's edge router. The big service providers like Sprint, AT&T, and Level 3 migrate the core routers they purchased a few years ago to the edge as they upgrade their cores to take account of recent advances in the technology. The opportunity to do so extends the life of equipment purchases, which reduces the economic cost.

[48] Many of these issues are being addressed as part of the FIA Nebula project in the design of the NCore, a hyper-reliable core routing architecture (see Smith et al, 2011).

[49] Taleb (2007).

event of a failure) or reducing the likelihood of a failure (making a rare event rarer still) is even more difficult to estimate. With such a large potential space for valuation differences, it may be difficult to contract over and credibly commit to appropriate ex ante or ex post actions directed at ensuring system reliability.

Second, and in light of these information problems, there are strong incentives to free-ride or otherwise defect from commitments. For example, one way to achieve such high reliability is to provide better than 1+1 redundancy, but it is less expensive to simply claim you have secured the additional redundancy than to actually make those investments.[50] An analogous phenomenon incurred in financial markets when bankers sold each other portfolio insurance that ultimately failed to provide the risk protection that it was represented to provide. Ensuring that all of the participants and components correctly implement the procedures and processes required to ensure the high degree of reliability is very difficult.

Coglianese and Mendelson (2010) discuss some of the regulatory approaches that may be required to help ensure reliability of components that require hyper-reliability. One of the things that policymakers can do is audit the processes employed by firms – do they have architectures and business operational plans that might reasonably be expected to deliver the anticipated levels of reliability? Do they have audit processes to make sure that the plans and architectures are being followed? Do they have processes in place to learn and adapt if they find earlier assumptions/plans require modification? In the event that a failure occurs, do they have credible response plans? Do they have the resources they need to address this? And so on. The focus of regulation cannot be on outcome performance since bad outcomes (failures) are not supposed to be observed (except extremely rarely), and enforcement triggered by such outcomes is not credible and so cannot influence ex ante behavior. The focus of regulation needs to be on the inputs and processes used.

In the legacy PSTN, this problem was confronted in the 1980s when a series of well-publicized software bugs in SS7 resulted in widespread outages. In response, the industry engaged in industry self-regulation efforts, forming the Network Reliability Council (NRC). The NRC provided a forum for industry participants to formulate and share best-practice plans and coordinate on reliability (and recovery) planning. Such industry self-regulation efforts are a common response to a heightened perception of risk from "Black Swan" catastrophes. Other examples include the Bhopal Chemical Plant disaster in 1984 and the Three Mile Island Nuclear Reactor meltdown in 1979. In the former case, significant loss of life occurred, while in the latter the risk of significant loss of life was seriously threatened. Interestingly, the subsequent efforts at industry self-regulation coordinated by policymakers in response was more successful in the case of nuclear power than for the chemical industry.[51] One potential reason for this was the presence of a shared sense of fate with nuclear power (i.e., another disaster would spell doom for the

[50] Varian (2001) presents a simple model that shows how incentives to invest in the optimal level of reliability often fail to exist.

[51] See Coglianese and Mendelson (2010).

entire industry) that was lacking in the chemical industry (i.e., chemical firms operate in distinct markets that would allow those not affected by a disaster to avoid responsibility). An interesting question is whether the Internet cloud might look more like the chemical industry than the nuclear power industry.[52] Addressing these and other issues will present policymakers with difficult challenges in the years to come.

5. Conclusions

The Internet is morphing into a cloud computing utility. It has evolved from a packet transport network riding on top of the PSTN to become they platform for all electronic communications, and increasingly the host to general purpose computing and storage resources. In assuming this role, the Internet becomes basic infrastructure and thereby attracts an enduring public interest in ensuring its universal accessibility and availability.

Relative to the PSTN and Internet of old, the new cloud utility Internet is substantially more complex and important to our economy and society. As we evolve toward a world of pervasive computing, toward the world of the Internet of Things, the Internet cloud will host an ever growing range of services that are mission critical to our daily lives and business interests.

This new environment will entail a heightened interest in ensuring reliability of Internet services analogous to the earlier public policy interest in ensuring reliability in the legacy telephone network. By contrast, however, control is much more decentralized and distributed at the technical, business, and regulatory levels in the Internet cloud.

It is clear that principal responsibility for managing Internet reliability will depend on market forces, with regulators potentially helping to steer markets when failures are detected. While prospects for market failures may increase with the Internet cloud, identifying these and managing them will pose strong challenges.

The policy challenge of ensuring Internet reliability may be devided into the customer protection challenge of ensuring adequate reliability in retail services, and the need to ensure against systemic failures in a hyper-reliable core. These two focus areas are equally complex but distinctly different in the types of challenges they confront.

Addressing both reliability challenges will require on-going collaboration across the value chain and will require new forms of performance monitoring and metrics. We are only now beginning the difficult task of engaging the multidisciplinary expertise required to address these significant challenges.[53]

[52] Just as the PSTN of old was more heavily regulated than the Internet cloud of today, the nuclear power industry is significantly more heavily regulated than the chemical industry.

[53] See Lehr (2012).

6. References

Armbrust, M., A. Fox, R. Griffith, A.D. Joseph, R.H. Katz, A. Konwinski, G. Lee, D.A. Patterson, A. Rabkin and I. Stoica (2009) "Above the Clouds: A Berkeley View of Cloud Computing," Technical Report UCB/EECS-2009-28, EECS Department, University of California, Berkeley.

Bauer, S., D. Clark, and W. Lehr (2011) "Powerboost," Proceedings of IEEE HomeNets'11, Toronto, Canada, August 2011. (pdf=http://people.csail.mit.edu/wlehr/Lehr-Papers_files/home22-bauer clark lehr.pdf)

Bauer, S., D. Clark, and W. Lehr (2010), "Understanding Broadband Speed Measurements," MITAS Working Paper, June 2010, available at: http://mitas.csail.mit.edu/papers/Bauer_Clark_Lehr_Broadband_Speed_Measurements.pdf

Botterman, M. (2009), "Internet of Things: an early reality of the Future Internet Workshop Report," prepared for the European Commission, Information Society and Media Directorate, 10 May 2009 (available at: http://ec.europa.eu/information_society/policy/rfid/documents/iotprague2009.pdf).

Coglianese, C. and E. Mendelson (2010) "Meta-Regulation and Self-Regulation," M. Cave, R. Baldwin and M. L. (eds), *The Oxford Handbook on Regulation.* Oxford: Oxford University Press.

FCC (2010) "Connecting America: The National Broadband Plan," Federal Communications Commission, Washington, DC, March 2010 (available at: http://www.broadband.gov/).

Jianli, P., S. Paul and R. Jain (2011) "A Survey of the Research on Future Internet Architectures," *Communications Magazine, IEEE*, 49(7), 26-36.

Isenberg, D. (1997), "Rise of the Stupid Network," manuscript, available at: http://www.rageboy.com/stupidnet.html.

ITU (2005) "The Internet of Things," International Telecommunication Union, *7th ITU Internet Report.* Geneva, Switzerland.

Labovitz, C., S. Lekel-Johnson, D. McPherson, F. Jahanian, J. Oberheide and M. Karir (2009) "Atlas Internet Observatory Report (Arbor Networks, University of Michigan, and Merit 2-Year Internet Traffic Study)," North American Network Operators Group Meeting 47 (NANOG47), October 2009, available at: http://www.nanog.org/meetings/nanog47/abstracts.php?pt=MTQ1MyZuYW5vZzQ3&nm=nanog47.

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, and D. Clark (2012), "Measuring Broadband Performance when Broadband is the New PSTN," MITAS Working Paper, May 2012.

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.

Lehr, W., D. Clark, P. Faratin, R. Sami, and J. Wroclawksi (2006) "Overlay Networks and Future of the Internet," Communications and Strategies, no. 63 (3rd Quarter 2006) 1-21.

Leibenstein, H., "Allocative Efficiency vs. X-Inefficiency," *American Economic Review*, vol. 56 (June 1966) 392-415

Lucky, R. (1997) "When Is Dumb Smart?" IEE Spectrum, 34(11), 21-21.

Mell, P. and T. Grance (2009) "The NIST Definition of Cloud Computing," *http://www.nist.gov/itl/cloud/upload/cloud-def-v15.pdf.*

Odlyzko, A. (1998) "'Smart' and 'Stupid' Networks: Why the Internet Is Like Microsoft," *netWorker*, 2(5), 38-46.

Rimal, B., E. Choi and I. Lumb (2010) "A Taxonomy, Survey, and Issues of Cloud Computing Ecosystems Cloud Computing," N. Antonopoulos and L. Gillam, Springer London, 21-46.

Smith, J., R. Broberg, A. Agapi, K. Birman, D. Comer, C. Cotton, T. Kielmann, W. Lehr, R. VanRenesse, and R. Surton (2011) “Clouds, Cable And Connectivity: Future Internets And Router Requirements,” Proc. 2011 Cable Connection Spring Technical Conference, June 14-16, Chicago, IL.

Taleb, N. (2007), *The Black Swan: The Impact of Highly Improbable Events*, New York: Random House.

Thierer, A. (2006) "Are 'Dumb Pipe' Mandates Smart Public Policy? Vertical Integration, Net Neutrality, and the Network Layers Model," in *Net Neutrality or Net Neutering: Should Broadband Internet Services Be Regulated*, edited by T. M. Lenard and R. J. May, Springer US, 73-108.

Varian, H (2001), “System Reliability and Free Riding”
(available at: http://people.ischool.berkeley.edu/~hal/Papers/2004/reliability).

Zhu, J. (2010) "Cloud Computing Technologies and Applications Handbook of Cloud Computing," B. Furht and A. Escalante, Springer US, 21-45.