



Editors: **Scott Bradner** • sob@harvard.edu
kc claffy • kc@caida.org

The (un)Economic Internet?

The Internet Economics track will address how economic and policy issues relate to the emergence of the Internet as critical infrastructure. Here, the authors provide a historical overview of internetworking, identifying key transitions that have contributed to the Internet's development and penetration. Its core architecture wasn't designed to serve as critical communications infrastructure for society, but developed far beyond the expectations of the original funding agencies, architects, developers, and early users. The incongruence between the Internet's underlying architecture and current use and expectations, however, means we can no longer study Internet technology in isolation from the political and economic context in which it is deployed.

kc claffy
and Sascha D. Meinrath
*Cooperative Association for Internet
 Data Analysis*

Scott O. Bradner
Harvard University

This article kicks off *IC*'s new series on policy, regulatory, and business-model issues relating to the Internet and its economic viability. These articles will explore a range of topics shaping both today's Internet and the discourse in legislatures and deliberative bodies at the local, state, national, and international levels in pursuit of enlightened stewardship of the Internet in the future.

Mindful of Internet connectivity's fundamental import for advanced as well as emerging economies and its day-to-day irrelevance for the unconnected vast majority of human beings, pieces for this series will cover technology as well as political, economic, social, and historical

issues relevant to *IC*'s international readership. In this inaugural article, we provide a historical overview of internetworking and identify topics that need further exploration – topics we particularly encourage authors to cover in future articles in this series.

A History of Internet (un)Economics

The modern Internet began as a relatively restricted US government-funded research network. One of the most revolutionary incarnations of this network, the early Arpanet, was limited in scope – at its peak, it provided data connectivity for roughly 100 universities and govern-

ment research sites. In the decades since, a few key transitions have been critical in radically transforming this communications medium. One of the most important of these critical junctures occurred in 1983, when the Arpanet switched from the Network Control Program (NCP) to the (now ubiquitous) Transmission Control Protocol and Internet Protocol (TCP/IP). This switch helped change the Arpanet's basic architectural concept from a single specialized infrastructure built and operated by a single organization to the "network of networks" we know today. Dave Clark discusses this architectural shift in his 1988 *Computer Communications Review* paper, "The Design Philosophy of the DARPA Internet Protocols."¹ He wrote that the top-level goal for TCP/IP was "to develop an effective technique for multiplexed utilization of existing interconnected networks."

During this same period, network developers chose to support data connectivity across multiple diverse networks using gateways (now called routers) as the network-interconnection points. Preceding communications networks, such as the telephone system, used circuit switching, allocating an exclusive path or circuit with a predefined capacity across the network for the duration of its use, regardless of whether it efficiently used the circuit capacity. Breaking with traditional circuit-switching network design, early internetworking focused on packet switching as the core transport mechanism, facilitating far more economically as well as technically efficient multiplexing of existing networking resources. In packet-switching networks, nonexclusive access to circuits is normative (although companies still sometimes buy dedicated lines to run the packet traffic over); thus, no specific capacity is granted for specific applications or users. Instead, data is commingled with packet delivery occurring on a "best effort" basis. Each carrier is expected to do its best to ensure that packets get delivered to their designated recipients, but no guarantee exists that a particular user will be able to achieve any particular end-to-end capacity. In packet-switching networks, capacity is more probability-based than statically guaranteed. Internet data transport's best-effort nature has caused growing tension in regulatory and traditional telephony circles. Likewise, as the Internet becomes an increasingly critical communications infrastructure for business, education, democratic discourse, and civil society in general, the need to systematically analyze core functionality and potential prob-

lem areas becomes progressively more important.

Early developers couldn't have foreseen the level to which the Internet and private networks using Internet technologies have displaced other telecommunications infrastructures. It wasn't until the mid 1990s that visionaries such as Hans-Werner Braun started warning protocol developers that they needed to view the future Internet as a global telecommunications system that would support essentially all computer-mediated communications. This view was eerily prescient, yet core Internet protocols haven't evolved to meet increasing demands and are essentially the same as they were in the late 1980s.

A growing number of researchers are convinced that without significant improvements and upgrades, the Internet might be facing serious challenges that could undermine its future viability. Features such as network-based security, detailed accounting, and reliable quality-of-service (QoS) control mechanisms are all under exploration to help alleviate perceived problems. In response to these concerns, the International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) Next Generation Networks study group (NGN; www.itu.int/ITU-T/ngn/) is working to define a very different set of protocols that would include these and other features.

Security: Not the Network's Job

Various people have offered explanations regarding the lack of security protocols in the Internet's initial design. Clark's seminal paper doesn't mention security, nor does the protocol specification for IP.² Because the network itself doesn't contain security support, the onus has fallen to those who manage individual computers connected to the Internet, to network operators to protect Internet-connected hosts and servers, and to ISP operators to protect their routers and other infrastructure services. Services such as user or end-system authentication, data-integrity verification, and encryption weren't built into the core Internet protocols, so they're now layered on an infrastructure that isn't intrinsically secure. Currently, few existing studies examine the potential economic rationales for this current and continuing state of affairs and the ramifications for the infrastructure's efficiency, performance, and sustainability.

QoS: Too Easy to Go Without

The original IP packet header included a `type of`

service field to be used as “an indication of the abstract parameters of the quality of service desired.”² This field, later updated by Differentiated Services,³ can define priority or special handling of some traffic in some enterprise networks and within some ISP networks, but it’s never seen significant deployment as a way to provide QoS across the public Internet. Thus, the QoS a user gets from the Internet is typically the result of ISP design and provisioning decisions rather than any differential handling of different traffic types. Thus far, “throwing bandwidth at the problem” has proven to be a far more cost-effective method for achieving good quality than introducing QoS controls.⁴

Yet, what happens if conditions change so that overprovisioning is no longer a panacea? The day-to-day quality most users experience from their broadband Internet service is good enough, for example, to enable voice-over-IP (VoIP) services such as Skype and Vonage, which compete favorably with plain old telephone services. However, the projected explosive growth of video and other high-bandwidth applications might increase congestion on parts of the current infrastructure to the point that special QoS mechanisms could be required to maintain usable performance of even the most basic services.

Accounting: A Missing Goal

In their first paper on TCP/IP, Vint Cerf and Robert Kahn felt that accounting would be required to enable proper payments to Internet transport providers.⁵ More than a decade later, Clark echoed this requirement in his Design Philosophy paper. In his listing of second-level goals affecting the TCP/IP protocol suite’s design, the seventh and final goal was that “the resources used in the Internet architecture must be accountable.”¹ As with security, however, no evidence exists that accounting was ever an operational goal for DARPA in developing and running the Arpanet, nor is there any indication that accounting was a goal for the US National Science Foundation (NSF) in the follow-on NSFnet. Indeed, if a government agency is paying in bulk for the entire system, accounting itself is a technical as well as economic inefficiency. Consequently, today’s Internet has no built-in accounting mechanisms, making it fundamentally different from previous circuit-switched networks and creating substantial debate as to how to fairly meter and charge for broadband infrastructure and usage.

The End-to-End Model’s Impact

The Internet’s architecture and initial deployment used an end-to-end (e2e) model of connectivity. Jerome Saltzer, Darren Reed, and Dave Clark first discussed elements of this model in their 1981 paper, “End to End Arguments in System Design.”⁶ The general rationale behind the e2e model is that the network doesn’t have to know the applications running on it because it’s simply a neutral transport medium. This neutral traffic handling has enabled explosive innovation in edge services and applications over the past several decades. For example, an application developer doesn’t need to get permission from ISPs, or pay them anything other than normal service fees, to deploy a new application. Likewise, network operators don’t know what applications are running on their networks, nor can they participate in the value chain for these applications.

Clark once said in an Internet Research Task Force (IRTF) presentation that the Internet “did not know how to route money.” He held that there was no efficient way for an independent service provider to cost/profit share with an ISP so that the ISP would provide better service to users who weren’t direct customers. The Internet economic model has always been “sender keeps all” – an ISP serving a particular customer keeps all of the revenue from that customer without regard to where his or her traffic is going. In many countries, no regulations covering peering relationships among providers exist, thus leaving ISPs on their own to decide whether to peer. Typically, especially in the commercial sector, these decisions are based primarily on immediate business interests with little public policy input.

Telephone Regulation

Many parts of the world have well-developed telephone networks. However, this robustness often comes at a cost to the networks’ users. Regulations requiring that telephone carriers ensure the reliability and price controls that they themselves demand in order to guarantee a rate of return on their investment boost service prices. A less regulated and price-controlled future for telephone carriers seems inevitable. It remains to be seen if they will be as willing to put significant resources into reliable infrastructures and the personnel needed to run them if competition sets the prices rather than regulation. Likewise, the intersections among regulatory structures, pricing, service quality, and

interconnectivity with other data communications services are still wide open for exploration.

Internet (non)Regulation

Although open-access regulations on PSTN trunks were essential to the development of the Internet, regulation of Internet service itself has remained largely laissez-faire. For example, until recently US ISPs didn't have to register with the government before offering services, and governments typically haven't regulated either ISPs' service offerings or their service quality. Yet, government attitudes toward the Internet are beginning to change. For example, the first major US regulation covering ISPs – the Communications Aid to Law Enforcement Act (CALEA) – goes into effect in May 2007 and requires ISPs to register with the government and be able to track users. Already, numerous regulators have begun investigating the viability of mandating that ISPs install QoS mechanisms because they believe that this is required to ensure that the Internet can reliably help emergency workers respond to natural or man-made disasters. Unless the network research community fundamentally changes our approach, future regulations will be considered, ratified, and implemented with little peer-reviewed empirical research documenting their likely technical and economic effects.

Internet Measurement

Because no systemic measurement activities exist for collecting rigorous empirical Internet data, in many ways, we don't really know what the Internet actually is. Thus, we don't know the total amounts and patterns of data traffic, the Internet's growth rate, the extent and locations of congestion, patterns and distribution of ISP interconnectivity, and many other things that are critical if we're to understand what actually works in the Internet. These data are hidden because ISPs consider such information proprietary and worry that competitors could use it to steal customers or otherwise harm their business. The information might not even be collected at all because no economic incentive exists to do so, nor do any regulations require its collection.

The Changing ISP Community

The original Internet was provided for "free" by governments and government-supported research institutes. In the US, direct federal government support for the backbone and attached regional

networks ended in the mid 1990s, although tax incentives continued to promote private as well as public infrastructure development. However, complete private ownership of the entire US Internet infrastructure hasn't yet occurred. Today, many states and consortiums continue to run their own networks, usually restricting who can use them in some way – most often to educational and research constituencies – and connecting these research networks into the global Internet.

Historically, most telephone carriers weren't interested in offering Internet service to individual homes or to the business community. Even when a telephone carrier did offer such services, it was usually through a separate division that company management often viewed as outside its basic mission. Instead, commercial ISPs often provided Internet service by leasing telephone carrier facil-

Today's Internet has no built-in accounting mechanisms, making it fundamentally different from previous circuit-switched networks.

ities or by setting up dial-up modem banks to interconnect with the plain old telephone system.

After Internet infrastructure commercialization began, the Internet service provision business model was predicated on making a profit by charging customers more than it cost an ISP to run the service. This business model is problematic given that Internet connectivity is a commodity service, with most customers caring more about low prices than claims of better quality or advanced services. Thus, competition, along with undefined accounting mechanisms for the new technology, drove prices below sustainable levels for most ISPs. The resulting massive provider consolidation is still in play, but customers are no more willing to pay high prices for Internet service in the new environment. A survey quoted in a 2002 US Federal Communications Commission (FCC) report determined that only 12 percent of customers would be willing to spend US\$40 per month for broadband Internet service.⁷

Meanwhile, telephone carriers began to offer broadband Internet service directly over their own

facilities, particularly in higher-income, urban residential markets, directly competing with commercial ISPs who had been offering service via overlays on the telephone carriers' facilities. Paralleling telephone carriers' entry into the broadband market, cable TV companies also began providing broadband Internet service over their own facilities. Today, most residential customers get Internet access service from telephone carriers or cable TV companies, in which the Internet business is only part of their service offerings. Although standardized, "cookie cutter" service packages have hampered what customers can do with their network services, no one has yet studied how the shift of broadband service provision from ISPs to phone and cable TV companies impacts Internet service quality and dynamism.

The (un)Economic Internet

All of these factors form the background to the current debates on the Internet's future, often lumped under the heading of "network neutrality" – a discussion with far wider and deeper implications than that label conveys. The key question at the root of the debate is whether viable economic models exist for Internet service provision, given the high cost of deploying physical infrastructure and operating the network, coupled with ISPs' current inability to participate in the much more profitable application value chain. Further complicating analyses of these factors are the internally conflicted regulatory agencies, tasked with ensuring both that the general public's best interests are kept foremost and that the "free market" be allowed to innovate and police itself.

Many first-generation ISPs went out of business because they couldn't find a successful business model given constraints from both the Independent Local Exchange Carriers (ILECs) and their own customer base. The current generation of telephone-carrier-based ISPs is asking regulators for the ability to charge differentially based on the applications used and content consumed. These companies claim that they won't be able to afford to deploy the necessary infrastructure upgrades without this type of discriminatory pricing. Their opponents worry that letting ISPs decide which applications can use their facilities and at what cost would destroy the very environment that enabled the creation of today's Internet.

Meanwhile, a growing number of communities have decided that they aren't well-served by exist-

ing ISPs (generally meaning the telephone carriers) and have decided to build their own Internet infrastructures. This is similar to what the academic community undertook immediately after NSF retired its (NSFnet) backbone, and to what many state education networks – such as California's Corporation for Education Network Initiatives (CENIC), Florida's Lambda Rail, and New Mexico's LambdaRail – are doing. There is a growing, though far from universal, view that basic Internet connectivity is a fundamental civil society requirement (much like roads, schools, and so on) and that governments should thus ensure universal access to this valuable resource.

Another scenario that will deeply alter the economics is commercial ISPs' leasing of government-funded infrastructure. These public-private partnerships are currently being developed in thousands of communities around the globe. Objective empirical analyses of the various business models for providing Internet infrastructure access, including empirical validation of inputs, outputs, and interacting technological factors, is one of the least understood yet vital aspects of this emerging critical infrastructure.

The Internet Economics series will focus on the ongoing debates surrounding issues of economics and policy, and how they are influenced by, and should influence, science and engineering research. We are heading into another decade of tremendous innovations, not only in wireless connectivity and high-bandwidth applications and services that use it but in the business models that will lead to their success or failure. Gaining a better understanding of the tussles (known outside our field as "economics and politics") among providers, users, and regulators of Internet access, services, and applications will help to ensure enlightened progress on security, scalability, sustainability, and stewardship of the global Internet in the 21st century and beyond. □

References

1. D. Clark, "The Design Philosophy of the DARPA Internet Protocols," *Proc. Symp. on Comm. Architectures and Protocols*, 1988, pp. 106–114.
2. J. Postel, *Internet Protocol*, RFC 791, Sept. 1981; www.ietf.org/rfc/rfc791.txt.
3. T. Li and Y. Rekhter, "A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)," IETF RFC

2430, Oct. 1998; www.ietf.org/rfc/rfc2430.txt.

4. P.C. Fishburn and A. Odlyzko, "The Economics of the Internet: Utility, Utilization, Pricing, and Quality of Service," *Proc. 1st Int'l Conf. Information and Computation Economics* (ICE 98), ACM Press, 1998, pp. 128–139.
5. V. Cerf and R. Kahn, "A Protocol for Packet Network Interconnection," *IEEE Trans. Communications*, vol. 22, no. 5, 1974, pp. 637–648.
6. J. Saltzer, D. Reed, and D. Clark, "End-to-End Arguments in System Design," *Proc 2nd Int'l Conf. Distributed Computing Systems*, ACM Press, 1981, pp. 509–512.
7. US Federal Communications Commission, *Third Report on the Availability of Advanced Telecommunications Capability Services*, tech. report, Feb. 2002; http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-02-33A1.pdf.

kc claffy is founder and director of the Cooperative Association for Internet Data Analysis (CAIDA) and adjunct associate professor in the Department of Computer Science and Engineering at the University of California, San Diego. claffy has a PhD in computer science from UCSD. Contact her at kc@caida.org.

Scott O. Bradner is senior technical consultant at the Harvard University Office of the Assistant Provost for Information Systems. He's also a member of the Internet Engineering Steering Group, vice president for standards for the Internet Society, and a member of the IEEE and the ACM. Contact him at sob@harvard.edu.

Sascha D. Meinrath is the Director for Municipal and Community Networking for the CAIDA COMMONS project and a telecommunications fellow at the University of Illinois, Institute for Communications Research, where he is finishing his PhD. His research focuses on community empowerment and the impacts of participatory media, communications infrastructures, and emergent technologies. Meinrath has a masters in psychology from the University of Illinois, Urbana-Champaign. He is the cofounder and executive director of CUWin, one of the world's leading open source wireless projects. Contact him at sascha@caida.org.