

**Bandwidth is Political:
Reachability in the Public Internet**

Nancy Paterson

A DISSERTATION SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

JOINT GRADUATE PROGRAM IN COMMUNICATION & CULTURE
YORK UNIVERSITY,
TORONTO, ONTARIO

DECEMBER 2009

BANDWIDTH IS POLITICAL

By Nancy Paterson

A dissertation submitted to the Faculty of Graduate Studies of York University in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Copyright © 2009 by Nancy Paterson.

Permission has been granted to: a) YORK UNIVERSITY LIBRARIES to lend or sell copies of this dissertation in paper, microform or electronic formats, and b) LIBRARY AND ARCHIVES CANADA to reproduce, lend, distribute, or sell copies of this dissertation anywhere in the world in microform, paper, or electronic formats and to authorize or procure the reproduction, loan, distribution or sale of copies of this dissertation anywhere in the world in microform, paper, or electronic formats.

The author reserves other publication rights, and neither the dissertation nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

BANDWIDTH IS POLITICAL

By Nancy Paterson

By virtue of submitting this document electronically, the author certifies that this is a true electronic equivalent of the copy of the dissertation approved by York University for the award of the degree. No alteration of the content has occurred and if there are any minor variations in formatting, they are the result of the conversion to Adobe Acrobat format (or similar software application).

Examination Committee members:

1. Dr. Michael Murphy
2. Dr. Jerome Durlak
3. Dr. Fred Fetcher

Bandwidth is Political

Abstract - Summary thesis statement

The global public Internet faces a growing but little studied threat from the use of intrusive traffic management practices by both wholesale and retail Internet service providers. Unlike research concerned with bandwidth and traffic growth, this study shifts the risk analysis away from capacity issues to focus on performance standards for interconnection and data reachability.

The long-term health of the Internet is framed in terms of “data reachability” – the principle that any end-user can reach any part of the Internet without encountering arbitrary actions on the part of a network operator that might block or degrade transmission. Risks to reachability are framed in terms of both systematic traffic management practices and “de-peering,” a more aggressive tactic practised by Tier-1 network operators to resolve disputes or punish rivals.

De-peering is examined as an extension of retail network management practices that include the growing use of deep packet inspection (DPI) technology for traffic-shaping. De-peering can also be viewed as a close relative of Net Neutrality, to the extent that both concepts reflect arbitrary practices that interfere with the reliable flow of data packets across the Internet. In jurisdictional terms,

however, de-peering poses a qualitatively different set of risks to stakeholders and end-users, as well as qualitatively different challenges to policymakers.

It is argued here that risks to data unreachability represent the next stage in debates about the health and sustainability of the global Internet. The study includes a detailed examination of the development of the Internet's enabling technologies; the evolution of telecommunications regulation in Canada and the United States, and its impact on Internet governance; and an analysis of the role played by commercialization and privatization in the growth of risks to data reachability.

DEDICATION

“Not all those who wander are lost. J. R. R. Tolkien” is the author’s notation added to the chapter one title *Networking and Network Routing: An Introduction* in D. Medhi & K. Ramasamy’s 2007 book titled *Network Routing*. I could not have said it better.

This is dedicated to my committee Dr. Jerome Durlak, Dr. Fred Fletcher and Dr. Michael Murphy. Assistance from Tracey Bickford, Julie Birkle, Karen Brophy, Dr. David Ellis, Dr. Caitlin Fisher, Diane Jenner, Ian Lumb, Sal Panudero, Eriks Rugelis, Elvira Sanchez de Malicki and Lynn Walker is gratefully appreciated. Without their support this work would not have been possible.

Table of Contents

Abstract - Summary thesis statement	iv
List of Figures	x
Introduction: A New Internet Era.....	1
The Commercialization of Bandwidth Stakeholders.....	3
Reachability and Net Neutrality.....	7
Dominant Role of the United States.....	9
Literature Review	11
Chapter I. Peering, Transit and Bandwidth Stakeholders.....	18
Peering and Transit in Practice	18
Peering	21
Transit.....	22
Interconnection: Border Gateway Protocol & Autonomous System Numbers.....	22
Emergence and Development of Bandwidth Stakeholders.....	30
Regional Peering Practices.....	34
Chapter II. Conflict and Confluence: Legacy Telephone Networks and the Emergence of the Internet.....	40
U.S. Role and Influence	42
Gatekeeping and Net Neutrality	44

Clash of Two Cultures.....	46
Rival Solutions to the Problem of Internetworking.....	55
Deregulation and Privatization: the Internet Becomes a Business.....	66
The FCC and Internet Backbone Policy	75
Telecommunications Services vs Information Services	81
Ad Hoc Oversight of the U.S. Backbone Marketplace.....	91
Chapter III. Internet Governance Before and After Commercialization.....	98
Development of the Technical Governing Bodies	101
ARPANET and After: Managing the Internet.....	104
The National Science Foundation and NSFNET.....	109
A New Architecture: Network Access Points.....	114
Commercialization: The NSF Withdraws and Transit Begins.....	118
The Impact of Commercialization on IANA: Old Functions, New Problems..	122
ICANN: the Domain Name System and other Conflicts	128
The IANA Function and Autonomous System Numbers (ASNs).....	132
Governance in a New Light.....	135
Chapter IV. Risks to Reachability.....	146
Peering, Interconnection and Market Concentration	146
Reachability as a Net Neutrality Issue	153

Traffic Management Practices Affecting Reachability	162
Packet Filtering	163
Traffic-shaping	167
AS Path Filtering & De-peering	173
MPLS: Expanding the Scope of Traffic Management	185
Chapter V. Conclusions and Future Directions	189
Lessons of History	190
Disclosure as an Instrument of Reform	195
Appendix A – IXmaps	205
Appendix B – Tier-1 Networks	207
Appendix C – Sample Weekly Routing Table Report	214
Appendix D – Autonomous Systems	227
References	237

List of Figures

Figure 1. Internet timeline. p. 55.

Figure 2. OSI Layered Model. p. 61.

Figure 3. MPLS logical topology. p. 187.

Introduction: A New Internet Era

The reliability of data transmission across the global public Internet is being compromised by the expanding use of traffic filtering, traffic-shaping and de-peering (techniques for managing network traffic), used by network operators at both wholesale and retail levels. These and other traffic engineering practices are contributing to a phenomenon known as data unreachability, namely a loss of connectivity that unexpectedly disrupts the ability of an Internet service provider (ISP) or end-user to communicate in sanctioned ways over the public Internet. Data unreachability has important and growing political, commercial and technical ramifications that to date have received little systematic attention from the academic research community.

Traffic management practices are a mixed blessing, as evidenced by the voluminous filings and testimony that became part of the record in the proceeding on ISP traffic management practices held by the Canadian Radio-television and Telecommunications Commission (CRTC) over the course of 2009 (CRTC, 2008b). In many situations, ISPs, especially those offering retail broadband connectivity to residential customers, use traffic management controls to enhance network security by eliminating or reducing unwanted traffic, such as spam and malware.

On the other hand, traffic management controls have become contentious, because many incumbent ISPs have made routine use of deep packet inspection (DPI) and similar tools to slow down or block altogether traffic they allege is adding to congestion of their networks, often by targeting particular protocols, especially the BitTorrent peer-to-peer protocol. De-peering is another action that may be undertaken by network operators that puts reachability at risk. De-peering is the practice by commercial backbone providers of unilaterally severing interconnections with other commercial backbone providers as a means of enforcing contracts, punishing rivals and expanding market share. De-peering does not fall within the generally accepted meaning of traffic management, but produces the same end result.

Although not yet commonplace, data reachability is an important object of study for reasons related to the promotion of both the public interest and scholarly inquiry. Except under network topologies that provide for redundant interconnections, the abrupt and arbitrary decision by one bandwidth provider to refuse data traffic being exchanged with another bandwidth provider can have damaging and disruptive effects on many other bandwidth stakeholders, including thousands or even millions of end-users.

Several persistent structural factors make the risks involved in de-peering, as well as disruptive traffic management practices, of more than passing interest to

industry stakeholders and the research community. First, de-peering and traffic management are planned and deliberate actions, not the result of accidental or unforeseen circumstances. Second, the commercial activities of backbone providers, while dependent on standards overseen by technical bodies such as the Internet Engineering Task Force (IETF) and Internet Architecture Board (IAB), are entirely unregulated, and fall outside the jurisdiction of international organizations otherwise involved in Internet governance, including the Internet Corporation for Assigned Names and Numbers (ICANN) and the International Telecommunications Union (ITU).¹ Third, the formal business arrangements made among backbone providers through contracts or service level agreements (SLAs) are rarely open to outside scrutiny and usually subject to rigorous non-disclosure agreements (see Appendix B – Tier-1 networks – Definition).

The Commercialization of Bandwidth Stakeholders

The issues examined in this chapter flow from a series of historical developments that saw the Internet grow from a publicly funded, non-commercial service to a global platform whose bandwidth facilities are now controlled by large private corporations. From 1988 to 1995, NSFNET, the network operated on behalf of the National Science Foundation, acted as the principal Internet backbone, eventually replacing the original Advanced Research Projects Agency Network or

¹ Many of the commercial entities that own and/or operate large backbone facilities are telecommunications carriers, some of whose activities are subject to scrutiny by regulatory bodies such as the FCC.

ARPANET. By the time NSFNET was shut down in April 1995, Internet connectivity was being provided by a growing number of commercial operators, using new software protocols and interconnection methods. The commercialization of the Internet was also accelerated by related events such as the Netscape initial public offering (IPO), which took place in August 1995, just months after the closure of NSFNET. These historical events are examined in detail below, in Chapter III.

For purposes of this investigation, the most significant shift precipitated by commercialization of the public Internet was from universal reliance on settlement-free peering (i.e. exchange of traffic without payment) to a mixed system in which paid transit came to play a crucial role. In general terms, peering is the “voluntary interconnection of two independent networks that exchange traffic on a settlement-free basis” (Cukier, 1997, p. 1), whereas transit is the provision of Internet access by one provider to another for payment. Although the distinction may appear to be straightforward (what was once free now comes at a cost), we will demonstrate in this chapter that the introduction of paid interconnection alongside traditional peering marked a radical shift in the behavior of the bandwidth marketplace. This shift was in part a result of the sheer growth and mainstreaming of the Internet over the course of the 1990s. It reflected a new set of priorities under which large firms, especially incumbent telecommunications carriers, were more likely to be concerned with protecting

their financial interests than with maintaining reliable interconnections across the Internet. As we explain at length in later chapters, the issue of reachability raises concerns of accountability and transparency because business arrangements made between networks are almost invariably subject to non-disclosure agreements and thus inaccessible to third parties.

One of the most compelling signs of lack of transparency in this market lies in the difficulties involved in merely identifying the major bandwidth stakeholders, especially the Tier-1 networks (as distinct from Tier-2 and Tier-3 networks). A Tier-1 network is commonly taken to be a network that fulfills two conditions: i) it can reach every other network on the Internet; and ii) it can do so without purchasing IP transit or paying settlements (Norton, 2001, p. 2). Operating on a settlement-free basis - a type of bill-and-keep arrangement familiar from inter-carrier compensation models - is thus a necessary but not sufficient condition for claiming status as a Tier-1 network. Tier-1s are often loosely termed "ISPs" - Internet service providers - as a result of the fact that the Tier-3 resellers of residential Internet access are primarily owned and operated by the incumbent carriers that own the Tier-1 networks.

A Tier-2 network typically operates on a mix of settlement-free peering with certain networks, while paying one or more other networks for transit facilities to reach at least some parts of the Internet. A Tier-3 network is one that operates

strictly on a transit basis, paying one or more upstream providers for Internet access. As for the Tier-1 networks, their number and identity change from time to time, and there are typically discrepancies from one industry list of Tier-1 networks to another. Tier-1 network operators are typically based in the United States, among them AT&T, Level 3 Communications, Qwest, Sprint, Verizon, XO, Abovenet and Cogent. The few non-US-based Tier-1 networks include TeliaSonera International, based in Stockholm; NTT in Japan; Global Crossing, based in Bermuda; Teleglobe, owned by India's Tata Group; and China Telecom (Brown, Hepner & Popescu, 2009, p. 5; Zmijewski, 2008b).

The stakeholder hierarchy just described poses serious problems for both participants and researchers - and describing these problems is one of the principal purposes of this chapter. These problems can be described as follows. First, it is difficult to establish definitively which firms are operating true Tier-1 networks. Second, networks that have not achieved Tier-1 status may be motivated for marketing purposes to claim they are Tier-1 operators - and such claims can be made with impunity, particularly since Tier-1 status is not determined simply by reach, bandwidth capacity or number of customers. Third, since every Tier-1 network must by definition peer with all other Tier-1 networks, individual networks have no means of establishing redundancy or backup arrangements for their peering interconnections with other Tier-1 networks. This top-level peering structure means that a de-peering incident may abruptly

partition off one section of the Internet from other sections, and when such an incident takes place, there are currently no international provisions in place for either regulatory oversight or third-party dispute resolution.

Reachability and Net Neutrality

One of the goals of the research presented here is to shift the locus of analysis away from capacity-related issues (i.e. the bandwidth available to ISPs and end-users in given markets), in favor of an emphasis on performance standards for interconnection and data reachability. In the last two to three years, much of the public debate and research literature on the health of the Internet has been confined to capacity issues. Regulators and lawmakers in a number of countries, including Canada and the United States, have wrestled with the question of what rules, if any, should be put in place to manage both the sanctioned network management activities of ISPs and the unsanctioned activities of file-sharing sites that may facilitate copyright infringement. The common denominator in all these discussions is the fear, real or misplaced, that both Internet backbones and last-mile connections are running out of capacity, raising the possibility of “brownouts.”

Current discussions on network capacity and ISP gatekeeping have crystallized around the concept of Net Neutrality - the case made by public-interest advocates and others that broadband ISPs should be required to offer non-

discriminatory access to their last-mile facilities, particularly for the benefit of paying subscribers in jurisdictions where incumbent providers can practise monopoly or near-monopoly leveraging of their facilities. ISPs operating in North America have responded by claiming that intrusive network management practices are necessary because the volume of Internet traffic has been rising too sharply for bandwidth providers to keep up (Hunter, 2006a, p. 11).

It is argued here that data unreachability especially created through traffic engineering practices, including de-peering, represents the next stage in debates about the health and sustainability of the global Internet. De-peering can be viewed as an extension of retail network management practices that include the growing use of deep packet inspection (DPI) technology for traffic-shaping. De-peering can also be viewed as a close relative of Net Neutrality, to the extent that both concepts reflect arbitrary ISP practices that may interfere with the reliable flow of data packets across the Internet. In jurisdictional terms, however, de-peering poses a qualitatively different set of risks to stakeholders and end-users, as well as qualitatively different challenges to policymakers. Whereas Net Neutrality is treated as a national issue in countries such as Canada and the United States, data reachability is international in its scope.

Dominant Role of the United States

One of the few ways in which information about market share and the business dealings of Tier-1 providers has become public is through the ad hoc review of mergers and acquisitions by legal authorities, often as part of the enforcement of anti-trust legislation. Over the last decade, the US Department of Justice (DoJ) has issued orders preventing or amending certain mergers and acquisitions, including the proposed merger of MCI-WorldCom and Sprint, which was disallowed in 2000. Prior to that, the DoJ had ordered WorldCom to divest InternetMCI, the backbone provider owned by MCI, as a condition the acquisition of MCI by WorldCom (MCI, WorldCom merger gets green light from DoJ, 1998, p. 1).

Several aspects of this legal environment as it affects large bandwidth stakeholders are explored in this chapter. The first of these concerns the fact that the enforcement actions of the DoJ have been carried out with very little regulatory guidance from the Federal Communications Commission (FCC). Second, because of the global reach of US companies, the European Union (EU) has also taken legal action in anti-trust cases such as those just cited. Third, and notwithstanding the involvement of competition regulators in jurisdictions such as the EU, the United States continues to have a unique and compelling interest in the behavior of Tier-1 and other bandwidth stakeholders.

This interest flows from the level of Internet-related activity in the US and the unusual political role of the US government in Internet governance.

Notwithstanding claims made by US government officials, the administration of the domain name system (DNS) was not “privatized” with the establishment of ICANN in 1998. Despite the global nature of its technical mandate, the international deployment of the DNS root servers and debates over the years about its multilateral responsibilities, ICANN has always operated under the aegis of the US Department of Commerce - a relationship that has been the source of much conflict at international forums such as the World Summit on the Information Society (WSIS). On the other hand, we argue in this paper that the strong focus in public debates on the Internet's *naming* functions and the role of ICANN should not be allowed to divert attention from the growing problems associated with data *routing* and reachability across the global Internet.

Literature Review

Relatively few scholars have addressed the issue of unreachability of data in the Internet. We begin with a glance at the best known studies.

Peering and the consequences of severing Internet interconnections were examined in Neil Cukier's 1997 paper, *Peering and fearing: ISP interconnection and regulatory issues* presented at the Harvard Information Infrastructure Project Conference on the Impact of the Internet on Communication Policy. This paper was an early warning on the topic of network peering policies and Internet traffic control. Cukier raised the issue immediately after passage of the 1996 US Telecommunications Act, suggesting that large networks seemed to have changed peer relationships into supplier-customer relationships.

The 2002 FCC Network Reliability and Interoperability Council V *Report to the Nation (Network Reliability and Interoperability Council V: The future of our nation's communications infrastructure)* is an important document regarding peering policy primarily because Focus Group 4 – Broadband, recommended that peering policies be made public information (Network Reliability and Interoperability Council, 2002a). Additionally, J. Scott Marcus is a former FCC senior adviser on Internet technology who authored a 2006 declaration in *Hepting et al. v. AT&T*, on behalf of the Electronic Frontier Foundation's class-action lawsuit against AT&T for collaboration with the warrantless domestic

spying program. Marcus included the same FCC NRIC report as exhibit R in his 2006 declaration, as the report provides a clear definition of peering and infrastructure (Marcus, 2006).

William Norton is an author and technical specialist on Internet infrastructure and peering. His 2003 paper, "The Evolution of the U.S. Internet Peering Ecosystem," provides excellent background for this discussion. Norton went on to work for Equinix, an Internet interconnection company funded by Cisco, Microsoft and private equity funds. The problems caused by de-peering were addressed by Brown, Hepner and Popescu of Renesys Corporation in a presentation entitled *Internet Captivity and the De-peering Menace* at the 2009 meeting of NANOG, the North American Network Operators Group. The NANOG email listserv (and archives) provide a valuable record of topical issues of concern to the system administrators who help run the Internet. Although these discussions are limited by mandate to technical issues, the participants often address the larger dilemmas associated with their work as Internet technologists.

Ars Technica is a Conde Nast Web site which describes itself as "a premier destination for technology news, analysis, and in-depth information. With 6 million monthly readers, *Ars Technica* is in the Top 20 of all IT News and Media sites online according to Hitwise, and ranks in Technorati's Top 10" (*Ars*

Technica, 2009). Rudolph van der Berg's information on connectivity and van Beijnum's information on de-peering are particularly good.

Technological, economic and regulatory factors that shape this field are surveyed in great detail by Nuechterlein and Weiser in *Digital Crossroads: American Telecommunications Policy in the Internet Age*.

Two research projects, one Canadian and one American, address the issue of filtering data crossing the public Internet, but neither of them examines reachability failures caused by de-peering. For example Anderson, Katz-Bassett, Krishnamurthy and Madhyastha from the Hubble project at the Department of Computer Science and Engineering, University of Washington are actively studying unreachability in the public Internet. In their paper "Studying Black Holes in the Internet with Hubble" presented at the *2008 USENIX Symposium on Networked Systems Design & Implementation (NSDI)* in San Francisco, they discussed their software tool entitled *Hubble* which finds Internet reachability problems resulting in "blackholes", where routes exist to a destination but packets are unable to reach the destination.

The OpenNet Initiative at the Munk Centre, University of Toronto, is focussed on research into Internet practices and policies of Internet filtering and surveillance. In a 2008 publication entitled *Access denied: The Practice and Policy of Global*

Internet Filtering, the editors include an essay by Murdoch and Anderson on Border Gateway Protocol (BGP) filtering practices and technology (Deibert, Palfrey, Rohozinski, & Zittrain, Ed. 2008, p. 57). The OpenNet Initiative research concentrates on state political controls designed to block or restrict Internet traffic. The present study concentrates more strictly on failures of reachability of data by examining: traffic filtering, routing interconnection policies of networks, de-peering (or the severing of connectivity between networks) and the resulting effects upon reliability and persistence of data for end-users.

Balakrishnan and Feamster's *Interdomain Internet Routing* describes autonomous systems, peering, transit and importing and exporting routes. The complexity of software and hardware used in interdomain routing is explained. They also touch on the difficulties faced by small entities to multi-home their connectivity. Yochai Benkler's 2006 book, *The Wealth of Networks*, explores the idea that policies employed in the technical administration of the Internet are related to other areas of public policy.

Reflections on Internet Transparency is the title of a 2007 Internet Architecture Board report edited by Aboba and Davies. This report discusses filtering of Internet traffic in the context of upgrading the Internet to make more IP addresses available. The report reviewed previous IAB statements on Internet transparency, discussed new transparency issues, noting that IETF documents

have touched on the issues of transparency as well. “Far from having lessened in relevance, technical implications of intentionally or inadvertently impeding network transparency play a critical role in the Internet’s ability to support innovation and global communication” (Aboba & Davies, 2007, p. 1). The report is an assessment of Internet traffic engineering policies. The report emphasizes the importance of the end-to-end principle, which stipulates that in the design of the Internet, intelligence is concentrated at the edges of the network, while the network itself is dedicated solely to delivering packets:

A network that does not filter or transform the data that it carries may be said to be “transparent” or “oblivious” to the content of packets. Networks that provide oblivious transport enable the deployment of new services without requiring changes to the core. It is this flexibility that is perhaps both the Internet’s most essential characteristic as well as one of the most important contributors to its success. (Aboba & Davies, 2007, p. 1)

In his Rochester Institute of Technology masters dissertation, *The Evolution of Internet Interconnections*, Sean Butler claims that industry self-regulation will fail. Robert Cannon echoes Sean Butler’s claim in “The Legacy of the Federal Communications Commission’s Computer Inquiries” (*Federal Communications Law Journal*, 2003). Cannon provides an overview of the *Computer I*, *II* and *III Inquiries*. The *Computer Inquiry II* established the important “basic” versus

“enhanced” dichotomy and created the regulatory provision called “comparatively efficient interconnection” or CEI. CEI established a set of rules whereby a telco offering basic voice service would be permitted to enter into enhanced (i.e. data) services markets without first creating a structurally separate entity, conditional upon their making available to competitors the same terms of interconnection as they enjoyed. This was based on the idea that the telcos would voluntarily divulge their network plans to competitors. Cannon notes that the commercial Internet exchange (CIX), which was set up in 1991 as the first exchange point for traffic between commercial Internet backbones, was skeptical about the feasibility of voluntary compliance.

[CIX...] objected to the movement to non-structural safeguards in *Computer Inquiry III* arguing that this created a problem with enforcement. Recognizing validity to the CIX objection, the Commission stated: ‘We believe that competitive ISPs will themselves monitor CEI compliance vigilantly, and will call the Commission’s attention to any failure by a BOC to follow through on its CEI responsibilities ... The Commission will not hesitate to use its enforcement authority, including the Accelerated Docket or revised complaint procedures, to review and adjudicate allegations that a BOC is falling short of fulfilling any of its CEI obligations’. Note, however, that this does create certain structural oddities. First, an unregulated industry, with little knowledge of the FCC, is asked to watch a regulated

industry. Second, small companies are asked to watch the largest corporations in the United States. Third, ISPs are placed in a position of filing complaints against their sole supplier of a crucial facility. Fourth, contrary to normal jurisprudence, the party that lacks the information has the burden of moving (normally, all things being equal, the party with the information has the burden of moving—in this case, the burden is on the ISPs, because the information is held by the BOCs. (Cannon, 2003, p. 203)

Adam Thierer, previous director of telecommunications studies for the CATO Institute in Washington DC in June 2007, provided the present author with an unpublished CATO research paper by Jennifer DePalma, “Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry.” The paper provided a useful survey of commercial de-peering concluding that the level of unreachability produced was not sufficient for alarm. Other important resources for this research include *Who Controls the Internet?* (Goldsmith and Wu, 2006), which addresses tiered levels of service on the Internet by distinguishing between content and capacity-based traffic controls. Legal scholar Rob Frieden raises the problem of Internet balkanization, but does not directly address peering or Internet infrastructure (Frieden, 1998).

Chapter I. Peering, Transit and Bandwidth Stakeholders

This chapter is fundamentally concerned with the technical and commercial arrangements that ensure data traffic is transmitted between the administratively separate networks that make up the public Internet. Reduced to the simplest terms, these networks exchange traffic with other networks in one or both of two ways: unpaid, what is termed a settlement-free arrangement, or paid, at negotiated commercial rates. The first of these is known as peering, the second as transit.

Peering and Transit in Practice

The single most important consideration in determining whether data transmission will be settlement-free or paid concerns the amount of traffic handled by the networks in question. In this respect networks operate like many other businesses that find it advantageous to exchange goods or services on a “contra” basis, without money changing hands - an arrangement that presupposes one party has something to offer roughly equal in value to what the other party has to offer. In the telecommunications industry, where compensation paid by one carrier to another for terminating voice calls may take a number of forms, such deals are referred to as “bill-and-keep.”

In the business of long-haul data transmission over the Internet, networks enter into peering relationships with other networks for the purpose of exchanging symmetric or near-symmetric volumes of data traffic between the customers of each network.² On the other hand, most networks, except the largest bandwidth providers, must pay for access to at least some of the routes on the Internet, under the commercial arrangement known as transit. In a typical transit relationship, one network operator (usually a smaller one) pays a fee to another network operator (usually a larger one) for the asymmetric exchange of data. The smaller entity pays to have its traffic handled by the larger entity because it has more to gain from interconnection of the two networks in question.

The main operational bandwidth stakeholders can be defined according to whether they use peering or transit, or a mix of the two, to service their customers. The key stakeholders are the Tier-1, Tier-2 and Tier-3 networks. A Tier-1 network operates strictly by peering with other Tier-1 networks, meaning it can reach every other network on the Internet without paying for transit. A Tier-2

² The terms “symmetric” and “asymmetric” are used commonly in networking parlance, especially to describe a two-way connection. In a symmetric network or connection, the bandwidth available in one direction is equal or nearly equal to the bandwidth available in the other direction; in an asymmetric network or connection, the bandwidth available in one direction is significantly greater than the bandwidth available in the other direction. Residential broadband connections (i.e. high-speed Internet connections) have historically been designed by both telephone and cable providers to be highly asymmetric, i.e. provisioned to offer much more bandwidth in the downlink than in the uplink. The ramifications of this design principle are discussed elsewhere in this paper. By extension, network engineers and others in the business of providing bandwidth use “symmetric” and “asymmetric” more generally, as in the case of “symmetric or near-symmetric volumes of data traffic between the customers of each network.” This phrase refers to peering relationships in which the amount of data transmitted by carrier X and accepted by carrier Y is roughly the same as the amount of data transmitted by carrier Y and accepted by carrier X.

network typically operates on a mix of settlement-free peering and paid transit while Tier-3 networks, which usually operate generically as xSPs (Internet service providers, Internet access providers, Internet storage providers, Internet application providers, Internet hosting providers and other services for end-users) rarely if ever have peering relationships unless there is a need to overcome the vagaries of geographical boomerang in connecting with a higher level network. This means they purchase access or bandwidth for their customers strictly on a transit basis.

In standard arrangements, the high-level networks known as Internet backbones or Tier-1 networks are transit-free (see Appendix B). These high-level networks do not have direct contact with end-users except insofar as they own local Tier-3 ISPs. Local, retail-level Internet service providers – Tier-3 ISPs, are the resellers that provide their end-users transit-based access to the global Internet.

Historically, Internet service providers have always paid for upstream transit and began operating long before the deployment of commercial network backbones.

In the early days of the Internet, retail Internet service providers operated banks of modems that provided local dial-up access. Today, most local Internet service providers purchase backbone access from a single upstream provider. Local ISPs tend to serve only residential users, while national and regional providers (Tier-2) serve businesses as well as residential users.

Peering

As noted above, peering is typically settlement-free, meaning that neither network operator pays the other for exchanging traffic. It must be noted, however, that such cost-free arrangements apply only to the data exchanged between customers of the two networks in question. Furthermore, peering arrangements between two networks are not always settlement-free. There is, in other words, a paid variant of unpaid peering.

In one type of situation, paid peering may come about as the result of changes made to a previously settlement-free peering relationship, in particular when the larger of two networks determines on the basis of a cost-benefit analysis that it deserves to be paid for handling the other network's traffic (van der Berg, 2009). In other situations, "near high-level networks" may choose to buy their way into the top echelon of Internet backbone networks by paying to peer with a larger network, since paid peering tends to be less expensive than transit. There are numerous "near high-level networks" that offer service combining some transit and some full or partial paid peering (van der Berg, 2009). For some authorities, the use of the term "paid peering" is a misnomer often applied in marketing strategies and distinctly different from "true" - i.e. settlement-free - peering.

Transit

Technically speaking, a transit agreement means the purchaser is paying another provider for connectivity to Internet locations outside of its own network. Transit may be based on a flat monthly charge or on a volume basis (Mbit/s per month). Full transit includes access to all the routes on the supplier network, as well as to those of the supplier's peers and the supplier's own upstream transit suppliers (van der Berg, 2009). As a variant of full transit, partial transit arrangements include access to a supplier's peers but exclude its transit suppliers. Partial transit has been traditionally more common in Europe than in the US, where peering and full transit have been more dominant (Williamson, 2003, pp.16-17). Residential end-user customers are by definition transit purchasers since they buy Internet access or IP transit from an ISP that in turn purchases transit from a larger upstream provider. That larger upstream provider may purchase full transit from its own provider or partial transit and partial peering combined - or alternatively it may be in a position to acquire access upstream from one or more high-level networks as a full-fledged peer.

Interconnection: Border Gateway Protocol & Autonomous System Numbers

Over and above the business efficiencies achieved by peering, network operators have an important technical goal in wishing to offer peering as a solution to routing, the process of selecting paths in a network along which to

send traffic. Routing is an especially complex engineering challenge in a packet-switched network such as the public Internet, since the underlying rationale for the use of packet switching technology lies in its ability to offer many different paths for a given transmission, with a view to maximizing overall network efficiency and promoting best-effort delivery of packet transmission. The crucial task of routing packets between networks is accomplished by a core Internet routing technology – Border Gateway Protocol or BGP – which has been in widespread use by network operators since the mid-1990s. End-users do not interact directly with BGP, despite its important role in determining the reachability of data packets in peering and transit interconnections (Medhi & Ramasamy, 2007, p. 239).

Each entry in the BGP routing table lists a distinct path from which the router (a programmable device to route Internet traffic) will choose one path to use (Medhi & Ramasamy, 2007, pp. 31-2, 59). Routing advertisements “pull” or attract data traffic rather than traffic being “pushed” around the Internet.

BGP works by maintaining global routing tables of IP networks, which designate network reachability among networks called “autonomous systems” or ASes. An AS is an “Internet domain” or collection of IP sub-networks and routers under the control of one entity (such as a university, a commercial enterprise, government, or other enterprise) that presents a single border gateway routing policy to the

global Internet. This common policy is the “administrative boundary” for an aggregated collection of blocks of IP addresses (Barrolli, Durrese, Iyengar, Kannan & Paruchuri, 2004, p. 6). Blocks of Internet addresses associated with an AS number are aggregated into “prefixes” which contain a range of one or more IP addresses (Huston, 2003, p. 3). A routing table contains all the routable IP prefixes, each of which is followed by a series of AS numbers indicating the AS path to that prefix:

A Global Routing Table lists every single prefix on the Internet, the different available paths to that prefix, and other information that lets the AS make forwarding choices based on the available paths. (Faratin, Clark, Gilmore, Bauer, Berger & Lehr, 2007, p. 6)

When an originating AS announces prefixes that have been allocated to it by a regional Internet registry (RIR), the announcement ordinarily propagates to all other ASes in the Internet. The Internet Corporation for Assigned Names and Numbers (ICANN) allocates AS numbers to the RIRs through the Internet Assigned Numbers Authority (IANA), which then sub-assigns the AS numbers (as well as sub-assigning blocks or prefixes of IP numbers) to ASes within their region. The five global five regional Internet registries are: Asia - Asia Pacific Network Information Center (APNIC); North America – American Registry for Internet Numbers (ARIN); EU and Middle East - Réseaux IP Européens Network

Coordination (RIPE NCC); Africa - African Network Information Coordination (AfrinIC); Latin American and Caribbean - Latin American and Caribbean Network Interconnection Coordination (LACNIC).

As an illustration of the scope of the technical data associated with an autonomous system (AS) number, French network operator Bouygues Telecom states on its Web site (Bouygues Telecom, n.d.) that its “AS Number is AS5410 registered at the RIPE NCC authority, is a French network operator and a RIPE Local Internet Registry (LIR).” This statement means that a large-scale IP network in Europe, called Bouygues Telecom, has the autonomous system number AS5410; is registered at the EU and Middle East regional Internet registry RIPE NCC; operates a network; and can assign IP numbers to customers that have been allocated to it by RIPE NCC.

Autonomous system numbers are divided into two ranges. Public AS numbers ranging from 1 to 64511 are used on the Internet. Private numbers in the 64512 to 65534 range may only be used internally within an organization. On a daily basis the Asia-Pacific Network Information Center (APNIC) automatically sends a fairly accurate report to all the other regional Internet registries regarding Internet connectivity from its point of view. This report is distributed weekly on the NANOG mailing list (see example, Appendix C, the Weekly Routing Table Report). The Analysis Summary of this particular report states, “Total

autonomous systems present in the Internet Routing Table: 31449". This means that of 65536 possible 16-bit AS numbers (actually 64511 possible public IP numbers) 31449 AS numbers are actively in use, as far as the Asia - Asia Pacific Network Information Center (APNIC) can see (Routing Analysis Role Account, 2008).

BGP manages the routing among the AS routers by selecting the best path for the traffic flow based on many factors. Peering information is typically described using the Routing Policy Specification Language (RPSL) in preparing commands used to configure the routers. Primarily the management of the routing system by the Border Gateway Protocol needs only an autonomous system set (AS-SET) defining the ASes which an AS peers with, or which purchases partial peering or transit. The AS-SET is called a "peer set"; it lists the peers that the network exchanges route announcements with, and is somewhat closely guarded information in the commercial context of contracts and agreements. Even though this peering information is closely guarded, it is available on the Internet in technical form (Medhi & Ramasamy, 2007, p. 313). Some public posting of this AS-SET peering information on the Internet is voluntary and the format of the information is difficult for a lay person to understand.

Traffic performance is optimized for customers by reducing the routing hops taken by a data packet or set of packets. A hop is an intermediate connection in

a string of connections linking two network devices, typically routers, whose task is to send packets along paths it has selected by means of routing algorithms (McPherson, Sangli & White, 2005, p. 15). As Medhi & Ramasamy put it: “[A] router in the Internet is also a host [*read server*] and is assigned an IP address” (Medhi & Ramasamy, 2007, p. 18) emphasis added. Each time a packet is forwarded from one router to the next, a hop occurs, and on the Internet, most transmissions go through several routers - and therefore several hops - before they reach their final destination (Medhi & Ramasamy, 2007, p. 52). The more hops, the longer it takes for data to go from source to destination. When networks interconnect, they agree to exchange traffic and routing information called “routing advertisements” - information that allows each network to compute traffic routes (Balakrishnan & Feamster, 2005, p. 5, 6). When two networks interconnect as peers this produces fewer hops than arrangements with other networks would otherwise have permitted.

Many other factors affect level of service in the hand-off of data traffic from one network provider to another. One of the most important of these concerns the various facilities around the globe at which networks are physically brought into contact with other networks. Bandwidth providers usually have dedicated connections at one of two kinds of public peering points: either metropolitan-area exchanges (MAEs) or network access points (NAPs) (Halabi, 1997, p. 8). Since bandwidth stakeholders of any size can exchange traffic through these publicly

accessible facilities, all or most of the Internet access service providers in a given area may end up putting all their traffic through the same interconnection point, with correspondingly negative effects on level of service (Halabi,1997, p. 13).

The original government-funded network access points are now usually referred to as Internet Exchanges (IX's). A few exchange points, particularly in the United States, are operated by facilities carrier hotel operators such as Telx, and there are also carrier-neutral third parties, such as Equinix (Blake, 1999, p.18; Telx opens the most network rich, interconnected colocation center in NYC metro area, 2009). Operators of these carrier-neutral facilities often go to great lengths to promote their services and encourage new peering.

As a result of such efforts, groups of otherwise distinct and unrelated networks may all agree to interconnect at an Internet Exchange under identical multi-lateral peering agreements (MLPAs), which are paying commercial arrangements (Medhi & Ramasamy, 2007, p. 293). These arrangements may be seen in sharp contrast to the pre-1996 practice of open, settlement-free peering at public interconnection points. Even though such multi-party agreements have been in place for many years, there is no general consensus on how to calculate an economic value for the working partnerships created by these agreements. This represents another example of how commercialization, in addition to creating potential technical efficiencies, has also created an uncertain business

environment for networking and the potential for conflicts that have no easy or obvious framework for third-party resolution.

While most networks have a single AS number, a limited number of large entities such as Google have multiple AS numbers or Internet domains in order to have administratively different policies regarding routing reachability (see Appendix D, RFC:1930 - Guidelines for creation, selection, and registration of an Autonomous System.) An AS is assigned a number that uniquely identifies the particular network or system on the global Internet. Until 2007, AS numbers were 16-bit integers, which allowed for a maximum of 65536 autonomous systems globally. As AS numbers are seen as a diminishing resource, in early 2007, in order to provide for more autonomous systems in the future, the RIRs started to issue 32-bit AS numbers. Commencing 1 January 2010, all RIRs will cease to make any distinction between 16-bit AS Numbers and 32-bit AS Numbers (ARIN Number Resource Policy Manual, 2009).

Whether or not there is a real need to move to 32-bit AS numbers is not completely clear. In a March 2009 discussion on the NANOG mailing list it became apparent that the level of assigned AS numbers, reaching 53,000 numbers at the time, was largely due to many of them being stagnant or not in use. This issue has arisen because of takeovers and other developments (Brilus, 2009). Discussions also suggested that a mere \$100-per-year charge for an AS

number levied by ARIN dissuaded the return of unused allocations (Schiller, 2009) and that there is no mechanism in place to decide if an AS number is not used or needed (Huston, 2009). Thus in March 2009 there was a significant number of ASNs allocated but not visible in the public Internet. It is not known to what extent these numbers may be still in use in various forms of private or semi-private networks (Huston, 2009).

Emergence and Development of Bandwidth Stakeholders

The predecessor to the Internet, the ARPANET, was originally designed in the late 1960s on the basis of peer-to-peer architectures, a concept which over the last decade has drawn a great deal of attention, although more for legal and social reasons than for technical ones. In December 1969, in one of its earliest implementations, four computers, located at UCLA, UC Santa Barbara, the Stanford Research Institute (SRI) and the University of Utah were connected and the first packet-based messages were shared among them (Halabi, 1997, p. 3). The primary motivation behind the funding of the ARPANET by the US Department of Defense was two-fold: to allow academic and military researchers to share computing resources with their peers across the country; and to do so through the use of innovative hardware and software that would enable communication among machines using entirely different operating systems and architectures.

The twin models for the exchange of Internet traffic under discussion here - peering and transit - began to emerge over the decade from 1986 to 1995. This was the period during which the NSFNET, the backbone network created by the National Science Foundation, gradually opened up the cloistered and strictly non-commercial world of networking and military research to hundreds of other networks, and eventually, millions of mainstream users.

Thus, for a quarter of a century, the platform that became the public Internet ran on a primary backbone: first ARPANET and subsequently NSFNET. The prodigious growth of the Internet over this period precipitated changes of many kinds. The visible signs of growth included sharp increases in the number of networks, hosts, backbones and users that were becoming part of the new Internet. For example, from 1986, when NSFNET went online, to April 1995, when it shut down, the number of Internet hosts (computers with a registered IP address) grew from about 4,000 to over four million (Zakon, n.d.). Other visible and later, politically important changes, included commercialization, globalization and mainstreaming, consisting of widespread adoption of the Internet by members of the general public.

Less visible were the changes going on in enabling technologies, particularly the way networks connected with each other and the way data packets were delivered. These technical changes also turned out to have significant political

and commercial ramifications. First of all, in the early period of the Internet's development, the architecture was designed around a single backbone, using the resources of the two government agencies just described. Under this design, all other networks allowed to participate were connected with one another via the single Internet backbone. At the same time, routing information was conveyed between the backbone and the other smaller networks using a long outmoded routing technology called the exterior gateway protocol, or EGP (Halabi, 1997, p. 9).

Today the Internet no longer has a single backbone in the traditional sense. Rather, it has many, comprised of the individual backbones of Tier-1 commercial bandwidth stakeholders (see Appendix B). These modern backbones represent a significant shift in the number and type of stakeholders; their commercial priorities; and the number and types of locations where they interconnect, called peering points. In addition, the EGP has long since been replaced by the crucial routing technology known as the Border Gateway Protocol, or BGP. One of the most compelling reasons for deployment of the BGP is that it facilitates completely decentralized routing - an important feature as the Internet moved away from the central authority model with the decommissioning of the NSFNET (Halabi, 1997, p. 5).

Meanwhile, a parallel transition to commercialization was taking place at the retail level, in the local loop. During the 1980s, small groups of Internet hobbyists dotted around North America developed and ran computer bulletin board systems (BBSs) which were typically free, local operations that used the telephone network to share access. Their most striking feature in the context being described here was that BBSs did not depend for their connectivity on any third-party provider, in particular the commercial ISPs owned by the large telephone and cable companies (which in turn owned and controlled the physical access facilities in the residential last mile). Gradually, beginning in the late 1980s, ISPs began to offer dial-up access to the Internet. From the early 1990s onward, the thousands of small ISPs that sprang up across North America, often referred to colloquially as “Mom and Pop” operations - were steadily bought up by larger ISPs, or forced out of business (Keshav, 1997, p. 43).

In this chapter, several perspectives are developed on the structure of the bandwidth industry and its stakeholders, especially the classification of networks as Tier-1, Tier-2 or Tier-3 operators, as described earlier. Tier-1 operators run the Internet’s backbone networks; Tier-2 operators (in the United States) are mostly comprised of the old RBOCs (Regional Bell Operating Companies) now called ILECs (incumbent local exchange carriers); and Tier-3 operators, a heterogeneous group comprised of CLECs (competitive local exchange carriers) and other digital service providers – a category that includes not only

conventional Internet service providers, but also companies that provide a wide range of access, content, applications and hosting services.

Regional Peering Practices

Outside of the US, most major peering points operate on well established and transparent policies. These peering points include Amsterdam-based AMS-IX (the world's largest); London's LINX; Frankfurt's DE-CIX; and PaNAP, Paris (Bouygues Telecom, n.d.). In Canada, TORIX (Toronto), which has grown substantially; VANIX (Vancouver); and QIX (Montreal) all utilize the multi-lateral peering approach to engage in public peering interconnection (PPI).

Within these Internet exchanges, several different types of network operators function alongside the Tier-1 networks; these include hosting providers (Web hosting services), access providers and content providers (cache servers and application providers). Tier-1 networks impose contractual requirements and bandwidth minimums at a level that usually prevents smaller entities from peering with them. For those stakeholders that can meet such requirements, arrangements are typically made through a private peering link known as a private network interconnection (PNI); such entities are then considered to be Tier-1 operators.

The physical facilities that house major Internet exchange points, known as carrier hotels, operate on two different connection principles: exchanges that utilize the large, purpose-built switches intended for multiple users; and the private network interconnections, or private peering just described. Often one network that another network operator wants to peer with might not be available at the exchange switch even though they are in the carrier hotel. The only way for such a network to interconnect is to peer privately with the other network (Davidson, 2008).

Here too there are differences in regional practices. Writing about a visit to the Palo Alto Internet Exchange (PAIX), Paul Vixie noted that private network interconnections at that exchange outnumbered switch connections by more than 100 to 1. He went on to say that “at exchanges where the Internet Systems Consortium (ISC) has a presence, there will generally be between two and twenty private network interconnections, compared to only one or two switch connections” (Vixie, 2008). The balance between the two connection types is reversed in other global regions, such as Africa. Users of the Kenyan Internet Exchange Point (KIXP), for example, are required to maintain a BGP-compliant router at the facility and all connections are carried out through the exchange switches. As one observer noted in 2008, “at this facility there is no demand for cross-connects or private network interconnections, and no such services are offered as yet” (Mwangi, 2008).

The most important considerations from the perspectives of both the provider and the end-user concerning how interconnection is negotiated are: the cost of interconnection in terms of network resources (the bandwidth expended to carry traffic from another network to another); the perceived value of the connectivity; and latency. For most stakeholders, the motivations for pursuing peering in addition to transit arrangements are financial and technical.

The financial benefits of peering are compelling. Peering decreases reliance on and therefore the cost of purchased Internet transit (this is the single greatest operating expense which networks seek to minimize in telecommunications costs). Internet transit is expensive, so networks seek out peering relationships with others (at zero or reduced cost) that provide more direct traffic exchange and reduce the load on expensive upstream transit services. Packet loss and latency slow traffic consumption, so a service such as peering which improves either of these will increase traffic efficiency. Engineering or technical considerations also encourage peering. Peering lowers inter-network traffic latency, because traffic exchanged between two peering entities is necessarily taking the lowest latency path (Halabi, 1997, p. 44).

Following the 1995 divestiture of the National Science Foundation network and the passage of the 1996 Telecommunications Act, large backbone networks had begun to alter their interconnection terms. The goal for most was to change peer

relationships into supplier-customer relationships (Aronson & Cowhey, 2009, p. 226). This change meant that many peering agreements which had been reciprocal and free became transit agreements requiring legal contracts and fees. Paying for transit from an upstream provider gives the customer access to all network routes in the upstream provider's routing table and in return the upstream provider receives and announces the customer's IP routes on all of its peering and transit interconnections.

Transit may be based on a flat monthly charge or on a volume basis (Mbit/s per month). As traffic volume rises, this type of arrangement can become costly, although some service providers prefer a transit relationship with their upstream provider since service level agreements (SLAs) often accompany transit interconnection contracts rather than peering. These service maintenance agreements offer some guarantees of performance above a simple "best effort" (Medhi & Ramasamy, 2007, p. 294). Regardless of the size of the network, in a peering relationship there are often no SLAs to guarantee rapid resolution of problems; a transit customer relationship generally has more contractual teeth, since money is changing hands.

Some networks began private peering agreements in 1996, shortly after control of the public peering sites (metropolitan-area exchanges and network access points) were divested by the US government. Private peering is based on

dedicated circuit connections established between two networks, each managing one end of the connection. Two networks in a private peering arrangement will typically enjoy improvements in service since reductions in the number of router hops means that fewer packets will be dropped. Fees are generally not charged when these connections are set up because they are mutually beneficial.

Providers with large traffic volumes tend to peer without charge with other large providers, while charging for interconnection with smaller entities.

Why might a network decline to participate in peering? One consideration is traffic asymmetry, i.e. if the traffic exchanged by two networks is not roughly equivalent in volume, the network that has to bear higher costs is likely to decline to participate. A second factor is the need for Tier-1 networks to compete on the basis of performance. Peering with other networks improves the performance of any given network, possibly making it a more powerful competitor. As commercial entities, most large networks are unlikely to peer if they believe they can instead sell transit to other, smaller entities. As large networks began to realize their value as peering partners for smaller players, they began to charge for transit. In 1999, *Telephony Online* reported that “[s]ome charge that the dominant backbone providers – UUNet, Sprint and Cable & Wireless – have used peering to retain what amounts to an oligopoly. Combined, those three providers control approximately 75% of backbone traffic” (Blake, 1999, p. 15).

This completes a brief overview of the basic concepts and industry players that will be referred to frequently through the rest of this paper. In the next chapter, a detailed examination is undertaken of the crucial contrasts between the highly centralized, legacy networks built by the incumbent carriers, and the packet-switched, decentralized networks that make up the global Internet. The chapter looks not only at the underlying technologies, but also at the regulatory treatment of networks in both Canada and the United States. The main purpose of this examination is to assess the effects of deregulation on data reachability and the behavior of the companies that provide both retail and wholesale access services.

Chapter II. Conflict and Confluence: Legacy Telephone Networks and the Emergence of the Internet

The main purpose of this chapter is to describe the historical development of communications networks, especially in the United States. The focus is on explaining how technological convergence and concentration of ownership, combined with deregulation, have allowed a small number of firms to exercise wide-ranging control, i.e. gatekeeping power, over Internet backbone facilities once funded and operated by public-sector institutions, particularly universities and government agencies such as ARPA (later DARPA, the Defense Advanced Research Projects Agency) and the NSF, National Science Foundation. It is argued in the following pages that the risks to data reachability can be traced directly to these historical developments. It is also argued that ISP gatekeeping powers have been exercised in recent years in the residential Internet access business and that ISP behaviors at both wholesale and retail levels pose essentially the same risk to end-users - data unreachability - though usually for different reasons.

The outline of events offered here centres on two distinctly different and originally competing types of communications network: the circuit-switched telephone networks built and operated by telecommunications carriers beginning in the late 19th century, based on engineering solutions conceived for telegraph networks

but optimized for voice telephony; and IP-based, packet-switched networks, which were optimized for data transmission and first developed under the auspices of the US government in the late 1960s, later forming the foundation of what became the global public Internet.

We distinguish throughout this paper between “voice-centric networks,” the principal purpose of which is the carriage of voice traffic, both nationally and internationally; and “data-centric networks,” whose principal purpose today is the provision of Internet bandwidth through peering and transit arrangements made among Tier-1, Tier-2 and Tier-3 networks (referred to collectively as bandwidth providers).

Several historical factors played a role in shaping the contemporary telecommunications marketplace. We have singled out three that are of particular significance to the present discussion: the dominant role played by the United States in the development of networking technologies; ISP gatekeeping powers and the debate over Net Neutrality; and the cultural clash between “Netheads” (the engineering and academic community associated with development of the Internet) and “Bellheads” (the engineering and business community associated with the interests of the developed world’s major telecommunications carriers).

U.S. Role and Influence

First, despite the presence of both legacy telephone networks and access to the Internet in most countries around the globe, the history that concerns us unfolded primarily in the United States - although we also touch on related events in Canada. This consideration applies not only to the Internet itself, whose enabling technologies were originally developed by American researchers funded by the US government, but also to the commercial carriers that have come to dominate the Internet's physical transport infrastructure. At this writing, six of the top carriers considered to be Tier-1 providers are American: AT&T - Dallas, Texas; Level 3 Communications - Broomfield, Colorado; Qwest - Denver, Colorado; SAVVIS - Town and Country, Missouri; Sprint - Overland Park, Kansas; and Verizon Business - Ashburn, Virginia (Thussu, 2006, p. 92; Brown, Hepner & Popescu, 2009, p. 5).

Furthermore, American carriers do business in many other parts of the world, both directly through owned and operated corporate divisions, and indirectly through affiliations and agreements with other carriers. Some have made efforts to influence policymaking in other jurisdictions, one recent example in March 2009 being the use of material originally drafted by AT&T in regulatory documents prepared by the European Commission:

Six MEPs [Members European Parliament] have taken text supplied by the American telecoms multi-national, AT&T, and pasted it directly into amendments tabled to the Universal Services directive in the Telecoms Package. The six are Syed Kamall, Erika Mann, Edit Herczog, Zita Pleštinská , Andreas Schwab, and Jacques Toubon.

AT&T and its partner Verizon, want the regulators in Europe to keep their hands-off new network technologies which will provide the capability for broadband providers to restrict or limit users access to the Internet. They have got together with a group of other telecoms companies to lobby on this issue. Their demands pose a threat to the neutrality of the network, and at another level, to millions of web businesses in Europe.

The Universal Services directive is supposed to set out the rights of users and consumers of telecommunications services in Europe. It should seek to guarantee their rights to access content, services and applications on the Internet, and to a clear channel connection to the open Internet. These damaging amendments seek to do the opposite, and close off the Internet by limiting rights.

As previously reported on *iptegrity.com*, and also in the *International Herald Tribune*, a lobbyist for AT&T has been pushing these

amendments around Brussels for several months. The text has been widely circulated. It therefore does astonish me that MEPs would so blatantly table it. (Horten, 2009)

The disproportionate transnational influence of American-based backbone providers has an important bearing on global interconnection, because the wave of deregulation that followed passage of the 1996 US Telecommunications Act combined with increased concentration of ownership, provided the major US carriers opportunities to exercise their market power.

Gatekeeping and Net Neutrality

There are significant parallels between the ability of the Tier-1 carriers to exercise gatekeeping powers at the backbone or wholesale level, and the ability of carriers that control last-mile facilities to exercise retail gatekeeping powers. In the retail sector, North America's incumbent ISPs enjoy a strong advantage over unaffiliated service providers in competing for customers at the application layer. Moreover, because retail Internet access rates are not regulated in the US and Canada, and there is little or no facilities-based intramodal competition in either jurisdiction, the incumbents (both telcos and cable MSOs) are able to manipulate

both the amount and type of data traffic initiated by their broadband subscribers,³ by means of the practice known as traffic-shaping.

The most salient parallel between the behavior of Tier-1 providers and incumbent residential ISPs concerns the risks posed by gatekeeping to the fundamental principle of free, open access to the public Internet and the reasonable expectation by end-users that, whatever the nature of their data or purpose of their transmission, they will not be prevented from communicating over the Internet by willful, arbitrary actions taken by bandwidth providers that are harmful to end-users. In the US, there is another factor that ties together wholesale and residential gatekeeping: common ownership. Thus among the American firms assumed to be Tier-1 providers, three of them - AT&T, Qwest and Verizon - also have corporate divisions that provide residential Internet access (DSL over copper lines, as well as high bandwidth services over fiber-optic lines; Verizon calls its FTTx (Fiber to the home, office or other location) service "FiOS" (Fiber-

³ In this research, "narrowband" and "broadband" refer to the rate at which bits of data are physically transferred through a network per unit of time, typically one second. This measure is known as the bitrate, informally as the channel "speed." Narrowband refers to the bandwidth of a dial-up Internet connection, restricted to 56,000 bits per second (56 Kbit/s). The term "broadband" is used to refer to a family of retail access technologies, particularly digital subscriber line (DSL), high-speed cable (DOCSIS or Data Over Cable Service Interface Specifications), and wireless platforms such as WiMAX (Worldwide Interoperability for Microwave Access, providing wireless transmission of data). The most important feature shared by these platforms is they provide a persistent connection, i.e. they eliminate the need to dial in to an ISP for each online session. A prominent issue for regulators is the speed at which a connection is deemed to be "broadband." In June of 2008, the FCC raised the minimum downlink speed qualifying a service as "broadband" from 200 Kbit/s to 768 kbit/s. In Canada, the CRTC continues to maintain a distinction between high-speed Internet access service, which includes speeds at or above 128 Kbit/s and broadband service, which includes speeds at or above 1.5 Mbit/s (CRTC, 2009c, p. 213, notes 229, 230).

Optic Service), while AT&T operates a residential service called “U-verse” that is a mix of VDSL (Very high bit-rate DSL) and FTTx.

Retail gatekeeping issues of the kind just noted were first analyzed in depth by Tim Wu in the paper “Network Neutrality, Broadband Discrimination” (Wu, 2003). The principle of Net (or Internet) neutrality rests on the assumption (contested by certain interested parties, including many large ISPs) that a useful public network - typically understood to be the Internet - should carry all data traffic in a non-discriminatory fashion, while users of the network, especially the customers of last-mile broadband ISPs, should enjoy non-discriminatory access to the content and services of their choosing. This principle reflects both the traditional behavior expected of true common carriers and the end-to-end design principle underlying the structure of the Internet. Although Net Neutrality has been actively promoted and discussed over the last three years by a wide range of advocacy groups, the FCC and CRTC have only recently begun to pay close attention to the public interest issues at stake. As evidence for Canada, the CRTC “traffic-shaping” hearings scheduled for summer 2009 (CRTC Telecom Public Notice CRTC 2008-19, 2008b).

Clash of Two Cultures

A third theme that emerges from the history of network development is what can be described as a clash of cultures between engineers who built telephone

company networks and those who built networks based on packet switching and internetworking protocols, culminating in the Internet protocol suite.

Several attributes of early packet-switched networks constituted a clear rejection of basic design principles followed by AT&T's engineers in the creation of the local and long-haul telephone networks that served the continental United States for many decades. Telephone networks were designed and built to operate in a highly centralized fashion, using expensive, "intelligent" components in the network itself to link what were essentially dumb terminals, and to do so in a way that guaranteed a high degree of reliability.

By contrast, most of the design work on the TCP/IP protocols that began in the early 1970s was based on the "end-to-end" principle, according to which the communications control functions of a network are, whenever feasible, pushed out to "smart" devices and applications at the edges of the network. IP-based networks, including the public Internet, are for this reason referred to as "dumb" - i.e. the intelligence is at the edges of the network and the Internet protocol itself (IP) does not provide for reliable transmission. The end-to-end design has significant advantages over the traditional design of the circuit-switched telephone network, such as lower capital and operating costs and much greater possibilities for innovation. On the other hand, these advantages were only achieved by making certain trade-offs: thus, the high quality of service attained

by AT&T's voice networks was sacrificed for the "best effort" standard that is built into the public Internet (Clark, Reed & Saltzer, 1984, pp. 277-288).

These differing approaches to network design were reflected in two distinct and competing sets of values and attitudes not merely towards technical details, but also towards the social and economic applications of these technologies. These differences have been encapsulated in the monikers "Netheads" and "Bellheads," terms which appeared in the popular press as long ago as 1996, in a *Wired* magazine article "Netheads vs Bellheads" (Steinberg, 1996, p. 1).⁴

In broad strokes, Bellheads are the original telephone people. They are the engineers and managers who grew up under the watchful eye of Ma Bell and who continue to abide by Bell System practices out of respect for Her legacy. They believe in solving problems with dependable hardware techniques and in rigorous quality control - ideals that form the basis of our robust phone system and that are incorporated in the ATM protocol.

Opposed to the Bellheads are the Netheads, the young Turks who connected the world's computers to form the Internet. These engineers see the telecom industry as one more relic that will be overturned by the

⁴ The feuding between these two camps went all the way back to the early 1960s, when Paul Baran set out to improve on what he mockingly described as Ma Bell's costly, "gold-plated" equipment - a feeling reciprocated by certain AT&T engineers, who felt Baran simply didn't understand good network design (Roland, 2000, p. 827).

march of digital computing. The Netheads believe in intelligent software rather than brute-force hardware, in flexible and adaptive routing instead of fixed traffic control. It is these ideals, after all, that have allowed the Internet to grow so quickly and that are incorporated into IP - the Internet Protocol. (Steinberg, 1996, p. 1)

This theme subsequently made its way into the scholarly literature, notably in Rob Frieden's 2001 paper, *Revenge of the Bellheads: How the Netheads Lost Control of the Internet* (Frieden, 2001). The general conclusions drawn in these two articles, and the date of their publication, are revealing. The date of the Wired article, 1996, marked the beginning of both the commercialization and mainstreaming of the Internet. The NSFNET had been shut down the previous year (April 1995), followed a few months later by the Netscape initial public offering (August 1995), widely credited with triggering broad public interest in the Internet and the Web for the first time.

No less significantly, 1996 was also the year that the ISO, the International Organization for Standardization, and the ITU-T, Telecommunication Standardization Sector (which administers telecommunications standards for the ITU) finally abandoned their two-decade attempt to develop and implement the OSI (Open System Interconnection) standard. This effort, which had wide international support from the telecommunications industry, was launched in

1977 and, as explained below, shared a general goal with TCP/IP: to create a common internetworking platform for computer and telephone networks that were otherwise not interoperable.

In his 1996 article, Steinberg framed the cultural battle between the telephone company engineers and those committed to the Internet protocol suite in narrow technical terms: the attempt by the RBOCs (Regional Bell Operating Companies) to introduce a controversial telecommunications technology called ATM (asynchronous transfer mode), a packet-switched protocol for data networking that has failed to make serious inroads into next-generation networking technologies. Steinberg's narrative was set in the context of a two-day meeting of NANOG, the North American Network Operators' Group, which had drawn senior engineering staff from America's largest ISPs:

The battle over whether to adopt ATM or to extend IP is likely to be the deciding fight between the Bellheads and the Netheads. The two protocols embody very different visions of communications leading to connected worlds with different social patterns, commerce, and even politics. In extreme terms, think of the difference between the chaotic world of the Web and the rigorously controlled, financially lucrative world of 900 numbers. The first reflects the technology of the Netheads, the second the technology of the Bellheads. (Steinberg, 1996, p. 2)

From the vantage point of 1996, Steinberg and many of his interviewees took two observations for granted. First, the debate over ATM was going to be the main point of contention among the engineers responsible for designing the networks of the future: "If the two sides agree on anything, it is this: As ATM goes, so goes cyberspace" (Steinberg, 1996, p. 2). Second, it was assumed by TCP/IP devotees that ATM was an inferior technology, and that the RBOCs were too committed to their legacy voice networks to adapt to newer - and better - internetworking principles. Opponents of ATM and the telephone company culture with which it was associated predicted - more or less correctly - that IP-based networks would eventually win out over ATM and the assumptions that went with it.

Five years after the *Wired* feature, Prof. Rob Frieden made a persuasive case that the clash of two cultures as described by Steinberg in *Wired* had led to a very different and improbable set of outcomes. Frieden's overall conclusion is reflected in the title of his paper: *Revenge of the Bellheads: How the Netheads Lost Control of the Internet*.⁵

⁵ The five-year period from 1996 to 2001 was marked by prodigious growth of the Internet. For example, in October 1996, the month Steinberg's *Wired* article was published, there were 460,000 sites on the Web. By October 2001, when Frieden's paper was published, the number of sites had grown by over 7,000%, to 33 million (Zakon, 2006).

As Frieden shows, narrow technical arguments such as that over ATM were eclipsed by much broader regulatory, economic and industrial issues.

Furthermore, the once sharp divergence between the technologies favored by Netheads and Bellheads respectively was on its way to extinction, as the onward march of digital convergence pushed all forms of network traffic into the IP camp - including voice. And most importantly, the large American carriers, including the former Baby Bells, were beginning to exercise a high degree of control over Internet access and imposing their own corporate view on how the Internet should be allowed to develop:

[This] paper identifies commercial developments in the Internet to support the view that the Internet has become more hierarchical and more like telecommunication networks, i.e., that both markets and technologies have converged. The paper concludes that despite such convergence, not all stakeholders have adjusted their perspective on regulation and the likely scope of governmental oversight. Such failure to adjust creates opportunities for stakeholders, well-versed in the telecommunications regulatory environment, to exploit this expertise if, as anticipated, Internet services become subject to some degree of regulatory oversight as a result of technological and marketplace convergence. Telecommunication carriers have the resources and wherewithal to dominate the Internet, primarily through ownership of most local and long haul transmission

conduits and by owning the largest Tier-1 ISPs that control most of the broadband backbone facilities providing Internet data transport.

Additionally the Bellhead culture has required incumbent telecommunications service providers to develop superior skills in working the regulatory process to their advantage. The paper suggests that Bellheads will outmaneuver Netheads particularly if the revenue siphoning effect of Internet-mediated services offsets the revenues generated from ISP leases of telecommunication transmission capacity. (Frieden, 2001, p. 3)

Frieden goes on to cast the rise of Bellhead control in terms of a change in the Internet's structure from flat to hierarchical. The practice of peering among ISPs was the norm for many years, not only because that was the tradition inherited from ARPA and the NSF, but also because "[t]he relatively small number and homogeneity of ISPs in terms of size, bandwidth, subscribership and geographical coverage made it plausible to assume that traffic flows were symmetrical" (Frieden, 2001, p. 9). The sheer growth of the Internet during the 1990s, along with significant changes in the US telecommunications policy framework, meant a growing gap between small, local ISPs and the large, national ISPs, many of which were owned by the Baby Bells and their affiliates, and which absorbed or put out of business hundreds of their smaller rivals (Cukier, 1997).

As a consequence, notes Frieden (p.10), business practices changed at both the retail and wholesale levels. The biggest Tier-1 networks, “having made the largest investments in response to demand, no longer could rely on governments for reimbursement” - a compelling reason to move from settlement-free peering to paid transit, while at the same time creating private peering points, in part as a response to congestion at the older public peering points. Frieden points to three major consequences arising from the decision by Tier-1 networks to move to private peering:

- 1) management, governance and control of the Internet has become more hierarchical;
- 2) Tier-1 ISPs have leveraged their superior bargaining power to assume primary control over the Internet; and
- 3) As telecommunications service provider incumbents own most of the Tier-1 ISPs, the Bellhead culture and mode of operation have become dominant.

Chronology

1962 - Paul Baran publishes first RAND Corp. research on packet switching, funded by the US Air Force.

1969 - First working 4-computer implementation of ARPANET.

1974 - Vint Cerf and Bob Kahn publish their early work on TCP, the transmission control protocol, specifying a method for linking different networks - internetworking - using packet-switching technology developed by Paul Baran and Leonard Kleinrock. Later expanded to TCP/IP, more accurately the Internet protocol suite.

1977 - ISO and ITU-T begin work on Open System Interconnection (OSI) standard for internetworking, in competition with TCP/IP.

1982 - Consent decree initiates breakup of AT&T, which is ordered to divest its local exchange operations. In 1984, they are spun off into the seven original RBOCs (Regional Bell Operating Companies; also known as the Baby Bells). The decree becomes known as the Modification of Final Judgment (MFJ), supervised by Judge Harold Greene.

1983 - ARPANET fully converted to run on TCP/IP.

1986 - NSFNET launched by National Science Foundation, replacing ARPANET as main Internet backbone.

1991 - World Wide Web becomes publicly available on the Internet. NSF lifts restrictions on commercial use of the Internet.

1995 - NSFNET is closed down; a growing number of ISPs provide Internet backbone facilities; Netscape IPO creates wide public interest in the Web.

1996 - The ISO and ITU-T close down development work on the OSI standard; the OSI 7-layer reference model remains in wide use for teaching.

Figure 1. Internet timeline from Zakon, R. H. (2006). Hobbes' Internet Timeline v8.2.

Rival Solutions to the Problem of Internetworking

Up until the mid-1970s, communications networks around the world were built and used mainly by telephone companies, and apart from some data applications

like facsimile transmissions (faxes), these networks were optimized for voice traffic and operated using circuit switching and analog technology. These latter two characteristics were a fundamental part of the prevailing network design, and the two areas where emerging data-centric (eventually IP-based) networks parted company with their predecessors. Although operators in different regions of the world used different and often incompatible signaling technologies, a number of systems were developed and maintained for making them interoperable, and for permitting international calling. This work was carried out by the ITU-T, Telecommunication Standardization Sector, which coordinates international standards-setting on behalf of the ITU, International Telecommunications Union. The ITU-T continues to be responsible for many different telecommunications standards known as “Recommendations” with no enforcement power (Noam, 2001, p. 118).

The most important single attribute of legacy voice-centric networks is their reliance on circuit switching. When a call is set up between two points on a circuit-switched network, a dedicated circuit is opened for that call and that call only, and the full bandwidth originally made available is kept in service even if the two parties are not communicating. A dedicated circuit used to set up telephone calls allocates a fixed amount of bandwidth: 4 kilohertz (kHz) in an analog circuit and 64 Kbps in a digital circuit. No other users can share a circuit until the

original call is torn down and the network resources are put at the disposal of other users.

If the circuit is cut or disrupted, the call is dropped until another dedicated circuit is opened. This design limitation offered corresponding benefits for telephone subscribers using circuit-switched networks, especially the quality of service (QoS) it provides. In Canada and the United States, AT&T and its sister companies, notably Bell Canada (owned by AT&T until 1956), built telephone systems that achieved notable degrees of reliability in voice quality, uptime and the availability of dialtone.

Despite the benefits associated with telephone network QoS, and the well-entrenched engineering assumptions behind the design of voice-centric networks, the 1970s marked the beginning of a growing unease about interoperability among many of the world's telecommunications carriers, equipment vendors and standards-setting bodies (Wu, 2007, p. 6). Established carriers were faced with many of the same problems that American academic and military researchers had begun to address in the 1960s. Chief among these telecommunications problems was how to bring order - and interoperability - to the scores of networking protocols developed and supported by both national governments and private firms, a list that included Appletalk, DECnet and NetWare among many others.

In broad terms, telecommunications carriers faced three challenges. One, their networks though highly reliable, were inefficient and costly, because of their reliance on circuit switching. Two, they were purpose-built, optimized for a single application - setting up, carrying and tearing down voice telephone calls - and not readily adaptable to other applications. Three, the international bodies responsible for standards-setting - the CCITT (from the French: Comité Consultatif International Téléphonique et Télégraphique), formed in 1956, then reconstituted as the ITU-T in 1993 - were notoriously slow at developing and approving standards, a major barrier to innovation (ITU World Telecommunication Standardization Assembly, 2000).

The established national telephone companies shared a further challenge. They were running production networks, not research networks, and provided services that hundreds of millions of telephone subscribers depended on every day. Thus, any pressure for changes or improvements in areas such as interoperability had to be balanced against the need to ensure that dialtone was ever-present. The telephone companies also had to find methods to ensure that calls made in one country could be terminated in any other country that offered international service - a challenge not only technical in nature but financial as well, since settlement procedures had to be put in place to ensure carriers were compensated for handling traffic on behalf of other carriers.

The global network of telephone company networks, termed the PSTN, for public switched telephone network, went through many design and standards changes at both national and international levels throughout the 20th century, with many of the innovations in switching and routing technologies stemming from work undertaken at Bell Labs, AT&T's R&D division. Two of the major engineering challenges facing telephone companies and standards-setting bodies were addressing systems and signaling protocols. The E.164 addressing scheme was developed to provide PSTN subscribers in every participating country with a unique code identifying their telephone sets (known in the industry as CPE, customer premises equipment).

It was in the area of signaling that the industry made the most significant advances in the efficiency and functionality of the PSTN, advances that borrowed heavily from the renegades who had developed packet switching and TCP. Within three years of the initial work completed by Vint Cerf and Bob Kahn on TCP, the CCITT was at work on a radically new method of signaling (beginning in 1975). Two years later, the ISO launched a development process that would lead to two decades of work on the OSI standard. Both of these projects borrowed from work on packet switching and digitization. In the case of OSI, the development effort was also increasingly seen as a direct competitor to the Internet protocol suite (Russell, 2006, p. 53). The competition between the Internet protocol suite and OSI was about organizational ethos as much as it was

about the technologies themselves - echoing the cultural conflict between Netheads and Bellheads. Andrew L. Russell describes the Internet-OSI standards war in the following terms:

During the 1980s and early 1990s, OSI enjoyed widespread support from national governments, particularly in Western Europe, North America, and the Far East. OSI enjoyed this level of support due in part to the strategic position of its sponsor, ISO. ISO was an “official” international standards body, meaning that it was populated by representatives from national governments who, in most cases, acted on behalf of the interests of their national telecommunications and computer firms. ISO’s organizational culture—concerned with defining and controlling the future of information and telecommunication services on behalf of its representatives from national governments—resembled contemporary democratic bodies insofar as it featured voting, partisan compromises, and rule-making behavior designed to protect financial interests. Such processes stand in stark contrast to the research and military orientation of the people and institutions that developed Internet protocols. (Russell, 2006, pp. 52-53)

The Bellhead standards culture achieved significant successes. In 1980, the CCITT finalized the specification for Signaling System #7 (usually referred to as SS7), marking a dramatic improvement in functionality over prior telephone

signaling protocols, used to open and close voice telephone calls. The SS7 protocol suite was of great importance to the industry in financial and business terms, because the out-of-channel signaling technology that formed the basis for SS7 allowed for the creation of most of the enhanced calling features developed by telephone companies over the last quarter century. In the coming years, however, telephone companies will be facing ever-increasing challenges from Internet-based telephony, usually referred to as voice over IP, or VoIP.

OSI Model			
	Data unit	Layer	Function
Host layers	Data	7. Application	Network process to application
		6. Presentation	Data representation and encryption
		5. Session	Interhost communication
	Segment	4. Transport	End-to-end connections and reliability
Media layers	Packet	3. Network	Path determination and logical addressing
	Frame	2. Data Link	Physical addressing
	Bit	1. Physical	Media, signal and binary transmission

Figure 2. OSI Layered Model from Black, U. (2000). IP Routing Protocols, Appendix A, p. 218.

The single most revolutionary feature shared by early versions of the Internet, SS7 and OSI was the notion of logical “layers” in a network, which together form a protocol “stack.” The rationale for incorporating layers into network design is that the resulting structure makes network applications, functions and hardware mutually independent, which in turn affords a far greater degree of connectivity and interoperability among different kinds of hardware and software. This was the primary goal behind the development of protocols for internetworking such as that undertaken in the ARPA-funded research that led to the Internet protocol suite. Whereas a lack of flexibility was one of the leading disadvantages of telecommunications networks optimized for voice, Internet protocols allowed applications to request network functions independently of the characteristics of the underlying physical network (Keshav, 1997, p. 69). Each protocol layer - which are “logical” or abstract, rather than physical - calls for services from the layer immediately below and provides services for the layer immediately above.

A layer does not define a single protocol; it defines a data communications function that may be performed by any number of protocols. Therefore, each layer may contain multiple protocols, each providing a service suitable to the function of that layer. The individual layers do not need to know how the layers above and below them function; they only need to know how to pass information to them. Isolating network communications functions in different layers minimizes the impact of technological change on the entire protocol suite. Consequently,

new applications can be added without changing the physical network, and new network hardware can be installed without rewriting the application software.

Some confusion exists over the TCP/IP nomenclature and the number of layers used by the Internet. This confusion arose in part because the OSI reference model incorporates a seven-layer design, while the Internet protocol suite comprises four layers, which together cover the functions reflected in the OSI seven-layer model. The Internet protocol suite layers are: the application layer (for example BGP, HTTP, FTP); the transport layer (for example TCP, UDP); the Internet or internetworking layer (for example IP) and the data link layer, incorporating the physical layer (for example PPP, Ethernet).

In simple terms, TCP ensures reliable end-to-end transport of data packets, through functions such as error correction, whereas IP handles the addressing and routing of data packets. The Internet Protocol itself enables packets to travel by many different routes to their destination, based on the addressing encoded into packet headers. It is a connectionless protocol, meaning IP can begin to deliver packets before any connection or circuit is established between two communicating host computers. And it does so by encapsulating data packets in a way that makes them able to traverse networks of almost any type. The

corresponding disadvantage is that IP is a best-effort delivery method and is not considered reliable on its own.⁶

As the TCP/IP suite was refined through the 1970s and early 1980s, work proceeded independently on the OSI protocol suite. Although they were both intended to accomplish the same major goals, and both utilized the modular or layered stack design, the two approaches differed in one critical respect. The TCP/IP suite was intended from the outset to act as a bridging technology that would connect existing network protocols, whereas the ultimate purpose of the OSI suite was to eliminate internetworking problems by replacing other platforms. The development of OSI generated controversy throughout the 1980s, as part of the longstanding feud between the TCP/IP community - the Netheads - and the engineering community whose affiliations lay with the established telecommunications carriers and their standards-setting bodies - the Bellheads.

Those opposed to the deployment of OSI pointed to two particular issues. One was the high costs associated with the goal of replacing existing network software and hardware on a massive scale. The other concerned the glacial pace

⁶ It should be noted that for certain kinds of transmissions, TCP is not the preferred protocol, particularly for time-sensitive delivery in which the priority is minimizing delays rather than error correction. In the transmission of messages like emails, which are not time-sensitive, TCP may resend dropped packets and perform other functions that introduce latency. These actions do not affect the integrity of the message and will normally not be noticed by the end-user. By contrast, audio and video streams, which are isochronous transmissions, cannot tolerate delays, meaning it is preferable to lose packets rather than have the recipient wait for retransmitted packets. In such situations, the network will use an alternative to TCP, such as UDP, user datagram protocol.

at which the ISO developed specifications. Nevertheless, even though OSI was considered by many in the field to be too complicated and to a large extent unworkable, it enjoyed widespread governmental support, including the support of the US government. Here is how computer scientist Gary C. Kessler describes the see-saw battle in *An Overview of TCP/IP protocols and the Internet*:

In 1988 ... the DoD and most of the U.S. Government chose to adopt OSI protocols. TCP/IP was now viewed as an interim, proprietary solution since it ran only on limited hardware platforms and OSI products were only a couple of years away. The DoD mandated that all computer communications products would have to use OSI protocols by August 1990 and use of TCP/IP would be phased out. Subsequently, the U.S. Government OSI Profile (GOSIP) defined the set of protocols that would have to be supported by products sold to the federal government and TCP/IP was not included.

Despite this mandate, development of TCP/IP continued during the late 1980s as the Internet grew. TCP/IP development had always been carried out in an open environment (although the size of this open community was small due to the small number of ARPA/NSF sites), based upon the creed "We reject kings, presidents, and voting. We believe in rough consensus and running code." (Dave Clark, M.I.T.) OSI products were still a couple of

years away while TCP/IP became, in the minds of many, the real open systems interconnection protocol suite. (Kessler, 2007)

By 1994, thanks to intense pressure from the main Internet engineering bodies, particularly the IETF (Internet Engineering Task Force), and the ISO's inability to move the OSI suite from theory to implementation, "the grand future planned for OSI was on the rapid decline" (Russell, 2006, p. 56). As we will see below in chapter 3, Russell's thesis that the bureaucratic innovations of the Internet pioneers were as important as their technical achievements does much to explain the long, difficult struggle to find workable solutions to Internet governance:

The Internet standards community not only introduced technological innovations; it also pioneered organizational innovations such as the use of electronic mailing lists to build consensus around technical work, a process open to all interested stakeholders, and the free and unrestricted distribution of standards. (Russell, 2006, p. 56)

Deregulation and Privatization: the Internet Becomes a Business

By 1996, the two-decade attempt by the ISO and the developed world's telephone carriers had ground to a halt, having finally lost the "religious war" (Russell, 2006, p. 56) over a global internetworking standard to the proponents of

TCP/IP. But thanks to a series of gradual changes that had unfolded in parallel, the American carriers in particular were beginning to consolidate their market power - and most importantly for our purposes, move into the data networking and ISP businesses. These developments were to have a profound effect on the next Internet war: that between powerful corporate gatekeepers and proponents of a free, open and robust public Internet. The developments in question were of three different kinds:

- The adoption by incumbent carriers of packet-switched, layered networks for most purposes, including their traditional voice business.
- The collapse of the content/carrier distinction as part of the deregulatory climate that was sweeping the developed world.
- And the privatization of Internet backbone resources that began in earnest with the closing of NSFNET in 1995.

The development of protocol layers in network design had an important long-term impact on the behavior of telecommunications carriers. In the pure common carrier model of legacy telephone networks, the core application - local and long distance voice telephony - is inseparable from the physical facilities carrying the application. However, the separation of the physical and data link layers of a network from the applications layer means that many services other than traditional voice telephony can be offered independently of the underlying

physical network and the carrier that controls the physical network. These fundamental technical changes would not have had a significant impact - such as creating opportunities for leveraging monopoly control of last-mile facilities - without equally fundamental changes to the framework for regulating telecommunications carriers, in both the US and Canada.

The deregulatory environment in American telecommunications is widely associated with passage of the 1996 Telecommunications Act. Yet the policy goals associated with deregulation, and the corresponding reliance on market forces to generate competition, predate the 1996 Act. A watershed divide in the early 1980s brought radical changes to the market for voice telephony and related services in the US: the 1982 consent decree that led to the breakup of AT&T and the introduction of sanctioned competition into the telephone market for the first time in modern US history. A few years later, Canada made its first move away from the regulated monopoly model for telephony to the competitive model, not through legislation but the decision rendered by the CRTC in 1992 to dismantle the long-distance monopoly then enjoyed by Bell Canada and its sister incumbents across Canada. In Telecom Decision CRTC 92-12, issued June 12, 1992, the federally-regulated telephone companies' monopoly in the provision of public long-distance voice telecommunication services was ended (CRTC Telecom Decision 92-12, 1992).

The US telecommunications market was built, like most, on the assumption that services such as voice telephony constituted a natural monopoly, making direct competition in telecommunications uneconomic, especially in the local loop, where transport facilities constitute a “bottleneck.” Since consumers and businesses could turn to only one provider for service, and basic telephone service was considered essential, that provider - AT&T in the US - had to meet several far-reaching requirements.

Services had to be priced on the basis of approved tariffs; AT&T had to meet public service obligations, such as making service available at affordable prices in high-cost areas; and treat all customers in a non-discriminatory fashion by being designated a common carrier, the scope of whose business was restricted to transmission. AT&T was also explicitly forbidden from manipulating or in any way interfering with the content or payload transmitted by customers over their facilities. Nor could AT&T create or disseminate content, or carry on any content-related business, such as broadcasting or publishing. In exchange for operating under these restrictions, AT&T enjoyed not only a monopoly on public telecommunications services in the US, but also the right to operate as a vertically integrated company combining basic research (Bell Labs), equipment manufacturing (Western Electric), and local and long-distance telephony.

AT&T's monopoly was brought to an end in one of the largest and most protracted anti-trust suits in US history, initiated by the US Department of Justice (DoJ) in 1974 (the year Vint Cerf and Bob Kahn published their first work on a protocol for linking packet-switched networks). The suit was settled in 1982 under the consent decree approved by Judge Harold Greene, and finally came into effect in 1984. AT&T was ordered to divest its local exchange operations, which were spun off into the seven original RBOCs also known as the Baby Bells.

Although the DoJ anti-trust suit was launched to punish AT&T for breaking the law, the larger rationale included the overarching policy goal of opening the American long-distance market to competition. For voice telephony or any service running over a public network, meaningful competition requires that new entrants be allowed to interconnect with the facilities owned by incumbent providers on reasonable terms. The power of incumbents to dictate or influence the terms governing interconnection has had a major impact not only on competition in the local loop, but also on the evolution of the Internet as a hierarchical structure - i.e. the tendency of bandwidth providers to develop vertical relationships in which paid transit has gradually replaced unpaid peering. As Eli Noam explains, AT&T's resistance to allowing local interconnection with rival firms gave the US government a persuasive argument for forcing the telecommunications monopolist to divest its local operations:

The government's main argument for splitting up AT&T was that it was inherently incapable of providing its long-distance competitors with equal interconnection to its local network. To buttress its case, the Justice Department introduced evidence of AT&T's resistance to competition through its unequal interconnection arrangements. Since regulatory requirements did not work in the face of AT&T persistence it was necessary, the government argued, and the court basically agreed, to split off the company's local operations, which had been the source of its bottleneck power. Thus the resistance to interconnection brought down the world's foremost telecommunications provider. (Noam, 2001, p. 37)

Meanwhile, throughout the entire period during which the AT&T suit was being litigated, the FCC was carrying on its own proceeding aimed at modernizing certain provisions of the 1934 Communications Act. This was the protracted series of investigations known as the *Computer Inquiries*, which began in 1966 and were divided into three proceedings, designated *I*, *II* and *III* respectively. The various orders issued as part of these proceedings were "designed to control the physical layer monopoly power that AT&T's Bell System and its progeny then exercised in providing the links over which distant computers could 'talk' to each other" (Nuechterlein and Weiser, 2007, p.151). In other words, despite important differences between them, the FCC and DoJ proceedings shared the general

goal of containing the market power of AT&T, and subsequently of the seven RBOCs spun off as a consequence of the divestiture.

Nevertheless, the *Computer Inquiries* had a more proactive purpose, and that was to reduce barriers to the growth of the fledgling data processing industry, which made growing use of telecommunications networks (Zarkin, 2003, p. 288). Given the vested interest AT&T and eventually other carriers had in data communications, the Commission's task was to promote competition in data services, while keeping incumbent carriers from exercising undue preference in the provision of such services. In 1981, as part of the *Computer II* orders, the Commission established a crucial distinction between conventional services as offered by common carriers, with their attendant obligations, and a new class of services that were to be exempt from such obligations - dubbed basic and enhanced services respectively. The FCC defined a basic service as a pure transmission capability over a communications path that is virtually transparent in terms of its interaction with customer supplied information - what is normally meant by a common carrier transmission service. The Commission then set out a contrasting concept it called "enhanced" services, i.e. services, offered over common carrier transmission facilities used in interstate communications, which involve the processing or alteration of the content of the message in some way by a computer (Zarkin, 2003, p. 294).

In Canada, a similar distinction has been created in the regulatory framework, though for entirely different reasons. Unlike the US, the Canadian communications industries are regulated under two entirely separate pieces of enabling legislation. The *Telecommunications Act* (which in 1993 replaced a patchwork quilt of statutes and orders) governs the activities of undertakings that offer telecommunications services, which includes traditional services like voice telephony and non-traditional ones like Internet access. The *Broadcasting Act*, which last underwent a major overhaul in 1991, governs the activities of undertakings that offer television, radio and similar services, as well as an important category known as broadcasting distribution undertakings, or BDUs. Although cable television was not even mentioned in the 1967 version of the *Broadcasting Act*, it came to play a special role in the promotion and protection of Canada's cultural sovereignty.

Although both these statutes are administered, as in the US, by a single regulatory tribunal - the Canadian Radio-television and Telecommunications Commission or CRTC - they have very different policy goals. The *Telecommunications Act* is primarily concerned with common carriers and the various protections and responsibilities, such as universal service and affordability, typically associated with core telecommunications services like voice telephony. The *Broadcasting Act*, on the other hand, spells out a broadcasting policy for Canada that is concerned with social goals and the

various ways in which broadcasters are expected to contribute the nation's cultural well-being.

Despite major differences in the Canadian and American policy frameworks, both nations are now struggling with how to adapt outdated provisions, intended for a world of conventional media, to the post-convergence communications culture brought about by digital technology. Indeed, on June 4, 2009, in releasing its decision on its new media proceeding, the CRTC took the unusual step of asking the Federal Court to settle a highly contentious issue involving the interaction of the two enabling statutes and the role of ISPs in supporting Canada's broadcasting system. In the review of broadcasting in new media - Broadcasting Regulatory Policy CRTC 2009-329, issued June 4, 2009; this is how the Commission framed the problem in its decision:

The Commission notes that, pursuant to subsection 4(4) of the *Broadcasting Act*, a telecommunications common carrier, as defined in the *Telecommunications Act*, when acting solely in that capacity, is not subject to the *Broadcasting Act*. Likewise, pursuant to section 4 of the *Telecommunications Act*, that statute does not apply in respect of broadcasting by a broadcasting undertaking. The legal issue as to whether ISPs are subject to the *Broadcasting Act* raises fundamental questions regarding the distinction, for the purpose of the *Broadcasting Act* and the

Telecommunications Act, between telecommunications common carriers and broadcasting undertakings. (CRTC, 2009a, para 68)

The issue raised here by the CRTC is an allusion to the asymmetric treatment of telecommunications carriers under Canada's *Telecommunications Act* and cable-TV providers under the *Broadcasting Act* – a dilemma that has its origins in the use of cable providers by federal authorities as gatekeepers to the broadcasting system (which allowed the CRTC to mandate preferential treatment of Canadian programming, as reflected in tiers of service and must-carry provisions in the regulatory framework for broadcasting).

The FCC and Internet Backbone Policy

The three proceedings conducted by the FCC during the 1960s, 1970s and 1980s - known as *Computer Inquiry I*, *Computer Inquiry II* and *Computer Inquiry III* respectively - were to have a significant impact on the behavior of telecommunications carriers and ISPs (Cannon, 2003, pp. 188, 190, 199). The goal of these proceedings was to determine the extent and nature of federal regulatory involvement in the converging computer and telecommunications industries, in particular, the extent to which common carriers could employ data processing techniques in furnishing services and what regulations might be required to manage the changing marketplace (Noam, 2001, p. 176). *Computer Inquiry I*, which began in 1966 and wound up in 1971, focused on preventing

AT&T from using its dominant market power to discriminate against newly emerging service providers. In this Inquiry, the FCC concluded that it was not necessary to regulate data processing services. With limited exceptions, carriers other than AT&T were permitted to set up a separate subsidiary if they wanted to provide data processing services. A consent decree then in force barred AT&T from offering any unregulated services or products (Cannon, 2003, pp. 178-9).

Computer Inquiry II, completed in 1980, was significant for the crucial distinction it established between what were termed “basic” and “enhanced” services. The main challenge for the FCC in this proceeding was how to deal with the rapidly expanding market for data services and the desire on the part of AT&T and other carriers to offer such services, now known as enhanced services, by contrast with the basic services associated with the transmission of information on a common carriage basis (which are subject to the Title II provisions of the 1936 Communications Act). The FCC ruling provided that telecommunications companies were allowed to provide enhanced services but only through structurally separate subsidiary companies (Noam, 2001, p. 176). While the Commission recognized that it was not feasible to bar AT&T, and later the divested Bell Operating Companies, from developing enhanced services, it also wanted to protect ratepayers from having to subsidize any new, unregulated equipment and services. *Computer Inquiry II* also deregulated customer

premises equipment (such as telephone sets and PBXs), allowing purchase from other vendors (Cannon, 2003, p. 185).

In *Computer Inquiry III* (1985 to 1989), the “basic” versus “enhanced” division was preserved; but the structural separation requirement (offering enhanced services through a subsidiary) was replaced by other protections related to the unbundling of network elements, namely comparatively efficient interconnection (CEI) and open network architecture (ONA). These obligations, which applied to telephone companies only (and not other kinds of carriers, like cable companies), were designed to allow competitors to access particular network functions (Cannon, 2003, p. 200; Steier, 1986, p. 1). The purpose of the ONA provisions, which dealt with network design and policy, was to prevent anti-competitive conduct based on the incumbents’ control of local networks, such as the imposition of unreasonable access charges imposed on private networks wishing to connect to the PSTN. ONA mandated that the carriers provide competitors offering enhanced service access to basic communications services on an equal basis and at an equal cost to those enjoyed by the carriers’ own enhanced service subsidiary operations (Noam, 2001, p. 177). The FCC did not, however, define standards for the ONA and the onus was on the carriers to define their compliance. Carrier’s initial plans were filed with the FCC by 1988 and were continually amended afterwards (Olsen & Tebbutt, 1995, p. 3). Additionally, there

was some conflict between the FCC inquiries and policies of the Antitrust Division of the US Department of Justice.

One goal of the 1996 Telecommunications Act was to open the local exchange market to competition with a continuation of the policy of ONA. Sections 251 and 252 of the Act imposed market-opening mechanisms, such as mandatory interconnection, unbundling and resale requirements on the ILECs (Noam, 2001, pp. 182-3). It was also expected that the Act would push local carriers into broadband technologies, with provisions for ensuring competition and deployment of advanced communications technologies. The problem was how to persuade the ILECs to give up their monopoly on local markets. The solution created by the 1996 Act was that the ILECs were permitted to enter the lucrative long distance telephone market only when their local markets were opened up to CLECs (competitive local exchange carriers).

Thus, by the end of the 1990s, the concept of network unbundling was firmly established in the United States. It was adopted by the WTO, and the European Union which provided for unbundled local loops by 2001. Though unbundling is a significant regulatory intervention, it had become a key tool for governments pursuing deregulatory policies. But it was also possible to anticipate its reduction in scope if various network elements became increasingly competitive. In time, one could expect the unbundled

elements to shrivel down to the last and most expensive to enter – the last part of call termination on the local loop beyond the switch, or even farther downstream to the sub-loop. (Noam, 2001, pp. 183-4)

Section 706 of the Telecommunications Act of 1996 was specifically directed at periodically examining if broadband enhanced services were being offered at a reasonable rate and efficiency to all citizens regardless of their location. This section of the Act inquires into whether “advanced telecommunications capability is being deployed to all Americans in a reasonable and timely fashion” (Wolf, 1999, p. 1). Since passage of the 1996 Telecommunications Act there have been five section 706 Inquiries conducted by the FCC (FCC, 1998b). In the first 706 inquiry (1998), the FCC asked in its Notice of Inquiry (FCC 98-187) whether it would have the authority to monitor peering arrangements and whether it would be in the public interest for the Commission to do so:

What can and should the Commission do to preserve efficient peering arrangements among Internet companies, especially in the face of consolidations of large proprietary gateways? We ask for comment whether the Commission should monitor or have authority over peering arrangements to assure that the public interest is served. (FCC, 1998b, para 79)

Further, this first Section 706 Inquiry asked:

Is it unrealistic to expect companies, many of whom have possessed and exercised market powers for decades, to behave like the non-network part of the Internet industry? Will interconnection occur, naturally or by operation of the antitrust laws, among advanced networks...? (FCC, 1998b, par 82)

Additionally the Commission asked “what, if any, system of regulation might best fit the market for advanced telecommunications capability” (FCC, 1998b, para 80). Many intervenors responded, including America Online, Northern Telecom, PSINet and SBC, all opposing federal involvement in peering and transit. One incumbent, Bell Atlantic, suggested that the regulator lower barriers for new entrants, in particular currently precluded entrants (Cybertelecom, 2007). In its second 706 Notice of Inquiry (2000), the FCC did not address the issue of Internet backbone peering. This did not prevent the ILECs from warning in their responses that there was an impending bandwidth shortage in the Internet backbone – in spite of clear evidence that, with the beginning of the “dotcom meltdown,” there was a growing bandwidth glut at the backbone level. Subsequent Section 706 Reports (Third, Fourth and Fifth Reports) have made only brief reference to Internet backbone providers and related regulatory issues.

Telecommunications Services vs Information Services

The regulatory distinction between basic and enhanced services, as created by the FCC, while intended to contain incumbent market power, would come to play an important role in enhancing such market power, especially in the local loop. One reason for this can be found in the provisions written by Congress into the 1996 Telecommunications Act, which, despite many changes in the marketplace in the intervening 15 years, essentially preserved the distinction between basic and enhanced services, albeit with a change of terminology: a “telecommunications service” became the functional equivalent of a “basic service,” while the newly coined “information service” became the functional equivalent of “enhanced service.”

The creation of this class of “information services” in 1996 was to have unintended consequences that, coinciding with the privatization and commercialization of the Internet, severely reduced competition and forced many ISPs out of the retail access market. Under provisions of the 1996 Act, competition was to be based largely on stimulating intramodal competition - competition between ILECs and CLECs, rather than on intermodal competition (competition between wireline telephone carriers and new entrants using other platforms, such as cable and wireless companies) (Noam, 2001, p. 54 - 59). This was a framework designed for the era of dialup access, in a time when subscribers could, by making the equivalent of a voice telephone call, connect

with any number of independent ISPs, unaffiliated with the incumbent in control of the last-mile bottleneck facilities.⁷ Unfortunately, what worked for dialup was entirely unsuited for broadband:

Although Congress took limited steps to ensure competitive parity between telephone and cable companies in the provision of voice and video services, the real competition between wireline and cable platforms arose in a market that hardly existed in 1996: the market for broadband Internet access. Because Congress did not foresee that cable and telephone companies would compete in this market, it did not set forth a clear regulatory framework for that market - let alone contemplate how to ensure regulatory parity between these competing platforms.... More generally, Congress left in place the arbitrarily compartmentalized regulation of the industry reflected in the multiple "Titles" of the [1934] Communications Act. (Nuechterlein & Weiser, 2007, p. 73)

The ultimate effect of some three decades of FCC orders, court rulings and lawmaking by Congress has been to place both high-speed cable and DSL service - by far the two most important residential broadband platforms - both out

⁷ Title II of the 1934 Act sets out provisions concerning common carriers, and the literature accordingly makes reference to "Title II obligations" in discussions touching on the regulation of common carriers. In many such discussions, classification as a Title II service is distinguished from Title I services, a reference to the general powers accorded the FCC in the 1934 legislation. A "Title II service" is thus used synonymously with "telecommunications service," while "Title I service" typically refers to an "information service."

of the reach of regulation and under the control of the largest ILECs and cable MSOs. The importance of how broadband services are classified lies in the complex set of rules governing interconnection - and in particular, whether a provider in offering a service is required to unbundle elements of the physical platform that service uses and lease them to competitors. By classifying cable and DSL broadband platforms as information services, rather than telecommunications services, such interconnection obligations are vacated (see Nuechterlein & Weiser, 2007, pp. 186-87). This is particularly relevant in the present paper insofar as interconnection obligations do not generally apply to bandwidth providers.

There are several schools of thought in the scholarly literature about the role of the FCC and Congress in creating a broadband marketplace that in most large American cities can be characterized as a duopoly. One of the most persistent criticisms of the Act is that it failed to take into account how much the market would be changed by the Internet and broadband technology. The FCC has been criticized even more sharply for taking an active role in classifying both high-speed cable and DSL as information services (Noam, 2001, p. 230). These regulatory actions have also been linked with the issue of Net Neutrality, which has become increasingly prominent over the last three years:

From a regulatory perspective, it is worth noting that Neutrality principles were never enshrined in law or regulatory practice. However, the Federal Communications Commission did rule that the underlying transmission components which were required for narrowband Internet service provider service were a “telecommunications service” and so subject to regulation and had to be made available to all on a non-discriminatory basis. In the broadband era, the Federal Communications Commission faced the question of how to classify the new cable modem and Digital Subscriber line services. The Federal Communications Commission in its Cable Modem Order [FCC 2002] declared broadband service over cable to be an “information service” and thus exempt from the telecommunications regulations imposed in the 1996 Telecommunications Act requiring local end-user level competition. In the *Brand X Internet Servs. v. FCC* decision, the Supreme Court upheld the Federal Communications Commission’s order in 2005 and the FCC swiftly followed with a ruling that Digital Subscriber line service was now also an “information service”. The US entered a new regime in which no one knows what the legal limits of discriminating practices are, thus priming the Net Neutrality debate. (Lehr, Peha & Wilkie, 2007, p. 709)

Much of this debate was centred on the degree to which certain aspects of the telecommunications market should or should not have been deregulated, and

whether the scope of common carriage was unduly narrowed by Congress and the FCC. It is worth noting, however, that arguments about deregulation may make unwarranted assumptions - as is the case with the 1996 Telecommunications Act, passed by the US Congress to expand competition from the long-distance arena to the local loop. The 1996 Act appeared to have been written to accomplish these goals by deregulating the telecommunications industry. As Nuechterlein and Weiser argue, however, the Act not only failed in its attempt to achieve robust competition, it also created new and even more complex levels of regulation:

[T]he 1996 Act advertised itself as the “most deregulatory [law] in history”. Anyone who believed that characterization at the time, however, was quickly disillusioned. The Act is not at all deregulatory in the straightforward sense of “tending to abolish regulation”. To the contrary, it adds an entirely new *dimension* to pre-1996 regulation by creating a broad new set of wholesale rules ... to the existing edifice of retail regulation. (Nuechterlein & Weiser, 2007, p. 407; emphasis original)

Following this logic, we would argue that the expansion of statutory provisions covering transactions at the wholesale level gave an advantage to the ILECs in their battle to use control of physical-layer facilities to gain a competitive advantage in services offered at the applications layer. That advantage flows not

merely from greater financial resources, but also from the much greater degree of comfort Bellheads have with regulation and government intervention. This is how Rob Frieden describes it:

As much as they might disparage regulation, Bellheads actually have benefited far more than they have suffered. Government served as guarantor of a stable revenue flow, even when insisting on long depreciation schedules and prescribing rates of return. Regulation pervades the Bellhead mindset as a necessary evil, but also as a mutually beneficial mechanism for both regulator and regulatee. Regulation offers a check against some of the most perverse marketplace forces making it possible for Bellheads to live in a safe, cautious and unremarkable environment. (Frieden, 2001, pp. 5-6)

In other words, incumbent carriers in control of the physical telecommunications infrastructure are likely to have competitive advantages, especially over new entrants, regardless of the degree to which any particular service offerings are subject to regulation.

In the wake of the 1996 Act, the confluence of several developments changed the telecommunications marketplace in ways that had a major influence on the structure of the Internet and the provision of bandwidth. These were: the lifting of

line-of-business restrictions; concentration of ownership among the Baby Bells; and the integration of retail and Tier-1 operations.

Under the 1984 divestiture, AT&T was confined to providing long distance service, while the RBOCs, which provided local service in their operating territories, were explicitly forbidden from getting into the long distance market until they proved to the court that they had opened their local territory to full competition. The 1996 Act went further in the direction of lifting line-of-business restrictions by allowing both long distance providers and the RBOCs to offer service bundles that included both local and long distance services. The dramatic shift precipitated by the new rules was accompanied by another, equally dramatic change in the landscape: over the next decade, the seven original RBOCs and AT&T were reduced to just three regional telephone companies through a series of mergers, acquisitions and name changes - most of which had occurred by 2000.⁸

A milestone was reached in late 1999, when Bell Atlantic became the first Baby Bell to be granted regulatory approval for offering long distance service in its operating territory. Bell Atlantic's modern history encapsulates two of the major

⁸ The seven RBOCs were Ameritech, acquired by SBC in 1999; Bell Atlantic, which acquired GTE in 2000 and changed its name to Verizon; BellSouth, acquired by AT&T Inc. in 2006; NYNEX, acquired by Bell Atlantic in 1996; Pacific Telesis, acquired by SBC in 1997; Southwestern Bell, which changed its name to SBC in 1995, then acquired AT&T Corp. in 2005 and changed its name to AT&T Inc.; and US West, acquired by Qwest in 2000.

trends of the era: consolidation and integration of retail and wholesale ISP businesses. Bell Atlantic was the first of the Baby Bells to acquire one of its sister companies - NYNEX, in 1996-97. It then merged in 2000 with GTE, the largest of the independent telephone companies (GTE held a controlling interest in BC TEL from 1955 to 2004) to become Verizon Communications. Five years later, in 2005, Verizon announced it was going to acquire long distance company MCI. The merger closed on January 6, 2006, in the wake of which the company created Verizon Business, a Tier-1 Internet bandwidth provider. Verizon was only one of several US carriers during this period that sought to expand from the long distance business to become Internet backbone providers. In 1998 alone, this group included Qwest, Level 3, ICG, Metromedia, Global Crossing and Williams (Noam, 2001, p. 38).

With the expansion of US influence across the Internet bandwidth business, and the burgeoning market power of American ISPs, concerns began to grow in the late 1990s about the extent to which US domestic regulatory issues may affect the welfare of Internet users in other countries. The United States and other interested parties have objected, arguing that global Internet interconnection markets should not be subject to traditional common carrier measures - and that these markets fall outside the scope of WTO (World Trade Organization) and ITU rules on dominant carriers. These issues have been addressed instead through bilateral antitrust agreements concluded between the US and the EU - e.g., the

1995 Agreement between the European Communities and the Government of the United States of America regarding the Application of their Competition Laws (Ungerer, 2000, p. 22 note 85).

By 1998, authorities such as Hal Varian were already pointing to the “balkanization” of the Internet, and to the close connection between technological advances, commercialization and a regulatory framework unsuited to the intermingling of telephone carriers and Internet resources:

The Internet already has begun to disaggregate into a hierarchy of networks based on available bandwidth, financial resources, number of Points of Presence and subscribership. This balkanization means that not all ISPs will have direct and seamless interconnection with all other ISPs, primarily because commercial interests favor disconnection of lesser ISPs unless and until they agree to one-way transfer payments upstream to larger ISPs. Market pressures have pushed the Internet toward balkanization and so far no legislative or regulatory edict has required interconnection like that imposed on common carriers. (Varian, 1998, p. A22)

Three years later, in 2001, Rob Frieden described a community of Internet bandwidth providers that had become thoroughly commercialized and hierarchical:

What had been a bill and keep arrangement with no transfer payment, became one where funds flowed from small ISP to larger ISP. The Internet became more hierarchical when peering remained an option only for the major Tier-1 ISPs at private peering points, with all other ISPs now bearing a financial obligation to compensate Tier-1 ISPs for the use of their networks. In essence a traditional financial settlement, metered access charge arrangement replaced the previous rough justice unmetered regime. The terms and conditions of such settlements, while untariffed and blocked from widespread scrutiny by nondisclosure agreements, appear quite like a telecommunications service arrangement.... Now Tier-1 ISPs operate at the top of a more hierarchical pyramid with a larger set of smaller ISPs operating lower in the pecking order. The Tier-1 ISPs have consolidated their control of the backbone facilities that constitute the Internet and in the process have converted former peers into customers of Tier-1 ISP facilities. (Frieden, 2001, pp. 11-12)

The FCC ruled in March 2002 that cable modem broadband service is an unregulated information service. As a result, without FCC regulation, cable companies were not required to share their broadband networks. In October 2003 the US Court of Appeals for the Ninth District ruled against the FCC, opening up cable modem networks and in March 2004 the court denied an FCC request for a rehearing of the case. However, the FCC and National Cable and Telecommunications Association (NCTA) were granted a stay of the court's decision pending a request for the Supreme Court to consider the case. In the 2005 "Brand X" decision, the Supreme Court upheld the original FCC order and the FCC swiftly followed with a ruling in August 2005 that digital subscriber line service was also now an information service – not a basic service.

Ad Hoc Oversight of the U.S. Backbone Marketplace

The Network Reliability Council was established by the FCC in 1992 following a series of major service outages in various local exchange and inter-exchange wireline telephone networks. The mandate of the Council was to study the causes of serious service outages and develop recommendations to reduce their number and effects on consumers. The Council was an advisory body comprising a forum of experts from the telecommunications industry, academe and consumer organizations, charged with developing measures to enhance network reliability. Since 1992, the Council has been convened under a series of changing mandates, reflected in its nomenclature: NRIC I through VII.

The Council's Charter was revised and its title changed to the *Network Reliability and Interoperability Council* by the FCC in April of 1996, after the passage of the Telecommunications Act. The Council was charged with advising the FCC on how section 256 of the Telecommunications Act – Coordination for Interconnectivity - should be implemented. Section 256 of the new Act required the FCC to establish procedures to oversee coordinated network planning by telecommunications carriers and other providers of telecommunications service. It permits the FCC to participate in the development of public network interconnectivity standards by appropriate industry standards-setting bodies. The charter for NRIC IV included assessing the impact of the year 2000 date change on networks and studying the current status of network reliability.

In January 2000, NRIC expanded the scope of its work to include Internet communications. In March 2000, FCC Chairman William Kennard and Commissioner Michael Powell announced that James Crowe, President and CEO of Level 3 Communications, Inc., would chair the next term of NRIC V. NRIC V was to provide the FCC with advice and recommendations on issues of reliability, interoperability and security arising in a multi-provider, multi-technology environment:

For the first time since the inception of the Network Reliability and Interoperability Council (NRIC or the Council), the FCC included in the

charter of the Council the mandate to address the unique issues arising from the interconnection of circuit-switched and packet-switched networks. (Network Reliability and Interoperability Council, 2002b, p. 2)

NRIC V met on a quarterly basis for two years, consulted with over 200 technical experts and established four focus groups: one, to continue work relating to the year 2000 (Y2K) date rollover on telecommunications networks; two, to evaluate the reliability of public telecommunications network services in the United States, including the reliability of packet-switched networks; three, to make recommendations concerning wireline-spectral compatibility and the development of spectrum management in wireline networks and facilitate the deployment of digital subscriber line and associated technologies; and four, to provide recommendations that when implemented would facilitate and ensure interoperability among public data networks.

Stopping short of recommending outright regulation, the Final Report of NRIC V Focus Group 4 produced two outputs: a short statement recommending that ISPs, especially the largest ISPs, consider, consistent with their business practices, publication of their criteria for peering (i.e. the terms and conditions under which they would peer with other networks for various types of traffic); and an informational paper discussing IP service provider interconnection, peering and transit service (Network Reliability and Interoperability Council, 2002c, d).

The informational paper detailed considerations that were taken into account in releasing the appended recommendation on publishing peering agreements:

In the United States, the decision to connect, how to connect, or to decline to connect, is driven by competitive market forces, rather than by government regulation. Because of the competitive nature of these arrangements, there is no legal obligation to disclose these decisions, terms, or to whom one connects. Decisions about which connection arrangement: peering, paid peering, or transit, or a hybrid arrangement, are determined by the competitive conditions of the market. Peering and transit are established pursuant to contracts between the parties. These contracts are usually treated as confidential business information. However, many would argue that the conditions under which providers are willing to enter into discussions regarding such contracts need not, and perhaps should not, be treated as confidential information. (Network Reliability and Interoperability Council, 2002d)

In an extraordinary letter, Focus Group IV wrote to Jim Crowe, NRIC V Chair, recommending that Internet backbones should publish their peering policies. While this was merely a recommendation, it is noteworthy that an FCC advisory body even considered the subject-matter as falling within its jurisdiction. This recommendation was not pursued and subsequent NRIC

broadband sub-groups have never revisited the topic. In July 2003, under a Memorandum of Understanding with the Department of Homeland Security, NRIC began to work in collaboration with the DHS (FCC Homeland Security Action Plan, 2003).

Over the last decade, merger proceedings have provided the FCC with further opportunities to address backbone-related issues, in particular those where approvals have been conditional on carriers divesting themselves of portions of their backbone facilities. This issue first came to wide attention with the MCI/WorldCom merger, to that date the largest merger in US history, valued at US\$37 billion. The merger proceedings began in 1997, and ended with MCI being required to divest its Internet backbone holdings. WorldCom was allowed to retain its backbone holdings, administered through its subsidiary UUNet. The MCI backbone facilities were sold off to Cable & Wireless (MCI, WorldCom merger gets green light from DoJ, 1998, p. 1). Prior to the closing of the merger at the end of the summer of 1998, both WorldCom and MCI disclosed their peering criteria in FCC public statements, but the merged company did not share such information freely once the regulatory proceeding was completed.

In 1999, following the WorldCom/MCI merger, the newly formed company sought approval to merge with Sprint, a deal which, if consummated, would have been valued at US\$129 billion and made the new company the largest

communications company in the US. In this case, however, the Department of Justice and the European Union Competition Bureau indicated they would block the merger because of the potentially anti-competitive impact of the merger on the Internet backbone market. In August 2000 the merger proceedings were suspended (Reuters, 2000, p. 1). The US Department of Justice made the following statement in connection with the proposed merger:

WorldCom's wholly owned subsidiary, UUNet, was the largest Tier-1 IBP by any relevant measure and is already approaching a dominant position in the Internet backbone market. Based upon a study conducted in February 2000, UUNet 's share of all Internet traffic sent to or received from the customers of the 15 largest Internet backbones in the United States was 37%, more than twice the share of Sprint, the next-largest Tier-1 IBP, which had a 16% share. These 15 backbones represent approximately 95% of all U. S. dedicated Internet access revenues. UUNet 's and Sprint's 53% combined share of Internet traffic is at least five times larger than that of the next-largest backbone provider. (US Department of Justice Complaint, 2000)

In reviewing MCI's merger with WorldCom, the FCC had noted that “the difficulties new entrants have encountered in connecting with [Internet backbone providers] are likely to continue after the merger. Therefore, we conclude that

peering is likely to remain an issue that warrants monitoring” (FCC, 1998a, par. 155). The Bell Atlantic/GTE merger proceeding in 2000 was approved by the FCC on the condition that GTE spin off its Internet backbone service, which was later renamed Genuity. This divestiture was required in order for Bell Atlantic to remain in compliance with its Sec. 271 obligations. The merged Bell Atlantic and GTE later became Verizon (Labaton, 2000, p. C1).

In 2006, as part of both the AT&T/BellSouth and Verizon/MCI mergers, the firms pledged to maintain at least as many open-peering, settlement-free arrangements as existed on the day of the merger closing, for a period of three years thereafter (Cybertelecom, 2006). They were also required to post their peering policies on publicly accessible Web sites for a period of two years. During this time the applicants were to post any revisions to their peering policies on a timely basis, as they occurred (Wigfield, 2005).

We turn now, in Chapter 3, to a close examination of the international bodies involved in various aspects of Internet governance - and the lessons they may provide for the issues of interconnection and data reachability.

Chapter III. Internet Governance Before and After Commercialization

The goal of this chapter is to examine some of the significant events in the history of Internet governance, beginning with the implementation of ARPANET, and to describe and assess those initiatives that have a bearing on the governance of ISP relationships. As explained below, peer-oriented and consensus-based oversight characterized technical innovation and standards-setting during the Internet's first 25 years. Since the mid 1990s, however, certain crucial facets of Internet governance have been compromised by commercialization, as well as by conflicts on the world stage over the role of the United States in governance.

One of the fundamental problems in the analysis of Internet governance stems from widespread disagreement over its definition and scope among scholars, activists, incumbent officials and other stakeholders. Jeanette Hofmann of the Internet Governance Project addressed this problem in her 2007 paper "Internet Governance: A Regulative Idea in Flux":

Although the term Internet governance has been in use about ten years now, there is not yet general consensus on its meaning. First of all, the concept 'governance' and its relationship to government is unclear;

second, it is unclear what extent and form of authority Internet governance has and in future should have. The concept became a popular catchphrase around the mid-1990s (Kleinwächter 2004a). According to William J. Drake, Internet governance proved to be a “heavily contested concept...from the very moment it entered into our collective lexicon”. (Hofmann, 2007, p. 1)

In order to better understand the complexities of governance in this context, we divide the Internet’s history into two periods, the first beginning in 1969 with the initial implementation of the ARPANET and ending in the mid-1990s, with the second period extending to the present day. The events dividing these two periods are central to our analysis, coinciding with the timeline noted by Hofmann concerning the use of Internet governance as a “popular catchphrase” from the mid-1990s onwards. These events mark a profound shift in the scale, scope and global significance of the Internet, and especially in its treatment by regulatory and standards setting bodies.⁹

There was no single set of events in the mid-1990s that formed a natural division of Internet history into two distinct phases. For ease of exposition, we can

⁹ One useful metric of Internet growth is the increase over the last 40 years in the number of host computers on the Internet (“host” refers to a machine with a registered IP address. Statistics from Hobbes Internet Timeline: 1969: 4 hosts; 1979: 188; 1989: 80,000; 1999: 56,218,000; 2006: 439,286,364 (Zakon, 2006).

summarize the changes most relevant to governance in terms of three distinct trends: commercialization, privatization and mainstreaming.

Commercialization. This concerns the *uses* to which early Internet facilities were put. Academic goals and content gave way gradually to commercial goals and content. This change was already underway in the early 1990s when the National Science Foundation relaxed its AUP (acceptable use policy) and allowed certain forms of commercial activity on the networks to which it controlled access. Commercialization was taking place while NSFNET was still the primary national backbone provider.

Privatization. This concerns the *ownership* and *control* of Internet facilities. Public ownership gave way to private ownership. Privatization took place at the same time as commercialization. Infrastructure resources developed with public funds were privatized, in large part (though not exclusively) because of the shutdown of NSFNET in 1995, as private firms assumed control of backbone facilities and peering points.

Mainstreaming. Use and control by a limited professional community of interest gave way to mainstream use by the general public. Again, many factors are implicated, but the single most important development was undoubtedly the introduction of the graphical Web browser in 1993, which provided easy access

to many useful applications for consumers with little or no technical expertise (although this crucial tool did not come to wide public attention until after the Netscape IPO in 1995). Today, 1.6 billion people are considered to be Internet “users”. (Internet World Stats, 2009)

Development of the Technical Governing Bodies

The pre-commercial period of Internet development (i.e. up to about 1995) was noteworthy for what Russell describes as “innovative” governance, based on peer review, open standards, consensus-based decision-making and a willingness to implement technical solutions before exhaustive testing (see above, chapter 2).¹⁰ This governance style is still very much part of the Internet’s technical governance, forming part of the philosophy of the Internet Society (ISOC), which was created in 1992 by the US authorities to handle most Internet-related technical oversight functions. ISOC is the organizational home of the Internet Engineering Task Force (IETF), the Internet Architecture Board (IAB), the Internet Engineering Steering Group (IESG) and the Internet Research Task Force (IRTF). The IETF is an all-volunteer group of experts that sets the underlying technical standards for the Internet. It describes itself as a loosely self-organized group of people who make technical and other contributions to the engineering and evolution of the Internet and its technologies. Membership in

¹⁰ This style of governance was captured in the oft-quoted phrase coined by Internet pioneer David D. Clark as part of a 1992 presentation to the IETF: “We reject: kings, presidents and voting. We believe in: rough consensus and running code.” See Russell, 2006, “Rough consensus and running code and the Internet-OSI standards war,” p. 1.

IETF working groups is open to anyone who chooses to participate. Its mission is described as follows:

The Internet Society (ISOC) is a nonprofit organization founded in 1992 to provide leadership in Internet related standards, education, and policy. With offices in Washington D.C., USA, and Geneva, Switzerland, it is dedicated to ensuring the open development, evolution and use of the Internet for the benefit of people throughout the world.

The Internet Society provides leadership in addressing issues that confront the future of the Internet, and is the organizational home for the groups responsible for Internet infrastructure standards, including the Internet Engineering Task Force (IETF) and the Internet Architecture Board (IAB).
(Internet Society, 2009)

Despite the common origins of the various US research and oversight bodies, and their shared predilection for peer-oriented governance, the creation of ISOC in 1992 marked the beginning of a sharp division of responsibilities into two streams. The first of these concerns the technical coordination of routing and internetworking standards and protocols, and the various other matters for which ISOC is responsible through its constituent bodies, such as the IETF. ISOC's creation did not, of course, signify that the Internet research community was

looking at design, policy and management issues for the first time, far from it. ISOC's primary mission since 1992 has been to provide a more formal, legally constituted home for the informal working groups whose antecedents stretch back to the initial implementation of ARPANET.

The second stream of technical issues that needed continuing oversight relates to Internet addressing - the registry or directory function allowing unique alphabetical addresses to be resolved as numeric Internet protocol (IP) addresses, operating as part of the Domain Name System (DNS). The DNS was created in 1983, so that all device addresses on the Internet could be recorded authoritatively in one place, using a structure able to scale to global proportions. In 1988, Washington created the Internet Assigned Numbers Authority (IANA) to look after the DNS. It was more a loose collection of functions than an actual organization - and mostly the responsibility of the highly admired Jon Postel, who worked alongside Vint Cerf and other Internet pioneers from the early days of ARPANET's development. As explained below, the informal nature of IANA led to a power struggle that resulted in the creation in 1998 of ICANN, the Internet Corporation for Assigned Names and Numbers.

While the ISOC engineering bodies have continued to operate with an open, peer-oriented approach to standards-setting, and do so on an international scale free of direct government control, ICANN has operated under the scrutiny - and

veto power - of the US Dept of Commerce, despite ICANN's mandate to represent the interests of Internet stakeholders from around the globe. The debates surrounding ICANN's role have been particularly strident over the course of 2009, because the Joint Project Agreement (JPA) entered into by the US Dept of Commerce with ICANN was set to expire in September 2009.

ARPANET and After: Managing the Internet

The experimental work that led to the ARPANET arose during the 1960s from a vision of interactive computing shared by a small group of researchers that included Robert Taylor, J. C. R. Licklider and Leonard Kleinrock, whose work was funded by the US Department of Defense through the Information Processing Techniques Office (IPTO) of its Advanced Research Projects Agency (ARPA). Although their work drew heavily on the research conducted by Paul Baran and others on packet switching, it was not directly concerned with achieving redundancy in a distributed network (a design concept aimed at overcoming the problem posed by having single points of failure in a communications network).¹¹

Instead, the impetus for creating the world's first operational packet-switched network was to foster a sense of community among university-based computer

¹¹ The original ARPANET was built around network nodes in four different locations: UCLA, Stanford Research Institute, UC Santa Barbara and the University of Utah. The first permanent link was established on November 21, 1969, while the full four-node network was completed on December 5 of that year (Zittrain, 2006, pp. 1975, 1989).

scientists and allow them to share resources that were scattered across the country, in an economical fashion. Those involved were interested in the challenge posed by trying to communicate across network links that relied on many different operating systems, software languages and hardware devices. By Robert Taylor's account, the immediate project trigger was his own sense of frustration at having three computer terminals in his office, each connected to networks that couldn't communicate with one other:

We had in my office three terminals to three different programs that ARPA was supporting.... For each of these three terminals, I had three different sets of user commands. So if I was talking online with someone at S.D.C. and I wanted to talk to someone I knew at Berkeley or M.I.T. about this, I had to get up from the S.D.C. terminal, go over and log into the other terminal and get in touch with them.

I said, oh, man, it's obvious what to do: If you have these three terminals, there ought to be one terminal that goes anywhere you want to go where you have interactive computing. That idea is the ARPAnet. (Markoff, 1999, p. 53)

This observation about the goal of sharing computing resources for a wide range of research programs has been made by others involved in the early stages of the ARPANET, such as Charles Herzfeld, ARPA director from 1965 to 1967:

The ARPANET was not started to create a Command and Control System that would survive a nuclear attack, as many now claim. To build such a system was clearly a major military need, but it was not ARPA's mission to do this; in fact, we would have been severely criticized had we tried. Rather, the ARPANET came out of our frustration that there were only a limited number of large, powerful research computers in the country, and that many research investigators who should have access to them were geographically separated from them. (Bellis, 2009)

What is clear from these first-hand accounts is that the early Internet developed as part of a free, open exchange of resources and ideas among a very small peer group of researchers. ARPA-supported scientists were funded to do basic research, and the mandate of this particular group was to enhance and expand time-sharing resources on large, expensive university-based computers. Nevertheless, despite the collegial nature of this research, the ARPANET collaborators turned their attention to the need for effective coordination and management at a very early stage.

In April of 1969, several months before completion of the initial network build-out, the ARPANET team launched what would arguably become the single most effective and important tool for technical coordination of the Internet: the Request for Comments (RFC) series. RFCs are the documents used by the engineering community in developing standards for the Internet. They had their modest beginning on April 7, 1969 with the release of RFC 1, a typewritten document prepared by ARPANET researcher Steve Crocker of UCLA and entitled "Host Software." Today numbering well over 5,000, they are the principal means whereby the IETF, created in 1986, establishes draft, then final specifications for Internet protocols and standards. The RFCs operate very much like peer-reviewed academic journals, in that they are judged on their merits by fellow professionals, according to established criteria (not all RFCs are intended to lead to a proposed standard). Although the RFCs have become more formal in style and must pass muster with the RFC Editorial Board, they have retained much of the pragmatic and democratic ethos of the early ARPANET (Crocker, 2009).

There is no better illustration of the common roots of the two principal "streams" of Internet governance than the work of computer scientist and Internet pioneer Jon Postel. Postel worked in a full-time position as Director of the Computer Networks Division at the USC Information Sciences Institute, which had been involved in the development of ARPANET from its inception. In 1988, Postel was named director of the newly formed Internet Assigned Numbers Authority (IANA),

although this appointment merely formalized the role he had been playing in development of the Internet namespace, or address system, since the early 1970s. In addition to running IANA, Postel was also the RFC Editor until 1998, the year of his death, and in that position wrote or co-authored over 200 RFCs, including RFCs 791-793, which defined the basic protocols of the Internet protocol suite. He was deeply involved with Internet governance in still other respects, including membership in ISOC and the Internet Architecture Board, as well as its predecessors.

Despite the prodigious growth of the Internet and Internet-related research during the late 1980s and early 1990s, the work of pioneers like Postel, Cerf and others was largely ignored by government regulators and telecommunications standards development organizations (SDOs), as well as by commercial interests - at least until the World Wide Web provided a mainstream platform for e-commerce and mass media content (and private firms began to provide ISP resources in place of publicly funded backbones and peering points). This lack of attention can be ascribed to several factors. One of these would certainly be the highly specialized nature of the research activities on which American computer scientists and engineers had focussed their professional attention.

But it was not just the nature of their work that insulated the Internet research community from outside scrutiny and allowed them to make such great strides.

The growth of this community and its successes were a function of *how* they communicated, not just *what* they communicated. In its early incarnations, the Internet was designed to enable the sharing of resources among co-equals. The Internet's early end-users were not only peers in a professional sense but also part of a peer-to-peer system in the technical sense, working on a system eminently suited for distributed, packet-switched networks, and for resource sharing within a research community.

The National Science Foundation and NSFNET

The National Science Foundation became involved in Internet-related research as early as 1979, when NSF representatives met with ARPA scientists and university researchers to develop a data communications network for use by computer science departments across the country that had no access to ARPANET. Then in 1983, the Department of Defense split off parts of ARPANET to form MILNET, at which time the NSF assumed most of the financial responsibility for ARPANET.

The NSF became even more directly involved in supporting Internet research and operations in 1986, when it agreed to fund NSFNET, which became the primary national Internet backbone (and four years later, in 1990, it replaced ARPANET altogether). Although NSFNET's initial mandate and capacity were quite limited, it would eventually act as a catalyst in the transformation of the

Internet into a global, commercially-based infrastructure. NSFNET, which operated for nearly a decade (1986 to 1995), underwent changes on several dimensions: purpose, acceptable uses, size, transmission capacity and architecture. In its position as the primary Internet backbone, NSFNET effectively governed users of the network and determined how they were expected to behave.

In the year prior to the deployment of NSFNET, the NSF funded the creation of five (later six) national supercomputer centers, which were then linked via NSFNET to promote resource sharing for academic research. In keeping with this academic focus, numerous university campus networks were connected to the NSFNET backbone through a series of regional networks (the network topology was thus based on a simple three-layer hierarchy). Interconnection was by no means confined to American networks. The first international connection to ARPANET, with University College, London, dated back to 1973. Similarly, NSFNET encouraged IP networks in other countries to interconnect, including several of Canada's regional IP networks in 1988. As the number of interconnecting networks and amount of traffic increased, NSFNET underwent two major upgrades from its original 56 Kbit/s channel capacity (equivalent to a residential dialup modem): it was increased to 1.5 Mbit/s (T1 capacity) in the summer of 1988 and again to 45 Mbit/s (T3 capacity) in 1991.

NSFNET was in an unusual position as a publicly funded operation whose mission ostensibly precluded commercial involvement. For one thing, it was obliged to carry out the work of maintaining and operating NSFNET with the help of private-sector partners. It had to make arrangements with AT&T to lease the long-haul lines it used to carry data, as it did not own the necessary physical infrastructure. In 1987, the NSF contracted with a non-profit consortium - Merit Network, Inc. - to build and maintain an upgraded network, completed the following year. The consortium participants included the State of Michigan and several universities, as well as MCI and IBM, whose role in this project would later give them an advantage in pursuing Internet-related business as privatization took hold.

NSFNET was also in an unusual position regarding its AUP (acceptable use policy), a short, relatively simple document that enshrined the network rules of governance. Between 1988 and mid-1990, the US government required that all NSFNET backbone traffic be part of initiatives whose chief purpose was to promote scientific research and other scholarly activities. The AUP's preamble ("General Principle") read as follows:

NSFNET Backbone services are provided to support open research and education in and among US research and instructional institutions, plus research arms of for-profit firms when engaged in open scholarly

communication and research. Use for other purposes is not acceptable.

(NSFNET backbone services Acceptable Use Policy, 1992)

Commercial activity was not banned outright; it was a matter of balance and priorities. The only two explicitly unacceptable uses were “[u]se for for-profit activities, unless covered by the General Principle or as a specifically acceptable use;” and “[e]xtensive use for private or personal business.” Networks and institutions that could demonstrate compliance with the AUP were granted interconnection rights to the backbone (Halabi, 1997, p. 8). This arrangement meant considerable economies to groups allowed access. Rather than paying for Internet connections, many eligible entities secured government grants that covered the costs of interconnection, as well as the development and maintenance of network systems. Until 1995, the NSF Connections Program provided access funding not only for American colleges and universities, but also for organizations such as libraries, museums and public health facilities.

The very year that NSFNET put its AUP into effect, a number of commercial activities got underway that signaled important long-term changes in the role of both public- and private-sector bodies. In 1988, the US Federal Networking Council approved the interconnection of the NSFNET to MCI Mail, the world’s first commercial email system. Other commercial electronic email services were soon connected, including Telemail and CompuServe. Email would prove to be

extremely popular and it expanded in conjunction with the rollout of ISP services aimed at the general public. Before 1988 was out, three commercial Internet service providers (ISPs) had sprung into existence: UUNET, PSINet and CERFNET.

As the networks comprising the Internet grew in number, size and volume of traffic, changes of a different kind were taking place in the way individual networks were interconnected with other networks across the United States. Whereas the NSFNET AUP was concerned with the nature and purpose of the packets being transmitted, interconnection policy was concerned with the physical and organizational arrangements made among individual networks for exchanging traffic. Certain provisions of the AUP left room for interpretation, but the document did make NSFNET policy on network uses explicit. Interconnection policy, on the other hand, did not receive the same attention and was left to the market to decide.

The privatization of Internet interconnection facilities was a gradual process that unfolded from inside the publicly-funded research establishment, as well as among private firms anxious to take advantage of new opportunities. The particular issue at stake for many network providers was whether they were going to be able to exchange traffic with other networks on a settlement-free basis and avoid paying for transit. Two events in 1990 raised concerns about

how the interconnection market would function once the NSF withdrew its support.

The first of these events, a workshop on commercialization and privatization, took place in March of 1990 and was co-sponsored by the NSF. A few months later, the organizations that won the 1987 bid to upgrade NSFNET to T1 capacity (IBM, MCI and Merit Network) created a new firm, incorporated as Advanced Network and Services (ANS). The following year, 1991, ANS was given the task of building and operating the next incarnation of NSFNET, now upgraded to T3 capacity and connecting 3,500 networks. NSF's relationship with IBM and MCI created concern among other private-sector firms about which ones would later be in a position to replace NSFNET (IBM and MCI each gave \$4 million to the NSF while acting as contractors). And indeed, when the NSF discontinued funding for NSFNET in 1995, it transferred its network operations to ANS.

A New Architecture: Network Access Points

The challenge facing the stakeholders through the whole transitional period from 1990 to 1995 was not simply a matter of transferring control of the Internet backbone to private sector firms. Two changes took place in the architecture of the Internet that made the transition more complicated. First, other providers began to offer backbone services as an alternative to NSFNET. Second, the explosive growth of IP networks and data traffic created the need for a new access architecture. Under the original three-tier hierarchy, campus networks

linked to regional networks and regional networks in turn linked to the NSFNET backbone. This architecture was soon unable to handle the demands being made on it. Moreover, many networks carrying commercial traffic were unable to gain access to the Internet as a whole because of the NSF's AUP. The solution was to create network access points, i.e. switching facilities that allowed multiple backbone providers to exchange traffic. Thus began many years of wrangling over the design and control of access points, compounded by the conflict developing over whether interconnection would continue to be free, and if so, on what basis.

As early as 1989, the federal government foresaw the need for exchange facilities because of the impending closure of ARPANET and the expanded role for the NSFNET backbone. To ease this transition, two publicly funded facilities were built in that year: the Federal Internet Exchanges located on the East and West Coasts respectively (FIX East and FIX West). This initiative was far from a lasting or comprehensive solution. By 1991, certain private-sector network operators were not merely concerned about NSFNET's non-commercial policy. They were also concerned that, after assuming full control of the primary Internet backbone, ANS would begin charging other network operators for access. In other words, ANS would require settlement-based interconnections to its Internet backbone (Butler, 2000, p.19). One early response to this perceived threat to the status quo was the creation of CIX, the Commercial Internet Exchange.

The Commercial Internet Exchange began as a non-profit trade association to facilitate public commercial interconnection to the Internet. It constituted the first attempt to ensure Internet access for commercial networks on a non-commercial basis, i.e. its members were required to exchange traffic on a settlement-free basis (Halabi, 1997, p. 12). CIX was formed by General Atomics (CERFNET), Performance Systems International (PSINet) and UUNET Technologies (AlterNet), after the NSF lifted restrictions on commercial uses of the Internet in early 1991. The CIX alliance was a reaction to the uncertainties surrounding the status of commercial traffic on the Internet. In 1992, as if to underline these uncertainties, the US House Subcommittee on Science, Research, and Technology, chaired by Representative Rick Boucher, asked the Office of the Inspector General to conduct an investigation into whether the NSF was authorized to relax its network AUP. An investigation concluded that the NSF did in fact have the power to relax the AUP, a finding reflected in an amendment to the NSF Authorization Act of 1993.

Under the CIX model, members were required to offer settlement-free, multilateral peering, in other words, agree to a peering relationship with all other CIX members. CIX attracted a considerable number of participants, many of whom had been ineligible for connection to the Internet because they carried commercial traffic. In the meantime, however, CIX became embroiled in a dispute with ANS, which had initially declined to participate in CIX, insisting it was

entitled to provide access on a transit rather than settlement-free basis. Although the ANS transit model was highly controversial, the idea of offering numerous access points rather than having individual networks connect directly to the backbone had been gaining favor as the Internet grew by leaps and bounds. ANS did eventually back down and joined CIX on a “trial” basis, after agreeing to abide by the rule of settlement-free interconnection. This change of approach on the part of ANS came about after the Congressional hearings in 1992 that led to the NSF Authorization Act commercial uses amendment (Noam, 2001, p.64).

CIX was not the only alternative to NSFNET to spring up at this time. In 1992, Metropolitan Fiber Systems (MFS), a large provider of network services, built a network exchange facility in the Washington DC area, known as an MAE, or metropolitan area exchange (the acronym was also used to reference a family of high-speed, fiber-based switching technologies - metropolitan area Ethernet - deployed in such facilities).

The Washington-area MAE, known as MAE-East, attracted many commercial networks seeking interconnection while bypassing NSFNET. Similar facilities were also built in San Jose and Los Angeles (MAE-West), and eventually in Dallas, Texas (MAE-Central). In 1996, four years after the launch of MAE-East, MFS acquired UUNET, which had been one of the original founders of CIX. In the acquisition of MFS, UUNet was purchased a few months later by WorldCom

(which trademarked “MAE”), and went on to become the world’s largest Tier-1 backbone provider (Butler, 2000, p. 39). In a foreshadowing of issues that would be raised in coming years by concentration of ownership in the telecommunications market, WorldCom’s proposed merger with Sprint was blocked by the Dept of Justice in 1999 on the grounds that UUNet and Sprint’s Tier-1 operations would, if combined, account for over 50% of the US backbone market and thereby reduce competition in the sector to an unacceptable degree. This action by the DoJ was one of the few systematic attempts by an agency of the US government to “govern” the Tier-1 market. As the Complaint explained:

WorldCom has attained (primarily through a series of acquisitions) a commanding position in the ownership and operation of the ‘backbone’ networks that connect the thousands of smaller networks that constitute the Internet, and Sprint is WorldCom's largest competitor in that market.
(US Department of Justice Complaint, 2000)

Commercialization: The NSF Withdraws and Transit Begins

While this flurry of activity was underway in the private sector, the NSF continued to make preparations for its eventual withdrawal from the backbone market. In May 1993, the NSF issued a solicitation for proposals (NSF 93-52) to build and operate four commercial network access points, or NAPs. They were intended to achieve the same goal as the interconnection facilities offered by the Federal

Internet Exchanges, CIX and the MAEs: allow major networks to interconnect with one other as part of an Internet topology that now included many backbones other than NSFNET.

Contracts for the NAPs were awarded to four different telecommunications companies. As part of its strategy, the NSF provided transition funding for a period of one year to regional network operators that wished to connect either directly to the NAPs or to providers connected to the NAPs (Halabi, 1997, p. 20). Two of the winning firms were RBOCs (Regional Bell Operating Companies): Pacific Bell (at one time the largest of the AT&T operating companies), which was awarded the contract for the San Francisco NAP; and Ameritech, which was awarded the contract for the Chicago NAP, along with Bellcore, formerly Bell Communications Research (SBC acquired PacBell in 1997 and merged with Ameritech in 1998.) The contract for the New York NAP, physically located in Pennsauken, NJ, was awarded to a division of Sprint, which was at the time the only major competitor to the incumbent telecommunications companies, along with MCI. The contract for the Washington DC access point was awarded to Metropolitan Fiber Systems, which was already running the MAE-East exchange point - an award that would become especially contentious the following year (1996) after MFS acquired UUNET (Halabi, 1997, p. 9). UUNET, with its dominant position in the backbone market, like other large ISPs, now had both

the incentive and the opportunity to stop peering with other networks, in favor of paid transit.

In 1997, UUNET and Sprint announced the formal end to public peering with their backbone facilities. Up to this point, ISPs of all sizes had been able to connect their networks at public peering points where they were assured of 24/7 access to the global Internet. After the closure of NSFNET, most peering points were still freely accessible. By 1996, however, many smaller networks faced a dual problem. On one hand, they were being asked to pay for transit. On the other hand, tremendous increases in traffic volumes were putting a strain on interconnection facilities and causing congestion, with associated packet loss and increased latency. The idea of providing access points based on mandatory multilateral peering had been popular, but most large stakeholders were increasingly opting for private bilateral agreements (Blake, 1999, p.15). The major problem arising from the closure of NSFNET was not simply the lack of financial and technical support, but more importantly, the lack of interconnection guidelines for access facilities. As Butler points out, getting access to an interconnection point was not the same as getting access to other networks at that facility:

It was and still is relatively simple to get a connection to any NAP, but it is an entirely different and much more difficult thing to get other networks to

peer with you there. There was much debate on this subject just prior to and just after the dissolution [of NSFNET], but no agreements were ever reached. (Butler, 2000, p.34)

The push for private, bilateral agreements and more paid transit owed a great deal to the fact that the four NAPs funded by the NSF were owned and controlled by four of the largest American telecommunications carriers - two RBOCs (PacBell and Ameritech) and what would become the two largest American backbone providers, Sprint and MFS (acquired by WorldCom in 1996). As Eli Noam put it: "Such arrangements began in 1995, partly due to the growing congestion of the public exchange points, partly due to the desire to bypass these exchanges and partly in order to establish control" (Noam, 2001, p.66). The largest ISPs were able to exercise control not only because of their sheer size and volume of traffic, but also because they owned access points and often the circuits needed to gain access to them.

These developments singled the end of open peering as a universal practice. Because private peering was feasible for large incumbent providers and not for their smaller competitors, the backbone market by the late 1990s was very restricted and had high barriers to entry, factors which put smaller ISPs at a distinct disadvantage. As Butler observes (writing about the state of the market in 1999), private peering had become a "good old boys club":

The large providers have established private peering with the other large providers and have effectively shut out the smaller providers from the “good” bandwidth. The private interconnects that offer faster transit and less packet loss can certainly be considered “better” bandwidth than the bandwidth at the NAPs, where congestion results in latency and loss.

A smaller ISP could buy transit on discrete circuits from the larger ISPs rather than peering at the NAPs but would then have difficulty competing as a peer due to this added costs. This allows large ISPs to exclude new entrants and smaller players because private interconnects are certainly not offered to all comers and are not really based on any publicly disclosed criteria. Over time, it is quite possible that this private interconnect space will be seen as anti-competitive and ultimately collusive due to these factors (Butler, 2000, p. 42).

The Impact of Commercialization on IANA: Old Functions, New Problems

Earlier in this chapter, we discussed the “innovative” style of governance typical of Internet management in its early days and which has to a large degree been carried over in the work of the ISOC technical bodies, such as the IETF. We noted that technical research and standards setting was highly collegial and based on open, consensus-based decision-making - an ethos captured in David

Clark's phrase "We reject: kings, presidents and voting. We believe in: rough consensus and running code" (Clark, 1992).

The work of the IETF and its sister bodies was given a more formal status with the creation of ISOC in 1992. One important technical function that was not brought under the wing of ISOC at its creation was IANA, the Internet Assigned Numbers Authority. IANA was more a function than a formal organization (e.g. it did not have the authority to sign legally binding contracts). Its home was the Information Sciences Institute of USC, where much of the early research on ARPANET was carried out, under the direction of Jon Postel. Postel and IANA, like much of the establishment that sustained ARPANET and its successors, worked on contract with the Department of Defense, via DARPA.

Postel, working through IANA, looked after an increasingly complex series of Internet-related functions, in close coordination with other technical groups that were formed in the wake of ARPANET. The functions for which Postel was responsible were centered on the two Internet "namespaces" - the IP numbers and domain names that provide the Internet's globally unique addressing system. These responsibilities are closely linked to the work of the IETF in the administration and implementation of Internet protocols and standards. In several respects, therefore, the IANA functions were and still are inseparable from functions for which ISOC is responsible.

In other respects, however, IANA has to maintain complex international relationships with other bodies, both national and international, whose interests are not always aligned with those of IANA. The most important of these are the Regional Internet Registries (RIRs), to which IANA assigns large blocks of IP addresses for use in their respective regions of the globe. As it is responsible for the Domain Name System (DNS), formally constituted in 1988, IANA also works in close conjunction with national top level domain (TLD) administrators, as well as with a number of different organizations that operate the “root nameservers,” the set of servers that hold the authoritative record of all the top level domains on the Internet (these are of two kinds: the 248 country code top level domains or ccTLDs, and the 20 generic TLDs or gTLDs).

As we have seen, the early 1990s was a period in Internet history which saw the beginning of inroads by private firms, as well as the first signs of sanctioned commercial activities. In a series of events that closely paralleled the activities of the NSF in funding NSFNET and related Internet resources, the US government decided to outsource certain IANA functions to the private sector, primarily those associated with the registration of domain names, which was becoming unwieldy (in January 1992, there were 727,000 host computers on the Internet; by January 1995, there were 4,852,000). This transfer of responsibilities took place in two main phases. In 1991, the DoD outsourced registration services to a private contractor; shortly thereafter, the military passed responsibility for Internet-related

services to the NSF. By this point, as we have seen, the NSF was deeply involved in supporting development of the Internet through its relationships with private sector service providers. Thus, in 1993, the NSF contracted with three firms to run the Internet Network Information Center, known as InterNIC, to manage the allocation of Internet addresses.

The firm that would stand to benefit most from this arrangement - and become embroiled in controversy as a result - was Network Solutions, Inc. (NSI), which had already been contracted to provide domain name services for the military prior to winning the NSF contract (along with two other high-tech firms). One particular provision of the original contract, as noted in the related NSF press release dated January 5, 1993, sowed the seeds for what would become a tumultuous change in how Internet services were provided, and as a direct result, in the creation of ICANN:

Consistent with FNC guidelines on obtaining reasonable cost recovery from users of NREN networks, the NSF has determined that the INTERNIC Information Services provider may charge users beyond the U.S. research and education community for any services provided (Estrada 1993).

In 1996, one year after the NSF closed down NSFNET, it instructed NSI to begin charging for domain name registrations, a reflection of the continuing commercialization and growing size of the Internet. The immediate policy issue for the US government was that, by providing free registrations, the NSF was subsidizing business and other activities which were not within its purview.

Over the next two years, Internet governance went through a wrenching series of disputes. They revolved in essence around the clumsy attempts by Washington to privatize the handling of domain name registrations and other IANA functions; the controversial corporate behavior of NSI, which was handed a monopoly in what would become a very lucrative business as registrar for the .com, .net and .org TLDs; and the battle that engulfed Jon Postel and the Clinton White House in 1997 and continued through 1998, when Postel passed away.

In April of 1997, as a sign of the growing interest being taken by business in the World Wide Web, the number of Web sites passed the one million mark (having begun with the launch of the very first Web site by CERN in December 1990). Before long, one of the most pressing concerns, one that could scarcely have been imagined by the Internet pioneers even a decade earlier, was “cybersquatting” - registration of a domain name in alleged violation of trademark rights. This interest on the part of business, especially large multinational corporations, coincided with the shifts in international trading arrangements

brought about by the creation of the WTO in 1995 - one of the roots of globalization, which was symptomatic of the increasingly international nature of the Internet and the Web. In the meantime, the introduction of a payment system for domain name registration - in itself a relatively small part of the overall activity involved in management of the Internet - generated a backlash against not only NSI but the NSF as well (for levying what was found by the courts to be an illegal tax). NSI was also creating controversy of another kind through its clumsy attempts to “filter” names it felt were inappropriate, which led to conflicts over First Amendment claims and counter-claims.

Over the course of 1997, as these conflicts unfolded, a number of actors became involved that had not previously had a role or interest in Internet governance, in particular the DNS and management of other, now controversial IANA functions, several of which Postel was still looking after in his role as de facto “head” of IANA. One of these groups was the short-lived IAHC or International Ad Hoc Committee, which took up the task of developing a new framework for allocating and managing domain names. The organization produced an influential document known as the Generic Top Level Domain Memorandum of Understanding or gTLD-MoU, which was signed by over 200 organizations. While the IAHC itself was formally dissolved in May of 1997, it set an important precedent: it was the first broad-based organization without ties to either government or private sector institutions that played an activist role in promoting

solutions to Internet governance based on public interest considerations. Many such organizations would follow in its footsteps, especially after the creation of ICANN, the Internet Corporation for Assigned Names and Numbers, in September 1998.

ICANN: the Domain Name System and other Conflicts

The About page of the ICANN Web site provides the following introduction to this high-profile body, undoubtedly the best known and most controversial body with responsibility for Internet-related functions:

To reach another person on the Internet you have to type an address into your computer - a name or a number. That address has to be unique so computers know where to find each other. ICANN coordinates these unique identifiers across the world. Without that coordination we wouldn't have one global Internet. (ICANN, 2009)

ICANN was formed in 1998. It is a not-for-profit public-benefit corporation with participants from all over the world dedicated to keeping the Internet secure, stable and interoperable. It promotes competition and develops policy on the Internet's unique identifiers. ICANN doesn't control content on the Internet. It cannot stop spam and it doesn't deal with access to the Internet. But through its

coordination role of the Internet's naming system, it does have an important impact on the expansion and evolution of the Internet.

According to many scholars, ICANN is a troubled organization that has not lived up to widely held expectations of accountability, international representation and transparency. Whatever the basis for judging ICANN's recent achievements, it was born in circumstances rife with conflicting interests, at a time of great upheaval across the global Internet. In late 1997 and early 1998, Jon Postel became embroiled with the Clinton White House over technical and policy changes to IANA-related functions. This dispute, along with the problems that had developed around the role of NSI and the domain registration process, prompted the Clinton administration to issue what would become the basis for a new governance framework: the famous Green Paper issued by the US Dept of Commerce in January 1998. This document proposed the creation of a private, not-for-profit corporation to take over management of key Internet functions, notably those performed by IANA (National Telecommunications and Information Administration, 1998).

Under the heading The Need for Change, the discussion draft listed seven compelling factors behind the White House proposals, beginning with the "widespread dissatisfaction about the absence of competition in domain name registration." Other factors noted in the document included the growing conflict

between trademark holders and domain name holders, and calls for “a more formal and robust management structure.” This was an obvious reference to the informal style in which Jon Postel had managed IANA since its inception. For all that Postel was the object of great admiration among the scientific community within which he had worked so effectively, the surge of interest in e-commerce in the business community, along with the concerns of policymakers, meant that Postel was viewed with increasing suspicion and distrust.

In some respects, ICANN was supposed to operate like the other Internet technical bodies, operating on consensus and being internationally representative. Furthermore, it was intended to spur competition in the registration business. The paper also made it clear that the US government was, after a suitable transition period, “seeking to end its role in the domain name system.” Very little of this grand design would be implemented in the way it was originally conceived - especially in the matter of US dominance as exercised through ICANN. It did not bode well that ICANN was created in secret, under an arrangement with the Dept of Commerce, which has always had effective control over ICANN through the contract made between ICANN and the Department. In 2003 the Department of Commerce renewed ICANN’s 3-year contract; in August of 2006 the contract was renewed again until 2011, subject to annual review (McCarthy, 2006). At this writing, the contract - now called the Joint Project Agreement (JPA) - is due for its annual renewal in September, and many groups

have called for changes that would reduce or eliminate the dominant role played by the US government.

The theme of US dominance is a common one in the literature. In his 2004 study, Cukier puts the emphasis on US control of the Internet as exercised through the technical functions associated with allocating large address blocks:

Due to its role as the first country online because it invented the Internet, the US has special control over certain names and numbers. As alluded to earlier, the US has a dominant position in the IP address number space in the current version of Internet Protocol, version 4. (Cukier, 2004, p. 42)

Writing in 2005, Morgenstern quoted remarks made by Michael Gallagher, the then head of the NTIA (National Telecommunications and Information Administration), and the official responsible for ICANN under the Bush administration. They are reminiscent of other official statements that assert plainly the American intention to maintain ultimate control over ICANN and its functions:

[T]he United States will continue to work with ICANN.... However, ICANN is the “appropriate technical manager of the Internet DNS,” not the final word. The United States will “continue to provide oversight so that the

Internet Corporation for Assigned Names and Numbers maintains its focus and meets its core technical mission” ... While the agency recognizes that other governments have “public policy and sovereignty concerns” relating to the Internet and domain services, Gallagher said, those interests should be focused on the ccTLD (country code top level domains). He said that the United States is committed to working with the international community to address these concerns, bearing in mind the fundamental need to ensure stability and security of the Internet’s DNS. (Morgenstern, 2005)

The IANA Function and Autonomous System Numbers (ASNs)

The technical work of ICANN has been carried over as what is termed the “IANA function.” This comprises overall responsibility for global Internet protocol address space allocation; top-level domains; domain name system management; root server system management functions; and other related administrative functions that were originally performed by the US government. ICANN carries out these tasks through management of IANA, which still executes the actual technical work (Drake & Wilson, 2008, p. 44). ICANN deals with policy through three supporting organizations. These are the Generic Names Supporting Organization (GNSO); the Country Code Names Supporting Organization (ccNSO); and the Address Supporting Organization (ASO).

One of the most common observations made about ICANN's struggle to achieve consensus and stability is that it has long failed to draw a clear line between purely technical functions and broader policymaking, such as the proposed creation of new gTLDs. As Milton Mueller from the Internet Governance Project at Syracuse University has described it:

The "IANA function" involves coordination of the root zone of the domain name system with 24 hour-a-day/7 days-a-week coverage. It includes receiving requests for and making routine updates of the country code top level domain (ccTLD) technical and administrative contacts and name server information. This function also includes receiving delegation and redelegation requests, [and] investigating the circumstances pertinent to those requests.... It also involves overall responsibility for delegating allocated and unallocated Internet protocol version four (IPv4) and Internet protocol version six (IPv6) address space and Autonomous System number space. These functions are crucial to the operation of the Internet, and their delegation to the ICANN by the US government is what gives the ICANN all of its policy leverage over the Internet (Mueller, 2006).

ICANN's own mission statements, and the huge volume of commentary on ICANN, indicate clearly that its role in managing the DNS, and the issues attaching to the two principal namespace functions, IP addressing and domain

names, have dominated public and scholarly debate over Internet governance. The DNS is a critical piece of the Internet's infrastructure. Moreover, Internet and Web address names, and even numeric IP addresses, are familiar to end-users, even if few realize that the coordination of the DNS to ensure universal resolvability is the primary function of ICANN.

IANA allocates IP addresses and Autonomous System routing numbers (ASNs) to the five global Regional Internet Registries (RIRs). These are: for Asia, the Asia Pacific Network Information Center (APNIC); for North America, the American Registry for Internet Numbers (ARIN); for the EU and Middle East, the Réseaux IP Européens Network Coordination (RIPE NCC); for Africa, the African Network Information Coordination (AfrinIC); and for Latin America and the Caribbean, the Latin American and Caribbean Network Interconnection Coordination (LACNIC) (Medhi & Ramasamy, 2007, p. 301). Each of the five global Internet registries is responsible for allocating both IP address blocks and autonomous system numbers within their region.

What has received little or no attention in the literature, or public debates about ICANN, is its role in the assignment of autonomous system numbers, or ASNs to entities such as universities, corporate, governmental, or other enterprises. An AS is a collection of IP sub-networks and routers that presents a single border gateway routing policy to the global Internet. This fact goes to two very important

points in our argument. First, the debates surrounding Internet governance have been focussed almost exclusively on the DNS and IANA namespace functions, at the expense of understanding the role of ASNs and routing policy in the maintenance of a robust global Internet. Second, little is known about ICANN's policy positions on the role that might be played by ASN management in a new governance framework for ISP relationships. The ICANN Web site maintains a posting of all ASNs as assigned in blocks to the regional authorities. (Internet Assigned Numbers Authority, 2009).

Governance in a New Light

In the last several years, a number of proposals have been put forward to make ICANN more accountable or effective, by changing its structure, requiring it to share authority with other bodies or replacing it altogether. To date, it appears that no one proposal is likely to hold sway. This assessment is based on two factors: one, the deep reluctance of the US government to relinquish control; and two, the inability of other organizations to overcome the very difficulties that allegedly plague ICANN. The most far-reaching of these have come from the United Nations initiative known as the World Summit on the Information Society (WSIS).

In December 2001, the United Nations General Assembly passed Resolution 56/183, which endorsed a two-phase structure for the proposed WSIS, to be

organized by the ITU, International Telecommunications Union (Hofmann, 2007, p. 14). The first meeting took place in Geneva in December 2003; the second phase took place in Tunis in November 2005. One crucial outcome of the first phase was the creation of a United Nations Committee to first define “Internet governance” and then make recommendations as to which organization should have responsibility for implementation (Drake & Wilson, 2008, p. 3). At the WSIS in 2006, the focus on Internet governance was perceived by representatives from many nations as an issue that had come about because of serious problems in ICANN’s management of the DNS. There was widespread dissatisfaction with the fact that ICANN is headquartered in California, and governed under US law through the MoU with the US Department of Commerce, administered by the National Telecommunications and Information Administration (NTIA).

In the wake of the WSIS, the ITU decided to enter the fray with ideas of its own on how to adjust ICANN’s role and become more involved in management of the Internet. The rationale for giving the ITU some responsibility in this area rested on the observation that conventional communications traffic and Internet services were converging, to the point of becoming indistinguishable. Despite the sound logic behind this view, some commentators have expressed skepticism. Some (e.g. McCullagh, 2005) have taken the position that the IETF has done an adequate job to date in managing Internet standards, whereas the ITU would be inclined to impose a centralized, broadcast-type model on the Internet,

particularly with respect to content - an outcome that would not be in the best interests of consumers or producers of content. Others, such as Oram, concede that there is widespread dissatisfaction with ICANN, but view the prospect of ITU involvement in either technical or policy matters relating to the Internet as highly problematic:

While the Internet Corporation for Assigned Names and Numbers has bumbled many tasks and exceeded its authority on others, its leaders have a sense of the fragility of the Internet ecology. The International Telecommunications Union, in contrast, is tromping all over the grounds, just in the process of mapping it. (Oram, 2003, p. 1)

The shortcomings for which ICANN is blamed should not, in our view, be allowed to obscure the very real challenges it faces, along with other important stakeholders. For example, while there has been controversy about ICANN's creation of new (official) gTLDs, serious threats to the integrity of the Internet have developed because of the creation of alternative, unauthorized root systems. When such systems overlap the authoritative global DNS by using the unique root information while adding new pseudo-TLDs, there is potential for serious long-term harm. A topical example is the Chinese root system, which is not available in the IANA root. As Joe Baptista of the Public Root Consortium has

pointed out, 300 million people see TLDs that are simply unavailable in the IANA root and therefore to the rest of the Internet population (Baptista, 2009).

Whatever the outcome of situations such as the unauthorized Chinese root, nothing is likely to stop the perennial debate over American dominance of the Internet, particularly through ICANN. On June 26, 2009, news was released that entrepreneur, technologist and author Rod Beckstrom had been named ICANN's new Chief Executive Officer and President. In a New York Times article published shortly after this announcement, Beckstrom indicated that he sees American control of ICANN as both positive and necessary. Entitled *New Chief Defends U.S. Base for Agency That Manages Web*, Beckstrom calls ICANN the best guardian of a "single, unified, global Internet." Beckstrom added that he is also opposed to any attempt to "fragment" ICANN by creating an "international subsidiary" of the organization (Pfanner, 2009).

In this same article, Pfanner presents the other side of the debate, in the person of ICANN critic Viviane Reding, the European Union media and telecommunications commissioner:

[Reding] recently called for a severing of Icann's [*sic*] links with the U.S. government when the current agreement with the Commerce Department expires this autumn. Instead, she proposed the creation of a "G-12 for

Internet governance” to oversee an independent Icann [sic]. “In the long run, it is not defensible that the government department of only one country has oversight of an Internet function which is used by hundreds of millions of people in countries all over the world,” Ms. Reding said in May. (Pfanner, 2009)

An important outcome of the WSIS meetings in 2003 and 2006 was to create a greatly heightened awareness of the issues surrounding Internet governance, despite the fact that the participants were unable to reach any significant degree of consensus on the term “Internet governance” itself. There was general agreement on the notion that the Internet governance debate is best conceived as a multi-stakeholder forum encouraging trilateralism, i.e. the participation of Civil Society representatives alongside commercial interests and United Nations member states. The WSIS concluded with the Tunis Agenda document and turned over the Internet governance debate to a new body, the Internet Governance Forum (IGF), which met in November 2006 (WSIS, n.d.).

In an effort to restore meaningful debate, the Internet Governance Project (IGP), housed at Syracuse University, has brought together several scholars with a special interest in the IGF and related multi-stakeholder discussion forums. The IGP, whose members include Jeanette Hofmann, Milton Mueller, Lee McKnight and John Mathiason, published a paper in 2004 *Making Sense of Internet*

Governance: Defining Principles and Norms in a Policy Context, in which they proposed a simple yet compelling definition of Internet governance:

... collective action, by governments and/or the private sector operators of TCP/IP networks, to establish rules and procedures to enforce public policies and resolve disputes that involve multiple jurisdictions.

(Mathiason, McKnight, Mueller, 2004, p. 4)

Up to this point, efforts at refining the conceptual apparatus for analyzing governance issues have remained somewhat vague and abstract, with few concrete applications. This is especially true of international forums such as the IGF, where national interests often get in the way of substantive discussion, despite general agreement that the Internet is an important global resource that should be governed in an international setting. Nevertheless, some commentators have begun to shift their focus away from the long-standing concerns over the DNS, recognizing that a stable global Internet will likely depend coming years on other factors, including the traffic engineering practices of ISPs. One such commentator is Jonathan Zittrain, who sees international oversight of the DNS in a new light:

[T]he focus on the management of domain names among those participating in dialogues about Internet governance is ... unfortunate. Too

much scholarly effort has been devoted to the question of institutional governance of this small and shrinking aspect of the Internet landscape. (Zittrain, 2006, p. 1979)

One set of issues which has captured increasing attention in the last three or four years, especially in North America, concerns Net Neutrality, along with other issues related to end-user access and welfare, including the problem of the Digital Divide. In the perspective presented here, we argue that connectivity and reachability issues in the last mile deserve to be considered alongside connectivity and reachability issues at the other end of the bandwidth scale, i.e. among Tier-1 networks. In our view, it is now time to see all such bandwidth issues within the same frame of reference - and in doing so, shift our energies away from the political issues associated with ICANN.

This change of perspective in no way means that we should confine ourselves to the strictly national treatment of Net Neutrality, in line with the traditional regulatory approach. On the contrary, emerging technologies and business practices point in the opposite direction - to the need for global solutions to global problems. This notion has captured the attention of numerous commentators, such as Richard Collins, writing in *The International Journal of Communication*. In his 2007 paper, part of a collection focussed on Net Neutrality, Collins makes reference to the United Nations Cardoso Panel report, which promotes the idea

that while the substance of politics is fast globalizing, the process of politics is not. The report, and Collins, argue for the participation of civil society institutions in “global governance” and “global policy networks” (Collins, 2007, p. 15). Exactly which civil society organizations and institutions are necessary for an adequate discussion of stakeholder interest is being debated at the Internet Governance Forum. A similar sentiment is voiced by Goldsmith and Wu, with more explicit reference to the Internet:

The bordered Internet does not imply that ... global Internet rules have no place, any more than our bordered world implies that there is no place for international law. On the contrary, many aspects of the Net will be governed on a global scale. (Goldsmith & Wu, 2006, p.164)

The forces of globalization and the inability of national regulatory regimes to keep pace with new communications technologies are blurring the boundary lines between domestic and international policy issues such as digital rights management, privacy, surveillance and p2p filesharing, as well as end-user access rights and the use of traffic engineering practices by ISPs. International bodies such as the WTO, trade associations such as the RIAA and CRIA, and regulatory tribunals such as the FCC and CRTC now share jurisdiction with courts and other national agencies in the effort to resolve Internet-related issues such as unauthorized file sharing and copyright infringement. Moreover, the long-

term trends we described at the beginning of this chapter - commercialization and privatization - have created a hybrid form of private-sector “governance” that allows both retail and wholesale ISPs to act as gatekeepers in ways that may have distorted the bandwidth marketplace.

We end this chapter by noting that in 2000, with the all-important distinction between a telecommunications service and an information service having been enshrined in the 1996 Telecommunications Act, the FCC seemed to take a hands-off position on the backbone provider market. As explained above in chapter 2, the classification of Internet service as an information service meant it was exempt from traditional common carriage regulation. In September 2000, the FCC’s Office of Plans and Policy issued a working paper by Michael Kende, Director of Internet Policy Analysis, entitled *The Digital Handshake: Connecting Internet Backbones* (Kende, 2000). While not a reflection of official policy, the paper in question suggested that the FCC would continue to allow market forces to dictate interconnection agreements. One of the arguments offered in the paper for refraining from regulation was that both domestic and foreign carriers had established a presence in the US backbone market, and it would be unwise to attempt to impose regulations on transactions involving one or more foreign carriers. This rationale speaks directly to the vexing problem noted above, namely that the Internet market does not fit neatly into the jurisdictional categories once typical of the regulated national telephone monopolies.

The paper argued that the FCC should not require that Internet backbones provide interconnection, going so far as to suggest that the Commission should not even investigate peering and pricing practices in the backbone market. The paper did concede that a dominant Internet backbone might engage in anti-competitive activities, including charging excessive prices for interconnection, engaging in predatory pricing, or discriminating in the quality of interconnection offered to competitors. It nevertheless concluded that the five nationwide Internet backbones constituted a sufficiently competitive market, and, moreover, that any reluctance on the part of incumbents to peer with new entrants into the market was no cause for concern because the backbones competed with each other for transit arrangements. As a result, the backbones would not be able to charge extraordinary prices for interconnection. On the basis of these findings, the paper concluded that traditional international telecommunications cost-sharing settlement was not an appropriate model for the Internet backbone market.

Two weeks after the FCC paper was released, the ITU issued its Recommendations at the World Telecom Standardization Assembly in Montreal. It approved a Recommendation (D.50) on Internet cost-sharing which appeared to open the door to a governance model for Internet interconnection, intended to encourage operators to adopt symmetric peering agreements (Johnson, 2000, p. 5; Jensen, 2005, p. 3). In this forum, as elsewhere, developing countries wanted a regulatory approach, whereas Canada, Europe and the United States preferred

a pro-market approach. At the Montreal Assembly in 2000, ITU delegates from developing countries raised concerns over an international Internet divide and the heavy concentration of Internet backbones and content in the United States.

The anti-regulatory position taken by the world's developed countries at this time did not square with the concentration of ownership and incumbent market power that characterized both the American and international telecommunications markets. Tier-1 backbone providers deserted public exchanges because of the technical and financial benefits of private peering. Internet backbone providers also began to change their terms of traffic exchange with smaller networks in order to recover their infrastructure costs as quickly as possible. These developments helped the largest players consolidate their market positions, while closing off peering opportunities for smaller operators. And all of this activity in the backbone market was paralleled in the retail market structure, where the ILECs and largest cable MSOs were successfully leveraging their control of last-mile facilities.

In the next chapter, we examine the relationship between this hierarchical market structure and the growing issues associated with data reachability.

Chapter IV. Risks to Reachability

Peering, Interconnection and Market Concentration

In the previous chapter, we examined a development that took place in the United States in parallel with the privatization of the Internet backbone: the creation and spread of interconnection points whose purpose was to accommodate the rapidly growing number of networks forming part of the public Internet. Until the early 1990s, interconnection points were relatively few in number and accessible to all eligible networks on the basis of free peering. They were public and open, i.e. there was no commercial hierarchy that encouraged larger networks to dictate the terms of interconnection or to charge for it as transit rather than settlement-free peering.

The US government foresaw the need for more substantial interconnection facilities by 1989, because of the expanding role being played by the NSFNET backbone. Public funds were earmarked for two facilities: the federal Internet exchanges located on the East and West coasts respectively (FIX East and FIX West). Within two years, however, several network operators had become concerned that once the NSF lifted restrictions on commercial uses of the Internet, and turned over operation of its backbone to ANS (American Network Services), the new operator would require settlement-based interconnections to its Internet backbone (Butler, 2000, p.19). An early response to this perceived

threat was the creation of CIX, the Commercial Internet Exchange, a non-profit trade association formed to ensure Internet access for commercial networks on a non-commercial basis (Halabi, 1997, p. 12).

In the midst of these developments, the NSF contracted for the creation of four Internet network access points or NAPs: in Washington DC; San Francisco; Pennsauken, NJ; and Chicago (Halabi, 1997, p. 9). For a period of one year, extending from 1993 through 1994, the NSF subsidized the migration of numerous networks to this new architecture, eventually withdrawing its support for the network access points, on the understanding that the original goal of public peering had been achieved (Halabi, 1997, p. 20). But other factors were by then in play that made interconnection complex and contentious. One of the most important was the commercial ambitions of large network operators that had become involved in supplying infrastructure for the Internet (these operators included, for example, the two RBOCs that had won the right to run two of the four original NAPs, Pacific Bell and Ameritech). Another crucial factor was the product of the Internet's prodigious growth: mounting problems with traffic congestion, with the resulting lost packets, longer latency times and so on.

The need for more dependable delivery across the Internet prompted many ISPs to begin looking for alternatives to open peering (Blake, 1999, p. 15). Until this point, public peering made the Internet highly accessible because ISPs of all

sizes were allowed to interconnect at carrier-neutral network access points and metropolitan-area exchanges. This arrangement was not, however, perceived as beneficial by large bandwidth providers with many customers to serve. Some began to express a preference for exchanging traffic through private peering agreements with networks of comparable size, rather than working through a system in which their customers were given no more priority than those of smaller networks (Blake, 1999, p.15). As Eli Noam explained it:

In contrast to the public interconnection of NAPs and MAEs, the third model of Internet interconnection is “private” and bilateral - peering agreements between backbone ISPs to exchange traffic among themselves or with smaller ISPs. International, national, regional and local ISPs exchange packets via peering centers. A service provider that connects its network to peering centers agrees to set up its routers so they can exchange traffic with routers on peer networks. Such arrangements began in 1995, partly due to the growing congestion of the public exchange points, partly due to the desire to bypass these exchanges and partly in order to establish control. (Noam, 2001, p. 66)

Generally speaking, network operators use peering coordinators to establish and manage interconnections with other networks, and often set up steering committees to evaluate peering requests. However, because network

interconnection agreements are commercial and may contain sensitive competitive information, they are usually confidential, making the full scope and implications of such arrangements difficult to gauge. Therefore, despite the increasing number of peering policies being made available online, a great deal of tension has developed around the problem of transparency, or non-transparency, in peering arrangements around the globe. Peering policies are often posted on network Web sites with a bare minimum of information concerning the terms that a network seeking connection must comply with. William Norton, chief technical officer at Equinix, called peering as “the black art of the Internet” because of the obscurity in which it is shrouded – a sentiment shared by other peering coordinators in the NANOG (Telecommunication Society of Australia Meeting, July 30, 2004).

Before the privatization and commercialization of the Internet began to take hold, fledgling networks found free, open peering interconnection advantageous, since they were able to scale up in reach and redundancy - and the US government encouraged them by providing ample subsidies. In public peering, a network connects to a shared exchange point, where in principle it can connect with all of the other networks at that point. After the mid-1990s, however, network operators began to interconnect not only at public exchange points but at private exchange points as well (Blake, 1999, p.15). In private peering, two networks are physically connected in a bilateral arrangement. Private peering is also used at the Tier-2

and Tier-3 levels to address the geographical problems associated with how many networks can communicate with each other – meaning they must connect indirectly through a national backbone provider. Tier-3 networks typically do not engage in peering since they operate on too small a scale, although they may seek out peering arrangements in order to increase efficiency and reduce costs:

The Internet does not have a purely hierarchical structure. It allows for flexibility around the basic hierarchical structure. First, two ISP's may exchange their local traffic without sending it all the way up and down the hierarchical structure. This "secondary peering" may be efficient in certain circumstances. Second, a participant (ISP, content provider, etc.) may be the customer of multiple backbones or ISP's, a practice labeled "multi-homing". Third, there exist isolated instances of transit contracts that are not pure customer relationships. (Laffont, Marcus, Rey, & Tirole, 2001, p. 288)

Nevertheless, most Tier-2 or Tier-3 networks that wish to access the rest of the global Internet need to pay to transit on a Tier-1 backbone network (Norton, 2001, p. 2). This is a departure from the circumstances in which most networks found themselves prior to widespread commercialization. The original rationale for open peering - known as "sender keep all" or SKA - was based on the high number of similar-sized networks. But the SKA model is viable only in situations

where the traffic exchanged by two networks is balanced (i.e. symmetric) and, as one author put it, where “the cost of terminating the traffic is low compared to the cost of metering it” (Cukier, 1997, p. 3).

The shift to privatization made the backbone market highly concentrated in a short period of time. The record indicates that by November 1997 the four largest US-based networks (UUNET, MCI, BBN and Sprint) controlled between 85% and 95% of total Internet backbone traffic (Cukier, 1997, p. 5). As communication across the Internet has come to depend more and more on a small number of large backbone providers, concerns have grown that these providers might charge prices above fair market value, especially for transit; practice price discrimination; or set unfair peering terms. These potential market distortions raise the risk that less powerful national backbones, a necessary part of a smoothly functioning Internet, may encounter difficulties in maintaining their peering agreements with the more powerful backbones. Similarly, start-up services may find it difficult to enter the market because they cannot afford to pay for transit (Blake, 1999, p.15).

The issues at stake here have become even more problematic with the widespread use of non-disclosure agreements as part of peering contracts, since smaller stakeholders and new entrants are prevented from seeing what terms

other stakeholders have been able to negotiate. In the words of one commentator:

The players with the biggest networks get to call the shots. The largest and oldest Internet service providers set up direct peering links with one another and share the cost. But smaller Internet service providers either have to buy their way in to this old boy's club, at an exorbitant price, or send their traffic through congested public peering points. (Gareiss, 1999, p. 1)

And as Neil Weinberg wrote in *Forbes* magazine:

[T]he practice [of private peering] is a stark departure from how the Internet worked its first three decades, when networks handled one another's data for free under a communal love-in known as "peering". It threatens to balkanize the Net into haves and have-nots (Weinberg, 2000, p. 236).

In recent years, a number of contentious social issues related to the Internet has captured the attention of many journalists and analysts, even in the mainstream press: redefining copyright and intellectual property (an issue receiving a great deal of attention in Canada at the time of writing); articulating the legal rights and

obligations of cultural consumers; and defining the role and liability of the ISPs over whose networks the sharing or “pirating” of content takes place. While these are important issues, they may be distracting us from a deeper and more pervasive issue: confidential agreements among large bandwidth providers on access and interconnection that are undermining the principles on which the Internet was built, thereby putting its long-term viability at risk.

Reachability as a Net Neutrality Issue

Over the past several years, international debates on Internet governance have focused on two areas. The first of these is Net Neutrality, which for the most part concerns broadband subscriber welfare within the regulatory frameworks established by national tribunals such as the FCC and CRTC. The second is the set of transnational issues associated with the Domain Name System (DNS) and its management by ICANN. On the other hand, little attention has been paid to the politics of interconnection and the various ways in which data reachability is being compromised by large bandwidth providers.

As it turns out, the principle of Net Neutrality as it applies to the behavior of Tier-3 (i.e. retail) ISPs has a great deal in common with the goal of data reachability as it applies to Tier-1 backbone providers. In both cases, the policy and social principles at stake spring from widespread concerns about the impact on consumer welfare of privatization, concentration of ownership, and regulatory

reliance on market forces and managed competition. In both cases, the perceived risk is that the ungoverned, unilateral and unpredictable behavior of large networks will prevent end-users from enjoying undisturbed access to the entire global Internet – that is to say, undisturbed by undue discrimination on the part of ISPs supplying access to the Internet on a commercial basis. Moreover, many of the traffic management practices applied by ISPs at the retail level that may discriminate against particular applications, Web sites or users, are the same practices Tier-1 backbone ISPs apply to manipulate or interfere with the traffic carried by other ISPs, whether Tier-1 or lower down the hierarchy. Despite these apparent similarities, however, Net Neutrality remains a controversial concept, one that resists easy definition and stirs vigorous disagreement even among its proponents.

Columbia law professor Tim Wu is often credited with introducing the term “network neutrality” in academic circles through his paper “Network Neutrality, Broadband Discrimination,” written in 2002 (Wu, 2003). Wu’s early formulations of the principle emphasized the benefits of a neutral public network, borrowing from the end-to-end principle inherent in the design of the Internet, and in a more general sense from the obligations associated with acting as a common carrier. As the debate on Net Neutrality moved into the public domain (the mainstream press paid little attention to the topic before 2006), the idea that ISPs should never discriminate became more nuanced and open to interpretation. Wu and

others have revised their positions, noting for one thing that the Internet has never been entirely neutral. Graham Longford describes Wu's more recent thinking in the following way:

Wu proposes a set of criteria for distinguishing between permissible and prohibited forms and grounds of discrimination.... Wu makes a distinction between "capacity-based" and "content-based" discrimination. "Capacity based" discrimination involves placing limits on the amount of traffic users can generate at a given time, without regard to the content of said traffic, while "content-based" discrimination involves blocking, degrading, or, alternatively, enhancing, users' access to certain content for commercial reasons. (Longford, 2007, p. 46)

In 2006, the debate became more public and vociferous. A large coalition of bloggers, educators and consumer-oriented advocacy groups, including Free Press, the American Library Association, Consumers Union and MoveOn, created a grassroots coalition called the Save the Internet campaign. In the space of just two months the coalition collected over a million signatures on a petition calling for network neutrality regulations and delivered it to Congress.

Support for Net Neutrality in some form has also been voiced by content providers such as Google, eBay, Microsoft, Yahoo! and Amazon.¹²

During this time, a number of American and Canadian incumbent telecommunications carriers and cable operators argued *against* the principle of Net Neutrality and especially against the idea of having regulations or legislation brought into force that would govern their behavior as ISPs. Some of these firms, notably AT&T and Verizon, funded It's our Net and Hands off the Internet, dubbed "Astroturf" advocacy groups because, despite accepting money from the carriers, they tried to create the appearance of grass-roots support for the carriers in their fight against Net Neutrality.

In the last three years, several bills have been introduced in both the US Congress and Canadian Parliament with a view to enshrining protections for Net Neutrality and residential broadband consumers. One of the most recent such bills is the Internet Freedom Preservation Act of 2009, introduced in July 2009 by Congressman Ed Markey (D-MA), intended to "establish a national broadband policy, safeguard consumer rights, spur investment and innovation," among other things. Other bills were introduced as far back as 2006, including COPE, the

¹² One of the first systematic analyses of the use of traffic-shaping by an ISP arose from a complaint filed in 2007 by Free Press with the Federal Communications Commission (FCC), alleging that Comcast Corp. had engaged in discriminatory conduct that violated the FCC's Internet Policy Statement. The FCC Order finding against Comcast was released August 20, 2008. At this writing, Comcast has filed suit against the FCC in the United States Court of Appeals for the District of Columbia Circuit, contending that the FCC order had no basis in law.

Communications Opportunity, Promotion and Enhancement Act of 2006 (HR 5252), and the Internet Non-Discrimination Act of 2006 (S. 2360). And in an unusual move for a state government, New York State has established Net Neutrality as a telecommunications standard within its jurisdiction (New York Department of State, 2009).

Wu stated in his 2003 paper that regulations governing Internet access networks must allow broadband operators to make reasonable tradeoffs between the requirements of different applications. This challenge to ISPs, and their critics, has complicated the neutrality debate, since it rests on the assumption that there are some intrusive actions made necessary by good network management practices, while there are others, based for example on undue preference, which should be forbidden:

[B]roadband operators should have full freedom to “police what they own” (the local network) while restrictions based on inter-network indicia should be viewed with suspicion.... [T]he concept of a total ban on network discrimination is counterproductive. Rather, we need distinguish between forbidden grounds of discrimination, those that distort secondary markets, and permissible grounds, those necessary to network administration and harm to the network. (Wu, 2003, pp. 168, 170)

Most of the attention paid to the Net Neutrality debate has been garnered in Canada and the United States, for reasons related to telecommunications infrastructure, digital convergence, and vertical and horizontal integration. These issues are reflected in recent attempts by regulators to provide operational distinctions between acceptable and unacceptable ISP practices, as in the CRTC proceeding conducted in 2009 on ISP traffic management (CRTC, 2008b). The Commission has been examining the behavior and role of Canada's ISPs in several proceedings, including that concerning broadcasting in new media (CRTC, 2009a).¹³ For several years now, engineers have been trying in a parallel effort to gauge the effects of technical and policy decisions made by network operators. Many such decisions have a direct impact on the reachability of Internet service, i.e. the ability of a network to establish and maintain connectivity so that traffic is successfully transmitted. Network engineers use the term "path robustness" to refer to persistence of connectivity between networks, a critical consideration in the arguments presented in this research.

In 2003, for example, an analysis was presented to an IEEE workshop (Institute of Electrical and Electronics Engineers) on Internet applications in which the authors attempted to gauge the impact of interconnection failures. The authors

¹³ As part of its findings in the June 4 new media broadcasting decision (Broadcasting Regulatory Policy CRTC 2009-329), the CRTC elected to refer to the Federal Court of Appeal "the question of whether ISPs, when they provide access to broadcasting content, are broadcasting undertakings within the meaning of the *Broadcasting Act* and are thus subject to the New Media Exemption Order" (para 69). The Commission filed its reference on July 28, 2009. The court's decision, if affirmative, could bring significant changes to how ISPs are regulated in Canada.

reported on observations made during one week of Internet activity using BGP tables¹⁴ from three different, internationally well-connected locations. They defined reachability “as a measure of path robustness over time, and thus as a significant measure of the general quality of the infrastructure” (Salido, Nakahara & Wang, 2003, p. 1). The authors were interested in analyzing the end-user’s experience of the Internet as a function of the reachability or quality of connections between the end-user and other locations on the Internet. The study concluded that single, relatively small incidences of networks failing to interconnect had major repercussions for Internet reachability and thus for end-users:

While overall trends are very stable, it is important to point out that individual incidents were responsible for a good deal of the major observed reductions in reachability. While small in percentage, these incidents are significant in the sense that they affected mostly Internet service providers and telecom companies. (Salido, Nakahara, Wang, 2003, p. 8)

Moreover, while the authors acknowledged that reasons for unreachability included misconfiguration and human error, they wrote a section entitled *Policy*

¹⁴ A routing table at a node provides a look up entry for each destination by identifying an outgoing link/interface or path. (Medhi & Ramasamy, 2007, p. 59) Also see page 24.

Effects, in which they noted that arbitrary network policies also have repercussions for the end-user:

Border Gateway Protocol assumes that the Internet is an arbitrarily interconnected set of autonomous systems. Hence, it allows each AS to independently formulate its own routing policies, and it allows these to override distance metrics in favor of policy concerns. BGP regards issues such as, which routes to accept from a neighbor and the preference with which those routes should be treated, as a local decision based on its routing policy. An important part of this routing policy is to decide which set of paths should be advertised to each BGP neighbor. The decision on which routes to accept from and advertise to various BGP neighbors, has a profound impact on what traffic crosses a network, and hence affects reachability.... [T]herefore, routing policy plays an important role in reachability determination. (Salido, Nakahara & Wang, 2003, p. 8)

More recently, researchers at the University of Washington's Department of Computer Science and Engineering have been using a software system dubbed Hubble to track reachability problems. Hubble monitors Internet transmission activity on a continuous basis to identify situations in which routes exist to a given destination but traffic is unable to reach the destination in question. For purposes of this project, reachability is defined as follows:

Global reachability - when every address is reachable from every other address - is the most basic goal of the Internet. It was specified as a top priority in the original design of the Internet protocols, ahead of high performance or good quality of service, with the philosophy that “there is only one failure, and it is complete partition”. (Anderson, Katz-Bassett, Krishnamurthy & Madhyastha, 2008, p. 1)

As we have seen from recent policy and legislative initiatives in Canada and the US, it is no easy task to reconcile the interests of end-users with those of the ISPs that supply them with last-mile bandwidth. This problem is compounded by concentration of ownership in the ISP business, which confers significant market power on the industry leaders. In Canada, where 500 companies are counted as Internet access providers, the top five companies captured 76% of total access revenues in 2008, namely, Bell Canada, TELUS Communications Corp, Rogers Communications Inc., Vidéotron Ltd., and Shaw Cablesystems G.P. and their affiliates (CRTC, 2009b, pp. 213, 214). Some traffic management practices are intended to improve the overall quality of service on a network, by reducing spam and malware, as well as protecting end-users from aggressive actions like distributed denial of service (DDOS) attacks.

But as testimony at the CRTC’s 2009 proceeding on ISP practices has shown, there is a fine line between this kind of initiative and ISP efforts to block or reduce

certain kinds of traffic in order to reduce what they allege is otherwise unmanageable congestion. Many of the large ISPs claim that they interfere with certain transmissions in order to improve the quality of service for customers – even though such actions inevitably mean that some customers will experience an unwanted failure in reachability as a result.

The problem of data unreachability, which now extends from the last mile to the Tier-1 level, is compounded by other factors, chief among these being the wall of secrecy behind which ISPs operate. Moreover, as we now discuss in detail, ISPs have access to an ever-expanding set of tools with which to manage traffic – and thus more ways to put data reachability at risk.

Traffic Management Practices Affecting Reachability

The traffic management practices of concern to us in this paper are packet filtering, traffic-shaping, autonomous system or AS list filtering, and de-peering.

We begin with a brief definition of each:

- **Packet filtering**, sometimes termed “blackholing,” refers to the rejection of a single IP address or a number of IP addresses (McPherson, Sangli & White, 2005, p. 84, 185).

- **Traffic-shaping** is the imposition of a delay to a set of packets according to some convention or policy established or agreed to by a network operator.
- **Autonomous system (AS) list filtering** is the suppression of certain routing options by removing one or more AS numbers from the routing path (Keshav, 1997, p. 342).
- **De-peering** is the action of severing the physical interconnection between two peered networks or autonomous systems, such that no traffic can flow from one to the other.

These practices are not the sole means by which network operators can precipitate data unreachability. But they are significant in that they scale from the barely detectable dropping of packets, all the way to severing major network connections, with the potential to disrupt thousands or even millions of end-users.

Packet Filtering

Filtering packets was introduced as a traffic management practice relatively early in the history of the modern Internet. One of the first major vendors to make equipment explicitly designed to manipulate data routing was Cisco Systems, which began promoting this idea in 1999:

In 1999, Cisco Systems issued a technical white paper, which described a new router that the company planned to sell to cable broadband providers.... In plain English, the broadband provider could inspect the packets flowing to and from a customer, and decide which packets would go through faster and more reliably, and which would slow down or be lost. Its engineering purpose was to improve quality of service. However, it could readily be used to make it harder for individual users to receive information that they want to subscribe to, and easier for them to receive information from sites preferred by the provider—for example, the provider's own site, or sites of those who pay the cable operator for using this function to help “encourage” users to adopt their services. There are no reports of broadband providers using these capabilities systematically. But occasional events, such as when Canada's second largest telecommunications company blocked access for all its subscribers and those of smaller Internet service providers that relied on its network to the website of the Telecommunications Workers Union in 2005, suggest that the concern is far from imaginary. (Benkler, 2006, p. 147)

Cisco's interest in developing and selling networking equipment that can manage data traffic has continued over the last decade:

Moreover, prioritizing and inspecting traffic (for security reasons) were important tools for building new equipment markets working from the router out through the rest of the network. Cisco, for example, is buying into service application companies that feature traffic prioritization and security schemes based on capabilities installed in Cisco routers.

(Aronson & Cowhey, 2009, p. 114)

Benkler's reference above to the TELUS dispute indicates the extent of the harm that can be caused when an ISP blocks a single IP address. In 2005, when TELUS was in the midst of a labor dispute with the Telecommunications Workers Union, the company blocked public access to the union's Web site by filtering out the IP address of the server on which the site was hosted. In doing so, it also blocked "more than six hundred additional websites hosted at the same IP address and cut off entire communities from the controversial content" (Geist, 2005). Researchers at Harvard, Cambridge and the University of Toronto's OpenNet Initiative later found that TELUS had actually blocked access to 766 unrelated sites (OpenNet Initiative Bulletin 010, 2005).

In May 2009, the state of Minnesota Alcohol and Gambling Enforcement Division of the Department of Public Safety invoked the Interstate Wire Act of 1961 to block access to online gambling by issuing to 11 ISPs, including Comcast, Qwest and Sprint, the names and IP addresses of nearly 200 gambling sites, requesting

that they be blocked so that end-users in the state could not access them (Walsh, 2009, p. A1). John Levine, a NANOG engineer, noted how the IP filtering would occur in response to a query on the listserv:

Notwithstanding the legality of such an order, how would one operationally enforce that order? Answer: The order has a list of IP addresses, so I expect the ISPs will just block those IPs in routers somewhere. (Levine, 2009)

According to an “issue brief” posted online at the time by NETCompetition.org, there was no sustainable rationale for this action by the Minnesota authorities, legally or ethically speaking:

[T]he aggressive move by Minnesota’s Alcohol and Gambling Enforcement Division ... marks an unprecedented erosion of net neutrality. Savetheinternet.com, an outspoken net neutrality organization, has said “net neutrality means no discrimination. Net Neutrality prevents internet providers from blocking, speeding up or slowing down web content based on its source, ownership or destination.” Under this conception, AGED’s move smacks of infraction by blocking Minnesota internet customers from accessing certain websites or “destinations,” which at this point includes the 200 online poker sites. Of the nearly 200 sites, only 44 even accept

U.S. players. Until a Minnesota statute makes playing online poker an explicit crime, these companies are having their web traffic disabled for what can only be deemed no legal reason. The Minnesota AGED simply has no leg to stand on. (NETCompetition, 2009)

Traffic-shaping

Although the emphasis in this chapter is on the practices of ISP businesses owned by telephone carriers, as opposed to cable MSOs, traffic-shaping is practised by all major ISPs, regardless of the technical platform in question. The difficulties created for end-users by traffic-shaping are compounded by the oligopolistic structure of the ISP industry in both Canada and North America. Few consumers have an access option other than their incumbent telephone carrier or cable supplier. Furthermore, despite attempts by the CRTC to create viable conditions for resellers, Bell Canada does its traffic-shaping at the wholesale level, meaning that all DSL resellers in Ontario and Quebec that depend on Bell's GAS (Gateway Access Service) are unable to offer a truly competitive service, free of intrusive practices like traffic-shaping (see below for comments on related CRTC proceedings).

Until recently, concerns with quality of service (QoS) were largely confined to the business sector. As part of their agreements with business customers, especially large ones or those with a significant Web presence, ISPs provide guarantees

about service defined by such parameters as low latency, high through-put and low cost. When QoS terms apply, the ISP in question is contractually expected to meter usage and set transport priorities for particular kinds of data. Traffic-shaping provides a means for controlling the volume of traffic being sent through a network in a given period of time by imposing delays on traffic so that the traffic conforms to some predetermined network flow constraint. QoS may be practised intradomain (within the AS or network) or interdomain (between ASes or networks). Here again, the argument is complicated by the difficulties of determining whether a particular traffic management practice is, in some objective sense, “intrusive” or contributes to better overall service. As Eli Noam writes:

Given the option, many customers could well select lower technical quality if the price is right. Millions of users prefer a jalopy to a Rolls-Royce.

To complicate things still further, one must recognize that quality to users is not a static concept but a relation between performance and requirements. Since the latter are shifting, what constitutes good quality is a moving target. What was good enough yesterday may not be enough today, and not just because we tend to take past luxuries soon for granted but also because past standards move from being merely convenient to being vital.... By becoming increasingly dependent on high-tech

communications flows, advanced societies also put themselves at risk. In consequence, demands on service quality increase because failure becomes unacceptable. (Noam, 2001, p. 203)

Traffic-shaping and Net Neutrality have become the subject of public debate in Canada, as well as a growing preoccupation for Canada's national regulator, the CRTC. In 2007, CRTC chairman Konrad von Finckenstein touched on Net Neutrality in his address to the Broadcasting Invitational Summit, as he did again in 2008 at the Canadian Telecom Summit (von Finckenstein, 2008). At about the same time, traffic-shaping practices became the subject of a Part VII complaint filed with the CRTC in April 2008 by the Canadian Association of Internet Providers (CAIP) against Bell Canada, requesting that the Commission order Bell Canada to cease and desist from using traffic-shaping of its wholesale ADSL service known as the GAS, Gateway Access Service (CRTC, 2008a). On November 20, 2008 the CTRC issued its decision regarding the CAIP complaint. Although the Commission denied CAIP's application, it did so after deciding that the issue of traffic-shaping (and related traffic management practices) should be examined in a full public inquiry. Thus, on the same day that it released its CAIP decision, the CRTC issued a "Notice of consultation and hearing to review the Internet traffic management practices of Internet service providers" (Telecom Public Notice CRTC 2008-19; CRTC, 2008b).

In the summer 2009, the CRTC held public hearings as part of its proceeding on traffic management practices. Michael Geist took a critical stance on the use of traffic-shaping and similar practices by Canada's largest broadband ISPs:

I think that the consumer groups rightly focused on who should bear the burden of demonstrating that DPI and other Internet traffic controls are consistent with current Canadian law. The groups argued that these are prima facie violations of Section 36 of the *Telecommunications Act* and that the onus therefore should fall on the carriers to show that there is a serious problem, the solution minimally impairs users' rights, and is proportional to harm.... Unfortunately, the questions that followed suggest that the CRTC Commissioners started these hearings having accepted the carriers' claims that congestion is a problem and that inhibiting the use of deep packet inspection could result in increased consumer costs for Internet access. (Geist, 2009a)

One of the issues that surfaced during the hearings was the CAIP complaint and whether, as a practical matter, Bell Canada in particular was technically able to refrain from traffic-shaping at the wholesale level (i.e. as part of its wholesale Gateway Access Service, or GAS). During questioning by Commissioner Denton on the last day of the hearing (July 14, 2009b), Bell representatives indicated that

they were obliged to treat all DSL resellers in exactly the same fashion in respect of wholesale provisioning:

6774 [MR. DANIELS] Our problems are, number one, that we can't distinguish in our network between the various different wholesale providers. We can distinguish all wholesale but we can't make a distinction, Oh! Execulink is managing it correctly, someone else isn't.

6775 COMMISSIONER DENTON: Why is that? Could not the code be written whereby the origin could be distinguished? Is this just a software problem in terms of recognition of these sources of traffic?

6776 MR. CONDON: It is a protocol problem, Commissioner. The tunnelling that is used is dynamic.

6777 COMMISSIONER DENTON: Say that again. The tunnelling...?

6778 MR. CONDON: The tunnel IDs are dynamic. We can't tell the difference between them.

6779 COMMISSIONER DENTON: Tunnel IDs are dynamic and therefore they can be spoofed or hidden or just basically anonymized?

6780 MR. CONDON: No, that is not what I meant, Commissioner. I wasn't suggesting anything of that nature. I was just suggesting that at that level of the network they are not identified, they are not a permanent identifier to, say, Execulink.

6781 COMMISSIONER DENTON: Okay. (CRTC, 2009b)

The CRTC proceedings dealing with traffic-shaping and other network engineering issues have focussed attention on the use of deep packet inspection technologies by both telephone carriers and cable companies in their wholesale and retail provisioning. In digital networks, data and increasingly voice communications are broken up into discrete packets that travel along independent routes between point of origin and destination where these fragments are then reassembled into the original whole message ("packet-based"). Not only is there no longer a dedicated analog circuit, but individual packets from the same communication may take completely different paths to their destination. To intercept or eavesdrop, "packet-sniffing" or "deep-packet inspection" hardware is employed:

Sandvine Inc., of Waterloo, Ontario, combines its DPI technology with equipment that lets Internet-service providers offer individual users more

finely tuned subscription packages. The company made headlines last year when it supplied Comcast Corp. with the ability to identify and block certain types of traffic shared between users - a high-profile case that led to sanctions on Comcast by the Federal Communications Commission. Sandvine's earnings and stock price plunged after the incident.

A Comcast spokeswoman says the company has since adopted a program that manages network traffic by the capacity used by subscribers, rather than by the type of traffic. Sandvine, which provided Comcast's redesigned network-management system, has begun diversifying beyond the U.S. cable market, according to Tom Donnelly, a co-founder and executive vice president.

More-powerful routers from companies like Cisco Systems Inc. often have DPI baked in. A Cisco spokesman says that in most cases the equipment is for enhanced Internet-security. (Rhoads, 2009, p. B6)

AS Path Filtering & De-peering

As explained earlier in this paper, peering is the voluntary, settlement-free exchange of traffic between two global networks. Peering presupposes that both parties are benefitting from the arrangement; if one or both networks believe they are no longer realizing a net benefit, they may decide to de-peer. A network may

wish to de-peer with another because it believes the other network is profiting unfairly from the free interconnection. This situation typically arises when the traffic ratio between the two peers becomes asymmetric, meaning costs are not being shared fairly. De-peering may also be precipitated by an abuse of the service connection, network instability, repeated routing leaks or one party's unwillingness to provide more peering capacity. In some cases, a network may wish to de-peer in order to peer with a different, more successful global bandwidth entity.

De-peering can be viewed as an extension of retail network management practices that include packet filtering and traffic-shaping. De-peering involves breaking the path between backbone networks by unplugging physical connections. Although it usually accompanies the severing of such a connection, autonomous system number (ASN) filtering is a separate operation from de-peering. AS path filtering occurs when network A disallows any data traffic from network B that must pass through intermediate networks. The incident described below, which took place in March 2008 at York University, details the effects of unreachability that can result from a network de-peering.

An instructor using the on-campus network, for which commercial Internet connectivity was supplied by Cogent, attempted to reach a particular Web site located in Europe that was important to his research. After repeated attempts, he

was still unable to reach the site, which he had never had difficulty reaching in the past. He contacted York's Information Technology department. Staff were mystified as to the cause of the connectivity failure. When the instructor went home, however, he encountered no problem reaching the Web site in question. Several days later, the IT staff involved discovered that York's transit supplier, Cogent, had discontinued its peering relationship with the European Tier-1 network, TeliaSonera, the bandwidth provider for the site the instructor had been trying to reach. The instructor was able to reach the Web site in question from his home, because Rogers, which supplied his Internet access, had a stable relationship with TeliaSonera through its upstream network providers.

Because Cogent did not take any steps to inform York University of the dispute and loss of connectivity, IT staff had to open a trouble ticket with Cogent in order to have the incident and its origins explained (Paterson, 2008). Cogent's actions, or failure to act, resulted in an outage that affected a significant section of the Internet:

[T]he list of impacted networks is too long to be included here, but they include a wide range of commercial, educational and government clients. On the Telia side, the victims include the Swedish Defense Data Agency, the Finnish State Computer Center, and broadband customers in St. Petersburg. With regard to Cogent, Blue Cross and Blue Shield of

Delaware, Kansas State University and Reuters America were all collateral damage. (Zmijewski, 2008a)

De-peering in itself is not typically a sufficient condition to cause data unreachability – unless the de-peering is accompanied by another step, Autonomous System or AS path filtering. When one AS sees incoming traffic that has either originated from the offending AS or has that AS number in its routing path, it simply disallows all the related data traffic. To one high-level network, the other network simply does not exist (Halabi, 1997, p. 171). Here is how one commentator described the Cogent-TeliaSonera dispute and its implications:

Cogent is one of the five largest networks in the world in terms of the number of peers with which it works and more than 95% of Cogent's traffic goes across private peering connections. The peering dispute between Cogent and TeliaSonera left many networks in the U. S. and Europe unable to connect with one another. Renesys, which tracks Internet routing, had some additional details on the impact of the dispute. Renesys said many networks were unable to simply route around the impasse, perhaps because one party (probably Cogent) had taken steps to block alternate traffic paths. Their analysis found that 2,383 TeliaSonera network prefixes could not reach Cogent at all, while 1,573 Cogent network prefixes were completely cut off from TeliaSonera. The impact

was most widely felt in Europe, but more than 1,900 U.S. network segments were affected as well. What was surprising was that networks in the U.S. were actually cut off from each another given that a largely U.S. provider was de-peering with a largely Swedish one.... Renesys noted that Flag and SingTel discontinued peering shortly after the Cogent-TeliaSonera peering was restored in late March, but were allowing customers to find one another via alternate routes. (Miller, 2008a, p. 1)

In cases where a de-peering is not accompanied by AS path filtering (as with Flag and SingTel in the quotation above), a network operator can take steps to mitigate its impact by looking for alternative routing possibilities for customers affected by suspension of service. Whether or not such alternatives exist, many commentators believe that ISPs should make their customers aware that a de-peering has occurred:

...[I]t is possible for providers to provide subscribers with information about the nature of the services being provided. Subscribers need to be aware of whether they are receiving oblivious transport, and if not, how the service affects their traffic. (Aboba & Davies, 2007, p. 3)

A strategy open to ISPs that wish to minimize the impact of loss of connectivity to one upstream bandwidth supplier is redundancy, i.e. arranging for connectivity

with two or more suppliers. These topologies are known as dual-homing (connectivity via two upstream suppliers) and multi-homing (connectivity via more than two upstream suppliers). This form of redundancy raises costs significantly, and also requires that network operators follow certain technical procedures concerning the allocation of numeric IP addresses.¹⁵

While present-day de-peering incidents tend to be abrupt and unpredictable, de-peering was carried out as a systematic business strategy in the 1990s, as Internet resources became privatized. Earlier, bandwidth providers had operated under open peering policies. As private bandwidth providers entered the game, however, they were lured by the profits and market share to be gained by systematic de-peering – charging erstwhile peers for interconnection privileges and for carrying their traffic. This trend tipped the scales against open peering, until it finally collapsed with MFS's (later Worldcom) 1996 purchase of UUNET. UUNET announced that smaller networks with whom it had previously peered would have to start paying to connect to its backbone. New, bilateral transit agreements would have to be negotiated that transformed peers into customers (Cukier, 1997).

¹⁵ Typically, purchasers of transit bandwidth use “provider-dependent” IP addresses. If an ISP that purchases transit wishes to be dual- or multi-homed, it must obtain provider-independent IP addresses in one or more blocks from their regional Internet registry, in a process similar to how AS numbers are obtained (McPherson, Sangli & White, 2005, p. 43). Under this formula, provider-independent IP addresses are advertised through more than one upstream transit bandwidth provider (McPherson, Sangli & White, 2005, p. 56). In principle, these strategies make the entire Internet reachable from two or more completely separate sources.

UUnet said it would stop peering with small carriers; they would have to pay. Sprint and AT&T followed suit within months. The modern Net began to emerge. Titans swap traffic free and charge others; those who can't pay take the back roads of unreliable public exchanges. (Weinberg, 2000, p. 236)

Twelve entities - including GeoNet Communications Inc., NetRail and Whole Earth Networks - balked at paying monthly interconnection fees in the tens of thousands of dollars to UUNET. After a barrage of negative publicity, however, UUNET agreed to grandfather its prior peering agreements. UUNET insisted that entities seeking peering had to sign a non-disclosure agreement (NDA) as well as a peering agreement. Other bandwidth providers, such as MCI and Sprint, followed UUNET's lead and began operating under non-disclosure agreements. As backbone providers were bought out by larger stakeholders, market consolidation decreased the number of backbones and widened the margin between Tier-1 providers and others (Blake, 1999, p. 15). This shift in the marketplace was notable for another reason, namely that the backbone business was becoming a desirable line of business for facilities-based telecommunications carriers.

Bolt, Beranek and Newman (later BBN Technologies) pioneered many Internet technologies, having been involved in the earliest stages of development work on

ARPANET. During the 1980s, BBN became one of the first private firms to operate an ISP division, known as BBN Planet. In 1997, GTE, a company with a long history of aggressive acquisitions in telecommunications, bought BBN Planet.¹⁶ The following year, the newly acquired company informed two firms with which it had peering agreements - Exodus and AboveNet - that it would no longer peer with them. Although this decision was not officially made public, it was openly discussed on the North American Network Operators Group listserv (Butler, 2000, p. 47). The dispute that broke out between Exodus and BBN Planet (GTE Internetworking), dubbed the “peering wars of 1998,” were symptomatic of the radical changes brought about by the privatization of backbone facilities (Fusco, 2000, p. 1). The peering wars led to major connectivity outages in 1998.

Meanwhile, Tier-1 backbone providers pursued the strategy of peering with fewer entities by converting peers into transit customers. By 2000, the practice of de-peering had expanded from a method for changing or discontinuing a business relationship, to a much more aggressive style of competition. Thus, in March 2000, when PSINet severed its peering relationship with Exodus, Exodus CEO Ellen Hancock commented publicly that, in addition to making her company a transit customer, PSINet was taking the unusual step of refusing to exchange traffic at public points (Fusco, 2000, p. 1). Her comments indicated that PSINet’s

¹⁶ In June 2000, GTE merged with Bell Atlantic to form Verizon Communications.

owners, not content with the initial de-peering, had also begun to apply active filtering of AS numbers, effectively making other networks invisible.

PSINet at this time enjoyed considerable market power, having been a major participant in the commercialization of the Internet. In 1991 the company acquired or merged with other ISPs and began expanding in Europe, while also investing heavily in fiberoptic technology. PSINet's goal was to become a single-source provider of emerging communications technologies. Despite some pioneering achievements and ambitious plans, however, the company was never profitable. PSINet's over-expansion left it vulnerable to more sophisticated competitors in the Internet service market. Its crippling debt load finally forced the company to file for bankruptcy protection in June, 2001, along with four of its subsidiaries. Cable & Wireless, which had stopped peering with PSINet and was charging for transit, was unable to collect fees it was owed during PSINet's bankruptcy proceedings. In another sign of changing times, C&W responded by cutting off its interconnection with PSINet for four days (Burton, 2001, p. 1).

Numerous other incidents took place in the following years which demonstrated that free-market competition and privatization were having adverse effects on data reachability.

In 2002, AOL Transit Data Network (ATDN) and Cogent Communications de-peered (Noguchi, 2002, p. 1). Cogent was at that time a three-year-old firm offering a single service: high-speed Internet access to customers such as schools, universities and some 6,000 other large bandwidth users. Cogent lost the peer connection it had maintained with ATDN in December 2002, after the business relationship between the two companies soured. ATDN's decision was prompted by a cost-benefit analysis which revealed that Cogent was delivering three times more data onto ATDN's network than it was carrying back to its own customers. The extent of this asymmetry in their traffic exchanges was having an unacceptable impact on ATDN's profitability.

In 2003, state-owned France Télécom (Wanadoo) & Proxad (now called "Free") de-peered in a dispute over size, capacity and congestion at peering points (Le Bouder, 2003, p. 1). In 2005, France Télécom severed all links between its network and that of its largest competitor, Cogent, in a dispute over Cogent's rapid growth on the continent (Jürgen & Smith, 2005). As a result, all parties linked via Cogent, including a number of German customers, were unable to reach a majority of France Télécom's customers. France Télécom accused Cogent of having violated its peering policy after Cogent retaliated by filtering all France Télécom IP addresses – another example of how de-peering was escalating into more aggressive and damaging corporate behavior.

Level 3 Communications, one of the world's largest bandwidth providers and a Tier-1 operator, has been involved in several disputes over the years with peering partners. In 2005, Level 3 Communications and Cogent Communications de-peered after Level 3 maintained it was carrying the bulk of the traffic in its deal with Cogent and thus providing free capacity (Brown, Hepner & Popescu, 2009, p. 14). As Level 3 spokeswoman Jennifer Daumler suggested, the arrangement was not commercially viable (Cowley, 2005). Three days into the standoff Level 3 backed down and restored its peering connection to Cogent (Ricknäs, 2008).

Several years earlier, in 1998, Level 3 had entered into an agreement with XO Communications to collaborate on development of a fiberoptic network. Subsequently, XO Communications agreed to purchase transit services from Level 3 that reduced its position as a full peer. In 2007, after a long-standing dispute between the two companies over their mutual rights and obligations, the court found in favor of XO Communications, forcing Level 3 to abide by the original agreement. Despite the contentious nature of their legal battle, neither of the two firms resorted to de-peering and no network outages took place as a result (XO wins ruling against Level 3 Communications, 2007, p. 1).

In 2008 Cogent Communications & TeliaSonera de-peered (Brown, Hepner & Popescu, 2009, p. 14). Telia made the following statement in a letter to its customers:

Cogent has decided not to exchange traffic directly with TeliaSonera's AS1299 or indirectly with AS1299 through a third-party provider. As a result, Cogent has partitioned the Internet and disrupted the flow of traffic between Cogent and TeliaSonera customers (Malik, 2008).

According to an article that appeared in *PC World*, the conflict with TeliaSonera developed over the cost for upgrading a peering point in the US (Ricknäs, 2008). This de-peering affected Northern and Central Europe, served by Telia, and the United States and Canada, served by Cogent (Brown, Hepner & Popescu, 2009, p. 16). This de-peering affected York University, among many customers. As one commentator noted about the effects of this de-peering:

This is the way the Internet works and sometimes doesn't work. If the businesses that run the show don't play nice with one another, their customers can pay the price of being cut off from parts of the 'net.... The Cogent/Telia tiff has been going on for 4 days now and only they can resolve their differences. The rest of the world can only hope for full connectivity to be restored. (Zmijewski, 2008a)

In 2008 Sprint Nextel and Cogent Communications de-peered, in what at the time of writing was the most recent high-level de-peering (Brown, Hepner & Popescu, 2009, p. 14). As one commentator wrote at the time:

At the heart of it, peering disputes are really loud business negotiations, and angry customers can be used as leverage by either side. This one will end as they always do, with one side agreeing to pay up or manage their traffic differently. (Miller, 2008b)

MPLS: Expanding the Scope of Traffic Management

Like the other network engineering tools and protocols just described, multiprotocol label switching (MPLS) is a technology used by network operators to manage data traffic. However, MPLS is in a class by itself and, as we shall see, has far-reaching implications for the Internet's open architecture.

MPLS is employed extensively in the core of global backbone networks to separate different types of traffic into distinct logical layers running on the IP infrastructure. In an MPLS-based network, incoming packets are assigned a "label." Packets are then forwarded along a label switch path (LSP) and each label switch router (LSR) makes forwarding decisions based solely on the contents of the label. At each hop, the LSR strips off the existing label and applies a new label, which tells the next routing hop how to forward the packet (DeGeest, 2001). MPLS represents a significant change in TCP/IP architectures because *it effectively replaces IP routing* – in other words, packet forwarding decisions are no longer based on fields in the IP header and routing table, but on labels that are attached to packets (Medhi & Ramasamy, 2007, p. 614).

The protocols supporting MPLS networking have been in development since the end of the 1990s, and by late 2002, MPLS networks were becoming widely implemented in Internet backbone networks (King-Guillaume, 2003; Guofeng, Hong, & Yi, 2004).¹⁷ In recent years network operators have been migrating MPLS from private networks to the public Internet, giving a broad new meaning to the concept of the Internet “cloud.” This trend means that end-user networks are becoming more closely entwined with and governed by the backbone network’s traffic control system.

MPLS is especially attractive in the transmission of audio and video data, or what is known as isochronous data transmission, because such transmissions are sensitive to latency or time delays, unlike traffic that is not isochronous, such as email transfers. AT&T's U-verse and Verizon's FiOS services are prominent examples of commercial networks that serve high-bandwidth video content and have MPLS networks at their backbone to maintain quality of service (Perez, 2008). Video, especially in high definition formats, has to be transmitted at very high bitrates in order to preserve quality. This goal is difficult to achieve over the Internet because even if both the content originator and content consumer use the same network, the packets they exchange typically have to traverse intervening networks. Thus, although data packets may get priority service while

¹⁷ The original patent application for the “MPLS packet” was filed on December 29, 2003 and the patent was assigned to AT&T Intellectual Property I, LP, Reno, Nevada, on October 11, 2008 (US Fed News, 2008).

traversing the sending network, they may lose their priority level before reaching their intended destination. Some commentators see this as an important shift in underlying assumptions about the Internet and how it functions: "... Internet TV services will require sophisticated monitoring and feedback software running in the video servers in order to respond quickly to changing quality of service" (Jamie, 2009; Hunter, 2006b, p. 32). In addition to a physical topology of routers and other hardware for interconnection, networks have a logical topology. Different routing software establishes the logical topology in different ways. For some backbone networks the core of the network (IP) has been replaced by MPLS, as represented in Figure 3.

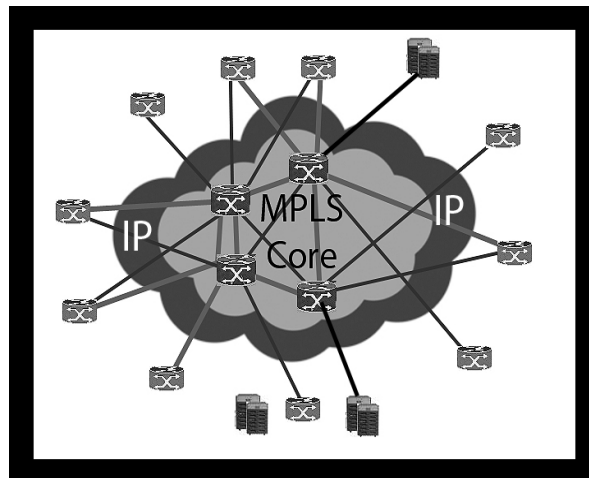


Figure 3. MPLS logical topology. Adapted from: Passmore, D. (2004, November 2). "Strategic Networking Overview: Major Trends in Broadband Networking". In *Proceedings of Next Generation Networks*. Boston, MA. p. 24.

The most controversial aspect of MPLS deployments concerns its potential use at the edge of the Internet, i.e. close to end-users, last-mile ISPs and ISP upstream bandwidth providers. Although the use of deep packet inspection (DPI) technologies is not dependent on the accompanying use of MPLS, these technologies are being incorporated more and more into network platforms that also use MPLS. To take one recent illustration:

XO Communications today unveiled a service that uses deep packet inspection (DPI) at the edge of its network to tell enterprise customers how their applications are performing over their Multi-Protocol Label Switching (MPLS) networks. The applications performance management service uses Fluke's Visual Uptime Select probes that sit on the customers' premises at the edge of XO's wide area network and use DPI to examine each packet in real time and identify what it is carrying.... (Wilson, 2009)

Traffic engineering, QoS policies, and practices such as filtering and de-peering, may eventually create network interoperability issues that will change the nature of the Internet. Traffic prioritization or traffic-shaping is a leading concern in the larger debate about data reachability and the concept of a neutral network. In theory, traffic engineering for a neutral network should ensure as a best practice that data is reachable, although there is much debate within the engineering community as to what constitutes an appropriately neutral network.

Chapter V. Conclusions and Future Directions

The preceding chapters offered a detailed analysis of a little studied but growing set of risks to the long-term sustainability of the global public Internet. We identified sustainability in terms of “data reachability” – the principle that any end-user on the Internet can transmit data to any other end-user without encountering arbitrary actions on the part of an intervening network operator that might block or degrade transmission of the data in question.

We identified risks to reachability in terms of several network engineering practices whose purpose is to manipulate data traffic: packet filtering, traffic-shaping and AS path filtering, along with the concurrent use of MPLS. The policy and business challenge here is that these practices are perceived in different and often contradictory ways by the ISPs that implement them and the customers who are subjected to them (although few retail customers would be able to identify the technical reasons for a disruption in their service). We suggested that while it is not strictly a network management tool, de-peering – the physical severance of connectivity between two Tier-1 networks – should be grouped with the other traffic interventions, since de-peerings are also based on unilateral decisions taken by network operators and have similar end results, namely, disruptions in reachability.

Lessons of History

In order to better understand the long-term implications of these developments, we devoted much of this study to the historical development of the Internet.

First, we wished to explain why the privatization of Internet resources (bandwidth, interconnection points, etc.), and the commercialization of Internet-related activities (as seen in the rise of e-commerce), eventually created a crisis of international governance and, in Canada and the US, a crisis of national regulation as well. Historically, we attributed this crisis to the clash of two cultures: the entrenched Bellhead culture that typified AT&T and other monopoly telephone carriers; and the Nethead culture that typified the community of engineers and computer scientists who created the Internet.

The Bellheads, who had long relied on closed, circuit-switched networks optimized for voice telephony, were unable to develop new networking models that could match the innovative possibilities of packet-switching and the Internet protocol suite. A two-decade protocol war, running from the mid-1970s to the mid-1990s, culminated in a decisive victory for TCP/IP over the Open System Interconnection (OSI) model that telephone monopolies and standards-setting bodies had struggled to implement. The Internet community carried the day by creating an open platform that would prove exceptionally resilient, scalable and adaptable to new technologies.

Even as the open Nethead culture prevailed in the mid-1990s, allowing for prodigious achievements such as the World Wide Web, the Bellheads – and Bell companies – had been steadily co-opting the Internet technologies they once vigorously opposed. Ironically, the system of governance that served the Internet bodies so well left a vacuum that the telecommunications carriers were well equipped and eager to fill, thanks to their extensive experience as supplicants before the FCC. The peer-oriented approach taken by the Netheads was ill-suited to the world of traditional regulation.

Internet governance was further altered by the confluence of several important events falling into the three-year period from 1995 to 1998: the creation of the WTO and the international push for deregulation; withdrawal of the NSF's support for Internet-related activities; passage of the 1996 Telecommunications Act and the thwarted efforts to introduce more competition into the marketplace; concentration of ownership that ultimately reduced the original seven Baby Bells (plus AT&T) to three incumbent carriers; the mainstreaming and commercialization of the Web; and the creation of ICANN in 1998. Many of these events had a direct relationship to the changes taking place in the Internet backbone market – particularly consolidation and the dominant role of American-based carriers.

The larger point here is that the whole concept of Internet governance became deeply fractured by the events of the mid-1990s. On the national level, the FCC and CRTC inherited regulatory frameworks that were ill-equipped for managing the social and economic ramifications of digital communications technologies, increasingly identified with the computer browser interface and the Web (the flurry of regulatory activity in 2009, including the CRTC's June 4 call for a national digital strategy, suggests just how unmanageable the Internet has been under regulations crafted for telecommunications carriers and broadcasters). On the international level, governance has not fared much better – with the notable exception of the technical governance provided by the ISOC-led groups that continue to manage Internet protocols, architecture, security and a wide range of other issues.

This crisis in governance formed the backdrop to the narrower set of issues we have concentrated on in this study: the technical and business behaviors of bandwidth providers.

In our discussions of the role played in recent years by the Tier-1 networks, we noted that the business relationships they enter into are problematic in two fundamental ways. First, these relationships are not subject to outside scrutiny on the part of any governing body with public-interest or fiduciary responsibilities. The second problem is that of transparency. The Tier-1 networks and their

customers operate in near-total secrecy, behind the non-disclosure agreements that have become routine in the industry, a practice that compounds their lack of accountability to any disinterested third party. The result is that the largest providers – and for that matter many smaller ISPs as well – can operate with impunity, making overly restrictive or harmful traffic management practices, along with de-peering incidents, more likely.

At this writing, many groups in the United States, and to a much lesser extent Canada, are engaged in a lively debate about the social and economic benefits of broadband access to the Internet. Generally speaking, for organizations which are concerned with the welfare of residential end-users, a robust and sustainable public Internet must be fast, affordable and ubiquitous. While our perspective is not inconsistent with these aims, the notion of data reachability shifts the emphasis to a slightly different set of concerns. To be sure, the use of traffic-shaping received a thorough airing before the CRTC in 2009, and is likely to come under close scrutiny at the FCC, under newly appointed chair Julius Genachowski. Nevertheless, our surmise is that for most end users, the sheer ability to reach any and every address on the global Internet is a remote abstraction, compared to a) being able to get broadband, b) being able to afford it, and c) having some assurance that once a message or request is sent to a particular address, the response time will be reasonably consonant with the network throughput the end-user has become accustomed to.

To the extent that attention has been paid to reachability in popular and academic discussions, the concept has been framed in terms of the risk of brownouts and other wide-area losses of service quality. It has been carefully positioned by the incumbent ISPs (e.g. in the July 2009 public hearings before the CRTC) as the unwanted yet potentially uncontrollable consequence of limited network capacity and rising demand on the part of subscribers for bandwidth-intensive content, especially video. This scenario, as the ISPs would have it, is compounded by the allegedly overwhelming costs of capital expenditures to upgrade middle- and last-mile capacity; and the profligate use, by a small number of customers, of peer-to-peer platforms like BitTorrent to facilitate the downloading of very large files. This framing of events forms the backdrop to current public discussions of what traffic-shaping is and why it is necessary.

The result of careful positioning by the incumbent ISPs is that reachability is almost universally treated as a *capacity* issue rather than as a *connectivity* issue. This distinction has important implications for scholarship, public policy and advocacy. Regardless of the degree to which capacity issues are addressed by policymakers and regulators, no amount of available bandwidth will keep the global Internet functioning properly if smaller ISPs and end-users must tolerate practices whose intention is to degrade connectivity, if not cut it off entirely.

In addition to moving the emphasis from capacity to connectivity, we have also emphasized the importance of a shift in perspective away from *naming* (along with the much disputed role of ICANN), in favor of a more concentrated focus on the engineering and policy issues related to Internet *routing*. We do not mean to imply that naming issues are not worthy of serious attention, especially given that the DNS root has already been put in jeopardy by actions such as attempts to partition off parts of the Internet by the government of China. Nevertheless, we contend that routing deserves more attention, precisely because it goes largely unnoticed outside the Internet engineering community, and yet has been undergoing changes that may gradually compromise the end-to-end principle.

Disclosure as an Instrument of Reform

The question, then, is - can anything to be done to alleviate the problems described in this study?

The first step is to be clear that accountability (to a higher authority) and disclosure (of technical and business practices) are separate and distinct issues. The Internet's problems are international problems and it is tempting to look to established international institutions for relief. To begin with, ICANN is poorly positioned to play any role, directly or indirectly, in addressing the business practices that affect routing and data reachability. Aside from the scope of its legal mandate, ICANN is embroiled in controversy on several fronts – and at this

writing it appears that the US government, through the Department of Commerce, will insist on maintaining effective control over ICANN's activities upon expiry of the Joint Project Agreement on September 30, 2009. If that is indeed the outcome, the renewal is likely to further undermine ICANN's credibility among stakeholders in the Internet community.

At the same time, the tenor of WSIS and IGF discussions in the past lends little credence to the prospect that the United Nations or one of its bodies, particularly the ITU, might be an appropriate forum for management of international peering agreements (an IGF meeting is scheduled for November 2009 in Egypt). For their part, the ISOC technical bodies have done an exemplary job of technical governance. On the other hand, they have no mandate to govern or control the *business* terms under which Internet resources are used by ISPs or other firms.

Two further distinctions apply here. First, technical practices (such as whether a network uses multi-homing) are not the same as business practices (such as how much a network *pays* for multi-homing); and they will certainly be treated in very different ways in any process of reform. Second, the rights and responsibilities of a Tier-2 or Tier-3 network are quite different from those of a residential end-user, which are different in turn from those of business customers of lower-tier ISPs. Nevertheless, what they all have in common may begin to

offer insights into how reachability can be promoted as a goal among Internet stakeholders.

We believe that focussing on disclosure is the best way to approach the issues attaching to reachability. There are several compelling reasons for this. All point to a set of goals related to increased understanding and awareness of what reachability is, why it is important, and how risks should be classified and assessed.

First, the information needed to make routing and reachability more secure is of a kind that many vendors already provide to their customers. Second, making large multinational ISPs accountable presupposes a hierarchical structure and trade-offs that would rival WTO agreements in their complexity. Pressuring ISPs for disclosure, on the other hand, is a distributed task that can be shared among many stakeholders. Third, disclosure, unlike accountability, is not an all-or-nothing solution; it can be achieved in very small, incremental stages. Fourth, disclosure – or rather non-disclosure – of crucial operating information (such as what traffic management practices are in use) has become an issue at all levels of the bandwidth hierarchy, from the last mile to the Tier-1 networks. Fifth, stakeholders, advocacy groups and end-users have much to gain from understanding that ample bandwidth alone, i.e. network “speed,” is not a sufficient condition of robust Internet access. Although it is certainly very

important, especially in Canada and the United States, that targets for typical last-mile bandwidth be set far higher than they are today, it is equally important that policymakers place interconnection and interconnection disclosures on the same footing as bandwidth – especially since bandwidth is increasingly treated as an undifferentiated commodity.

There are many problems arising from what is not disclosed about peering agreements, beginning with those entered into by Tier-1 providers. Any entity that peers or interconnects in any way with a Tier-1 network must sign a non-disclosure agreement, under which interconnection and peering arrangements can be cancelled on little notice by larger networks. Moreover, Tier-1 providers and other large networks meet regularly to discuss engineering issues - in private. This ingrained resistance to public disclosure encourages arbitrage; tends to favor transit over peering, despite the technical and public-interest advantages of peering; and solidifies the market power of the largest ISPs.

Although this is not a treatise on contract law, it seems reasonable to suggest that when any firm purchases bandwidth from an ISP, that party should have the right to know not merely what throughput can be sustained, but also how upstream interconnection arrangements will affect their real-world connectivity. Ideally, ISPs should also be required to communicate full and timely information on de-peering to their downstream users, to minimize any harmful impact on

data reachability. Clearly such requirements would apply in quite different ways to a) Tier-1 operators and their peering arrangements; b) Tier-2 and Tier-3 ISPs that purchase transit for some or all of their bandwidth needs; and c) residential end-users. Nevertheless, ISPs should generally bear a responsibility for identifying – at least to their customers - all those networks with which they peer or from which they purchase transit. Being in possession of this information is indispensable if individual ISPs, especially smaller ones, are to have a sense of both control and responsibility over their Internet connectivity.

It is also clear that national regulators could extend these obligations still further, to include, for example, a framework to ensure that networks meeting certain minimum criteria could not be refused a peering arrangement by an upper-tier network, with the onus on the larger network to demonstrate that any refusal was not anti-competitive. On the other hand, there is in existence a wide range of contract models in use by businesses of all types that incorporate a service level agreement (SLA), with appropriate stipulations as to quality of service (QoS). Guarantees related to factors such as uptime are almost exclusively found in business-to-business contracts, since they are much more costly than best-effort retail arrangements. But at the very least, they provide a huge body of practical experience in the provision of connectivity.

It must be recognized that many barriers stand in the way of implementing disclosure practices, even at the national level, under the auspices of an existing regulatory framework. Let us mention two. First, there is a substantial difference between disclosing information about connectivity to a customer on a confidential basis, and disclosing information *publicly* about connectivity arrangements and traffic management practices. Second, in both the Canadian and American jurisdictions, the retail broadband business has been deregulated. This would suggest that the CRTC and FCC would normally find it extremely difficult to re-regulate residential broadband with a view to reining in anti-competitive practices.

The relationship between market power and the ability of ISPs to hide their terms of trade even from their own customers has become a hallmark of the incumbent residential broadband providers in Canada and the US. Indeed, one of the most significant features of the CRTC's ISP hearings, held from July 6 to July 14, 2009, was the sheer amount of information about traffic management practices that was discussed – and how much of it had been withheld from customers by Canada's incumbent broadband ISPs. Michael Geist was among those who found that disclosure of terms and practices by Canada's broadband providers leaves a great deal to be desired:

Each day brought new and surprising revelations about how little ISPs tell their customers about their traffic management practices. By far the most

egregious was Rogers, which admitted that it charges tiered pricing for faster upload speeds but that all tiers were throttled to the same speed when using P2P. In other words, the Extreme subscriber who pays \$59.99 per month and is promised fast upload speeds (1 Mbps) actually gets the same upload speed as the Express subscriber who pays \$46.99 per month and is promised upload speeds of 512 kbps. There were similar stories from many other ISPs, who disclosed actual speeds that bring P2P down to a virtual crawl. *Disclosure has improved over the past year as the issue has gained prominence, but there clearly is a long way to go.* (Geist, 2009b, emphasis added)

A close reading of the transcripts for the seven days of public hearings indicates not only how much information is not disclosed to Canadians about their Internet access, but also how difficult it is even for the Commission to get complete, intelligible answers to the most basic questions about traffic management practices. As discussed above in Chapter IV, one of the most pressing issues about broadband in Canada concerns whether the Commission needs to intervene in the DSL wholesale market to ensure that resellers are not put at a serious competitive disadvantage by the actions of incumbent Bell Canada. During the July hearing, Bell officials were questioned closely about why the company's Gateway Access Service (GAS) is configured in such a way that all the resellers in Ontario and Quebec that depend on it are forced to have their

own customers' data traffic-shaped, in lockstep with Bell's retail Sympatico customers.

This issue first came to light in 2008, and fittingly enough it did so because of a failure of disclosure on Bell's part – i.e., a failure to disclose that it was traffic-shaping downstream to its wholesale customers. This revelation became the subject of a complaint by the Canadian Association of Internet Providers that is still unresolved. But the merits of that particular complaint aside, the July public hearings have raised new questions about the role of incumbent networks in a competitive broadband market, and in particular about both the origins and effects of the kind of traffic-shaping being practised by Bell. On the basis of certain comments made at the hearings (which provide only indirect evidence), we hypothesize that Bell may have deployed MPLS in its networks in such a way as to make it impossible for any of its GAS resellers to offer access to its customers without traffic-shaping as a way to achieve market differentiation. At one point in his questioning of Bell, Commissioner Denton was trying to understand whether there were any circumstances in which the resellers would be able to offer access to their customers and either do their own traffic-shaping, or refrain from it, at their option:

6793 COMMISSIONER DENTON: Well, if they wanted to spend the Cap-Ex, would they be prevented from doing so by your rules?

6794 MR. DANIELS: In terms of if they have their own DPI equipment?

6795 COMMISSIONER DENTON: Yes. If they wanted to go ahead and make the investments to conform to reasonable requirements, what would prevent them from doing so?

6796 MR. DANIELS: Nothing would prevent them, but the truth is we wouldn't -- today, as I say, every single one of them would have to do exactly that investment exactly the same way for it to be reliable, and what's the point of that to ask them all to do that when we're doing it for them?

6797 They can't -- this is the key, they're coming to you and saying, I want to be able to distinguish my traffic and have different ways of going about doing it better, equally reasonable.

6798 They can do that if they buy HSA, they can do that if they unbundle loop, but if they're on a shared GAS network where their traffic is going to impact our retail customers and being priced accordingly, then they can't have a different solution, they'd have to have the exact same solution as us and every single one of them would have to have that exact same

solution as us because we can't distinguish between one ISP who does it and another one who says they're going to do it but they didn't.

6799 COMMISSIONER DENTON: Well, we'll watch that statement -- hold that statement for the future. (CRTC, 2009b)


In our view, when competition in a broadband market is almost entirely dependent on resale of access by non-facilities-based ISPs, it is not good public policy to allow the incumbent to undermine market differentiation, especially by means of service features that provoke customer dissatisfaction and frustration. Looking ahead, however, we sense that a more useful and intriguing challenge lies in trying to determine whether the deployment of MPLS is now occurring this close to the edges of the Internet, in this kind of delicate competitive relationship.

And if that turns out to be the case, it would certainly be worthwhile exploring whether the use of MPLS has any anti-competitive ramifications for downstream users, and whether such ramifications are or are not technically avoidable.

The answers to questions such as these will require a combination of continuing scholarly research, informed public interest advocacy and enlightened policymaking. In the meantime, we will, paraphrasing Commissioner Denton, hold these tasks for the future.

Appendix A – IXmaps

IXmaps is an interactive tool under development that will permit Internet users to see the route that their Internet data packets take across North America. When the user enters a destination URL in their browser they see in an adjacent window, a map of North America on which the packets' routes will be displayed. Each 'exchange point' or 'carrier hotel' along that route is displayed in Google Earth as a labeled highlighted icon. When the icon is clicked, a balloon appears containing an HTML document with information about that interchange point and the building where it is physically located. This includes a photo of the building, its ownership, ISPs and other facilities within it and known links or alliances among the building owners, clients, and government or corporate entities. Examples are the building at 151 Front Street, Toronto or the multi-story tower called One Wilshire in Los Angeles. Clickable icons are color-coded and an overlay with a legend is shown which contains a brief summary of the route and exchange points traversed. In the Google Earth 'Places panel', an ordered list of traceroute hops with IP addresses and AS numbers is shown.



IXmaps employs a unique traceroute visualization from the user's system to the destination (displayed in Google Earth) and presents unique information about Internet exchange points transited along the way. These developments are being implemented with funding from the SSHRC ITST program. IXmaps is part of the New Transparency Project at the Faculty of Information, University of Toronto. The project team consists of Dr. David J. Phillips, Associate Professor and Chair of Doctoral Studies, Faculty of Information and Dr. Andrew Clement, Professor, Faculty of Information and Coordinator, Information Policy Research Program as well as Nancy Paterson, PhD candidate Communication & Culture, YorkU and Associate Professor, Ontario College of Art & Design.

Appendix B – Tier-1 Networks

From Wikipedia, the free encyclopedia -
http://en.wikipedia.org/wiki/Tier_1_network

Jump to: navigation, search

This article may contain original research or unverified claims. Please improve the article by adding references. See the talk page for details. (July 2008)

A Tier-1 Network is an IP network (typically but not necessarily an Internet Service Provider) which connects to the entire Internet solely via Settlement Free Interconnection, also known as settlement free peering.

Contents:

1. Definition
2. Politics
3. Routing issues
4. Marketing issues
5. Global issues
6. Telecom Providers Tier-1 & 2
7. See also
8. References

Definition

Although there is no authority which has defined the "tiers" of Internet networks, the most common definition is:

- A network that can reach every other network on the Internet without purchasing IP transit or paying settlements.⁽¹⁾

By this definition, a Tier-1 Network is a Transit-Free network. But not all Transit-Free Networks are Tier-1 Networks. It is possible to become transit free by paying for peering or agreeing to settlements.

It is trivial to objectively prove (or disprove) a network is transit free. The fourteen (14) networks listed below, and only those fourteen, are transit free (as of July 2008). The most widely quoted source is Renesys Corporation, but the base information to prove the claim is publicly accessible from many locations, such as the RIPE RIS database, the Oregon Route Views servers, the Packet Clearing House, and others.

It is impossible for an outside authority to confirm that a network is not paying settlements of any type because such business agreements are frequently not public information, or even covered under a Non-Disclosure Agreement. The information presented here is the best collective knowledge of the Internet peering community. There is little disagreement amongst the community itself,

even though there is no quotable source for the information. (For clarity, here we will define the "peering community" as the set of peering coordinators for networks which are present at Internet Exchanges on at least two continents.) It is commonly believed [citation?] that observing this definition strictly would result in every network being disqualified. For instance, many large telephone companies who are also Tier-1 Networks buy, sell, or swap fiber amongst themselves. Even if it were possible to list every transaction, it is not possible to know if some of those transactions were required for or in payment of a peering connection.

As a result, the term Tier-1 Network is used in the industry to mean a network with no overt settlements. An overt settlement would be a monetary charge for the amount, direction, or type of traffic sent between networks.

Common definitions of Tier-2 and Tier-3 networks:

- Tier-2 - A network that peers with some networks, but still purchases IP transit or pays settlements to reach at least some portion of the Internet.
- Tier-3 - A network that solely purchases transit from other networks to reach the Internet.

Politics

There are many reasons why networking professionals use the "Tier Hierarchy" to describe networks, but the most important one is better understanding of a particular network's political and economic motivations in relationship to how and with whom it peers.

By definition, a Tier-1 network does not purchase IP transit from any other network or pay settlements to any other network to reach any other portion of the Internet. Therefore, in order to be a Tier-1, a network must peer with every other Tier-1 network. A new network cannot become a Tier-1 without the implicit approval of every other Tier-1 network, since any one network's refusal to peer with it will prevent the new network from being considered a Tier-1.

Routing issues

Because a Tier-1 does not have any alternate transit paths, Internet traffic between any two Tier-1 networks is critically dependent on the peering relationship. If two Tier-1 networks arrive at an impasse and discontinue peering with each other (usually in a unilateral decision by one side), single-homed customers of each network will not be able to reach the customers of the other network. This effectively "partitions" the Internet, so that one portion cannot talk to another portion, which has happened several times during the history of the Internet. Those portions of the Internet typically remain partitioned until one side purchases transit (thus losing its "Tier-1" status), or until the collective pain of the outage and/or threat of litigation motivates the two networks to resume voluntary peering.

It is important to remark here that Tier-2 (and lower) ISPs and their customers are normally unaffected by these partitions because they can have traffic with more than one tier-1 provider.

Marketing issues

Because there is no formal definition or authoritative body which determines who is and is not a Tier-1, the term is often misused as a marketing slogan rather than an accurate technical description of a network. Frequent misconceptions of the "tier hierarchy" include:

- Tier-1 networks are closer to the "center" of the Internet.
In reality, Tier-1 networks usually have only a small number of peers (typically only other Tier-1s and very large Tier-2s), while Tier-2 networks are motivated to peer with many other Tier-2 and end-user networks. Thus a Tier-2 network with good peering is frequently much "closer" to most end-users or content than a Tier-1.
- Tier-1 networks by definition offer "better" quality Internet connectivity.
By definition, there are networks which Tier-1 networks have only one path to, and if they lose that path, they have no "backup transit" which would preserve their full connectivity.
Some Tier-2 networks are significantly larger than some Tier-1 networks, and are often able to provide more or better connectivity.
- Tier-2 networks are "resellers" of Tier-1 networks.
Only Tier-3 networks (who provide Internet access) are true "resellers", while many large Tier-2 networks peer with the majority or even vast majority of the Internet directly except for a small portion of the Internet which is reached via a transit provider.

Because the "tier" ranking system is used in marketing and sales, a long-held though generally misguided view among customers is that they should "only purchase from a Tier-1". Because of this, many networks claim to be Tier-1 even though they are not, while honest networks may lose business to those who only wish to purchase from a Tier-1. The frequent misuse of the term has led to a corruption of the meaning, whereby almost every network claims to be a Tier-1 even though it is not. The issue is further complicated by the almost universal use of non-disclosure agreements among Tier-1 networks, which prevent the disclosure of details regarding their settlement-free interconnections. Some of the incorrect measurements which are commonly cited include numbers of routers, route miles of fiber optic cable, or number of customers using a particular network. These are all valid ways to measure the size, scope, capacity, and importance of a network, but they have no direct relationship to Tier-1 status. Another common area of debate is whether it is possible to become a Tier-1 through the purchase of "paid peering", or settlement-based interconnections, whereby a network "buys" the status of Tier-1 rather than achieving it through settlement-free means. While this may simulate the routing behaviors of a Tier-1

network, it does not simulate the financial or political peering motivations, and is thus considered by most Peering Coordinators to not be a true Tier-1 for most discussions.

Global issues

See also: Internet Exchange Point

A common point of contention among people discussing Tier-1 networks is the concept of a "regional Tier-1". A regional Tier-1 network is a network which is not transit free globally, but which maintains many of the classic behaviors and motivations of a Tier-1 network within a specific region.

A typical scenario for this behavior involves a network that was the incumbent telecommunications company in a specific country or region, usually tied to some level of government-supported monopoly. Within their specific countries or regions of origin, these networks maintain peering policies which mimic those of Tier-1 networks (such as lack of openness to new peering relationships and having existing peering with every other major network in that region). However, this network may then extend to another country, region, or continent outside of its core region of operations, where it may purchase transit or peer openly like a Tier-2 network.

A commonly cited example of these behaviors involves the incumbent carriers within Australia, who will not peer with new networks in Australia under any circumstances, but who will extend their networks to the United States and peer openly with many networks. Less extreme examples of much less restrictive peering requirements being set for regions in which a network peers, but does not sell services or have a significant market share, are relatively common among many networks, not just "regional Tier-1"s.

While the classification of "regional Tier-1" does hold some merit for understanding the peering motivations of such a network within different regions, these networks do not meet the requirements of a true global Tier-1 because they are not transit free globally.

Telecom Providers Tier-1 & 2

The original Internet backbone was the ARPANET. It was replaced in 1989 by the NSFNET backbone. This was similar to a Tier-1 backbone. The Internet could be defined as anything able to send datagrams to this backbone.

When the Internet went private, a new network architecture based on decentralized routing (EGP/BGP) was developed. The Tier-1 ISPs and the peer connections made the NSFNET redundant and later obsolete. On April 30, 1995, the NSFNET backbone was shut down.

Currently, Tier-1 ISPs form the closest thing to a backbone.

This section needs additional citations for verification. Please help improve this article by adding reliable references (ideally, using inline citations). Unsourced material may be challenged and removed. (March 2008)

The following 10 networks are believed to be Tier-1 Networks (i.e. they do not have an overt settlement on any peering link with any other network) by the overwhelming majority of the peering community.

Name	AS#	Sept, 2007 (2) (3) degree	Peering policy
AT&T	7018	1382	AT&T Peering policy
Global Crossing (GBLX)	3549	499	Peering policy (2003)
Level 3(L3)	3356		
NTT Communications (Verio)(AS2914 was originally TLGnet; merged in 1999 after assets were bought)	2914	254	NTT Communications Routing Policy
Qwest	209	828	North America; Intl
Sprint	1239	880	
Tata Communications (formerly Teleglobe)	6453		
Verizon Business formerly UUNET	701	1452	Verizon UUNET Peering
SAVVIS	3561		
TeliaSonera International Carrier	1299		TSIC Peering Policy

Most Tier-1 networks are headquartered in the United States, except for Global Crossing, which is headquartered in Hamilton, Bermuda, TeliaSonera which is headquartered in Stockholm, Sweden and NTT, which purchased the US network Verio to become a Tier-1 Network and is headquartered in Tokyo, Japan. (NTT is partially owned by the Japanese government.)

The following networks were Tier-1 Networks and may still be, but there is some question in the community as to whether they are now paying settlements to one or more of their peers.

Name	AS#	Sept, 2007 (2) (3) degree	Peering policy
AOL Transit Data Network (ATDN)	1668		ATDN Peering Policy

The following networks are Transit-Free Networks, even though they have settlement based or paid peering with one or more other networks:

Name	AS#	Sept, 2007 (2) (3) degree	Settlement Peer
AboveNet	6461		Sprint/AS1239
Cogent Communications	174		Sprint/AS1239 & possibly Level 3 Communications (L3)/AS3356
XO Communications	2828		Sprint/AS1239 & Level 3 (L3)/AS3356

Due to the marketing considerations mentioned above, many people mistakenly believe that other networks are Tier-1 when they are not. Because of this, many online resources and forums incorrectly list several non-qualifying networks as Tier-1.

Below is a list of some of these Tier-2 networks which are often listed as Tier-1, along with their upstream providers:

- Allstream/AS15290 (Verizon Business/ AS701 transit, AT&T/ AS7018 transit, Level 3 Communications (L3)/ AS3356 transit)
- British Telecom/ AS5400 (Global Crossing (GBLX)/ AS549 transit, Level 3 Communications (L3)/ AS 356 transit, Sprint Nextel Corporation/AS1239 transit)
- Cable and Wireless/ AS273 (Level 3 Communications (L3)/ AS 356, SAVVIS/ AS561 transit)
- Deutsche Telekom/ AS320 (Sprint Nextel Corporation/ AS 239 transit)
- France Telecom/ AS511 aka OpenTransit (Sprint Nextel Corporation/ AS1239 transit)

- Hurricane Electric/ AS6939 (Global Crossing (GBLX)/ AS3549 transit, TeliaSonera/ AS1299 transit)
- PCCWGlobal/ AS3491 (Global Crossing (GBLX)/ AS3549 transit)
- Tele2/ AS1257 (Sprint Nextel Corporation/ AS1239 transit)
- Time Warner Telecom/ AS4323 (Sprint Nextel Corporation/ AS1239 transit)
- Tiscali International Network (TINet)/ AS3257 (Sprint Nextel Corporation/ AS1239 transit; Verizon Business (AS701) transit)

See also

- Tier-2 network
- Peering
- Internet transit
- Network access point

References

1. "How the 'Net works: an introduction to peering and transit: Page 4". <http://arstechnica.com/guides/other/peering-and-transit.ars/4>. Retrieved on 2008-11-04.
2. a b c CAIDA AS ranking
3. a b c Visualizing Internet Topology at a Macroscopic Scale April 2005

Appendix C – Sample Weekly Routing Table Report

North American Network Operators Group

Date Prev | Date Next | Date Index | Thread Index | Author Index | Historical

Weekly Routing Table Report

Date: Sat, 13 Jun 2009 04:11:14 +1000 (EST)
 From: Routing Analysis Role Account <cscora@apnic.net>
 To: apops@apops.net, nanog@nanog.org, routing-wg@ripe.net,
 afnog@afnog.org, ausnog@ausnog.net, sanog@sanog.org
 Reply-to: pfs@cisco.com
 Subject: Weekly Routing Table Report

*This is an automated weekly mailing describing the state of the Internet
 Routing Table as seen from APNIC's router in Japan.*

Daily listings are sent to bgp-stats@lists.apnic.net

For historical data, please see <http://thyme.apnic.net>.

If you have any comments please contact Philip Smith <pfs@cisco.com>.

Routing Table Report 04:00 +10GMT Sat 13 Jun, 2009

Report Website: <http://thyme.apnic.net>

Detailed Analysis: <http://thyme.apnic.net/current/>

Analysis Summary

BGP routing table entries examined:	287824
Prefixes after maximum aggregation:	137147
Deaggregation factor:	2.10
Unique aggregates announced to Internet:	142799
<i>Total ASs present in the [global] Internet Routing Table:</i>	<i>31449</i>
Prefixes per ASN:	9.15
Origin-only ASes present in the Internet Routing Table:	27360
Origin ASes announcing only one prefix:	13306
Transit ASes present in the Internet Routing Table:	4089
Transit-only ASes present in the Internet Routing Table:	96
Average AS path length visible in the Internet Routing Table:	3.6
Max AS path length visible:	29
Max AS path prepend of ASN (3816)	22
Prefixes from unregistered ASNs in the Routing Table:	500

Unregistered ASNs in the Routing Table:	147
Number of 32-bit ASNs allocated by the RIRs:	181
Prefixes from 32-bit ASNs in the Routing Table:	49
Special use prefixes present in the Routing Table:	0
Prefixes being announced from unallocated address space:	829
Number of addresses announced to Internet:	2050857808
Equivalent to 112 /8s, 235 /16s and 239 /24s	
Percentage of available address space announced:	55.3
Percentage of allocated address space announced:	64.0
Percentage of available address space allocated:	86.4
Percentage of address space in use by end-sites:	77.4
Total number of prefixes smaller than registry allocations:	142267

APNIC Region Analysis Summary

Prefixes being announced by APNIC Region ASes:	68490
Total APNIC prefixes after maximum aggregation:	24490
APNIC Deaggregation factor:	2.80
Prefixes being announced from the APNIC address blocks:	67896
Unique aggregates announced from the APNIC address blocks:	30776
APNIC Region origin ASes present in the Internet Routing Table:	3655
APNIC Prefixes per ASN:	18.53
APNIC Region origin ASes announcing only one prefix:	997
APNIC Region transit ASes present in the Internet Routing Table:	518
Average APNIC Region AS path length visible:	3.5
Max APNIC Region AS path length visible:	18
Number of APNIC addresses announced to Internet:	453620336
Equivalent to 21 /8s, 241 /16s and 36 /24s	
Percentage of available APNIC address space announced:	84.5

APNIC AS Blocks 4608-4864, 7467-7722, 9216-10239, 17408-18431
(pre-ERX allocations) 23552-24575, 37888-38911, 45056-46079
APNIC Address Blocks 58/8, 59/8, 60/8, 61/8, 110/8, 111/8, 112/8,
113/8, 114/8, 115/8, 116/8, 117/8, 118/8, 119/8,
120/8, 121/8, 122/8, 123/8, 124/8, 125/8, 126/8,
180/8, 183/8, 202/8, 203/8, 210/8, 211/8, 218/8,
219/8, 220/8, 221/8, 222/8

ARIN Region Analysis Summary

Prefixes being announced by ARIN Region ASes:	123257
Total ARIN prefixes after maximum aggregation:	65969

ARIN Deaggregation factor:	1.87
Prefixes being announced from the ARIN address blocks:	124046
Unique aggregates announced from the ARIN address blocks:	51834
ARIN Region origin ASs present in the Internet Routing Table:	13033
ARIN Prefixes per ASN:	9.52
ARIN Region origin ASs announcing only one prefix:	4995
ARIN Region transit ASs present in the Internet Routing Table:	1276
Average ARIN Region AS path length visible:	3.3
Max ARIN Region AS path length visible:	24
Number of ARIN addresses announced to Internet:	1009227328
Equivalent to 21 /8s, 6 /16s and 109 /24s	
Percentage of available ARIN address space announced:	194.0

ARIN AS Blocks 1-1876, 1902-2042, 2044-2046, 2048-2106
 (pre-ERX allocations) 2138-2584, 2615-2772, 2823-2829, 2880-3153
 3354-4607, 4865-5119, 5632-6655, 6912-7466
 7723-8191, 10240-12287, 13312-15359, 16384-17407
 18432-20479, 21504-23551, 25600-26591,
 26624-27647, 29696-30719, 31744-33791
 35840-36863, 39936-40959, 46080-47103
 53248-55295

ARIN Address Blocks 24/8, 63/8, 64/8, 65/8, 66/8, 67/8, 68/8,
 69/8, 70/8, 71/8, 72/8, 73/8, 74/8, 75/8,
 76/8, 96/8, 97/8, 98/8, 99/8, 108/8, 173/8,
 174/8, 184/8, 199/8, 204/8, 205/8, 206/8, 207/8,
 208/8, 209/8, 216/8,

RIPE Region Analysis Summary

Prefixes being announced by RIPE Region ASes:	65736
Total RIPE prefixes after maximum aggregation:	38944
RIPE Deaggregation factor:	1.69
Prefixes being announced from the RIPE address blocks:	64869
Unique aggregates announced from the RIPE address blocks:	43770
RIPE Region origin ASs present in the Internet Routing Table:	13116
RIPE Prefixes per ASN:	4.95
RIPE Region origin ASs announcing only one prefix:	6865
RIPE Region transit ASs present in the Internet Routing Table:	1966
Average RIPE Region AS path length visible:	4.0
Max RIPE Region AS path length visible:	28
Number of RIPE addresses announced to Internet:	482559648
Equivalent to 21 /8s, 221 /16s and 2 /24s	
Percentage of available RIPE address space announced:	102.7

RIPE AS Blocks 1877-1901, 2043, 2047, 2107-2136, 2585-2614
 (pre-ERX allocations) 2773-2822, 2830-2879, 3154-3353, 5377-5631
 6656-6911, 8192-9215, 12288-13311, 15360-16383
 20480-21503, 24576-25599, 28672-29695
 30720-31743, 33792-35839, 38912-39935
 40960-45055, 47104-52223

RIPE Address Blocks 62/8, 77/8, 78/8, 79/8, 80/8, 81/8, 82/8,
 83/8, 84/8, 85/8, 86/8, 87/8, 88/8, 89/8,
 90/8, 91/8, 92/8, 93/8, 94/8, 95/8, 109/8,
 178/8, 193/8, 194/8, 195/8, 212/8, 213/8, 217/8,

LACNIC Region Analysis Summary

Prefixes being announced by LACNIC Region ASes:	23862
Total LACNIC prefixes after maximum aggregation:	5946
LACNIC Deaggregation factor:	4.01
Prefixes being announced from the LACNIC address blocks:	23726
Unique aggregates announced from the LACNIC address blocks:	13297
LACNIC Region origin ASs present in the Internet Routing Table:	1119
LACNIC Prefixes per ASN:	21.20
LACNIC Region origin ASs announcing only one prefix:	361
LACNIC Region transit ASs present in the Internet Routing Table:	184
Average LACNIC Region AS path length visible:	4.0
Max LACNIC Region AS path length visible:	29
Number of LACNIC addresses announced to Internet:	71622272
Equivalent to 3 /8s, 50 /16s and 72 /24s	
Percentage of available LACNIC address space announced:	71.2

LACNIC AS Blocks 26592-26623, 27648-28671, plus ERX transfers
 LACNIC Address Blocks 186/8, 187/8, 189/8, 190/8, 200/8, 201/8,

AfriNIC Region Analysis Summary

Prefixes being announced by AfriNIC Region ASes:	6035
Total AfriNIC prefixes after maximum aggregation:	1453
AfriNIC Deaggregation factor:	4.15
Prefixes being announced from the AfriNIC address blocks:	6445
Unique aggregates announced from the AfriNIC address blocks:	2468
AfriNIC Region origin ASs present in the Internet Routing Table:	296
AfriNIC Prefixes per ASN:	21.77
AfriNIC Region origin ASs announcing only one prefix:	88

AfriNIC Region transit ASs present in the Internet Routing Table: 59
 Average AfriNIC Region AS path length visible: 3.8
 Max AfriNIC Region AS path length visible: 15
 Number of AfriNIC addresses announced to Internet: 19460608
 Equivalent to 0 /8s, 187 /16s and 219 /24s
 Percentage of available AfriNIC address space announced: 58.0

AfriNIC AS Blocks 36864-37887 & ERX transfers
 AfriNIC Address Blocks 41/8, 197/8,

APNIC Region per AS prefix count summary

ASN	No of nets	/20 equiv	MaxAgg	Description
4766	1709	6931	402	Korea Telecom (KIX)
17488	1600	130	101	Hathway IP Over Cable Internet
4755	1251	362	128	TATA Communications formerly
9583	1100	87	549	Sify Limited
4134	892	16925	377	CHINANET-BACKBONE
7545	793	198	101	TPG Internet Pty Ltd
23577	780	34	664	Korea Telecom (ATM-MPLS)
18101	751	216	30	Reliance Infocom Ltd Internet
24560	714	230	174	Bharti Airtel Ltd.
9829	683	572	18	BSNL National Internet Backbone

Complete listing at <http://thyme.apnic.net/current/data-ASnet-APNIC>

ARIN Region per AS prefix count summary

ASN	No of nets	/20 equiv	MaxAgg	Description
6389	4290	3647	324	bellsouth.net, inc.
4323	1863	1035	376	Time Warner Telecom
1785	1687	717	138	PaeTec Communications, Inc.
20115	1627	1447	733	Charter Communications
7018	1506	5924	1042	AT&T WorldNet Services
6478	1375	305	476	AT&T Worldnet Services
2386	1264	683	918	AT&T Data Communications Serv
3356	1204	10980	452	Level 3 Communications, LLC
11492	1114	208	12	Cable One
18566	1062	296	10	Covad Communications

Complete listing at <http://thyme.apnic.net/current/data-ASnet-ARIN>

RIPE Region per AS prefix count summary

ASN	No of nets	/20 equiv	MaxAgg	Description
3292	454	1902	392	TDC Tele Danmark
12479	450	578	6	Uni2 Autonomous System
702	43	1861	347	UUNET - Commercial IP

				service
30890	413	87	193	Evolva Telecom
35805	358	24	4	United Telecom of Georgia
8866	351	109	21	Bulgarian Telecommunication C
3301	344	1684	307	TeliaNet Sweden
3215	343	3041	108	France Telecom Transpac
3320	338	7066	296	Deutsche Telekom AG
9121	319	1442	25	TTnet Autonomous System

Complete listing at <http://thyme.apnic.net/current/data-ASnet-RIPE>

LACNIC Region per AS prefix count summary

ASN	No of nets	/20 equiv	MaxAgg	Description
8151	1466	2879	233	UniNet S.A. de C.V.
10620	906	206	115	TVCABLE BOGOTA
22047	591	302	14	VTR PUNTO NET S.A.
7303	56	298	85	Telecom Argentina Stet-France
28573	538	563	35	NET Servicios de Comunicacao S.A
11830	486	292	54	Instituto Costarricense de EI
6471	443	96	32	ENTEL CHILE S.A.
11172	443	102	70	Servicios Alestra S.A de C.V

7738	404	794	28	Telecomunicacoes da Bahia S.A
3816	363	188	81	Empresa Nacional de Telecomun

Complete listing at <http://thyme.apnic.net/current/data-ASnet-LACNIC>

AfriNIC Region per AS prefix count summary

ASN	No of nets	/20 equiv	MaxAgg	Description
8452	981	188	7	TEDATA
24863	884	82	40	LINKdotNET AS number
20858	324	34	5	EgyNet
3741	277	856	237	The Internet Solution
2018	243	215	143	Tertiary Education Network
6713	160	151	12	Itissalat Al-MAGHRIB
33783	152	10	8	EEPAD TISP TELECOM & INTERNET
29571	139	15	8	Ci Telecom Autonomous system
5536	123	8	9	Internet Egypt Network
5713	115	507	66	Telkom SA Ltd

Complete listing at <http://thyme.apnic.net/current/data-ASnet-AFRINIC>

Global Per AS prefix count summary

ASN	No of nets	/20 equiv	MaxAgg	Description
6389	4290	3647	324	bellsouth.net, inc.
4323	1863	1035	376	Time Warner Telecom
4766	1709	6931	402	Korea Telecom (KIX)
1785	1687	717	138	PaeTec Communications, Inc.
20115	1627	1447	733	Charter Communications
17488	1600	130	101	Hathway IP Over Cable Interne
7018	1506	5924	1042	AT&T WorldNet Services
8151	1466	2879	233	UniNet S.A. de C.V.
6478	1375	305	476	AT&T Worldnet Services
2386	1264	683	918	AT&T Data Communications Serv

Complete listing at <http://thyme.apnic.net/current/data-ASnet>

Global Per AS Maximum Aggr summary

ASN	No of nets	Net Savings	Description
1785	1687	1549	PaeTec Communications, Inc.
17488	1600	1499	Hathway IP Over Cable Interne
4323	1863	1487	Time Warner Telecom
4766	1709	1307	Korea Telecom (KIX)

8151	1466	1233	UniNet S.A. de C.V.
4755	1251	1123	TATA Communications formerly
11492	1114	1102	Cable One
18566	1062	1052	Covad Communications
22773	1062	996	Cox Communications, Inc.
8452	981	974	TEDATA

Complete listing at <http://thyme.apnic.net/current/data-CIDRnet>

List of Unregistered Origin ASNs (Global)

Bad AS	Designation	Network	Transit AS	Description
16927	UNALLOCATED	12.0.252.0/23	7018	AT&T WorldNet Service
15132	UNALLOCATED	12.9.150.0/24	7018	AT&T WorldNet Service
32567	UNALLOCATED	12.14.170.0/24	7018	AT&T WorldNet Service
13746	UNALLOCATED	12.24.56.0/24	7018	AT&T WorldNet Service
32567	UNALLOCATED	12.25.107.0/24	7018	AT&T WorldNet Service
26973	UNALLOCATED	12.39.152.0/24	7018	AT&T WorldNet Service
26973	UNALLOCATED	12.39.154.0/23	7018	AT&T WorldNet Service
26973	UNALLOCATED	12.39.159.0/24	7018	AT&T WorldNet Service
32326	UNALLOCATED	12.40.49.0/24	7018	AT&T WorldNet Service
25639	UNALLOCATED	12.41.169.0/24	7018	AT&T WorldNet Service

Complete listing at <http://thyme.apnic.net/current/data-badAS>

Advertised Unallocated Addresses

Network	Origin AS	Description
41.223.112.0/22	5713	Telkom SA Ltd
41.223.176.0/22	36981	>>UNKNOWN<<
41.223.188.0/24	22351	Intelsat
41.223.189.0/24	26452	Local Communications Networks
62.61.220.0/24	24974	Tachyon Europe BV - Wireless
62.61.221.0/24	24974	Tachyon Europe BV - Wireless
63.140.213.0/24	22555	Universal Talkware Corporatio
63.143.251.0/24	22555	Universal Talkware Corporatio
64.31.32.0/19	11955	ServiceCo LLC - Road Runner
64.31.59.0/24	7017	ServiceCo LLC - Road Runner

Complete listing at <http://thyme.apnic.net/current/data-add-IANA>

Number of prefixes announced per prefix length (Global)

/1:0 /2:0 /3:0 /4:0 /5:0 /6:0
 /7:0 /8:19 /9:10 /10:20 /11:58 /12:166
 /13:348 /14:601 /15:1154 /16:10514 /17:4723 /18:8087
 /19:16903 /20:20163 /21:19974 /22:25876 /23:25740 /24:150868
 /25:863 /26:1030 /27:538 /28:148 /29:8 /30:5
 /31:0 /32:8

Advertised prefixes smaller than registry allocations

ASN	No of nets	Total ann.	Description
6389	2796	4290	bellsouth.net, inc.
4766	140	1709	Korea Telecom (KIX)

17488	1313	1600	Hathway IP Over Cable Interne
1785	1163	1687	PaeTec Communications, Inc.
11492	1044	1114	Cable One
18566	1043	1062	Covad Communications
2386	977	1264	AT&T Data Communications Serv
4323	953	1863	Time Warner Telecom
9583	951	1100	Sify Limited
8452	914	981	TEDATA

Complete listing at <http://thyme.apnic.net/current/data/sXXas-nos>

Number of /24s announced per /8 block (Global)

4:13	8:206	12:2245	13:10	15:19	16:2
17:4	20:36	24:1081	32:52	34:2	38:571
40:97	41:1701	43:1	44:2	47:22	52:4
55:2	56:3	57:24	58:570	59:640	60:459
61:1073	62:1107	63:2013	64:3718	65:2384	66:3622
67:1623	68:824	69:2643	70:557	71:174	72:1676
73:2	74:1572	75:169	76:311	77:849	78:549
79:353	80:971	81:833	82:560	83:434	84:634
85:1030	86:397	87:650	88:354	89:1445	90:57
91:2312	92:344	93:1055	94:1228	95:1031	96:140
97:217	98:245	99:22	109:1	110:143	111:9
112:136	113:125	114:256	115:326	116:1163	117:530
118:288	119:692	120:145	121:754	122:1031	123:694
124:994	125:1336	128:224	129:236	130:127	131:408
132:74	133:9	134:186	135:38	136:224	137:153
138:161	139:78	140:433	141:119	142:385	143:345
144:362	145:48	146:380	147:164	148:519	149:233
150:177	151:190	152:147	153:141	154:2	155:276
156:170	157:302	158:115	159:312	160:284	161:151
162:271	163:169	164:482	165:502	166:275	167:359

168:674	169:164	170:471	171:39	172:10	173:287
174:226	178:1	180:1	183:1	186:22	187:96
188:22	189:432	190:2745	192:5780	193:4256	194:3307
195:2696	196:1101	198:3598	199:3360	200:5086	201:1278
202:7900	203:8170	204:3879	205:2153	206:2441	207:2731
208:3883	209:3400	210:2681	211:1122	212:1604	213:1657
214:80	215:31	216:4534	217:1306	218:387	219:443
220:1223	221:488	222:304			

End of report

- Prev by Date: Re: SBCglobal routing loop.
- Date Index
- Thread Index
- Author Index
- Historical

Appendix D – Autonomous Systems

BEST CURRENT PRACTICE

J. Hawkinson, BBN Planet

T. Bates, MCI

Network Working Group

Request for Comments: 1930

BCP: 6

Category: Best Current Practice

March 1996

Guidelines for creation, selection, and registration of an Autonomous System (AS)

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Abstract

This memo discusses when it is appropriate to register and utilize an Autonomous System (AS), and lists criteria for such. autonomous systems are the unit of routing policy in the modern world of exterior routing, and are specifically applicable to protocols like EGP (Exterior Gateway Protocol, now at historical status; see [EGP]), BGP (Border Gateway Protocol, the current de facto standard for inter-AS routing; see [BGP-4]), and IDRIP (The OSI Inter-Domain Routing Protocol, which the Internet is expected to adopt when BGP becomes obsolete; see [IDRP]). It should be noted that the IDRIP equivalent of an AS is the RDI, or Routing Domain Identifier.

Table of Contents

1.	Introduction	2
2.	Motivation	2
3.	Definitions	2
4.	Common errors in allocating autonomous systems	5
5.	Criteria for the decision -- do I need an AS	5
5.1	Sample Cases	6
5.2	Other Factors	7
6.	Speculation	7
7.	One prefix, one origin AS	8
8.	IGP issues	8
9.	AS Space exhaustion	8
10.	Reserved AS Numbers	9
11.	Security Considerations	9
12.	Acknowledgments	9
13.	References	9
14.	Authors' Addresses	10

1. Introduction

This memo discusses when it is appropriate to register and utilize an Autonomous System (AS), and lists criteria for such. Autonomous systems are the unit of routing policy in the modern world of exterior routing, and are specifically applicable to protocols like EGP (Exterior Gateway Protocol, now at historical status; see [EGP]), BGP (Border Gateway Protocol, the current de facto standard for inter-AS routing; see [BGP-4]), and IDRP (The OSI Inter-Domain Routing Protocol, which the Internet is expected to adopt when BGP becomes obsolete; see [IDRP]). It should be noted that the IDRP equivalent of an AS is the RDI, or Routing Domain Identifier.

2. Motivation

This memo is aimed at network operators and service providers who need to understand under what circumstances they should make use of an AS. It is expected that the reader is familiar with routing protocols and will be someone who configures and operates Internet networks. Unfortunately, there is a great deal of confusion in how autonomous systems should be used today; this memo attempts to clear up some of this confusion, as well as acting as a simple guide to today's exterior routing.

3. Definitions

This document refers to the term "prefix" throughout. In the current classless Internet (see [CIDR]), a block of class A, B, or C networks may be referred to by merely a prefix and a mask, so long as such a block of networks begins and ends on a power-of-two boundary. For example, the networks:

```
192.168.0.0/24
192.168.1.0/24
192.168.2.0/24
192.168.3.0/24
```

can be simply referred to as:

```
192.168.0.0/22
```

The term "prefix" as it is used here is equivalent to "CIDR block", and in simple terms may be thought of as a group of one or more networks. We use the term "network" to mean classful network, or "A, B, C network".

The definition of AS has been unclear and ambiguous for some time. [BGP-4] states:

RFC 1930 Guidelines for creation of an AS March 1996

The classic definition of an Autonomous System is a set of routers under a single technical administration, using an interior gateway protocol and common metrics to route packets within the AS, and using an exterior gateway protocol to route packets to other autonomous systems. Since this classic definition was developed, it has become common for a single AS to use several interior gateway protocols and sometimes several sets of metrics within an AS. The use of the term Autonomous System here stresses the fact that, even when multiple IGPs and metrics are used, the administration of an AS appears to other autonomous systems to have a single coherent interior routing plan and presents a consistent picture of what networks are reachable through it.

To rephrase succinctly:

An AS is a connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy.

Routing policy here is defined as how routing decisions are made in the Internet today. It is the exchange of routing information between autonomous systems that is subject to routing policies. Consider the case of two autonomous systems, X and Y exchanging routing information:

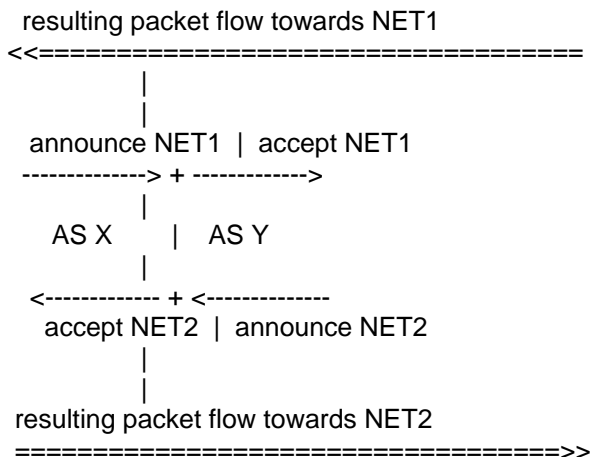
NET1 ASX <---> ASY NET2

ASX knows how to reach a prefix called NET1. It does not matter whether NET1 belongs to ASX or to some other AS which exchanges routing information with ASX, either directly or indirectly; we just assume that ASX knows how to direct packets towards NET1. Likewise ASY knows how to reach NET2.

In order for traffic from NET2 to NET1 to flow between ASX and ASY, ASX has to announce NET1 to ASY using an exterior routing protocol; this means that ASX is willing to accept traffic directed to NET1 from ASY. Policy comes into play when ASX decides to announce NET1 to ASY.

For traffic to flow, ASY has to accept this routing information and use it. It is ASY's privilege to either use or disregard the information that it receives from ASX about NET1's reachability. ASY might decide not to use this information if it does not want to send traffic to NET1 at all or if it considers another route more appropriate to reach NET1.

In order for traffic in the direction of NET1 to flow between ASX and ASY, ASX must announce that route to ASY and ASY must accept it from ASX:



Ideally, though seldom practically, the announcement and acceptance policies of ASX and ASY are symmetrical.

In order for traffic towards NET2 to flow, announcement and acceptance of NET2 must be in place (mirror image of NET1). For almost all applications connectivity in just one direction is not useful at all. It should be noted that, in more complex topologies than this example, traffic from NET1 to NET2 may not necessarily take the same path as traffic from NET2 to NET1; this is called asymmetrical routing. Asymmetrical routing is not inherently bad, but can often cause performance problems for higher level protocols, such as TCP, and should be used with caution and only when necessary. However, asymmetrical routing may be a requirement for mobile hosts and inherently asymmetrical situations, such as a satellite download and a modem upload connection.

Policies are not configured for each prefix separately but for groups of prefixes. These groups of prefixes are autonomous systems.

An AS has a globally unique number (sometimes referred to as an ASN, or Autonomous System Number) associated with it; this number is used in both the exchange of exterior routing information (between neighboring autonomous systems), and as an identifier of the AS itself.

In routing terms, an AS will normally use one or more interior gateway protocols (IGPs) when exchanging reachability information within its own AS. See "IGP Issues".

4. Common errors in allocating autonomous systems

The term AS is often confused or even misused as a convenient way of grouping together a set of prefixes which belong under the same administrative umbrella, even if within that group of prefixes there are various different routing policies. Without exception, an AS must have only one routing policy.

It is essential that careful consideration and coordination be applied during the creation of an AS. Using an AS merely for the sake of having an AS is to be avoided, as is the worst-case scenario of one AS per classful network (the IDEAL situation is to have one prefix, containing many longer prefixes, per AS). This may mean that some re-engineering may be required in order to apply the criteria and guidelines for creation and allocation of an AS that we list below; nevertheless, doing so is probably the only way to implement the desired routing policy.

If you are currently engineering an AS, careful thought should be taken to register appropriately sized CIDR blocks with your registration authority in order to minimize the number of advertised prefixes from your AS. In the perfect world that number can, and should, be as low as one.

Some router implementations use an AS number as a form of tagging to identify interior as well as exterior routing processes. This tag does not need to be unique unless routing information is indeed exchanged with other autonomous systems. See "IGP Issues".

5. Criteria for the decision -- do I need an AS?

- * Exchange of external routing information

An AS must be used for exchanging external routing information with other autonomous systems through an exterior routing protocol. The current recommended exterior routing protocol is BGP, the Border Gateway Protocol. However, the exchange of external routing information alone does not constitute the need for an AS. See "Sample Cases" below.

- * Many prefixes, one AS

As a general rule, one should try to place as many prefixes as possible within a given AS, provided all of them conform to the same routing policy.

- * Unique routing policy

An AS is only needed when you have a routing policy which is different from that of your border gateway peers. Here routing policy refers to how the rest of the Internet makes routing decisions based on information from your AS. See "Sample Cases" below to see exactly when this criteria will apply.

5.1 Sample Cases

- * Single-homed site, single prefix

A separate AS is not needed; the prefix should be placed in an AS of the provider. The site's prefix has exactly the same routing policy as the other customers of the site's service provider, and there is no need to make any distinction in routing information.

This idea may at first seem slightly alien to some, but it highlights the clear distinction in the use of the AS number as a representation of routing policy as opposed to some form of administrative use.

In some situations, a single site, or piece of a site, may find it necessary to have a policy different from that of its provider, or the rest of the site. In such an instance, a separate AS must be created for the affected prefixes. This situation is rare and should almost never happen. Very few stub sites require different routing policies than their parents. Because the AS is the unit of policy, however, this sometimes occurs.

- * Single-homed site, multiple prefixes

Again, a separate AS is not needed; the prefixes should be placed in an AS of the site's provider.

- * Multi-homed site

Here multi-homed is taken to mean a prefix or group of prefixes which connects to more than one service provider (i.e. more than one AS with its own routing policy). It does not mean a network multi-homed running an IGP for the purposes of resilience.

An AS is required; the site's prefixes should be part of a single AS, distinct from the autonomous systems of its service providers. This allows the customer the ability to have a different representation of policy and preference among the different service providers.

This is ALMOST THE ONLY case where a network operator should create its own AS number. In this case, the site should ensure that it has the necessary facilities to run appropriate routing protocols, such as BGP4.

5.2 Other factors

- * Topology

Routing policy decisions such as geography, AUP (Acceptable Use Policy) compliance and network topology can influence decisions of AS creation. However, all too often these are done without consideration of whether or not an AS is needed in terms of adding additional information for routing policy decisions by the rest of the Internet. Careful consideration should be taken when basing AS creation on these type of criteria.

- * Transition / "future-proofing"

Often a site will be connected to a single service provider but has plans to connect to another at some point in the future. This is not enough of a reason to create an AS before you really need it. The AS number space is finite and the limited amount of re-engineering needed when you connect to another service provider should be considered as a natural step in transition.

- * History

AS number application forms have historically made no reference to routing policy. All too often autonomous systems have been created purely because it was seen as "part of the process" of connecting to the Internet. The document should be used as a reference from future application forms to show clearly when an AS is needed.

6. Speculation

1) If provider A and provider B have a large presence in a geographical area (or other routing domain), and many customers are multi-homed between them, it makes sense for all of those customers to be placed within the same AS. However, it is noted that case should only be looked at if practical to do so and fully coordinated between customers and service providers involved.

2) Sites should not be forced to place themselves in a separate AS just so that someone else (externally) can make AS-based policy decisions. Nevertheless, it may occasionally be necessary to split up an AS or a prefix into two autonomous systems for policy reasons.

Those making external policy may request the network operators make such AS changes, but the final decision is up to those network operators who manage the prefixes in question, as well as the autonomous systems containing them. This is, of course, a trade off -- it will not always be possible to implement all desired routing policies.

7. One prefix, one origin AS

Generally, a prefix can should belong to only one AS. This is a direct consequence of the fact that at each point in the Internet there can be exactly one routing policy for traffic destined to each prefix. In the case of an prefix which is used in neighbor peering between two autonomous systems, a conscious decision should be made as to which AS this prefix actually resides in.

With the introduction of aggregation it should be noted that a prefix may be represented as residing in more than one AS, however, this is very much the exception rather than the rule. This happens when aggregating using the AS_SET attribute in BGP, wherein the concept of origin is lost. In some cases the origin AS is lost altogether if there is a less specific aggregate announcement setting the ATOMIC_AGGREGATE attribute.

8. IGP Issues

As stated above, many router vendors require an identifier for tagging their IGP processes. However, this tag does not need to be globally unique. In practice this information is never seen by exterior routing protocols. If already running an exterior routing protocol, it is perfectly reasonable to use your AS number as an IGP tag; if you do not, choosing from the private use range is also acceptable (see "Reserved AS Numbers"). Merely running an IGP is not grounds for registration of an AS number.

With the advent of BGP4 it becomes necessary to use an IGP that can carry classless routes. Examples include OSPF [OSPF] and ISIS [ISIS].

9. AS Space exhaustion

The AS number space is a finite amount of address space. It is currently defined as a 16 bit integer and hence limited to 65535 unique AS numbers. At the time of writing some 5,100 autonomous systems have been allocated and a little under 600 autonomous systems are actively routed in the global Internet. It is clear that this growth needs to be continually monitored. However, if the criteria applied above are adhered to, then there is no immediate danger of AS space exhaustion. It is expected that IDRP will be deployed before this becomes an issue. IDRP does not have a fixed limit on the size of an RDI.

RFC 1930 Guidelines for creation of an AS March 1996

10. Reserved AS Numbers

The Internet Assigned Numbers Authority (IANA) has reserved the following block of AS numbers for private use (not to be advertised on the global Internet):

64512 through 65535

11. Security Considerations

There are few security concerns regarding the selection of autonomous systems. AS number to owner mappings are public knowledge (in WHOIS), and attempting to change that would serve only to confuse those people attempting to route IP traffic on the Internet.

12. Acknowledgments

This document is largely based on [RIPE-109], and is intended to have a wider scope than purely the RIPE community; this document would not exist without [RIPE-109].

13. References

[RIPE-109]

Bates, T., Lord, A., "Autonomous System Number Application Form & Supporting Notes", RIPE 109, RIPE NCC, 1 March 1994.
URL: <ftp://ftp.ripe.net/ripe/docs/ripe-109.txt>.

[BGP-4]

Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1654, T.J. Watson Research Center, cisco Systems, July 1994.

[EGP]

Mills, D., "Exterior Gateway Protocol Formal Specifications", STD 18, RFC 904, International Telegraph and Telephone Co., April 1984.

[IDRP]

Kunzinger, C., Editor, "OSI Inter-Domain Routing Protocol (IDRP)", ISO/IEC 10747, Work In Progress, October 1993.

[CIDR]

Fuller, V., T. Li, J. Yu, and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, BARRnet, cisco, MERIT, OARnet, September 1993.

[OSPF]

Moy, J., "OSPF Version 2", RFC 1583, March 1994.

[ISIS]

Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Multi-Protocol Environments", RFC 1195, Digital Equipment Corporation, December 1990.

14. Authors' Addresses

John Hawkinson
BBN Planet Corporation
150 CambridgePark Drive
Cambridge, MA 02139
Phone: +1 617 873 3180
EMail: jhawk@bbnplanet.com

Tony Bates
MCI
2100 Reston Parkway
Reston, VA 22094
Phone: +1 703 715 7521
EMail: Tony.Bates@mci.net

References

Aboba, B. & Davies, E.B. (Ed.) (2007, July). *Reflections on Internet transparency*.

Internet Architecture Board. Retrieved online January 16, 2008, from:

<http://www.watersprings.org/pub/rfc/rfc4924.txt>

Anderson, A., John, J., Katz-Bassett, E., Krishnamurthy, A., & Madhyastha, H.

(2008). Studying Black Holes in the Internet with Hubble. *USENIX*

Symposium on Networked Systems Design & Implementation (NSDI). San

Francisco, California. Retrieved online May 1, 2009, from:

<http://www.cs.washington.edu/homes/ethan/hubble/nsdi08/index.html>

ARIN Number Resource Policy Manual. (2009). Retrieved online June 15, 2009,

from: <https://www.arin.net/policy/nrpm.html#five1>

Aronson, J.D. & Cowhey, P. F. (2009). *Transforming global information and*

communication markets: The political economy of innovation. Cambridge:

MIT Press.

Ars Technica. (n.d.). Retrieved online June 5, 2008, from:

<http://arstechnica.com/site/advertise.ars>

Balakrishnan, H. & Feamster, N. (2005) lecture. *Interdomain Internet routing*.

Retrieved online January 20, 2008, from:

<http://pages.cs.wisc.edu/~akella/CS740/F08/740-Papers/hari-bgp-notes.pdf>

Baptista, J. (2009, June 16). Can anyone surf to peking university? Governance

Listserv. Retrieved online June 16, 2009, from:

<http://lists.cpsr.org/lists/arc/governance/2009-06/msg00279.html>

Barrolli, L., Durresi, A., Iyengar, S.S., Kannan, R., & Paruchuri, V. (2004).

Efficient and secure autonomous system based traceback. *Journal of Interconnection Networks*, 5(2), 151-164.

Bellis, M. (2009). Inventors of the modern computer: ARPANET - the first

Internet. Retrieved online July 19, 2009, from:

<http://inventors.about.com/library/weekly/aa091598.htm>

Benkler, Y. (2006). *The wealth of networks*. New Haven: Yale University Press.

Black, U. (2000). *IP routing protocols*. New Jersey: Prentice Hall.

Blake, P. (1999, August 23). Can public peering survive. *Telephony*, 237(8).

Supplement. p.14. Retrieved online September 1, 2008, from:

http://telephonyonline.com/mag/telecom_public_peering_survive/

Bouygues Telecom. (n.d.). *Peering policy*. Retrieved online July 1, 2008, from:

<http://peering.t-online.fr/>

Brilus, S. (2009, March 17). Redundant AS's. *Message posted to NANOG*

mailing list. Archived at:

<http://www.merit.edu/mail.archives/nanog/msg16299.html>

Brown, M. A., Hepner, C., & Popescu, A.C. (2009, January). *Internet Captivity*

and the De-peering Menace. Paper presented at NANOG 45. Santo

Domingo, Dominican Republic.

Burton, G. (2001, June 7). PSINet-C&W dispute causes Internet blackout.

Information Age magazine. Retrieved online January 13, 2009, from:

<http://web.archive.org/web/20070927183732/http://www.information->

[age.com/article/2001/june/psinet-](http://www.information-age.com/article/2001/june/psinet-)

[c_and_w_dispute_causes_internet_blackout](http://www.information-age.com/article/2001/june/psinet-c_and_w_dispute_causes_internet_blackout)

Butler, S. (2000). *The evolution of Internet interconnections*. Unpublished masters dissertation. Capstone Project, Rochester Institute of Technology. Retrieved online September 1, 2008, from: <http://www.2sparrows.org/Old-2sparrows/rit/final%20thesis.pdf>

Cannon, R. (2003). The legacy of the Federal Communications Commission's Computer Inquiries. *Federal Communications Law Journal*. 55(2), p.167.

Canto, V. (2002, January 14). How Enron failed. *National Review Online*. Retrieved online September 1, 2008, from: http://www.nationalreview.com/nrof_canto/canto011402.shtml

Canadian Radio-television and Telecommunications Commission. Telecom Decision CRTC 92-12. (1992, June 12). Retrieved online June 30, 2009, from: <http://www.crtc.gc.ca/eng/archive/1992/DT92-12.htm>

Canadian Radio-television and Telecommunications Commission. Telecom Decision CRTC 2008-108. (2008a, November 20). Retrieved online July 15, 2009, from: <http://www.crtc.gc.ca/eng/archive/2008/dt2008-108.htm>

Canadian Radio-television and Telecommunications Commission. Telecom Public Notice CRTC 2008-19. (2008b, November 20). Review of the

Internet traffic management practices of Internet service providers.

Retrieved online June 30, 2009, from:

<http://www.crtc.gc.ca/eng/archive/2008/pt2008-19.htm>

Canadian Radio-television and Telecommunications Commission. Broadcasting

Regulatory Policy CRTC 2009-329. (2009a, June 4). Retrieved online

June 30, 2009, from: [http://www.crtc.gc.ca/eng/archive/2009/2009-](http://www.crtc.gc.ca/eng/archive/2009/2009-329.htm)

[329.htm](http://www.crtc.gc.ca/eng/archive/2009/2009-329.htm)

Canadian Radio-television and Telecommunications Commission. (2009b, July

14). Transcript of Proceedings before the Canadian Radio-television and

Telecommunications Commission. Subject: Review of the Internet traffic

management practices of Internet service providers. Retrieved online

September 1, 2009, from:

<http://www.crtc.gc.ca/eng/transcripts/2009/tt0714.htm>

Canadian Radio-television and Telecommunications Commission.

Communications Monitoring Report. (2009c, August). Retrieved online

August 30, 2009, from:

<http://www.crtc.gc.ca/eng/publications/reports/policymonitoring/2009/2009>

[MonitoringReportFinalEn.pdf](http://www.crtc.gc.ca/eng/publications/reports/policymonitoring/2009/2009MonitoringReportFinalEn.pdf)

Clark, D. (1992, July). A Cloudy Crystal Ball: Visions of the Future. Presentation to the IETF. Retrieved online August 1, 2009, from:

http://xys.ccert.edu.cn/reference/future_ietf_92.pdf

Clark, D., Reed, D. & Saltzer, J. (1984, November). End-to-End Arguments in System Design. *ACM Transactions on Computer Systems*. 2(4), pp. 277-288.

Collins, R. (2007). Rawls, Fraser, redistribution, recognition and the World Summit on the Information Society. *International Journal of Communication*. 1, pp.1-23.

Comcast Corporation, Petitioner, v. Federal Communications Commission and United States of America, Respondents. In the United States Court of Appeals for the District of Columbia Circuit. (2009, July 27). Opening Brief for Petitioner Comcast Corporation. Case: 08-1291. Retrieved online Sept 1, 2009, from: http://static.arstechnica.com/Comcast_Opening_Brief_%28as_filed%29.pdf.

Cowley, S. (2005, October 10). ISP spat leaves customers disconnected. *Network World*. Retrieved online January 11, 2009, from:

http://www.infoworld.com/article/05/10/06/HNisp spat_1.html

Crocker, S. (2009, April 7). How the Internet Got Its Rules. *The New York Times*.
p. 29.

Crowcroft, J. (2007). Net neutrality: The technical side of the debate - a white
paper. *International Journal of Communication*, 1, pp.567-579.

Cukier, K. (1997, December 3). *Peering and fearing: ISP interconnection and
regulatory issues*. Paper presented at the Harvard Information
Infrastructure Project Conference on the Impact of the Internet on
Communication Policy. Retrieved online September 12, 2008, from:
<http://www.cukier.com/writings/peering-cukier-dec97.html>

Cukier, K. (2004, October 2). *Multilateral control of Internet infrastructure and its
impact on US sovereignty*. Paper presented at Telecommunications Policy
Research Conference. Retrieved online November 20, 2008, from:
<http://www.cukier.com/writings/cukier-netgov-TPRC04.pdf>

Cybertelecom. (2006). AT&T / Bell South merger commitments. *Internet
Communications Regulation Digest*. Retrieved online January 11, 2009,
from: <http://www.cybertelecom.org/docs/attbsconditions.htm>

- Cybertelecom. (2007). Federal Internet law & policy misc issues. *Internet Communications Regulation Digest*. Retrieved online October 14, 2008, from: <http://www.cybertelecom.orgnotes/misc.htm#Peering>
- Davidson, A. (2008, February 14). *Re: question on the topology of Internet exchange points*. Message posted to NANOG mailing list. Archived at: <http://www.merit.edu/mail.archives/nanog/msg06093.html>
- de Sola Pool, I. (1983). *Technologies of freedom*. Cambridge: Harvard University Press.
- DeGeest, K. (2001). *What is a MPLS VPN anyway?* SANA Reading Room whitepaper. Retrieved online June 26, 2009, from: http://www.sans.org/reading_room/whitepapers/vpns/what_is_an_mpls_vpn_anyway_718
- Deibert, R., Palfrey, J., Rohozinski, R., & Zittrain, J. (Eds). (2008). *Access denied: The practice and policy of global Internet filtering*. Boston: MIT Press.
- DePalma, J. (n.d). *Maturation in a fee market: The changing dynamics of peering in the ISP industry*. Unpublished paper from the CATO Institute.

Estrada, S. (1993, January 5). FYI: US NIC changes or non-changes. Retrieved online August 29, 2009, from: <http://www.ripe.net/ripe/maillists/archives/lir-wg/1992/msg00028.html>

Faratin, P., Clark, D., Gilmore, P., Bauer, S., Berger, A. and Lehr, W. (2007, September 30). *Complexity of Internet interconnections: Technology, incentives and implications for policy*. Paper presented at The 35th Research Conference on Communication, Information and Internet Policy, George Mason University School of Law, Arlington, VA. Retrieved online July 19, 2009, from: http://people.csail.mit.edu/wlehr/Lehr-Papers_files/Clark%20Lehr%20Faratin%20Complexity%20Interconnection%20TPRC%202007.pdf

Federal Communications Commission. (1996). *Telecommunications Act of 1996*. Public Law 104-104, Washington, D.C. Retrieved online January 15, 2009, from: <http://www.fcc.gov/telecom.html>

Federal Communications Commission. (1998a, September 14). *In the Matter of Application of WorldCom, Inc. and MCI Communications Corporation for Transfer of Control of MCI Communications Corporation to WorldCom, Inc.*, FCC 98-225, CC Docket No. 97-211. Retrieved online April 21,

2008, from:

http://www.fcc.gov/Bureaus/Common_Carrier/Orders/1998/fcc98225.txt

Federal Communications Commission. (1998b, August 7). *Inquiry Concerning the Deployment of Advanced Telecommunications Capability to All Americans in a Reasonable and Timely Fashion, and Possible Steps to Accelerate Such Deployment Pursuant to Section 706 of the Telecommunications Act of 1996*. FCC 98 – 187, CC Docket 98-146.

Retrieved online July 16, 2008, from:

http://www.fcc.gov/Bureaus/Common_Carrier/Notices/1998/fcc98187.pdf

Federal Communications Commission. (2003, July). *Homeland Security Action Plan*. Department of Homeland Security. Retrieved online June 20, 2009, from: http://www.fcc.gov/homeland/#action_plan

Federal Communications Commission. (2008, August 20). FCC File No. EB-08-IH-1518. *Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications*.

Retrieved online September 1, 2009, from:

http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf

Frieden, R. (1998, Fall). Without public peer: The potential regulatory and universal service consequences of Internet balkanization. *Virginia Journal of Law and Technology*, 3 art.8.

Frieden, R. (2001, October). Revenge of the Bellheads: How the Netheads lost control of the Internet. Retrieved online June 30, 2009, from: <http://ssrn.com/abstract=290121> or DOI: 10.2139/ssrn.290121

Fusco, P. (2000, April 5). PSINet, Exodus terminate peering agreement. *InternetNews*. Retrieved online January 11, 2009, from: http://www.internetnews.com/xSP/article.php/8_334471

Gareiss, R. (1999, October 7). The old boy's network: Better net performance requires better peering. *Data Communications*, 28(14), p.36.

Geist, M. (2005, December 19). Dangers lurk in ISPs' bid for new tolls, two-tier Web. *Toronto Star*. p.C5.

Geist, M. (2009a, July 6). CRTC Network Management Hearings, Day One: Sandvine, Juniper, Consumer Group. *Michael Geist Blog*. Retrieved online July 7, 2009, from: <http://www.michaelgeist.ca/content/view/4103/125/>

Geist, M. (2009b, July 15). In Case You Missed It: Reflecting on the CRTC's Net Neutrality Hearing. *Michael Geist Blog*. Retrieved online August 1, 2009, from: <http://www.michaelgeist.ca/content/view/4135/125/>

Goldsmith, J., & Wu, T. (2006). *Who Controls the Internet?* New York: Oxford.

Gross, G. (2004, April 15). ISPs to form national lobbying group. *InfoWorld*. Retrieved online September 17, 2008, from: http://www.infoworld.com/article/04/04/15/HNisps_1.html

Guofeng, Z., Hong, T. & Yi, Z. (2004). Mapping edge-based traffic measurements onto the internal links in MPLS networks. In Huebner-Szabo de Bucs, F. & van der Mel, R. (ed). *Proceedings of SPIE, Performance, quality of service, and control of next-generation communications networks II*. Bellingham, WA. Vol. 5598.

Halabi, B. (1997). *Internet Routing Architectures*. Indianapolis: Cisco Press: New Riders Publishing.

Hofmann, J. (2007). Internet governance: A regulative idea in flux. In R. K. J. Bandamutha (Ed.), *Internet Governance: An Introduction*. (pp. 74-108). Hyderabad: Icfai University Press. Also Retrieved online January 25,

2008, from:

<http://duplox.wzb.eu/people/jeanette/texte/Internet%20Governance%20english%20version.pdf>

Horten, M. (2009, March 25). Six MEPs table AT&T's Internet-limiting proposals.

IPtegrity website. Retrieved online June 29, 2009, from:

http://www.iptegrity.com/index.php?option=com_content&task=view&id=290&Itemid=9

Hunter, P. (2006a). IPTV caught in vice between rising quality expectations and growth in traffic. *IPTV: Special supplement to Cable and Satellite International*.

Hunter, P. (2006b, March-April). Playout fades into origination as IP brings challenges and solutions. *Cable and Satellite International*.

Huston, G. (2003). Analyzing the Internet BGP routing table. *The Internet Protocol Journal* 4(1).

Huston, G. (2009). *Re: Redundant AS's. Message posted to NANOG mailing list.*

Archived at: <http://www.merit.edu/mail.archives/nanog/msg16368.html>

International Telecommunication Union. (2000, September 27 – October 6).

World Telecommunication Standardization Assembly (WTSA-2000)

Montreal, Canada. Retrieved online June 30, 2009, from:

<http://www.itu.int/itudoc/itu-t/circ/circ5/262.html>

International Telecommunication Union. (2004). *Chairman's Report, ITU New*

Initiatives Programme, Workshop on Internet Governance, 26-27 February

2004. Retrieved online September 1, 2007, from:

<http://www.itu.int/osg/spu/forum/intgov04/workshop-internet-governance-chairmans-report.pdf>

Internet Assigned Numbers Authority. (2009). Autonomous System (AS)

Numbers. Retrieved online July 25, 2009, from:

<http://www.iana.org/assignments/as-numbers/as-numbers.xml>

Internet Corporation for Assigned Names and Numbers. (2009). About page.

Retrieved online August 3, 2009, from: <http://www.icann.org/en/about/>

Internet Engineering Task Force (2009). Retrieved online May 25, 2009, from:

<http://www.ietf.org/tao.html>

Internet Society. (2009). Introduction to ISOC. Retrieved online May 25, 2009, from: <http://www.isoc.org/isoc/>

Internet World Stats. (2009). World Internet users and population stats. Retrieved online July 20, 2009, from: <http://www.internetworldstats.com/stats.htm>

Jamie. (2009, July 26). Re: AT&T. Layer 6-8 needed. Message posted to *NANOG mailing list*. Archived at: <http://www.merit.edu/mail.archives/nanog/msg19604.html>

Jensen, M. (2005, September). *Interconnection costs*. Association for Progressive Communications. Issues Paper Series. Retrieved online August 17, 2008, from: http://www.apc.org/en/system/files/interconnection_costs+en.pdf

Jürgen, K. & Smith, R. W. (2005, April 21). France Telecom severs all network links to competitor Cogent. *Heise online*. Retrieved online January 11, 2009, from: <http://morse.colorado.edu/~epperson/courses/routing-protocols/handouts/cogent-ft.html>

Kende, M. (2000). *The digital handshake: Connecting Internet backbones*. (OPP Working paper no 32). Office of Plans and Policy, Federal

Communications Commission. Washington, D.C. Retrieved online

September 12, 2007, from:

http://www.fcc.gov/Bureaus/OPP/working_papers/oppwp32.pdf

Keshav, S. (1997). *An engineering approach to computer networking*. Boston: Addison-Wesley.

Kessler, G. (2007). An Overview of TCP/IP Protocols and the Internet.

GaryKessler website. Retrieved online June 30, 2009, from:

<http://www.garykessler.net/library/tcpip.html>

King-Guillaume, L. (2003, January 15). Advanced peering: A better alternative to the traditional Internet peering model. *Telephony Online*. Retrieved online February 28, 2008 from:

http://telephonyonline.com/broadband/infocus/telecom_advanced_peering_better/index.html

Labaton, S. (2000, June 17). FCC approves Bell Atlantic-GTE merger. *New York Times*. p. C1.

Laffont, J.J., Marcus, S. Rey, P., & Tirole, J. (2001, January 5-7). *Internet Peering*. In Baldwin, J. D. & Oaxaca, R. L. (2001, May). Papers and

Proceedings of the 113 Annual Meeting of the American Economic Association. New Orleans, LA.

Le Boudier, G. (2003, January 11). Problème de peering entre Free et France Télécom. *LinuxFr*. Retrieved online January 13, 2009 from:
<http://linuxfr.org/2003/01/21/11058.html>

Lehr, W. H., Peha, J. M., & Wilkie, S. (2007). The state of the debate on network neutrality. *International Journal of Communication*. 1, pp.709-716.
Retrieved online November 15, 2007, from:
<http://ijoc.org/ojs/index.php/ijoc>

Levine, J. (2009, May 4). Re: Minnesota to block online gambling sites? *Message posted to NANOG mailing list*. Archived at:
<http://www.merit.edu/mail.archives/nanog/msg17723.html>

Longford, G. (2007). *Network neutrality vs network diversity: A survey of the debate, policy landscape and implications for broadband as an essential service for Ontarians*. Unpublished post-doctoral research. University of Toronto.

Marcus, J. S. (2006, June 8). *Declaration in support of Plaintiffs' motion for preliminary injunction*. Electronic Frontier Foundation. Retrieved online June 4, 2009, from: <http://www.eff.org/cases/att/attachments/unredacted-marcus-declaration> & http://www.eff.org/files/filenode/att/SER_marcus_decl.pdf

Markoff, J. (1999, December 20). An Internet pioneer ponders the next revolution. *The New York Times*. p. 38.

Mathiason, J., McKnight, L. & Mueller, M. (2004). Making sense of Internet governance: Defining principles and norms in a policy context. In MacLean, D. (Ed.) *Internet Governance: A grand collaboration*. New York: United Nations Information Communications Task Force Series 5.

McCarthy, K. (2006, August). ICANN awarded net administration until 2011. *The Register*. Retrieved online October 25, 2007, from: http://www.theregister.co.uk/2006/08/16/icann_awarded_iana/print.html

McCullagh, D. (2005, March 29). *Newsmaker: The U.N. thinks about tomorrow's cyberspace*. Cnet news. Retrieved online November 7, 2008, from: http://news.cnet.com/The-U.N.-thinks-about-tomorrows-cyberspace/2008-1028_3-5643972.html

MCI, WorldCom merger gets green light from DoJ. (1998, July 15).

InternetNews.com. Retrieved online January 11, 2009, from:

<http://www.internetnews.com/bus-news/article.php/21611>

McPherson, D., Sangli, S. & White, W. (2005). *Practical BGP*. New York:

Addison-Wesley.

Medhi, D. & Ramasamy, K. (2007). *Network Routing*. Boston: Elsevier.

Miller, R. (2008a, March 18). Cogent Telia dispute widely felt. *IDG TechNetwork*.

Retrieved online December 18, 2008, from:

http://www.datacenterknowledge.com/archives/2008/Mar/18/cogent-telia_peering_dispute_widely_felt.html

Miller, R. (2008b, October 31). Peering Dispute Between Cogent, Sprint. *Data*

Center Knowledge Web site. Retrieved online September 2, 2009, from:

<http://www.datacenterknowledge.com/archives/2008/10/31/peering-dispute-between-cogent-sprint/>

Morgenstern, D. (2005). Feds won't let go of Internet DNS. *eWeek Magazine*,

July 1, 2005. Retrieved online March 16, 2007, from:

<http://www.eweek.com/c/a/Infrastructure/Feds-Wont-Let-Go-of-Internet-DNS/>

Mueller, M. (2006). I aint'a giving up IANA. *Internet Governance Project blog*.

Retrieved November 23, 2007, from:

<http://blog.internetgovernance.org/blog/archives/2006/8/22/3340286.html>

Mwangi, M. (2008, February 15). *Re: Question on the topology of Internet exchange points. Message posted to NANOG mailing list. Archived at:*

<http://www.merit.edu/mail.archives/nanog/msg06105.html>

National Telecommunications and Information Administration. (1998, January 30). Proposal to improve technical management of Internet names and addresses. Retrieved online July 20, 2009, from:

<http://www.ntia.doc.gov/ntiahome/domainname/dnsdrft.htm>

NETCompetition. (2009, May 7). Calling the bluff. Retrieved online September 1, 2009, from: [http://netcompetition.org/index.php/go/in-the-news-](http://netcompetition.org/index.php/go/in-the-news-more/calling_the_bluff/)

[more/calling_the_bluff/](http://netcompetition.org/index.php/go/in-the-news-more/calling_the_bluff/)

Network Reliability and Interoperability Council. (2002a). *Network Reliability and Interoperability Council V: The future of our nation's communications*

infrastructure. Report to the Nation; Index. Focus Group 4. Retrieved online September 11, 2008, from: <http://www.nric.org/pubs/index.html>

Network Reliability and Interoperability Council. (2002b). *Network Reliability and Interoperability Council V: The future of our nation's communications infrastructure. Report to the Nation; Final Report.* Focus Group 4. Retrieved online September 1, 2007, from: <http://www.nric.org/pubs/nric5/2B4finalreport.doc>

Network Reliability and Interoperability Council. (2002c). *Network Reliability and Interoperability Council V: The future of our nation's communications infrastructure. Report to the Nation; Appendix A.* Focus Group 4. Retrieved online September 1, 2007, from: <http://www.nric.org/pubs/nric5/2B4appendixa.doc>

Network Reliability and Interoperability Council. (2002d). *Network Reliability and Interoperability Council V: The future of our nation's communications infrastructure. Report to the Nation; Appendix B.* Focus Group 4. Retrieved online September 1, 2007, from: <http://www.nric.org/pubs/nric5/2B4appendixb.doc>

Neuman, W. R., McKnight, L. Solomon, R. J. (1998). *The gordian knot*.

Cambridge: MIT Press.

New York Department of State (2009). Welcome to the Division of Administrative

Rules. Official Compilation of codes, rules and regulations of the State of

New York. Title 16. Department of Public Service. Ch VI. Telephone and

Telegraph Corporations. Subchapter A. Service. Part 605. Common

Carrier Rules. Retrieved online June 4, 2009, from:

<http://www.dos.state.ny.us/info/register.htm>

Noam, E. (2001). *Interconnecting the Network of Networks*. Boston: MIT Press.

Norton, W. (2001). *Internet service providers and peering. Draft 2.5*. Last

modified: 05/30/2001. Retrieved online June 15, 2007, from:

<http://pages.cs.wisc.edu/~akella/CS740/S08/740-Papers/Nor00.pdf>

Norton, W. (2003). *The evolution of the U.S. Internet peering ecosystem. Draft*

1.1. Last modified: 11/19/2003. Retrieved online June 15, 2007, from:

<http://www.nanog.org/meetings/nanog31/abstracts.php?pt=NjA2Jm5hbm9nMzE=&nm=nanog31>

NSFNET backbone services Acceptable Use Policy. (1992, June). Retrieved

online July 15, 2009, from:

http://w2.eff.org/Net_culture/Net_info/Technical/Policy/nsfnet.policy

Nuechterlein, J. E. & Weiser, P. J. (2007). *Digital crossroads: American telecommunications policy in the Internet age, 2nd edition*. Cambridge: MIT Press.

Olsen, K. & Tebbutt, J. (1995, February). *The impact of the FCC's Open Network Architecture on NS/NP telecommunications security*. NIST

Special Publication, 800. Retrieved online June 4, 2009, from:

<http://csrc.nist.gov/publications/nistpubs/800-11/>

OpenNet Initiative Bulletin 010. (2005, August 2). *Telus blocks consumer access to labour union web site*. Retrieved online July 1, 2008, from:

<http://opennet.net/bulletins/010>

Oram, A. (2003, December 12). When did we give away the Internet? *CircleID weblog*. Retrieved online July 16, 2008, from:

http://www.circleid.com/posts/when_did_we_give_away_the_internet_also
www.circleid.com/print/396_0_1_0/

Passmore, D. (2004, November 2). *Strategic networking overview: Major trends in broadband networking*. In Proceedings of Next Generation Networks. Boston, Ma.

Paterson, N. (2008, April 20). Is peering a net neutrality issue? *Message posted to Internet Governance Forum listserv*. Archived at:
<http://lists.cpsr.org/lists/arc/governance/2008-04/msg00295.html>

Perez, M. (2008, October 23). AT&T completes next-gen IP backbone network. *Information Week website*. Retrieved online July 15, 2009, from:
<http://www.informationweek.com/news/telecom/business/showArticle.jhtml?articleID=211600239>

Pfanner, E. (2009, July 12). *New chief defends U.S. base for agency that manages web*. Retrieved online July 25, 2009, from:
http://www.nytimes.com/2009/07/13/technology/internet/13iht-icann13.html?_r=1&scp=1&sq=beckstrom&st=cse

Remarks by the President on securing our nation's cyber infrastructure (2009, May 29). *Office of the Press Secretary. The White House*. Retrieved online May 30, 2009 from:
http://www.whitehouse.gov/the_press_office/Remarks-by-the-President-on-Securing-Our-Nations-Cyber-Infrastructure/

Reuters. (2000, June). Sprint-WorldCom: Dead on the vine. *Reuters*. Retrieved online January 13, 2009, from:

<http://www.wired.com/politics/law/news/2000/06/37250>

Rhoads, C. (2009, July 29). Tech Journal: Snooping on Web Traffic Gains Favor Amid Fears. *The Wall Street Journal*.

Ricknäs, M. (2008, October 31). Sprint-Cogent dispute puts small rip in fabric of Internet. *PC World*. Retrieved online January 11, 2009, from:

http://www.pcworld.com/businesscenter/article/153123/sprintcogent_dispute_puts_small_rip_in_fabric_of_internet.html

Roland, A. (2000). Review of Janet Abbate, *Inventing the Internet*. *Technology and Culture* 41(4). pp. 826-828.

Routing Analysis Role Account. (2008, July 18). *Weekly routing table report*.

Retrieved online July 18, 2008, from:

<http://www.merit.edu/mail.archives/nanog/msg09688.html>

- Russell, A. L. (2006, July-September). Rough consensus and running code and the Internet-OSI standards war. *IEEE Annals of the History of Computing*. 28(3). pp.48-61.
- Salido, J., Nakahara, M., Wang, Y. (2003, June 23-24). An analysis of network reachability using BGP data. *WIAPP 2003 Proceedings: The Third IEEE Workshop on Internet Applications*. Page(s): 10 – 18.
- Schiller, H. (2009). Re: Redundant AS's. *Message posted to NANOG mailing list*.
Archived at: <http://www.merit.edu/mail.archives/nanog/msg16364.html>
- Steier, R. (1986, January). Computer Inquiry III. *ACM: From Washington*. 29(1).
- Steinberg, S. (1996, October). Netheads-vs-Bellheads. *Wired Magazine*. 4(10).
Retrieved online June 30, 2009, from:
<http://www.wired.com/wired/archive/4.10/atm.html>
- Telx opens the most network rich, interconnected colocation center in NYC metro area. (2009, April 14). *Market Wire*.
- Thussu, D. K. (2006). *International Communication*, 2nd Edition. Oxford: New York.

US Department of Justice Complaint. (2000, June 26). *USA v. MCI & Sprint*. par.

32. Retrieved online July 20, 2009, from:

<http://www.usdoj.gov/atr/cases/f5000/5051.htm>

US Fed News. (2008, October 11). Georgia Inventor Develops Multiprotocol Label Switching Packet.

U.S. G.P.O. (2008). Foreign Intelligence Surveillance Act of 1978 Amendments Act of 2008. H.R. 6304. Public Law 110-261. July 10, 2008. Washington, D.C: Supt. of Docs. Congressional Record, Vol. 154 (2008). Retrieved online June 4, 2009, from:

<http://www.govtrack.us/congress/bill.xpd?bill=h110-6304>

U.S. Telecommunications Act of 1996. Pub. LA. No. 104-104, 110 Stat. 56 (1996). Retrieved online June 4, 2009, from:

<http://www.fcc.gov/Reports/tcom1996.txt>

Ungerer, H. (2000). Access issues under EU regulation and antitrust law: The case of telecommunications and Internet markets. *Weatherhead Center for International Affairs, Harvard University. Columbia International Affairs*

Online. Retrieved online December 8, 2008, from:

http://www.ijclp.net/ijclp_web-doc_4-5-2000.html

van Beijnum, I. (2008, February 25). Insecure routing redirects YouTube to

Pakistan. Ars Technica. Retrieved online May 12, 2009, from:

<http://arstechnica.com/old/content/2008/02/insecure-routing-redirects-youtube-to-pakistan.ars>

van der Berg, R. (2009). *How the 'Net works: an introduction to peering and*

transit. Ars Technica website. Retrieved online July 1, 2008, from:

<http://arstechnica.com/old/content/2008/09/peering-and-transit.ars>

Varian, H. R. (1998, June 8) How to Strengthen the Internet's Backbone, *The*

Wall Street Journal. p. A22.

Vixie, P. (2008, February 14). *Re: Question on the topology of Internet exchange*

points. Message posted to NANOG mailing list. Archived at:

<http://www.merit.edu/mail.archives/nanog/msg06094.html>

von Finckenstein, K. (2007). Address to the 2007 Broadcasting Invitational

Summit. Jackson's Point, Ontario. June 26, 2007. Retrieved online July 8,

2008, from:

<http://www.crtc.gc.ca/eng/NEWS/SPEECHES/2007/s070626.htm>

von Finckenstein, K. (2008). Address to the 2008 Canadian Telecom Summit.

Toronto, Ontario. June 17, 2008. Retrieved July 8, 2008, from:

<http://www.crtc.gc.ca/eng/NEWS/SPEECHES/2008/s080617.htm>

Walsh, P. (2009, April 30). New tactic in war on online gambling. *The Star Tribune*.

Weinberg, N. (2000, June 12). Backbone bullies. *Forbes*.

Wigfield, M. (2005, October 31). FCC approves SBC/AT&T and VERIZON/MCI

mergers. FCC News Release. Retrieved online May 12, 2009, from:

http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-261936A1.pdf

Williamson, J. (2003, August). Peering into the future. *Telecommunications - International Edition*. 37(8),p.16.

Wilson, C. (2009, January 21). XO uses DPI for app performance management.

Telephonyonline. Retrieved online May 12, 2009, from:

http://www.telephonyonline.com/business_services/news/xo-dpi-app-performance-0121/index.html

Wolf, C. (1999). Internet infrastructure issues: Regulation and un-regulation of the 'pipes' that provide the Internet. *FindLaw*. Retrieved online December 1, 2008, from: <http://library.lp.findlaw.com/articles/file/00086/002196>

Wu, T. (2003). Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*. Vol. 2, p. 141. Retrieved online June 30, 2009, from: <http://ssrn.com/abstract=388863> or DOI: 10.2139/ssrn.388863

Wu, T. (2007, February). Wireless net neutrality: Cellular carterfone and consumer choice in mobile broadband. New America Foundation, Wireless Future Program. Working Paper #17. Retrieved online June 30, 2009, from: http://www.newamerica.net/publications/policy/wireless_net_neutrality

XO wins ruling against Level 3 Communications. (2007, November). *Telecom News*. Retrieved online January 11, 2009, from: <http://vartips.com/carriers/level-3/xo-wins-ruling-against-level-3-communications.html>

Zakon, R. H. (n.d.). Hobbes' Internet Timeline v8.2. Retrieved online June 30, 2009, from: <http://www.zakon.org/robert/internet/timeline/>

Zarkin, M.J. (2003). Telecommunications policy learning: The case of the FCC's computer inquiries. *Telecommunications Policy* 27. p.283–299.

Zittrain, J. (2006, May). The generative Internet. *Harvard Law Review*, 119(7), pp.1974-2040. Retrieved online February 28, 2008, from: <http://www.harvardlawreview.org/issues/119/may06/zittrain.pdf>

Zmijewski, E. (2008a, March 17). You can't get there from here. *Renesisys Blog*. Retrieved online January 13, 2009, from <http://www.renesys.com/blog/2008/03/you-cant-get-there-from-here-1.shtml#more>

Zmijewski, E. (2008b, December 17). Rising to the Top: A Baker's Dozen. *Renesisys blog*. Retrieved online September 1, 2009, from: <http://www.renesys.com/blog/2008/12/winners-and-losers-for-2008.shtml>