NATIONAL ASSOCIATION OF STATE UTILITY CONSUMER ADVOCATES NASUCA

August 8, 2014

The Honorable Fred Upton, Chairman
The Honorable Henry Waxman, Ranking Member
Committee on Energy and Commerce
2125 Rayburn HOB, 2322A Rayburn HOB
Washington, D.C. 20515

The Honorable Greg Walden, Chairman
The Honorable Anna Eshoo, Ranking Member
Subcommittee on Communications, Communications Technology & the Internet
2125 Rayburn HOB, 2322A Rayburn HOB
Washington, D.C. 20515 Washington, D.C. 20515

Dear Chairmen Upton and Walden and Ranking Members Waxman and Eshoo:

Enclosed are National Association State Utility Consumer Advocates ("NASUCA') comments on White Paper #4 on interconnection. These were e-mailed on August 8, 2014 and we are now sending a hard copy to the committee.

Sincerely,

Charles Acquard, Executive Director
NASUCA
8380 Colesville Road, Suite 101
Silver Spring, MD 20910
[redacted]

NATIONAL ASSOCIATION OF STATE UTILITY CONSUMER ADVOCATES NASUCA

August 8, 2014

RESPONSE TO HOUSE COMMITTEE ON ENERGY AND COMMERCE
FOURTH WHITE PAPER

NASUCA[1] submits these comments to the House Committee on Energy and Commerce ("Committee") in response to the Committee's request. NASUCA very much appreciates the opportunity to comment on each of the eight "Questions for Stakeholder Comment" on interconnection..

As a general overview, interconnection between networks has always been necessary (even before the 1913 Kingsbury Compact). It is all the more necessary in this era of multiple networks where a consumer may have last-mile access to only one or two of these networks. The need to regulate interconnection remains because some network owners retain significant market power,[2] and because market forces have failed to deliver good quality services to all consumers. Consumers need the protection of such regulation, in order to ensure their access to the full Internet ecology and their receipt of good quality service.

QUESTIONS FOR STAKEHOLDER COMMENT

1. In light of the changes in technology and the voice traffic market, what role should Congress and the FCC play in the oversight of interconnection? Is there a role for states?

---

[1] NASUCA is a voluntary, national association of consumer advocates in more than forty states and the District of Columbia, organized in 1979. NASUCA's members are designated by the laws of their respective states to represent the interests of utility consumers before state and federal regulators and in the courts. Members operate independently from state utility commissions, as advocates primarily for residential ratepayers. Some NASUCA member offices are separately established advocate organizations while others are divisions of larger state agencies (e.g., the state Attorney General's office). Associate and affiliate NASUCA members also serve utility consumers, but have not been created by state law or do not have statewide authority.

[2] See NASUCA response to WP#3, accessible at http://nasuca.org/nwp/wp-content/uploads/2013/11/NASUCA-Response-to-House-Committee-Third-White-Paper.pdf, and to NASUCA's Comments to the Federal Communications Commission ("FCC") in WC Docket No. 09-191, accessible at http://apps.fcc.gov/ecfs/document/view?id=7020375001.

Congress should continue to maintain the interconnection obligations set forth in the 1996 Telecom Act. Those obligations should not be reduced below the current requirements, but should be unified across platforms so as to ensure an open and vibrant national broadband and communications infrastructure. The current state role in arbitrating interconnection agreements should be maintained.

2. Voice is rapidly becoming an application that transits a variety of network data platforms. How should intermodal competition factor into interconnection mandates? Does voice still require a separate interconnection regime?

Interconnection is not just a matter for voice, but for it and all of the other applications that transit the networks. The level of intermodal competition for voice – or for other services – does not minimize the need to regulate interconnection between networks. As indicated in the response to #1, a general interconnection regime for voice and broadband – covering **more** networks, not fewer – is needed.

3. How does the evolution of emergency communications beyond the use of traditional voice service impact interconnection mandates?

Interconnection for emergency communications -- regardless of format -- should always be part of any mandate.

4. Ensuring rural call completion has always been a challenge because of the traditionally high access charges for terminating calls to high-cost networks. Does IP interconnection alleviate or exacerbate existing rural call completion challenges?

Despite ongoing FCC action,[3] as well as ongoing action in a number of states, continuing oversight over rural call completion is and will remain essential. The collection of data required by the FCC's recently adopted rules has not occurred. Moreover, analysis regarding the causes of the difficulties from call completion failures and conclusions regarding their solutions remain uncertain.

IP interconnection seriously exacerbates the challenges of call completion.[4] Hundreds of carriers are involved in the transporting of calls. Insufficient network software and hardware, including overloaded switches and call paths, can and do cause slowdowns in processes and consequent call failures. Equipment failures can and do cause call failures. The alarming rise of rural call

---

[3] Rural Call Completion, WC Docket No. 13-39, Report and Order and Further Notice of Proposed Rulemaking, 28 FCC Rcd 16154 (2013) (RCC Order) (requiring certain collection and reporting data on how successfully calls are being delivered, especially to rural areas; prohibiting false audible ringing; and seeking comment on additional reforms pertaining to autodialer traffic, intermediate providers, and on other safe harbor options and reporting requirements).

[4] See NASUCA comments in FCC Docket No. 13-39 at 11-15, accessible at http://apps.fcc.gov/ecfs/document/view?id=7022313049.

completion failures in recent years, coincident with the increasing interconnection of intermediate carriers using IP protocols, followed by the efforts of the FCC and the states to address the difficulties, attest to the need for a regulatory presence in order to restore and maintain the integrity of the network as a whole.

5. Should we analyze interconnection policy differently for best-efforts services and managed services where quality-of-service is a desired feature? If so, what should be the differences in policy between these regimes, and how should communications services be categorized?

As the networks grow and become more important for consumer communication, it is appropriate that they grow in quality and capability. Networks with mere "best efforts" capabilities should be limited to uses that do not require quality.[5]

6. Much of the committee's focus in the #CommActUpdate process has been on technology-neutral solutions. Is a technology-neutral solution to interconnection appropriate and effective to ensure the delivery and exchange of traffic?

The unified across platforms approach described in response #1 above is, in fact, "technology-neutral." All technology is treated the same when all networks are required to interconnect on just and reasonable terms, for the protection of all the networks' consumers.

7. Wireless and Internet providers have long voluntarily interconnected without regulatory intervention. Is this regime adequate to ensure consumer benefit in an all-IP world?

See responses to #s 1 and 6, above.

8. Is contract law sufficient to manage interconnection agreements between networks? Is there a less onerous regulatory backstop or regime that could achieve the goals of section 251?

Contract law itself is not sufficient to manage interconnection agreements between networks, where the fully-interconnected nature of those networks is of such immense value to consumers, states, and the National economy, but where network owners have market dominance. Thus a regime less "onerous" than contract law should not be considered. If the question refers to a system less onerous than the current legal and regulatory mechanisms, those who claim that the current system is burdensome tend to be network owners, whose primary interest is in being able to exert their market power.

Conclusion

NASUCA again appreciates the opportunity to provide these comments to the Committee. As NASUCA has stated in many previous contexts, the public interest is best served when policy-

[5] See id. at 15.

makers are not swayed by the business plans and pecuniary interests of particular companies - or indeed, particular industries. A balanced approach that considers the interests of consumers is best.

Respectfully,

Charles Acquard, Executive Director
NASUCA
8380 Colesville Road, Suite 101
Silver Spring, MD 20910
[redacted]

THE RURAL BROADBAND ASSOCIATION

NTCA–The Rural Broadband Association Comments in Response to U. S. House of Representatives Energy & Commerce Committee White Paper 4: Network Interconnection (Released July 15, 2014)

August 8, 2014

4121 Wilson Boulevard, Suite 1000 • Arlington, VA 22203-1801
Phone/703-351-2000 • Fax/703-351-2001 • www.ntca.org

I. INTRODUCTION

NTCA–The Rural Broadband Association (NTCA) hereby submits comments on the Energy and Commerce Committee "Network Interconnection" white paper.[1] NTCA represents nearly 900 small, rate-of-return rural telecommunications providers (commonly called "RLECs"). RLECs serve about five percent (5%) of the US population and roughly 40 percent (40%) of its landmass. These companies operate in areas long ago left behind by larger providers because the markets were too high-cost – too sparsely populated, too far from larger towns and cities, and/or just too challenging to serve in terms of topography or terrain. As anchors in the communities in which they live and serve, these small businesses create jobs, fuel the economy, and connect rural Americans to the world. Moreover, these rural network operators have been at the forefront of the broadband and Internet Protocol ("IP") evolution for years, making every innovative effort to deploy advanced networks that respond to consumer and business demand for cutting-edge services and extracting greater efficiencies from network operations in the face of operating in hard-to-serve areas.

II. BACKGROUND

Policymakers are currently confronting the issue of "IP interconnection." This matter contemplates just one facet of the Nation's evolution from communications networks based upon time division multiplexing (TDM) to internet protocol (IP), but it is one that is tied inextricably to the Nation's commitment to universal access and availability. As noted in the Committee's white paper, "[T]he interconnection of networks has been at the heart of communications policy since the Kingsbury Commitment in 1913 when AT&T guaranteed interconnection with independent companies in exchange for a government-sanctioned monopoly on long-distance

---

[1] "Network Interconnection," Energy and Commerce Committee, U.S. House of Representatives (available at http://energycommerce.house.gov/sites/republicans.energycommerce.house.gov/files/analysis/CommActUpdate/20140715WhitePaper-Interconnection.pdf) (last viewed Aug. 4, 2014, 15:13) (E&C Paper).

service."[2] The significance of interconnection was underscored when Congress enacted sweeping legislation aimed at facilitating competition in the communications marketplace in the Telecommunications Act of 1996.[3] In particular, Sections 251 and 252 of the Act direct all telecommunications carriers to interconnect with other telecommunications carriers, and impose distinct obligations upon incumbent local exchange carriers (ILECs).[4]

In the instant discussion of IP interconnection, it would be easy, yet imprecise, to view IP interconnection through lenses different than those ground in the landmark 1996 Act or even the core principles captured by the Kingsbury Commitment. This is not surprising – IP-based services rely upon an entirely different technological protocol than the TDM services that reigned in the days of the 1996 legislation and certainly the technologies that underpinned early telephone service. Accordingly, many are led, or might be misled, into thinking generally that basic principles requiring reasonable and transparent interconnection terms have little applicability in today's world and markets, or more specifically that the standards set forth in sections 251 and 252 of the Act have naught to do with emerging IP interconnection issues. Both assumptions are incorrect.

As a general matter, universal service cannot be fulfilled in a "market" that compels all network operators, large and small, to come to distant, centralized interconnection points to meet one another. Such interconnection arrangements may work well if one has a nationwide footprint. For a small operator serving several thousand consumers in rural New Mexico or Texas, however, the costs of interconnecting with other operators in Dallas or Los Angeles can

---

[2] E&C Paper, at 1.

[3] Telecommunications Act of 1996, Pub. L. No. 104-104, 110 Stat. 56 ("1996 Act"). The 1996 Act amended the Communications Act of 1934. Hereinafter, the Communications Act of 1934, as amended by the 1996 Act, will be referred to as "the Act," and all references to the Act will be to the Act as it is codified in the United States Code.

[4] *See*, 47 U.S.C. §§ 251, 252. These sections of the Act will be cited frequently in these comments. Each such citation may rely upon this footnote.

be cost-prohibitive, making the services delivered to the small company's customers similarly cost-prohibitive, as well. This result defies a core universal service principle that all parties and policymakers tout in considering any reforms or rewrites. Indeed, AT&T – one of the largest nationwide providers – has emphasized that, even in an IP world, transport networks are hardly "cost-free."[5] AT&T's correct observation is especially important when considering the situation of smaller providers that have but a few thousand consumers across whom to spread such costs. Universal service simply cannot and will not be served in the absence of rules that ensure interconnection nearer to the rural and remote markets they serve. Therefore, reasonable interconnection rules to govern the exchange of increasing amounts of data over IP networks are just as, if not more, important as they were a century ago when policymakers encouraged the Kingsbury Commitment to keep rural areas connected.

Beyond this general overview, an argument that sections 251 and 252 are irrelevant to interconnection of IP networks is incorrect, as illustrated more directly by recent Federal Communications Commission (FCC) action. Specifically, in 2011, the FCC executed a watershed decision that, *inter alia,* gathered all telecommunications beneath the statutory umbrella of section 251(b)(5),[6] which addresses inter-carrier compensation (ICC) for telecommunications exchanged among carriers.[7] Section 251(b), generally, describes the

---

[5] *Connect America Fund; Developing a Unified Intercarrier Compensation Regime; Protecting and Promoting the Open Internet: Ex Parte Presentation of AT&T,* Docket Nos. 10-90, 01-92, 14-28, attachment at 14, (filed Jul. 10, 2014) (*AT&T Ex Parte*). *See, also,* "Who Should Pay for Netflix?" Jim Cicconi, AT&T Public Policy Blog (Mar. 21, 2014) (available at http://www.attpublicpolicy.com/consumers-2/who-should-pay-for-netflix/) (last viewed Aug. 4, 2014, 15:28).

[6] 47 U.S.C. § 251(b)(5).

[7] *Connect America Fund; A National Broadband Plan for Our Future; Establishing Just and Reasonable Rates for Local Exchange Carriers; High-Cost Universal Service Support; Developing a Unified Intercarrier Compensation Regime; Federal-State Joint Board on Universal Service; Lifeline and Link-Up; Universal Service Fund – Mobility Fund: Report and Order and Further Notice of Proposed Rulemaking,* Federal Communications Commission Docket Nos. 10-90; 09-51; 07-135; 05-337; 01-92; 96-45; 03-109; 10-208, FCC 11-161, at paras. 771, 772 (2011) (emphasis added).

obligations of local exchange carriers (LECs) as they interconnect with other carriers. Elements governed by section 251(b) are discrete parts of the general obligation of *all* telecommunications carriers to interconnect. By way of illustration, *all* telecommunications carriers are subject to section 251(a); LECs (a subset of "telecommunications carriers") are subject to the general requirements of section 251(a) *and* the specific requirements of section 251(b), and; incumbent LECs (a subset of "LECs") are subject to the general requirements of section 251(a), the specific requirements of section 251(b), and additional obligations of section 251(c). To the extent that any element is included in the specific parameters of section 251(b) or 251(c), it is by definition included in the general scope of section 251(a).

By asserting interconnection authority over transport and termination functions of telecommunications in section 251(b)(5), the FCC acknowledged that such interconnection is governed overall by the general requirements of section 251(a). Section 251(a) imposes a "general duty of telecommunications carriers" to "interconnect directly with the facilities and equipment of other telecommunications carriers." The statute does not distinguish or even indicate that any different treatment may arise based upon the *type* of traffic that is exchanged. Rather, section 251 (and the following section 252, which addresses the task of negotiation agreements to govern traffic exchanged by telecommunications carriers) simply ascribes the duty to interconnect, and all the components of that obligation, upon telecommunications carriers. There is no mention that any obligation is subordinate or otherwise affected by the type of technology underlying the traffic that is to be exchanged. (Indeed, in the 2011 order, the FCC expressly *included* Voice over Internet Protocol (VoIP) traffic within the Section 251(b)(5) framework[8] – it could not have done so had the statute been technology-specific and treated treatment of IP traffic as beyond its scope.) Inasmuch as the FCC brought all

---

[8] *Transformation Order*, at para. 736, *et seq. See, also, id.* at note 1443.

telecommunications traffic beneath Section 251(b)(5) for the specific purpose of addressing ICC rates, there is a clear statutory path leading to the conclusion that all such traffic is necessarily subject to the overarching requirements affecting interconnection, generally.

This conclusion is supported not only by logic, but also by the legislative history of the 1996 Act. The Senate comment explains,

> [n]ew subsection 251(a) imposes a duty on local exchange carriers possessing market power . . . to provide interconnection with other telecommunications carriers that have requested interconnection for the purpose of providing telephone exchange service or exchange access service . . . [t]he relevant market shall include all providers of telephone exchange service or exchange access service in a local service area, *regardless of the technology used to provide such service.*[9]

That adjuration is uncontested in the conference agreement, and moreover is uncontested in the statute.

Finally, in addition to the plain reading of the statute and the legislative history, the conclusion that the obligation to interconnect among telecommunications carriers arises regardless of their underlying technology is consistent with the intent of the 1996 Act and the statutory provisions aimed at fulfilling those goals. The pro-competitive stance of the 1996 Act intends to open communications markets to new entrants and competition. This is evident from the manifest obligations surrounding interconnection, collocation, resale, number portability, negotiation and arbitration, and other aspects of sections 251 and 252. Collectively, these are aimed at bringing to consumers options in their selection of telecommunications services, and none of those revolve, rely, or are otherwise revised by the technology that underlies the telecommunications service.

From these reasons emerge: (1) the threshold determination that IP interconnection is simply “interconnection” as the industry has addressed it for nearly 20 years and even decades

---

[9] H.R. Conf. Rep. 104-458, p.117.

prior, and (2) a further conclusion that universal service cannot be fulfilled in the absence of sound and clear rules that create a "regulatory backdrop" against which just and reasonable interconnection not only may, but must, occur.

III. SPECIFIC INQUIRES AND RESPONSES

The Committee asks several questions:

1. *In light of the changes in technology and the voice traffic market, what role should Congress and the FCC play in the oversight of interconnection? Is there a role for the states?*

Congress and the FCC must play a role in the oversight of interconnection, both with respect to voice and data. As other questions posed by the Committee herein make clear, voice is increasingly becoming no more than a flavor of data or "an application on the network," albeit an important one to which, as described below in response to those other questions, additional importance must be attached. As noted above, however, universal service will be frustrated, if not defeated, if an entirely unregulated market compels rural consumers to bear the full cost of what AT&T has noted is costly IP transport. Universal service and interconnection must be coordinated – and Congress and the FCC, with input from the states, must assume a leadership role in ensuring that such coordination occurs and is observed on an ongoing basis.

With respect specifically to voice, the importance of voice service is recognized by the FCC, and governing laws and regulation must (a) ensure those services remain available and (b) proscribe actions that may compromise such availability. The FCC noted the primacy of voice service, and the focus on functionality, rather than the underlying technology, when it stated, "Given that consumers are increasingly obtaining voice services over broadband networks as well as over traditional circuit switched telephone networks, we agree with commenters that urge the Commission to *focus on the functionality offered, not the specific technology used* to provide the supported service."[10] As Congress recognized the need to interconnect for telecommunications services, there is no reasoned basis to infer that its treatment would differ depending on the technical protocols beneath it.

As noted above, the FCC brought all telecommunications traffic under section 251(b)(5), including VoIP traffic. Accordingly, the Commission has acknowledged that the state commissions must be part of the process by which resultant interconnection agreements are crafted. The clear statutory provisions that involve state commissions in sections 251 and 252 proceedings provide a useful platform for such involvement. State involvement is imperative in order to ensure that agreements meet local conditions and promote the public interest and universal service.

A recent decision in Michigan illustrates the interest of the states. Addressing a disagreement between Sprint and AT&T Michigan, the state commission disallowed a provision that stated:

> All traffic that Sprint exchanges with AT&T Michigan pursuant to this Agreement will be delivered in TDM format. Sprint may exchange traffic with AT&T Michigan pursuant to a separate agreement, and nothing herein prohibits Sprint from exchanging traffic with AT&T Michigan in IP format pursuant to such an agreement.[11]

The state commission rejected this "contingent resolution," stating it had found previously that pursuant to 251(c)(2), AT&T *must* provide Sprint with IP-to-IP interconnection. Therefore, the parties "must file, for Commission approval or rejection, the agreement by which AT&T Michigan shall provide Sprint with IP-to-IP interconnection." The commission characterized

---

[10] *Transformation Order*, at para. 77 (emphasis added).

[11] *I/M/O Petition of Sprint Spectrum LP for Arbitration Pursuant to Section 252(b) of the Telecommunications Act of 1996 to Establish Interconnection Agreements with Michigan Bell Telephone Company d/b/a AT&T Michigan; I/M/O Joint Submission of Sprint Spectrum LP and Michigan Bell Telephone Company d/b/a AT&T Michigan for Approval of an Interconnection Agreement*, Michigan Public Service Commission, Case Nos. U-17349, U-17569, at 2 (Mar. 18, 2014).

AT&T Michigan's duty to interconnect, no matter the protocol, as "an *on-going interconnection obligation* under Section 251(c)(2)."[12]

Parties who seek to exclude IP interconnection from the purview of any regulatory oversight portend hazard for consumers whose ability to obtain services at fair and reasonable rates, terms and conditions, could be compromised if commercially-negotiated interconnection agreements without a regulatory backstop become the standard. In these regards, and even as the FCC may explore an evolution of obligations for interconnection, the following principles must adhere, regardless of the technology by which interconnection is achieved:

1. Any carrier (including RLECs and other ILECs) should be entitled to initiate interconnection negotiations and arbitrations with any other carrier pursuant to Sections 251 and 252, and without regard to the technical protocols underlying the service.

2. State commissions must invigorate reviews with special attention to Section 251(f)(1) and 252(f)(2) exemptions to ensure that agreements and their impacts are not economically burdensome, are technically feasible, and are consistent with the universal service principles of Section 254.

3. State commissions must ensure a feasible and manageable definition of "network edges," in a manner that is consistent with the standards of Section 251(f)(2) and recognizes the intrinsic, inseverable link between interconnection and universal service.

4. The duty to interconnect shall not impose new onerous obligations on any RLEC with regard to determining the network edge implicated by the "rural transport rule" that is currently part of the interconnection requirements captured by Section 251 and 252.[13]

More broadly, with respect to interconnection for the exchange of data other than voice, Congress and the FCC must consider what interconnection rules and/or support mechanisms are necessary to preserve and advance universal service. If left to "the market" alone and in the absence of any guidance with respect to the contours of reasonable interconnection, there is risk

---

[12] *Id.* at 4 (emphasis in original).

[13] *See, Transformation Order* at paras. 998, 999.

that interconnection arrangements could undermine, rather than preserve and advance, universal service in an IP world.

2. ***Voice is rapidly becoming an application that transits a variety of network data platforms. How should intermodal competition factor into interconnection mandates? Does voice still require a separate interconnection regime?***

In many instances, voice is becoming an application that transmits across a variety of data platforms. Voice over Internet Protocol (VoIP), however, has yet to be classified by the FCC as either an information or telecommunications service. Rather, the FCC has rendered VoIP subject to certain regulatory obligations when such specialized application was determined to be in the public interest. Therefore, VoIP is subject to various regulations of universal service, E-911, and CALEA.[14] The obligations attaching to VoIP in these instances arise without *per se* FCC classification of the service.

These targeted FCC treatments of VoIP demonstrate that the essential element of the service is not the technological protocol underlying it but rather the use, form, and features of the service as delivered to and utilized by the consumer. Therefore, the Committee question might be reformed to ask whether consumer requirements justify or demand regulatory processes to ensure access to voice services. In that regard, and as noted by the Committee, the emergence and existence of intermodal competition may be recognized, but ultimately some structure must be in place to ensure that voice services are available.

---

[14] *See, e.g., Universal Service Fund Contribution Methodology,* WC Docket No. 06-122; *Federal-State Joint Board on Universal Service,* CC Docket No. 96-45; *1998 Biennial Review – Streamlined Contributor Reporting Requirements Associated with Administration of Telecommunications Relay Service, North American Numbering Plan, Local Number Portability, and Universal Service Support Mechanisms,* CC Docket No. 98-171, *Telecommunications Services for Individuals with Hearing and Speech Disabilities, and the Americans with Disabilities Act of 1990,* CC Docket No. 90-571, *Administration of the North American Numbering Plan and North American Numbering Plan Cost Recovery Contribution Factor and Fund Size,* CC Docket No. 92-237, *Number Resource Optimization,* CC Docket No. 99-200, *Telephone Number Portability,* CC Docket No. 95-116, *Truth in Billing Format,* CC Docket No. 98-170, *IP-Enabled Services,* WC Docket No. 04-36, Report and Order and Notice of Proposed Rulemaking, 21 FCC Rcd 7518 (2006) at para. 2, and *Communications Assistance for Law Enforcement Act and Broadband Access and Services (CALEA),* ET Docket No. 04-295, RM-10865, First Report and Order and Further Notice of Proposed Rulemaking, 20 FCC Rcd 14989 (2005) at para. 8.

Absent such regulatory oversight, there is grave risk that voice services may be either (a) unavailable or (b) available but at diminished quality to certain consumers. The threat that services could be rendered unavailable when commercial negotiations break down can be discerned from the various blackouts that attend retransmission consent negotiations in the video market. Notwithstanding the obligations and interests among the parties, the threat of a blackout may be perceived as leverage. Where video entertainment is concerned, outages may be inconvenient but not crucial for life and safety. In the realm of critical voice services, however, Congress and the FCC cannot abide processes that could result in a "voice services blackout."

The need for regulatory oversight in voice is evidenced in recent experience with telephone calls that do not complete in rural areas. Absent firm and enforceable regulatory mechanisms, the market left to itself may well not view rural areas as important enough to serve or may be encouraged to extract unfair advantage in negotiations with small carriers serving even smaller customer bases. The call completion issue is especially troubling because the problems persist even in the face of a declaratory ruling and new rules (the latter yet awaiting OMB approval).[15]

Therefore, the threshold question might not be whether voice requires a separate interconnection regime, but whether voice, as a whole, warrants a comprehensive system of treatment that ensures the delivery of reliable quality voice services to all who desire it, and whether within that comprehensive system interconnection is required on reasonable terms and conditions. And, yet, an expansion of this question warrants mention, as well.

Notwithstanding the unique standing of voice within the universe of communications services, generally, it cannot be argued that other data services relying on IP protocols are

---

[15] *See*, *Rural Call Completion: Report and Order and Further Notice of Proposed Rulemaking*, Docket No. 13-39, FCC 13-135 (2013).

becoming increasingly critical for personal, commercial, and public safety communications. In light of those developments, it is essential that public policy recognizes the need to ensure interconnection that supports the availability of those services to all areas of the Nation and for all who desire it. There is, certainly, sound reason to pay special attention to voice communications, but the need to ensure interconnection for voice should not be mistaken as license to disavow the need for just and reasonable interconnection that enables and supports other services.

3. *How does the evolution of emergency communications beyond the use of traditional voice service impact interconnection mandates?*

As noted above, voice services occupy a unique place in the Parthenon of communications. Therefore, as initiatives to expand the forms of communications between users and public safety answering points (PSAPs) move forward,[16] two points bear recognition: (1) the on-going functionality of emergency voice communications must be assured, and (2) interconnection mandates must ensure users' ability to communicate with emergency personnel in the widest range available. The critical role of voice services in emergency communications, however, must be recognized.

In many instances, voice communications afford people the most instantly accessible and easiest form of communication unburdened by the need to depress keys, buttons, or other mechanisms. A voice call to 911 may be not only be the quickest way to disseminate information, but also a most accurate way for the call center to use the caller's vocal tone, inflection, breathing, and other indicators to shed light on the situation that might not be

[16] *See*, *Facilitating the Deployment of Text-to-911 and Other Next Generation 911 Applications; Framework for Next Generation 911 Deployment: Policy Statement and Second Further Notice of Proposed Rulemaking*, Docket Nos. 11-153, 10-255, FCC 14-6 (2014). Additionally, the FCC is scheduled to consider on August 8, 2014, a *Second Report and Order and Third Further Notice of Proposed Rulemaking* that (a) establishes deadlines for covered text providers to be capable of delivering texts to appropriate 911 PSAPs and (b) seeks comment on potential improvement to current text-to-911 technology (*see*, "FCC Announces Tentative Agenda for August Open Meeting," FCC Press Release dated July 18, 2014).

illuminated in a text or other form of message. One media report quotes a PSAP manager explaining how a person's voice can provide details a text message cannot: "We can keep a person on the line in a domestic situation or something like that and hear things in the background – with text messaging you don't get that. Much is lost in the written word. You don't have voice inflection, background noise or any of that."[17] Moreover, voice is the most broadly available form of communications; text and photos, while capable of being transmitted via many devices, may be characterized fairly as subordinate to voice services that are the predominant function of both (a) wired and (b) fixed and mobile wireless services. That hierarchy, however, does not diminish the value of text, photos, and other forms of information delivered to emergency responders. As noted by an emergency management coordinator, texting may be a preferable alternative "[i]f you are voice impaired, or you are in a hostile environment . . . and you don't feel comfortable making a sound."[18] Therefore, while interconnection mandates must, in the first instance, ensure quality voice connection between users and voice providers, they must recognize as well the growing incorporation of and reliance on such services as text-to-911. Any evolution of interconnection requirements must accommodate and promote evolutions in critical public safety and public interest applications relied upon by users.

4. *Ensuring rural call completion has always been a challenge because of the traditionally high access charges for terminating calls to high-cost networks. Does IP interconnection alleviate or exacerbate existing rural call completion challenges?*

At the outset, it must be noted the rural call completion is not difficult "because of the traditionally high access charges for terminating calls to high-cost networks." Rural call

---

[17] "911 Call Centers Consider Impact of FCC Texting Proposal," Emergency Management (Feb. 12, 2014) (http://www.emergencymgmt.com/safety/911-call-centers-FCC-Texting-Proposal.hmtl) (last viewed August 1, 2014, 14:07).

[18] "Text-to-911 Coming, But Dispatchers Much Prefer Voice Calls," NW News Network (May 16, 2014) http://nwnewsnetwork.org/post/text-911-coming-dispatchers-much-prefer-voice-calls) (last viewed Aug. 1, 2014, 14:21).

completion issues arise even in places where access charges are relatively low, and also with respect to certain kinds of traffic that are no longer subject to access charges. Rather, rural call completion has arisen as a challenge in the past several years because least cost routers identify incentives to *avoid unlawfully* the costs – access charges and/or otherwise – of terminating calls to rural areas. Put another way, it is simply easier and cheaper for certain providers to avoid completing calls to rural areas independent of access charges, and a lack of clear rules and sustained enforcement only enables such conduct. Identifying the root problem as "high access charges" is roughly akin to blaming dairy farmers when a shoplifter takes a gallon of milk. The problem is not in the pricing, but in the users who seek to avoid those costs. As AT&T has noted in recent filings with the FCC,[19] cost issues do not evaporate in an IP environment regardless of whether the ultimate product delivered to the consumer is a per-minute TDM service or a dedicated IP service. The rub of the issue is to identify where costs are incurred, and to ensure due compensation for them.

IP interconnection will not solve call completion issues. Mere changes to the underlying technological protocols used to hand-off of data between networks will not magically improve call quality or completion. Moreover, while some may assert that "public Internet" routing of voice calls could help, all this does in fact is to shift the costs of carrying all voice traffic to rural consumers (once again putting universal service at risk). Additionally, the quality and capabilities of best-efforts "public Internet" routing do not, at least as of yet, meet the standards that consumers expect and need for their voice services. This is particularly true in that "public Internet" routing of traffic does not and cannot achieve the kind of "HD Voice" that appears to be "the promised land" of those who tout the benefits of IP-based connectivity.[20] Instead, call

---

[19] *AT&T Ex Parte* (*see* n.4, above).

[20] See *VCXC Petition for Notice of Inquiry on the Migration to HD Voice* (filed Feb. 25, 2014). On March 7, 2014, VCXC's Petition was associated with GN Docket No. 13-5, and can may be found at

completion, no matter the underlying protocols, must be resolved primarily through rigorous enforcement of existing law and regulation, and the formulation (as the FCC has undertaken) of specialized measures to ensure adherence to statutory principles in the exchange of data of all kinds.[21] These principles include (1) compensation for use of the network (which in an IP environment may transition from a minutes-of-use basis to flat-rate or packet-based pricing structures)[22], and (2) access to services by all users throughout the Nation. Those principles must remain immutable regardless of the technological protocols that underlie the interconnection. Any rewrite of the Communications Act should reinforce statutory authority for the FCC to enforce completion requirements – for data of *all* kinds – to rural areas. In short, rural call completion will persist and continue to serve as "the canary in the coal mine" for what a lack of interconnection and data exchange obligations can mean until firm rules are put into place and enforced against *every* actor/network operator in the chain.

The current call completion issue, in fact, appears to be a harbinger of hazards that may await in a world without rules. Rural carriers are noting an increase in "robo calls." In numerous exchanges among NTCA membership, detailed records of the calls, numbers, and times have been traced in both switches and tandems. Informal conversation with FCC staff confirmed numerous complaints in the industry are being noted, and indications are that these calls are coming over the IP network from overseas. Federal agencies, however, face difficulty controlling and enforcing the problem. The FCC has rigorous tools to enforce against harassment and fraud - the same consumer-oriented focus that is enabled for the public switched telephone network (PSTN) should be present to ensure similar safeguards in an IP-interconnected environment.

---

http://apps.fcc.gov/ecfs/document/view?id=7521089152.

[21] *See*, *Rural Call Completion: Report and Order and Further Notice of Proposed Rulemaking*, Docket No. 13-39, FCC 13-135 (2013).

[22] *AT&T Ex Parte* (*see* n.4, above).

Historically, when the industry transitioned from trunk connections, to T-1 connections, and then to ATM, no sweeping overhaul was required. Similarly, none should be necessary as the industry evolves to IP. Policies should ensure the many entrances and exits in the IP network do not mask users' ability to abuse.

5. ***Should we analyze interconnection policy differently for best-efforts services and managed services where quality-of-service is a desired feature? If so, what should be the differences in policy between these regimes, and how should communications services be categorized?***

Interconnection policy may recognize differences between "best-effort" and "managed services," but public policy should ensure two outcomes: (1) universal consumer access to quality service, and (2) carrier access to quality interconnection on just and reasonable terms. Policies should avoid outcomes in which a service provider that seeks to provide high quality service to the end user is denied just and reasonable interconnection to enable that result. That sort of outcome could skew competition and produce results adverse to consumer interests. By way of example, if Company A manages facilities at both the "end-user" and "network" segments, and Company B seeks interconnection to provide only "end user" service in the same market, Company A should be precluded from offering Company B only a level of interconnection that would enable Company B to provide "best effort," rather than "managed," services. Indeed, in an IP environment, anything that is not managed (*i.e.*, managed with priority) becomes a "best effort" service.

This preclusion, however, should be subject to the same sort of section 251(f) exemptions currently set forth in the Communications Act. Congress was mindful in section 251(f) to exempt, modify, or suspend certain interconnection obligations that might otherwise apply to RLECs. Notwithstanding the prevailing interest in promoting competition throughout the Nation, Congress recognized wisely that smaller rural carriers face special circumstances that would warrant mandatory interconnection measures more harmful than helpful.[23] Therefore, Congress stipulated that RLECs are exempt from interconnection and other obligations under section 251(c) until a bona fide request has been received and the state commission confirms that the request is "not unduly economically burdensome, is technically feasible, and is consistent with section 254 . . ."[24] Similarly, section 251(f)(2) of the Act permits an RLEC or any other smaller incumbent to seek suspension or modification of obligations under subsections (b) and (c) of Section 251 where it can be shown that such a request is necessary in order to avoid, *inter alia,* "a significant adverse impact on users of telecommunications services, generally; [or] a requirement that is unduly economically burdensome."[25] Although Section 251(f)(2) specifies an adverse use on "users of telecommunications services," NTCA submits that a developed record may justify targeted application of similar principles to avoid adverse impacts on users of information services.

6. *Much of the committee's focus in the #CommActUpdate process has been on technology-neutral solutions. Is a technology-neutral solution to interconnection appropriate and effective to ensure the delivery and exchange of traffic?*

The solution to ensuring the delivery and exchange of traffic should apply regardless of the technology utilized. Toward that end, "technology neutral" approaches that look toward securing fulfillment of public policy principles are appropriate. By way of example, in a recent

---

[23] *See, e.g.,* Sen. Rep. 104-23 at 22 ("The Committee intends that the FCC or a State shall, consistent with the protection of consumers and allowing for competition, use this authority to create a level playing field, particularly when a company or carrier to which this subsection applies faces competition from a telecommunications carrier that is a large global or nationwide entity that has financial or technological resources that are significantly greater than the resources of the company or the carrier"); 142 Cong. Rec. S 709 (Feb. 1, 1996) (Comments of Sen. Daschle of South Dakota) ("The bill before us also recognizes the important role that must be played by Public Utilities Commissions (PUCs) in rural States. PUC's are the best entities to judge whether a given market within their State can support competition. That's not a judgment we should make from Washington. Nor is it something we can or should leave to the unbridled, unspecified judgment of the private sector. Those who have taken the risks and made investments to extend cable or phone services to smaller rural communities should not be placed at risk of being overwhelmed by larger, better-financed companies.")

[24] 47 U.S.C. § 251(f)(1).

[25] 47 U.S.C. § 251(f)(2).

FCC proceeding, diverse parties, including RLECs, cable companies, large ILECs, CLECs, and state public service commissions,[26] all concur that sections 251 and 252 provide no basis to conclude – or even opportunity to argue – that the technological means of interconnection is relevant to a determination of whether the statutory framework applies. Section 251 does not distinguish between IP networks and other networks, nor does anything in section 252 indicate that the just and reasonable rate for interconnection of facilities and equipment should not or cannot be determined in instances where the equipment happens to be IP-enabled. Rather, section 251(a) plainly deals with interconnection between the "facilities and equipment of . . . telecommunications carriers" generally, and section 251(c) refers generically to interconnection of "facilities and equipment" and a "network." Under current conditions, then, the issue of technological neutrality is answered by the prevailing evaluation as to whether the parties are telecommunications carriers – since, if they are, the duty to interconnect applies regardless of the underlying technology.

In 2012, NTCA, along with other trade associations, urged the FCC to develop a complete record before issuing rules specific to IP interconnection, but noted that in the absence of IP-specific rules, the "existing statutory framework and the associated time-tested rules should (and indeed, must) continue" to apply to interconnection between carriers, whether IP or otherwise.[27] In summary, a policy can hardly be characterized as "technology-neutral" if IP-

---

[26] *Connect America Fund; A National Broadband Plan for Our Future; Establishing Just and Reasonable Rates for Local Exchange Carriers; High-Cost Universal Service Support; Developing a Unified Intercarrier Compensation Regime; Federal-State Joint Board on Universal Service; Lifeline and Link-Up; Universal Service Fund – Mobility Fund:* Federal Communications Commission Docket Nos. 10-90; 09-51; 07-135; 05-337; 01-92; 96-45; 03-109; 10-208, *Comments of Windstream* at 8; *Comments of Frontier Communications* at 12-14; *Comments of ITTA* at 8, 9; *Comments of CompTel* at 13-13; *Comments of XO* at 12; *Comments of Regulatory Commission of Alaska* at 11; *Comments of California Public Utility Commission* at 9; *Comments of Public Service Commission of Wisconsin* at 8-10; *Comments of Indiana Regulatory Commission* at 7, 8 (hereinafter *Connect America Fund, et al.*) (Feb. 24, 2012).

[27] *Connect America Fund, et al.* Comments of NECA, NTCA, OPASTCO, and WTA, at 23 (Mar. 30, 2012).

based networks are, merely because of their processing protocol, treated differently than every other network that came before them.

7. *Wireless and internet providers have long voluntarily interconnected without regulatory intervention. Is this regime adequate to ensure consumer benefit in an all IP world?*

IP interconnection is, within the broad scope it is anticipated to achieve, at a nascent state. The past or current state of events may prove insufficiently instructive to the structure that ultimately should be imposed when broadband and data traffic surge to the forefront of what consumers and network operators face on a prevalent, if not all encompassing, basis. In all respects, however, experience demonstrates that the market yet requires regulatory backstops to ensure equitable access to interconnection across the industry. This result is not a failure of the market, but is rather reflective of natural incentives and tendencies that permeate nearly every marketplace. The plethora of federal and state bodies regulating highways, pharmaceuticals, food, power, aviation, and even weights and measures evidence the reality that, free from regulatory oversight, the incentives to embark on routes harmful to consumers and the public interest may be too enticing for certain players to resist.

In the world of communications, the regulations are not necessarily aimed at preserving health and safety (though access to emergency services may fall within that ambit) but may be more accurately defined as ensuring the promotion of policies aimed at bettering the public interest. Toward this end, certain natural market incentives and tendencies may discourage voluntary actions in the marketplace. But, where access to services, or competition, is identified as a core public interest, a regulatory process that ensure fairness across the field (*i.e.*, access for one party and fair compensation for the other) may be necessary to help craft the marketplace best suited to the recognized public interest.

There is no compelling reason to believe that a change in the underlying protocol from TDM to IP would affect the incentives of various participants and encourage negotiations that result naturally in agreements which promote public interest policies. Rather, although in some circumstances, and perhaps after some further development of the marketplace and technology, the universe may evolve to a form in which solely voluntary agreements are sufficient, in the immediate term the methodologies set forth by the Act are appropriate. These methodologies blend the best of several interests – they begin with an opportunity for voluntary negotiations, and are attended by a process to "fall back" to arbitration if negotiations fail; the statute also sets discernible standards that balance the various interests of marketplace competition with the real situation of small rural carriers, and permit those closest to the parties (the state commissions) authority to adjudge the agreements. Finally, these processes are both time- and litigation-tested, so that the parameters of the process, *without regard to the technical protocols underlying the agreements*, are known and can inform well the market transition to an all IP environment. If as that transition occurs the invocation of mediation and/or arbitration diminishes and that process becomes a dusty, unused relic, then the law may conclude that voluntary agreements are sufficient. However, until experience and a well-formed evidentiary record present otherwise, IP interconnection should remain subject to the existing standards and processes of sections 251 and 252 of the Act.

8. *Is contract law sufficient to manage interconnection agreements between networks? Is there a less onerous regulatory backstop or regime that could achieve the goals of section 251?*

In a "perfect market," market forces would be sufficient to promote beneficial activity while constraining either inefficient or inappropriately discriminatory interactions. Given the respective market positions of large and small telecommunications carriers, Congress wisely crafted an approach to mandatory interconnection that offers opportunity for negotiations free of regulatory intervention, yet with several regulatory "backstops" to ensure the implementation of important public policy initiatives. These measures operate to both open and protect various markets in order to promote beneficial competition while bolstering incentives to ensure universal service in rural and other challenging areas. On the one side, the negotiation provisions enable free and open commercial negotiations. On the other side, if, the rates, terms, and conditions are unpalatable to either party, then the contract is presented for objective third-party review via arbitration in which the regulatory process can ensure important public policy objectives are fulfilled. Moreover, the Section 251(f) exemptions provide a backstop to preempt potential actions that could disable fundamental universal service initiatives in rural and other high-cost local areas.

The inclusion of IP traffic in interconnection obligations, as noted above, has no bearing on the existing regulatory construct that surrounds interconnection. Moreover, the same challenges and opportunities, benefits and potential pratfalls, and market-power driven incentives that validate the current statutory construct remain as applicable in an IP infused environment as in a TDM-based world. As noted above, the principles promoted by the 1996 Act are aimed at improving the consumer experience and ensuring access to services throughout the Nation. Consumers care that quality service is delivered, with functions that are useful, beneficial, and affordable. NTCA submits that the kernel of the consumer experience is the actual end-user experience when using the service, and that contemplation of the technical protocols underlying the service are essentially immaterial so long as access and quality of service are maintained. In that regard, the regulatory processes aimed at bringing that experience to consumers remain valid to the extent they operate independent of the underlying technology. Therefore, if commercial incentives, rural market conditions, and trends in interconnection relationships remain essentially unchanged whether the underlying protocols are TDM or IP, then the legal obligations of the parties and governing regulatory processes surrounding interconnection should remain unchanged, as well. In these regards, the use of tariffs and interconnection agreements subject to regulatory oversight can be viewed as appropriate measures. Tariffs, specifically, eliminate certain administrative burdens that are associated with individually-negotiated interconnection agreements, especially where the amount of traffic exchanged may not justify the cost of negotiations.

By contrast, tariffing mechanisms have been utilized for decades and are well understood within the industry. Ultimately, any manner of flexible and creative approach might be appropriate – for example, tariffs that blend separate rates for connections and bandwidth, or which differentiate between "managed services" and those implicating quality-of-service (QoS) pricing arrangements. In short, "tariffs" need not mean only what they did in the past, but can rather become a useful means of making known to the market what terms and conditions are available and deemed just and reasonable, with there being an opportunity to negotiate where the market desires to do so.

Ultimately, the technological efficiency of IP interconnection should be reflected in interconnection and compensation agreements that govern it. Therefore, wholesale changes that eliminate beneficial policies should be avoided. If in fact the market evolves to reflect the technological efficiencies of IP, then companion administrative benefits should accrue naturally. If, however, that reflective evolution does not occur, then a regulatory backstop will ensure critical access to services for consumers and providers.

IV. CONCLUSION

The Act mandates that interconnection among telecommunications carriers occur pursuant to sections 251 and 252, regardless of the underlying technology. Even if Congress determines ultimately that IP interconnect may be suited to a different form of regulatory oversight, existing mechanisms that promote universal service principles provide a sound base for current action. This regulatory backstop is necessary because naturally-arising interests and market forces do not create sufficient ground to ensure interconnection at just and reasonable rates, terms and conditions. Congress and the FCC are commended to construct a record of market development and base any future revisions upon data collected as IP traffic becomes the prevalent form throughout the market. Until that time, however, existing section 251 and 252 provisions apply to IP interconnection and should be maintained in order to fulfill universal service goals.

August 8, 2014

Hon. Fred Upton
Chairman
Energy and Commerce Committee
US House of Representatives
2125 Rayburn House Office Building
Washington, DC 20515

Hon. Greg Walden
Chairman
Communication and Technology Subcommittee
Energy and Commerce Committee
US House of Representatives
2125 Rayburn House Office Building
Washington, DC 20515

Octoshape Comments to Communications Act Update

Chairmen Upton and Walden,

First and foremost thank you for your interest, efforts, and request for comments in the investigation of this very complex topic.

Octoshape innovates in the area of audio and video transmissions over IP networks. We focus on multiscreen or connected device broadband content delivery, broadcast quality contribution video to cable, satellite and IPTV headends, and video distribution to connected devices in the enterprise, all via the public Internet. This video data flows across Internet Service Provider infrastructure that is likely directly or indirectly impacted by legacy rulings or new amendments to the Communications Act.

It is clear that consumer media consumption is shifting from traditional methods like broadcast TV over cable, satellite and IPTV, to broadband device consumption models. The most recent Cisco VNI report[1] suggests that 79% of all consumer Internet traffic will be IP Video by 2018. The report also suggests that mobile traffic will grow three times faster than fixed IP traffic from 2013 to 2018[1]. Traffic growth in mobile networks is particularly more concerning to operators as the capacity on these networks is far more constrained than fixed networks. The rate of high quality content consumption is clearly outpacing capacity capital investment in the last mile. Octoshape was asked to contribute a white paper on the topic to the European Broadcast Union [2].

We find ourselves standing in a moment very similar to the conversion of Analog to Digital. A time where spectrum cost, availability, and efficiency is crucial to businesses that will utilize the infrastructure. A time where innovation is required for business to evolve and break through its current natural boundaries defined by the underlying technology and infrastructure. A time where a breakthrough in technology naturally creates and serves new consumer demand that had been previously hidden and sometimes blocked by traditional constraints. In these times ecosystems often require disruptive technologies that are not burdened with the weight of legacy deployment and mindshare. In some cases legislation has created these flourishing environments, usually in the infancy of an infrastructure deployment, but often legislation cannot keep pace with the technology advancement, thus producing a chilling effect in investment and innovation as a market matures.

Octoshape creates technologies that bridge the evolution of Broadcast TV to Broadband TV. These technologies focus in two main areas, video quality and video efficiency, over first, middle and last mile broadband networks which include fixed, wireless, and mobile implementations. Octoshape technologies increase the consumer quality experience, while at the same time expand the deployment lifetime, and reduce the traffic loads on existing capital deployments of the Internet Service Providers. These efficiencies give way to new business models that mimic the economic structures of traditional Broadcast TV.

Perhaps the most important topic relevant value proposition resulting from Octoshape innovations is the network efficiency gains to the operators:

1) Core Protocols: Octoshape technologies ingest standard streaming video formats that are inherently inefficient and transparently convert them for transport over the Internet in a more efficient form. Independent testing by our customers has shown up to 25% less data transmitted in comparison with traditional technologies over best effort networks with latency and packet loss. Efficiencies of this type are important to wireless providers that need to save precious mobile spectrum, or cable, satellite and IPTV providers that need to make very tough decisions about spectrum allocation between TV or broadband services.

2) Built for Cloud and Virtualized environments: Traditional content distribution technologies are heavily restricted by distance or latency over the Internet. In order to achieve high quality video experiences, server infrastructure must be distributed very close to the end user. Therefore content distribution service providers or last mile operators need to deploy many servers at the edge of their respective networks often leading to inefficient clustering of resources, and ultimately overspending on capital infrastructure. Cloud technologies are deployed in a more centralized architecture and can facilitate elasticity on-demand over multiple products. The Octoshape technologies are structured and arranged in a way to be dynamically deployed over these virtualized environments. Resiliency mechanisms built into the core protocols provide the end consumer a TV quality video experience regardless if the underlying infrastructure resources are expanding or contracting. Resources no longer need to be “close” from a geographic or network latency perspective giving operators greater flexibility in designing more efficient networks. Operators can now get products to market faster, make smarter capital investments that have a longer lifespan, leverage other companies capital investments in cloud infrastructure, and implement architectures that have far less impact on the environment from an energy consumption perspective.

3) Leveraging Existing Infrastructure: Octoshape has also designed the video distribution systems to leverage IETF (Internet Engineering Task Force) standards for efficient delivery of data. Octoshape enables operators to leverage non-proprietary mechanisms in their existing infrastructure like Native and AMT (Automatic Multicast Tunneling) Multicast to deliver video data in a dramatically more efficient way over the Internet. These mechanisms which have been deployed in production with a major telecommunication operator in the United States[3] since late 2010 result in as much as 85% traffic reduction across the majority of the operator network infrastructure.

Due to the unique components of our technology that eliminate the relationship between server distance and video quality, Octoshape enjoys a global perspective on the video distribution industry with broadcast and telecommunication customers in all the major regions, North America, LATAM, EMEA, and APAC. We were fortunate enough to be invited by the President of France to the eG8 forum in May of 2011 where entrepreneurs and executives were gathered to advise politicians attending if regulation was necessary on the Internet. Massive increases in broadband video consumption are driving public policy discussion globally on every topic from consumer privacy to traffic discrimination. As more users are connected, and content distribution and consumption becomes less geographically bound, we see that the world becomes a much smaller place. We believe this “era of accessibility” will ultimately weigh into our own domestic communications policy in ways that we cannot fathom today.

Octoshape recognizes that the industry of communication has and will continue to experience exponential change in technological advancement. Advancements in some cases are so dramatic that they fundamentally change operational economics, consumer behavior, and consequently business strategy in ways that could not be foreseen even from the most recent amendments to the Communications Act. We recognize that policy makers have created regulation intending to spur infrastructure investment or to protect against monopolistic or consumer predatory pricing behavior. In the drive toward the next evolution of the ecosystem, we would respectfully urge the Commission to lean toward less regulation in the market to promote free market capitalism, which will drive investment and thus further innovation. Seen from our vantage point a regulation-free interconnection marketplace in this phase of the ecosystem works quite well today and does not require further regulatory intervention.

Respectfully,

[redacted]

Scott Brown
VP Strategic Technology Partnerships
Octoshape

References:

[1] Cisco Visual Networking Index: Forecast and Methodology, 2013–2018
http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html

[2] EBU White Paper Submission: Video Consumption Dramatically Outpaces Internet Capacity
http://bit.ly/1oBr6WC

[3] AT&T Innovation Space Blog: First Ever Open-Web Multicast Event
http://bit.ly/1yjN4Dq

RESPONSE OF THE OREGON TELECOMMUNICATIONS ASSOCIATION (OTA) TO HOUSE ENERGY AND COMMERCE COMMITTEE

Modernizing the Communications Act
NETWORK INTERCONNECTION QUESTIONS FOR STAKEHOLDER COMMENT

Due Date of August 8, 2014

We offer responses to the questions posed by the Committee by focusing on four major points, and reference individual questions as appropriate in each section. Our four sections for this fourth white paper focusing on network interconnection issues are as follows:

THE ROLE FOR THE FCC IS TO REGULATE PURSUANT TO TITLE II ALL TRANSPORT AND TRANSMISSION CAPACITY OFFERED ON UNDERLYING NETWORKS

RURAL CALL COMPLETION ISSUES WILL STILL REQUIRE REGULATORY ATTENTION IN AN IP WORLD

CHANGES IN NETWORK INTERCONNECTION SHOULD NOT BE USED AS A REASON FOR LARGE URBAN CARRIERS TO IMPOSE UNREASONABLE COSTS ON RURAL CARRIERS

EMERGENCY COMMUNICATIONS REQUIRE SPECIAL ATTENTION IN POLICY FORMATION

We appreciate the opportunity to offer input on these network interconnection issues and look forward to the remaining white papers that the Committee intends to release during 2014.

THE ROLE FOR THE FCC IS TO REGULATE PURSUANT TO TITLE II ALL TRANSPORT AND TRANSMISSION CAPACITY OFFERED ON UNDERLYING NETWORKS

This first section of our response addresses White Paper 4 questions[1] one, two, five and six. The Committee must determine in its effort to modernize the Communications Act whether the desire to protect consumers in a broadband paradigm is still a relevant public policy concept. We suggest that it remains an important foundation for any attempt to enact a forward-looking **national** public policy.

In order to ensure seamless interconnection across the entire network of networks, the OTA endorses the concept introduced on July 18, 2014 by NTCA – The Rural Broadband Association in the Federal Communications Commission (FCC) Open Internet proceeding (GN Docket No. 14-28). NTCA suggested correctly that one of the roles that the FCC should continue to play is to ensure that with respect to the **transport and transmission capacity on all networks over which data travel,** the networks will 1) be interconnected on just and reasonable terms; 2) that the Commission has the clear and unquestionable ability to step in when those networks do not interconnect seamlessly, and 3) that the important public policy goals of consumer protection, universal service, competition, and public safety are not threatened by the unjust and unreasonable acts or omissions of any given network operator.

[1] Q1: In light of the changes in technology and the voice traffic market, what role should Congress and the FCC play in the oversight of interconnection? Is there a role for the states? Q2: Voice is rapidly becoming an application that transits a variety of network data platforms. How should intermodal competition factor into interconnection mandates? Does voice still require a separate interconnection regime? Q5: Should we analyze interconnection policy differently for best-efforts services and managed services where quality-of-service is a desired feature? If so, what should be the differences in policy between these regimes, and how should communications services be categorized? Q6: Much of the committee's focus in the #CommActUpdate process has been on technology-neutral solutions. Is a technology-neutral solution to interconnection appropriate and effective to ensure the delivery and exchange of traffic?

NTCA offered a very narrow application of Title II as a means to this important end. Since some members of this Committee and the House as a whole have expressed concern about Title II, let's be specific about what is **not** proposed in this regard. **What is not proposed is any overarching Title II authority**. In the NTCA proposal, Title II would not be appropriately applied to any services that are offered atop of these regulated networks unless those services would qualify for that treatment on their own merits. There is no hidden agenda to "regulate the Internet", but rather a need to avoid a problem larger than rural call completion presents today. As NTCA stated at page 11 of its FCC filing: "*Applying Title II to networks that merely transmit data between points does not hinder innovation in 'the Internet ecosystem' or saddle new services with legacy regulations. Rather, it simply ensures that the networks upon which Internet data travel will be interconnected on reasonable terms, that the Commission can step in when those networks do* **not** *interconnect seamlessly, and that important public policy goals of consumer protection, universal service, competition, and public safety are not threatened by the unjust and unreasonable acts or omissions of any given network operator.*"

What is proposed by NTCA is a **limited and targeted** application of Title II regulation **specifically and only with respect to** the transport and transmission capacity of all networks and content delivery networks (CDNs) involved in the transmittal of data between points (the network layer as opposed to the service layer). What does this accomplish? Several positives accrue to such an approach. Sections 201 and 202 provide a solid basis for carefully crafted rules to protect consumers. In Section 201, rules are in place that require service to be provided upon reasonable request and details a carrier's duty to interconnect. Section 202 provides rules that prohibit unjust and unreasonable discrimination. Sections 206, 207, and 208 provide a backdrop for the resolution of complaints and offer enforcement mechanisms. Title II as applied in this context would ensure that consumer connectivity is not lost or impaired due to a disagreement or dispute between underlying network operators.

RURAL CALL COMPLETION ISSUES WILL STILL REQUIRE REGULATORY ATTENTION IN AN IP WORLD

Question 4 in this white paper[2] focuses on the troubling issue of rural call completion. At the end of the background section of this 4th White Paper, the Committee offers the observation that "*the federal government has been reluctant to engage in disputes regarding IP interconnection*." While it can be argued that some progress has indeed occurred, the FCC has been reluctant to bring to an immediate END the problem of rural call completion for several OTA members. It has been over three years since the FCC was made aware of the magnitude of this problem in a series of ex parte[3] letters, and reports from operating companies indicate the problem continues today.

The very fact that the problem is still not solved despite FCC attempts to enforce existing rules is instructive for this Energy and Commerce Committee inquiry. The "market" is not producing an equitable solution for rural customers in the nation. A need remains for targeted and reasonable regulation to protect both the rural consumer experience as well as promote universal service and public safety. Even with rules in place, problems persist. Removing basic rules would lead on a path to even larger problems and serve to thwart progress to achieving a truly national broadband platform.

---

[2] Ensuring rural call completion has always been a challenge because of the traditionally high access charges for terminating calls to high-cost networks. Does IP interconnection alleviate or exacerbate existing rural call completion challenges?
[3] Ex parte letter from NTCA CEO Bloomfield to FCC Chair Genachowski dated September 20, 2011, following up on a June 13, 2011 letter to FCC Enforcement Bureau from NTCA counsel and other parties.

As the Committee seeks to modernize the Communications Act, we respectfully request that public safety and universal service issues not be a secondary thought to how fast large carriers can shed regulatory oversight. Such an approach is a vital piece of a balanced public policy solution that will meet the customer needs of OTA's members.

CHANGES IN NETWORK INTERCONNECTION SHOULD NOT BE USED AS A REASON FOR LARGE URBAN CARRIERS TO IMPOSE UNREASONABLE COSTS ON RURAL CARRIERS

With the eighth question[4] posed in this White Paper 4, the Committee addresses the concept of allowing means other than current FCC rules to oversee the interconnection process.

In determining how best to modify federal legislative policy, it is important to review recent interconnection location proposals by AT&T and the implications that such a change would have on rural carriers and customers. In a January 24, 2014 ex parte[5] with the Federal Communications Commission, AT&T suggests that the model for both Tier 1 IP voice and peering interconnection is 5 to 8 interconnection points in total for the entire country. Under this proposed scenario, the use of fewer interconnection points covering much larger geographic areas would result in a significant increase in costs on rural ISPs and ultimately rural consumers and business customers. This increase is caused by the smaller providers having the full responsibility for transporting traffic to interconnection points a great distance from their facilities in such a proposed arrangement, in many cases over facilities owned by large carriers such as AT&T.

[4] Is contract law sufficient to manage interconnection agreements between networks? Is there a less onerous regulatory backstop or regime that could achieve the goals of Section 251?

[5] AT&T's Director – Federal Regulatory filed an ex parte letter in GN Docket No. 13-5, WC Docket Nos. 13-97 and 10-90 that showed the essence of its proposal at presentation slide 11.

We submit to this Committee that underlying networks are not "free" in an IP-enabled paradigm any more than they are in a legacy TDM world. Small or rural ISPs possess little or in most situations NO bargaining power with respect to negotiating interconnection terms with large national operators. Without something resembling "rules for the IP road" in place, such cost shifts will not be borne equitably across the networks and such an outcome will serve to drastically impact the goal of universal service in our emerging broadband world.

EMERGENCY COMMUNICATIONS REQUIRE SPECIAL ATTENTION IN POLICY FORMATION

Question 3 in this fourth white paper[6] states: "How does the evolution of emergency communications beyond the use of traditional voice support impact interconnection mandates?"

Just like in real estate, it's location, location, location. The more remote the location, the more crucial the access to emergency services. In parts of Oregon, several OTA members offer service to customers in rugged, isolated terrain. They are the lifeline for these customers.

We encourage this Committee to keep in mind the needs of both urban residents and voters that live in the rural regions of the country as emergency communications network interconnection is reviewed.

---

[6] Question 7 is addressed tangentially here as well: "Wireless and Internet providers have long voluntarily interconnected without regulatory intervention. Is this regime adequate to ensure consumer benefit in an all-IP world?"

Packet Clearing House
572 B Ruger Street, Box 29920
The Presidio of San Francisco
San Francisco, California
94129-0920 USA
+1 415 831 3100 main
+1 415 831 3101 fax

Chairman Fred Upton
House Energy and Commerce Committee

Chairman Greg Walden
House Communications and Technology Subcommittee

Friday, August 8, 2014

Packet Clearing House Response to the Network Interconnection Inquiry

Chairmen Upton and Walden:

Packet Clearing House is the international organization responsible for providing operational support and security to critical Internet infrastructure, including Internet exchange points and the core of the domain name system. PCH was formed in 1994, and is supported by more than five hundred companies in the Internet industry and more than twenty governments, including that of the United States.

PCH's first and most important task is the building and support of Internet Exchange Points. IXPs are the places where Internet networks interconnect, to exchange traffic between their customers and create Internet bandwidth. PCH assists in the construction of IXPs, and provides equipment, training, and ongoing support to IXPs, their operators, and their Internet Service Provider participants. Over the past twenty years, PCH has helped build more than 150 IXPs, and has provided assistance to more than half of the world's 400 IXPs. PCH also maintains the global directory of IXPs and publishes statistics about their use and growth.

PCH has also conducted the only global survey of Internet network interconnection agreements, analyzing more than 142,000 such agreements, representing 86% of all the world's Internet carrier networks. I've attached the report, *Survey of Characteristics of Internet Carrier Interconnection Agreements*, also available online here: https://www.pch.net/resources/papers//peering-survey/PCH-Peering-Survey-2011.pdf. That survey was input to the OECD's 2013 Digital Economy paper *Internet Traffic Exchange: Market Developments and Policy Challenges*, which I co-authored, also attached and available online here: http://www.oecd-ilibrary.org/internet-traffic-exchange_5k918gpt130q.pdf.

Our twenty years of working with more than 100 governments on Internet regulation leads us to conclude that regulating Internet network interconnection, or "peering," would only address a symptom, rather than the underlying problem. **Internet interconnection regulation would vastly increase the overhead expense of interconnection** (critical, because the number of interconnections is so much higher than in voice networks) and would hamper the industry's ability to limit the so-called "bulletproof" networks that are created to serve spammers and online criminals, **while doing nothing to solve the underlying problem: terminating monopolies abusing their customers**, which is fundamentally not a problem of interconnection, but of lack of competition. Only the discipline of competition in the marketplace will reward good actors and punish bad actors. **Instead of focusing your efforts on the symptoms, we ask you to address the problem: lack of competition for terminating last-mile monopolies.**

Thank you for your attention to this critical issue, [redacted]

Bill Woodcock
Executive Director
Packet Clearing House

1a) In light of the changes in technology and the voice traffic market, what role should Congress and the FCC play in the oversight of interconnection?

Congress and the FCC should advocate for consumer interests, but that should take the form of ensuring that every consumer has a choice between competitive providers of last mile bandwidth. *Real* last mile competition means the ability for new market entrants to make direct infrastructure investment, laying new fiber and creating new bandwidth, not just renting leftover capacity from incumbents and hoping to mark it up and resell it. That's not competition, that doesn't build new infrastructure, and it doesn't lower prices. Regulating interconnection is a sideshow, patching a symptom while diverting attention from the actual problem: lack of last mile competition.

1b) Is there a role for states?

No. The Internet and most Internet service providers operate across state or national borders. To the degree that universal regulatory harmonization can be advanced, Internet service providers enjoy greater latitude to focus on providing good service at a competitive price and waste less time chasing diverse regulatory requirements. This is an area where Federal preemption is absolutely necessary; Internet regulation state-by-state would create expensive chaos, and make the Internet more expensive and less accessible.

2a) Voice is rapidly becoming an application that transits a variety of network data platforms. How should intermodal competition factor into interconnection mandates?

Any interconnection mandates, in the unfortunate case that they become necessary, should be absolutely neutral, and should not attempt to distinguish between or classify networks based upon technology, business model, service area, customer base, or any other arbitrary criterion. In the Internet, a network is a network; how it's used is often outside the direct control of its operator, and is subject to rapid change as user preferences and demands evolve. "Modes" should not be distinguished between, nor should any "mode" receive special, preferential, or different treatment.

2b) Does voice still require a separate interconnection regime?

Absolutely not. Voice, like video, email, and web browsing, is an application, which runs over the top of, and transparently to, the network. Priority voice access to emergency services is the only area in which technology-specific mandates applying to voice services are appropriate.

3) How does the evolution of emergency communications beyond the use of traditional voice service impact interconnection mandates?

It does not. These are independent issues. Interactive public access to emergency services may need to evolve to keep pace with users' expectations. Emergency responders should stop relying solely on twentieth-century communications technologies, while at the same time, providers of innovative communications applications may need to make special accommodation for gateways to legacy emergency services PSTN numbers. These issues are very real and require attention, but do not intersect network interconnection issues.

4) Ensuring rural call completion has always been a challenge because of the traditionally high access charges for terminating calls to high-cost networks. Does IP interconnection alleviate or exacerbate existing rural call completion challenges?

Internet services have not historically had access to universal service funds, and Internet service providers do not pay each other settlements, so Internet access in rural areas has been largely unsubsidized: slower and more expensive than in dense urban areas which can be more efficiently served. The pre-industrial-revolution public policy goal of incentivizing urban-to-agricultural migration is at least a century overdue for revision, and the Internet has not, to date, been burdened with the task of subsidizing people to turn farmland into suburbs. As farmland is now in much shorter supply than farmhands, it's probably best that we don't so burden the Internet in the future.

5a) Should we analyze interconnection policy differently for best-efforts services and managed services where quality-of-service is a desired feature?

"Quality of Service" is the last resort of an underbuilt network: discarding packets that have already been paid for by customers. Intentional failure to deliver a paid-for service is a matter of commercial fraud rather than of communications technology. This deserves serious attention, but cannot be solved through regulation of interconnection, which is not the location of the problem.

5b) If so, what should be the differences in policy between these regimes, and how should communications services be categorized?

Communications services should not be categorized. A network is a network. How customers choose to use that network should not afford a network special preferential treatment, nor result in its penalization.

6) Much of the committee's focus has been on technology-neutral solutions. Is a technology-neutral solution to interconnection appropriate and effective to ensure the delivery and exchange of traffic?

Absolutely. Global best-practice in Internet service regulation is generally thought to be an Internet service class license term to the effect that a regulated entity *under normal circumstances, deliver any packet addressed to a domestic destination without passing it across the national border*. This is a simple, technology-neutral regulation, which does not attempt to accord special privileges to networks based upon categorizations like "access" or "content" or "edge," applying equally to all. It gives a regulatory lever to resolve disputes and depeerings, while simultaneously encouraging a number of business and economic best-practices, and creating no additional burden whatsoever on any rationally-operated network. An implementation of this practice in a country of large geographic scope like the United States would probably look much like Sweden's implementation, reducing the boundary from the national borders to LATA-like subdivisions of the country.

7) Wireless and Internet providers have long voluntarily interconnected without regulatory intervention. Is this regime adequate to ensure consumer benefit in an all- IP world?

Yes. Wireless is not the problem area; wireless networks have typically interconnected with the rest of the Internet without difficulty. Market-dominant wireline access networks are the problem area, where consumers are being victimized.

8a) Is contract law sufficient to manage interconnection agreements between networks?

Yes, and further, the unregulated market's drive toward globally-uniform connection terms has been spectacularly successful, far beyond the scope or capacity of any regulatory body to effect. National regulators typically have jurisdiction over a few regulated telecommunications carriers within their country, but this constitutes only a minority of the Internet networks that interconnect… Academic, governmental, not-for-profit, commercial networks outside the telecommunications industry, and providers of unregulated services all interconnect as well. The most important finding of the 2011 survey was that 99.73% of *all* of the world's networks, not just regulated telecommunications service providers, enjoy an entirely uniform understanding of the commercial and legal terms of interconnection. This understanding is so well harmonized that 99.51% of interconnection agreements are not even formalized in written form. Regulation cannot hope to achieve a result with this degree of uniformity, since no regulator has authority over any significant share of the global industry. In this regard, the unregulated market has produced an outstanding success.

8b) Is there a less onerous regulatory backstop or regime that could achieve the goals of section 251?

As mentioned in the answer to (7), there is already exists a global regulatory best-practice, which the United States could follow, which would serve consumer and national security interests while avoiding technology-specific regulation or special categorization and treatment of different networks.

Packet Clearing House
572 B Ruger Street, Box 29920
The Presidio of San Francisco
San Francisco, California
94129-0920 USA
+1 415 831 3100 main
+1 415 831 3101 fax

Survey of Characteristics of Internet Carrier Interconnection Agreements

Bill Woodcock & Vijay Adhikari[1]
Packet Clearing House
May 2, 2011

Introduction

The Internet, or network of networks, consists of 5,039 Internet Service Provider (ISP) or carrier networks, which are interconnected with one another in a sparse mesh.[2] Each of the interconnecting links takes one of two forms: transit or peering. Transit agreements are commercial contracts in which, typically, a customer pays a service provider for access to the Internet; these agreements are most common at the edges of the Internet. Transit agreements have been widely studied and are not the subject of this report. Peering agreements – the value-creation engine of the Internet – are the carrier interconnection agreements that allow carriers to exchange traffic bound for one another's customers; they are most common in the core of the Internet. This report examines and quantifies a few of the characteristics of Internet peering agreements.

The Survey

In preparing this report, we analyzed 142,210 Internet carrier interconnection agreements. We collected our data by voluntary survey, distributed globally through all of the regional Network Operators Groups between October 2010 and March 2011. The responses we received represented 4,331 different ISP networks, or approximately 86% of the world's Internet carriers, incorporated in 96 countries, including all 34 OECD member countries and seven of the 48 UN Least Developed Countries. For each agreement, in addition to the identities of the carriers party to the agreement, we asked the following three questions:

- Is the agreement formalized in a written document, or is it a "handshake" agreement?
- Does the agreement have symmetric terms, or do the parties exchange different things?
- What is the country of governing law of the agreement?

In addition, we made the following determination for each agreement:

- Is the agreement bilateral or multilateral?

In 1,032 cases, both parties to the same agreement responded to our survey, and in 99.52% of those cases, both parties' answers to each of the three questions were identical. We believe that, among other things, this indicates that respondents understood the questions clearly and were able to answer unambiguously and accurately.

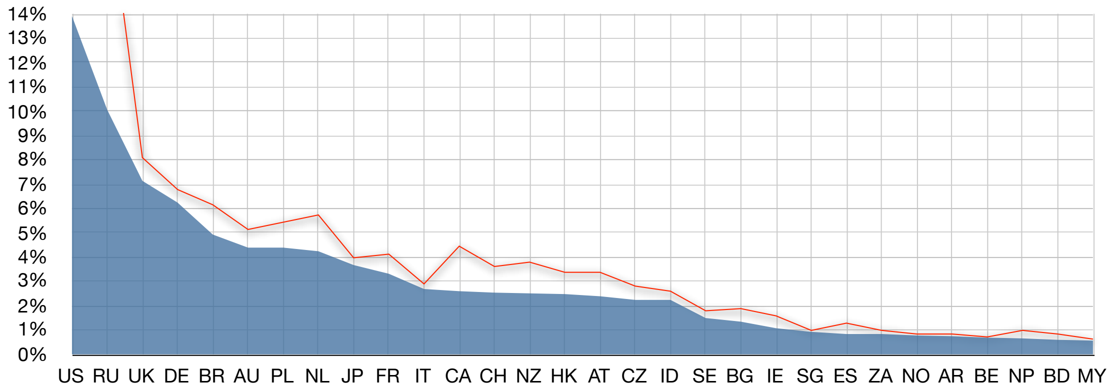

Figure 1: Top thirty countries of incorporation of the represented networks, as a percentage of those in the dataset.

The largest number of networks represented in the dataset were incorporated in the United States (466), followed by Russia (337), the United Kingdom (239), Germany (209), and Brazil (165). On the long tail of the curve, 45, or nearly half, of the countries were represented by three or fewer networks. The red line in Figure 1 indicates the total number of networks incorporated in each country; the blue area indicates those represented in the responses to our survey. In most countries, a significant and relatively uniform majority of the networks are represented in our data, but our coverage in the United States (30%) and Russia (52%) was disproportionately small relative to other countries, and this does slightly affect the results of some of our country-specific analyses of these two countries, as we discuss later.

Informal Agreements

Of the total analyzed agreements, 698 (0.49%) were formalized in written contracts. The remaining 141,512 (99.51%) were "handshake" agreements in which the parties agreed to informal or commonly understood terms without creating a written document. The common understanding is that only routes to customer networks are exchanged, that BGP version 4 is used to communicate those routes, and that each network will exercise a reasonable duty of care in cooperating to prevent abusive or criminal misuse of the network.[3] This huge number of informal agreements are arrived at by the "peering coordinators" or carrier-interconnection negotiation staff of the networks, often at self-organized regional or global "peering forums" that take place many times each year.[4]

Symmetric Terms

Of the agreements we analyzed, 141,836 (99.73%) had symmetric terms, in which each party gave and received the same conditions as the other; only 374 (0.27%) had asymmetric terms, in which the parties gave and received conditions with specifically defined differences. Typical examples of asymmetric agreements are ones in which one of the parties compensates the other for routes that it would not otherwise receive (known as "paid peering"), [5] or in which one party is required to meet terms or requirements imposed by the other ("minimum peering requirements"). [6] In the more common symmetric relationship, the parties to the agreement simply exchange customer routes with each other, without settlements or other requirements.[7]

Governing Law

No interconnection agreements were reported that utilized a country of governing law that was not also the country of incorporation as well as the location of primary operation of one of the two carriers party to the agreement. Stated another way, in no case did the parties choose a country of governing law that was not one of their own countries of incorporation and primary operation. This indicates that there is, as yet, no country that has such compelling rule of law in the field of carrier interconnection as to incentivize this behavior. Contrast this with other areas of commerce in which countries tailor regulatory or legislative environments to attract business as, for example, the registration of much maritime shipping in Panama or banks in Switzerland.

Nonetheless, clear preferences were expressed in the data, with the distribution of countries of governing law being sparser than the distribution of countries of incorporation and operation. In other words, some countries' governing law was preferred to a greater degree than their frequency as a country of incorporation would suggest, whereas others were preferred for governing law less frequently than they appeared as a country of incorporation.

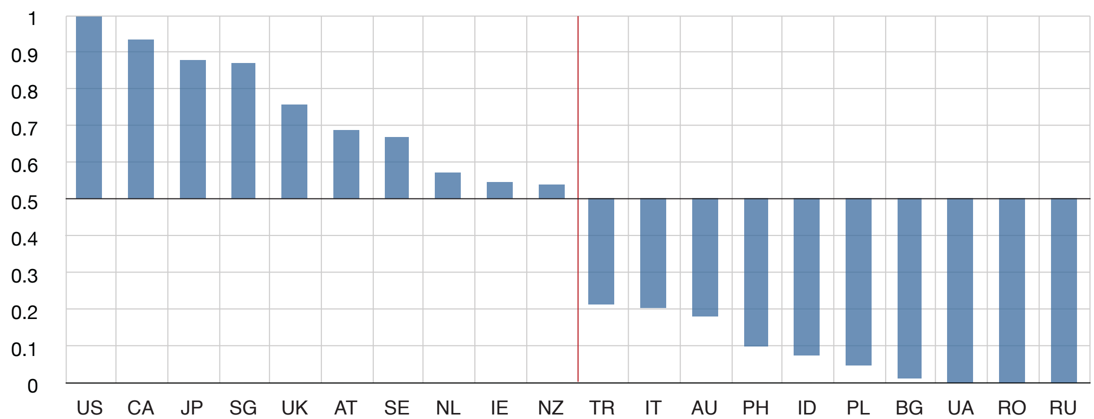

Figure 2: Probability of selection as a country of governing law, ten most-likely and ten least-likely countries

When we compare the frequency of appearance as a country of incorporation to the frequency of selection as a country of governing law (Figure 2), in nearly every interconnection agreement in which one of the two parties is incorporated in the United States or Canada that country is selected as the country of governing law in preference to the country of incorporation of the other party to the agreement. At the opposite end of the spectrum, there were no agreements in the dataset in which Russia, Romania, or the Ukraine was selected to supply governing law for an agreement with a country outside this group of three, even though 337 Russian, eighteen Ukrainian, and eight Romanian networks are represented in the dataset. Each time a Russian, Romanian, or Ukrainian network interconnected with a foreign network, the parties elected to use the other country's governing law.

National Interconnection Partners

Looking solely at the frequencies with which pairs of countries of incorporation appear within the dataset, it is possible to chart the relative number of connections between any country and all others. By way of example we chart the most frequent interconnection partners (those consisting of more than 1%) of each of the four countries that are most frequently represented in our dataset – the United States, Russia, the United Kingdom, and Germany (Figure 3).

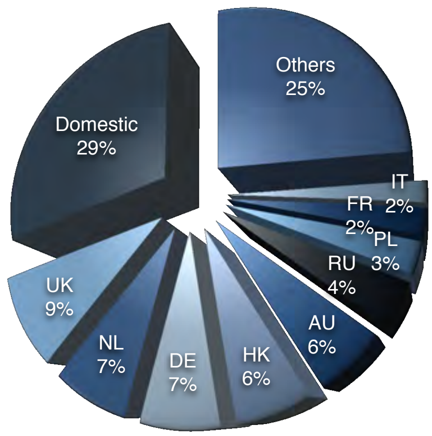

United States

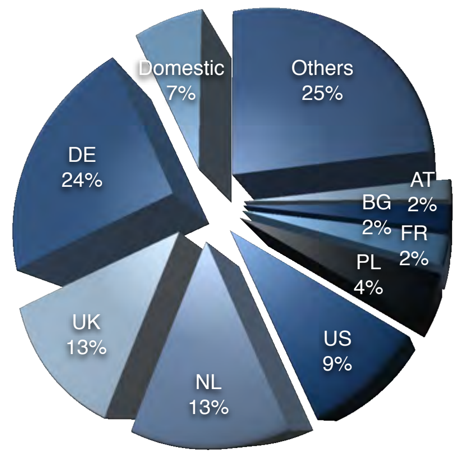

Russia

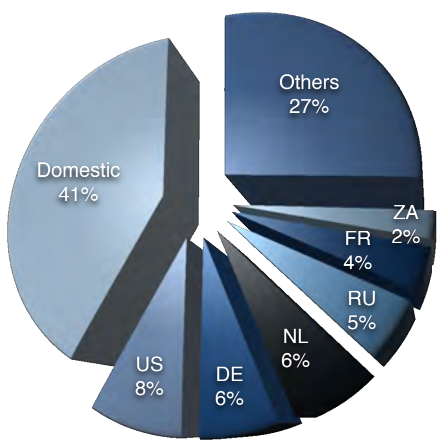

United Kingdom

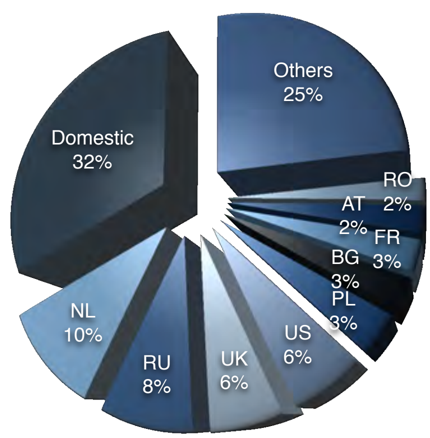

Germany

Figure 3: Trends in peering partners, selected countries

Among these partners, linguistic cohorts, geographically proximal neighbors, and frequent commercial trading partners are favored. The only real surprise is the relatively small share of domestic interconnection agreements observed within Russia, and we believe that this can be attributed to a selection bias in the dataset rather than to actual conditions on the ground; though we received many survey responses from networks that interconnect with U.S. and Russian networks, fewer were received from U.S. and Russian networks themselves, which would account for their relatively low shares of domestic interconnections.

Degree of Interconnection

Most of the networks represented have small numbers of interconnection partners. Of the 4,331 networks, 2,696 (62%) have ten or fewer interconnection agreements, and only twelve of the represented networks have more than 700 interconnection agreements.

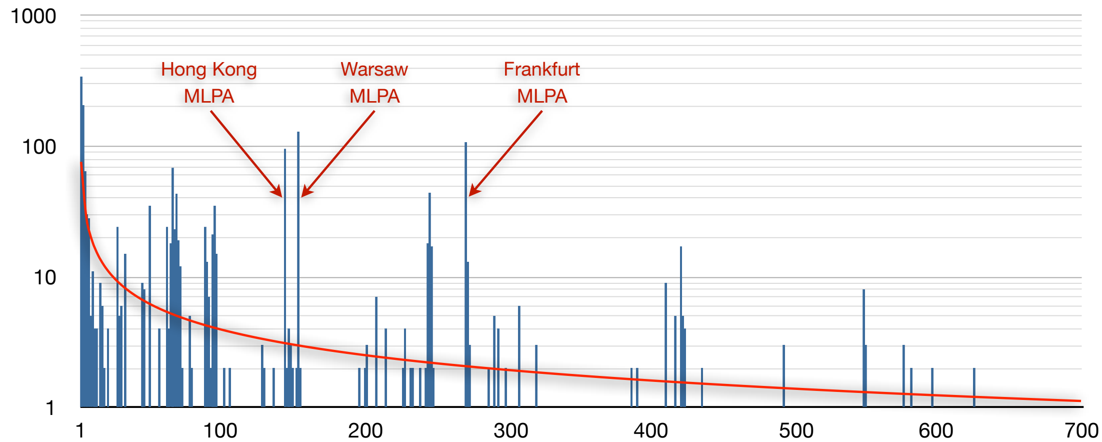

Figure 4: Distribution of number of networks (X axis) with each quantity of interconnection partners (Y axis)

A number of "spikes" are visible in the distribution graph (Figure 4), with major ones appearing clustered around the values 144, 154, and 271. These are the effect of large multilateral peering agreements (MLPAs), specifically the ones associated with the Hong Kong, Warsaw, and Frankfurt Internet exchange points.

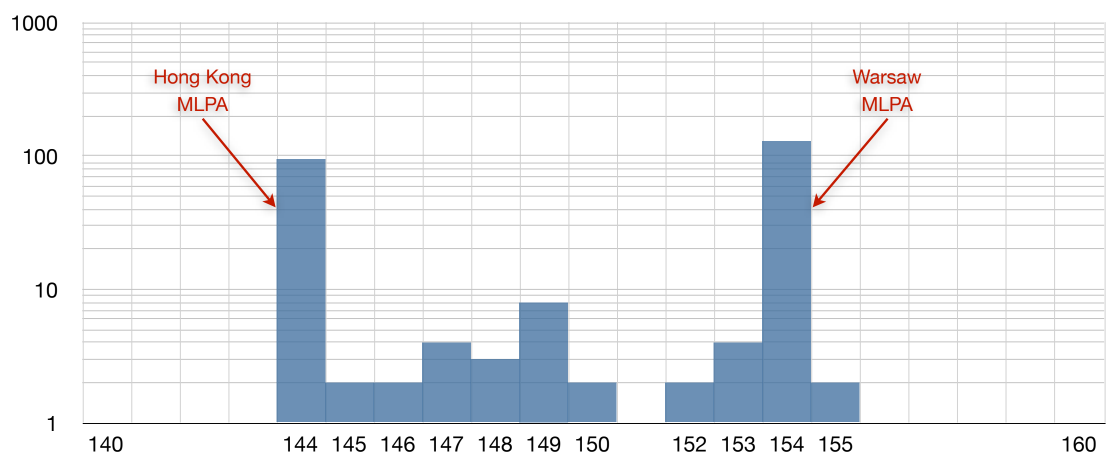

Figure 5: Expanded view of the Y axis range 140 - 160 from Figure 4, detailing the Hong Kong and Warsaw MLPAs

In each case, there exist a large number of networks that all peer with each other, creating a spike at that value, which trails off as a function of the portion of those networks that also have other interconnection agreements. To some degree, the volume of the tail to the right of the spike varies with the age of the MLPA, since MLPAs that have existed longer generally include members who have had more time to also form bilateral agreements outside the MLPA. Generally speaking, multilateral peering agreements are identifiable as spikes that have similar values in both X and Y axes in figures 4 and 5.

Unexpected Results

One unexpected result of this survey is a new understanding of the prevalence of multilateral peering. Multilateral peering, the exchange of customer routes within groups of more than two parties, has long been characterized as a practice principally engaged in by smaller networks. It has been commonly assumed that large networks decline to participate in multilateral peering agreements, and that multilateral agreements are therefore outside of the mainstream of peering practice. Although the method by which we collected our survey data does not allow us to compare absolute quantities of bilateral agreements to multilateral agreements, the majority of the Autonomous System pairs we observed were connected through multilateral agreements, and many of those agreements were very large, with dozens or hundreds of participants.[8] With the exception of the cluster circled in red (which consists of "tier-1" ISPs), each of the other vertical clusters in figures 6 and 7 represents a multilateral agreement, similar to the spikes in figures 4 and 5.

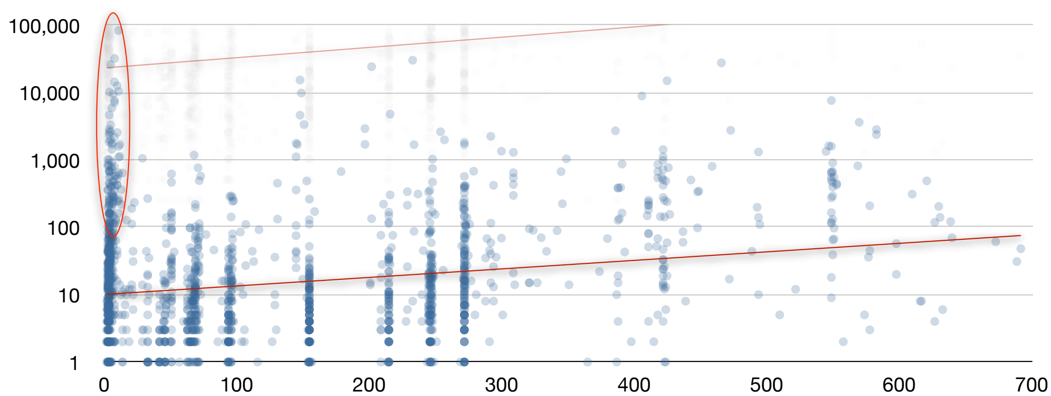

Figure 6: Number of advertised prefixes (Y axis) over number of interconnection partners (X axis) per carrier

It seems possible that, just as "donut peering" overtook "tier-1" peering in the late 1990s, multilateral peering may now be overtaking bilateral peering, at least in sheer numbers, if not necessarily in volume of traffic.[9] In both cases, market-dominant networks loudly derided as "peripheral" a practice that sought to render them irrelevant, but that practice slowly gained prevalence over time, becoming mainstream without ever receiving much notice. As an example, the 144 participants in the Hong Kong Internet Exchange multilateral peering agreement represent 10,296 AS-pair adjacencies, and *each one* of those participants individually exceeds the average "tier-1" carrier in degree of interconnection. When articulated in writing, multilateral peering agreements tend to follow the same general form and terms as other peering agreements, with the sole exception of having more than two parties.[10]

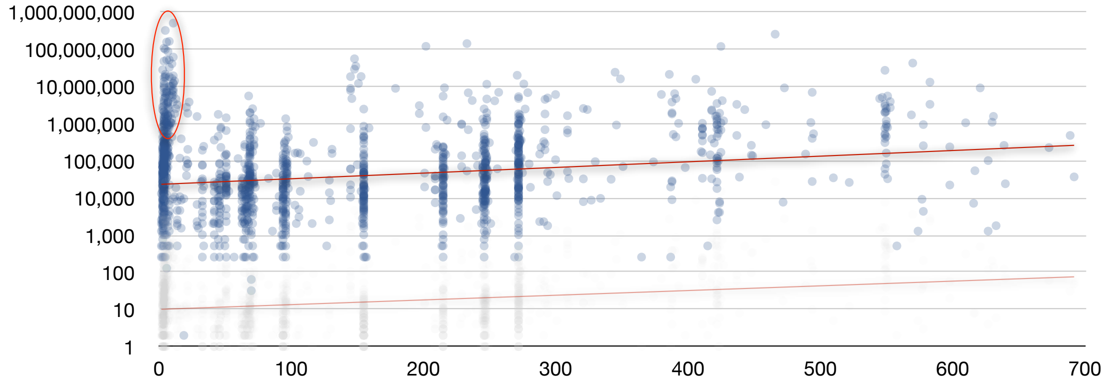

Figure 7: Number of advertised IPv4 addresses (Y axis) over number of interconnection partners (X axis) per carrier

Another finding of this survey, predictable on the face of it but to an unexpected degree, is how far the mainstream trend in number of interconnection agreements has left behind the legacy "tier-1" networks, which tend to rely upon very small numbers of interconnection agreements. The circled clusters in figures 6 and 7 represent the "tier-1" ISPs, each of which has a large number of advertised IPv4 prefixes, and consequently a larger number of actual IPv4 addresses, yet very few interconnection partners. Following the red line that indicates the average correspondence between size and number of interconnection partners to the right, most of the "tier-1" ISPs would have several thousand peers, if they were within mainstream ratios. By contrast, large content-distribution networks ("CDNs"), which have similar scale and degree of infrastructural investment tend to be exemplars of mainstream trends in our data, with very broad interconnection, both in absolute numbers and in geographic diversity. Although this may be self-evident, we expected to see a single order of magnitude difference between the number of agreements held by similarly sized networks in those two categories, whereas the actual difference was of two orders of magnitude.

Further Work Necessary

One weakness of this study, which provides reason for future work, is that we had relatively few mechanisms by which to compare the distribution of the responses we received to an objective "ground truth," or to preexisting datasets, in order to determine how statistically representative our survey respondents were to the Internet as a whole. Because previous studies of carrier interconnection agreements have been many orders of magnitude more narrowly focused than this one, they do not provide a statistically useful baseline against which we can characterize our dataset. A comparison against our own internal interconnection agreement data would have shed little light on how the survey dataset compares to the Internet as a whole and would have precluded including our own network's data in the survey. Furthermore, there is no mechanism for directly observing all of the peering agreements that exist in the Internet, and thus no ground truth to compare to. We hope that our foray into characterization of carrier interconnection agreements encourages researchers in the academic community to follow up with further work on the subject.

This paper and future versions may be found at http://pch.net/resources/papers/peering-survey

[1] Bill Woodcock, research director, woody@pch.net. Vijay Kumar Adhikari, on summer internship at PCH during his doctoral research at the University of Minnesota under professor Zhi-Li Zhang, vijay@pch.net.

[2] Smith, Philip, *Weekly Routing Table Report*, April 15, 2011, Transit ASes: http://thyme.apnic.net/rv3-data/2011/04/15/mail-global

[3] For a discussion of standard symmetric peering terms and conditions, read Chris Hall's http://www.highwayman.com/peering/peering_agreement.html. Although much more long-winded, the London Internet Exchange's model peering agreement also encapsulates the generally accepted terms of a symmetric peering agreement: https://www.linx.net/good/bcp/peeringagreement_draftv4.html.

[4] For a global schedule of Internet governance meetings, including many peering forums, see http://internetmeetings.org. For specific examples, see the Global Peering Forum website, http://peeringforum.net or the European Peering Forum website, http://www.peering-forum.edu.

[5] A discussion of MWEB, a South African ISP, transitioning from paid peering to normal peering can be read at http://mybroadband.co.za/news/broadband/16313-MWEB-peering-link-cuts-How-impacts-you.html. Specific solicitations of paid peering can be found on the websites of the AOL Transit Data Network, http://www.atdn.net/paid_peering.shtml; Cox Communications, http://www.cox.com/peering/paid-peering.asp; and Verizon Business, http://www22.verizon.com/wholesale/productguide/partnerportprogram.

[6] Bill Norton discusses the barriers to entry often contained in "minimum peering requirements" in his *Study of 28 Peering Policies*: http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html. Original documents can be found on the websites of Comcast, http://www.comcast.com/peering; Tiscali, http://www.as3257.net/peering-policy; AT&T http://www.corp.att.com/peering; and Internet Solutions, ftp://ftp.is.co.za/tech/peering.pdf.

[7] Definitions and discussions of peering and its general terms can be found on the Packet Clearing House website https://www.pch.net/wiki/pch:public:glossary#p; Wikipedia, http://en.wikipedia.org/wiki/Peering; and Bill Norton's website, http://drpeering.net/white-papers/Ecosystems/Internet-Peering.html.

[8] An "Autonomous System" (AS) is a uniquely-identified Internet network. Autonomous System Numbers (ASNs) are the numeric identifiers assigned by the Regional Internet Registries (RIRs) and used within the Internet routing system to define a specific bounded network that has its own uniquely defined routing policies. An AS-pair is a pair of networks that interconnect with each other.

[9] "Donut peering" is the practice of small and medium-size networks peering with each other aggressively in order to reduce the detrimental impact of a larger network refusing to peer with them. This results in a "donut" of densely interconnected networks surrounding a self-proclaimed "tier-1" network – the "donut hole" that is poorly interconnected with the networks around it. For a further discussion of donut peering, see the Cook Report's November 2002 *Economics of IP Network Interconnection*, http://www.cookreport.com/backissues/nov-dec2002cookrep.pdf; or my own January 2003 lecture to the University of Minnesota Digital Technology Center, *Internet Topology and Economics: How Supply and Demand Influence the Changing Shape of the Global Network* http://www.pch.net/resources/papers/topology-and-economics. "Tier-1" is the moniker some carriers in the mid-1990s gave themselves as they attempted to form a cartel, peering with each other but nominally refusing to peer with any networks outside the cartel. Their misunderstanding of Internet growth rates led them to become irrelevant, as the portion of the market held outside the cartel grew exponentially while that inside the cartel grew in linear fashion.

[10] A range of typical multilateral peering agreements can be found on the websites of the Open Peering Initiative http://www.openpeering.nl/mlparegistry.shtml; the Kansas City Network Access Point http://www.kcnap.net/peering-policy.html; Red Bus Internet Exchange http://www.rbiex.net/assets/joining/mlpa.pdf; and the Indonesia Internet Exchange http://www.iix.net.id/library/iix-peering-agreement_ind.pdf. Note that their specific terms differ little if at all from those of the bilateral agreements discussed in note 2.

OECDpublishing

Please cite this paper as:

Weller, D. and B. Woodcock (2012), "Internet Traffic Exchange: Market Developments and Policy Challenges", *OECD Digital Economy Papers*, No. 207, OECD Publishing.
http://dx.doi.org/10.1787/5k918gpt130q-en

OECD Digital Economy Papers
No. 207

Internet Traffic Exchange

MARKET DEVELOPMENTS AND POLICY CHALLENGES

Dennis Weller, Bill Woodcock

OECD

Unclassified **DSTI/ICCP/CISP(2011)2/FINAL**

Organisation de Coopération et de Développement Économiques
Organisation for Economic Co-operation and Development **17-Oct-2012**

English - Or. English

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INDUSTRY
COMMITTEE FOR INFORMATION, COMPUTER AND COMMUNICATIONS POLICY

DSTI/ICCP/CISP(2011)2/FINAL
Unclassified

Working Party on Communication Infrastructures and Services Policy

INTERNET TRAFFIC EXCHANGE

MARKET DEVELOPMENTS AND POLICY CHALLENGES

R. van der Berg

JT03329031

Complete document available on OLIS in its original format

English - Or. English

FOREWORD

The report was prepared by Dennis Weller of Navigant Economics and Bill Woodcock of Packet Clearing House. It is published on the responsibility of the Secretary General of the OECD.

The statistical data for Israel are supplied by and under the responsibility of the relevant Israeli authorities. The use of such data by the OECD is without prejudice to the status of the Golan Heights, East Jerusalem and Israeli settlements in the West Bank under the terms of international law.

1. Note by Turkey

The information in this document with reference to « Cyprus » relates to the southern part of the Island. There is no single authority representing both Turkish and Greek Cypriot people on the Island. Turkey recognizes the Turkish Republic of Northern Cyprus (TRNC). Until a lasting and equitable solution is found within the context of the United Nations, Turkey shall preserve its position concerning the ―Cyprus issue".

2. Note by all the European Union Member States of the OECD and the European Commission

The Republic of Cyprus is recognised by all members of the United Nations with the exception of Turkey. The information in this document relates to the area under the effective control of the Government of the Republic of Cyprus.

MAIN POINTS

The Internet has continued to develop at a spectacular rate over the past five years. It now includes two billion users, and traffic has grown eightfold in that time. Yet debate continues about the efficacy of the Internet‘s economic model of traffic exchange. This report seeks to make information available for policy makers and other stakeholders.

Since the Internet was commercialised in the early 1990s, it has developed an efficient market for connectivity based on voluntary contractual agreements. Operating in a highly competitive environment, largely without regulation or central organisation, the Internet model of traffic exchange has produced low prices, promoted efficiency and innovation, and attracted the investment necessary to keep pace with demand.

The performance of the Internet market model contrasts sharply with that of traditional regulated forms of voice traffic exchange. If the price of Internet transit were stated in the form of an equivalent voice minute rate, it would be about USD 0.0000008 per minute—five orders of magnitude lower than typical voice rates. This is a remarkable and under-recognised endorsement of the multi-stakeholder, market driven nature of the Internet.

A survey of 142 000 peering agreements conducted for this report shows that the terms and conditions of the Internet interconnection model are so generally agreed upon that 99.5% of interconnection agreements are concluded without a written contract. That these ―rules of the game" are so ubiquitous and serviceable indicates a degree of public unanimity that an external regulator would be hard-pressed to create. The parties to these agreements include not only Internet backbone, access, and content distribution networks, but also universities, NGOs, branches of government, individuals, businesses and enterprises of all sorts—a universality of the constituents of the Internet that extends far beyond the reach of any regulatory body‘s influence.

As incumbent networks adopt IP technology, there is a risk of conflict between legacy pricing and regulatory models and the more efficient Internet model of traffic exchange. By drawing a ―bright line" between the two models, regulatory authorities can ensure that the inefficiencies of traditional voice markets will not take hold on the Internet.

The Internet has expanded to cover the globe, with many emerging economies growing at a faster pace and closing the ―digital divide" gap with OECD countries; yet some emerging economies still suffer from the effects of lack of competition or regulatory liberalisation. Evidence shows that, when allowed to do so, market participants will self-organise efficient Internet exchange points, producing Internet bandwidth to the benefit of the local economy and significantly reducing their costs, including in foreign currency. This course of action is strongly recommended in economies that do not yet have abundant domestic means of Internet bandwidth production.

An unbroken chain of basic physics research, development, and production of new technologies allowed the Internet‘s growth to keep pace with demand for the first thirty years of its existence, but investment in basic optoelectronic physics fell during the economic downturn in 2001. Consequently, the speed of network interfaces has stalled, and this has led to a transition from exponential growth to linear growth. Investment in basic research needs to be reinstated to return to a level of growth that will meet the economic and social development goals OECD countries expect of the Internet economy.

TABLE OF CONTENTS

MAIN POINTS

The Internet has continued to develop at a spectacular rate over the past five years. It now includes two billion users, and traffic has grown eightfold in that time. Yet debate continues about the efficacy of the Internet's economic model of traffic exchange. This report seeks to make information available for policy makers and other stakeholders.

Since the Internet was commercialised in the early 1990s, it has developed an efficient market for connectivity based on voluntary contractual agreements. Operating in a highly competitive environment, largely without regulation or central organisation, the Internet model of traffic exchange has produced low prices, promoted efficiency and innovation, and attracted the investment necessary to keep pace with demand.

The performance of the Internet market model contrasts sharply with that of traditional regulated forms of voice traffic exchange. If the price of Internet transit were stated in the form of an equivalent voice minute rate, it would be about USD 0.0000008 per minute—five orders of magnitude lower than typical voice rates. This is a remarkable and under-recognized endorsement of the multi-stakeholder, market driven nature of the Internet.

A survey of 142,000 peering agreements conducted for this report shows that the terms and conditions of the Internet interconnection model are so generally agreed upon that 99.5% of interconnection agreements are concluded without a written contract. That these "rules of the game" are so ubiquitous and serviceable indicates a degree of public unanimity that an external regulator would be hard-pressed to create. The parties to these agreements include not only Internet backbone, access, and content distribution networks, but also universities, NGOs, branches of government, individuals, businesses and enterprises of all sorts—a universality of the constituents of the Internet that extends far beyond the reach of any regulatory body's influence.

As incumbent networks adopt IP technology, there is a risk of conflict between legacy pricing and regulatory models and the more efficient Internet model of traffic exchange. By drawing a "bright line" between the two models, regulatory authorities can ensure that the inefficiencies of traditional voice markets will not take hold on the Internet.

The Internet has expanded to cover the globe, with many emerging economies growing at a faster pace and closing the "digital divide" gap with OECD countries; yet some emerging economies still suffer from the effects of lack of competition or regulatory liberalisation. Evidence shows that, when allowed to do so, market participants will self-organize efficient Internet exchange points, producing Internet bandwidth to the benefit of the local economy and significantly reducing their costs, including in foreign currency. This course of action is strongly recommended in economies that do not yet have abundant domestic means of Internet bandwidth production.

An unbroken chain of basic physics research, development, and production of new technologies allowed the Internet's growth to keep pace with demand for the first thirty years of its existence, but investment in basic optoelectronic physics fell during the economic downturn in 2001. Consequently, the speed of network interfaces has stalled, and this has led to a transition from exponential growth to linear growth. Investment in basic research needs to be reinstated to return to a level of growth that will meet the economic and social development goals OECD countries expect of the Internet economy.

INTRODUCTION AND EXECUTIVE SUMMARY

Since the Internet was privatised in the early 1990s, the growth of Internet traffic has been explosive. What was once a research network serving a limited community is now a global phenomenon of immense economic and social importance. Three factors have made this dramatic growth in traffic possible:

- The basic architecture of the Internet has proved to be remarkably adaptable. The independence of the seven layers of the architecture has allowed innovation to proceed separately in each. There has not, however, been a dramatic change in the architecture of the Internet over the period of rapid growth since privatisation. It has been the adaptability of the structure, rather than fundamental change in it, that has accommodated growth so far.

- Innovation by equipment manufacturers has made each new generation of transport facilities, routers, and storage less expensive and more efficient. Without these advances, no amount of investment would have been sufficient to accommodate the growth in traffic. Even with those benefits, technical improvements have not always been able to keep pace with demand, and, as explained further below, they have sometimes been a limiting factor.

- The commercial agreements that have evolved over the past twenty years have created an efficient global market for connectivity. The progress of this market has been described in previous OECD reports in 2002[1] and again in 2006[2]. The Internet traffic exchange market has ensured universal connectivity worldwide, supplied the investment to meet growing needs, and provided an environment that encourages diversity and innovation among all the participants in the Internet.

The global success of the Internet is thus a product of commercial as well as technical innovation. Business models have evolved through a process of experimentation, in the voluntary agreements between parties. In terms of the volume of information transmitted, the Internet market now vastly exceeds the pre-existing arrangements through which traditional networks have exchanged circuit-switched (TDM) voice traffic. While national regulatory authorities have closely regulated TDM traffic exchange to achieve such policy goals as universal connectivity and competition, the Internet market has attained those same goals with very little regulatory intervention, while performing much better than the older markets in terms of prices, efficiency, and innovation. As Internet traffic continues to grow, and TDM traffic shrinks, the Internet model for traffic exchange has become the global norm.

This report reviews developments in the market for Internet traffic exchange, and the performance of that market, since the 2006OECD report on this subject. It also considers the challenges the Internet market is likely to face and policy issues that governments may be called upon to consider in the next few years.

Performance of the Market

Since the last OECD report on this subject in 2006, the volume of traffic on the Internet has continued to grow at an average rate of about 50% per year. As a result, the total volume of traffic is now about eight times greater than it was in 2006. Growth has been intensive, as the usage of each subscriber has increased. Today, twenty households with average broadband usage generate as much traffic as the entire Internet carried in 1995. Growth has also been extensive, as the Internet has expanded around the world to reach two billion users. The rate of growth has varied across countries; details are provided below under ―Regional and National Trends."

The market model of Internet traffic exchange, based on voluntary commercial agreements for peering and transit, has performed very well. The market has generated the large investments needed to sustain growth, and has guided investment to the most productive uses. It has also sustained very low prices. Rates for Internet transit have declined dramatically over the last fifteen years, although the rate of change has leveled off in the most recent period. For large volume wholesale agreements for traffic exchanged at major Internet Exchange Points (IXPs), transit is available for USD 2 to USD 3 per megabit per month; some networks making vary large commitments have been reported to pay less than 1 USD. Retail prices charged by ISPs to corporate clients at smaller volumes are somewhat higher, and vary depending on location and market conditions. The section of this report entitled ―Market Performance Measures" discusses how Internet bandwidth is produced and how different measures of performance should be interpreted.

The basic market models of Internet transit and peering agreements have been developed through experience over the last fifteen years, to a point where the norms of these arrangements are now widely understood and accepted within the Internet community. A survey of peering agreements conducted for this report gathered responses from 4 331 ISP networks, representing 86% of the world‗s Internet carriers, incorporated in 96 countries. They reported 142 210 peering agreements. Of these, 141 512, or 99.51%, were ―handshake" agreements in which the parties agreed to commonly understood terms without creating a written document. Because this model is so well-developed, peering has been able to proliferate among Internet entities with very low transaction costs. Another mechanism for reducing transaction costs is the widespread use of multilateral peering agreements, in which many networks meeting at an IXP join a single agreement rather than conclude separate bilateral arrangements. The average number of agreements reported was 32.8 for each of the entities in the survey. The distribution is long-tailed, with 62% of the respondents having ten or fewer agreements. It has therefore been possible to ensure global connectivity among two billion users by means of a relatively small number of agreements, less than 1% of a full mesh. An important function of the Internet market for traffic exchange is to determine which of the many possible exchange arrangements should actually be implemented. A full description of the survey and its results is presented in Annex 1 to this report.

Structural Evolution of the Market

As the Internet has expanded, it has also followed a continual process of reinvention. Roles, relationships, and business models of Internet participants have shifted over time as the Internet has expanded:

- The Internet has expanded geographically across different regions of the world, with faster growth in emerging economies than in more established ones. The ―Regional Survey" section of this report provides a county-by-country review of Internet development since 2006.

- Smaller networks have developed more peering relationships with one another and are less dependent on transit services provided by larger backbone networks.

- IXPs have been established in many more areas and have in turn attracted greater Internet resources to those areas, reducing the need for ―tromboning" of traffic out of the country or region. The availability of in-region points of exchange allows for more direct routing of traffic, increasing service quality. It creates better conditions for investment in Internet assets within the region. By reducing the need to use long-distance transmission capacity for in-region traffic, it frees up existing long-haul capacity to handle out-of-region traffic.

- New investment has reduced concentration and relieved bottlenecks, such as for undersea cable capacity. Much of this investment has come from new entrants and from firms in emerging economies, spreading ownership geographically and among more participants.

- New patterns of usage have developed among Internet users. Voice services have declined in relative importance. Peer-to-peer delivery, which had previously led to concerns that it might overwhelm the Internet with volume, is still significant but it is becoming a smaller part over overall usage.[3] Streaming and direct downloads, especially of video content, have become the most rapidly growing category among applications on the Internet. The transition of applications from desktop to cloud has also contributed to the growth in traffic.

- New categories of participants have invested to improve quality and create new alternatives to transit. These include self-supply by online service and content providers such as Google, as well as intermediary content delivery network (CDN) service providers such as Akamai and Limelight. CDN services have supported, and have grown in parallel with, the growing demand for applications such as video streaming and download. Taken together, many of the structural changes summarised here - reduced reliance on transit, local availability of IXPs, direct delivery of traffic by CDNs, and caching of content closer to the user - have all contributed to make routing more direct, reduce latency, and improve quality.

These developments have made the structure of the Internet flatter and broader, and reduced its dependence on any one player or group. Today, only a small percentage of the traffic on the Internet ever touches any of the old backbone networks. Google is now ranked third among networks in global traffic carried, behind only Level 3 and Global Crossing, and it‗s notable that each of these largest networks are born of the Internet era, rather than evolved from incumbent TDM predecessors. In general, the growth of the Internet over the past five years has increased the effectiveness of competition in the market for Internet traffic exchange.

The reasons behind the very good performance of the Internet market model of traffic exchange are explored in the section entitled ―Why has the Internet market performed so well?" in the Annex to this report. The contrast between the results observed in the Internet market and comparable markets for exchange of traditional circuit-switched voice (time-domain multiplexed, or TDM) traffic is striking. For example, Internet transit service provides what is effectively, in TDM terms, global transport and termination. The price of USD 2 to USD 3 per megabit per month therefore includes a traffic-weighted average of transport costs to all the possible destinations in the world, as well as the costs of terminating on local access networks in each country. Stated in terms of an equivalent per-minute price for delivery of voice traffic, this is less than USD 0.0000008, five orders of magnitude less than wholesale rates for services providing comparable functions in TDM markets. The reasons for this performance include the efficiency of packet-switched technology, competition in Internet markets, and the flexibility of routing arrangements among Internet networks. The market has also benefitted from the policy environment, in which governments have refrained, in most OECD countries, from regulation of the market for Internet traffic exchange.

Legal and Policy Frameworks

Every market exists within a legal and policy framework established by government. Markets for network communications in OECD countries have benefitted from liberalisation policies which have sought to open markets and promote competition. These have included specific measures to ensure availability of leased lines and access to rights of way, for example. The success of the Internet IP market for traffic exchange would not be possible without this broad policy framework of liberalisation. Indeed, in regions where the development of Internet IP traffic exchange has been less satisfactory, the cause has generally been a lack of sufficient liberalisation within the country or region, rather than a lack of performance by the global Internet market as a whole. The close relationship between market-opening policies and Internet development was discussed in the most recent OECD report in 2006, and is also evident in the regional survey provided in this report. Some recommendations for features of a broad policy framework of liberalisation which would provide a sound basis for Internet development are provided in Annex 6 to this report.

Governments have generally found it necessary to regulate traditional markets for TDM traffic exchange, given the characteristics of TDM networks and the competitive conditions in those markets. Regulation has in many cases been adopted to address observed market failures. In contrast, it has been possible for the Internet market for traffic exchange to achieve more favourable results, in most cases without regulatory intervention. This has been achieved for the reasons noted above, despite the fact that many of the same firms participate in both the TDM market and the Internet market. The paradox presented by regulation of interconnection and traffic exchange is that while it can provide a remedy in some circumstances where market failures have occurred, in other cases it can interfere with the successful operation of the market and confer market power on entities which would not otherwise have it. For reasons discussed in detail in the section ―Why has the Internet market performed so well” in the Annex to this report, this risk is particularly powerful in the market for Internet traffic exchange, where the market has produced very good results with little regulation, and where the potential for regulatory intervention to damage the market is particularly strong. Some discussion of how the Internet traffic exchange market fits within existing regulatory frameworks, and a review of limited cases where OECD governments have intervened in this market, is provided in Annex 7 to this report.

Challenges for the Future

The rapid growth and continuing evolution of the Internet market will present new challenges for all of the three enabling factors listed in the Introduction: the architecture of the Internet, technological innovation, and the ongoing development of the Internet market model for traffic exchange. These in turn are likely to give rise to new policy challenges for governments.

Challenges for Internet Architecture

The most immediate challenge for the architecture of the Internet is the transition of addressing from IPv4 to IPv6. In February 20111 the Internet Assigned Numbers Authority issued the last IPv4 address blocks to the Internet Registries. The Asia-Pacific Regional Internet Registry, (APNIC) ran out of unreserved IPv4 addresses in April 2011, and each of the other four regional Registries will run out over the coming months, or perhaps a year or two in the case of AfriNIC. While the number of networks employing IPv6 addresses has grown more than sixfold in the last five years, thousands more have yet to begin their transition. This issue is discussed under ―Challenges for the future” in this report, and has been addressed in greater detail in a 2010 OECD report.[4]

Challenges for Technical Innovation

Rapid innovation in technology, leading to new generations of equipment providing greater capacity at lower unit cost, is one of the most important factors that has allowed the growth of the Internet infrastructure to keep pace with growth in demand. In particular, the interfaces that allow telecommunications signals to be driven through fibre optic cable have stepped from 10 megabits per second to 100 megabits, 1 gigabit, and 10 gigabits in successive generations of equipment. This has been made possible by basic research in optoelectronic physics. However, when the dot-com bubble of the 1990s burst, investment in such research was significantly reduced. Given a lag of about three years between research developments in the laboratory and the availability of new equipment in the market, by 2005 the rate of increase in fibre optic interface speeds began to deviate from trend. Additions to capacity have shifted from an exponential pattern of growth through increases in interface speeds, to a linear one of bundling together additional 10-gigabit interfaces.

The exponential growth that allowed Internet bandwidth production to keep pace with increasing demand, with lower costs over time allowing corresponding reductions in consumer prices, has unfortunately transitioned into linear growth that is not keeping pace with increased demand and is not producing greater economies at larger scales; thus, wholesale and retail prices have not, by and large, fallen significantly beyond the introduction of those final Nx10 gbps technologies in any market.

This issue is described in detail in the section ―Why supply of Internet bandwidth is no longer meeting demand" in this report, which also explains why such linear additions to capacity lead to diminishing, and ultimately negative returns. Unless research in basic optoelectronic physics can be renewed, limitations in interface speed will present a significant challenge to the ability of the Internet to handle increasing demand.

Challenges for the Internet model of traffic exchange

The efficiency and adaptability of the commercial arrangements for Internet traffic exchange has played a key role in the expansion of the Internet over the last fifteen years. This market has evolved rapidly over time. The ongoing process of structural change in the Internet market has helped to extend the Internet around the world, make it more competitive, and supported the widening range of services available to its users.

As this evolution continues over the next few years, it will present challenges to networks on the Internet to renew their business models, and in the process to adapt the Internet market model for traffic exchange to support these new approaches. For example, the growth of online delivery of video through streaming and direct download has led content companies, such as movie studios and broadcasters, to renegotiate their relationships with online content distributors like Google, Apple, and Netflix. The Internet arrangements for delivering this content have driven evolution of relationships among these content distributors, CDN providers, backbone networks, and local access networks. The distinctions among these categories have blurred as networks have provided different combinations of these functions. This process of development is described in ―Structural change in the Internet" in this report.

This evolution has led to changes in the terms of trade among the different parties, and to experimentation with different variants on the existing agreements for peering and transit. These variants have included regional peering, multilateral peering, single-hop access, and paid peering, and are reviewed in more detail in ―Adaptations of the Internet model" in this report.

Challenges for Policy

The evolution of market structure, patterns of use, and business arrangements on the Internet will present corresponding challenges for policy makers. They must maintain an underlying legal and regulatory framework that provides a liberalized market setting, within which Internet participants have the freedom to continue to innovate and grow. Some basic elements of such a framework are discussed in the ―Prescriptive Annex," in the annex to this report. Many aspects of policy, such as Internet governance, security and privacy, are outside the scope of this report. Since the focus of the report is on the market for Internet traffic exchange, the discussion in the section of the report entitled ―New models for policy" is also focused on regulation of that market. The issues that are likely to confront policy makers in the next few years include the following:

- *Establishing a bright line*. Convergence between traditional voice networks and the Internet has the potential to produce conflicts between the existing framework for the exchange of voice traffic, which is largely regulated, and the Internet model of traffic exchange based on voluntary agreements. As Internet traffic has grown, and TDM volumes have diminished, the Internet market model has become the global norm for traffic exchange. There is a risk that parties who have seen the TDM regulatory regime as working to their advantage might lobby governments to extend elements of traditional TDM regulation into the Internet space. National regulatory authorities (NRAs) might be tempted to do so in order to preserve norms and procedures to which they are accustomed, to maintain policy goals that are embedded in the older regulation, or simply in the name of technological neutrality. Similarly, in the context of the larger discussion over global Internet governance, some countries or parties might see advantage in re-establishing international regulation of Internet traffic exchange under treaty.

  While governments will always have the option of establishing regulation if it should become necessary, in the case of the market for Internet traffic exchange the results in the absence of regulation have been very favourable, and the risk that regulation would undermine the efficient working of the market is great, for reasons discussed under ―Why the Internet market has performed so well," in the Annex to this report. This report therefore recommends that NRAs should establish a ―bright line" between the TDM market, which might remain subject to regulatory status-quo, and the Internet market for traffic exchange, which would continue to operate under commercial agreements. A very high threshold of market failure should be established to justify intervention in the Internet market. Similarly, the growth of the Internet, together with efforts to liberalise international markets for telecommunications and promote competition, has produced enormous benefits for users, as well as for economic and social development, around the world. A treaty-based return to a regulated framework - in effect, a new settlements regime --for Internet traffic exchange would risk undoing some of those gains, while offering little potential benefit.[5]

  Specific policy concerns may arise in the case of voice traffic where established numbering systems tie the ability to route terminating traffic to the carrier to whom the number is assigned. The report offers suggestions for policies to address this concern while minimising intervention in the market.

- *Reform of TDM interconnection.* Several OECD member countries have undertaken efforts to reform their existing regulation of TDM interconnection. These efforts offer the potential to increase the efficiency of the TDM market, while at the same time smoothing the transition to the use of IP interfaces for the exchange of voice traffic.

- *Allowing for disagreement.* In a market based on voluntary agreements, it is inevitable that parties will not always agree. In the Internet, where less than 1% of the possible bilateral arrangements are actually in effect, this is not a cause for concern. Occasionally disputes rise to the attention of policy makers, and more rarely may actually lead to some disruption. While there must be some limit to the amount of disruption that a private dispute can be allowed to cause, governments should generally resist the temptation to intervene.

- *Pricing models to support future investment.* The rapid growth of Internet traffic creates a challenge for local access networks to provision increased capacity in middle mile facilities. In particular, online delivery of video content is a challenge for access networks not designed with that in mind. Some parties have recently suggested a series of pricing options for access networks seeking to fund investment to meet increased demand. These include the imposition of a broad-based, mandatory termination charge. These proposals are reviewed in detail in the report. While investment by all stakeholders across networks is welcome, as a general matter mandatory charges which are not accepted voluntarily by the other party to the transaction, or which would require either government intervention or collusion to enforce, are not consistent with the Internet model of traffic exchange, and should not be permitted.

- Traffic exchange and network neutrality. As noted above, the structural evolution of the Internet has led to realignment of the roles of Internet participants, including creators of content, online distributors of content, CDNs, backbone networks, and access networks. This has led to negotiation of new agreements for Internet traffic exchange. As agreements among online service and content providers, CDNs, and access networks are negotiated, a balance is being struck on the extent to which quality-enhancing resources will be brought to bear, who will provide those resources, and on what terms. In the process, the market for Internet traffic exchange is generating answers to many of the questions raised in recent debates over network neutrality.

The report reviews these market developments, as well as some disputes which have arisen within the last year, including those between Cogent and Orange in France, and between Level 3 and Comcast in the United States. In general, the market appears to be developing in an orderly way, with the outcomes falling within a relatively limited range that appears to be reasonable. The best course for regulators at this point may be to monitor and observe this process, as the NRAs in these two cases have chosen to do. It is not clear how intervention in favour of any one party or group (content providers, CDNs, access networks) would improve the outcome.

Market Performance Measures

In 2011 the Internet, or network of networks, consisted of 5 039 independent Internet service provider (ISP) networks, which are interconnected in a sparse mesh. Each of the interconnecting links takes one of two forms: transit or peering. Transit agreements are commercial contracts in which a customer (which may itself be an ISP) pays a service provider for access to the Internet; these agreements are most common at the edges of the Internet. Peering agreements are the carrier interconnection agreements that are within Internet exchange points (IXPs) and allow carriers to exchange traffic bound for one another‘s customers; they are most common in the core of the Internet and are the source of the Internet bandwidth commodity.

There are several approaches to understanding the performance of the Internet market. There are measures of supply, measures of demand, and measures of price. Because currency is convertible between markets, price is often used as a way of comparing Internet services on a global basis. This is not as simple or reliable a mechanism is it might seem; it is easy to wind up with an apples-to-oranges comparison, even if the prices of two markets seem directly comparable. Internet bandwidth is not a simple commodity like kilowatt-hours of electricity. Comparisons are complicated by its aggregation properties and performance characteristics, so a basic understanding of these issues is essential to understanding price comparisons.

Internet bandwidth is produced at the sites of interconnection between networks, the IXPs. There are presently approximately 350 IXPs, more than half of which are in Europe and the United States.[6] ISPs transport the bandwidth from the IXP to the consumer, often reselling it through a chain of intermediaries before it arrives at its ultimate user. The ISPs directly adjacent to the IXP and participating in the bandwidth-production process are most often referred to as ―backbone” providers; those adjacent to the customer are most often referred to as ―access” providers. This chain of providers between the IXP and the consumer may be as short as one, or occasionally as long as three or four or more, particularly in small or rural markets or in markets where regulatory policy has allowed a ―national gateway” provider a monopoly. Each of these intermediary providers has operational costs, which are passed along in prices to the consumer, and each intervening router, switch, or kilometre of network takes its toll on the performance and reliability of the connection. Thus, shorter paths between IXP and consumer yield higher performance at a lower price. The concept is familiar in other contexts: a consumer living adjacent to a farm can buy fresher agricultural produce at a lower price, whereas one who lives far away pays a higher price for produce that is more bruised and has a shorter remaining shelf-life. Internet bandwidth is very much a perishable commodity, like agricultural produce.

Thus, the first thing to understand when making price comparisons between Internet service offerings is how far the bandwidth has been transported from an IXP. This is itself a complicated question, because it relates not to a single specific IXP but to the mix of IXPs used to deliver traffic, which varies with user preferences and habits. Generally, an ISP that connects to many IXPs, and specifically to many IXPs near the site of both a customer and the parties the customer wishes to communicate with, is able to provide a high-quality service at a low price. An ISP that does not connect to any IXPs, or connects only to IXPs that are distant from the customer and their correspondence partners, cannot provide high-quality service or offer a low price.

Quality of service, in the Internet as in other networks, is measured in four characteristics: loss, latency, out-of-order delivery, and jitter.[7] A discussion of the specifics of these characteristics is outside the scope of this report, but it is important to understand that each is specific, has a clear and universally accepted definition, and is subject to unambiguous measurement and quantification. Internet service quality is objective and does not rely upon panels of human judges coming to subjective ―mean opinion scores.” Two different ISPs may quite legitimately engineer their networks in very different ways, in order to optimise the quality of one or another of those four parametres, each of which has different value to different consumers. Someone using their connection primarily for gaming values low latency dearly, whereas someone transmitting video values low out-of-order delivery and jitter to a far greater degree. Competitors in a mature market offer differentiation in the optimization of these different and unrelated axes of quality, and this frustrates simple comparisons based on bits-per-second versus cost.

To understand price comparisons, it is necessary to understand the ISP‘s function as an aggregator. The commodity of Internet transit is, essentially, the right to modulate signal onto passing unmodulated bits, paired with the right to receive inbound signals. This is equivalent to a service in which railway boxcars pass the customer at a predetermined rate that is proportional to the amount the customer pays. Some of these boxcars arrive containing goods bound for the customer, who may also put their own outbound goods on any of the boxcars as they go by. The total amount of goods a customer can send or receive is governed by the number of boxcars that pass in any period of time. This is the ―bit per second,” or ―bps,” rate. A thousand bits per second is a kilobit per second (kbps), a million bits per second is a megabit per second (mbps), a billion bits per second is a gigabit (gbps), and a trillion bits per second is a terabit (tbps). Although they have the right to do so, consumers purchasing gigabit broadband service will not in all likelihood actually fill all of those ―boxcars,” or unmodulated bits. A gigabit is the equivalent of seventy thousand pages of text, sixty-five high-resolution photographs of the sort produced by modern digital cameras, or two full-length record albums. Although it is quite possible that someone may wish to transmit or receive at least that much data, and it is of great benefit to be able to do so in the span of a second, it is unlikely that anyone would wish to *both send and receive* that much *every second* of every day. It is, in fact, unlikely that anyone would wish to do so more than a tiny fraction of a percent of the time.

Specifically, an average broadband user consumes about 15 gigabytes per month.[8] This constitutes an average usage of 45 kilobits per second, or 0.005% utilszation of that broadband service. (It would, however, be 78% of an analog modem dial-up service, such as was common fifteen years ago.) In a perfect world, an ISP could aggregate more than 22 million such customers into a single gigabit of fully utilized bandwidth, at an IXP or purchased as transit from a ―backbone” ISP adjacent to an IXP. At wholesale rates, purchased near a large IXP such as Frankfurt, Amsterdam, London, or Washington, DC, that gigabit of bandwidth would cost about USD 2 000 per month (or USD 2 per megabyte per month), which is not substantially more than it would cost to participate directly in the IXP and generate it oneself. Generally, the wholesale markets near IXPs are very competitive and extract little excess rent, and this benefits consumers. This would seem to indicate that each consumer‗s fraction of that cost is about one hundredth of a cent per month. And though this is notionally true, again the real world is not quite so simple, because users do not aggregate perfectly. Each user does not wait his or her turn and send or receive only when others are not. In fact, the protocols the Internet is built upon tend to break down as networks approach 100% utilisation; they depend upon there being a few ―empty boxcars" in the passing train to estimate degrees of contention and probability of loss and to make other critical decisions. Thus ISPs tend to operate their networks at 20%–70% average utilisation. This still represents a wholesale cost of only one twentieth of a cent per month per broadband customer.

One lesson to be drawn here is that the local broadband access network, and backhaul to the IXP, is a far greater component of the cost of retail broadband service than the costs of the IXP or the competitive backbone ISPs. Transporting bandwidth from IXPs to customers‗ sites does have a very real cost. What is striking about these prices, however, is how little they have changed in the past five years. The root cause of that stagnation had not yet become clearly apparent at the time of the 2006 OECD report on Internet traffic exchange.

The long-term, exponential growth of the Internet, the increasing performance and decreasing prices, have always been dependent upon exponential gains in technological development. From the standpoint of physics, from the standpoint of research and development, and from the standpoint of producing equipment and bringing it to market, there is no apparent end to the regular advances that can be made, if the work is performed. Since the telecommunication investment collapse in 2001, that work has not, by and large, been funded or performed. This has led to the transition from exponential growth, which held from the origin of the Internet in 1969 until the last of the optoelectronic technologies researched prior to 2001, made their way to market in 2003 and 2004, to linear growth, which has obtained since then. The earlier exponential growth allowed Internet bandwidth production to keep pace with increasing demand, with lower costs over time allowing corresponding reductions in consumer prices; the current linear growth is not keeping pace with increased demand and is not producing greater economies at larger scales; thus, wholesale and retail prices have not, by and large, fallen significantly beyond the introduction of those final Nx10 gbps technologies in any market.[9]

CHALLENGES FOR THE FUTURE

This section examines the challenges that are likely to develop over the next few years, given the current trends documented above. It discusses possible responses and adaptations by the participants in the Internet, outlines policy issues that OECD governments may face in the near future as a result of these market developments, and considers some factors they might usefully consider in addressing them.

Architecture

The major architectural change occurring in the Internet over the past five years is the transition from IPv4 to IPv6. On 3 February 2011, the Internet Assigned Numbers Authority issued the last IPv4 address blocks to the regional Internet Registries, and on 15 April 2011, APNIC, the Asia-Pacific Regional Internet Registry ran out of unreserved IPv4 addresses. Each of the other four Registries will run out over the coming year or two. At the same time, the number of Autonomous Systems, or networks, that are actually employing IPv6 addresses has grown more than six fold in the past five years, from 600 in May of 2006 to more than 3 800 in May of 2011 (Figure 1). This is a positive sign with regard to the transition from IPv4 to IPv6. There are nonetheless many thousands more networks that have yet to begin their transition.[10]

Figure 1. Number of autonomous systems advertising IPv6 addresses in the global routing table (Y axis) over time (X axis)

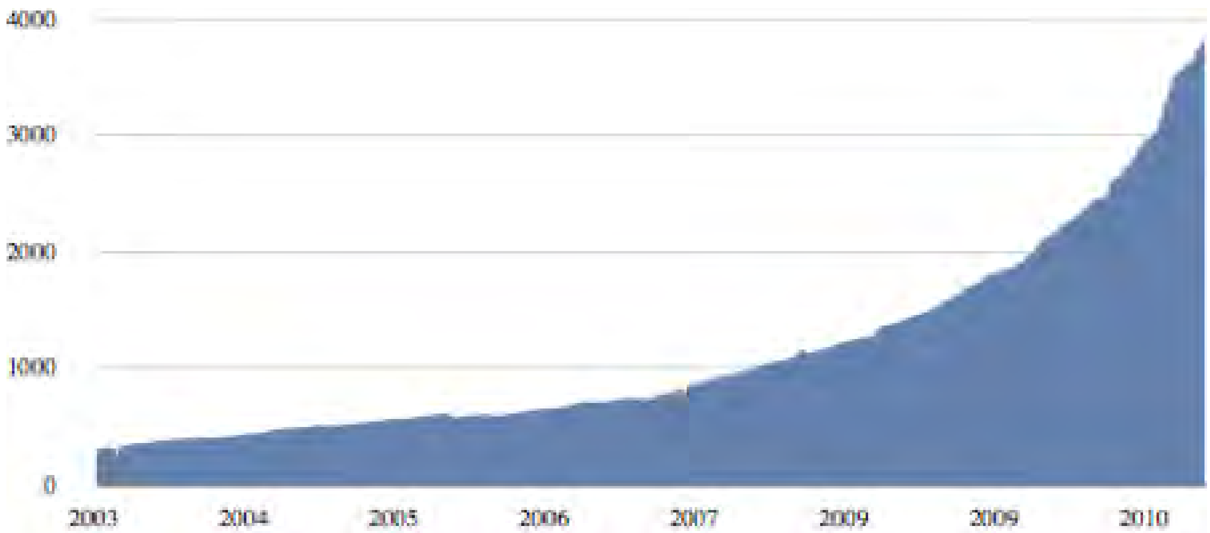

Source: Packet Clearing House.

One cause for concern is that some IXPs have not yet secured an IPv6 subnet from their regional Internet registry to facilitate IPv6 peering by their participant ISPs. This process is frustrated by the lack of IXP-specific policy in four of the five Registries. Only AfriNIC has specific policy for allocation of IPv6 addresses to IXPs.[11] Further discussion of this issue is provided in Annex 5.

Why supply of Internet bandwidth is no longer meeting demand

The dot-com boom of the late 1990s necessitated a follow-on investment boom in the telecommunications infrastructure that supported the growing use of new Internet products and services. In the wake of the dot-com investment collapse in March 2000, there was a similar collapse in funding for telecommunications infrastructure, in which investors rapidly pulled out of fibre-laying projects, the telecommunications companies that would have used them, and the equipment vendors that built high-speed routers and switches. What has been discussed far less is the fact that, at the same time investors were pulling out of all of those highly visible sectors of technology, they were also pulling funding for basic optoelectronic physics, the basic research that provides the technological means to drive telecommunications signals through fibre optic cable at ever-higher speeds. There is about a three-year lag between when a new technology is researched in the lab and when it makes its first appearance in commercially available network interfaces. Historically, each new generation of optoelectronic interface has been designed to be one order of magnitude faster than its predecessor; 10-megabit interfaces gave way to 100-megabit interfaces, which were in turn replaced by 1-gigabit interfaces, and those were replaced by 10-gigabit interfaces. When each new speed of interface is introduced, it is quite expensive, but it provides new “headroom” for growth. This headroom is consumed, and additional capacity is generally needed prior to the introduction of the next speed of interface, so “link aggregation,” or “LAG,” is performed, bundling two, and then three, interfaces of a given speed together, to provide some linear growth until the next order-of-magnitude faster interface becomes available and cost-effective. Conceptually, then, the growth of individual links in the Internet follows the pattern 1, 2, 3, 10, 20, 30, 100, 200, 300, 1000, 2000, 3000, etc. When graphed with a log scale on the Y axis, this upgrade pattern forms a nearly straight line. (Figure 2).

Figure 2: Conceptual diagram of Internet growth, showing 115% annualized growth in the form of annual 1, 2, 3, 10, 20, 30 pattern upgrades

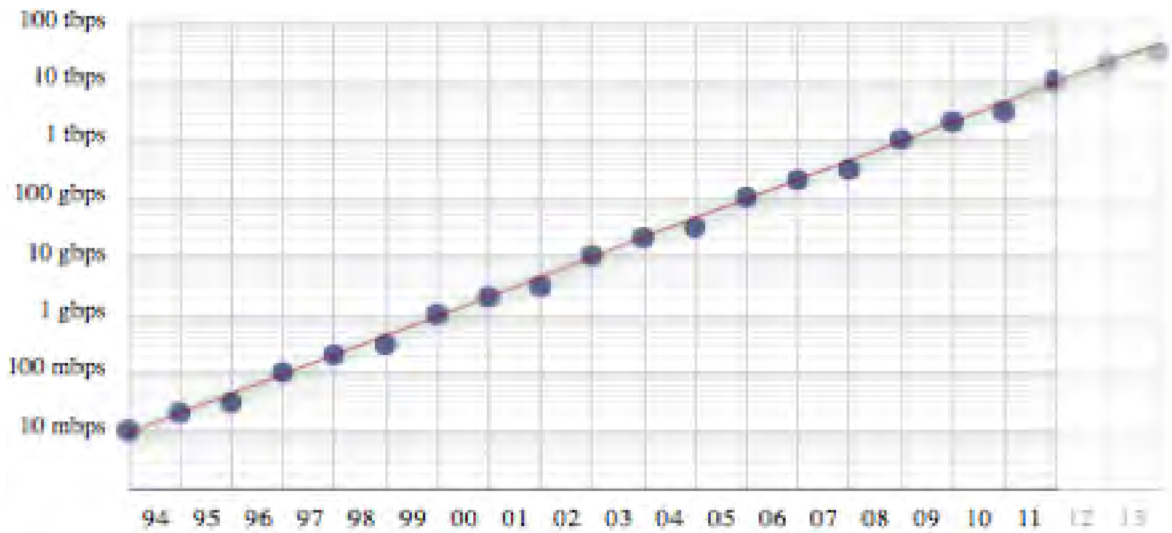

Source: Packet Clearing House.

Figure 3: Optoelectronic interface speeds used in the core of the Internet (Y axis) by year of deployment (X axis)

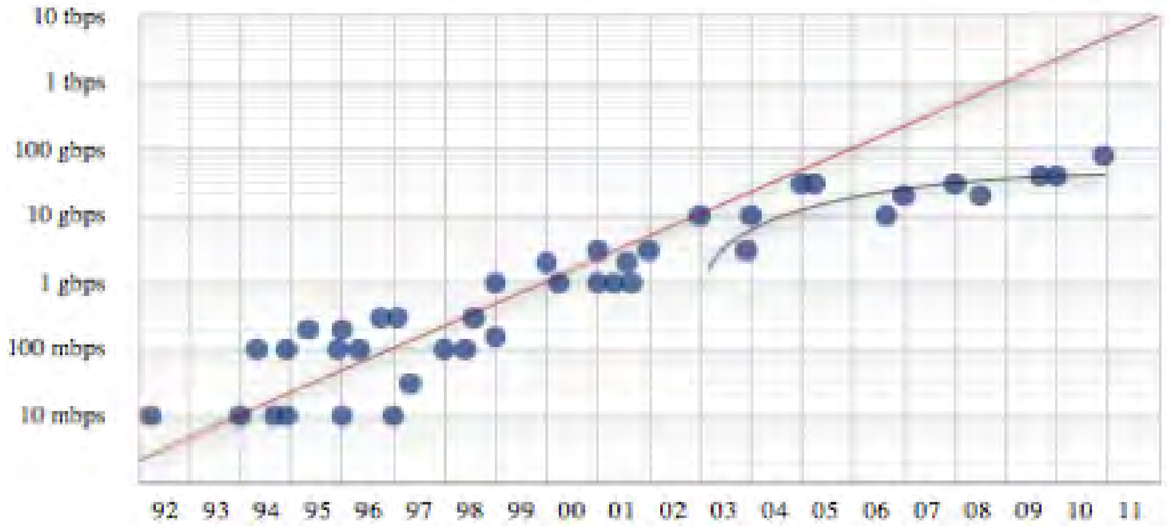

Source: Packet Clearing House.

In Figure 3, each dot represents the deployment of a new interface speed within a major IXP, and the straight line is the same 115% annual growth rate taken from our theoretical model in Figure 2. The data express a trend that begins to change slope after 2001 and hits what is essentially a hard cap after the introduction of 10 gbps interfaces at the end of 2003. Those 10 gbps interfaces are bundled together in groups of two and three. When they should be making the exponential jump to 100 gbps, 1 tbps, and 10 tbps, instead they continue on in linear fashion, with LAG groups of four, five, six, seven, and now even eight. For reference, the black line shows a linear growth rate of 500 mbps per month, which appears to better characterise the Internet's more recent trend.[12]

The inability of linear addition of components to meet exponential growth in demand may seem self-evident to mathematicians and provisioning engineers, but there are some corollary facts that make the problem a little easier for others to understand. First, each tenfold increase in interface performance has been accompanied by a minor and temporary increase in price. Essentially, the 10 gbps switches of today cost no more than their 10 mbps equivalents did fifteen years ago, despite carrying a thousand times more data. Second, it is not possible to simply continue adding on more switches to provide more capacity, even if it were possible to spend exponentially more money on switches, because the number of ports on each switch is limited, and switches must be interconnected with each other in a mesh, which consumes more ports for interconnection. Specifically, it requires n – 1 ports on each of n switches. A point of diminishing returns is quickly reached, at which the addition of one more switch actually *decreases* the amount of available bandwidth rather than increasing it.

Figure 4. **Available bandwidth (Y axis) per number of meshed switches (X axis) for two combinations of switch size and number of participants**

Source: Packet Clearing House.

In Figure 4, the point of diminishing returns is illustrated for two examples. One example is taken from DE-CIX, the Frankfurt exchange, which uses the largest available switches, 288 ports of 10 gbps each.[13] This exchange has 325 participating ISPs, each of which could hypothetically be allocated 80 gbps in the form of eight 10 gbps ports in a LAG bundle, using a mesh of fifteen switches. Adding a sixteenth switch would not increase total capacity, and adding a twenty-third switch would decrease total capacity. In the second example, using 48-port switches, which are the most common size, and twenty-five participants, which is a reasonable average at exchange points, each ISP can be allocated 50 gbps using four switches. Adding a fifth switch would not increase total capacity, and adding an eighth switch would decrease total capacity. The development of switches with more ports would superficially appear to ameliorate this problem, but in fact, it merely moves the same problem from the externally-visible network topology, to the backplane of the switch, without resolving it. Only the development of faster optoelectronic interfaces, and the resultant faster network interfaces, can solve the problem and allow the Internet to resume economically-efficient growth.

As a side note, several mathematicians were asked, in the preparation of this report, to examine the problem of modelling the complex relationship illustrated above between number of switches, number of ports per switch, number of participants, and number of ports allocated to each participant. Although the "recurrence relations" mathematics of reducing the specific cases to a general case have proved difficult, the mathematicians' outsider viewpoint did result in a startlingly "outside the box" suggestion of a new and potentially more scalable topology for multiswitch IXPs, in which each participant is connected once to each of a growing set of stand-alone switches rather than given multiple connections to a single switch that is in turn meshed with others.[14] Such a model would pose its own new complexities and challenges, and thus will not provide a "magic bullet" solution to the larger problem, but is nonetheless now being investigated further by IXP operators.

Challenges for the Internet mode of traffic exchange

In earlier stages of Internet development, Internet networks and those in the TDM space were largely separate. It was therefore relatively straightforward for them to develop separately, with TDM networks subject to traditional regulation and Internet networks free to develop a new model of traffic exchange. Now, however, convergence is creating a collision between the two market and regulatory models. For at least the past decade, voice telephony service providers have been using IP technology for their internal transport. Although many interfaces with customers and other networks have remained TDM, essentially all of the transport and switching within the networks is now done using IP. Similarly, traditional services such as voice and video have increasingly been provided using applications that run over IP, both by ―over-the-top" third party service providers and by the traditional networks themselves.

Experience suggests that, when the two models meet in the real world, the Internet or packet-centric model tends to displace the legacy or circuit-centric model. This process of displacement may lead to tensions on several levels, and to appeals for new intervention by regulatory authorities. Although, as discussed above, the Internet market has generally produced more favourable outcomes, it may also disrupt the status quo and alter the terms of trade in ways that are not seen as advantageous by some market participants. At the same time, the policy frameworks that were developed in earlier stages of liberalisation to promote goals such as fairness, competition, and universal connectivity have been built around the traditional regulatory structure. Policy makers may therefore see displacement of the older market model by the Internet as disruptive of policy mechanisms that were valuable in the legacy context. These disruptions will, of course, look different from the perspectives of different participants.

Incumbent perspective

From the standpoint of incumbent operators, the public‘s transition to Internet uses and applications can mean a loss of revenues from inter-carrier payments as TDM traffic declines. For many operators, access or termination charges are already a relatively small component of their revenues; this group includes mobile operators in markets such as the United States.[15] For others, the net flow of inter-carrier payments is small or negative; this group includes landline incumbents in many major markets, including Europe and the United States. For these players, the gradual decline of TDM traffic exchange is unlikely to pose a threat to revenues, and in some cases it may actually provide some benefits.

However, other traditional networks still experience significant net inflows of inter-carrier payments today. These include mobile carriers in Europe, although the trend in mobile termination rates has been downward in recent years. Smaller incumbent telecommunication operators in the United States fall into this category; though their rates remain high, traffic volumes subject to access charges have been declining.[16] A third group are some national incumbents in emerging economies. Many of these have garnered high settlement rates in the past and have seen revenues from this market segment decline as traffic migrates to the Internet.

Nevertheless, as previous OECD work has shown, reliance on high settlement charges and limits to entry have restricted opportunities for growth. Where national policy has facilitated the growth of Internet and mobile markets, increased consumption of those services has created social and economic benefits while also generating revenue growth to offset declining revenues from traditional settlements.[17]

Entrant perspective

Competitors whose entry has been facilitated by the legacy regulatory framework in countries that have liberalised may have a variety of concerns about the transition to Internet interconnection modes. In general, Internet transit and peering arrangements are likely to give an entrant greater connectivity at lower cost. Further, in some cases the transition may reduce an entrant‘s payments to larger players, for example, for mobile termination charges.

In other cases, however, the entrant may have a business model that takes advantage of the traditional payment scheme to produce a net flow of revenue. Within the traditional regulatory framework, in most OECD countries, the entrant has enjoyed certain guarantees, which may have established effective property rights. These may include the right to demand interconnection, to choose the point of interconnection, and to seek intervention by the regulator to impose terms on the incumbent. In the Internet model of voluntary agreements, these are not applicable.

Policy maker perspective

The gradual demise of TDM voice traffic is creating a delicate challenge for regulatory authorities. The ability of the Internet model to produce efficient results and disrupt older models argues in favour of allowing it to continue to develop without undue interference. At the same time, regulators may face calls for intervention from parties whose interests have been affected. Applying the existing TDM framework to Internet traffic will almost certainly be harmful, but policy goals of promoting competition remain valid. The challenge, then, is to determine what limited mix of regulatory measures going forward will best ensure that markets remain open to competition without interfering with the successful operation of that market.

The transition to Internet models may also create mechanical and administrative complexities for governments. For example, OECD countries have developed, and the OECD has relied upon, reporting systems in which service providers furnish data on the state of telecommunication markets. As voice traffic shifts to unregulated entities and service provision is scattered among a wider variety of entities and business models, it may become more difficult for policy makers to maintain a clear view of the state of the market. As a result, policy makers across the OECD are keen to explore which metrics are needed to inform stakeholders while not being overly burdensome to market participants. New tools, made possible by the Internet, may assist in that process. For instance, the OECD has worked on standardised reporting mechanisms for IXP traffic levels.[18]

Existing frameworks and applications of regulation

As noted above, a legal and regulatory framework to provide for the rule of law, open markets, and promote competition, is essential for the success of the Internet, as it is for other markets. As the Regional Survey in this report makes clear, the development of the Internet in some regions has been hampered by the lack of a fully liberalised legal and regulatory setting. This section focuses more specifically on the regulatory frameworks in OECD countries, as they pertain to the exchange of IP traffic.

To date, policy makers in OECD countries, recognising the value generated by the Internet model, have generally refrained from applying legacy regulatory frameworks to the exchange of IP traffic. To some extent, this result has flowed from the inherent structure of the legal framework. However, a few very limited actions have been taken to intervene in markets for Internet traffic exchange, often with unintended collateral outcomes, and many issues have been raised by parties or by the regulators themselves.

The European Commission has designated a limited number of relevant markets in which *ex ante* regulation may be warranted. Wholesale markets for IP traffic exchange are not among those identified as susceptible to *ex ante* regulation. An NRA may identify a market other than those designated by the Commission, based on national circumstances, but must do so under a three-criterion test.[19] However, voice call termination for fixed and mobile networks are among the designated markets, without regard to the technology; thus the provision of voice termination over an IP interface could be susceptible to *ex ante* regulation. This, combined with other factors, may have created an incentive for the continuing use of TDM interfaces to deliver VOIP traffic to incumbent networks.

In 2006 the Polish regulator, UKE, adopted a draft decision that would impose on the Polish incumbent Telekomunikacja Polska (TP) an obligation of nondiscrimination and transparency as regards transit of IP traffic. The draft contained a number of specific provisions, including an obligation to prepare and submit a Reference offer related to IP peering and an obligation to set prices based on costs incurred. The European Commission responded in 2010 by expressing serious doubts as to the compatibility of the notified draft measure with Community law.[20]

UKE had adopted its draft decision based on concern that TP would selectively degrade traffic arriving via third-party transit providers in order to force transit customers to purchase more expensive services directly from TP. The Commission noted that the two markets that UKE had identified, peering with TP and transit for exchange with TP, were not included in the list of markets identified as susceptible to *ex ante* regulation in the Annex to the Recommendation on relevant markets. The Commission was not convinced that UKE had met the three criteria for identifying markets not on the list.

In particular, the Commission found evidence that peering and transit were effective substitutes and could therefore be classified as separate markets. The Commission noted that indirect interconnection appeared to be an effective alternative, and that only about 15% of the total exchange of Internet traffic in Poland transited TP.

The Commission also expressed doubts that, even if the markets identified by UKE were separate, TP was capable of exercising significant market power in those markets. The Commission noted that TP had incentives to interconnect to meet its own needs for connectivity. The Commission also observed that TP‗s prices for Internet transit were trending downward, which the Commission attributed to low entry barriers and competitive constraints. The Commission‗s action is currently under appeal by UKE.[21] In the announcement of the Commission‗s decision, Digital Agenda Commissioner Neelie Kroes said: ―The Commission fully shares the objectives of the Polish regulator in seeking competitive markets, but our assessment is that regulation of these particular markets for Internet traffic exchange services is not necessary to protect consumers or competition. If the market itself is able to provide for fair competition, don‗t disturb it with unnecessary regulations."[22]

In the United States, the FCC has generally refrained from regulation of Internet interconnection. However, as a condition for FCC approval of a merger between AT&T and BellSouth in 2006, AT&T agreed to maintain, for a period of three years, at least as many settlement-free peering agreements with domestic operating entities within the United States as it had in effect at the time of the merger.[23] Although this was a limited intervention in the Internet market, it is not clear what practical effect the condition had, and it has not been translated into any more general policy.

Much discussion in the United States has focused on whether services such as broadband or VOIP should be categorised as ―telecommunication" or ―information" services. The FCC has recently sought comment on whether regulation should be applied to the exchange of voice traffic over IP interfaces, but has yet to take any action.[24]

In Canada, the NRA, the Canadian Radio-television and Telecommunications Commission (CRTC), has recently adopted a number of new regulatory provisions with respect to network interconnection for voice services.[25] The CRTC found that IP interconnection for voice traffic should continue to be carried out under bilateral commercial arrangements. It also found it unnecessary to mandate a default tariff for IP voice network interconnection. However, if a carrier is providing IP voice network interconnection to an affiliate, a division of its operations or an unrelated service provider, then it must provide similar arrangements with other carriers. The carriers are to complete the negotiation process within six months of a request for interconnection. Either party may request mediation by the CRTC staff, or apply to the CRTC for intervention if an arrangement is not concluded within the six-month period.

Traditionally, TDM traffic has been exchanged locally, often in multiple locations close to local switching centres. Regulatory frameworks for TDM traffic have reflected this pattern, giving parties requesting interconnection from an incumbent the ability to choose the point of interconnection (“POIs”).[26] In contrast, Internet networks typically find it efficient to interconnect at a few common IXPs, and in the Internet model these are mutually determined through commercial agreement. This has occasionally led to disagreement between incumbents and entrants regarding the number of points of interconnection and how these should be determined.[27]

In Canada, the CRTC recognised in its recent decision that it would be efficient for IP voice interconnection to be carried out over “significantly fewer” points of interconnection than currently exist under the TDM regime. It estimated that two such points per province might be reasonable. However, it declined to prescribe a set of POIs, leaving the determination of these points to negotiations between parties.[28]

In Australia, where the government is making a large investment in a new National Broadband Network (NBN) the issue of the number of points of interconnection in the new network has had to be examined. Larger players, such as Telstra and Optus, favour many points of interconnection, while smaller ISPs tend to favour fewer, more centralised points, which would minimize their overhead and transport costs. The government therefore involved the independent regulator in this matter.[29] In 1997 the Australian Competition and Consumer Commission required the incumbent Telstra to peer with three other large players—Optus, Ozemail, and Connect.com—thus creating what is sometimes referred to as “the gang of four.” The net result of this intervention is difficult to evaluate. It may have had the unintended effect of limiting opportunities for smaller players outside the group to pursue alternative arrangements. Since the policy was put in place, Optus, Ozemail, and Connect.com have all been acquired by foreign carriers.[30]

It is also the case that regulations that are not explicitly intended to apply to Internet traffic exchange may have that effect. For example, restrictions on the ability to export certain data, such as customer profiles, intended to protect security and privacy, may also limit the development of Internet topology and the growth of Internet assets in some regions. Similarly, tax policies in each country toward broadband and Internet businesses are likely to affect the choice of the locations for investment in Internet assets. In the United States, for example, federal law prohibits taxation of Internet services. Some leaders have recently called for international agreements to restrict the ability of individual countries to compete on the basis of their tax policy.[31]

ONGOING DEVELOPMENT OF THE MARKET

Just as the basic architecture of the Internet has proven to be remarkably adaptable, so too has the basic model of Internet traffic exchange. After a period of experimentation in the early 1990s, the model of peering and transit has converged to a stable framework that has been able to accommodate the Internet‘s growth.

The Internet itself is in a perpetual process of reinvention and transition. Traffic continues to grow exponentially. Large new investments are needed in backbone networks to support this growth. Operators of local access networks are also investing to transform the technology in their networks, provide faster access speeds, handle the increase in traffic, and offer new services; the roles and business models of all the players are evolving over time.

All of these developments lead to adaptations in the model for Internet traffic exchange. In turn, ongoing developments in traffic exchange will influence business models and play a part in determining the terms of trade among the participants. The extent to which policy makers need to become involved in this circle of change depends, in part, on how virtuous it will be in the absence of intervention. So far, the market has produced excellent results without much intervention. How is this process of adaptation likely to play out in the future, and what challenges will it pose for policy?

Adaptations of the Internet model

Regional peering

Liberalisation and the resulting growth in Internet activity in emerging economies have led to the development of regional networks with substantial volumes of traffic. With IXPs in more countries, there is less need for tromboning traffic to Europe or the United States, and a greater portion of traffic is exchanged within a region.[32] These developments raise the possibility that, when a large regional network and a global backbone network engage in peering negotiation, their perceptions of the value of the relationship may differ. The regional network may have more traffic, and offer better connectivity, within the region than the global backbone. The global network may have worldwide connectivity that the regional entity does not, and thus it may be unwilling to accept the regional network as a peer for all traffic.

One way to bridge this difference and find agreement is to treat the two types of traffic differently. In a regional peering agreement, traffic is exchanged only if bound for the region surrounding each IXP where the pair peer. As the regional network grows, it becomes able to peer with the global backbone network in more locations, potentially ―collecting the whole set" if it grows to achieve geographic parity with the global backbone network. In cases where the benefits of direct interconnection outweigh the costs, this compromise allows competitors to realise some of the potential gains from trade. On the other hand, it imposes inefficiencies in routing on both the participating carriers. For more on regional peering, see Annex 2.

Multilateral peering

Historically, most peering has taken the form of a bilateral relationship between two networks, and the fine granularity of control this provides has been viewed as desirable. But the work of managing hundreds or thousands of such bilateral relationships creates cost and friction within large networks that are present at many IXPs, and the complexity of managing more than even a single one may prove a challenge for the smallest networks, those that have only just arrived at their first IXP. By joining a single multilateral relationship with many peering partners present at an IXP rather than concluding bilateral agreements with each of the other parties, a network can reduce transaction costs. Among the most notable results of the PCH survey reported below are the large number of parties to some of these agreements and the resulting degree to which the number of relationships facilitated through multilateral agreements dwarfs the number embodied in bilateral agreements.

Structural change in the Internet

Share-shifting backbone carrier

UUNET, the largest ISP of the 1990s, was acquired by Metropolitan Fiber Systems in 1996, which in turn was acquired by WorldCom, which itself merged with MCI in 1998, raising concerns that this concentration would have a dominant position on the Internet. By the time of the peak of this concentration, immediately before its 2002 bankruptcy, it was providing some sort of service to nearly 9% of the Internet, either directly or indirectly. By 2005, when Verizon purchased the remains, UUNET had deteriorated substantially, and even with Verizon‘s careful husbanding of its investment Renesys estimates that the present-day network today ranks only eighth among backbone networks, behind carriers such as Telia and Tata, assuming that two recently announced mergers proceed as planned (Figure 5).[33] Clearly, UUNET and its successors have not achieved the dominant position that was predicted in 1996.

Figure 5: Global share of IPv4 addresses routed by each backbone carrier after proposed mergers.

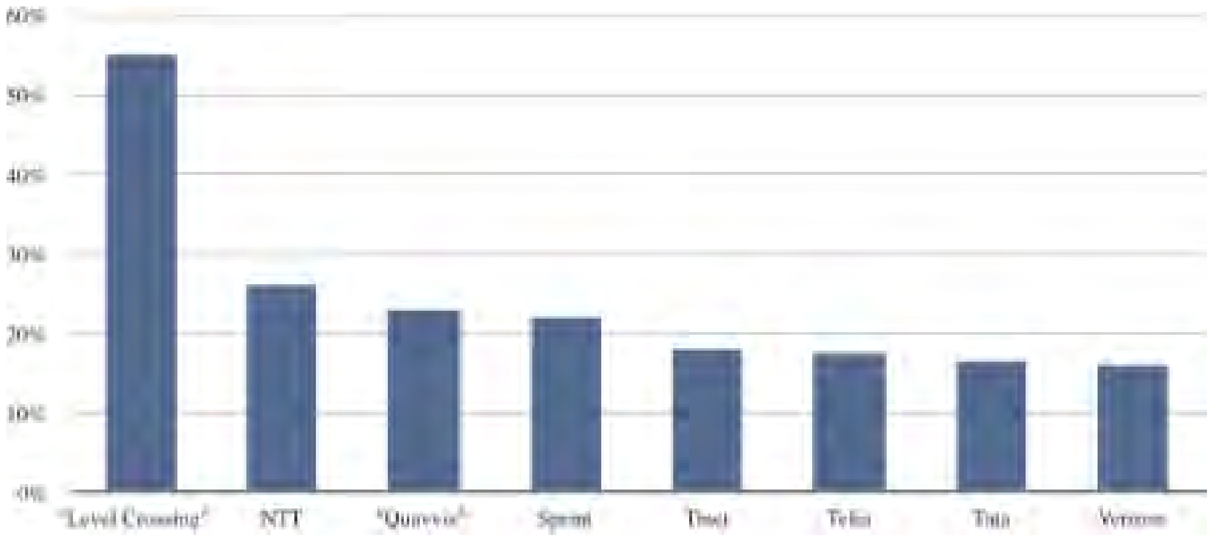

Source: Renesys

In April 2011, Level 3 announced plans to acquire Global Crossing.[34] By most measures, these are today the two largest backbone carriers. Two weeks later, a merger of Qwest and Savvis was announced. These announcements have renewed concerns about concentration at the core of the Internet. Renesys has estimated that the products of these mergers, which they have dubbed ―Level Crossing” and ―Qavvis,” would rank first and third, respectively, among backbone carriers. Although at first glance this might raise concern, the experience of UUNET indicates that nothing remains the same in this space for long. Indeed, within these rankings Sprint was until recently number two, having been overtaken in late 2010 by Global Crossing, and since then by NTT. Further, as noted earlier, the relative market share of the backbone carriers as a group has been falling over the past decade.[35]

There are several reasons for this evolution in market structure. The backbone carriers‘ market strategy sought to monetise the value their networks created, but in the process they sacrificed growth in comparison to smaller networks pursuing other strategies. Smaller networks began in the 1990s to peer more aggressively among themselves—so called ―donut peering”—to minimise their transit costs and improve quality of service, and thereby reduced their reliance on the large, legacy carriers. The geographic expansion of the Internet, and the growth of IXPs around the world, further reduced the role of the global backbones. The growth of CDNs created alternatives to transit for the delivery of traffic. Finally, some of the backbone carriers have given greater priority to expanding other aspects of their businesses, such as providing enterprise solutions to global business clients and entering the CDN market, and perhaps less priority to transit as the source of their growth.

Growth of CDNs

Over the past decade the topology and market structure of the Internet have changed significantly as a new category of networks has reached maturity. These content delivery networks (CDNs) serve as aggregators of content, systems for delivery of traffic directly to the terminating network, and providers of quality-enhancing inputs, such as caching of content close to the end user. Like any other network, CDNs often seek to enter into their own peering agreements.[36] Local caching also reduces the volume of traffic that needs to be delivered to the terminating network. But cost is not the only driver of the growth of CDNs. Providers of online services, such as the BBC, Google, Netflix, and Hulu, seek to improve the quality of the experience they provide to their customers. More direct delivery, fewer intermediate hops, and local caching reduce latency and improve the quality of service.

This market segment has grown rapidly. A 2009 study by Atlas Internet Observatory estimated that the top five ―pure play” CDNs—LimeLight, Akamai, Panther, BitGravity, and Highwinds—represented close to 10% of Internet traffic.[37] Akamai‘s revenues have quadrupled in the past five years.[38] This function is not, however, performed solely by stand-alone players. Online service providers purchase inputs from CDNs, but in many cases they also self-supply. Google, for example, carried about 6% of Internet traffic in 2009, according to Atlas. A substantial part of this is generated by YouTube, which Google owns. Google and Comcast, neither of which appeared on Atlas‘s list of the top ten Internet networks by volume in 2007, just two years later had risen to number three and number six on that list, respectively.[39] Backbone networks themselves have expanded their businesses to provide CDN services. Level 3, for example, has been building its participation in this market and recently entered into an agreement to provide CDN services to Netflix. Indeed, some analysts have speculated that an important consideration in Level 3‘s decision to acquire Global Crossing was the latter‘s CDN business, and TATA‘s acquisition of BitGravity serves the same function.

Share-shifting services

For decades, circuit-switched voice service was the primary offering of the world‘s telecommunication networks, accounting for most of their usage and revenue. Today, voice traffic accounts for only a microscopic share of the bandwidth exchanged among networks worldwide. This is true in part because voice is a low-bandwidth application and also because voice plays a much smaller role in the mix of services that people consume. Further, a substantial portion of voice traffic is already exchanged over IP. For all these reasons, the transition of global traffic exchange from TDM to packet switching is already largely complete. Nevertheless, legacy TDM pricing models and regulatory frameworks remain, so the exchange of TDM voice traffic still accounts for a proportion of carriers‘ revenues that is far greater than the share of traffic or costs it represents.

As consumers and firms make use of a wider array of services, voice is less prevalent in communication than it once was. Sequential age cohorts have embraced, first, voice, then email, short-messaging services (*e.g.*, SMS and Twitter), and social media sites (*e.g.* Facebook), and more recently video as their preferred means of communication. In 2010, average mobile voice usage per customer in the United States, which had grown rapidly over the past decade and is today the highest in the world, declined slightly for the first time on record. As with Internet bandwidth trends, this follows the lead of Japan, which experienced its peak voice-minutes use some ten years earlier. Meanwhile, according to Nielsen, the average teenager in the United States between the ages of 13 and 17 sent or received 3 339 texts per month in 2010—more than six for every waking hour (for teenage girls, the average is 4 050).[40] Voice calling declined across all age groups, but the drop was particularly steep for teenagers, at 14% in one year (2Q 2009 to 2Q 2010). Those customers who continue to use voice services are often willing to accept limited voice quality in order to obtain better mobility, lower price, or features such as follow-me services.

More broadly, the mix of services transmitted over the Internet has shifted significantly. Among the top ten global applications (by percentage of Internet traffic) identified by Atlas, video was the fastest growing between 2007 and 2009. The share of traffic related to peer-to-peer (P2P) applications, the category that had raised the most concern about straining Internet capacity, declined the most.[41] Part of this decline is due to changes in consumer behaviour as streaming and direct downloading of content become more popular. Part is attributable to improved P2P software and more efficient algorithms for minimizing the load that P2P presents to networks. For example, a voluntary industry coalition called P4P, which includes both network operators and P2P application providers, has developed a cooperative method for sharing routing information to allow P2P applications to identify the least costly location from which to retrieve any requested file.[42] Part of the decline is also the result of changes in legal requirements in many countries, and part reflects efforts of P2P providers to make their services less readily observable. Atlas found, for example, that the portion of traffic associated with P2P was only 0.85% in 2009 using their customary methods, though it believes the true figure is closer to 18%.[43]

The most rapidly growing categories of service on the Internet are video services and web or cloud applications that move content and software from the local desktop into datacenters. Users are increasingly obtaining video and other media through streaming or direct download, in preference to P2P applications. The annual survey published by Sandvine, an Internet traffic measurement company, suggests that, in North America, 45.7% of the downstream traffic on fixed networks is generated by what they categorise as Real-Time Entertainment (online video). The largest single contributor to this is Netflix, which by itself accounts for more than 20% of downstream traffic during peak hours (8 to 10 pm.) The growth of these services is mirrored by the growth of CDNs, since each has driven the other. However, P2P is still the largest driver of upstream traffic on North American fixed networks, accounting for 53.3% of total bytes in Sandvine‘s results.[44]

As the mix of services has shifted, and as those services as well as CDN functions have been supplied by different entities, the distinctions between backbone networks, access networks, and media companies have blurred. For example, Comcast‘s role in this universe has changed substantially in a short time. In 2007 it was primarily a local cable operator, lacking its own backbone facilities, mainly focused on residential video and broadband services and highly dependent on upstream transit suppliers. By 2009 it had become a major provider of voice services, a net exporter of traffic, the sixth largest network by traffic volume, and the largest user of IPv6 addresses on the Internet.[45]

Investment in access networks

Operators of local access networks around the world must invest substantially to keep up with the growth of traffic, to provide consumers and businesses with higher-speed broadband connections, and to develop new sources of revenue as older services, such as fixed-line rentals, decline.[46] Some must also satisfy the expectations of policy makers for infrastructure development to meet national goals.

Access networks around the world are upgrading their infrastructure and technology. For mobile networks this means fourth-generation standards like LTE; for wireline networks, faster local access technology such as fibre. Accommodating more traffic also means additional fibre backhaul facilities and greater connectivity to the Internet.

Both fixed and mobile networks are planning their transition paths around a set of industry standards called IP Multimedia Subsystem (IMS), originally developed to facilitate multimedia services over 3G mobile networks. In part, IMS is a platform that will allow operators of access networks, and their suppliers, to provide advanced, integrated services in competition with services provided by online service integrators. IMS provides a standard to which future equipment can be built, with standard interfaces that allow equipment from different suppliers to interoperate.[47]

As broadband subscription has increased over the past decade, this extensive growth generated additional revenues in subscription fees, even if some revenue was lost as legacy services declined. Now, however, as the rate of subscription growth tapers off, increased usage per broadband subscriber does not automatically generate additional revenue. Access network operators are challenged to develop business models that allow them to invest economically. One component of such a model is the provision of integrated, proprietary services such as video, games, integration between wired and wireless services, and quality-enhancing features. Competition between the IMS platform and service provision over the Internet is potentially beneficial, so long as it can take place within a reasonably competitive framework. Policy makers may have reasonable concerns about whether IMS architecture allows operators to affect this competition adversely. The new business model might also include new, more flexible pricing of broadband services to take account of the wide range of usage demands that different customers place on the system.

In several OECD countries, governments are now making significant investments in new, publicly owned, broadband facilities. While these investments have furthered national objectives in augmenting broadband capabilities, this ―deprivatisation" may create the potential for tension with the principles that have guided liberalization. Privatisation has been an important element of liberalisation, in order to avoid possible conflicts between government as policy maker and government as network operator. By re-introducing government in the role of network operator, these investments may recreate the potential for the kinds of conflicts liberalisation intended to avoid.

NEW MODELS FOR POLICY

The Internet environment is in constant transition, with the roles and relationships among its participants never remaining the same for long. As the transition from TDM to the Internet model continues, what mix of policies can OECD governments choose to promote the best possible market outcome? This section reviews some areas of concern and possible directions for policy.

Establishing a bright line

The application of the legacy regulatory framework to Internet traffic exchange is likely to cause significant harm to the development of the Internet. For this reason, while they retain the ability to impose regulation if necessary to address significant market failure, regulators should set a very high threshold to justify any interference in this market.[48] Still, the exchange of traffic on a TDM basis will continue to decline for years to come. Each network should be able to plan the transition away from TDM on its own schedule. Since many networks will exchange both TDM and IP traffic for the foreseeable future, the most reasonable course of action for a regulator is probably to make clear that exchange of traffic over an IP interface will continue to be carried out on the basis of voluntary commercial agreements without regulation (as long as the very high threshold is not met). At the same time, any market participant with standing to request a TDM interface from another provider would continue to have that right, within the existing regulatory framework.

This approach would draw a ―bright line” between the two policy models and establish clear expectations among market participants. Leaving the TDM framework in place would provide an alternative means for interconnecting voice traffic for any carrier that wishes to take advantage of it. It would allow each carrier to choose the time path for its own transition. Rather than establish a time frame for phasing out existing TDM obligations, the regulator can simply let them transition away as TDM interfaces themselves do.

One possible effect of this approach is that networks that benefit the most from the current arrangements for TDM traffic may seek to slow the transition from TDM to IP interfaces. Ultimately, though, as its IP data traffic grows, each network will need to negotiate agreements for the connectivity necessary to carry that traffic. The desire to reduce its own costs, and the possibility of obtaining better terms from other networks, will over time create countervailing incentives to eliminate TDM interfaces. Further, reform of TDM inter-carrier charges would help to minimize the incentive to slow the transition.

It should be emphasised here that regulation of the exchange of traffic, which is the subject of this discussion, is not the same as policy designed to open markets and assure access to facilities for the purpose of providing competitive services.

The question of voice

Some observers who agree that interconnection obligations should not apply to the Internet still raise concerns that, for voice service alone, some regulation may be required.[49] For VoIP services routed to a telephone number, the concern is that the terminating access network is ultimately the only one that can complete the call; hence the possibility that some terminating monopoly power may persist, since the routing flexibility discussed earlier would be constrained.

Certainly this argument does not apply to traffic exchanged over standard IP interfaces, without the use of a telephone number for routing. Further, the risk of recreating the current TDM pricing regime in the Internet should give regulatory authorities reason for extreme caution. Therefore, the most reasonable balance may be to adopt a series of policy measures that fall short of regulating interconnection but monitor the development of the market, with the possibility of applying remedies should they become necessary.

For example, if the traditional TDM regulatory model is maintained as discussed above, then a voice provider that dislikes the terms offered for an IP exchange of voice traffic can simply demand TDM interfaces for that traffic.[50] This solution is well proven; it is the way most VoIP traffic terminates to the PSTN today, and in most cases the necessary arrangements already exist (Box 1). Though the cost of converting traffic to TDM in order to present it to another carrier is not zero, it is not large enough to confer a significant advantage to the incumbent‘s own voice service. The inter-carrier charges that apply to this exchange may be more significant, depending on the level of the charges and the balance of traffic. Again, reform of existing TDM charges would mitigate that concern.

Box 1: Voice over IP and "VoIP Peering"

When VoIP calls are sent and received from computer to computer, without the use of a telephone number, the traffic generated by those calls is routed and exchanged in the same manner as any other Internet traffic. VoIP gateways such as Vonage, Jajah, SkypeOut, and Google Talk give the VoIP user the ability to call PSTN subscribers using their telephone numbers. In this case, the traffic is routed to an access network that hosts the recipient, either by the VoIP provider or an intermediary, and converted to TDM prior to delivery to the called party. TDM inter-carrier charges may apply to this exchange.

Some VoIP services also provide the VoIP customer with a telephone number, so that calls can be received from the PSTN or from another VoIP customer who dials the recipient's telephone number. In many cases, such calls are routed through the PSTN, because the PSTN is equipped to identify the correct routing for the call based on the telephone number. This may involve two conversions, from IP to TDM and back, or it may be entirely VoIP, but it always requires the payment of inter-carrier charges.

It is more efficient for the call to be routed directly to the terminating VoIP provider over the Internet exclusively. To do this, the originating carrier must be able to translate the telephone number into an IP address it knows how to reach. Similar database lookups have been used to route 800 calls and to provide number portability. But industry standards must be established, investments made, and arrangements put in place for such a system to work.

The ENUM standard, developed by the Internet Engineering Task Force (IETF) as a means of associating telephone numbers and IP addresses, was put into effect in 2000.[51] It was originally intended as a public capability that could be accessed by any Internet participant. In practice, however, implementation efforts have been constrained to produce arrangements in which only one carrier—the one through which the number has been assigned—is able to terminate a call.

Several industry groups have worked to develop frameworks for additional terms of reference for VoIP peering, but none has yet been widely adopted, since VoIP is just another application riding over the top of the layer 3 IP traffic, and peering negotiators are loath to introduce special terms or exceptions where this would require renegotiation, possibly to their detriment, of many agreements. Further, there is a fine line between industry standard-setting efforts and co-ordination to establish approaches for inter-carrier charges for the delivery of VoIP traffic.[52] The former is a necessary function in the further development of VoIP services; the latter may warrant careful monitoring.

It may also be possible for government to promote the development of efficient peering arrangements to facilitate the exchange of VoIP traffic The terminating carrier‘s ability to exercise control over voice termination exists only to the extent that no other provider can route calls directly to that number. Therefore, policies to promote the development of ENUM solutions that are not carrier-based, but are available to any voice provider, would help to alleviate the need to regulate the exchange of VoIP traffic. Similarly, competition law measures to discourage any joint efforts to administer pricing would help to minimize the risk posed by any terminating market power.

Finally, the existence of intermediaries who maintain interconnection arrangements with the incumbent could also mitigate the risk of anticompetitive behaviour toward smaller players. In the United States, the FCC has found that rural incumbents must accept traffic presented to them by other carriers, regardless of whether the traffic originates from a VoIP provider.

Once direct regulation of the exchange of VoIP traffic is adopted, it is likely to become permanent, even after the concerns that motivated it have vanished. Such regulation has the potential to do substantial harm, so regulators will be wise to exhaust all available alternatives before considering such action.

Reform of TDM interconnection

Several OECD countries have undertaken efforts to reform their existing frameworks for the exchange of TDM traffic, with a view toward moderating the level of inter-carrier charges and, in some cases, rationalising their structure. In the United States, the FCC adopted in October 2011 significant reforms for TDM inter-carrier compensation, as well as a related reforms to its universal service fund.[53] In Europe, regulators have established a glide path for a downward transition of mobile termination rates. In Mexico, recent court decisions have moved the regulator‘s efforts forward to establish rules for interconnection with Mexico‘s largest mobile carrier.[54]

Given the continued decline in TDM voice traffic, seeking a theoretically perfect framework for TDM charges may not be a good use of resources. But where some charges are relatively high, or where an inconsistent application of charges distorts incentives, action by regulators to moderate these disparities could help to reduce artificial incentives for carriers to preserve TDM arrangements when VoIP exchanges might be more efficient.

Allowing for disagreement

For a commercial agreement to be voluntary, parties to it must be able to decline or terminate it. The ability of each participant to do this is the fundamental source of the market discipline that guards against unreasonable behaviour such as selling transit to cybercriminals, or permitting excessive levels of spam or malware. Unfortunately, de-peering or the threat of de-peering may also be used as a lever to attempt to compel commercial behaviour.

When there is disagreement between commercial parties, there is an incentive for one or the other to seek intervention by regulatory authorities. Any real possibility of such intervention will overhang commercial negotiations. In 1995, for example, Level 3 notified Cogent of its intent to end their peering agreement. When the notice period specified in the contract expired, Level 3 terminated the arrangement. Because many of Cogent‘s customers had no alternative path to Level 3, their access to the subset of Level 3‘s customers who were single-homed was cut off. In the United States, Cogent made appeals to the FCC and Congress for intervention. Eventually a new agreement was reached. It is not clear to what extent the new terms reflected only commercial considerations and to what extent they were influenced by the visibility of being fought out in the press and in front of potential regulators.[55]

It is clear that a form of implicit regulation can be applied by allowing market participants to believe that intervention is possible. In a few cases, such as in an emerging economy where the market is not well established and the incumbent has a protected position, a certain amount of informal discussion may be beneficial. In well-functioning Internet markets, though, allowing the possibility of intervention to influence parties‘ negotiating positions runs the risk of distorting the outcome in ways that are not beneficial. For this reason, it may be advisable for the regulator, having drawn a bright line, to stick to it with a certain amount of clarity. To be sure, there must be some limit to the amount of disruption a private dispute can be allowed to cause, just as there is with a labour dispute, for example. But if this range of outcomes is too narrow, or applied inconsistently, it will have negative consequences for the market.

Pricing models to support future investment

Some analysts have suggested that the current model of Internet traffic exchange does not provide sufficient revenue for access networks to fund the investments needed to build high-speed local broadband networks. It is also argued that the model does not give online service providers the correct market incentives to optimise their use of network resources.

The Internet includes many platforms which intermediate among different economic actors. These include networks such as Level 3 or Telefonica, as well as content, service, and application aggregators such as Google, itunes, or Netflix. Platforms of this kind are sometimes referred to as creating "two-sided markets" in which revenue might be derived from entities on either "side" of the platform. Other examples include newspapers and magazines, credit card networks, gaming platforms such as Nintendo or X-Box, and operating systems such as Microsoft Windows.[56] For example, a magazine might derive revenue from readers, through sales or subscriptions, or from advertisers, or some combination of both. Google obtains most of its revenue from advertising, and relatively little from users. Much work has been done in recent years to examine network pricing and net neutrality issues within the framework of two-sided markets.[57]

One allegation is that, when CDNs establish peering arrangements with local access networks, they deprive them of transit revenue.[58] This may be true, although it is equally true that such a peering arrangement saves the access network from paying transit charges to receive the content the CDN wishes to deliver, and the CDN could just as easily claim that the access network was depriving it of transit revenue. The level of countervailing market power held by the CDN (or its client) is difficult to judge, but the fundamental point is that the agreement is voluntary, so, complaints notwithstanding, the access network agrees to peer only if it is made better off doing so. The CDN saves money on transit, but it also brings real resources into the system in the form of transport facilities, caching, and other inputs. Certainly the development of CDNs has displaced some existing transit revenues, but that is competition at work, and directly benefits the consumer. Ultimately, CDNs are just more networks, like backbone and access providers; although there are clearly defined categorical differences between different kinds of networks, they are orthogonal to the business models they choose to employ. Whether a network hosts more content, or more users, does not affect its need, or the legitimacy of its need, for peering and transit.

The authors of a recent AT Kearney study offer four pricing models that might, in their view, offer the operator of an access network the opportunity to maintain a sustainable business model while making investments necessary to cope with Internet traffic growth.[59]

One option is for the access network owner to adopt a range of pricing models, perhaps, as already occurs in several countries, incorporating variable usage pricing or a series of nonlinear offers to accommodate different levels of demand. This appears to be a reasonable approach—subject, of course, to the presence of sufficient competition to constrain rates to be reasonable and to the normal application of competition law.

A second option is to introduce what are termed "traffic-dependent charges for all traffic". It is presumed that this termination charge would be uniformly applied to all interconnecting carriers and to all traffic. While there is nothing in the economics of two-sided networks to preclude the payment of fees by interconnecting entities on the basis of voluntary agreements, the imposition of a non-voluntary, mandatory charge across all networks raises many concerns. This would require enforcement, either by the government or by a coalition of the access networks. This model appears to resurrect the legacy pricing methods of the past and to apply them to the Internet. In addition to raising costs for all interconnecting networks and their users, it would deprive the Internet of the ability to pursue the kind of evolution and discovery of new models discussed elsewhere in this report. Providing adequate investment for new networks is a worthwhile objective, but those networks should have to earn their revenue by providing value businesses, consumers, or interconnecting networks are willing to pay for. Governments should not support this approach, and they should prevent any collusive action to impose such a system.

A third option is to implement Quality of Service (QoS) over the public Internet. This would involve an end-to-end co-ordination of QoS across all, or at least a critical mass, of networks across the entire Internet. The possibility of this ―inter-provider QoS" has been discussed for many years but has failed each time implementation has been attempted, so it seems unlikely that the necessary coordination could be brought about now. Further, for revenue from this approach to reach the terminating carrier, a cascading system of charges is envisioned, which again sounds dangerously close to a legacy settlement system. Even if such a system could be created, it would impose tremendous costs, and it is doubtful that it could promote the kind of experimentation and evolution that have made the Internet so efficient.

The final option is that access networks could offer enhanced-quality services based on voluntary commercial agreements. It appears that this idea could be combined with the first option. In effect, this proposal is simply for the access network to become a CDN and offer services in competition with Akamai, just as other networks have done. These services could take the form of caching within the access network or provision of direct transport arrangements to online service providers. This option appears to be workable, provided that all agreements are voluntary, that competitive conditions are maintained, and that competition law enforcement is available if anticompetitive behaviour should develop. In fact, it appears that this option is being exercised by a range of networks today, though we have limited information about the details of these agreements or how they were negotiated..It is not clear how much demand there is for quality of service enhancements, and in what form. As noted above, most voice users seem to be willing to accept the level of quality that best-effort VoIP services provide today. One-way video services are not particularly sensitive to latency, although other factors may be important. As memory has become less expensive, more video content is cached close to the customer, who is unaffected by quality issues in the transmission between the content provider and the cache.[60] It is therefore uncertain how much additional revenue this line of business could generate to improve the business case for future investment.

Nonetheless, there is no obvious reason additional access networks should not be able to implement CDN technologies, just as many other networks have done.

Traffic exchange and network neutrality

Just as convergence is creating a collision between the TDM and Internet exchange models, so too is it likely to create a juxtaposition between the debate on network neutrality and the development of the Internet market. Network neutrality is, in effect, about the terms under which content is passed from one network to another.[61] As the structure of the Internet market changes, and new dimensions are added to peering agreements, the model of bilateral (or multilateral) agreements on which the Internet is based may be the vehicle through which important aspects of this question are answered.

The growth of CDNs has already created a market in which quality-enhancing resources are exchanged, using the same basic model of commercial agreement that has always been the basis for the Internet. One question is whether any particular group of networks should be excluded from participating in this market, such as local broadband access providers. It is not clear why any such restriction should apply.

In December 2010, a dispute arose between Level 3 and Comcast.[62] These two companies had an existing peering agreement. Comcast, in turn, had already established agreements with other CDNs, including Akamai, which involved some payment to Comcast.[63] Level 3 then entered into an agreement to provide CDN services to Netflix, delivering Netflix‗s video content. This would result in Level 3 delivering a much greater volume of outbound traffic toward Comcast than it had previously done; nearly a quarter of all households in the United States stream video from Netflix, and Netflix has 61% market share, compared to 8% for Comcast‗s own competing service.[64] This would oblige Comcast to build additional capacity to carry the traffic to Comcast‗s customers.

After first agreeing to continue to honour the existing peering agreement, Comcast then demanded that Level 3 purchase transit for the Netflix portion of the traffic it exchanged with Comcast. The other traffic between Comcast and Level 3 would remain under the terms of the peering agreement. Note that in some ways this would be similar to a regional peering agreement, except that in this case the distinction between the two groups of traffic would be based on one of the endpoints of the traffic rather than on geography.

Like the dispute between Cogent and Level 3 in 2005, this one was fought in the press and before the FCC.[65] Level 3 argued that Comcast‘s position was a violation of network neutrality and covered by the FCC‘s recently adopted policy on that subject. Comcast maintained that the dispute was simply a negotiation over a peering agreement. Level 3 in this case was on the other side of the dispute, compared with the 2005 case. Previously, it had defended its right to de-peer Cogent; now it was the one being de-peered.

In February 2011, in response to a question at a congressional hearing, the chairman of the FCC, Mr. Genachowski, expressed the view that the FCC‘s recent order on network neutrality was focused on protecting broadband consumers, not on peering disputes. The network neutrality rules ―don't change anything with existing peering agreements," he said.[66] Though his answer is not a formal statement of policy, it appears to indicate that the FCC does not intend to intervene in this dispute.[67]

This leaves the two parties to negotiate. Without presuming the outcome, it seems that both sides have something to gain from an agreement. Comcast would receive the video anyway, in the absence of an agreement with Level 3, but would likely have to pay transit to receive it. Level 3 may have made quality commitments to Netflix that it could fulfill only by delivering the video traffic to Comcast.

More broadly, although the FCC may not consider its network neutrality policy to apply to peering agreements, the market for peering agreements may be in the process of addressing at least some of the concerns raised in the debate. As agreements among online service and content providers, CDNs, and access networks are negotiated, a balance is being struck on the extent to which quality-enhancing resources will be brought to bear, who will provide those resources, and on what terms. The best course for regulators at this point may be to monitor and observe this process, as the FCC appears to be doing. It is not clear how intervention in favour of content providers, or access networks (or CDNs, for that matter), would improve the outcome.

Given the rapid growth of CDNs over the past few years, there is no market evidence that access providers have been able to prevent the delivery of services on this basis or to extract unreasonable terms. On the contrary, some observers, noting the concentration of large online service providers, have wondered whether their countervailing market position may be too strong.[68] In many cases, parties have agreed to exchange traffic between CDNs and access networks on a peering basis. Where payment has been made, it has evidently not been large enough to slow the growth of the CDN segment. Remembering that transit payments on the Internet are very small compared to traditional inter-carrier charges, it appears unlikely that any amount a CDN might agree to pay would materially affect the competition between online video services and the proprietary services of the access network. If a distortion of this kind should develop, intervention can be considered where necessary. At present, it appears that voluntary agreements between CDNs and other parties have helped to reduce costs and improve the quality of online content delivery.

REGIONAL SURVEY

The state of development of the Internet economy differs greatly by region. As a general matter, over the past five years the Internet has expanded geographically, bringing more Internet assets and better service to many markets around the world and supporting the dramatic growth of broadband subscription in emerging economies. However, significant differences in results remain. Some of these differences are due to features such as population density, geographic features, or language commonalities. Others have historically been thought of as following general trends in economic development, though the success of wireless mobile penetration has discounted this view. Telecommunications services are midway up the hierarchy of needs; education, clean water, and the like are necessarily a higher priority—though it is difficult to make progress in these areas without improving communication availability. Many other factors, however, are the effects of the regulatory environment. In nearly every region of the world, regulatory decisions constrain Internet development far more than technology, geography, economy, or demand. These regulatory decisions are rarely intended to benefit Internet users; they are nearly always defined by other communities of interest, at the expense of Internet users. In this section, the state of Internet development in each of a set of loosely defined regions of commonality is addressed, with considerations of production (IXPs), distribution (ISPs), and aspects of policy or regional constraints where relevant to IXP and Internet development.

Africa

Networking in Africa tends to radiate out from the hubs of greater affluence or regulatory liberality that exist in South Africa, Egypt, and the East African Community. Central Africa has few resources and little attention to turn to networking, being more immediately concerned with matters closer to basic needs. North Africa has the natural advantage of being adjacent to the fiber-filled Mediterranean but, with the shining exception of Egypt, this has been counteracted by the disadvantage of regulatory protected incumbent monopolies and conservative governmental policy. The largest and most prosperous of the West African countries, Nigeria and Ghana, have active Internet markets, and Nigeria has even had IXPs for some years, but they do not thrive to the degree that their population and spending power would predict. This is largely due to governmental policies that do not always provide confidence for investment and promote competition, including in areas such as undersea cables.

Egypt, despite a brief interruption of Internet service during its January 2011 revolution, has long been one of the bright spots on the African Internet map. Egypt was the first Arab League country to build an IXP, ten years ago, and it was also among the first African countries to host root and top-level domain nameservers, making itself more resilient to connectivity failures. It now has a second IXP and active competition at all levels of the ISP market. Egypt‘s regulatory and communications ministry have long been among the most active in promoting technical initiatives, in large part because of their active exchange of staff and know-how with the ISP community, the Library of Alexandria, and the Regional Internet Registries. The Egyptian Ministry of Communications and Information Technology was among the leaders of the very successful effort to introduce Arabic script to the domain name system, and it also hosts the disaster-recovery site for AfriNIC, the African Regional Internet Registry. The Library of Alexandria deserves special mention, as the single largest consumer of Internet bandwidth in Africa. The Bibliotheca Alexandrina, or Maktabat al-Iskandariyah, is a modern twenty-first-century library, built on the site of the Ptolemaic ancient library and centered on the notion that knowledge should be shared; it has vast electronic collections, most of which are freely copyable. It is the first library to harness the strength of Internet-based information sharing, in a way that is understood intuitively by a generation that has grown up with peer-to-peer file sharing. It is trilingual, with collections in Arabic, French, and English, and it is one of the three sites of the Internet Archive.

South Africa, after an uneven period of transition — during which preference was given to monopoly provision of services which considerably hampered development of Internet access — has managed a return to its status as the largest Internet market on the African continent, though it may lose that title to the more populous Nigeria in time. Nonetheless, competitive ISPs persevered, building an IXP early, in Johannesburg in 1996, and maintaining it despite many obstacles. One of those ISPs, Internet Solutions, was the first African network of any kind to grow to international backbone carrier status, establishing routing hubs in London, New York, and Hong Kong, China and participating actively in international Internet governance processes. The Cape Town IXP was one of the first and most dramatic failures of an IXP, ten years ago, but was successfully resurrected in 2009. The only African country with more than two active IXPs, South Africa also has an exchange in Grahamstown and will shortly have a fourth in Durban.

The member states of the East African Community—Kenya, Uganda, Tanzania, Rwanda, and Burundi—have successfully acted together to aggregate Internet traffic and commercial opportunities, despite their relatively sparsely distributed populations. All five countries have IXPs (the state of the Burundi exchange is unknown to the authors), and Kenya and Tanzania each have two. Kenya, Tanzania, and Uganda have active cross-border competition, which should serve as a springboard for some of their ISPs to make the leap into intercontinental backbone network status. All were hampered by a lack of international fibre connectivity until recently, when three cable systems, SEACOM, TEAMS, and EASSy, arrived virtually simultaneously in 2010, providing both redundancy and competition. An unfortunate limitation of the new cable systems is that they ―T‖ into major east-west cable systems in the Gulf of Aden and the Red Sea, but costs remain high for East African carriers to reach Europe and Asia on those other cable systems.

Mauritius also deserves mention, as the host of AfriNIC and several other African Internet governance institutions. Despite being the corporate home of the SEACOM East African cable system, Mauritius has been bypassed by this and the other two new cable systems, leaving its competitive situation unimproved, since its domestic incumbent was sold, in much the same way South Africa‘s was, to a foreign carrier that maintains monopoly hold on both of the prior cable systems that land there, SAT-3 and SAFE.

Latin America

The growth of the Internet market in Latin America in the past five years has been dramatic, led by the success of a long-term programme of new exchange point development by ―Comitê Gestor da Internet no Brasil" (CGI), the Brazilian Internet Steering Committee, a public-private partnership funded in large measure by revenue from domain name registrations within the .BR country-code top-level domain. Between 2006 and 2011, Brazil has grown from four IXPs to nineteen, maintaining their position of leadership in the region; in addition to having the region‘s largest exchange, in São Paulo, Brazil has hosted more than half of Latin America‘s IXPs for most of the history of the Internet‘s expansion in the region. Brazil‘s success has been a product of several factors coinciding: the CGI takes the long view, with a consistent programme of economic development, rather than short-term one-off projects. Their IXPs are among only a handful in the world that are the product of a considered and intentional economic model. Before beginning, they made a careful investigation of CityLink, the New Zealand IXP system, and of the SeattleIX, including site visits, observation of the annual governance meeting, and interviews with the founders, board of directors, and Internet exchange participants. This feedback loop has continued to the present day, with CGI management staff and board members actively investigating the successes and failure of other IXPs and participating in the international IXP operations community.

Argentina‘s Cámara Argentina de Internet (CABASE), the Argentine Internet services industry association, an entirely private-sector organization, founded the first IXP in Latin America, in Buenos Aires in 1998, and has recently been inspired to further growth by the example of its Brazilian neighbuor, with three new IXPs founded in Neuquén, Rosario, and Bahía Blancain 2010 and four more planned for this year.

Mexico continues to lag, being by far the largest nation in the world, and the only OECD member country, to continue without any domestic Internet exchange capacity. Mexico‘s traffic continues to be exchanged largely on the east coast of the United States, and to a lesser degree in the exchanges of its Latin American neighbours and European trading partners. This is a consequence of the near-universal dominance of Mexico‘s incumbent, Telmex. This situation may be on the brink of reform, as COFETEL, the Mexican regulator, has opened access to competitive long-haul circuits, has licensed a second national carrier, and is investigating the establishment of an IXP. The lack of domestic traffic exchange has had a dramatically-visible effect on Mexican transit pricing, relative to other economies of similar size and development:

Figure 6: Retail price of enterprise Internet transit, mbps/month, in USD

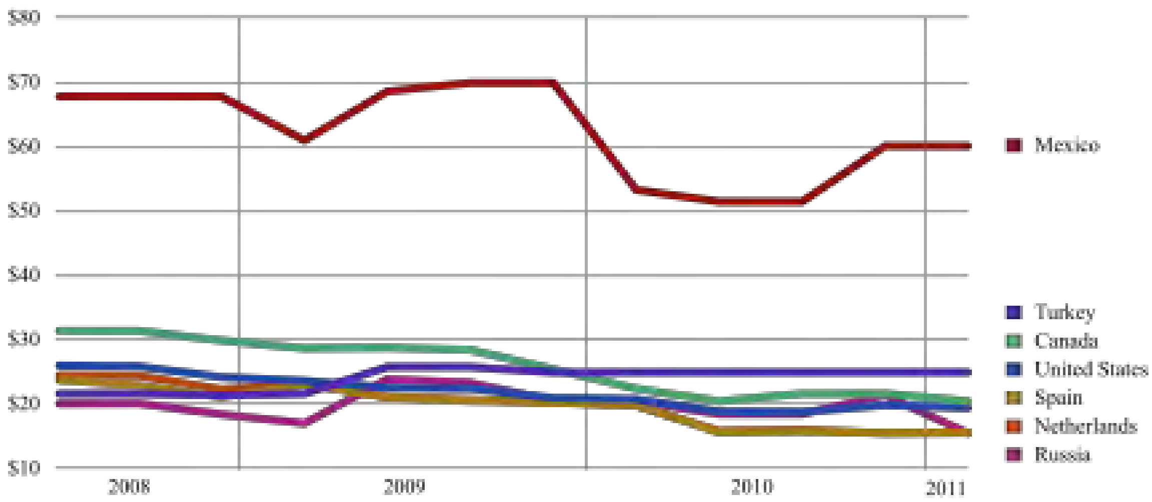

Source: Telegeography

The relatively few Mesoamerican attempts at local traffic exchange have met with mixed success thus far. Panama, Nicaragua, and El Salvador have formed IXPs that are either no longer active or are not growing, and Belize, Guatemala, Honduras, and Costa Rica have not yet attempted to form exchanges.

Columbia and Peru each have one functioning IXP, and Chile has three, but these countries have been hampered by oligopolistic behaviour on the parts of their ISPs. Generally, there are high fees or barriers-to-entry set at the IXPs, or the IXPs are not sufficiently neutral to attract wide participation.

Belize is unique in the hemisphere in that peering there is not only obstructed but actively prohibited by the combined actions of the government and the dominant ISP. Only the incumbent is allowed to offer Internet services. VoIP is illegal and actively blocked. The arrival of satellite Internet has begun to change the landscape, and both Direct TV and Starband offer satellite Internet access. Though an improvement, this is hardly conducive to the kind of local applications that promote tangible economic development.

Speaking very generally, the Latin American region is characterised by a relatively small number of relatively large ISPs, many of which are national incumbents that have made the transition to multinational regional carriers. There is a relatively low rate of formation of new, small ISPs within the region, and this lack of new market entrants hampers the growth rate of the industry overall relative to other regions. While

South America had relatively few IXPs and dominant incumbents fought intense peering battles, international connectivity, principally to Miami, seemed scarce. Now that South American IXPs are burgeoning and peering is more common, the fibre that was being used to trombone domestic traffic through Miami has been freed up for actual international traffic, and continuous construction of new fibre systems and upgrades of existing ones have made international connectivity from South America to North America and Europe relatively plentiful.

The Caribbean

The pace of growth in the Caribbean has increased sharply over the past five years as these island nations liberalised their markets, extracting themselves from long-held monopolies. The creation of a competitive undersea carrier in Columbus Communications, which acquired the ARCOS-1 ring in 2006, and the formation of IXPs in Puerto Rico, Haiti, Curaçao, St Maarten, Grenada, the Dominican Republic, and the British Virgin Islands have been major drivers in this increased growth, as have the efforts of the reinvigorated Caribbean Telecommunications Union, which has been conducting a continuous travelling ―ICT Roadshow" for the past two years, promoting telecommunications self-sufficiency and local production in each of the countries of the region. Initiatives such as these have aided governments and regulators in understanding the nexus of routing economics and regulatory policy. It has also resulted in governments taking a more informed and active role in development of the industry.

The region has suffered several major natural disasters during the past five years, which have hampered domestic investment in telecommunications infrastructure, but in general the Caribbean is making efficient use of scarce resources to improve a market that faces significant hurdles due to low population density, low per-capita income, and the high expense of having to lay undersea, rather than terrestrial, long-haul cable. It does have the natural advantage of lying on the path between the two much larger markets of South America and the east coast of the United States. Efforts to form IXPs have begun in most countries in the region, and many of these are likely to come to fruition in the next two years, which will reduce tromboning and free up existing capacity on the undersea fibre systems for more efficient use for purely international carriage.

The Middle East

The Middle East is the region that lags farthest behind in regional and national Internet production capacity, smaller than the Caribbean, which it far exceeds in other measures of economic size. Cairo was, for many years, the sole exception, since 2001 hosting the only IXP of any Middle Eastern or Arab League country. In 2007, Lebanon joined the ranks with the BeirutIX, which led to an unprecedented 425% growth in Internet traffic over the course of the following year and an additional 300% growth the year after, as pent-up demand was fulfilled by ISPs that were able to provide vastly more service at a lower cost. Bahrain has also had an exchange since 2009, housed in a state-of-the-art datacenter facility, but it carries unfortunately little traffic; in addition, a transit co-operative named the ―Bahrain Internet Exchange" has more recently made the transition to also providing actual IXP services. Saudi Arabia has had a government-led effort to turn up IXPs in three cities, which proceeded as far as a successful trial in Riyadh in 2010, but no further; no actual production traffic followed the trial. No other Middle Eastern countries have yet succeeded in setting up IXPs, though the UAE, Qatar, Jordan, Iraq, and Iran have each had public discussions of the possibility.

Within the region, Iran is notable for having made an announcement in April 2011 of a ‗halal' network aimed at Muslims on an ethical and moral level" and to ―increase Iran and the Farsi language's presence" on the Internet.[69] This has drawn considerable attention and some criticism, due to the perception that what was being announced was a new regime of domestic censorship and an attempt to export that to neighbouring countries. A closer reading of the announcement, however—along with the much less heralded announcement of the Europe Persia Express Gateway, a joint investment between Iran, Oman, Cable & Wireless, and Rostelecom, to build a new cable system connecting Tehran to the world‘s largest Internet exchange in Frankfurt—indicates that Iran‘s actual project has much more to do with increasing the operational and economic efficiency and profitability of its network and bringing it into closer conformity with international norms of best-practice.[70] Making it possible for content produced or stored domestically to compete effectively with content stored elsewhere is part and parcel of good practice in making a more efficient network.

Europe

Europe as a region is far more cognizant of the value of Internet exchange points than other regions are, and of the value of attracting new participants to its exchanges to promote economic growth. This has led to a degree of competition between European nations, each desiring to have the largest IXP, that‘s not really visible in other regions. While undoubtedly beneficial, it has also arguably led to a structuring of IXPs in European countries to maximise their visible or apparent size, potentially at the cost of their efficiency, and particularly by comparison with those in neighbouring countries with which they are competing for ISP infrastructural investment spending. It has also led to some measure of contention over instrumentation and the accuracy of reporting of some IXP growth metrics. Thus, while the relative sizes of European exchanges are discussed in this report, it is worthy of note that the exact figures change frequently, and are often the subject of debate.

Figure 7. Progression of largest European exchange toward the center of Europe population density

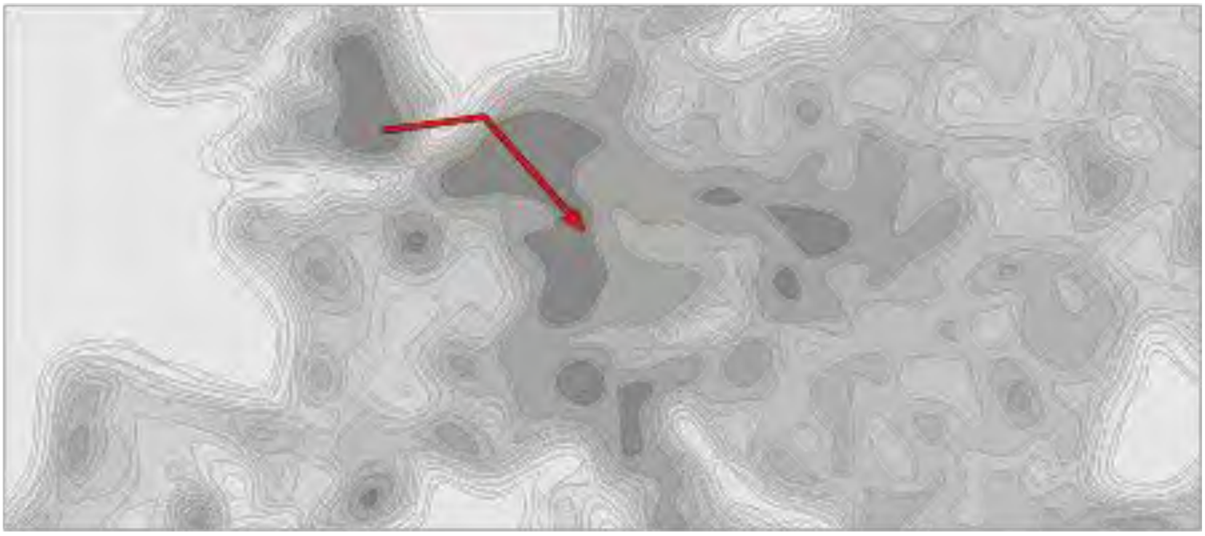

Source: Packet Clearing House.

A trend that has been playing out over the past twenty years in Europe is the progression of the title of ―Europe‘s largest exchange" gradually toward the centre of European population density (Figure 7). To those who work with exchanges, there is a certain inevitability to this, though usually it occurs on a much smaller scale. The average distance between an exchange and all of the members of the population it serves is a major measure of the efficiency, and thus the long-term success, of any exchange. Therefore, all else aside, exchanges nearest to centres of population density are the most successful, in that the bandwidth they generate can be carried to their constituents at the lowest possible cost while retaining the highest possible performance. Speed times distance equals cost. Thus Amsterdam gradually overtook London, and Frankfurt gradually overtook Amsterdam, in bandwidth produced, in part because each was, in turn, proximal to a larger population. Stated another way, the number of people who could reach Amsterdam at low cost was greater than the number who could reach London at that same cost, and the number who could reach Frankfurt at that cost was lower still. In recent years, the larger exchanges of Eastern Europe, in Prague and Sofia, have been garnering a significant share of the European market, although it seems very unlikely that either would overtake Frankfurt in the foreseeable future.

In general, though, Eastern Europe has been moving up in many different rankings of Internet growth. The Moscow and Kiev exchanges are now the fourth and fifth largest in the world, by bandwidth, and Russia and the Ukraine are now the second and fifth largest countries by number of international network adjacencies, with annualized growth rates of 40% and 44%, both considerably higher than the other countries of the top ten (Table 1, and see also Appendixes 1 and 3).

Table 1. Top ten countries by network adjacencies, showing high growth rates in Eastern Europe

	September 2009	September 2010	Net Change	Percent Change
United States	63230	72871	+9,641	+15%
Russia	11894	16633	+4,739	+40%
United Kingdom	7263	8474	+1,211	+17%
Canada	4179	5327	+1,148	+27%
Ukraine	3662	5272	+1,610	+44%
Germany	4851	5191	+340	+7%
Poland	3884	4646	+762	+20%
Australia	2342	3101	+759	+32%
Japan	2843	3092	+249	+9%
Korea	2795	2906	+111	+4%

Source: Packet Clearing House.

Europe has more IXPs than any other region, and seven of the top ten exchanges both by number of participants and volume of traffic. If any grounds exist to criticise the performance of European IXPs, they tend to be more expensive to participate in than IXPs in other regions, since they tend to have employees and budgets for marketing and public relations and policy process participation, and their negotiation power with colocation facilities tends to be relatively weak by comparison with their counterparts in other regions, which leads to additional expenses and being spread thinly across many datacenters. This latter is a uniquely European trait, and to the degree that European IXPs can surmount it they will benefit greatly, since their growth will no longer be bound by the cost and capacity limitations of the inter-switch trunks between colocation facilities.

Canada and the United States

Although Canada‘s Toronto IXP is a success relative to the size of the served community, most of Canada‘s population is located on the more temperate climate of their southern border with the United States, and this has resulted in the vast majority of their traffic being sourced in, and accounting for a portion of the growth of, the larger northern United States exchanges in Seattle, Chicago, and New York at the expense of Canada‘s own smaller exchange in Ottawa. There have been other attempts to create exchanges in most other major cities in Canada, but these have failed to reach fruition. Recently CIRA, the registry for Canada‘s .CA country-code top-level domain, has been discussing funding a more co-ordinated effort to bring IXPs to more Canadian cities, after the model of their Brazilian counterpart. If Canada pursues this effort carefully, it would likely succeed.

The United States has had slow but steady growth over the past five years, going from 74 IXPs producing 118 gigabits per second of observable bandwidth in 2006 to 85 IXPs producing 826 gigabits per second today, 47.6% annualized growth over the period.[71] The caveat ―observable" is relevant because IXPs in the United States utilise crossconnects much more heavily than those in Europe, offloading exchanged traffic from the Ethernet switches at the cores of the exchanges. When traffic takes this ―express lane" bypassing the central switch, it is no longer observable in the statistics. Although it would be nice to have directly comparable statistics between regions, the technical and economic advantages of using crossconnects wherever feasible far outweigh the desire for statistics. This contrast between European and American IXPs plays out in several other specifics as well: the relatively inward-looking private-sector exchanges of the United States generally do not view their traffic statistics as a measure of competitive success versus exchanges in other countries or regions, allowing them to focus much more directly on the primary goal of creating valuable bandwidth at the lowest possible cost. This focus on a single goal has also given them the resolve to negotiate more critically with building (and more particularly colocation facility) owners, so switch fabric extension policies in the United States tend to mirror the planned policy in Brazil much more closely than the happenstance situation of most of the larger European exchanges.

The backbone carriers that connect to these exchanges and sell wholesale transit in the United States and Canada continue to do relatively well in a relatively competitive marketplace. Wholesales prices for high-volume transit have remained between about USD 1.40 and USD 3 per megabit per second per month range for quite some time, but there are not substantially better prices to be found in Western Europe, which is another good indicator that the actual size of United States exchanges, if one were to include crossconnect traffic, is as large as those of London, Amsterdam, and Frankfurt. The lack of significant wholesale price drops over the past five years is largely due to the continued lack of standardised network components at speeds faster than 10 gbps, as discussed elsewhere in this report. Mergers and acquisitions among these large backbone carriers have been closely scrutinised by regulatory authorities but generally allowed to proceed. Tests of that regulatory permissiveness in 2011 were Level 3 and Global Crossing, by many measures the two largest backbone carriers globally, and Qwest and Saavis, two medium-sized global carriers. Both announced planned mergers considered in 2011. Even discounting customers who currently purchase from both Level 3 and Global Crossing and would likely reduce that to a single connection to the merged entity, the union of the two would still be more than twice the size of NTT, its next nearest competitor.[72]

The Pacific

The Pacific islands combine many of the greatest challenges in Internet deployment: small populations, low income, and vast intervening distances across deep seabeds. Hawaii and Guam are on the path of the Japan-US and TPC-5 cables, among others, which have given them high-speed but expensive connectivity for many years. The commissioning of the Southern Cross cable system in 2001 brought fibre to Fiji, and Samoa and American Samoa got their first fibre data service in 2010, largely due to the efforts of their privately funded competitive carrier, Blue Sky Communications. In the Pacific, only Hawaii and Guam (excluding Australia and New Zealand, addressed below) have Internet exchange points, though there have been ongoing discussions in Fiji and Samoa for some years. Fiji is notable for having formed a public-private cybersecurity working group to discuss and publish cybersecurity practices and recommendations to the Fijian network community. Along with the formation of a Pacific CERT, hosted by the University of the South Pacific and assisted by AusCERT and JPCERT, this marks the beginning of an awareness of cybersecurity and cybercrime legal issues in a region that is not an attractive target for cybercrime but has been used as a staging area for some kinds of attacks, particularly those requiring a cooperative or unwitting host ccTLD registry.

Australia and New Zealand

Australia has eleven IXPs, the first two established in Perth and Melbourne in 1997 and 1998, followed by others on the east coast between 2001 and 2004. The large number reflects the great distances across Australia‘s interior and the consequent cost of crossing them with telecommunications signals more than it does the relatively sparse population. Historically, one factor that set Australia apart, in the exchange of Internet traffic, was a decision by the regulator in 1997 to require the largest ISP to peer with the next three largest ISPs (considered elsewhere in this report). This may have prevented the earlier development of competitive infrastructure by small players since it was some time before additional competitive carriers, such as Pipe Networks and Vocus Communications, emerged and further increased competition and market access for smaller players.

The exchange in Perth is of particular note because it has had, since its establishment, a *mandatory* multilateral peering agreement; that is, every participant in the IXP is obligated to peer with each and every other participant. This differs from the majority of multilateral agreements, which are available but not mandatory. This relatively extreme position has strongly influenced the composition of the membership of the Perth exchange, making it at once more populous and more populist, drawing many more small and non-ISP participants while giving Telstra and Optus reason not to participate. Though this tradeoff is of mixed benefit to the Perth exchange, it is one they have been happy to continue for the past fifteen years, and it serves as a didactic real-world example of an uncommon policy for the rest of the world.

Australia is today building a government-owned national broadband network (NBN) that aims to provide fibre to 93% of the nation‘s households. The remaining population will be covered by wireless and satellite services. The layer 2 fibre network provides several points for interconnection. ISPs will sell retail services but still rely on their own network resources, or those of other providers, for backbone and international connectivity to the rest of the global Internet. They will also continue to exchange traffic among themselves at IXPs or through private interconnection. A contentious element of Australian developments has been the decision taken by the regulator on the number of points for interconnection. This decision fell to the independent regulator, as opposed to being taken by market participants as in other countries, due to the fact that the NBN will hold a near-total monopoly on fixed network access. The competition regulator was asked by the Australian Government to provide advice on the initial number and location of the points of interconnection considering the long term interests of end-users. In applying this principle, the competition regulator had regard to the promotion of competition, achieving any-to-any connectivity between end-users and encouraging the economically efficient use of and investment in infrastructure. The competition regulator was asked by the Australian Government to provide advice on the initial number and location of the points of interconnection considering the long term interests of end-users. In applying this principle, the competition regulator had regard to the promotion of competition, achieving any-to-any connectivity between end-users and encouraging the economically efficient use of and investment in infrastructure. [73]

New Zealand was one of the first OECD countries to liberalise its telecommunication market, though initially without an industry regulator. Slow progress in the development of competition, and particularly interconnection between networks, led to a change in policy direction at the turn of the century. In 2001, a Telecommunication Commissioner was established to apply industry-specific regulation and improve interconnection arrangements and other procompetitive policies. In the following decade saw the development of a successful private-sector fibre and broadband industry in the country's main population centres. As in Australia, a decision has been taken by the government to fund a fibre network in association with private-sector providers. These fibre providers will tender to provide local access networks that will have virtual monopolies with vertically separated retail competition for Internet access. To bid as a participant in the new access networks, Telecom New Zealand will have to split off its retail services.

As in Australia, questions have arisen over existing private-sector investment. CityLink, the operator of competitive municipal fibre networks in Wellington and Auckland as well as all of New Zealand's seven IXPs, sees its future as very uncertain. At the same time, FX Networks, New Zealand's primary competitive carrier, which has invested significantly in fibre deployment, could be unfairly disadvantaged by competitors receiving benefits from public investment.

South Asia

India is by far the largest of the countries of the South Asian subcontinent in terms of population and the size of its economy. India has, however, lagged in achieving self-sufficiency in bandwidth production. Nepal, Bangladesh, Sri Lanka, and Pakistan all have successful IXPs of some years' standing and have organised the popular South Asian Network Operators Group. India made an unfortunate choice by selecting Enron Broadband Services, and its local subsidiary Enron India, to construct a state-mandated exchange, which Enron turned into an experiment in market speculation and the creation of routing-based investment derivatives.[74] A second government-formed organisation called NIXI, the National Internet Exchanges of India, was then established. After initially granting monopoly status to NIXI and placing regulatory requirements upon ISPs to pay fees, India has liberalized these requirements to some extent. In the meantime, exchanges have still not been constructed commensurate with demand, and NIXI charges growth-inhibiting usage-based fees. Recent estimates by Reliance Industries indicate that NIXI's entire current domestic exchange capacity could support less than 10% of the needs of a single national broadband provider.

Because of the domestic regulatory roadblocks to investment and competition, India's backbone networks have turned to international markets, and they have been quite successful. In January 2004, Reliance Communications purchased the global FLAG undersea cable system; in January 2007, they followed this with the purchase of Yipes, a United States backbone network. TATA responded by buying Tyco, the largest undersea cable construction company, in November 2004, and followed this with the purchase of Teleglobe, Canada's Intelsat signatory and a major global backbone network in February 2006; in June 2008 it acquired Neotel, South Africa's second largest carrier; in January 2010 it acquired British Telecom's cloud services division; and in 2011 it completed its acquisition of BitGravity, a large CDN and cyberdefense service provider. Together, these acquisitions make India one of the world's largest holders of international undersea cable systems and technology. In 2010, Bharti, the world's fifth-largest mobile operator, after failing twice to merge with MTN, acquired Warid in Bangadesh and all of Zain's African operations. It also lit its portion of the Japan–United States Unity cable as well as its portion of the East African EASSy cable. In 2011, it has already lit its share of the new Europe-India Gateway cable.

East Asia

After a long period of leading the world in domestic Internet production, Korea, Japan, and Hong Kong, China have slowed somewhat in the past five years, but they continue to provide an example for other countries in the efficient distribution of bandwidth from IXPs to consumers at high speed and low cost. It is commonly observed by Internet infrastructure specialists that improvement of the Internet depends upon a circular path of improvement of each component of the Internet‘s infrastructure: IXPs, international connectivity, content, backbone networks, and access networks. One circumnavigates this circle endlessly, upgrading each in turn, lest any become a bottleneck that restrains the growth of the others. Overinvestment in large IXPs avails a country little if it lacks competitive access to a fibre local loop across which to deliver the bandwidth, for example. Korea, Japan, and Hong Kong, China have all been exemplary in striking a balance between each of these components, upgrading them all at a comparable pace.

Seven years of Chinese governmental efforts to establish a mainland-style managed exchange in Hong Kong, China have still not found many takers, but the original HKIX, the second exchange in the region and in many ways the first example of an exchange in the modern model, continues to be one of the primary locations for Asian peering. Mainland China, in many ways the opposite of Hong Kong, China remains the world‘s largest population bottlenecked behind a monopoly international gateway. Interestingly, the monolithic negotiating power of this single point of entry, China Telecom, quietly but significantly contributed to the dissolution of the ―Tier-1” cartel.

SingTel, the incumbent telecommunication operator in Singapore, is majority owned by the investment arm of the Singapore government, leading to predictable conflicts-of-interest with the regulator. Participants believe that there were obstacles placed in the path of the Singapore Open Exchange, after it was opened in 2001, and that a similar set of disadvantages were applied to Equinix‘s peering facility after their merger with Pihana in 2002. Some of these disadvantages have been overcome with recent direct access to an international cable landing, but bandwidth production lags regional rival Hong Kong, China. More recently Singapore has funded a national fibre access network, with a virtual monopoly, and a separation between infrastructure and the competitive provision of services. Singapore is viewed as a desirable country for both its centrality to Asian undersea cable systems and its rule of law (see Annex 1, Figure A1.2), and with greater market openness and competition could see significantly higher rates of growth.

Korea‘s experience in the late 1990s, as its many broadband providers began having to bundle large groups of 1 gbps circuits together prior to the arrival of 10 gbps network interfaces, has foreshadowed what the rest of the world has been experiencing more recently with the lack of 100 gbps and 1 tbps interfaces. This drove them away from switched IXPs and toward massive colocation facilities with dense meshes of direct crossconnects, a model the United States soon followed. The consequent reduction in visibility of their network growth statistics, and the fact that Korean network research is less presented at conferences in the rest of the world, has made Korea something of an enigma in networking for some foreign observers. Its domestic network is clearly very large, but with more than 96% domestic traffic at last count it is nearly impossible to judge the size accurately from the outside.

Japan continues to lead the Asian region in most Internet statistics, due to a long head start and a combination of relatively aggressive investment in research and development and a high demand for Internet services in their domestic population. In addition, Japan‘s high concentration of transpacific cables made it the initial ―jumping off” point for United States networks entering Asia, which meant that its domestic IXPs often had more than half international carriers exchanging non-Japanese traffic, in its early years. Since then, the massive Japanese retail broadband market has overtaken international participation in their exchanges. Many Japanese exchanges are hosted in datacenter facilities that are in buildings and served by fibre systems belonging to related companies. The Japanese exchanges have a well-deserved reputation for being among the most reliable and incident-free in the world.

Regional and National Trends

When one examines Internet statistics, nearly all numbers are —up and to the right,” defining growth. The rare exceptions tend to be short-term in nature and are more often indicative of a shift in growth from a well-instrumented protocol or piece of infrastructure to a new and incompletely understood one. It is also helpful to remember that there are ways in which the Internet is intensely geographically specific and dependent—limited by fibre pathways, the speed of light, and the per-cubic-meter cost of excavating trenches—and other ways in which it is entirely free of geographic constraint—enabling entrepreneurs in Nairobi to upstage those in San Jose, and programmers in Kiev to compete on an equal footing with those in Bangalore or Manila. Generally, the degree of geographic dependency tracks with position in the OSI seven-layer protocol model. Physical infrastructure is tied to geography, and applications work equally well everywhere to the degree that they do not require services from the layers below. Location-aware applications, like turn-by-turn driving directions, require information as well as connectivity from layers far down in the protocol stack and are thus much more likely to function as intended in regions with omnipresent and highly available lower-layer network services; applications like solitaire games, which have no external dependencies, function as intended anywhere they have a host and power.

When comparing statistics between countries, it is often useful to look at both the growth in absolute numbers, as in Figure 8, and the annualised growth rates, as in Figure 9, since the latter often give more insight into what the future will bring.

Figure 8. Annualized absolute growth in domestic Internet bandwidth production, top thirty countries, 2005–2010

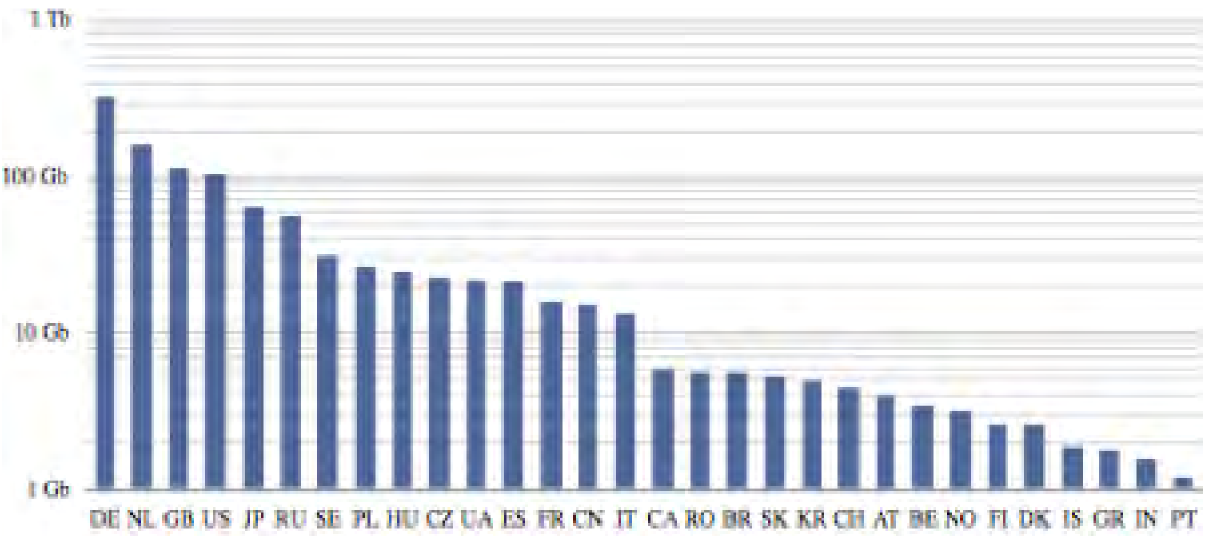

Source: Packet Clearing House.

The visible statistics for Internet bandwidth production in many Western European countries are high relative to the unobservable total, whereas the visible statistics for the United States significantly underrepresent the total. Private-sector investment does compensate for this to some degree by accounting for the availability of crossconnects, but investment and economic growth favour transparency, so in the long run both markets will probably be better served if they find ways of better exposing numbers that more accurately reflect total bandwidth production. In the mean time, the large visible statistics do draw attention and investment to Germany, the Netherlands, and the United Kingdom.

Figure 9. Annualised percentage growth in domestic Internet bandwidth production, grouped by region, 2005–2010

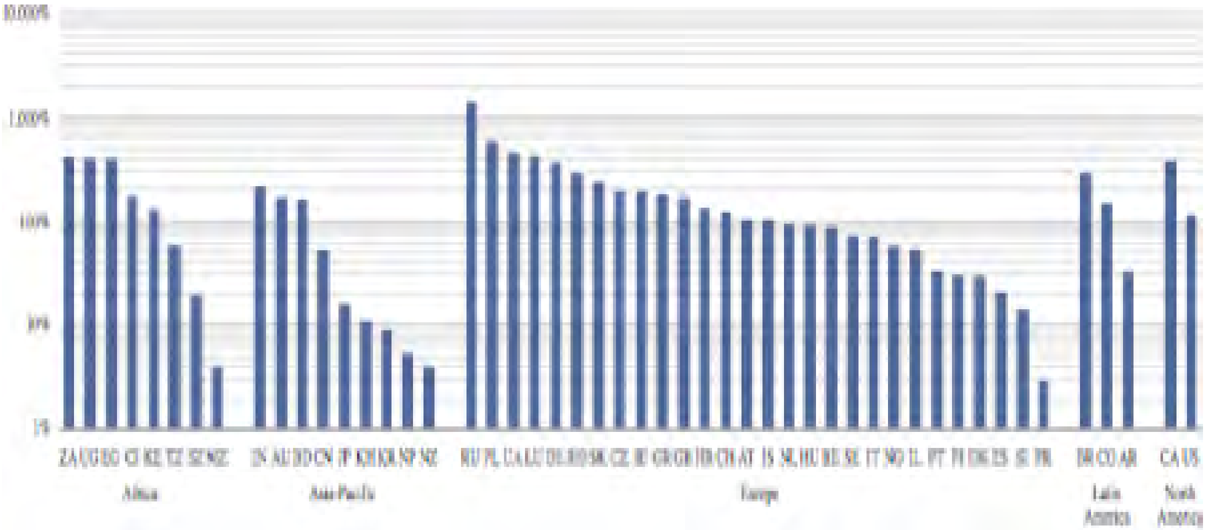

Source: Packet Clearing House.

The five-year annualized growth rates in Figure 9 show that each region embodies a spread between fast-growing countries (more than 200% annualized growth) and slow-growing countries (as little as 20% annualized growth). As we discuss below, a long-term growth rate of 115% has been a reasonable target, but that target is becoming more difficult to reach as the 2001 halt to basic research has resulted in a corresponding stoppage of the flow of new and faster network components. As would be expected over this five-year term, somewhat less than half of the countries for which figures are available meet or exceed this 115% growth rate. The countries of Eastern Europe, led by Russia at an astounding 1 470% five-year annualized growth rate, are the emerging stars of the past five years, though South Africa, Uganda, and Egypt also exceed 400% annualized growth.

Figure 10. Average distance of participants from IXPs within each economy, grouped by region

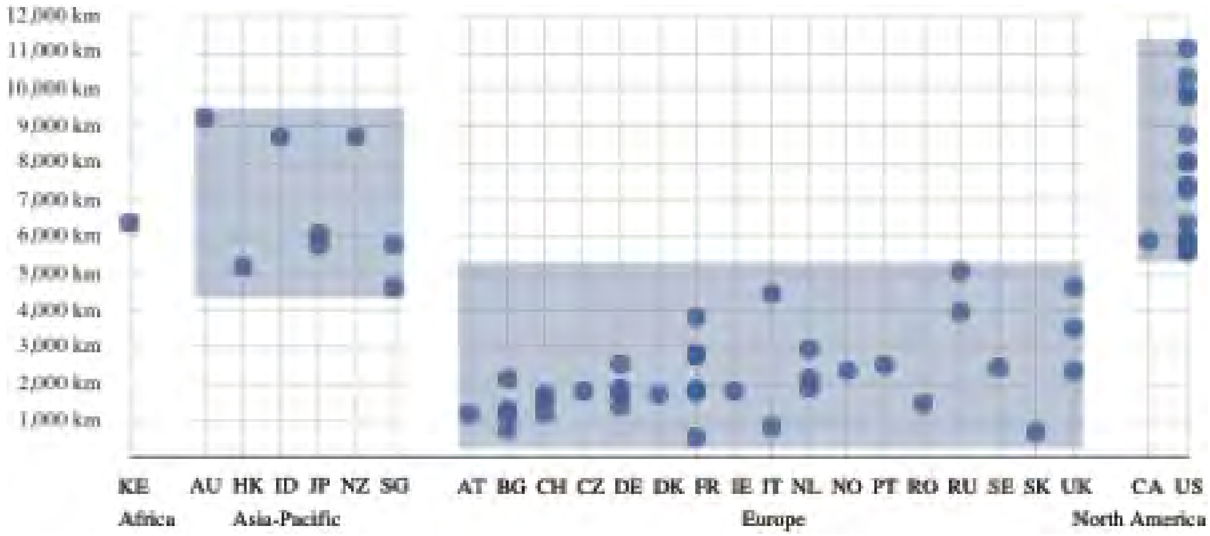

Source: Packet Clearing House

When the average distance between each IXP and the countries-of-incorporation of its member ISPs (Figure 10) is examined, we see distinct regional clusters. Participants in North American IXPs have, on average, extended their networks the greatest distance in order to connect to the IXP, followed by those in the Asia-Pacific region, whereas European IXPs have the smallest portion of long-distance participants. Because only the IXPs for which we have sufficient data to yield a clear and meaningful average are shown, many of the smaller IXPs in all regions are not represented on this chart; only Nairobi is included from Africa, and none at all from Latin America. Small IXPs do not show different trends than large ones, but they are excluded from the graph because of their lower confidence level. Two lessons that can be taken from this graph: if one builds a local IXP, no matter how remote, ISPs from other regions will build infrastructure to reach it; and exchanges (like those in Europe) can become very successful even if they do not attract a substantial portion of intercontinental participants.

APPENDIX 1: NATIONAL INTERNET STATISTICS

	IPv4 Addresses				IPv6 Addresses				Autonomous System Numbers				Network Adjacencies			
Country	Sep 2009	Sep 2010	Net Change	Percent Change	Sep 2009	Sep 2010	Net Change	Percent Change	Sep 2009	Sep 2010	Net Change	Percent Change	Sep 2009	Sep 2010	Net Change	Percent Change
Afghanistan	/16	/16	+/19	+13%					8	16	+8	+100%	7	18	+11	+157%
Albania	/15	/15	+/18	+8.8%					13	16	+3	+23%	30	47	+17	+57%
Algeria	/12	/11	+/13	+41%	/32	/31	+/32	+100%	11	12	+1	+9.1%	31	35	+4	+13%
American Samoa	/20	/20				/32	+/32	new	2	2			4	8	+4	+100%
Andorra	/17	/17			/32	/32			1	1			3	3		
Angola	/16	/15	+/17	+55%	/32	/32			16	20	+4	+25%	35	47	+12	+34%
Anguilla	/20	/20							1	1			1	1		
Antigua and Barbuda	/16	/16	+/22	+1.9%					2	4	+2	+100%	2	3	+1	+50%
Argentina	/9	/9	+/12	+16%	/27	/27	+/32	+2.4%	236	270	+34	+14%	700	776	+76	+11%
Armenia	/15	/13	+/14	+155%	/30	/29	+/30	+75%	36	47	+11	+31%	99	196	+97	+98%
Aruba	/18	/17	+/18	+111%					2	3	+1	+50%	9	9		
Australia	/7	/7	+/10	+15%	/19	/19	+/25	+1.1%	1004	1156	+152	+15%	2342	3101	+759	+32%
Austria	/9	/9	+/12	+12%	/26	/26	+/28	+31%	330	366	+36	+11%	1362	1611	+249	+18%
Azerbaijan	/14	/14	+/16	+28%					21	26	+5	+24%	47	69	+22	+47%
Bahamas	/16	/15	+/18	+22%	/31	/31			4	4			4	4		
Bahrain	/14	/13	+/16	+30%	/31	/30	+/32	+50%	21	24	+3	+14%	38	46	+8	+21%
Bangladesh	/13	/12	+/15	+24%	/29	/28	+/29	+137%	83	99	+16	+19%	282	354	+72	+26%
Barbados	/15	/15	+/18	+15%					6	7	+1	+17%	20	30	+10	+50%
Belarus	/14	/12	+/12	+219%	/32	/30	+/31	+200%	54	71	+17	+31%	104	162	+58	+56%
Belgium	/9	/9	+/13	+7.2%	/27	/24	+/24	+993%	147	159	+12	+8.2%	425	475	+50	+12%
Belize	/16	/16	+/21	+4.1%	/32	/32			5	5			2	4	+2	+100%
Benin	/19	/17	+/18	+233%	/32	/32			2	4	+2	+100%	3	5	+2	+67%
Bermuda	/16	/16	+/20	+5%	/32	/31	+/32	+100%	13	13			30	31	+1	+3.3%
Bhutan	/18	/18			/32	/31	+/32	+100%	4	4			17	21	+4	+24%
Bolivia	/13	/13	+/15	+22%	/30	/30			12	14	+2	+17%	32	34	+2	+6.3%
Bosnia and Herzegovina	/13	/13	+/15	+22%	/31	/29	+/29	+300%	21	24	+3	+14%	66	92	+26	+39%
Botswana	/16	/15	+/18	+23%		/32	+/32	new	7	8	+1	+14%				
Brazil	/7	/7	+/9	+19%	/16	/16			920	1020	+100	+11%	1848	2647	+799	+43%
British Indian Ocean Territory	/20	/20							1	1			4	2	-2	-50%
British Virgin Islands	/21	/19	+/19	+300%					1	1			1	1		
Brunei Darussalam	/14	/14	+/24	+0.13%	/31	/31			4	4			10	12	+2	+20%
Bulgaria	/10	/10	+/14	+7.6%	/28	/28	+/30	+33%	435	480	+45	+10%	1250	1707	+457	+37%
Burkina Faso	/17	/17	+/22	+3.4%					3	4	+1	+33%	7	9	+2	+29%
Burundi	/21	/21														
Cambodia	/15	/14	+/17	+31%		/30	+/30	new	35	40	+5	+14%	122	205	+83	+68%
Cameroon	/16	/15	+/17	+29%	/32	/32			7	9	+2	+29%	21	23	+2	+9.5%
Canada	/6	/6	+/11	+3.5%	/26	/25	+/27	+59%	1295	1393	+98	+7.6%	4179	5327	+1148	+27%
Cape Verde	/19	/18	+/20	+50%												
Cayman Islands	/17	/17	+/20	+13%					6	6			5	12	+7	+140%
Central African Republic	/20	/20	+/22	+31%					1	1						
Chad		/20	+/20	new												

	IPv4 Addresses				IPv6 Addresses				Autonomous System Numbers				Network Adjacencies			
Chile	/10	/10	+/12	+17%	/28	/28	+/29	+50%	121	134	+13	+11%	283	368	+85	+30%
China	/4	/4	+/7	+20%	/26	/23	+/24	+535%	449	467	+18	+4%	862	898	+36	+4.2%
Colombia	/10	/9	+/12	+15%	/28	/28	+/30	+23%	93	98	+5	+5.4%	202	229	+27	+13%
Comoros									1	1			1	2	+1	+100%
Congo-Brazzaville	/22	/19	+/19	+700%					1	2	+1	+100%				
Congo-Kinshasa	/18	/18	+/21	+13%					2	3	+1	+50%	2	2		
Cook Islands	/19	/19			/32	/32			1	1			2	2		
Costa Rica	/11	/11	+/16	+3.8%	/29	/29	+/31	+33%	14	19	+5	+36%	25	38	+13	+52%
Cote D'Ivoire	/15	/15	+/19	+7%	/31	/31			7	8	+1	+14%	11	18	+7	+64%
Croatia	/11	/11	+/15	+6.3%	/30	/29	+/30	+60%	72	78	+6	+8.3%	211	240	+29	+14%
Cuba	/15	/15	+/20	+3.9%	/30	/30			6	6			7	7		
Cyprus	/13	/13	+/15	+26%	/29	/29	+/31	+33%	52	61	+9	+17%	119	165	+46	+39%
Czech Republic	/9	/9	+/12	+14%	/26	/26	+/27	+49%	422	671	+249	+59%	916	1443	+527	+58%
Denmark	/9	/9	+/13	+7%	/27	/27	+/28	+80%	186	211	+25	+13%	500	510	+10	+2%
Djibouti	/18	/18			/32	/32			1	1			2	4	+2	+100%
Dominica									2	2			9	11	+2	+22%
Dominican Republic	/13	/13	+/15	+24%	/30	/30			9	12	+3	+33%	24	23	-1	-4.20%
Ecuador	/12	/12	+/14	+28%	/29	/29	+/30	+50%	39	43	+4	+10%	136	157	+21	+15%
Egypt	/10	/10	+/12	+31%	/29	/29	+/32	+17%	51	60	+9	+18%	185	232	+47	+25%
El Salvador	/13	/13	+/17	+8%	/32	/32			16	16			42	25	-17	-40%
Equatorial Guinea	/21	/21								1	+1	new				
Eritrea	/20	/20							1	1						
Estonia	/12	/12	+/16	+9.6%	/29	/29	+/32	+10%	30	36	+6	+20%	96	101	+5	+5.2%
Ethiopia	/18	/18							1	1			4	6	+2	+50%
Faroe Islands	/17	/17	+/20	+12%	/32	/32			5	5			8	14	+6	+75%
Fiji	/15	/15	+/18	+15%	/30	/30	+/31	+67%	6	8	+2	+33%	5	8	+3	+60%
Finland	/9	/9	+/14	+3%	/27	/27	+/28	+87%	141	158	+17	+12%	390	486	+96	+25%
France	/6	/6	+/10	+4.9%	/19	/19	+/26	+0.6%	594	674	+80	+13%	1926	2192	+266	+14%
French Guiana	/21	/21							2	2				1	+1	new
French Polynesia	/17	/17				/32	+/32	new	1	1			4	3	-1	-25%
Gabon	/15	/15	+/18	+8%					2	4	+2	+100%	4	4		
Gambia	/18	/18							1	1			1	2	+1	+100%
Georgia	/13	/12	+/15	+17%		/31	+/31	new	29	31	+2	+6.9%	88	156	+68	+77%
Germany	/6	/6	+/9	+7.4%	/19	/19	+/23	+6.1%	1221	1335	+114	+9.3%	4851	5191	+340	+7%
Ghana	/15	/14	+/16	+35%	/31	/30	+/32	+50%	20	29	+9	+45%	53	89	+36	+68%
Gibraltar	/16	/16	+/19	+17%	/32	/32			10	11	+1	+10%	33	42	+9	+27%
Greece	/10	/10	+/13	+14%	/29	/28	+/29	+78%	140	150	+10	+7.1%	297	332	+35	+12%
Greenland	/18	/18				/32	+/32	new	1	1			2	2		
Grenada	/21	/20	+/21	+89%		/32	+/32	new	2	2			1	1		
Guadeloupe	/20	/20				/32	+/32	new	2	2			8	8		
Guam	/15	/15	+/17	+17%		/30	+/30	new	5	6	+1	+20%	32	24	-8	-25%
Guatemala	/13	/13	+/17	+5.9%	/30	/30			24	27	+3	+13%	89	90	+1	+1.1%
Guernsey	/20	/19	+/21	+50%					1	1			3	2	-1	-33%
Guinea	/22	/22							1	1			1	1		
Guinea-Bissau	/22	/22														
Guyana	/17	/17	+/20	+13%	/32	/32			1	2	+1	+100%	5	5		
Haiti	/16	/15	+/16	+101%	/32	/31	+/32	+100%	8	10	+2	+25%	8	13	+5	+63%
Honduras	/15	/15	+/18	+22%	/48	/48			20	23	+3	+15%	35	37	+2	+5.7%
Hong Kong	/9	/9	+/12	+15%	/27	/26	+/27	+100%	281	328	+47	+17%	1656	2183	+527	+32%
Hungary	/10	/10	+/14	+8.3%	/28	/27	+/29	+30%	188	198	+10	+5.3%	472	602	+130	+28%
Iceland	/13	/13	+/18	+2.6%	/29	/29	+/32	+12%	29	36	+7	+24%	104	125	+21	+20%
India	/8	/7	+/9	+45%	/27	/26	+/27	+154%	333	417	+84	+25%	1352	1905	+553	+41%
Indonesia	/9	/9	+/11	+22%	/27	/26	+/27	+137%	411	470	+59	+14%	1104	2133	+1029	+93%
Iran	/11	/10	+/11	+81%	/28	/27	+/29	+42%	102	148	+46	+45%	271	479	+208	+77%

	IPv4 Addresses				IPv6 Addresses			
Iraq	/16	/15	+/16	+66%		/31	+/31	new
Ireland	/10	/10	+/14	+7.2%	/27	/27	+/30	+19%
Isle of Man	/17	/17	+/21	+9.2%	/32	/32	+/48	+0%
Israel[75]	/10	/10	+/13	+16%	/29	/29	+/31	+29%
Italy	/7	/7	+/11	+8.2%	/20	/20	+/27	+0.55%
Jamaica	/15	/15	+/17	+19%	/32	/32		
Japan	/5	/5	+/9	+5.5%	/19	/19	+/21	+27%
Jersey	/17	/17			/32	/32		
Jordan	/13	/13	+/16	+18%	/31	/29	+/30	+250%
Kazakhstan	/12	/11	+/13	+49%	/32	/31	+/32	+100%
Kenya	/13	/12	+/12	+202%	/30	/29	+/30	+100%
Kiribati	/20	/20						
Korea	/6	/5	+/7	+33%	/20	/20	+/29	+0.12%
Kuwait	/12	/12	+/14	+25%	/32	/30	+/31	+200%
Kyrgyzstan	/15	/15	+/19	+6.3%	/32	/32		
Laos	/17	/16	+/19	+21%		/30	+/30	new
Latvia	/12	/12	+/15	+6.9%	/30	/28	+/29	+140%
Lebanon	/14	/13	+/15	+33%	/30	/29	+/32	+20%
Lesotho	/18	/18				/32	+/32	new
Liberia	/22	/19	+/19	+800%				
Libya	/14	/14	+/20	+1.4%				
Liechtenstein	/16	/16	+/21	+3.8%	/31	/31		
Lithuania	/11	/11	+/17	+1.2%	/30	/29	+/30	+125%
Luxembourg	/13	/13	+/17	+5.8%	/28	/28	+/30	+23%
Macao	/14	/14	+/16	+23%	/31	/31		
FYR Macedonia	/13	/13	+/20	+0.92%	/30	/30		
Madagascar	/17	/16	+/18	+42%	/32	/32		
Malawi	/18	/16	+/17	+319%				
Malaysia	/10	/10	+/12	+26%	/27	/27	+/28	+48%
Maldives	/17	/16	+/18	+44%		/32	+/32	new
Mali	/17	/17	+/18	+50%	/32	/32		
Malta	/13	/13	+/15	+24%	/30	/29	+/32	+20%
Mariana Islands	/18	/18						
Marshall Islands	/21	/21				/32	+/32	new
Mauritania	/17	/17						
Mauritius	/14	/13	+/14	+100%	/30	/29	+/30	+75%
Mexico	/8	/7	+/10	+23%	/28	/27	+/28	+143%
Micronesia	/21	/19	+/20	+200%	/32	/32		
Moldova	/13	/13	+/15	+39%	/30	/29	+/30	+75%
Monaco	/16	/16			/32	/32		
Mongolia	/15	/14	+/16	+50%		/30	+/30	new
Montenegro	/15	/15	+/17	+45%				
Montserrat	/22	/22						
Morocco	/12	/12	+/14	+27%	/31	/31		
Mozambique	/15	/15	+/19	+5.1%	/32	/32		
Myanmar	/18	/18				/32	+/32	new
Namibia	/15	/15	+/18	+11%	/48	/48		
Nauru	/19	/19				/32	+/32	new
Nepal	/15	/14	+/15	+121%	/31	/29	+/29	+300%
Netherlands	/8	/7	+/12	+5.3%	/23	/23	+/26	+9.4%
Netherlands Antilles	/14	/14	+/18	+5%	/30	/29	+/30	+100%
New Caledonia	/16	/15	+/17	+56%	/32	/30	+/30	+400%
New Zealand	/9	/9	+/13	+12%	/27	/26	+/27	+82%
Nicaragua	/14	/14	+/18	+9.6%	/32	/32		

	Autonomous System Numbers				Network Adjacencies			
Iraq	5	11	+6	+120%	9	24	+15	+167%
Ireland	107	114	+7	+6.5%	289	333	+44	+15%
Isle of Man	16	17	+1	+6.3%	119	123	+4	+3.4%
Israel[75]	229	250	+21	+9.2%	738	929	+191	+26%
Italy	565	617	+52	+9.2%	1910	2306	+396	+21%
Jamaica	9	9			27	38	+11	+41%
Japan	762	795	+33	+4.3%	2843	3092	+249	+8.8%
Jersey	4	4			22	23	+1	+4.5%
Jordan	24	26	+2	+8.3%	64	77	+13	+20%
Kazakhstan	56	67	+11	+20%	222	353	+131	+59%
Kenya	35	49	+14	+40%	83	171	+88	+106%
Kiribati								
Korea	862	916	+54	+6.3%	2795	2906	+111	+4%
Kuwait	37	41	+4	+11%	137	149	+12	+8.8%
Kyrgyzstan	13	14	+1	+7.7%	33	62	+29	+88%
Laos	6	9	+3	+50%	8	12	+4	+50%
Latvia	198	219	+21	+11%	494	554	+60	+12%
Lebanon	39	45	+6	+15%	108	150	+42	+39%
Lesotho	4	4			6	5	-1	-17%
Liberia	1	2	+1	+100%	1	2	+1	+100%
Libya	1	1			5	3	-2	-40%
Liechtenstein	7	8	+1	+14%	36	44	+8	+22%
Lithuania	108	118	+10	+9.3%	295	314	+19	+6.4%
Luxembourg	30	35	+5	+17%	109	142	+33	+30%
Macao	4	4			9	9		
FYR Macedonia	23	24	+1	+4.3%	100	109	+9	+9%
Madagascar	3	3			4	6	+2	+50%
Malawi	4	7	+3	+75%	3	9	+6	+200%
Malaysia	89	107	+18	+20%	316	421	+105	+33%
Maldives	3	3			4	6	+2	+50%
Mali	4	4			5	8	+3	+60%
Malta	22	25	+3	+14%	64	94	+30	+47%
Mariana Islands	1	1			5	3	-2	-40%
Marshall Islands	1	1			5	2	-3	-60%
Mauritania	1	1			4	2	-2	-50%
Mauritius	13	16	+3	+23%	7	11	+4	+57%
Mexico	238	253	+15	+6.3%	663	740	+77	+12%
Micronesia	1	3	+2	+200%	6	5	-1	-17%
Moldova	28	37	+9	+32%	106	155	+49	+46%
Monaco	1	1			5	3	-2	-40%
Mongolia	27	34	+7	+26%	56	108	+52	+93%
Montenegro	9	11	+2	+22%	26	32	+6	+23%
Montserrat								
Morocco	7	7			20	19	-1	-5%
Mozambique	14	16	+2	+14%	14	75	+61	+436%
Myanmar	2	2			3	7	+4	+133%
Namibia	9	11	+2	+22%	34	44	+10	+29%
Nauru		2	+2	new				
Nepal	28	32	+4	+14%	73	137	+64	+88%
Netherlands	444	502	+58	+13%	2299	2530	+231	+10%
Netherlands Antilles	23	26	+3	+13%	41	47	+6	+15%
New Caledonia	4	4			6	6		
New Zealand	268	306	+38	+14%	607	1096	+489	+81%
Nicaragua	14	16	+2	+14%	36	23	-13	-36%

	IPv4 Addresses				IPv6 Addresses				Autonomous System Numbers				Network Adjacencies			
Niger	/18	/18	+/22	+5.9%					1	3	+2	+200%		2	+2	new
Nigeria	/13	/13	+/14	+55%	/31	/30	+/30	+150%	49	68	+19	+39%	182	234	+52	+29%
Niue	/22	/21	+/22	+100%		/48	+/48	new		1	+1	new				
Norfolk Island	/21	/21							2	2			1	1		
North Korea		/22	+/22	new						1	+1	new				
Norway	/9	/9	+/12	+10%	/24	/24	+/27	+12%	145	163	+18	+12%	347	373	+26	+7.5%
Oman	/14	/14	+/18	+7.1%	/31	/30	+/32	+50%	3	3			1	4	+3	+300%
Pakistan	/11	/11	+/12	+69%	/29	/28	+/29	+100%	50	63	+13	+26%	188	215	+27	+14%
Palau	/20	/20			/32	/32			1	1			1	3	+2	+200%
Palestine	/14	/13	+/16	+19%		/32	+/32	new	20	27	+7	+35%	37	50	+13	+35%
Panama	/12	/12	+/15	+8%	/30	/30			78	86	+8	+10%	189	214	+25	+13%
Papua New Guinea	/17	/17	+/20	+14%	/31	/30	+/32	+50%	3	6	+3	+100%	2	5	+3	+150%
Paraguay	/15	/14	+/15	+68%	/31	/31			10	13	+3	+30%	28	48	+20	+71%
Peru	/11	/11	+/13	+22%	/29	/29	+/32	+14%	19	21	+2	+11%	32	46	+14	+44%
Philippines	/10	/10	+/12	+19%	/28	/27	+/28	+100%	196	224	+28	+14%	543	727	+184	+34%
Poland	/8	/8	+/11	+13%	/21	/21	+/27	+1.4%	1022	1236	+214	+21%	3884	4646	+762	+20%
Portugal	/10	/10	+/14	+7%	/28	/28	+/30	+22%	62	65	+3	+4.8%	205	249	+44	+21%
Puerto Rico	/12	/12	+/16	+9.9%	/32	/29	+/30	+500%	43	51	+8	+19%	139	133	-6	-4.30%
Qatar	/13	/13	+/15	+35%	/32	/32			6	6			12	12		
Romania	/9	/9	+/13	+8%	/29	/27	+/28	+191%	578	630	+52	+9%	1587	2401	+814	+51%
Reunion	/17	/17							1	1			2	2		
Russia	/7	/7	+/9	+24%	/25	/25	+/26	+98%	2849	3300	+451	+16%	11894	16633	+4739	+40%
Rwanda	/15	/15	+/20	+2.8%	/32	/30	+/31	+200%	5	7	+2	+40%	9	24	+15	+167%
Saint Kitts and Nevis	/17	/17	+/22	+4%	/32	/32			2	2			5	6	+1	+20%
Saint Lucia	/24	/24														
Saint Martin	/19	/19							1	1			1	3	+2	+200%
Saint Pierre and Miquelon	/20	/20							1	1				1	+1	new
Saint Vincent and the Grenadines	/19	/18	+/20	+40%					1	1			2	2		
Samoa	/18	/18	+/21	+15%	/32	/31	+/32	+100%	3	4	+1	+33%	5	6	+1	+20%
San Marino	/18	/17	+/20	+18%					3	4	+1	+33%	5	4	-1	-20%
Saudi Arabia	/10	/10	+/13	+17%	/29	/28	+/29	+55%	81	90	+9	+11%	229	331	+102	+45%
Senegal	/16	/14	+/14	+213%	/32	/32				1	+1	new				
Serbia	/12	/11	+/13	+53%	/29	/29	+/32	+11%	83	100	+17	+20%	306	396	+90	+29%
Seychelles	/18	/17	+/19	+47%	/32	/31	+/32	+100%	4	4			1	1		
Sierra Leone	/18	/18	+/21	+13%	/48	/48			6	8	+2	+33%	10	8	-2	-20%
Singapore	/10	/10	+/13	+11%	/27	/27	+/28	+56%	162	192	+30	+19%	651	869	+218	+33%
Slovakia	/11	/11	+/14	+16%	/28	/28	+/32	+6.7%	77	84	+7	+9.1%	220	283	+63	+29%
Slovenia	/12	/11	+/14	+23%	/28	/27	+/29	+50%	174	204	+30	+17%	557	647	+90	+16%
Solomon Islands	/19	/19			/32	/32				2	+2	new		3	+3	new
Somalia		/22	+/22	new												
South Africa	/8	/8	+/11	+12%	/28	/27	+/29	+67%	162	194	+32	+20%	401	939	+538	+134%
Spain	/8	/7	+/11	+8.9%	/27	/26	+/27	+124%	325	358	+33	+10%	967	1143	+176	+18%
Sri Lanka	/13	/13	+/15	+27%	/30	/29	+/30	+167%	15	15			57	65	+8	+14%
Sudan	/15	/15	+/17	+26%	/32	/31	+/32	+100%	3	7	+4	+133%	20	26	+6	+30%
Suriname	/17	/17							1	1			2	2		
Swaziland	/17	/17	+/19	+36%	/32	/32			2	4	+2	+100%	6	9	+3	+50%
Sweden	/8	/8	+/12	+5.5%	/24	/24	+/25	+48%	419	458	+39	+9.3%	1325	1245	-80	-6%
Switzerland	/9	/9	+/12	+12%	/25	/25	+/26	+44%	506	544	+38	+7.5%	2289	2314	+25	+1.1%
Syria	/14	/13	+/14	+64%		/32	+/32	new	4	4			11	11		
Chinese Taipei	/7	/7	+/10	+14%	/21	/21	+/29	+0.39%	201	205	+4	+2%	517	622	+105	+20%
Tajikistan	/16	/16	+/20	+8.3%					8	8			12	29	+17	+142%
Tanzania	/15	/14	+/15	+118%	/30	/30	+/32	+33%	29	33	+4	+14%	42	62	+20	+48%
Thailand	/10	/9	+/11	+29%	/28	/27	+/29	+56%	243	280	+37	+15%	805	1011	+206	+26%

	IPv4 Addresses				IPv6 Addresses				Autonomous System Numbers				Network Adjacencies			
Timor-Leste	/20	/20							1	1			5	3	-2	-40%
Togo	/19	/18	+/20	+63%						2	+2	new	3	3		
Tokelau		/21	+/21	new		/32	+/32	new		1	+1	new				
Tonga	/20	/19	+/21	+47%	/32	/32			2	2			3	3		
Trinidad and Tobago	/13	/13			/29	/29	+/32	+17%	8	8			18	13	-5	-28%
Tunisia	/13	/11	+/11	+332%	/32	/32			2	2			1	1		
Turkey	/9	/9	+/12	+9.7%	/28	/27	+/29	+47%	339	373	+34	+10%	638	850	+212	+33%
Turkmenistan	/19	/19							2	2			4	8	+4	+100%
Turks and Caicos	/19	/19							2	2				2	+2	new
Tuvalu	/19	/19														
Uganda	/15	/15	+/18	+7.4%	/32	/30	+/30	+300%	15	20	+5	+33%	24	46	+22	+92%
Ukraine	/10	/9	+/11	+37%	/28	/27	+/28	+110%	1523	1721	+198	+13%	3662	5272	+1610	+44%
United Arab Emirates	/11	/11	+/14	+12%	/30	/30	+/31	+67%	13	18	+5	+38%	68	116	+48	+71%
United Kingdom	/6	/6	+/9	+9.5%	/22	/22	+/26	+6.1%	1710	1863	+153	+8.9%	7263	8474	+1211	+17%
United States	/2	/2	+/7	+2.5%	/18	/18	+/23	+3.1%	19169	20429	+1260	+6.6%	63230	72871	+9641	+15%
Uruguay	/12	/12	+/19	+1.1%	/28	/28	+/45	+0%	20	26	+6	+30%	58	58		
US Virgin Islands	/16	/16			/32	/32			6	6			14	18	+4	+29%
Uzbekistan	/14	/14	+/20	+2.1%	/31	/31			29	29			56	96	+40	+71%
Vanuatu	/19	/19			/32	/32			3	3			4	3	-1	-25%
Vatican City	/19	/19			/32	/32										
Venezuela	/10	/10	+/14	+7.2%	/28	/28	+/32	+6.7%	56	58	+2	+3.6%	112	120	+8	+7.1%
Viet Nam	/9	/8	+/10	+82%	/29	/28	+/31	+20%	87	107	+20	+23%	250	366	+116	+46%
Wallis and Futuna	/21	/20	+/22	+50%					1	1			1	1		
Yemen	/17	/17	+/18	+57%		/32	+/32	new	2	2			7	7		
Zambia	/16	/15	+/16	+163%		/31	+/31	new	7	12	+5	+71%	9	19	+10	+111%
Zimbabwe	/17	/17	+/21	+5.3%	/32	/31	+/32	+100%	7	12	+5	+71%	14	28	+14	+100%
Aland Islands	/20	/20							1	1			2	2		
Total	**/1**	**/1**	**+/4**	**+8.2%**	**/15**	**/15**	**+/20**	**+3.9%**	**44995**	**49667**	**+4,672**	**+10%**	**147633**	**179697**	**+32064**	**+22%**

APPENDIX 2: INTERNET EXCHANGE POINT REGIONAL FIVE-YEAR STATISTICS

	Internet Exchange Points				Domestic Bandwidth Production			
Region	**Apr 2006**	**Apr 2011**	**Net Change**	**Percent Change**	**Apr 2006**	**Apr 2011**	**Net Change**	**Percent Change**
Africa	18	22	+4	+22%	159M	3.22G	+3.06G	+1921%
Asia-Pacific	60	76	+16	+27%	636G	1.13T	+497G	+78%
Europe	85	137	+52	+61%	797G	6.28T	+5.49T	+688%
Latin America	20	34	+14	+70%	4.81G	62.3G	+57.4G	+1193%
North America	76	88	+12	+16%	121G	885G	+764G	+634%
Total	**259**	**357**	**+98**	**+27%**	**1.56T**	**8.37T**	**+6.81T**	**+81%**

APPENDIX 3: INTERNET EXCHANGE POINT NATIONAL FIVE-YEAR STATISTICS

	Internet Exchange Points				Domestic Bandwidth Production			
Country	Apr 2006	Apr 2011	Net Change	Percent Change	Apr 2006	Apr 2011	Net Change	Percent Change
Angola	1	1				3M	+3M	new
Argentina	1	1			529M	1.42G	+893M	+169%
Australia	10	11	+1	+10%	358M	3.52G	+3.17G	+885%
Austria	1	2	+1	+100%	8G	51.1G	+43.1G	+539%
Bahrain		1	+1	new				
Bangladesh	2	2			10M	94.3M	+84.3M	+843%
Belgium	2	2			6.22G	34.8G	+28.5G	+458%
Botswana	1	1						
Brazil	9	19	+10	+111%	3.47G	54.7G	+51.2G	+1478%
Bulgaria		2	+2	new		53.2G	+53.2G	new
Cambodia	1	1			1.9M	3M	+1.1M	+58%
Canada	2	2			3.03G	63.5G	+60.5G	+1998%
Chile	1	1						
China	2	3	+1	+50%	35G	131G	+96.5G	+276%
Colombia	1	1			690M	6G	+5.31G	+770%
Congo-Kinshasa	1	1						
Cote D'Ivoire	1	1			399K	4M	+3.6M	+903%
Croatia	1	1			117M	922M	+805M	+689%
Cuba	1	1			50M	50M		
Cyprus	1	1			10M	10M		
Czech Republic	1	3	+2	+200%	15G	169G	+154G	+1024%

	Internet Exchange Points				Domestic Bandwidth Production			
Denmark	1	1			8.14G	20.6G	+12.5G	+153%
Dominican Republic		1	+1	new				
Ecuador	2	2			40M	40M		
Egypt	1	2	+1	+100%	12M	261M	+249M	+2076%
Estonia	1	2	+1	+100%	1.9G	1.42G	-479M	-25%
Finland	2	2			11.6G	29.5G	+18G	+155%
France	9	15	+6	+67%	39.5G	45.4G	+5.9G	+15%
Germany	11	14	+3	+27%	74.6G	1.49T	+1.42T	+1901%
Ghana	1	1						
Greece	1	1			949M	9.88G	+8.93G	+941%
Haiti		1	+1	new		2K	+2K	new
Hungary	1	1			23.5G	134G	+111G	+472%
Iceland	1	1			1.07G	6.8G	+5.74G	+539%
India	4	7	+3	+75%	686M	8.33G	+7.64G	+1115%
Indonesia	5	7	+2	+40%	3.4G	3.4G	+20K	
Ireland	1	3	+2	+200%	1G	11.1G	+10.1G	+1009%
Israel	1	1			669M	2.55G	+1.88G	+280%
Italy	4	7	+3	+75%	23.2G	108G	+84.4G	+364%
Japan	14	16	+2	+14%	289G	535G	+247G	+85%
Kenya	1	1			12.7M	98.8M	+86.1M	+675%
Korea	4	4			302G	444G	+142G	+47%
Laos		1	+1	new		1M	+1M	new
Lebanon		1	+1	new		8.34M	+8.34M	new
Lithuania		1	+1	new				
Luxembourg	1	2	+1	+100%	102M	2.38G	+2.28G	+2235%
Malawi		1	+1	new				
Malaysia		1	+1	new		838M	+838M	new
Malta	1	1						
Mauritius	1	1						

	Internet Exchange Points				Domestic Bandwidth Production			
Mongolia	1	1						
Mozambique	1	1			4.14M	5M	+864K	+21%
Nepal	1	1			15M	19.2M	+4.23M	+28%
Netherlands	5	5			230G	1.35T	+1.12T	+486%
Netherlands Antilles		2	+2	new		1.49M	+1.49M	new
New Zealand	5	7	+2	+40%	447M	540M	+92.7M	+21%
Nicaragua	1	1						
Nigeria	1	2	+1	+100%		196K	+196K	new
Norway	1	2	+1	+100%	8.61G	34G	+25.4G	+295%
Pakistan		1	+1	new		3.6M	+3.6M	new
Panama	1	1						
Paraguay	1	1			40M	40M		
Peru	2	2						
Philippines	1	2	+1	+100%		3.3M	+3.3M	new
Poland	2	5	+3	+150%	6.57G	206G	+200G	+3036%
Portugal	1	1			2.77G	7.52G	+4.75G	+171%
Puerto Rico		1	+1	new				
Romania	1	2	+1	+100%	1.94G	31G	+29G	+1501%
Russia	4	14	+10	+250%	9.23G	687G	+678G	+7350%
Saudi Arabia		1	+1	new				
Singapore	2	3	+1	+50%		218M	+218M	new
Slovakia	1	2	+1	+100%	3.12G	42.3G	+39.2G	+1257%
Slovenia	1	1			13.9G	24.3G	+10.4G	+75%
South Africa	3	3			126M	2.81G	+2.68G	+2126%
Spain	6	6			152G	318G	+166G	+109%
Sri Lanka	1	1						
Swaziland	1	1			126K	256K	+130K	+103%
Sweden	8	12	+4	+50%	50.6G	241G	+190G	+377%
Switzerland	3	3			3.4G	25.1G	+21.7G	+637%

	Internet Exchange Points				Domestic Bandwidth Production			
Taiwan	4	4			3.56G	3.22G	-345M	-10%
Tanzania	1	2	+1	+100%	2.4M	9.56M	+7.16M	+298%
Thailand	1	1			1.7G	1.7G		
Turkey		1	+1	new				
Uganda	1	1			1.34M	29.4M	+28.1M	+2088%
Ukraine	2	5	+3	+150%	11.9G	289G	+277G	+2326%
United Kingdom	9	12	+3	+33%	87G	853G	+766G	+880%
United States	74	85	+11	+15%	118G	821G	+704G	+599%
Viet Nam	2	2			300M	300M		
Zambia	1	1						
Zimbabwe	1	1						
Total	**259**	**357**	**+98**	**+27%**	**1.56T**	**8.37T**	**+6.81T**	**+81%**

APPENDIX 4: COUNTRIES AND TERRITORIES STILL LACKING AN INTERNET EXCHANGE POINT

Afghanistan
Albania
Algeria
American Samoa
Andorra
Anguilla
Antigua and Barbuda
Armenia
Aruba
Azerbaijan
Bahamas
Barbados
Belarus
Belize
Benin
Bermuda
Bhutan
Bolivia
Bosnia and Herzegovina
Bouvet Island
British Virgin Islands
Brunei
Burkina Faso
Burundi
Cameroon
Cape Verde
Cayman Islands
Central African Republic
Chad
Christmas Island
Cocos Islands
Comoros
Congo Brazzaville
Cook Islands
Costa Rica
Djibouti
Dominica
El Salvador
Equatorial Guinea
Eritrea
Ethiopia
Greenland
Grenada
Guadeloupe
Guatemala
Guernsey
Guinea
Guinea-Bissau
Guyana
Heard and McDonald Islands
Honduras
Iran
Iraq
Isle of Man
Jamaica
Jersey
Jordan
Kazakhstan
Kiribati
Kuwait
Kyrgyzstan
Lesotho
Liberia
Libya
Liechtenstein
Macao
FYR Macedonia
Madagascar
Maldives
Mali
Mariana Islands
Marshall Islands
Martinique
Mauritania
Mayotte
Mexico
Micronesia
Moldova
Monaco
Montenegro
Montserrat
Morocco
Palau
Palestine
Papua New Guinea
Pitcairn
Qatar
Reunion
Saint Barthelemy
Saint Helena
Saint Kitts and Nevis
Saint Lucia
Saint Martin
Saint Pierre and Miquelon
Saint Vincent and the Grenadines
Samoa
San Marino
Sao Tome and Principe
Senegal
Serbia
Seychelles
Sierra Leone
Solomon Islands
Somalia
South Georgia and South Sandwich
Sudan
Suriname
Svalbard and Jan Mayen
Syria
Tajikistan
Timor-Leste
Togo
Tokelau
Tonga
Trinidad and Tobago
Tunisia
Turkmenistan
Turks and Caicos
Tuvalu
United Arab Emirates
Uruguay
US Virgin Islands
Uzbekistan

Falkland Islands
Faroe Islands
Fiji
French Guiana
French Polynesia
Gabon
Gambia
Georgia
Gibraltar
Myanmar
Namibia
Nauru
New Caledonia
Niger
Niue
Norfolk Island
North Korea
Oman
Vanuatu
Vatican City
Venezuela
Wallis and Futuna
Western Sahara
Yemen
Yugoslavia
Aland Islands

ANNEX 1

SURVEY OF CHARACTERISTICS OF INTERNET CARRIER INTERCONNECTION AGREEMENTS [76]

The Internet, or network of networks, consists of 5,039 Internet Service Provider (ISP) or carrier networks, which are interconnected with one another in a sparse mesh.[77] Each of the interconnecting links takes one of two forms: transit or peering. Transit agreements are commercial contracts in which, typically, a customer pays a service provider for access to the Internet; these agreements are most common at the edges of the Internet. Transit agreements have been widely studied and are not the subject of this report. Peering agreements—the value-creation engine of the Internet—are the carrier interconnection agreements that allow carriers to exchange traffic bound for one another‘s customers; they are most common in the core of the Internet. This report examines and quantifies a few of the characteristics of Internet peering agreements.

The Survey

In preparing this report, we analysed 142,210 Internet carrier interconnection agreements. We collected our data by voluntary survey, distributed globally through all of the regional Network Operators Groups between October 2010 and March 2011. The responses we received represented 4 331 different ISP networks, or approximately 86% of the world‘s Internet carriers, incorporated in 96 countries, including all 34 OECD member countries and seven of the 48 UN Least Developed Countries. For each agreement, in addition to the identities of the carriers party to the agreement, we asked the following three questions:

- Is the agreement formalized in a written document, or is it a ―handshake" agreement?
- Does the agreement have symmetric terms, or do the parties exchange different things?
- What is the country of governing law of the agreement?
- In addition, we made the following determination for each agreement:
- Is the agreement bilateral or multilateral?

In 1,032 cases, both parties to the same agreement responded to our survey, and in 99.52% of those cases, both parties‘ answers to each of the three questions were identical. We believe that, among other things, this indicates that respondents understood the questions clearly and were able to answer unambiguously and accurately.

Figure A1.1. Top thirty countries of incorporation of the represented networks, as a percentage of those in the dataset.

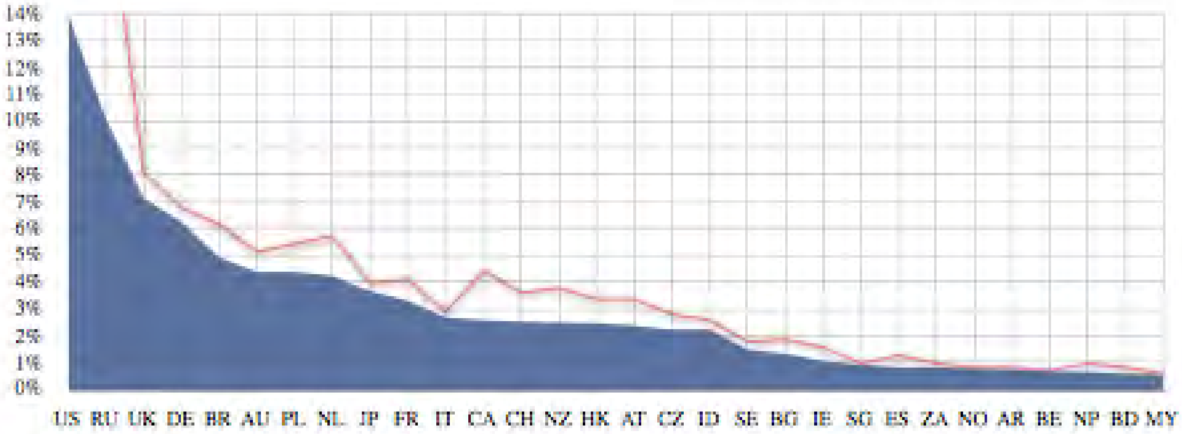

Source: Packet Clearing House.

The largest number of networks represented in the dataset were incorporated in the United States (466), followed by Russia (337), the United Kingdom (239), Germany (209), and Brazil (165). On the long tail of the curve, 45, or nearly half, of the countries were represented by three or fewer networks. The line in Figure A indicates the total number of networks incorporated in each country; the area indicates those represented in the responses to our survey. In most countries, a significant and relatively uniform majority of the networks are represented in our data, but our coverage in the United States (30%) and Russia (52%) was disproportionately small relative to other countries, and this does slightly affect the results of some of our country-specific analyses of these two countries, as we discuss later.

Informal agreements

Of the total analysed agreements, 698 (0.49%) were formalized in written contracts. The remaining 141 512 (99.51%) were ―handshake" agreements in which the parties agreed to informal or commonly understood terms without creating a written document. The common understanding is that only routes to customer networks are exchanged, that BGP version 4 is used to communicate those routes, and that each network will exercise a reasonable duty of care in co-operating to prevent abusive or criminal misuse of the network.[78] This huge number of informal agreements are arrived at by the ―peering co-ordinators" or carrier-interconnection negotiation staff of the networks, often at self-organised regional or global ―peering forums" that take place many times each year.[79]

Symmetric terms

Of the agreements we analysed, 141 836 (99.73%) had symmetric terms, in which each party gave and received the same conditions as the other; only 374 (0.27%) had asymmetric terms, in which the parties gave and received conditions with specifically defined differences. Typical examples of asymmetric agreements are ones in which one of the parties compensates the other for routes that it would not otherwise receive (known as ―paid peering"),[80] or in which one party is required to meet terms or requirements imposed by the other (―minimum peering requirements").[81] In the more common symmetric relationship, the parties to the agreement simply exchange customer routes with each other, without settlements or other requirements.[82]

Governing law

No interconnection agreements were reported that utilised a country of governing law that was not also the country of incorporation as well as the location of primary operation of one of the two carriers party to the agreement. Stated another way, in no case did the parties choose a country of governing law that was not one of their own countries of incorporation and primary operation. This indicates that there is, as yet, no country that has such compelling rule of law in the field of carrier interconnection as to incentivize this behaviour. Contrast this with other areas of commerce in which countries tailor regulatory or legislative environments to attract business as, for example, the registration of much maritime shipping in Panama or banks in Switzerland.

Nonetheless, clear preferences were expressed in the data, with the distribution of countries of governing law being sparser than the distribution of countries of incorporation and operation. In other words, some countries' governing law was preferred to a greater degree than their frequency as a country of incorporation would suggest, whereas others were preferred for governing law less frequently than they appeared as a country of incorporation.

Figure A1.2. Probability of selection as an economy of governing law, ten most-likely and ten least-likely economies

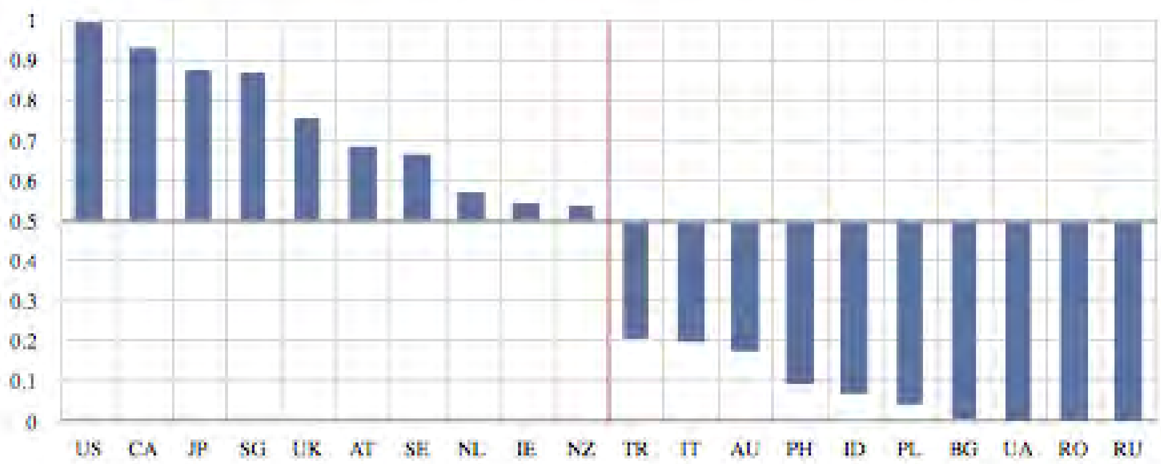

Source: Packet Clearing House.

When we compare the frequency of appearance as a country of incorporation to the frequency of selection as a country of governing law (Figure A1.2), in nearly every interconnection agreement in which one of the two parties is incorporated in the United States or Canada that country is selected as the country of governing law in preference to the country of incorporation of the other party to the agreement. At the opposite end of the spectrum, there were no agreements in the dataset in which Russia, Romania, or the Ukraine was selected to supply governing law for an agreement with a country outside this group of three, even though 337 Russian, eighteen Ukrainian, and eight Romanian networks are represented in the dataset. Each time a Russian, Romanian, or Ukrainian network interconnected with a foreign network, the parties elected to use the other country's governing law.

National interconnection partners

Looking solely at the frequencies with which pairs of countries of incorporation appear within the dataset, it is possible to chart the relative number of connections between any country and all others. By way of example we chart the most frequent interconnection partners (those consisting of more than 1%) of each of the four countries that are most frequently represented in our dataset—the United States, Russia, the United Kingdom, and Germany (Figure A1.3).

Figure A1.3. Trends in peering partners, selected countries

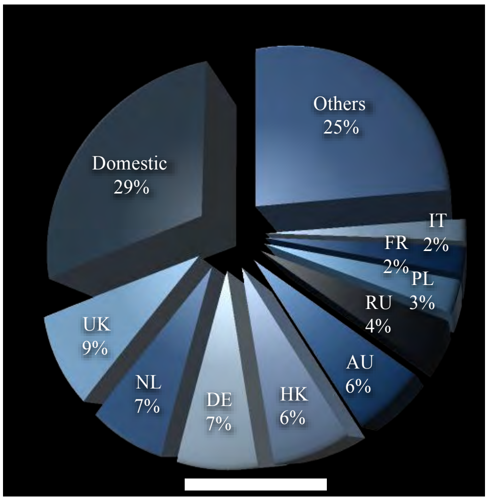

United States

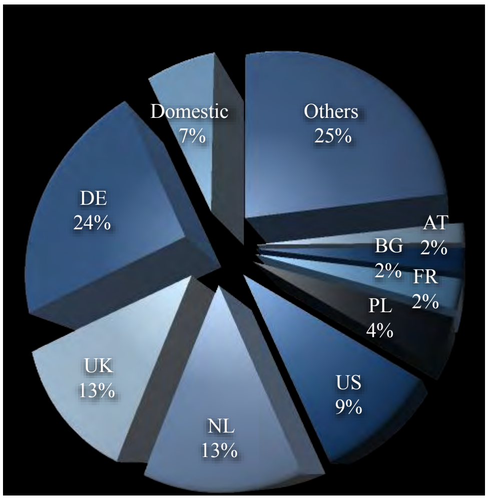

Russia

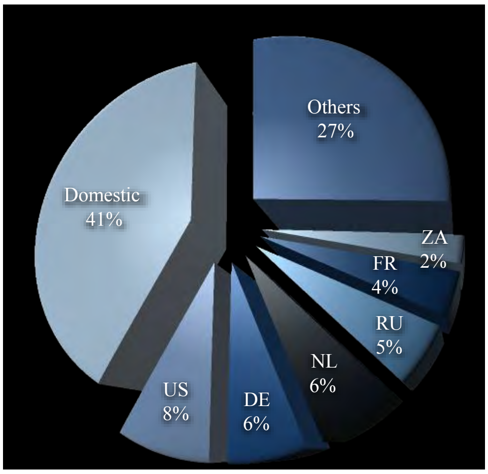

United Kingdom

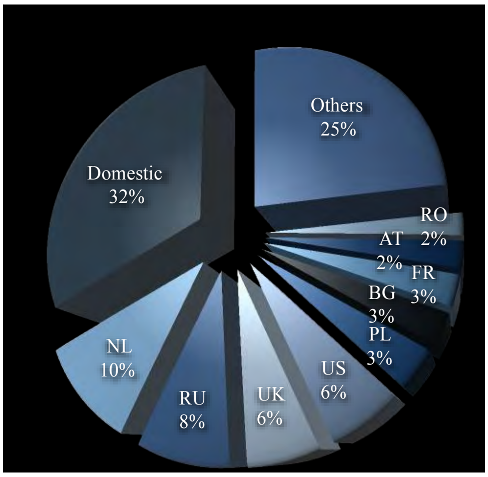

Germany

Source: Packet Clearing House.

Among these partners, linguistic cohorts, geographically proximal neighbours, and frequent commercial trading partners are favoured. The only real surprise is the relatively small share of domestic interconnection agreements observed within Russia, and we believe that this can be attributed to a selection bias in the dataset rather than to actual conditions on the ground; though we received many survey responses from networks that interconnect with United States and Russian networks, fewer were received from United States and Russian networks themselves, which would account for their relatively low shares of domestic interconnections.

Degree of interconnection

Most of the networks represented have small numbers of interconnection partners. Of the 4 331 networks, 2 696 (62%) have ten or fewer interconnection agreements, and only twelve of the represented networks have more than 700 interconnection agreements.

Figure A1.4. Distribution of number of networks (X axis) with each quantity of interconnection partners (Y axis)

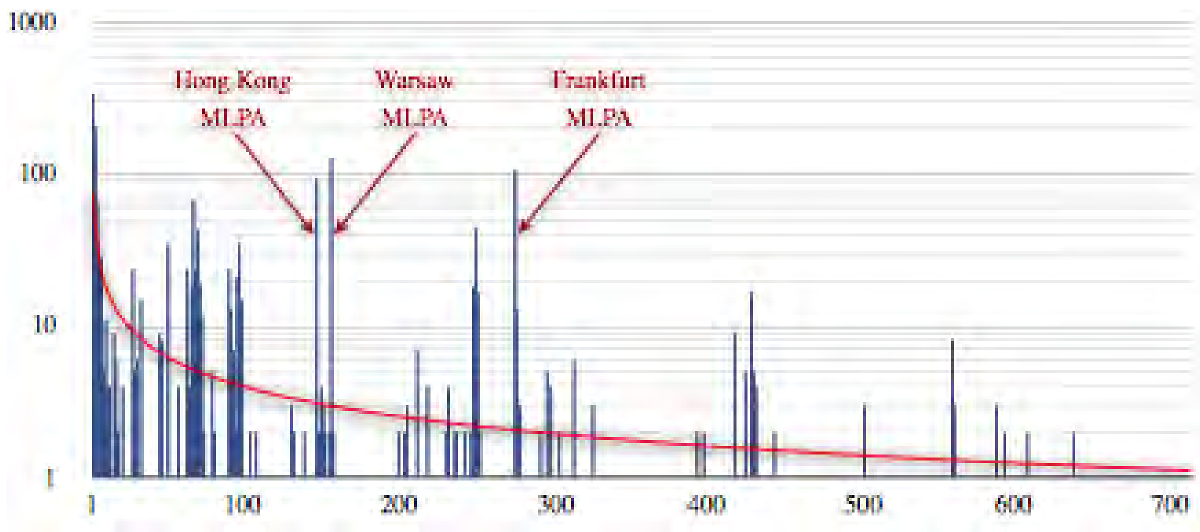

Source: Packet Clearing House.

A number of ―spikes" are visible in the distribution graph (Figure A1.4), with major ones appearing clustered around the values 144, 154, and 271. These are the effect of large multilateral peering agreements (MLPAs), specifically the ones associated with the Hong Kong, China, Warsaw, and Frankfurt Internet exchange points.

Figure A1.5. Expanded view of the Y axis range 140–160 from Figure 4, detailing the Hong Kong, China and Warsaw MLPAs

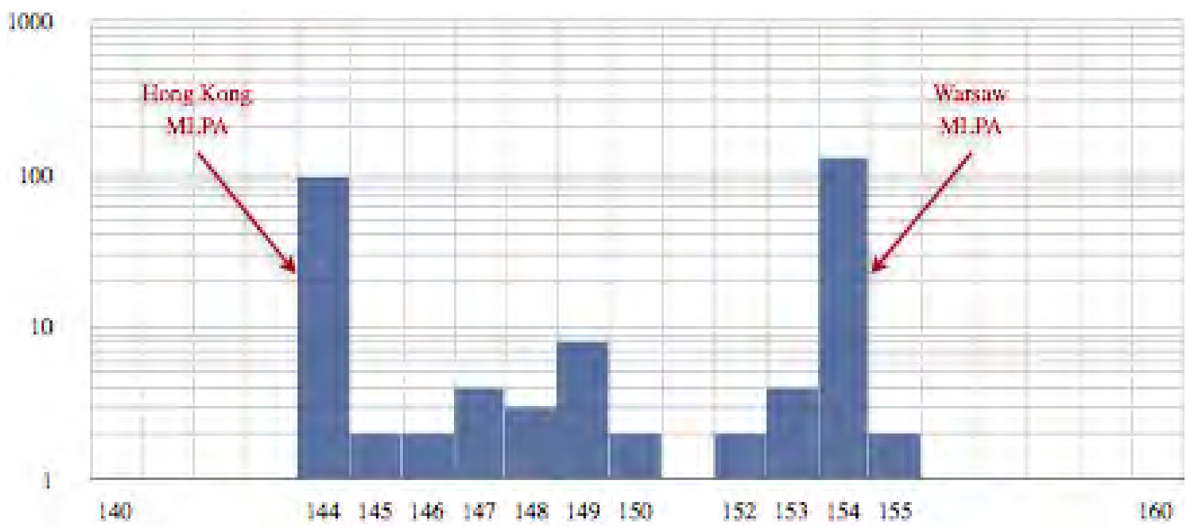

Source: Packet Clearing House.

In each case, there exist a large number of networks that all peer with each other, creating a spike at that value, which trails off as a function of the portion of those networks that also have other interconnection agreements. To some degree, the volume of the tail to the right of the spike varies with the age of the MLPA, since MLPAs that have existed longer generally include members who have had more time to also form bilateral agreements outside the MLPA. Generally speaking, multilateral peering agreements are identifiable as spikes that have similar values in both X and Y axes in Figures A1.4 and A1.5.

Unexpected results

One unexpected result of this survey is a new understanding of the prevalence of multilateral peering. Multilateral peering, the exchange of customer routes within groups of more than two parties, has long been characterised as a practice principally engaged in by smaller networks. It has been commonly assumed that large networks decline to participate in multilateral peering agreements, and that multilateral agreements are therefore outside of the mainstream of peering practice. Although the method by which we collected our survey data does not allow us to compare absolute quantities of bilateral agreements to multilateral agreements, the majority of the Autonomous System pairs we observed were connected through multilateral agreements, and many of those agreements were very large, with dozens or hundreds of participants.[83] With the exception of the cluster circled in red (which consists of “Tier-1” ISPs), each of the other vertical clusters in Figures A1.6 and A1.7 represents a multilateral agreement, similar to the spikes in Figures A1.4 and A1.5.

Figure A1.6: Number of advertised prefixes (Y axis) over number of interconnection partners (X axis) per carrier

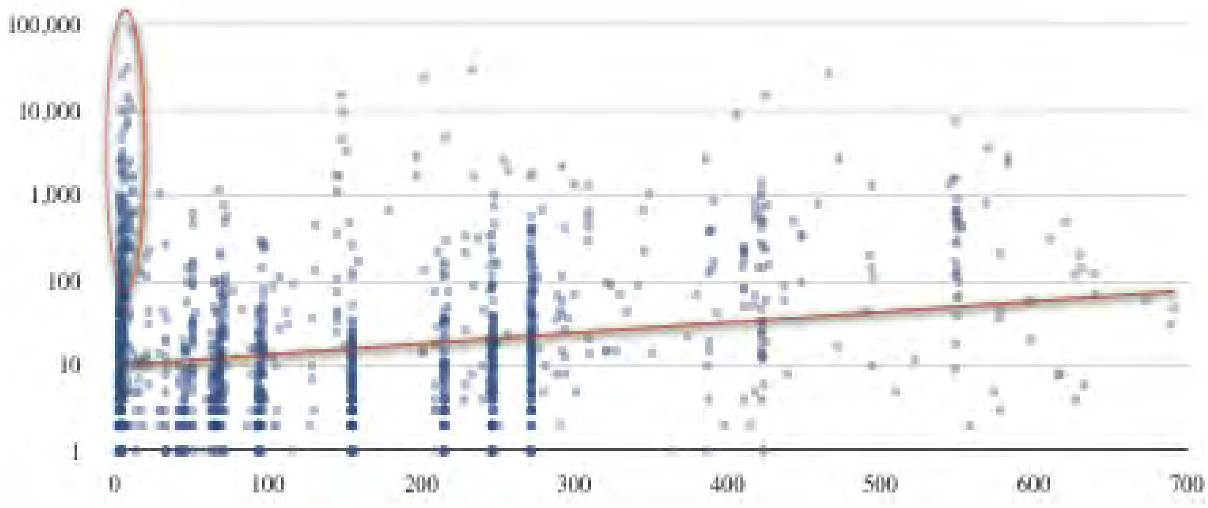

Source: Packet Clearing House.

It seems possible that, just as “donut peering” overtook “Tier-1” peering in the late 1990s, multilateral peering may now be overtaking bilateral peering, at least in sheer numbers, if not necessarily in volume of traffic.[84] In both cases, market-dominant networks loudly derided as “peripheral” a practice that sought to render them irrelevant, but that practice slowly gained prevalence over time, becoming mainstream without ever receiving much notice. As an example, the 144 participants in the Hong Kong, China Internet Exchange multilateral peering agreement represent 10 296 AS-pair adjacencies, and *each one* of those participants individually exceeds the average “Tier-1” carrier in degree of interconnection. When articulated in writing, multilateral peering agreements tend to follow the same general form and terms as other peering agreements, with the sole exception of having more than two parties.[85]

Figure A1.6. Number of advertised IPv4 addresses (Y axis) over number of interconnection partners (X axis) per carrier

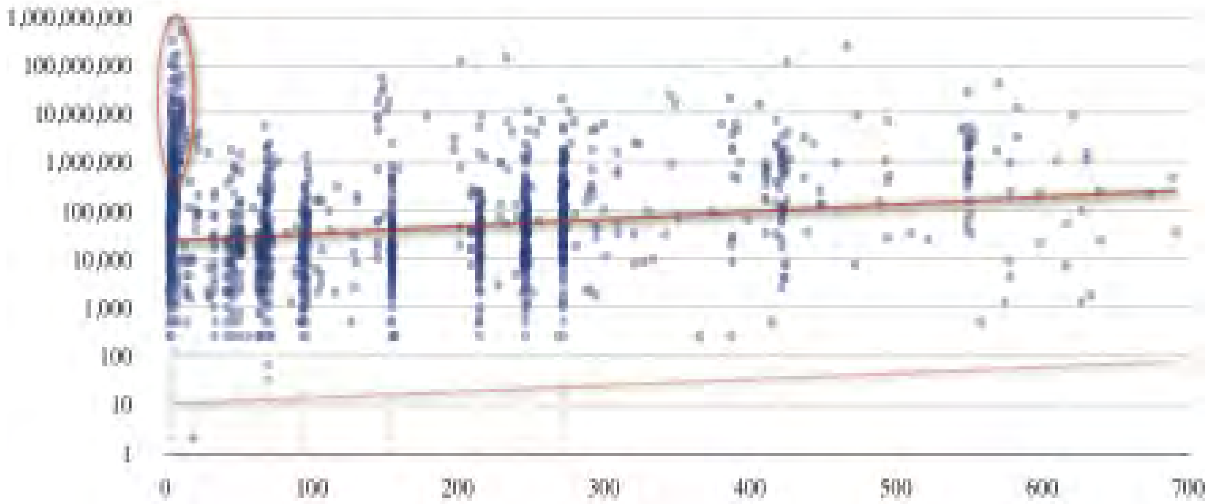

Source: Packet Clearing House.

Another finding of this survey, predictable on the face of it but to an unexpected degree, is how far the mainstream trend in number of interconnection agreements has left behind the legacy “Tier-1” networks, which tend to rely upon very small numbers of interconnection agreements. The circled clusters in Figures A1.5 and A1.6 represent the “Tier-1” ISPs, each of which has a large number of advertised IPv4 prefixes, and consequently a larger number of actual IPv4 addresses, yet very few interconnection partners. Following the red line that indicates the average correspondence between size and number of interconnection partners to the right, most of the “Tier-1” ISPs would have several thousand peers, if they were within mainstream ratios. By contrast, large content-distribution networks (“CDNs”), which have similar scale and degree of infrastructural investment tend to be exemplars of mainstream trends in our data, with very broad interconnection, both in absolute numbers and in geographic diversity. Although this may be self-evident, we expected to see a single order of magnitude difference between the number of agreements held by similarly sized networks in those two categories, whereas the actual difference was of two orders of magnitude.

Further work necessary

One weakness of this study, which provides reason for future work, is that we had relatively few mechanisms by which to compare the distribution of the responses we received to an objective “ground truth,” or to pre-existing datasets, in order to determine how statistically representative our survey respondents were to the Internet as a whole. Because previous studies of carrier interconnection agreements have been many orders of magnitude more narrowly focused than this one, they do not provide a statistically useful baseline against which we can characterise our dataset. A comparison against our own internal interconnection agreement data would have shed little light on how the survey dataset compares to the Internet as a whole and would have precluded including our own network's data in the survey. Furthermore, there is no mechanism for directly observing all of the peering agreements that exist in the Internet, and thus no ground truth to compare to. We hope that our foray into characterisation of carrier interconnection agreements encourages researchers in the academic community to follow up with further work on the subject.

ANNEX 2

REGIONAL PEERING

Large globe-spanning networks may decline to peer with smaller regional networks out of a concern over notional ―free riding.” The rationale expressed is that, since these small networks peer in fewer locations and cover a smaller territory, the average distance between the larger network‘s customer and the IXPs at which they would hand off traffic is greater than the average distance between the smaller network‘s customer and the IXPs. Although there is no issue of fairness involved since there are no externalities or shared costs to be divided, the larger network may feel that punitive behavior will, while hurting both parties, neutralize advantages conferred by the smaller network‘s greater efficiency.

Figure A2.A. Regional Peering

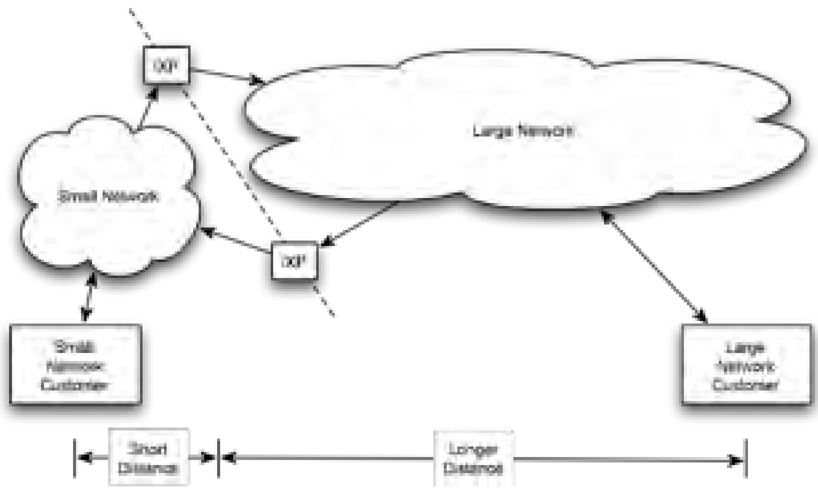

Source: Packet Clearing House

A compromise mechanism called ―regional peering” has existed for at least twelve years but is not widely practiced.[86] By partitioning the set of routes advertised at each IXP, the networks can agree to peer for traffic in-region, but not for traffic outside the region (Figure A2.1).

Figure A2.2. .Within region peering

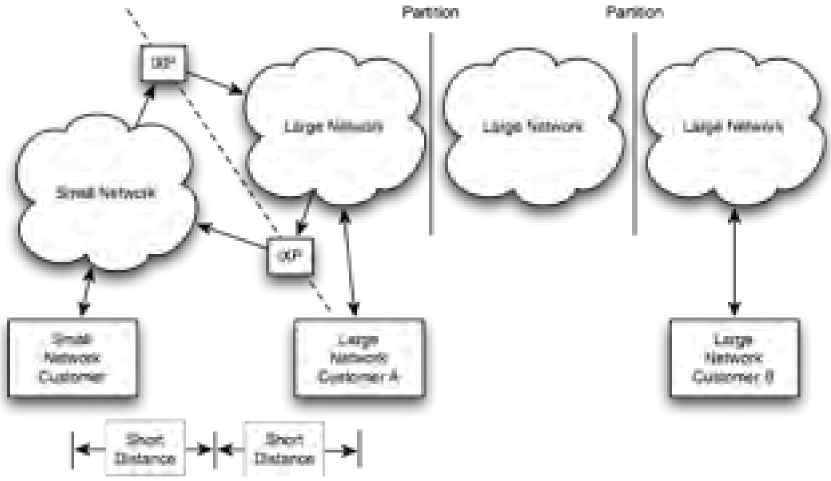

Source: Packet Clearing House.

In Figure A2.2, the large network has partitioned, making the routes to its Customer B unavailable to the small network at either IXP. This forces traffic between the small network and the large network‘s Customer B through less-efficient paths while allowing traffic ―within region" between the smaller network and the large network‘s Customer A to be exchanged through the efficient peering path. In all likelihood, it raises both carriers‘ costs relative to normal peering. It is, however, better than no peering at all.

Implementation of regional peering requires a more technically sophisticated routing configuration. It has been practiced by UUNet continuously since 1990 and implemented by a wide variety of other carriers since then.

Figure A2.3. Degree of overlap in announced prefixes (Y axis) by number of geographically distinct peering sessions (X axis) for 500 unique AS-pairs

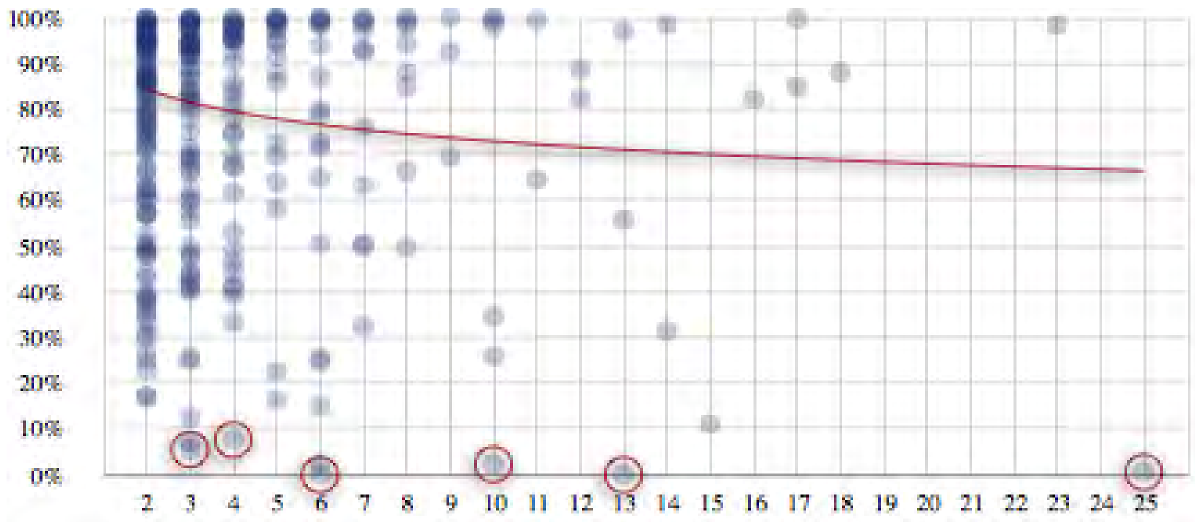

Source: Packet Clearing House.

In the chart of 500 unique AS-pairs which are peering at two or more geographically disparate locations (Figure A2.3), we see that the vast majority of ISPs advertise 100% congruent sets of prefixes across all locations, and those that differ between locations typically do so by only a small degree. Only a handful, perhaps 1%–2%, appear to be fully implementing a regional peering policy and, as it happens, none of those are actually legacy ―Tier-1" carriers.

ANNEX 3

CLOUD COMPUTING

Since it was popularized by the debut of Amazon‘s Elastic Cloud Computing service in August 2006, the broad distribution of computational resources into the ―cloud" of the undifferentiated Internet has become a required buzzword in any discussion of the future of computing and networking. Described simply, cloud computing refers to the outsourcing of computational cycles and storage to remote servers operated by a third party. Implicit in this notion is that the consumer of the resources is unaware of, and does not want to be bothered with, the details of how the service is provided, where it is located, or who specifically is providing the service.

One explanation for the rise of cloud computing is that it is a reaction to dot-com era ove-rinvestment in inefficiently utilised server hardware. As has been famously noted, a virtual inhabitant of the online game Second Life consumes the same amount of real-world electricity as a real-world inhabitant of Brazil, about 1.8m Wh/year, and produces 1.2 tons of CO_2.[87] Meanwhile, each Google search has been calculated to generate 5–10 g of CO_2.[88]

What is important about these figures is not the cost and impact of the successful services but that each unsuccessful service, each server sitting in a colocation centre idly spinning its disk while the dot-com company that owns it goes bankrupt, uses just as much power. The collective investment in idle and underutilised computational infrastructure by overly optimistic investors during the dot-com era sorely taxed the capacities of the supporting datacenter, electrical power, cooling, and server manufacturing industries and made it clear that a second round of such investment would not be possible.

This history spurred the success of server virtualisation technologies like Xen, VMware, Parallels, and QEMU, which allow the creation of multiple independently administered virtual servers running on the same physical server at the same time. This allows the consolidation of several lightly used servers onto a single much more efficiently utilised physical machine while reducing the cost and labour of hardware management as well as the power and environmental impacts of providing service. As well, and perhaps more important in the long run, it abstracts the logical tasks of administering the software services from the physical tasks of administering the server hardware. In doing so, it makes possible a new business model: the outsourced provision of generic computation and storage, decoupled from the specifics of operating systems and software.

The obvious providers of such a service? Those who had invested in the largest numbers of not-fully-utilized servers and had built out management infrastructure for all of that hardware. One might conclude that companies like Amazon and Google, which operate vast datacenters full of their own servers, should be able to predict their own needs accurately and not overprovision significantly. Yet servers are discrete quanta, and not available in an infinite variety of configurations. Thus the computational needs and the storage needs of large online content providers are unlikely to match up exactly, leading to excesses of one or the other. In addition, providing good service requires that capacity be provided in advance of demand, rather than too late, and the vicissitudes of physical facility construction and the logistics of server procurement and installation require that substantial slack be included in provisioning timelines. Thus, the average server must come online substantially in advance of need, so that the late server still arrives in time. All of this results in a window of overcapacity, typically weighted somewhat to one side or other of the compute cycles versus storage line. Since the abstraction advantages of server virtualization are of benefit within nearly any service provider network, this excess capacity is already packaged in such a way that it can be easily sold, recouping some value for the power and cooling it would consume anyway.

Ultimately, then, cloud computing serves both those who do operate a growing physical infrastructure plant, and have excess capacity, and those who do not need to purchase it on an outsourced basis.

From a network-topology perspective, cloud computing does not change traffic flows much. The servers (and their corresponding traffic flows) may be under more centralised ownership, but the actual locations are likely to still be large datacenters with good fibre connectivity to major IXPs. As with most developments in the use of the network, this one will increase network usage over time, as people store backups and content in remote datacenters rather than in local, directly attached storage.

Likewise, the reliability of services based on cloud computing infrastructure is unlikely to be significantly better or worse than that of services built on traditionally owned hardware. Although the people operating the hardware are likely to be doing so in a much more professional manner, the degree of abstraction and the reduction in the alignment of incentives between the operators of the hardware and software both counteract that improvement. If a customer organisation truly embraces virtualisation and cloud computing, however, long-term reliability gains may be realised through resilience in the face of the rare catastrophic events that take a whole datacenter offline in the long term: an aggressively virtualised party could just redeploy the same servers on a different cloud computing service.

The largest unresolved drawbacks to cloud computing are in the areas of privacy and security. Because cloud computing services may be provided from any location, without the knowledge of the client business, it is very likely that many smaller commercial consumers of cloud services are unwittingly violating national and regional privacy directives by transporting personally identifiable customer information into jurisdictions where it receives fewer protections. And because the cloud service provider has no insight into the nature of the content it is hosting, it also has little knowledge of the particular threats against it or who may or may not be authorised to observe or modify content or processes that handle it. Larger enterprises with internal regulatory-compliance assurance organisations are more likely to do the due-diligence necessary to understand that they will simply be unable to comply with privacy protection regulation if they utilise cloud computing services. As de-anonymisation attacks have demonstrated time and time again, customer data cannot be so stripped of relevant context as to be reasonably protected against reconstruction and correlation with other data sources.[89] Cryptographic protection of entire computational processes and subdivision of processes across multiple unrelated clouds have been proposed as possible means of mitigating these problems, but both pose substantial technical challenges and remain relatively vulnerable to attackers who have visibility of a large portion of a customer‘s processes. In other words, these defenses rely more upon security-through-obscurity than upon formally defensible mechanisms.

Because it aggregates supply efficiently and provides a high degree of liquidity in the marketplace for computational and storage resources, cloud computing is valuable and is unlikely to be a passing fad. At the same time, it poses substantial security challenges, so it is probably not reasonable to use cloud computing for the handling of personally identifiable information until substantial security advances have been made and tested by time.

ANNEX 4

WHO PAYS FOR WHAT?

Unlike the endless settlements and tariffs and regulatory maneuvering of the legacy circuit-switched voice market, the Internet has a simple method of allocating costs: where networks peer, each party pays its own costs and retains its own revenue. The simplicity of this cost model, and the assurance of fairness that it gives the market, is the single most important factor in the Internet‘s thirty-year history of exponential growth. Without this assurance that investments can be capitalized upon without fear of predatory behaviour on the part of the networks one interconnects with, or of unpredictable, external costs imposed by regulatory action, perhaps at the behest of one‘s competitors, the Internet could never have received the degree of investment that was required to grow and replace the circuit-switched network. This freedom from distortion is what has allowed the exponential growth that has made the Internet what it is.

Figure A4.1. Legacy circuit-switched settlements and financial externalities

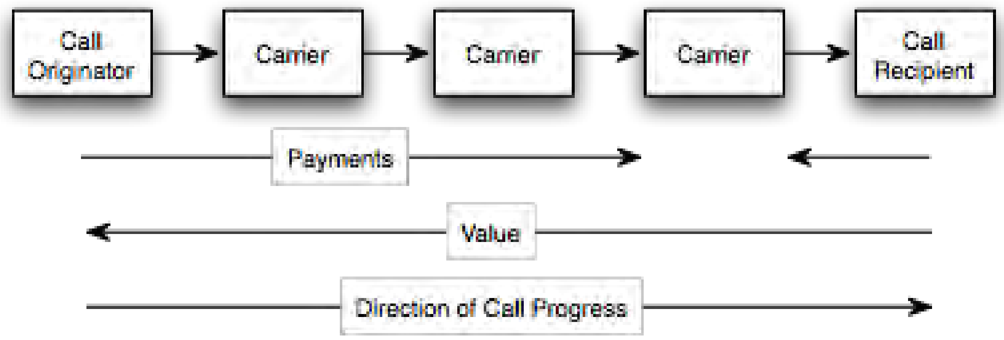

Source: Packet Clearing House.

In the legacy circuit-switched voice network under Calling Party Pays (Figure A4.1), calls were directional, in that they had an originator and a recipient, and payments followed the direction of call progress up to the recipient‘s carrier, who received payment from both directions. This was thus a privileged position and one that all carriers jockied to attain in as many cases as possible. Other frameworks have been used, such as one-sided access charges. Because each of these systems rewarded gaming of the market, exploitation of loopholes, and creation of burdens for one‘s competitors, they often distracted participants from the actual creation of value for a customer. This system neither rewarded efficiency nor directed investment to its best use, and the ultimate result was an expensive, linear-growth market that produced little value for the public.

Figure A4.2: Internet payments and value symmetric and compartmentalised

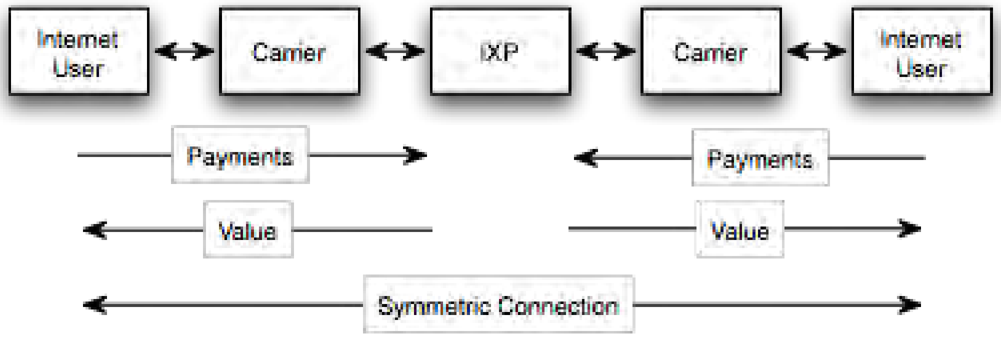

Source: Packet Clearing House.

By contrast, in Internet peering the connection is symmetric, having no higher value to the party at one end than the other, and each party pays its own way, retaining the value that it creates (Figure A4.2). There is no privileged position that expects a settlement from any other party, so no incentive to engage in gaming of the market. Since participation is voluntary, behaviour perceived as detrimental can be avoided by simply ending the agreement. Similarly, when money does change hands in a transit arrangement, it happens only because both sides are made better off. Protectionism is not rewarded, and in fact it carries a debilitating cost in the form of forgone growth. Because of this more straightforward business model and freedom from externally imposed costs, investment is rewarded, and thirty years of exponential growth have resulted from reinvestment of those profits.

Since speed distance = cost, more or less, in telecommunications systems, it is important to note that the distance between Internet users and the IXPs that their traffic passes through is the single largest factor in the cost of the service they receive. Thus, as illustrated in Figure A4.5, we see that Internet User A and Carrier A, who are farther from an IXP, pay more than Internet User B and Carrier B, who are nearer to that IXP.

Figure A4.3. Simplified model of Internet traffic exchange, showing difference in cost of service depending upon distance from IXP

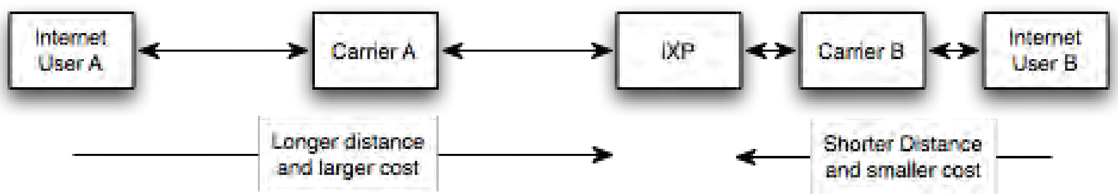

Source: Packet Clearing House.

Because exchange points would only rarely be exactly midway between two users, in most individual communications one user would wind up paying more than the other user, even if, on average, it all balanced out. But the simplified model shown here does not tell the whole story. It illustrates an exchange of traffic with mirror symmetry through a single IXP, which is actually rare. Normally, the path of traffic exchange has rotational symmetry, rather than mirror symmetry, and the traffic in each direction passes through a different exchange, as depicted in Figure A4.4.

Figure A4.4. Full model of Internet traffic exchange, showing normal symmetric retention of costs and revenue in traffic passing between two carriers

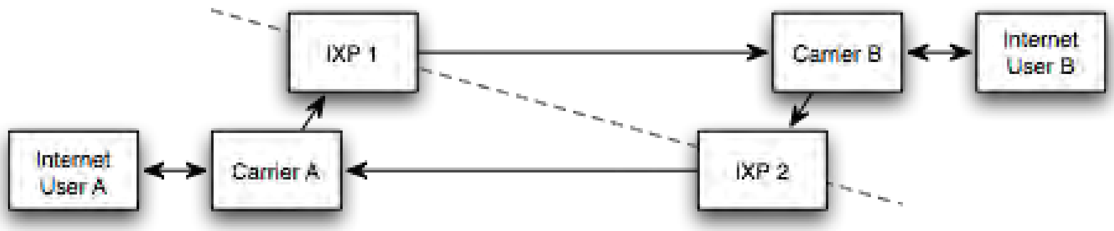

Source: Packet Clearing House.

Here, each carrier makes a ―hotpotato routing" decision, when faced with two possible paths to the other carrier, through IXP 1 and IXP 2. In every case, a rational carrier delivers an outbound packet through the shortest available path, because this minimises cost while maximising performance. For Carrier A, IXP 1 is nearer than IXP 2, thus Carrier A always delivers outbound packets toward Carrier B through IXP 1. This leaves Carrier B, which has received the packet at IXP 1, with the larger cost of hauling the packet the greater distance inbound from IXP 1 before delivering it to its User B and receiving payment. This division of costs is reversed, however, when User B replies to User A, and Carrier B makes the same hot-potato routing decision, delivering outbound packets through IXP 2. Now Carrier A pays more to haul inbound packets in from IXP 2 before delivering them to User A and receiving payment. Thus, provided there is an IXP near both Carrier A and Carrier B, and both carriers participate in both IXPs, there is an equitable sharing of costs and revenues in that each carrier retains all of its own revenue and bears only its own costs, and neither‘s costs are disproportionately higher than the other‘s.

The two caveats, however, are significant. First, there must be an IXP near a carrier for the carrier to be able to benefit from the lower costs of short outbound paths; second, the carrier must participate in distant IXPs in order for other carriers to be able to offload their own outbound traffic on similarly short paths.

In the relatively optimised markets of equal-sized carriers in regions densely populated with IXPs, like Western Europe and the coasts of North America, this equitable sharing is the norm. But when an ISP chooses to limit the entry points to its network, or in the opposite case when ISPs fail to establish and grow IXPs within their own region, this equitability cannot exist.

Figure A4.5. Locations of IXPs

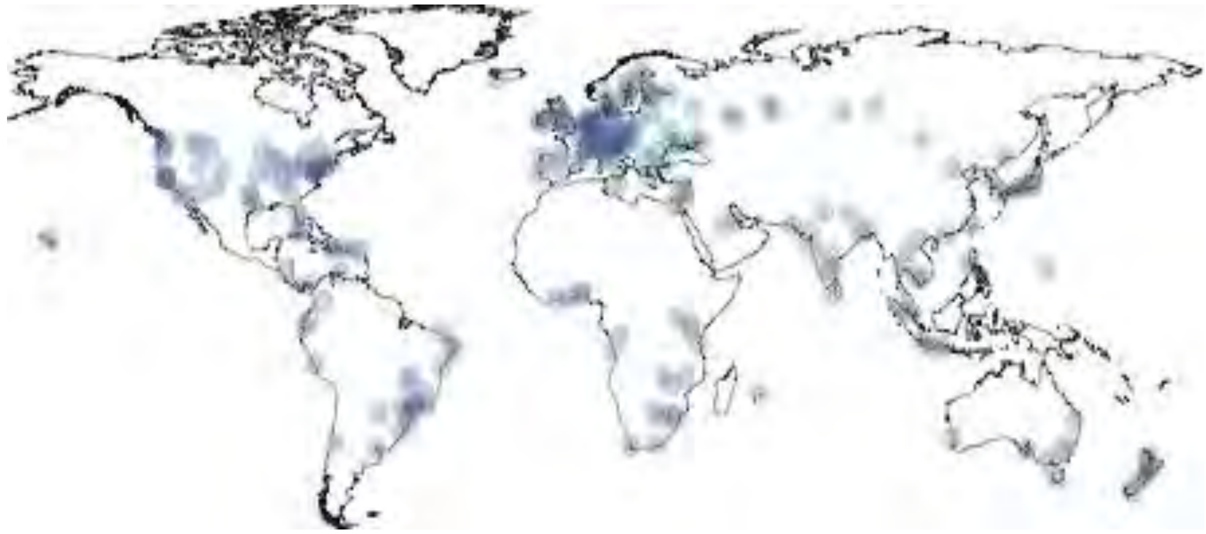

Source: Packet Clearing House.

At the opposite end of the spectrum, the density of IXPs varies greatly by region. Out of 357 IXPs in the world, African ISPs have established only 22, and Latin American ISPs only 34. There are substantial areas, like northern Africa and central Asia, where the ISPs have established none at all. In fact, 106 countries and territories presently have no IXP.[90] By contrast, European ISPs have established 137 and North American ISPs have established 88.[91] This means that, for instance, a Moroccan ISP can peer in Madrid, but a Spanish ISP cannot form a reciprocal connection in Rabat. By failing to establish an IXP near itself, the Moroccan ISP disadvantages itself, and consequently spends more money than necessary, hauling both inbound and outbound traffic to Madrid, rather than inbound only. This is only one of many substantial disadvantages to not having an IXP nearby.

ANNEX 5

IPV4 ADDRESSES AND THE FUTURE GROWTH OF IXPS

Figure A5.1. Distribution of Internet exchange points (Y axis) by number of participants (X axis)

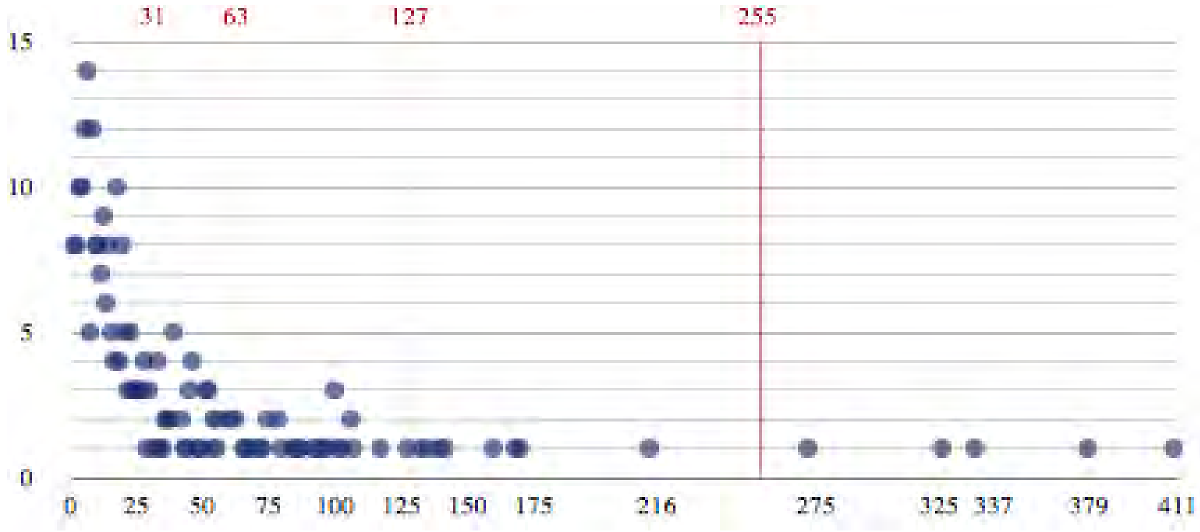

Source: Packet Clearing House.

The bulk of the world's IXPs currently have between ten and one hundred participants. The vertical red lines in Figure A5.1 indicate the boundaries of common sizes of IPv4 subnet used on IXP peering switch fabrics, with 256 being by far the most common size. As yet, only five IXPs have passed 256 participants, and each has been forced to renumber the IP addresses that their participants use when they had to expand from a block of 256 addresses to a larger block of 512. This renumbering process is extraordinarily laborious, since it requires the coordination of all of the 250-some member organisations in a complex technical procedure. We believe that if IXPs continue to be assigned blocks of 256 addresses, as more of them grow into maturity, we will see an artificial constraint begin to appear: IXPs hovering just below a 256-participant ceiling, growth halted while they defer the artificially imposed cost of renumbering. Policy proposals currently under consideration within the regional Internet registries would ensure the continued availability of IPv4 addresses to IXPs throughout the IPv4-to-IPv6 transition period, but the ratification of those proposals into public policy is far from assured.

ANNEX 6

PRACTICAL IMPLEMENTATION: MECHANISMS AND PRACTICES

This section is a guide to the most effective mechanisms and practices for implementing the suggestions made in the main body of this report.

In many ways, the process of continuously developing and making a more effective Internet economy is a matter of improving each portion of it in turn. This is because all of the parts are interdependent and each can serve as a bottleneck, halting the growth of the others. Thus, elements of the Internet economy must interact in an efficient manner if the economy as a whole is to continue to grow in the long term. Some of the key elements of a country‘s Internet economy are domestic bandwidth production, in the form of Internet exchange points; the local-loop infrastructure operated by the Internet service providers, to carry bandwidth to users‘ homes, schools, and workplaces; the international capacity that allows them to reach foreign destinations; and the publication of content that they find useful and relevant. If any of these are left unattended to, development will be much less than optimal.

Internet exchange points

Internet exchange points (IXPs) are the source of nearly all Internet bandwidth. A country that lacks IXPs must import Internet bandwidth from other countries that do possess them. Like factories and farms, they are a primary means of producing a commodity that‘s potentially quite expensive to import. As a general rule, the cost of telecommunication services is the product of the speed of the service multiplied by the distance covered (*i.e.* speed x distance = cost). The further afield you go for your bandwidth, the more expensive, and slower, it will be. Thus it‘s always preferable to use a local IXP, to one that‘s further away. As a result, IXPs proliferate in areas where Internet service providers, users, and policy makers are well informed on matters of telecommunication economics. A corollary is that a region which has many functioning IXPs and produces more bandwidth than it consumes, can export bandwidth to other regions at a profit; the Netherlands, for example, is a large net-exporter of Internet bandwidth, using only half of what it produces, domestically.[92]

The amortisation period of the investment required to construct an IXP is typically between two and twenty days. For example, if the cost of Internet transit in a country is USD 300 per megabit per second per month, an IXP is constructed at a cost of USD 20 000 and it produces one gigabit of bandwidth at inception, it‘s creating USD 300 000 per month in value, and would have an amortization period of two days.[93]

Internet exchange points typically cost between USD 2 000 and USD 50 000 to construct, depending how elaborate the physical facilities are, and each participating Internet service provider can expect to invest about that much again in upgrades to their own facilities to make efficient use of it. A more expensive facility is not necessarily a better one, and nearly always offers a lower rate of return on investment. It is invariably a better investment to build more small exchanges in more locations, than fewer larger ones in fewer locations.

Internet exchange points are best located coincident with centres of population density, because this minimizes the average distance traveled by packets that cross them. IXPs are the natural ―wire centres" or hubs for communications fibre infrastructure. IXPs need not be located within datacenters (which do not usually occur in high-population-density city centres), though some amount of datacenter space will naturally form immediately adjacent to any successful IXP, despite the cost of real estate being very high within a dense urban core.

Nearly any city is large enough to support an IXP; there is no set formula for determining a threshold of viability, but the elements to be considered are the size of the local economy and the cost of alternatives. The larger the local economy, the greater the demand for Internet bandwidth, and the more money there is to support growth. The further a city is from other IXPs, the higher the cost of Internet bandwidth, and the greater the beneficial impact of creating a new local source of bandwidth will be. As a rough indicator, there are approximately 140 IXPs in Europe, and approximately 90 in the United States; the 140th largest city in Europe and the 90th largest city in the United States each have about 225 000 inhabitants. Exchanges also fare perfectly well in towns of as few as 10 000 inhabitants if they are far removed from other population centres.[94]

It is important that IXPs publicise their success, particularly by publishing their aggregate traffic volume statistics in the commonly-accepted format.[95] This promotes visibility of the value they are creating, attracting additional participants who produce more traffic, in a virtuous cycle that further increases productivity and value.

Domain name service

The most critical service which runs atop the Internet is the Domain Name System. Domain Name Service (DNS, an acronym that does double-duty, standing for both System and Service) is the directory system that allows other services, like web sites and e-mail, to be located. Much like a telephone book, the DNS converts names to numeric addresses. Unlike a telephone book, the DNS is a distributed hierarchy, consisting of a “root,” operated by ICANN, the Internet Corporation for Assigned Names and Numbers, which delegates “Top Level Domains” or TLDs to national and commercial operators (for example, .ca, the country-code Top Level Domain, or ccTLD, for the country of Canada, is delegated by ICANN to CIRA, the Canadian Internet Registration Authority), which in turn delegate “second level domains,” and so on. Each level of this hierarchy contains information about the level below it, but generally no further.

The consequence of this is that in order to find services like web sites and email on the Internet, it is necessary to be able to communicate with the root nameservers and TLD nameservers, as well as the nameservers immediately responsible for the domain you wish to communicate with. If, for example, you wish to reach the web site “*www.cra.gc.ca*” to read about the policies of the Canada Revenue Agency, your computer must first communicate with a root nameserver, then a nameserver for .ca, the Canadian ccTLD, then one for gc.ca, the Canadian government, then one for cra.gc.ca, the Canada Revenue Agency. If any of these servers in this chain are unreachable, you will be unable to find and communicate with the Canada Revenue Agency, regardless of whether their own servers are up and running.

Thus it is critically important to have a reasonable selection of root and TLD nameservers well-connected within your country or cyber-defensive boundary. In practical terms, this means making connections to each of your Internet exchange points available to root and TLD nameserver operators, and recruiting a number of them to participate in your IXPs. Particular care should be taken to ensure that nameservers for one‘s own domains be located both within one‘s borders (to serve one‘s local constituency) and outside of one‘s borders (to absorb attacks and to ensure visibility of one‘s domestic services internationally). As an example, Kenya has servers for the .ke domain both in the IXP in Nairobi, and in many other IXPs around the world, and they have induced several root-server operators and nearly a hundred other ccTLD operators to bring servers to the IXP in Nairobi. This ensures that Kenyan Internet users are both well served locally and well defended internationally.

Local content

As discussed above, optimisation of price/performance of Internet services consists largely of reducing the distance between the source and destination of any packet of data, in order to allow it to travel more quickly, at a lower cost. The more Internet exchange points exist, the more nearly one will happen to fall on the axis between any two points that constitute a source and destination of traffic. However, even greater gains may be made if the source and destination can be brought closer to one another. If the source and destination are both individuals or organizations, it is likely impractical to move them about, but if one end of the transaction is a server, particularly a server of non-unique information, as is the case of most common web servers, the data on the server may be replicated to many locations, so as to be proximal to as many users as possible.

This is the principle behind Content Distribution Networks (CDNs), which replicate data to many servers, usually located in or adjacent to IXPs. Large CDNs like Google, Limelight, and Akamai, have servers in most major IXPs, and even many smaller ones, in an effort to give Internet users faster, cheaper access to the CDNs‘ customer‘s data.

A healthy Internet commercial ecosystem ensures that both locally-produced content and content served by foreign organisations can be easily hosted in close proximity to every IXP. This minimizes the cost and distance between the content and the users whose Internet traffic crosses that IXP.

As with Domain Name System data, other content, like local news stories, entertainment sites, and online ―radio stations” are best served both locally and remotely. By hosting their data locally, local users receive optimal performance at a low price; by hosting copies of the data in remote locations as well, they may be able to minimize international transit costs and hold some cyberattacks at bay.

Competition and regulation

The Internet ecosystem has evolved in the presence of thousands of networks which simultaneously co-operate to produce bandwidth and compete to sell it to consumers. In this environment, more competition invariably produces lower prices, better performance, and greater resilience. There is no quality of price or performance that is improved by lessening the degree of competition. There are no ―natural monopolies” in the Internet, and there is essentially no size of ISP that is too small to function efficiently and contribute positively to the Internet economy. Concentrations of bandwidth quickly become expensive bottlenecks, whereas broad distribution of bandwidth across many smaller competitors efficiently uses resources and creates a large market for infrastructural components.

There are two ways in which ineffective regulation could hinder or retard the growth of the Internet: by constraining the number of providers, and by constraining the number of customers. Most frequently, where this is present, policy and regulatory settings fail to support competition in the local loop, by constraining the number of network operators who are allowed access to the public rights-of-way, and thus reducing the amount of fibre infrastructure that can be built. As long as a monopoly or oligopoly network operator believes that the regulator will protect it from competition by disallowing competitors access to the right-of-way, that operator will create artificially constrained supply, by installing as little fibre as possible; only enough to meet its own needs, and never enough to support growth of the market as a whole. When policy or regulation constrains the number of network operators who can build physical facilities, those settings are directly reducing competition and the size of the market. Turning from providers to customers, policy or regulatory approaches that constrain the ―lower end” of the market by mandating minimum size, price, or quality of connections can put Internet access out of reach of less well-to-do potential users. It potentially does this by setting a floor on the price and making the Internet inaccessible to those who cannot afford that price; this likely does very little to promote or protect consumer interests, while constraining the number of potential customers, and thus the size of the market.

The primary goals of Internet regulation should be the promotion of competition and new market entrants, while ensuring that the underlying public rights-of-way can be inexpensively utilized by all. More ISPs, more IXPs, more content publishers, and more users, no matter how small each may be, is the surest and quickest path to growth, effectiveness, and prosperity for an Internet economy.

International connectivity

As a country‘s domestic Internet bandwidth production increases, so too will its use of international transit capacity. More users accessing more content will necessarily partake of both local and distant services and information, and a successful local economy will necessarily attract international customers and attention as well. Therefore, while a successful Internet economy will produce as much bandwidth locally as possible, it must also provide for parallel growth in the international facilities that keep it connected to the world.

This requires long-haul fibre cable systems be built, either laid across the sea-bed, or overland. Such fibre systems are extraordinarily expensive, often running to the hundreds of millions of USD. They are usually built by consortia of large international Internet and telecommunication companies, sometimes as few as two or three partners, and sometimes dozens. If access to these networks is made available to ISPs in an open and competitive environment it can provide a significant boost for growth in the overall Internet economy. At the same time, if access to such facilities does not exist or those facilities that do exist are constrained by monopoly power it can be a tremendous drag or hindrance on such growth. Where there is insufficient competition, policy makers and regulators need to examine what options are available to increase competitive access to international connectivity.

It is important to remember that a single cable system, regardless of the number of companies competing to sell service on it, may still provide very little in the way of technical redundancy or resiliency, and lands in a specific set of other countries, which may or may not be allies. As a matter of national policy, then, it‘s wise to encourage the construction of multiple independent systems of international cables, leading to as many different other countries as possible. Some of those countries may be able to supply bandwidth, others may be potential customers, and most will be both.

Voice-over-IP and ENUM

Voice-over-IP, or VoIP, refers collectively to the set of technologies used to pass voice communications over the Internet. This includes the protocols that allow PSTN calls to travel internationally, unbeknownst to caller or recipient, standards-based protocols that allow Internet-enabled telephones and private branch exchanges to interoperate, and even end-user applications like Skype that allow users to talk without considering the process as having anything to do with a telephone. VoIP has been embraced by OECD countries, as a technological step forward and a driver of private-sector innovation and commercial efficiency. VoIP has been outlawed or shunned in some countries, because of unwillingness of their incumbent monopolies to modernize, and the perception that VoIP use goes hand-in-hand with reduced international call termination revenues.

VoIP is a set of protocols and technologies, and is inherently policy-neutral. Some of the same telecommunication operators that strive to have VoIP illegalized for end-users take full advantage of the savings it affords on the wholesale side. It is a mistake to conflate the technology and protocols with tax and regulatory policies; they are fully independent and orthogonal. It is only when governments outlaw VoIP that they may lose regulatory levers over telephony. The tool by which the efficiency of VoIP is most easily harnessed while maintaining regulatory policy control is called ENUM, or E.164 NUmber Mapping. E.164 is the ITU-T standard for number assignment in the legacy PSTN network, while DNS, or the Domain Name System, is the Internet standard for service discovery in the modern network. ENUM allows a mapping of legacy telephone numbers into the modern Internet, with complete end-user control over number portability, in the same way that end-users control the disposition of domain names that they own.

Any country can, by operating an up-to-date ENUM delegation system for its E.164 country-code, give anyone, not just a handful of legacy telephone companies, access to telephone numbers within that country code, while maintaining full visibility of the inbound call disposition and length, and the ability to levy taxes on the inbound call termination revenue associated with those calls, rather than just those from a dwindling pool of legacy PSTN operators. In the long term, E.164 phone numbers are going away, because they don‘t add any value to the calling process, but support for ENUM can give countries that are dependent on call termination revenue a gradual path away from that dependency, rather than an abrupt and immediate loss of income.

IPv4 and IPv6 addresses and transition

The transition from Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6) addressing, which has been ongoing since 1995, reached a turning-point in 2011, as the stock of previously-unused IPv4 addresses has reached depletion in some regions.[96] Access to IP addresses is regulated by the Regional Internet Registries on the basis of need.[97] Organisations are qualified to receive allocations of IP addresses based upon demonstrated need and a history of responsible use. Unfortunately, that no longer guarantees a supply of IPv4 addresses, absent a transfer of unused addresses from a previous user. The transitional technologies (―NAT” and ―dual stack”) that will allow the Internet to move from IPv4 to IPv6 require each network to use some IPv4 addresses for backwards-compatibility, as well as the IPv6 addresses upon which their future growth depends. This is true of new market entrants as well as previously-established networks. This means that national regulators who wish to understand the IPv4 backwards-compatibility requirements of new market entrant network service providers may wish to communicate with their Regional Internet Registries, or even to participate in the IP address allocation policy development process.

ANNEX 7

WHY HAS THE INTERNET MARKET PERFORMED SO WELL?

The performance of the market for Internet traffic exchange stands in sharp contrast to that of inter-carrier markets for the exchange of traditional circuit-switched traffic. Over the past five years the Internet market has sustained prices for connectivity that are vanishingly low compared to the wholesale prices commonly seen for the exchange of circuit-switched voice (time-domain multiplexed, or TDM) traffic. For example, Internet transit service provides what is effectively, in TDM terms, global transport and termination. The price of USD 2 to USD 3 per megabit per month therefore includes a traffic-weighted average of the call-originator‗s transport costs to all the possible destinations around the world, and of the costs of terminating on local access networks in each country, some of which have very strong positions in their local markets. Stated in terms of an equivalent per-minute price for voice-over-IP (VoIP) traffic, the combined cost to caller and recipient is less than USD 0.0000008 per minute, five orders of magnitude less expensive than wholesale services providing comparable functions in TDM markets. This improvement has been achieved with little or no intervention by regulatory authorities. It is to be hoped that, as the Internet market continues to grow and new generations of equipment become more efficient, still lower unit prices are forthcoming. In any case, the performance already achieved is far better, by any measure, than anything observed in the market for circuit-switched traffic exchange.

In most countries the exchange of circuit-switched traffic has been closely regulated by national authorities. In some submarkets, such as among fixed networks in Europe, this has produced modest prices (although still far higher than those for IP exchanges). Still, TDM markets exhibit many anomalies and distortions. For example, if a customer in the United States were to call a user in France using Skype, not including the underlying Internet access and the recipient‗s costs, the call would be free if it terminated to a computer, USD 0.023 per minute if it terminated to a land line, and USD 0.209 per minute if it terminated to a mobile phone.[98] It is worthy of note that the cost to the *recipient* of a TDM call is likewise many orders of magnitude higher than to one who receives a call via the Internet; while the originator and recipient of an Internet call share a very small cost equally, the recipient of a TDM call pays a small fraction of a comparatively huge cost, whether by minute or through their monthly minimum flat-rate charges. The result is that TDM call recipients also pay vastly more than Internet call recipients. National regulatory authorities in OECD countries have sought to improve the results of the TDM traffic exchange market that have produced these anomalies. In the United States, for example, the Federal Communication Commission (FCC) recently launched a proceeding to reform its inter-compensation framework for TDM traffic.[99] European regulators have gradually pushed mobile termination rates down. At the same time, market participants have spent substantial resources in litigation and advocacy to influence regulation. In contrast, the majority of peering agreements in the Internet are carried out on a handshake basis, without even a written agreement (see Annex 1).

More substantially, the anomalies in TDM prices and access arrangements have influenced the behaviour of consumers and providers alike. Consumers have engaged in elaborate call-back procedures to avoid international settlement charges, turned off their mobile phones to avoid paying for received calls, and established multiple wireless accounts in different countries—with different numbers and handsets—to avoid roaming charges and inter-carrier surcharges. Carriers, too, have often based their business decisions on the details of tariff and settlement schemes rather than on underlying economic factors. Extreme examples include the use of mobile gateways in France prior to 2004 to reoriginate fixed-to-mobile calls as mobile-to-mobile calls.[100] In the United States, the FCC has received complaints about ―phantom traffic" designed to qualify for a lower inter-carrier rate by concealing the true end points of the call. ―Traffic pumping" schemes attract calls to small rural incumbents in the United States by offering consumers inducements such as free conference calls. The consumer's long-distance provider is then obliged to pay high access charges to the incumbent, who splits the proceeds with the conference call provider. Each of these abuses has typically been addressed and resolved through further regulatory action. Notably, no analogous schemes have developed on the Internet, since there is no settlement mechanism to be gamed.

Although some of these extreme examples are outliers with relatively small effects on the market overall, inefficiencies in the wholesale market for traffic exchange can lead to major losses in social welfare. For example, wireless customers in the United States, where inter-carrier payments for wireless traffic are low, use their phones an average of 691 minutes per month. The next highest usage level reported in the OECD is in Israel, with 210 minutes; thirteen OECD countries report average monthly usage of less than 100 minutes.[101] When these figures are adjusted to reflect structural differences across markets, the differences in usage, though narrowed somewhat, are still large. Differences in consumer tastes may explain some variation, but it is reasonable to believe that much of this difference in consumption is attributable to the inter-carrier payment arrangements and their effects on the marginal prices consumers face.

An inefficient market for traffic exchange can inhibit investment or direct it to less productive uses. Because TDM voice traffic is declining in most markets, and new carrier voice equipment has been predominantly VoIP for at least ten years, the question of new investment in circuit-switched equipment is largely moot.[102] But the rapid growth in IP traffic is creating an ongoing need for large investments to expand the capacity of Internet backbones as well as local broadband access networks. So far, the Internet model of traffic exchange has worked effectively to call forth the investment necessary to keep up with the explosive growth in traffic and to direct that investment to the areas where it is most productive; the forward-looking challenge of meeting these investment needs is discussed in the Report, under ―Challenges for the future." It is doubtful, though, that these results could have been achieved had the Internet market shared the pricing models and incentive structures of the TDM market.

This raises the question of what factors have allowed the Internet model of interconnection to produce such favourable outcomes. Some of these factors are:

Efficient packet-switched technology

The cost of handling traffic is lower for packet-switched networks than it has been for circuit-switched ones. This greater efficiency has made lower Internet connectivity prices possible. It is clear, however, that the anomalies noted above in the TDM market are not driven by differences in cost. Further, prices for TDM traffic exchange are higher even though in many cases most of the functions internal to both networks are carried out over IP, and the underlying fibre has always been the same. Thus, the differences in price performance observed between the two markets cannot be entirely explained by technical differences in unit cost; they are more completely explained by differences in the efficiency with which these markets have functioned. Further, market and technical efficiency are interrelated, in that the well-functioning Internet market, driven by a high degree of competition to deploy resources effectively, has played an important role in ensuring that costs are minimised.

Competition in Internet exchange markets

An important factor is the degree of competition. At various points in the development of the Internet over the past fifteen years, there were concerns that consolidation would reduce the effectiveness of competition. First, it was feared that UUNET would establish a dominant position; later, that the ―Tier-1" (backbone) carriers as a group would gain market power. In each case, the market shares of these large players subsequently declined as new and peripheral participants grew faster than those in the self-defined core. Competition in the Internet market has proved to be remarkably resilient. As discussed in the Report, the evolution of the Internet over the last five years has made its structure flatter and broader, and has reduced its dependence on any one player or group.

Maintaining competitive conditions will continue to be a primary policy goal of governments. As the regional survey below shows, market results have improved in most regions. Yet concerns remain in some countries, particularly those where policies have not yet been liberalized sufficiently to open opportunities for new participants. Concentration also remains a concern on some undersea cable routes.

It appears that even those players with relatively strong positions in the legacy TDM world have found it difficult to turn those positions into economic advantage in the exchange of Internet traffic. Wireless carriers who had been able to negotiate highly asymmetric termination charges for TDM voice traffic have not obtained similarly lop-sided terms for Internet traffic exchange. At the same time, national incumbents who have maintained high TDM settlement rates have more often found themselves purchasing transit at retail prices to route their international Internet traffic. Thus, the Internet market benefits not only from a high degree of competition but also from a greater resilience—the ability to obtain good results even where some players have a degree of market power.[103]

Flexibility in arrangements

The traditional regulatory mechanisms governing interconnection of circuit-switched networks were introduced early in the past century. In a world of national incumbent networks, domestic interconnection was usually not an issue, and the main focus was on creating international agreements on interconnection and settlements. In a few countries with regional incumbents, such as Canada and the United States, interconnection among those networks was also a concern, and domestic regimes of regulation were adopted.[104] The need for regulation was driven by the lack of competition, and by the limited routing flexibility of early TDM networks in the routing of traffic. Given those conditions, a refusal by one carrier to enter into a bilateral agreement with another could make it impossible for some subscribers to reach some others. The initial goal of regulation was therefore not so much to constrain prices - international settlement rates, for example, were very high - but to ensure universal connectivity. Decades later, as OECD countries opened their telecommunication sectors to competition, national regulators faced challenges in generalizing a closed system designed for incumbents to allow interconnection for new entrants, a process that is still unresolved in some emerging economies today. More recently, regulators in some OECD countries have adopted regulation to deal with specific outcomes that were deemed unsatisfactory, for example in the case of wireless termination rates in Europe.

As the commercial Internet has expanded, a new and larger universe has been created in which universal connectivity has also been obtained. No broadband user sending an email to another person on the other side of the world pauses to wonder whether her local broadband network is interconnected with the network of the recipient. There have been exceptions, but they generally reflect government action to block communication in some countries rather than failure of the market for traffic exchange. The very definition of the Internet is the ―network of interconnected networks." Whereas constant regulatory supervision has been necessary to ensure this connectivity in the circuit-switched world, it has been achieved largely without intervention in the Internet.

An important factor in this development is the inherent flexibility of routing on packet-switched networks. By design, individual packets within the same communication may take different paths to reach their destination. The Internet comprises more than five thousand networks. Although they are all interconnected, it would be both impractical and uneconomical for each pair of networks to have a bilateral interconnection arrangement. In the survey results reported in Annex 1, responses were received from 4 330 networks, about 86% of the networks on the Internet.[105] They reported 142 222 peering agreements, an average of 32.8 agreements for each of the entities in the survey. The actual distribution is long-tailed, with 62% of the respondents having ten or fewer agreements.[106] It has therefore been possible to assure global connectivity among two billion users by means of a relatively small number of agreements, less than 1% of a full mesh.

Thus, unlike the TDM world a century ago, where both competition and routing options were limited, in today's Internet a refusal to deal does not result in a loss of universal connectivity. Two networks decide whether a bilateral agreement between them will make both better off; the agreement is adopted only if both parties agree that it will. In the absence of an agreement, traffic between subscribers is completed through an alternative route, usually by means of transit arrangements with one or more additional networks.[107] The only question at issue when two networks negotiate is whether a bilateral arrangement will add value that more than justifies its cost to each party, which is itself quite low because of the informal nature of the agreement. The negotiation may affect the terms of trade between the parties but not whether traffic will be exchanged between the subscribers of the two networks. Further, compared to the TDM market, the range of possible outcomes with respect to the terms of trade is closely bounded, because transit is readily available at low prices that are themselves limited by the low alternative cost of peering.[108]

This flexible market has limited the ability of terminating networks—even those that might otherwise have strong market positions—to extract rents. This has been a key factor in the resilience of Internet exchange markets.

The market has also been a mechanism for deciding which interconnection relationships, among the millions of possible ones, should be implemented. It places a value on investment and innovation that increase the utility of a network, in the estimation of other entities on the Internet, and provides a means of monetizing that value through improved terms of trade in interconnection. It has thus created a virtuous circle of incentives for investment, innovation, and growth.

Different players have adopted different strategies in response to these incentives. Larger networks that have made substantial investments seek peering partners who in their view bring similar numbers or value of customers to the exchange. But the cost of this self-limiting strategy has been a loss of market share as smaller players and new arrangements have grown around them at a faster rate. This in large measure explains the collective and individual decline of the market positions of the legacy backbone carriers over the past decade.

At the other extreme, small players seek to reduce their transit costs by entering into as many peering agreements as possible, taking advantage of their ability to keep transaction costs to a minimum. In the PCH survey reported in Annex 1, 99.51% of the peering agreements reported were on a handshake basis, further reducing the cost of establishing and maintaining relationships.

The incentives established by the Internet market have also motivated structural innovations that have rearranged the market over the past few years, reducing costs, eroding established positions, and resetting incentives over time. For example, investments in IXPs have reduced costs and attracted network investments to more regions around the world. The availability of IXPs has in turn facilitated further structural developments, such as the rapid growth of multilateral agreements shown in the survey results, which have further reduced transaction costs. Investments by CDNs have improved quality, particularly in developing regions, but have also created efficient and cost-effective options for routing traffic.

Minimal regulation of Internet traffic exchange

In the history of circuit-switched networks, regulation has often played a necessary role in ensuring interconnection, controlling market power, and promoting competition. On the other hand, regulation can often have unintended consequences. As previous OECD work in this area has recognized, ―Using regulation to intervene in Internet interconnection may well distort a market outcome which is currently delivering greater provider and network diversity." Though regulation of traffic exchange can be a tool to constrain market power, it can also confer market power on entities that would not otherwise have it. In addition to the factors discussed above—the level of competition, the flexibility of IP routing arrangements—the Internet market for traffic exchange has been able to produce favourable outcomes while it has been relatively free of regulation. In particular, because interconnection agreements have been entered into voluntarily, Internet traffic exchange, unlike that of TDM traffic, has not been subject to an obligation to interconnect.

One analyst, Kevin Werbach, without addressing the inherent cost-benefit trade-off, has suggested that governments create a general obligation for Internet transit networks to interconnect, leaving specific terms and conditions to be negotiated. In a more limited recommendation, Packet Clearing House, in its work with Internet exchange points, notes that a regulatory best-practice is to merely require that any two carriers operating in a country ―ensure that traffic between any two endpoints within the country not cross the national border under normal operation." Given that universal endpoint connectivity is a policy goal, at first glance a requirement to interconnect would appear natural. There are two reasons why a general obligation as suggested by Werbach is problematic: First, it ignores the expense of constructing a full interconnection mesh. The Internet functions with a high degree of redundancy and resilience today, utilizing a surprisingly sparse interconnection mesh, with low costs as a result: Second, there is good reason to believe that imposing a universal obligation to interconnect would interfere with the successful operation of the market.

An essential factor in the success of any market is the ability of participants to accept or decline proposed transactions. The most basic defense that any market participant has against unreasonable terms offered by another participant is to refuse to enter into the proposed transaction. This is true whether the participant is a consumer faced with unattractive offers or a firm faced with unreasonable terms proposed by a supplier or customer. Two factors could deprive market participants of the ability to make such choices: monopoly power in the provision of an essential good or service, or a legal or regulatory requirement to deal with the other party. In fact, the ability of ISPs to refuse to interconnect with any cohort who fails to adhere to commonly accepted norms and practices has been the primary mechanism by which the market enforces reasonable terms, regulates its own quality, and limits abuses like spam, malware, and cybercrime.

This is the paradox presented by interconnection requirements. A refusal to deal with another entity can, in some circumstances, be an anticompetitive exercise of market power. In that case, a requirement to deal can provide a remedy. This has been the logical basis for the existing regulatory frameworks that govern TDM traffic exchange in OECD member countries today. On the other hand, the ability to refuse unreasonable terms is the essential mechanism through which markets discipline parties' behaviour and impose reasonable outcomes. In a study for the European Commission, Scott Marcus states, ―We do not advocate an interconnection obligation as regards IP data traffic in general, and we do not see a need to mandate any-to-any peering; however, National Regulatory Authorities must be able to intervene if interconnection breaks down, especially where this is a manifestation of some form of market power. To be effective, this power to intervene must include the ability to cap the price of IP data interconnection."

A case in the United States illustrates the risk that an obligation to interconnect may confer market power, even upon entities that would not otherwise have it. In the late 1990s AT&T petitioned the FCC for clarification of its obligations to interconnect—for TDM traffic—under United States law. AT&T, acting in this case as a long-distance provider, was attempting to negotiate an agreement with a competitive local provider (CLEC), which proposed terms that AT&T found unreasonable, including a high per-minute access rate. If AT&T declined these terms, it would not have an interconnection agreement with a carrier that had requested one. Was this permissible under United States law pertaining to TDM traffic exchange? After careful review, the FCC concluded in 2001 that AT&T was obliged to enter into an agreement, upon request, with the competitor. In the same order, the FCC concluded that this obligation would make it impossible for the market to function adequately without intervention. The FCC therefore decided, reluctantly, to impose regulation on the access rates charged by CLECs, carriers who otherwise lacked market power and were not subject to any other form of rate regulation.

In the current Internet market, if two networks do not agree on terms, then they simply do not enter into an agreement. But as the FCC example illustrates, if they must have an agreement, then there must also be some regulatory mechanism to resolve disagreements and establish terms. This may take the form of a tariff-setting mechanism, compulsory arbitration, or, in the FCC case cited above, comparison to a benchmark. Thus, an interconnection obligation does not provide the parsimonious solution that Werbach sought. The regulatory authority is likely instead to be drawn into the regulation of rates and terms to which Marcus refers. Even if this is done indirectly, such as through arbitration, the effect will be to overhang the process of negotiation between parties. The reaction function of each party will be based in large part on its expectation of the outcome if the matter goes to arbitration. Once enough precedent has been established to allow such expectations to converge, agreements tend to reflect terms established by the regulator rather than those that would be set by market competition.

Regulatory intervention in the Internet market, therefore, carries risks. It may involve the regulatory authority in recurring disputes and litigation. It would substitute the rates and terms produced by the regulatory mechanism for those determined by the market. By establishing property rights for existing players, it could actually confer market power. To the extent that it is built around existing structures and relationships, it may inhibit evolution of new approaches, a process that has been going on more or less continuously since the commercial Internet began. It is exactly this morass of regulation and counter-regulation, in which parties seek to saddle their competitors with externalities rather than focus upon innovation and growth, that limited the TDM market while the Internet surpassed it. Further, experience has shown that once regulation is adopted, it tends to be extremely persistent, and is very difficult to reform or remove.

The anomalies noted above in the TDM market have, to a large extent, been created by parties seeking to gain advantage by exploiting restricted competition and specific aspects of the regulatory structure which were undoubtedly originally intended to mitigate that lack of competition. The Internet market is largely free of this kind of behaviour because competition is relatively unrestricted, and every agreement is entered into voluntarily. If a party to a peering agreement believes that it is being disadvantaged by the behaviour of the other party, it can simply end the arrangement and deliver traffic by other means. If a policy structure were based on an obligation to maintain such arrangements, then this discipline would be lost.

Today‘s Internet market based on commercial agreements allows underlying economic factors to determine which subset of the millions of possible direct interconnection arrangements should actually be implemented. This market outcome has changed substantially over time through a dynamic process of evolution. If a regulatory obligation to interconnect were to be established, then this choice could be imposed unilaterally by one of the parties, which would dramatically change the dynamics of the market. It is extremely unlikely that the efficient outcomes produced so far would be maintained in this environment.

On the contrary, it is plausible to expect a form of adverse selection, since the greatest incentive to employ this right to demand interconnection would be in those instances where the shift in the terms of trade between the market outcome and the regulated one would be the greatest.

More generally, introducing the structures and incentives of the TDM market to the Internet environment would undermine the mechanisms that have allowed the market to call forth the substantial investments necessary to keep up with the dramatic growth of the Internet and to direct those resources where they can do the most good.

All of these reasons suggest that governments should approach any call for increased regulation of the Internet market with great caution, and with an appreciation for the risks involved. This does not mean that governments should accept uncompetitive outcomes, and in general this has not been necessary, since market performance has so far been very good.

ENDNOTES

1 OECD, Internet traffic exchange and the development of end-to-end international telecommunication competition, 2001, DSTI/ICCP/TISP(2001)5/Final

2 OECD, Internet traffic exchange and the development of end-to-end international telecommunication competition, 2001, DSTI/ICCP/TISP(2001)5/Final

3 See for instance Cisco Visual Networking Index 2010, referenced in an interview with the researchers *www.itbusinessedge.com/cm/community/features/interviews/blog/ciscos-vni-video-unseats-p2p-as-the-growth-king/?cs=44027*

4 OECD, Internet Addressing—Measuring Deployment of IPv6, April 2010, DSTI/ICCP/CISP(2009)17/FINAL

5 See, for example, "Internet Backbone Interconnection Agreements," Discussion Paper, International Chamber of Commerce, *www.iccwbo.org/uploadedFiles/ICC/policy/e-business/Statements/ITIS_IBIA.pdf*. See also closing ceremony speech by Herbert Heitmann, Chair of the ICC Commission on E-Business, IT and Telecoms, at the meeting of the Internet Governance Forum, Nairobi, 30 September 2011. *www.iccwbo.org/basis/index.html?id=45908*

6 Internet Exchange Point Directory: *http://pch.net/ixpdir*

7 OECD, Internet traffic prioritisation: an overview, 2006, DSTI/ICCP/TISP(2006)4/FINAL

8 Cisco Visual Networking Index Usage Study 2010: *www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/Cisco_VNI_Usage_WP.pdf.*

9 Internet traffic growth isn't doubling every year as it did at the start of the 21st century, however it is still growing briskly at about 40% per year according to Cisco's Visual Networking Index 2012.

10 OECD, Internet Addressing—Measuring Deployment of IPv6, April 2010, DSTI/ICCP/CISP(2009)17/FINAL.

11 AfriNIC policy IPv6 Provider Independent Assignment for End-Sites, *www.afrinic.net/docs/policies/AFPUB-2007-v6-001.htm*.

12 See also Henk Steenman and Ted Seely, Carrier Hurdles to Meeting 10GE Demand, at Page 13, *www.ieee802.org/3/hssg/public/mar07/seely_01_0307.pdf*.

13 As this report was being written, the equipment vendor Extreme Networks debuted a new switch, the ―BlackDiamond X Series," offering an unprecedented 768 ports of 10 gbps each. Seventeen such switches would hypothetically allow the Amsterdam exchange to offer their 411 participants 1.9 tbps of LAG-bundled bandwidth each, for example, albeit at astronomical cost to both the exchange and the participants. *www.extremenetworks.com/products/blackdiamond-x.aspx*.

14 Ron Avitzur, Lisa Lippincott, and Mat Marcus, personal correspondence with Bill Woodcock, April and May 2011.

15 Mobile carriers in the United States generally exchange TDM traffic with one another on the basis of voluntary bill-and-keep agreements. The rate for exchange with larger landline incumbents, such as AT&T

and Verizon, is symmetric and very low (USD 0.0007 per minute). The only significant net flow of funds is in the form of payments to small rural telecommunication operators, who still charge much higher rates.

16 In the United States, the access charges of these smaller incumbents are subject to regulation by the FCC and by state regulatory commissions. Although some of these are under price caps, many are set by a rate-of-return mechanism. In particular, the interstate charges are reset more or less automatically on a regular basis. As traffic declines, embedded costs do not decline in proportion. Therefore, there is a risk of a kind of ―death spiral," in which falling demand leads to rising prices, which further encourage avoidance by other parties.

17 OECD, Internet traffic exchange and the development of end-to-end international telecommunication competition, 2001, DSTI/ICCP/TISP(2001)5/Final, at Page 31.

18 OECD, Good Practices in Internet Exchange Point Documentation and Measurement, DSTI/ICCP/CISP(2007)9/FINAL.

19 Commission recommendation 2007/879/EC of 17 December 2007 on relevant product and services.markets, OJ L 344 (28/12/2007), pp.65-69. The three criteria are a) the presence of high and non-transitory barriers to entry. These may be of a structural, legal or regulatory nature; b) a market structure which does not tend towards effective competition within the relevant time horizon. The application of this criterion involves examining the state of competition behind the barriers to entry; and c) the insufficiency of competition law alone to adequately address the market failure(s) concerned.

20 See *http://europa.eu/rapid/pressReleasesAction.do?reference=IP/10/240&format=HTML&aged=0&language=EN&guiLanguage=en*.

Commission letter at *http://circa.europa.eu/Public/irc/infso/ecctf/library?l=/commissionsdecisions&vm=detailed&sb=Title*.

In 2006, UKE adopted the following obligations on TP: (i) nondiscrimination, to prevent TP from interfering with the quality and other parameters of IP transmission between TP‘s end users and other electronic communication operators; and (ii) a transparency obligation, consisting of making available on request the technical information relating to the configuration of equipment used for IP data transmission by TP.

In 2007, UKE added the following obligations: (i) access, regardless of the ratio of the exchanged traffic, in order to ensure communication between end users of TP‘s network and users of another operator (if exchange of IP traffic met certain criteria, TP would have to provide access and peering for free); (ii) non-discrimination, ensuring the same conditions of peering between TP‘s network and the network of other operators in order to ensure communication between end users of TP‘s network and the users of another operator‘s network, in particular by offering the same conditions in comparable circumstances, as well as offering services on conditions not worse than those used within its own undertaking or in relations to subsidiaries; (iii) an obligation to provide information on peering; (v) an obligation to prepare and submit a Reference offer related to peering; and (iv) an obligation to set prices based on costs incurred. UKE explained that it can assess whether TP‘s prices are based on costs incurred through a comparison with prices of other operators present on the Polish market or with prices of operators transferring IP traffic internationally, and that it may determine in a separate measure the level of prices, or their maximal or minimal level.

21 See appeal by UKE at *http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:C:2010:209:0041:0042:EN:PDF*

22 *http://europa.eu/rapid/pressReleasesAction.do?reference=IP/10/240&format=HTML&aged=0&language=EN&guiLanguage=en*.

23 See Letter from Robert W. Quinn Jr., dated December 28, 2006, Notice of Ex Parte Communication, In the Matter of Review of AT&T Inc. and BellSouth Corp. Application for Consent to Transfer of Control, WC Docket No. 06-74.

24 See Report and Order and Further Notice of Proposed Rulemaking, FCC 11-161, Adopted October 27, 2011. Released November 18, 2011. (FCC Universal Service and Access Reform Order) See also Comments filed February 24, 2012, and Reply Comments filed March 30, 2012.

25 See ―Network interconnection for voice services," File number 8643-C12-201105297, 19 January 2012. (CRTC 2012)

26 Under United States law, for example, the traffic may be exchanged at ―any technically feasible point" in the incumbent's network.

27 See, for example, ERG, Final report on IP Interconnection, *www.cmt.es/es/publicaciones/anexos/ERG(07)09_rept_on_ip_interconn.pdf*, and German Federal Network Agency, Key Elements for the Interconnection of IP-Based Networks, *www.bundesnetzagentur.de/SharedDocs/Downloads/EN/BNetzA/Areas/Telecommunications/TelecomRegulation/IPinterconnection/KeyElementsId14810pdf.pdf?__blob=publicationFile*.

28 ―Network interconnection for voice services," File number 8643-C12-201105297, 19 January 2012 at paras. 51-53.

29 *www.accc.gov.au/content/index.phtml/itemId/975609/fromItemId/142*.

30 See *www.pc.gov.au/__data/assets/pdf_file/0016/55222/intelmkt.pdf*. See also *http://drpeering.net/white-detail in a 2010 CISP detail in a 2010 CISPs/Ecosystems/Australia-Peering-Ecosystem.php*.

31 *www.reuters.com/article/2011/04/27/france-g8-internet-idUSLDE73Q1VU20110427* (Reuters). French President Nicolas Sarkzoy said on Wednesday he would invite Internet business leaders to an end-May G8 summit to discuss regulation of the web, and he called for more scope for governments to tax Internet firms.

32 ―Tromboning" is the expensive and inefficient practice of moving traffic between two adjacent points through a distant one. It most frequently occurs when the management cost of implementing the more efficient direct route outweighs the operational of doing so. For a discussion of the effect of regional IXP deployment on tromboning, see Michael Kende, "Overview of recent changes in the IP interconnection ecosystem," Analysys Mason, 23 January 2011.

33 *www.Renesys.com/blog/2011/04/qwavvis-the-battle-for-second.shtml*. Note that percentages on the chart sum to more than 100, because of multi-homing.

34 *www.level3.com/en/About-Us/Newsroom/Press-Release-Archive/2011/2011-04-11-globalcrossing.aspx*.

35 Indeed, the Level 3 – Global Crossing merger was approved by the FCC in 2011. In approving the Order, the FCC found that the backbone market is competitive, so much so, that they noted an increase in the number of Tier 1 backbones since a prior analysis in 2005.

36 Bill Krogfoss, Marcus Weldon, and Lev Sofman, Internet Architecture Evolution and the Complex Economies of Content Peering, Alcatel-Lucent 2011. (Kogfoss 2011).

37 Labovitz, et al, Atlas Internet Observatory, 2009 Annual Report at Page 15. *www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf*. The study was a joint project of Arbor Networks, The University of Michigan, and Merit Network.

38 Welbush Equity Research, September 2010.

39 Labovitz, et al, Atlas Internet Observatory, 2009 Annual Report, at Page 13.

40 *http://blog.nielsen.com/nielsenwire/online_mobile/u-s-teen-mobile-report-calling-yesterday-texting-today-using-apps-tomorrow/.*

41 Labovitz, et al, Atlas Internet Observatory, 2009 Annual Report at Pages 22-24.

42 *www.pandonetworks.com/p4p*. The group‘s work is based on research by authors at Yale and the University of Washington. See Haiyong Xie, Arvinf Krishnamurthy, Avi Siberschatz, and Richard Yang, P4P: Explicit Communications for Cooperative Control Between P2P and Network Providers.

43 Labovitz, et al, Atlas Internet Observatory, 2009 Annual Report at Pages 22–24.

44 Sandvine Fall 2010 Global Internet Phenomena Report at 9-10. *www.sandvine.com/downloads/documents/2010%20Global%20Internet%20Phenomena%20Report.pdf* Note that while streaming video usage is concentrated at peak hours, P2P traffic is highest in off-peak periods.

45 Labovitz, et al, Atlas Internet Observatory, 2009 Annual Report at Page 19.

46 A viable future model for the Internet, AT Kearney, 2011 *www.atkearney.com/index.php/Publications/a-viable-future-model-for-the-internet.html.*

47 Marcus, et al, The Future of IP interconnection: Technical, Economic, and Public Policy Aspects. WIK-Consult, 2008, prepared for the European Commission (WIK 2008), at Pages 114–120.

48 WIK (2008) recommends that no general obligation to interconnect be established but that the regulator should monitor the market to see if the largest network in a country maintains a reasonable number of traffic exchange agreements with other networks.

49 Marcus, et al, The Future of IP interconnection: Technical, Economic, and Public Policy Aspects. WIK-Consult, 2008, prepared for the European Commission (WIK 2008),, at Page 131.

50 In its recent decision on interconnection policy, the Canadian regulator, the CRTC, specifically affirmed the ongoing right of interconnecting parties to demand TDM interfaces. However, it also allowed such arrangements to be made under commercial negotiations where parties agree. See CRTC 2012 at paras. 78-79.

51 Patrik Fältström, RFC 2916, E.164 number and DNS, *www.ietf.org/rfc/rfc2916.txt*, and Carpenter, Baker, and Roberts, RFC 2860, Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority, *www.ietf.org/rfc/rfc2860.txt*.

52 For an example of a recent expression of concern by European competition law authorities over standard-setting activities in the telecom industry, see NY Times Global Edition, ―European Antitrust Regulators Question Phone Operators" March 14, 2012. *www.nytimes.com/2012/03/15/technology/european-antitrust-regulators-question-phone-operators.html*

53 See FCC Universal Service and Access Reform Order.

54 WSJ, May 5, 2011. ―Rare Headwinds in Mexico Buffet Slim," *http://online.wsj.com/article/SB10001424052748703730804576311352951033930.html?KEYWORDS=carlos+slim*, and May 5, 2011, ―Mexico Court Deals Carlos Slim a Setback on Phone Fees," *http://online.wsj.com/article/SB10001424052748703730804576311352951033930.html?KEYWORDS=carlos+slim*.

55 A viable future model for the Internet, AT Kearney, 2011 *www.atkearney.com/index.php/Publications/a-viable-future-model-for-the-internet.html.*

56 For a general discussion of two-sided markets, see Jean-Claude Rochet and Jean Tirole, ―Two-sided Markets: A Progress Report," November 29, 2005 (Rochet and Tirole 2005) *http://idei.fr/doc/wp/2005/2sided_markets.pdf* Defining what makes a market two-sided is actually difficult, and several alternative definitions are reviewed by Rochet and Tirole. One such definition is that a market is two-sided if the total volume of demand handled by the platform is sensitive not only to the total volume of the fees charged, but also to the structure of those fees. A related definition is that the market is two-sided if the Coase Theorem does not hold. A non-technical discussion of two-sided markets and strategies that apply to them is provided by Thomas Eisenmann, Geoffrey Parker, and Marshall W. Van Alstyne, ―Strategies for Two-Sided Markets," Harvard Business Review, October 2006.

57 Benjamin E. Hermalin and Michael J. Katz, ―The Economics of Product-Line Restrictions With an Application to the Network Neutrality Debate," AEI-Brookings Joint Center for Regulatory Studies, Working Paper 07-02, February 2007. Hermalin and Katz analyze a net neutrality requirement as a product line restriction which limits the ability of a two-sided network platform to offer a range of different quality levels. See also Mark Armstrong, ―Competition in Two-sided Markets," RAND Journal of Economics, Vol. 37, No. 3, Autumn 2006. *http://eprints.ucl.ac.uk/4324/1/competitionintwosidedmarkets.pdf*. The issue of platform selection in a market with multiple two-sided platforms is addressed in Benjamin E. Hermalin and Michael L. Katz, ―Your Network or Mine? The Economics of Routing Rules," February 2004. *http://faculty.haas.berkeley.edu/katz/Your%20Network%20or%20Mine%20posted.pdf*

58 Krogfoss (2011). It is not clear what the financial effects of an agreement with a CDN would be on an access network. The authors present a numerical example in which the effect is negative, but this is not a general proof that this will always be the case.

59 A viable future model for the Internet, AT Kearney, 2011 *www.atkearney.com/index.php/Publications/a-viable-future-model-for-the-internet.html.*, at Page 25.

60 Marcus, et al, The Future of IP interconnection: Technical, Economic, and Public Policy Aspects. WIK-Consult, 2008, prepared for the European Commission at Page 26.

61 See Dennis Weller, ―the Internet Market for Quality," Communications & Strategies, November 2011. See also David Clark, William Lehr, and Steven Bauer, ―Interconnection in the Internet:the Policy Challenge,"Massachusetts Institute of Technology, August 9, 2011, at page 15. Presented at the Telecommunications Policy Research Conference, September 2011. www.tprcweb.com (Clark et al 2011)

62 *www.cedmagazine.com/articles/2011/03/tier-1-party.aspx.*

63 *http://seekingalpha.com/article/248775-take-away-from-fcc-ruling-on-comcast-level-3-dispute-nothing-has-changed.*

64 NPD Group: *Six Out Of 10 Digital Movies Are Streamed via Netflix*, *www.npd.com/press/releases/press_110315b.html*.

65 See Wired, February 27, 2011. ―Peers or Not? Comcast and Level3 Slug It Out at FCC's Doorstep," *http://online.wsj.com/article/BT-CO-20110216-718576.html.*

66 See WSJ, ―FCC Chairman: Net Neutrality Rules Don‘t Cover Comcast-Level 3 dispute." February 16, 2011. *http://online.wsj.com/article/BT-CO-20110216-718576.html*.

67 At about the same time as the dispute between Level3 and Comcast, another dispute was developing in France, between Cogent and Orange. Like the FCC, ARCEP, the French NRA, has declined to intervene in this case. However, in August, 2011, Cogent filed a complaint with the French competition authority. See *www.telecompaper.com/news/cogent-files-competition-complaint-against-orange-france*. See also *www.latribune.fr/technos-medias/internet/20110829trib000645165/orange-veut-nous-faire-payer-pour-atteindre-ses-clients.html*

68 Bill Krogfoss, Marcus Weldon, and Lev Sofman, Internet Architecture Evolution and the Complex Economies of Content Peering, Alcatel-Lucent 2011.at Page 5.

69 Iran Times, Iran Plans Islamic Internet, *www.iran-times.com/english/index.php?option=com_content&view=article&id=1690:iran-plans-islamic-internet&catid=98:whats-left&Itemid=425*, April 21, 2011.

70 Oman Daily Observer, Omantel to participate in new international gateway system, March 29, 2011, *http://main.omanobserver.om/node/45486*.

71 Source: Packet Clearing House Internet exchange point statistical summary, *https://prefix.pch.net/applications/ixpdir/summary/growth*.

72 Earl Zmijewski, *A Closer Look at the 'Level 3 + Global Crossing' Union,* April 2011, *www.renesys.com/blog/2011/04/level-crossing.shtml*

73 ACCC Advice to Government National Broadband Network Points of Interconnect - Public Version, November 2010, pages 16-17

74 Enron India web site: *www.bspl.co.in*.

75 The statistical data for Israel are supplied by and under the responsibility of the relevant Israeli authorities. The use of such data by the OECD is without prejudice to the status of the Golan Heights, East Jerusalem and Israeli settlements in the West Bank under the terms of international law.

76 This report prepared by Bill Woodcock and Vijay Adhikari for Packet Clearing House, San Francisco, May 2, 2011, and repeated here with stylistic revisions by permission of the publisher. It and future versions may be found at *http://pch.net/resources/papers/peering-survey*.

77 Smith, Philip, Weekly Routing Table Report, April 15, 2011, Transit ASes: *http://thyme.apnic.net/rv3-data/2011/04/15/mail-global*.

78 For a discussion of standard symmetric peering terms and conditions, read Chris Hall‘s *www.highwayman.com/peering/peering_agreement.html*. Although much more long-winded, the London Internet Exchange‘s model peering agreement also encapsulates the generally accepted terms of a symmetric peering agreement: *https://www.linx.net/good/bcp/peeringagreement_draftv4.html*.

79 For a global schedule of Internet governance meetings, including many peering forums, see *http://internetmeetings.org*. For specific examples, see the Global Peering Forum website, *http://peeringforum.net* or the European Peering Forum website, *www.peering-forum.edu*.

80 A discussion of MWEB, a South African ISP, transitioning from paid peering to normal peering can be read at *http://mybroadband.co.za/news/broadband/16313-MWEB-peering-link-cuts-How-impacts-you.html*. Specific solicitations of paid peering can be found on the websites of the AOL Transit Data

Network, *www.atdn.net/paid_peering.shtml*; Cox Communications, *www.cox.com/peering/paid-peering.asp*; and Verizon Business, *http://www22.verizon.com/wholesale/productguide/partnerportprogram*.

81 Bill Norton discusses the barriers to entry often contained in ―minimum peering requirements" in his Study of 28 Peering Policies: *http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html.* Original documents can be found on the websites of Comcast, *www.comcast.com/peering*; Tiscali, *www.as3257.net/peering-policy*; AT&T *www.corp.att.com/peering*; and Internet Solutions, *ftp://ftp.is.co.za/tech/peering.pdf*.

82 Definitions and discussions of peering and its general terms can be found on the Packet Clearing House website *https://www.pch.net/wiki/pch:public:glossary#p*; Wikipedia, *http://en.wikipedia.org/wiki/Peering*; and Bill Norton's website, *http://drpeering.net/white-papers/Ecosystems/Internet-Peering.html*.

83 An ―Autonomous System" (AS) is a uniquely-identified Internet network. Autonomous System Numbers (ASNs) are the numeric identifiers assigned by the Regional Internet Registries (RIRs) and used within the Internet routing system to define a specific bounded network that has its own uniquely defined routing policies. An AS-pair is a pair of networks that interconnect with each other.

84 ―Donut peering" is the practice of small and medium-size networks peering with each other aggressively in order to reduce the detrimental impact of a larger network refusing to peer with them. This results in a ―donut" of densely interconnected networks surrounding a self-proclaimed ―Tier-1" network—the ―donut hole" that is poorly interconnected with the networks around it. For a further discussion of donut peering, see the Cook Report's November 2002 Economics of IP Network Interconnection, *www.cookreport.com/backissues/nov-dec2002cookrep.pdf*; or Bill Woodcock's January 2003 lecture to the University of Minnesota Digital Technology Center, Internet Topology and Economics: How Supply and Demand Influence the Changing Shape of the Global Network, *www.pch.net/resources/papers/topology-and-economics*. ―Tier-1" is the name of a mid-1990s attempted cartel, in which a group of carriers peered with each other but nominally refused to peer with any networks outside the cartel. Their misunderstanding of Internet growth rates led them to become irrelevant, as the portion of the market held outside the cartel grew exponentially while that inside the cartel grew in linear fashion. For a more detailed explanation, see the box entitled ―Confusing Terms of Art" in the main document.

85 A range of typical multilateral peering agreements can be found on the websites of the Open Peering Initiative *www.openpeering.nl/mlparegistry.shtml*; the Kansas City Network Access Point www.kcnap.net/peering-policy.html; Red Bus Internet Exchange *www.rbiex.net/assets/joining/mlpa.pdf;* and the Indonesia Internet Exchange *www.iix.net.id/library/iix-peering-agreement_ind.pdf*. Note that their specific terms differ little if at all from those of the bilateral agreements discussed in note iii.

86 Packet Clearing House, How Big Networks Can Peer Regionally, January 2007, *www.apricot.net/apricot2007/presentation/tutorial/scg-regional-peering.pdf*, and Barry Greene, L2 Internet eXchange Point (IXP) using a BGP Route Reflector, December 2000, *www.gixa.org.gh/pdf/L2_Route_Reflector_IXP_v0.4.pdf*, at Pages 35–38.

87 Nick Carr, *Avatars consume as much electricity as Brazilians,* December 5, 2006: *www.roughtype.com/archives/2006/12/avatars_consume.php*.

88 Alex Wissner-Gross, *How you can help reduce the footprint of the Web,* January 11, 2009: *www.timesonline.co.uk/tol/news/environment/article5488934.ece*.

89 See Michael Barbaro and Tom Zeller, *A Face Is Exposed for AOL Searcher* No. 4417749, the New York Times, August 9, 2006*, www.nytimes.com/2006/08/09/technology/09aol.html*; and Arvind Narayanan and Vitaly Shmatikov, Robust De-anonymization of Large Sparse Datasets, May 2008, *www.cs.utexas.edu/~shmat/shmat_oak08netflix.pdf*.

90 Packet Clearing House, *Report on Internet Exchange Point Locations*, *https://prefix.pch.net/applications/ixpdir/summary/*.

91 Internet exchange point statistics may be found at *http://pch.net/ixpdir*.

92 1.4tbps of domestic production, as compared with 708gbps of domestic consumption, assuming the Netherlands' 6.4M broadband customers use the same 112kbps average that other European broadband customers consume. *http://oecd.org/sti/ict/broadband* and *http://pch.net/ixpdir/summary/growth*

93 These numbers are approximately those of the Kuala Lumpur Internet Exchange, a relatively typical high-investment exchange in a relatively developed economy.

94 IXPs per region: *https://prefix.pch.net/applications/ixpdir/summary/growth-region/*

95 , *Good Practices in Internet Exchange Point Documentation and Measurement, OECD 2007,* DSTI/ICCP/CISP(2007)9

96 See OECD efforts in this area: OECD resources on Internet addressing: IPv4 and IPv6 *www.oecd.org/document/14/0,3746,en_2649_34223_44954318_1_1_1_1,00.html* and IETF RFC 1883, *Internet Protocol version 6 Specification*, December 1995.

97 IETF RFC 2050, *Internet Registry IP Allocation Guidelines*, November 1996.

98 See *www.skype.com/intl/en-us/prices/payg-rates/#cc=FR*. This example illustrates the effect of asymmetric termination rates charged by mobile carriers in Europe. It also shows that, as mobile termination rates have fallen substantially in recent years, Skype may not have passed along those reductions in their retail rates. One factor may be that reductions in some countries, although implemented, are still subject to legal proceedings. Contracts for transit and termination may require, if challenges to reductions are successful, that payments are backdated. In these situations, some ―over-the-top" operators hold rates at the previous level until the lower wholesale rates are confirmed.

99 FCC 11-13, Notice of Proposed Rulemaking and Further Notice of Proposed Rulemaking, Adopted February 8, 2011.

100 See Loutrel, Benoit, ―Bill and Keep in the French Mobile Industry: A Case Study", WIK Conference, April 4–5, 2006. *www.wik-consult.com/fileadmin/Konferenzbeitraege/2006/Bill_and_Keep/Loutrel.pdf*.

101 OECD Communications Outlook 2011. Structural differences across markets may include the way originating and terminating minutes are counted. Although penetration (the percentage of people having a mobile service) is comparable between the United States and Europe, the number of accounts per 100 population is much higher in Europe, reflecting a higher propensity to maintain multiple accounts.

102 See, for example, ―Top 10 Dying Industries," Wall Street Journal, March 28, 2011, citing analysis by IBIS World *at www.ibisworld.com/Common/MediaCenter/Dying%20Industries.pdf*. Wireline telecom carriers are at the top of the list, followed by record stores and film processing.

103 The term resilience is used here to mean the ability of the Internet market to achieve good results even where some market imperfections are present. Of course, the decentralized nature of the Internet has also made it more resilient with respect to physical disruption by natural disasters or other causes. See Inter -X: Resilience of the Internet Interconnection Ecosystem, European Network and Information Security Agency (ENISA), April 2011. *www.enisa.europa.eu/act/res/other-areas/inter-x/report/interx-report*.

104 The regulatory framework in the United Sates was established by an agreed settlement in a competition law case. The main concern in this case had been the possibility that AT&T, the largest incumbent, might

refuse interconnection with smaller regional incumbents. See *books.google.com/books?id=GrEZQHQ5-ncC&pg=PA130&lpg=PA130&dq=kingsbury+commitment+of+1913&source=bl&ots=-btNFC32T_&sig=uBwQjzSr7S6chOEbhBMFAjhS6ow&hl=en&ei=QXzQTa65GMLTgAeq_-TADA&sa=X&oi=book_result&ct=result&resnum=1&ved=0CBoQ6AEwADgK#v=onepage&q=kingsbury%20commitment%20of%201913&f=false*

105 For a detailed description of the survey, see Annex 1. The respondents include most of the largest Internet backbone networks and some of the largest CDNs, so it is reasonable to say that they represent some approximation of 86% of Internet traffic.

106 Note that some of these agreements are multilateral, so that the average number of entities with whom each network has agreements is somewhat higher. The survey includes only peering agreements, so the total number of agreements per network, including both transit and peering, would be higher.

107 OECD, Internet traffic exchange and the development of end-to-end international telecommunication competition, 2001, DSTI/ICCP/TISP(2001)5/Final, at Page 9: ―This raises the question of what happens if two Autonomous Systems can not agree to exchange traffic. The answer to this question is that traffic will still be exchanged between these networks but, instead of being direct, it will flow through one or more other networks via transit relationships. This is an extremely important point to bear in mind when considering Internet traffic exchange. To continue with the example of AS 15169, Deutsche Telekom, Telmex and many other large networks exchange traffic with Google via other networks. If such players did not believe this met their needs they would negotiate a more direct relationship. At the same time, Deutsche Telekom and Telmex have many other exchange relationships which, in turn, make them attractive to partners who do provide connectivity to networks such as Google."

108 The availability and price of transit depend in part on the progress in opening markets to entry and competition. In fairness, it should also be noted that the market for TDM traffic exchange today has much greater flexibility than it did years ago. Intermediaries such as long-distance carriers effectively provide a function similar to transit, accepting wholesale traffic from mobile carriers or regional fixed providers for delivery to terminating access networks. For domestic traffic within the United States, a typical wholesale rate is about two cents per minute, with roughly one cent for the cost of the long-distance network and the other cent for the average access charge paid to terminate the traffic.

Internet Peering has Obsolesced Interconnection Regulation

The old adage is true here; "if it ain't broke don't fix it." The Internet peering marketplace works exceptionally well and it has for its entire twenty year history. The unparalleled success, growth, and resiliency of the unregulated model for the Internet backbone peering marketplace has been nothing short of phenomenal in enabling and ensuring everyone reasonable access to the Internet.

Summary of What Congress Needs to Know

As Congress considers modernizing communications legislation concerning Internet peering and interconnection issues, it is imperative to understand how Internet peering and voice interconnection are fundamentally different from two key policymaking perspectives.

A. **How Internet networks are completely different from railroad and electricity networks.**
    1. Place-independent vs. place-dependent
    2. Software-dependent vs. hardware dependent
    3. Digital vs. analog
B. **How IP packet networking fundamentally differs from legacy voice interconnection.**
    1. Circuit technology vs. packet-switching technology
    2. Continuous vs. discontinuous transmission
    3. Predictable vs. unpredictable transmission
    4. Unitary service vs. multiple services
    5. Centralized vs. decentralized architectures
    6. Location-driven vs. location-agnostic
    7. Best efforts vs. guarantees
    8. Accounting simplicity vs. complexity
    9. Access vs. connection

Given the reality described below -- that inter-networked computer networks are effectively the opposite of railroad, electricity, and telephone networks -- trying to impose telephone interconnection rules on IP inter-networking is akin to forcing a square peg into a round hole. It predictably breaks both the peg and the hole.

A. How are Internet networks completely different than railroad and electricity networks?

1. Place-independent vs. place-dependent

Railroad and electricity interconnection is *place*-dependent, Internet "interconnection" is not place or physical-location-dependent.

This is a huge difference as physical-place-dependency can create a physical interconnection chokepoint in railroads or electricity. In contrast there are no physical-place-dependent chokepoints for the Internet because one can access/connect to the Internet from many different places, through many different entities, and via many different technologies, e.g. electrically via wires like copper or coax, optically via different fiber configurations, or wirelessly via many different licensed and unlicensed frequencies and providers.

Choice of place, facility, provider, and technology mean multiple dimensions of competition and no lasting chokepoints because if one encounters a temporary congestion problem in one part of the Internet, one has the choice to take their traffic and business elsewhere. No chokepoints mean no need for proscriptive regulation of Internet peering arrangements.

2. Software-dependent vs. hardware dependent

Railroad and electricity interconnection is hardware-dependent, whereas Internet interconnection is *software*-dependent. Railroads and electric networks require one universal physical standard of wheel gauge and axle width, or physical electrical transformers and wall plugs, to interconnect to these respective networks. In contrast, the software design and protocol of Internet connections make interconnection hardware-agnostic, seamless and automatic, and hence inherently competitive and choice-rich.

Simply, the genius of Internet packet-technology networks is that they do not require any interconnection, permission, or negotiation points, because inherent in Internet Protocol is that packets are automatically routed seamlessly between different internet networks to their destination by design. Inherently Internet packet technology makes the concept of telephone interconnection obsolete because the technology supplants what used to require hardware and regulation to achieve. Most simply, Internet protocol innovation inherently obviates an FCC role for regulating Internet peering arrangements.

3. Digital vs. analog

Railroad and electricity interconnection involved *analog* technology, whereas Internet peering must involve *digital* computer technology. Importantly, railroads required a set *continuous* physical path or circuit from point A to point B. Electricity networks require a *continuous* electrical circuit from origin to destination.

In contrast, digital technology in general, and Internet packet technology in particular, is *discontinuous* -- the antithesis of a telephone or electrical *continuous* circuit. It is this inherently *discontinuous* digital innovation that enables IP networks to be place-agnostic and hardware-agnostic, and inherently competitive and choice-rich.

More specifically, the innovation of digital IP packet networks subdivides information into many small packets to enable more efficient transmission. The packets get individually-routed unpredictably and comingled with other packets to minimize bandwidth waste. At the ultimate destination, the packets get immediately reassembled by any device anywhere. Internet Protocol is inherently a competitive technology, inherently made even more competitive by Moore's law, which ensures digital networks continually enjoy rapidly declining computing costs.

B. How does IP packet networking fundamentally differ from legacy voice interconnection?

1. Circuit technology vs. packet-switching technology

Interconnection is a telephone-technology-specific concept. It was created and designed for circuit-switching telephone technology that purposefully produces dedicated electrical circuits connecting specific telephones.

In stark contrast, inter-networking or "peering" is an Internet-technology-specific concept. It was created and designed for packet-switching Internet technology. It purposefully subdivides information into many small packets to enable much more efficient transmission.

2. Continuous vs. discontinuous transmission

Interconnection is all about creating and maintaining a continuous electrical circuit, while inter-networking is all about one computer network facilitating the forwarding of data packets to another computer network based on a discontinuous system of Internet protocol packet delivery.

3. Predictable vs. unpredictable transmission

The predictable transmission path of telephone circuit-switched connections is exact opposite of the unpredictable transmission path of Internet packet delivery.

4. Unitary service vs. multiple services

The PSTN was designed to provide one service -- voice. In stark contrast, the Internet was designed to enable any product or service one could invent. On the Internet voice is an app. Consequently, the Internet can mimic the PSTN easily and cheaply with Voice over IP (VoIP), just like the Internet can mimic broadcasting with video streaming.

5. Centralized vs. decentralized architectures

Telephone technology by design is centrally-organized in a hub-and-spoke manner utilizing central office switches to connect calls. Internet technology by design is decentralized utilizing routers which simultaneously forward commingled packets between computer networks via routing tables and “best efforts.”

6. Best efforts vs. guarantees

Since the predictable and centralized design of a telephone network naturally lends itself to predictable, 99.999% reliability and operation, it is naturally conducive to a centralized management regulatory construct like common carrier interconnection.

However, the unpredictable, “best efforts,” and decentralized Internet design is antithetical and inherently hostile to a centralized management regulatory construct like common carrier interconnection.

7. Location-driven vs. location-agnostic

Interconnection is important in a circuit-switched network because one must connect at the nearest geographic point, which inherently limits competitive interconnection choices. However, inter-networking or peering with other computer networks can largely happen without respect to location so consequently an Internet entity inherently has many competitive choices to inter-network and access the Internet.

8. Accounting simplicity vs. complexity

Lastly, interconnection is an economic concept focused on accounting and billing for originating and terminating voice traffic by seconds or minutes. By design Internet technology does not originate or terminate connections; it breaks information into packets, commingles and forwards them for reassembly on any device anywhere. Accounting for individual end-to-end bit routes would require pervasive deep packet inspection technology that could also be used simultaneously for surveillance.

9. Access vs. connection

One does not “connect” to the Internet, one accesses it. The PSTN “interconnects” different entities based on compulsory, regulated, interconnection-agreements. The Internet inter-networks different entities based on voluntary, unregulated “inter-networking” or “peering” agreements. Simply, circuit-switched technology can need regulation to ensure interconnection, whereas Internet technology does not because the design of the Internet naturally facilitates voluntary and competitive inter-networking or peering access to the Internet in a location-agnostic and technology-agnostic manner.

Mandatory Interconnection: Should the FCC Serve as Internet Traffic Cop?

progressive policy institute ppi

BY HAL J. SINGER

MAY 2014

Since the agreement between Comcast and Netflix was struck in February 2014, several parties have called on the Federal Communications Commission (FCC) to regulate dealings between networks that comprise the Internet generally, and to dictate the terms of interconnection by Internet service providers (ISPs) in particular. This Policy Brief considers the costs and benefits to consumers if the FCC interferes with the terms under which ISPs connect with transit providers, content providers, and others. A key lesson from the economics literature that informs this question is that antitrust enforcement acts as a substitute for sector-specific interconnection obligations in industries that have made sufficient progress along the "deregulatory arc." Because the communications sector was set on a deregulatory path nearly 20 years ago, has the time come to rely on antitrust to adjudicate interconnection disputes on the Internet?

Interconnection agreements between the networks that comprise the Internet have been privately negotiated without a regulatory backstop.

Introduction

To date, interconnection agreements between the networks that comprise the Internet have been privately negotiated without a regulatory backstop.[1] The vast majority of these negotiations have gone down without a hitch. Some notable interconnection disputes in the United States involved Cogent-AOL (2002), Cogent-Level 3 (2005), and Cogent-Sprint (2008).[2] While transit companies such as Cogent and Level 3 have complained about the quality of interconnection with certain Internet service providers (ISPs),[3] consumers have largely been unaffected; rarely does a dispute turn into a prolonged service disruption for customers. Yet the question of the FCC's role in dealings among these "core" networks is front and center inside the Beltway.

The interconnection controversy is playing out as the FCC grapples with new rules to "Protect and Promote an Open Internet,"[4] which are designed to protect "edge" providers such as content providers, application providers, and device makers. In its May 2014 Notice of Proposed Rulemaking, the FCC tried to distinguish interconnection from so-called "net neutrality" issues:

> Separate and apart from this connectivity [to the Internet by the ISP] is the question of interconnection ('peering') between the consumer's network provider and the various networks that deliver to that ISP. That is a different matter that is better addressed separately. Today's proposal is all about what happens on the broadband provider's network and how the consumer's connection to the Internet may not be interfered with or otherwise compromised.[5]

Although the Open Internet proposals are designed to address the management of traffic *within* an ISP's network, the FCC also seeks comment on how it can ensure that an ISP "would not be able to evade [its] open Internet rules by engaging in traffic exchange practices that would be outside the scope of the rules as proposed."[6] The issue is clearly timely and ripe for resolution.

Some scholars have advocated for greater FCC involvement in interconnection disputes. For example, Werbach (2014) suggests that the FCC's mobile-data-roaming order could serve as a regulatory template for compelling interconnection on the Internet.[7] Under this approach, networks could negotiate terms for interconnection; where conflicts arise, the FCC would provide a backstop for dispute resolution.[8] Narechana and Wu (2014) advocate that the FCC classify the ISP's transfer of data from content providers to consumers as a telecommunications service, subject to "common carrier" regulation.[9] The authors argue that "because such sender-side regulation focuses on incoming traffic, it also provides a useful framework for addressing interconnection disputes between broadband carriers and content providers."[10] This more invasive approach would give the FCC power to compel interconnection without need for voluntary negotiations, and interconnection rates could be set by regulatory fiat.

Missing from much of this debate is an analysis of the social costs and benefits associated with mandatory interconnection. This Policy Brief seeks to identify these effects from the consumers' vantage and offers an economic principle that may guide policymakers to a narrowly tailored solution. In their review of interconnection obligations across several network industries, Carlton and Picker (2006) explain that sector-specific interconnection obligations and antitrust enforcement serve as *complements* in partially deregulated industries; in fully deregulated industries, antitrust enforcement acts as a *substitute* for sector-specific interconnection obligation.[11] Because the communications sector was set on a deregulatory path nearly 20 years ago, has the time has come to rely on antitrust to adjudicate interconnection disputes on the Internet?

Is There a Case for Mandatory Interconnection?

The original basis for mandatory interconnection was to address a monopoly problem in long-distance service. In particular, interconnection was designed to provide ubiquitous coverage by supporting independent local networks in their dealings with a dominant long-distance network.[12] The 1934 Act required interconnection only if the FCC found that it would be in the public interest.[13] The general duty to interconnect did not arise until the modern era, and began with mandatory interconnection for complementary offerings such as equipment or long-distance services.[14] The 1996 Act imposed interconnection obligations on providers of local telephone services with horizontal rivals. During the late 1990s, regulators were concerned about interconnection among backbone providers, under the rationale that a large provider may have less incentive to interconnect with smaller rivals; without interconnection, customers may have an incentive to buy service from the largest provider with the best-connected network, risking monopolization of the industry.[15]

The 1996 Act set the communications industry on a deregulatory path, with the aim of spurring competitive entry into local voice and video services. As developed more fully in Part IV, whether a sector-specific interconnection obligation is still needed for communications depends on our progress along the "deregulatory arc."[16] To the extent that such an obligation was designed to address a monopoly problem, absence of a monopoly and evidence of competitive supply suggests mandatory interconnection regulations are unnecessary. And if that original basis is eroded, is there some alternative basis for mandatory interconnection not rooted in monopoly power?

The general duty to interconnect did not arise until the modern era, and began with mandatory interconnection for complementary offerings such as equipment or long-distance services.

Before exploring an alternative basis, let's quickly dispose of the monopoly justification. While we may not have arrived at some competitive nirvana, there is no debate as to whether the communications market may be fairly characterized as a monopoly nearly 20 years after the 1996 Act. The majority of residential voice service has shifted to wireless networks,[17] and prices for those voice services have been steadily declining over time. According to the FCC's most recent data, monthly average revenue *per unit* for wireless service declined from $48.04 in 2006 to $46.63 in 2011;[18] wireless voice revenue *per minute* has declined from $0.06 to $0.05 over the same period.[19] And voice revenue per minute in the United States ($0.033) is less than one third of the European average.[20] Residential consumers can choose among three technologies—wireless, telephone-based VOIP, or cable-based VOIP—for voice service, and can choose among four nationwide providers of wireless voice service.

Broadband data prices are harder to pin down, but the evidence is also inconsistent with monopoly. According to a 2010 survey by Pew Research, the average price for broadband service was roughly $41 per month.[21] A 2013 study by ITIF estimated the average price of a connection with 5 to 20 Mbps was $35.33 per month, which ranked favorably compared to other countries.[22] Cable modem providers compete with fiber-based telcos (including fast variants of DSL) capable of delivering download speeds in excess of 6 Mbps in the vast majority (about 71 percent) of U.S. households,[23] leaving about 27 percent households with a choice between cable modem and slow variants of DSL (and another two percent with none). And when wireless networks are overlaid on wireline networks, less than four percent of U.S. homes were beholden to a single provider of broadband service capable of delivering download speeds of 6 Mbps as of December 2012.[24] To the extent that mobile broadband or DSL restrains the price of cable modem service,[25] monopoly provision is moot.

Given the massive economies of scale in the supply of broadband, there should be no expectation of myriad suppliers. Fortunately, empirical evidence suggests that entry by a *single* broadband provider generates significant price effects. Using a regression model on an FCC dataset on residential broadband subscribership and speeds at the census tract level, Wallsten and Mallahan (2010) demonstrated that prices for cable modem service were up to $4.84 per month lower where cable faced an overbuilder (a firm that builds a rival broadband delivery system for the same set of consumers).[26] They also found that cable modem speeds were faster in the overbuild areas.[27] These results suggest the competitive outcomes are achievable with a modest degree of entry, which will likely be realized through a combination of fiber-based and wireless broadband.

Finally, the market for business customers appears to be increasingly competitive. Cable providers are making inroads in the Ethernet segment of the business broadband market. According to *Cable Industry Insider*, cable operators provided one quarter of the U.S. Ethernet services by the end of 2012, and cable's share is expected to reach one third (and even higher in metro areas) over the next few years.[28] According to another survey by Vertical Systems, by mid-2013, cable providers accounted for one fifth of the total U.S. Ethernet retail port base; indeed, cable operators installed more new retail Ethernet ports than the big telcos over the first six months of 2013.[29]

In light of this evidence, it is a stretch to defend an interconnection obligation as a means to address monopoly. But perhaps there is some other compelling basis for interconnection not rooted in monopoly? Werbach (2014) offers a number of alternative reasons for why mandatory interconnection is needed in the Internet era. Citing several notable interconnection disputes, he argues the mandatory interconnection serves as an "anti-fragmentation policy" that prevents service disruptions, reduces transaction costs, and fosters efficient integration.[30] He points out that the Comcast-Level 3 kerfuffle took three years to hammer out (although a standstill agreement preserved the flow of traffic in the interim),[31] and that some Verizon customers may have experienced degraded service before Verizon resolved its dispute with Cogent by upgrading port capacity on certain interconnection links.

In addition to these benefits, Werbach notes two other bases for mandatory interconnection that may exist even in reasonably competitive industries. The first, which he refers to as the "terminating-access-monopoly" problem,[32] suggests that even if there are choices for ISPs, conditional on a customer's choosing a particular ISP, a network seeking to deliver data to that ISP's customers needs access for the handoff. In the absence of an interconnection obligation at commercially reasonable rates, the ISP could exercise market power.[33] Second, Werbach explains that, even in the absence of monopoly, mandatory interconnection ensures universal connectivity, particularly in rural areas where some broadband access is lacking.[34]

It is a mistake to presume that regulator-driven interconnection arrangements are always more efficient than commercial ones, particularly when regulators have no way of knowing what solutions are most efficient.

Of these bases, the most compelling benefit of mandatory interconnection is the reduction in the length of service reductions.[35] For example, Cogent attributed the following disruption in Boston to its failure to reach an interconnection agreement with Comcast:

> The resulting traffic jam hurt Comcast subscribers and Cogent customers. For example, one of our business customers in the Boston area has many employees who telecommute from home. Those employees with Comcast Internet service at home experienced problems accessing and using their company's network because of the traffic jam.[36]

The Sprint-Cogent dispute in 2008 reportedly interfered with certain users' (whose ISPs relied on Cogent) ability to send emails to or access the websites of other users (whose ISP relied on Sprint) and vice versa.[37] The table below lists the major interconnection disputes in the United States, as well as the associated impact on Internet customers.

Major U.S. Interconnection Disputes

Parties	Year	Service Disruption	Service Outage
Cogent-AOL	2002	"left several educational institutions without service"	7 days
Cogent-Level 3	2005	"blacked out connections between their customers"	3 days
Sprint-Cogent	2008	"no longer possible for many Sprint customers and Cogent customers to directly communicate across the Internet"	4 days
Comcast-Level 3	2010	No evidence of customer impact	0 days
Verizon-Cogent	2013	"many Verizon customers had serious trouble connecting to websites that rely on Cogent for Internet connectivity"	0 days
Comcast-Cogent	2014	"employees with Comcast Internet service at home experienced problems accessing and using their company's network"	0 days

Sources: Karl Bode, AOL. Cogent Peering Spat, DSL Reports, Dec. 31, 2002; Level 3, Cogent resolve peering dispute, renew deal, Computer World, Oct. 28, 2005; Cogent becomes transit-free, renesys.com, June 26, 2008; Sprint, Cogent Resume Peering, Keep Arguing, Data Center Knowledge, Nov. 2, 2008; Brian Stelter, Netflix Partner Says Comcast 'Toll' Threatens Online Video Delivery, New York Times, Nov. 29, 2010; April Glaser & Seth Schoen, Peering into the Soft Underbelly of Net Neutrality, Electronic Frontier Foundation, Feb. 19, 2014; Schaeffer Testimony at 6.

Three of the major interconnection disputes did not lead to service outages, and even those that did were resolved within a week. There is no doubt, however, that such disruptions could be costly assuming they are not resolved quickly.

When assessing the purported benefits of mandatory interconnection, the relevant question is whether the likelihood that such a disruption will occur in the absence of a regulatory obligation is significantly greater than zero. Stated differently, of the 250-plus exabytes of projected U.S. IP traffic during 2014,[38] what fraction is at risk of being disrupted given the likelihood of an interconnection dispute? One estimate of that probability is the historical frequency of disputes that lead to service disruptions.[39] Based on the data in the above table, the historical disruption rate seems very small. If so, then no matter what the associated disruption costs, the *expected* cost of not imposing an interconnection obligation is likely small. Moreover, to the extent that having a regulatory backstop to air one's grievances causes access-seeking networks to drive a harder bargain, mandatory interconnection could perversely increase the likelihood of disputes.[40]

Finally, the cost of adjudicating these disputes must be netted out of the social benefits. The adjudication costs may not be trivial: For example, the dispute resolution process could get bogged down over the appropriate price for interconnection, raising the cost of adjudication. Should the rates be cost-based (like TDM interconnection) or inversely related to elasticities (a la Ramsey pricing)? Baseball-style arbitration of the kind endorsed by Werbach requires the existence of (competitive) market comparables to determine fair-market value. It is not clear whether any obvious comparables exist. And to the extent that state public utility commissions are involved in setting rates, as they were in the interconnection proceedings associated with implementing the 1996 Act,[41] the costs of adjudicating disputes could be even greater. It is a mistake to presume that regulator-driven interconnection arrangements are always more efficient than commercial ones, particularly when regulators have no way of knowing what solutions are most efficient.

The Social Costs of Government-Mandated Interconnection Obligations

Against these suggested benefits, one must weigh the social costs of imposing mandatory interconnection obligations on ISPs. It is beyond the scope of this Policy Brief to quantify all of these costs. Enumerating and categorizing them, however, may be helpful.

First, mandatory interconnection could undermine the incentive of ISPs to expand or enhance broadband networks. To use Werbach's Montana-roaming example,[42] if one mobile carrier covered the northern part of the state, and a second mobile carrier covered the southern part, the best outcome for consumers would be an invasion of each network into the other's territory—duopoly is better than monopoly. But if interconnection rates are mandated at zero (or sufficiently below the incremental cost of self-provision), it may not pay for either carrier to expand its network. Some have blamed mandatory roaming for Sprint's and T-Mobile's decision not to build out in high-cost areas (but rather rely on roaming), even in Sprint's home state of Kansas.

Applied to wireline broadband networks, if a telecom believed that it could not be compensated for upgrading its capacity (either from DSL to fiber or increase the density of a fiber network) due to restrictions on what it could charge for paid peering, then it might abandon or curtail the investment decision. The experience of mandatory unbundling in Europe has likewise shown the decreased incentive for network operators to invest in fiber. In a similar vein, interconnection-pricing rules determined on a case-by-case basis could create investment-deterring uncertainty relating to the implementation of regulations. This uncertainty could be exacerbated if multiple state regulators are involved in the rate-setting process.

Second, mandatory interconnection could undermine the incentive of transit providers to extend their reach into the last mile. Just as mandatory interconnection (and unbundling) undermined the CLECs' incentive to invest in their own facilities,[43] regulated interconnection rates could deter transit or even content providers from building the last-mile connections. Google has started laying its own fiber in select cities through the country. And Level 3 and Cogent both offer Internet access in addition to transit and content delivery network (CDN) services. It is therefore not a stretch to consider these access-seeking network owners as potential entrants in last-mile connectivity. Any interconnection regime should carefully consider an entrant's tradeoffs in making versus buying terminating access; if the terms of buying are too generous, then deploying last-mile networks become relatively less attractive.

Third, mandatory interconnection could unravel paid arrangements between content providers and ISPs if "better terms" could be secured via intermediary networks through regulation. The most likely explanation for why Netflix's CEO advocated for "strong net neutrality" protections for Cogent and Level 3 is that, to the extent those intermediaries can secure better interconnection terms via regulation than Netflix via negotiation, Netflix could reduce its transit costs. In other words, Netflix would not get zero-price connection but Cogent might, setting up arbitrage opportunities. Regulatory uncertainty about if and when mandatory interconnection is imposed could induce large content providers such as Google and Amazon to refrain from entering paid-peering arrangements with ISPs.

A duty to interconnect is arguably more invasive than a non-discrimination requirement.

Fourth, CDNs or transit providers might not contribute significant value to certain transactions, such as those involving large content providers that have vertically integrated into "middle-mile" services. There seems to be a continuing role, however, for these intermediaries for moving or augmenting the traffic of small to medium-sized content providers. Regulatory life support for intermediaries on transitions for which they bring questionable value could attract rent-seeking behavior. Inefficiencies could arise as the regulator caters to special interests or makes mistakes.

Fifth, mandatory interconnection could reduce the incentive of two parties to reach an agreement to minimize total costs. Besen and Israel give an example in which a CDN could be relieved of $2 million in costs so long as the ISP incurred a $1 million investment. If the mandatory interconnection rate were set below $1 million, the cost-saving technology would not be adopted.[44]

Sixth, policies that thwart negotiations between content providers and ISPs—whether they concern quality of service or involve paid peering—could lead to higher broadband access prices via the "seesaw principle." The economics of two-sided markets suggests that raising money from one side of the market (advertisers) puts downward pressure on prices for the other side of the market (end users). Preventing network providers from exercising pricing flexibility could increase the proportion of the network costs that providers must recover directly from end users.

Seventh, given that Netflix represents roughly one third of download traffic during peak hours,[45] the recently adopted arrangements between content providers and ISPs can be viewed as an efficient solution to the classic peak load pricing problem. While peak-load pricing is often associated with regulated utilities, it has frequently been applied in competitive industries with periodic demand fluctuations for a non-storable good. Peak-load pricing efficiently allocates the scarce resource (bandwidth at peak hours), by raising the price to those who demand it most (Friday night movie viewers). Preventing network providers from exercising pricing flexibility could force high-intensity users to be inefficiently subsidized by everyone else, causing ISPs to raise rates to end-users across the board.

Eighth, proponents of mandatory interconnection should be careful what they wish for. While they may be able to convince the FCC that it should require interconnection (at potentially a zero rate), it is unlikely that they could convince every regulator around the world of that. And once the FCC had set the precedent that these IP interconnection matters should not be left to commercial negotiations, other regulators around the globe would be free to determine the terms and rates for interconnection as they see fit, with no guarantee that regulators across the globe would successfully coordinate their policies. This, in turn, could dramatically increase regulatory uncertainty and the frequency of service disruptions.

Policy Implications

Assuming the social costs of mandatory interconnection exceed the benefits, what might some alternative, less-invasive policy look like? Two considerations are worth keeping in mind. First, the interconnection debate is not proceeding in a vacuum: The FCC is developing certain protections for content providers in its revised Open Internet rules, including non-discrimination and a no-blocking rule. A duty to interconnect is arguably more invasive than a non-discrimination requirement.[46] Non-discrimination would require ISPs to offer quality-of-service agreements to all similarly situated content providers at the same terms; by comparison, mandatory interconnection would require ISPs to deal with any content provider or transit provider that wished to terminate traffic on the ISP's network. In other words, interconnection is a duty to deal in the first instance.

When deciding whether to overlay an interconnection obligation on top of protections for content providers, one must articulate the *incremental* benefits that are derived from the added layer of protection. To the extent that content providers (as well as application, service providers, and device providers) could be protected by an effective Open Internet regime, including a "minimum level of access"[47] established by the no-blocking rule, the only remaining class of providers that could benefit from mandatory interconnection would consist of intermediaries that operate at the core of the network, such as standalone CDN or transit providers. Unlike content providers, who generate positive spillovers (information and artistic content can be viewed as "public goods") and thus cannot be expected to monetize their investment, the rationale for protecting these intermediaries is less compelling. Although these intermediaries might be marginalized in the absence of regulatory protection on transactions involving large content providers—Netflix and Google have developed their own CDNs[48]—it is not clear how consumers would benefit from rules that reinserted the presence of these intermediaries. It is not even clear whether these intermediaries need interconnection revenues to thrive. For example, Level 3 acknowledged in a first quarter 2014 earnings call that Netflix is "not even in our top 30 customers, so the revenue impact is relatively small."[49] Small and mid-sized content providers will continue to rely on third party CDNs and transit providers in the absence of mandatory interconnection.[50]

The communications sector was set on a deregulatory path nearly 20 years ago with the passage of the 1996 Telecommunications Act, perhaps the time has come to turn to antitrust to adjudicate interconnection disputes on the Internet.

Second, there is always the antitrust backstop for excluded networks. In their review of interconnection obligations across several network industries, Carlton and Picker explain that sector-specific "interconnection policies" and antitrust enforcement serve as complements in partially deregulated industries.[51] In fully deregulated industries, however, antitrust acts as a substitute for sector-specific interconnection obligations.[52] The role of mandatory interconnection has waned in network industries as they become competitive:

> The deregulated network industries that we examined all show a similar pattern: after deregulation, there is massive consolidation, *a lessening of the reliance on interconnection from other firms*, a decline in either wages or employment or both, and a fall in prices with a reduction or end to any cross subsidy. Consumers benefit, special interests are harmed.[53]

Because the communications sector was set on a deregulatory path nearly 20 years ago with the passage of the 1996 Telecommunications Act, perhaps the time has come to turn to antitrust to adjudicate interconnection disputes on the Internet.

So where does this leave us? In the absence of mandatory interconnection, content providers can deal with ISPs directly for enhanced quality of service pursuant to a paid-peering arrangement; they are free to turn down enhanced quality of service for standard treatment. As explained above, content providers may be able to avail themselves of non-discrimination protections that could come out of new Open Internet proceeding, if they can prove they are not receiving equal treatment for reasons relating to affiliation or favored status.

Transit providers, CDNs, and other intermediary networks can avail themselves of antitrust courts if ISPs refuse to deal as a means of extending their (alleged) market power into adjacent markets. In *Otter Tail*,[54] the Supreme Court found antitrust liability for an electric utility company for failure to interconnect with another utility even though the Federal Power Commission could order such interconnection. To be fair, antitrust cases do not proceed quickly. And with the exception of cases like *Aspen Skiing* and *AT&T*, antitrust rarely imposes mandatory obligations to interconnect, other than as a remedy for an independent antitrust violation. In *Trinko*, the Supreme Court recognized that antitrust has only weakly embraced affirmative duties to interconnect. Accordingly, while the path is not clear, excluded networks should have a reasonable chance of prevailing so long as they can establish monopoly power (presumably in terminating access) and antitrust impact (in the form of higher prices or reduced output in some relevant product market).

Finally, there are other ways to create a backstop without imposing a duty to interconnect. For example, Weiser (2009) proposes the development of a self-regulating organization (SRO) that wields decision-making authority and over which the FCC has authority. The SRO would act like a standard-setting body. Under this approach, the FCC must first develop norms such as requirements to provide some level of transparency over the terms of peering as well as pre-announced standards for how to “de-peer” an Internet backbone provider. Weiser suggests that matters that cannot be resolved in SRO can be appealed to the FCC.[55]

Conclusion

This Policy Brief examines the benefits and costs of dictating the terms of interconnection on the Internet. Based on a preliminary review of the evidence, the benefits (net of enforcement costs) appear to be slight, whereas the costs could be economically significant. In light of these tradeoffs, the FCC should be hesitant to dictate the terms of interconnection. Additional layers of protection, including non-discrimination requirements and antitrust enforcement, as well as the continued private negotiation of interconnection, should be more than sufficient to keep the Internet humming along on all cylinders and ensure we get to see the next installment of *House of Cards* without any delay.

Endnotes

1. For an early primer on interconnection, see General Accounting Office, Characteristics and Competitiveness of the Internet Backbone Market, GAO-02-16, Oct. 2001 [hereafter *GAO Study*].
2. Cogent has been involved in similar disputes outside of the United States. *See, e.g.*, Karl Bode, *Cogent Involved in Another Peering Dispute*, DSL Reports, Mar. 18, 2008 (describing dispute with Swedish telecom operator Telia). Other major interconnection disputes are described in Part II.

3. Cogent claimed that after it began carrying Netflix traffic in mid-2012, its relation with Comcast soured, at which point Comcast allegedly refused to augment capacity at Cogent's interconnection points and demanded that Cogent pay to connect to Comcast's network. *See* Written Statement of Dave Schaeffer Chairman and Chief Executive Officer Cogent Communications Group, Inc, Before the U.S. House of Representatives Committee on the Judiciary Subcommittee on Regulatory Reform, Commercial and Antitrust Law, May 8, 2014 [hereafter *Schaeffer Testimony*]. In a May 2014 blog post, Level 3 claimed that several ISPs were refusing to upgrade their peering connections for the past year. *See* Observations of an Internet Middleman, *available at* http://blog.level3.com/global-connectivity/observations-internet-middleman/.
4. Notice of Proposed Rulemaking, In the Matter of Protecting and Promoting the Open Internet, May 15, 2014 [hereafter *2014 Open Internet NPRM*].
5. *Id.* at 87.
6. *Id.* at 22. Some commenters argue that interconnection is related to "net neutrality" in the sense that ISPs "can block traffic, discriminate, or impose access fees either once traffic is within their network . . . or when the traffic is at the edge of their network (through interconnection)." *See, e.g.*, Marvin Ammori, Interconnection disputes are net neutrality issues, Apr. 7, 2014, *available at* http://ammori.org/2014/04/07/interconnection-disputes-are-network-neutrality-issues-of-netflix-comcast-and-the-fcc/. Others explain that the issues are in fact separate, and that the Netflix-Comcast dispute is highly specialized due to the unusually large amount of traffic created by a single content provider. *See, e.g.*, Marguerite Reardon, C*omcast vs. Netflix: Is this really about Net neutrality?,* CNET, May 15, 2014, *available at* http://www.cnet.com/news/comcast-vs-netflix-is-this-really-about-net-neutrality/?tag=nl.e703&s_cid=e703&ttag=e703&ftag=CAD090e536.
7. Kevin Werbach, *No Dialtone*: *The End of the Public Switched Telephone Network*, 66 FED. COMM. L. J. (2014).
8. As part of the "IP transition" from the public switched telephone network, ISPs could commit to offer interconnection on commercially reasonable terms, subject to a backstop arbitration mechanism. Werbach further suggests that ISPs be required to disclose terms of signed interconnection agreements, and that "baseball-style" arbitration (where each party offers last-and-final proposal and the arbitrator chooses the offer closest to fair-market value) could be used to adjudicate pricing disputes. *See id.* Although *No Dialtone* addresses the need for interconnection for VOIP only, Werbach has argued for a broader mandate for all Internet traffic in prior writings. *See, e.g.*, Kevin Werbach, *Only Connect*, 22 BERKELEY TECH. L.J. 1233 (2007).
9. Tejas N. Narechania & Tim Wu, Sender Side Transmission Rules for the Internet, Filed Comments with the FCC, Apr. 14, 2014.
10. *Id.* at 2. *See also* Brian Fung, *The decades-old idea that could break the net neutrality logjam*, WASH. POST, Apr. 21, 2014 ("It's a way of giving the commission interconnection and net neutrality at the same time; we have the magic formula and it'll solve all your problems.") (quoting Wu), *available at* http://www.washingtonpost.com/blogs/the-switch/wp/2014/04/21/the-decades-old-idea-that-could-break-the-net-neutrality-logjam/.
11. Dennis Carlton & Randal Picker, Antitrust and Regulation, John M. Olin Law & Economics Working Paper (2006).
12. *No Dialtone* at 35 ("In an era of regulated monopoly, the government mandates interconnection to ensure ubiquitous service, and regulates interconnection charges to allocate costs across the network."). Some scholars view the 1913 Kingsbury Commitment as the beginning of a compact between regulator and monopolist. Under this interpretation, in return for the government's agreement not to pursue its case against AT&T as a monopolist, AT&T agreed to allow independent telephone companies to interconnect with AT&T's long-distance network.
13. Christopher Yoo, *Is There a Role for Common Carriage in an Internet-Based World*, 51 HOUSTON LAW REVIEW 545-608 (2013).
14. In the 1960s, the FCC imposed a duty to interconnect with complementary services with *Hush-a-Phone* and *Carterfone* (customer-premise equipment), and *Execunet* I and II (long-distance service). These decisions were reinterpretations of the traditional common carrier obligations. Other enforcement episodes in the modern era highlight the nexus between mandatory interconnection and monopoly. For example, in 1978, as a remedy to a monopolization claim, the D.C. Circuit ordered AT&T to interconnect with MCI's long-distance network.
15. *GAO Study* at 19.

16. *No Dialtone* at 2 ("There are two mainstream views about how to handle the PSTN transition. One is that it represents the completion of a deregulatory arc begun at the AT&T divestiture and accelerated by the Telecommunications Act of 1996. The other is that longstanding regulatory obligations need only to be extended to a new world."). It bears noting that an IP interconnection is different from PSTN interconnection because of the nature of the Internet. For example, a large content provider might connect to an ISP via hundreds of Ethernet ports in myriad locations. Accordingly, there is no clear boundary around the ISP's network because there are so many paths between Internet exchanges and neighborhoods.
17. According to the CDC, nearly 40 percent of U.S. homes had only wireless telephones during the first half of 2013, and another 16% received all or almost all calls on wireless telephones despite also having a landline telephone. Stephen J. Blumberg, Ph.D., and Julian V. Luke, CDC, Wireless Substitution: Early Release of Estimates From the National Health Interview Survey, January–June 2013, available at http://www.cdc.gov/nchs/data/nhis/earlyrelease/wireless201312.pdf.
18. Federal Communications Commission, Annual Report and Analysis of Competitive Market Conditions With Respect to Mobile Wireless, Including Commercial Mobile Services, Sixteenth Report, Mar. 21, 2013, at 186.
19. *Id.* at 179.
20. *Id.* at 257.
21. Pew Research, Trends in Broadband Adoption, *available at* http://www.pewinternet.org/2010/08/11/trends-in-broadband-adoption/.
22. Richard Bennett, Luke A. Stewart, and Robert D. Atkinson, The Whole Picture: Where America's Broadband Networks Really Stand, ITIF, Feb. 2013, at 49-50, *available at* http://www2.itif.org/2013-whole-picture-america-broadband-networks.pdf.
23. Internet Access Services: Status as of December 31, 2012, Figure 5(a), *available at* http://transition.fcc.gov/Daily_Releases/Daily_Business/2013/db1224/DOC-324884A1.pdf.
24. *Id.* at Figure 5(b).
25. Aaron Smith, "Cell Internet Use 2012," Pew Internet and American Life Project, June 26, 2012 (www.pewinternet.org/Reports/2012/Cell-Internet-Use-2012.aspx) (showing that 17 percent of cell phone owners do most of their online browsing on their phone, rather than on a computer or other device); Joan Engebretson, "Pew: Smartphones Fueling Wireless Broadband Substitution?" *Telecompetitor*, July 11, 2011 (www.telecompetitor.com/pew-smartphones-fueling-wireless-broadband-substitution) (showing that a quarter of smartphone owners reported that for the most part they went online by using their phone rather than a computer and roughly one-third of those owners lacked a high-speed home broadband connection); Cisco, "To Prevent 15% of Customers from Cord-Cutting, Fixed Broadband SPs Consider WiFi Solutions to Deliver the Mobility Customers Seek, October 2011 (www.cisco.com/web/about/ac79/docs/FastFacts/FastFacts_WiFi_Defense_against_BB_Cord_Cutting_Oct2011.pdf) (estimating that up to 15 percent of U.S. consumers could cut their broadband wireline in favor of a mobile data connection by 2016).
26. Scott Wallsten and Colleen Mallahan, "Residential Broadband Competition in the United States," BE Press Working Paper, March 2010, at 32, table 7 (http://works.bepress.com/cgi/viewcontent.cgi?article=1105&context=scott_wallsten). The authors found that cable modem prices declined between $1.25 (cable speed tier 6) and $4.84 (cable speed tier 5) per month when cable modem providers faced an overbuilder. Coefficients were estimated at the 1 percent significance level. In contrast, the authors found that cable modem prices did not decline significantly when cable providers faced DSL *or* fiber to the home providers (their "two-provider" results), suggesting either that DSL did not constrain the price of cable modem service, thereby neutralizing the impact of fiber competition, or that neither DSL nor fiber constrained the price of cable modem service. Unfortunately, the authors did not estimate the incremental price-constraining effect of fiber only.
27. *Id.* 29.
28. Cable Industry Insider, Cable Operators & Ethernet: Serious Market Share (2013), *available at* http://www.heavyreading.com/cable/details.asp?sku_id=3060&skuitem_itemid=1505.
29. Vertical Systems Group, Mid-Year 2013 U.S. Carrier Ethernet Leaderboard, Aug. 20, 2013, *available at* http://www.verticalsystems.com/vsglb/mid-year-2013-u-s-carrier-ethernet-leaderboard/.

30. See, e.g., David A. Malueg & Marius Schwartz, Compatibility Incentives Of A Large Network Facing Multiple Rivals, JOURNAL OF INDUSTRIAL ECONOMICS (2006) ("An installed-base share of at least 50% is necessary but not sufficient to make autarky [non-compatibility with rival networks] unambiguously profitable.").
31. *Id.* at 40 ("Presumably, the companies had continued to exchange traffic the past three years under some interim arrangement, before agreeing to new terms.").
32. *Id.* at 41 ("In other words, a network seeking to deliver video or voice content to an AT&T U-verse broadband access subscriber needs to terminate that traffic on AT&T's network. The fact that AT&T has many broadband competitors is irrelevant once the customer has chosen a particular one. In the telecommunications market, this is known as the terminating access monopoly.").
33. A hypothetical retail monopoly does not by itself imply that market power will be exercised in upstream interconnection. Put another way, monopoly is a necessary condition for certain anticompetitive acts, but is not sufficient.
34. *Id.* at 45 ("Clearly, the FCC recognizes that as the PSTN migrates to IP technology, the need for interconnection to ensure universal connectivity does not evaporate."). The likely reason why PSTN interconnection terms and conditions needed to be regulated is because interconnection was mandated. An interconnection obligation can encourage certain parties to seek excessive access rates. Commercial negotiations for IP interconnection have worked well in the absence of a mandate, as no party wants to deny a commercially sensible agreement. *See, e.g.,* Dennis Weller & Bill Woodcock, Internet Traffic Exchange: Market Development and Policy Challenges, OECD Digital Economy Papers at 62 (2013) (finding that the vast majority of IP interconnection agreements are carried out based on a "handshake," compared with multi-thousand page PSTN interconnection tariffs that are often litigated).
35. The terminating-access monopoly seems to have no limiting principle: An ISP with 1 percent market share could be subjected to any number of regulatory obligations pursuant to this basis. And while ubiquitous coverage is laudable, there are more direct approaches to encouraging broadband deployment into rural areas, including subsidies.
36. Schaeffer Testimony at 6.
37. Mikael Ricknäs, *Sprint Reconnects Cogent, But Differences Are Unresolved*, NETWORK WORLD (Nov. 3, 2008), *available at* http://www.networkworld.com/news/2008/110308-sprint-reconnects-cogent-butdifferences.html?fsrc=netflash-rss.
38. Cisco Visual Networking Index, *available at* http://www.ciscovni.com/forecast-widget/index.html.
39. One economic model demonstrates that networks with a non-dominant market share have strong incentives to interconnect, suggesting that the probability of disruptions flowing from interconnection disputes is small. *See, e.g.*, David A. Malueg & Marius Schwartz, *Compatibility Incentives Of A Large Network Facing Multiple Rivals*, JOURNAL OF INDUSTRIAL ECONOMICS (2006) ("An installed-base share of at least 50% is necessary but not sufficient to make autarky [non-compatibility unambiguously profitable."). Their model was designed to address issues of interconnection between horizontal rivals such as instant-messaging providers or pure backbone providers; its policy implication may be limited in the case of interconnection between an ISP and a backbone provider.
40. Assume for simplicity there are two networks that have an interconnection agreement that must be renewed each year. Let p_1 and p_0 be the probability of a dispute with and without mandatory interconnection, respectively, and q be the conditional probability of a dispute leading to a service disruption. Assuming generously that an adjudication regime could drive the expected costs of service disruptions to zero, then the net benefit of mandating interconnection would be $p_0 \times q \times C - p_1 \times D$, where p_0 is the probability of a dispute in the absence of a regulatory backstop, p1 is the probability of a dispute in the presence of a regulatory backstop, and D is the cost to arbitrate a dispute. Crediting a regime with driving the costs to zero would require that the arbitrator resolve the dispute instantaneously and thereby head off any associated disruption costs. In reality, the arbitrator would have to gather facts, build a record, and then hold a hearing. And if the arbitrator's decision were subject to appeal, the process could take even longer. Moreover, the mandatory-interconnection regime should not be credited for heading off service disruptions associated with disputes that the regime itself created—these would be the "benefits" given by $(p_1 - p_0) \times q \times C$.
41. OECD, Working Party on Telecommunication and Information Services Policies Interconnection and Local Competition, Figure 4, *available at*

http://www.oecd.org/sti/ieconomy/1894706.pdf (showing that in 1999, the regulated interconnection charge for call termination on Bell Atlantic New York's network was roughly one cent per minute).

42. *No Dialtone* at 35 ("The [Montana] state troopers often had to drive 30 miles or more to get a usable signal. Public safety services were adversely affected for residents of that part of Montana. This example illustrates the power of interconnection."). Werbach presumes that the cause of the outage was a failed roaming agreement between AT&T and Verizon. But AT&T's conversion of Alltel's CDMA network to GSM was likely to blame, as old CDMA phones used by certain Verizon customers would not work on the GSM network.
43. Robert W. Crandall, Allan T. Ingraham & Hal Singer, *Do Unbundling Policies Discourage CLEC Facilities-Based Investment?*, 4 Topics In Economic Analysis and Policy (2004).
44. Stanley M. Besen & Mark A. Israel, The Evolution of Internet Interconnection from Hierarchy to "Mesh": Implications for Government Regulation, Tech Policy Institute (2012).
45. Peter Kafka, *Netflix Still Eats a Third of the Web Every Night; Amazon, HBO and Hulu Trail Behind*, *All Things D*, May 14, 2014.
46. Edward Wyatt, *F.C.C., in a Shift, Backs Fast Lanes for Web Traffic*, New York Times, Apr. 23, 2014, *available at* http://www.nytimes.com/2014/04/24/technology/fcc-new-net-neutrality-rules.html?_r=0.
47. *2014 Open Internet NPRM* at 37.
48. Richard Bennett, Netflix and Comcast Declare Peace, High Tech Forum, Feb. 23, 2014, *available at* http://www.hightechforum.org/netflix-and-comcast-declare-peace/ (explaining that Google and Netflix have built their own CDNs).
49. Level 3 Communications Management Discusses Q1 2014 Results - Earnings Call Transcript (quoting Jeffrey K. Storey).
50. *Id.* ("There's a new Netflix coming along. I don't know what the next one will be, but there are other companies out there trying to grow and trying to have the success that they've had. We will continue to serve those customers, use our infrastructure.").
51. Carlton & Picker at 40 ("The trends in network industries indicate that regulators, not antitrust courts, will bear the responsibility for formulating interconnection policies *in partially deregulated industries*, but antitrust will remain in the background as a club that firms can use if regulators allow incumbents to acquire market power either through merger or predatory acts." (emphasis added)).
52. *Id.* at 22 ("This recent history highlights a move away from regulation towards antitrust as a means to control competition and reveals how regulation and antitrust can be both substitutes and, in some settings, complements. The substitution involves the complete replacement of regulation with antitrust, *as occurs when industries become deregulated* (e.g., airlines and trucks)." (emphasis added)).
53. *Id.* at 39 (emphasis added).
54. 410 U.S. 366 (1973).
55. Phil Weiser, The Future of Internet Regulation, Colorado Law Legal Studies Research Paper Series, Working Paper 09-02, Feb. 2009.

About the Author

Hal Singer is a Senior Fellow at Progressive Policy Institute, Principal at Economists Inc. and Adjunct Professor at Georgetown's McDonough School of Business.

Acknowledgement

I would like to thank Richard Bennett, Kevin Caves, Larry Downes, Gerry Faulhaber and Robert Litan for valuable comments.

About the Progressive Policy Institute

The Progressive Policy Institute (PPI) is an independent research institution that seeks to define and promote a new progressive politics in the 21st century. Through research and policy analysis, PPI challenges the status quo and advocates for radical policy solutions.

Public Knowledge

Telecommunications Act Rewrite
Question Responses: Network Interconnection

Jodie Griffin, Senior Staff Attorney
Public Knowledge

U.S. House of Representatives
Committee on Energy and Commerce
Subcommittee on Communications and Technology

August 8, 2014

Interconnection is a hallmark of networks of any kind, and it is a fundamental part of communications policy in the United States. Interconnection and other competition policies lie at the heart of the development of a robust and competitive communications market. As we saw more than 100 years ago, without mandatory interconnection the phone network will slide inevitably toward monopoly as the largest carriers can gain anticompetitive advantages by withholding access to their customers from competitors. As carriers now move toward all-IP networks, policymakers must determine how they will best achieve interconnection and competition between providers post-transition. These policies are critical to creating and maintaining a functioning interconnected network and a competitive market for communications services.

The duty to interconnect with other networks was first a means of enabling universal service in rural areas in the days of the old AT&T monopoly so rural cooperatives, municipalities, and local businesses could bring service to places AT&T found too expensive to serve. Later, as amendments to the Act shifted national policy from regulated "natural monopoly" to encouraging competition among competing networks, interconnection became the *sine qua non* of fostering and developing competition. Unless we propose to return to the days of regulated natural monopoly, our national policies must absolutely guarantee that competing networks will continue to accept each other's traffic and terminate each other's calls in a manner that both preserves call quality throughout the country and actively promotes a robust and competitive environment.

The Communications Act imposes upon telecommunications carriers "the duty to interconnect directly or indirectly with the facilities and equipment of other telecommunications carriers."[1] But it remains unsettled what this obligation would really mean in a world where the FCC has treated broadband service as an information service and declined to classify interconnected VoIP at all.

Even as network technologies and uses change, certain fundamental characteristics of networks—and the problems that can arise in them—stay the same. Interconnection

[1] 47 U.S.C. § 251; *see also* 47 U.S.C. § 201(a).

policies that encourage build-out and enable competition are crucial to the network, regardless of which protocols that network is using. A small network provider cannot launch if it cannot obtain reasonable interconnection terms with other carriers. And an existing carrier cannot get around the fact that, no matter how many middlemen exist in the system, data or voice service being sent to a particular person must at some point travel over the infrastructure of the carrier serving that person.

As the phone network transitions to Internet Protocol and as the IP-to-IP interconnection market changes and develops, it is crucial that policymakers at the state and federal levels can collect information about recent developments and can set reasonable rules when doing so is necessary to protect competition and consumers.

1. **In light of the changes in technology and the voice traffic market, what role should Congress and the FCC play in the oversight of interconnection? Is there a role for states?**

Congress and the FCC remain important authorities as the voice traffic market develops, especially in this time of transition. Regardless of the protocols being used or the number of interconnection points, our social needs, the economics of networks, and the essential services (transporting information at the direction of the user) remain the same. Policymakers must be able to stay informed about the agreements that shape the basic operations of our communications networks. But knowing what is going on—while necessary—is not enough. Policymakers must also retain authority to ensure interconnection agreements are not being used to stifle competition or diminish service in hard-to-reach rural areas.

Additionally, even as the number of interconnection points may decrease, state authorities' expertise and interest in ensuring efficient connections for the networks located in their states remain the same. Therefore state agencies should also be able to collect information about interconnection agreements and ensure interconnection disputes do not harm consumers or competition.

As discussed above, the duty to interconnect first arose as a means of ensuring service in rural areas in the days of the AT&T monopoly when rural cooperatives, municipalities, and local businesses brought service to places AT&T found too expensive to serve. Later, as amendments to the Act shifted national policy from regulated "natural monopoly" to encouraging competition among competing networks, interconnection became the *sine qua non* of fostering and developing competition. Unless we propose to return to the days of regulated monopoly, regulators must absolutely guarantee that competing networks will continue to accept each other's traffic and terminate each other's calls in a manner that both preserves call quality throughout the country and actively promotes a robust and competitive environment.

In particular, subscribers to different networks must not find themselves the victims of "peering disputes" that cut off communications and vital services. If NBC and AT&T have a retransmission dispute and AT&T video subscribers temporarily lose NBC programs, it is annoying. But if Comcast and AT&T have a "peering dispute" and millions of AT&T wireless customers can't call Comcast landlines, it is a communications disaster. It is not enough to speculate that incentives will prevent such a thing from occurring. The relevant agencies must retain adequate authority to make sure that such an event continues to be impossible after the transition.

It is not simply idle speculation to imagine this happening to a post-transition PSTN. Already, some carriers are refusing to file IP-to-IP interconnection agreements at the state level.[2] Without adequate interconnection requirements, consumers may find themselves suffering from interconnection disputes between carriers that provide not just their video and internet access, but their basic voice service as well. And if the interconnections that have tied together our voice network unravel, dominant service providers will be able to leverage their customer bases against competitors and control increasingly large shares of the market, resulting in higher prices and fewer choices for consumers.

2. **Voice is rapidly becoming an application that transits a variety of network data platforms. How should intermodal competition factor into interconnection mandates? Does voice still require a separate interconnection regime?**

Nearly 96% of the country subscribes to a basic voice service.[3] It is true that many people now also communicate with other technologies, using a plethora of devices. People like choices. But the idea that the rise of these new technologies renders basic voice communication obsolete is demonstrably false. No other technology even comes close to the level of ubiquitous adoption that voice service has, and there remain several types of communications that overwhelmingly rely on voice service, like emergency services.

Even if we were to consider only traditional copper POTS (Plain Old Telephone Service), that still leaves 100 million people, as well as millions of small businesses, relying on that service.[4] Of those, 15 million people depend on POTS alone. A remaining 85 million people keep their POTS lines despite also having mobile phones. We are dealing with the technology that remains the foundation on which many of our additional communications technologies rest. To approach this exercise with the attitude that "the market" has chosen

---

[2] *See Petition for a Determination that Verizon IP-to-IP Interconnection Agreements Must Be Filed for Review and Approval and for Associated Relief*, D.T.C. No. 13-2, Commonwealth of Massachusetts Department of Telecommunications and Energy.

[3] *FCC Releases New Telephone Subscribership Report* (Aug. 13, 2013), http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-292758A1.doc.

[4] Anna-Maria Kovacs, *Telecommunications Competition: The Infrastructure-Investment Race*, INTERNET INNOVATION

ALLIANCE (Oct. 2013), http://internetinnovation.org/images/misc_content/study-telecommunications-competition-

09072013.pdf.

and we can safely shut down a system on which 100 million people still directly depend courts disaster.

And as voice service increasingly relies on different protocols, technologies, and physical infrastructure, it is even more important that users can depend on networks interconnecting reliably and seamlessly. Particularly as new technologies and services introduce more complexity into interconnection agreements, there will be more risk for the connections that bind our networks together to start unraveling at the edges, either because no market actor has the responsibility to ensure calls and other data transmissions go through, or because companies with the most last-mile connections find opportunities to leverage their customers against smaller or competing companies. Whatever the cause, the FCC and other authorities must be able to collect information about developments in the interconnection market and set rules where needed to ensure reliable, nondiscriminatory connections.

3. **How does the evolution of emergency communications beyond the use of traditional voice service impact interconnection mandates?**

As we have increasingly depended on competition among networks, we have increasingly depended on interconnection to ensure traffic can travel from one network to another. The absence of interconnection guarantees puts vital services at risk. Even without anti-competitive intent, Congress found it necessary to order interconnection between IP-based providers and ILECs to ensure reliable 9-1-1 access.[5]

One area of possible concern is the rise of VoIP and texting applications that use PSTN phone numbers as user IDs, but are not "interconnected" services because they do not interoperate with the PSTN or SMS networks. For example, an application like WhatsApp allows people to communicate with each other using their phone numbers as user identification, but it can be used only to communicate with other WhatsApp users. By contrast, a service like Google Voice is a true interconnected texting app since it allows users to communicate with any other SMS user, regardless of whether that user is a Google Voice subscriber.[6] As text-to-911 or other new emergency service mandates begin to be applied to wireless providers, consumers may be confused as to where such services work and where they do not. The FCC's move to require text-to-911 for native and interconnected SMS services[7] is a good first step but policymakers should consider approaches to emergency communications that do not rely on subtle technological distinctions that may elude average users, or other means to educate consumers as to the

---

[5] *See* Next And Emerging Technologies 911 Improvement Act of 2008, Pub. L. 110-283

[6] In the middle ground are applications like Apple's iMessage that are private networks like WhatsApp but have the ability to fall back to a phone's native SMS capabilities to communicate with non-iMessage users.

[7] Julian Hattem, *FCC Pushes for 911 Emergency Text Messages*, THE HILL (Aug. 8, 2014), http://thehill.com/policy/technology/214692-fcc-calls-for-texts-to-911.

distinction between interconnected services that offer emergency communications capabilities and private services that do not.

As customers rely on new technologies for access to critical emergency services, policymakers must ensure those technologies will continue to deliver on the core functions customers expect of them. If consumers are assuming that, for example, VoIP or wireless service will always connect them immediately to 9-1-1, we must make sure those consumers will not be left unconnected at the very moment they need the phone network most.

4. **Ensuring rural call completion has always been a challenge because of the traditionally high access charges for terminating calls to high-cost networks. Does IP interconnection alleviate or exacerbate existing rural call completion challenges?**

The recent rural call completion problem demonstrates how technological transitions will inevitably result in unforeseen complications, which highlights the importance of having authorities equipped and ready to gather information about the problem and resolve it with minimal disruption to consumers. A transition of this magnitude and level of complexity will inevitably result in service problems that may not be the fault of any one actor, but must nonetheless be assessed and dealt with in a timely manner by competent authorities.

Rural call completion problems increasingly occur when calls originate in IP networks and are routed to rural TDM-based systems.[8] Using IP enables providers to employ third party "least cost routers" to minimize the cost of transporting traffic. This normally useful function of minimizing cost creates problems with rural exchanges as providers employ multiple routes to arbitrage intercarrier compensation (ICC) rates. This extended routing introduces latency, which the target rural TDM system interprets as "dead air." Incoming calls are dropped, or may never get completed.

Any rules or regulatory system that regulates only on the basis of market power or consumer fraud cannot address rural call completion. There is no bad actor trying to use market power to force rural exchanges out of business or to pay some sort of monopoly rent to transmit the call. Nor is the carrier from whom the call originates engaged in an unfair practice. The problem comes from the natural market dynamic of carriers using new technology to save money.

Nor do market incentives fix the problem. Larger urban carriers do not lose customers due to rural call completion problems—at least not in sufficient numbers to make the additional expense of direct routing worthwhile. It is often said carriers have no

[8] *See* Harold Feld, *If the Phone Doesn't Ring, It's Rural America*, DAILY YONDER (Mar. 29, 2013), http://www.dailyyonder.com/if-phone-doesnt-ring/2013/03/22/5733; Harold Feld, *Rural Call Completion and Network Neuropathy*, PUBLIC KNOWLEDGE (Apr. 3, 2013), http://www.dailyyonder.com/if-phone-doesnt-ring/2013/03/22/5733.

incentive to give bad service. But the real question is whether carriers have adequate incentive to give good service.

Those unwilling to concede that the IP transition raises problems that only government oversight can solve are quick to point out that it is artificially high ICC rates that drive carriers to use least cost routing in the first place. This is true, but not terribly relevant. The FCC has already announced a plan to eliminate ICC altogether and shift to "bill and keep." That has not solved the rural call completion problem.

If we insist the FCC has no role in the IP transition except ongoing elimination of its "legacy regulation," then we have two choices for dealing with rural call completion. We can immediately eliminate ICC, with the resultant disastrous impact on rural carriers from the sudden elimination of revenue. Or we can tell rural America to wait several years and hope the problem clears up.[9] Alternatively, we can embrace the idea that the FCC still has a role to play in maintaining a stable phone system and that a rush to eliminate its power to ensure that calls originating on one network terminate on another are premature.

The other problem with the "ICC is the problem" approach to rural call completion is that it assumes that no such problem could ever happen again. This seems wildly optimistic. No one anticipated rural call completion would become an issue in the transition. It seems very likely that other unanticipated problems will arise, and we will want the FCC to address them.

As a practical matter, the recent rural call completion problem also reminds us that rural communities may bear the brunt of unexpected complications tied to the IP transition, with potentially devastating consequences. The Commission has rightly recognized that this issue speaks to our foundational expectation that the phone network will be reliable for all Americans, including those in rural areas, and has opened a proceeding to learn more about exactly why the rural call completion problem is getting worse.[10] But even so, the FCC has received some shockingly inadequate carrier responses to rural call completion complaints. For example, one carrier told the FCC: "We have contacted the [rural complainant] and have successfully resolved this matter by advising [her] that due to living in a rural area she will experience service issues."[11]

This is why we need rules of the road: problems will inevitably arise as old systems fade away and new ones arise, but carriers have clearly shown that we cannot simply

---

[9] A bipartisan coalition of U.S. Senators has urged the FCC to act expeditiously to resolve the rural call completion problem. S. Res. 157, 113th Congress. For some, at least, waiting for ICC to zero out does not seem the preferable option.

[10] *Rural Call Completion*, Report and Order and Further Notice of Proposed Rulemaking, WC Docket No. 13-39 (rel. Nov. 8, 2013).

[11] FCC Enforcement Advisory, Rural Call Completion: Long Distance Providers Must Take Consumer Complaints About Rural Call Completion Problems Seriously (July 19, 2013), http://transition.fcc.gov/Daily_Releases/Daily_Business/2013/db0719/DA-13-1605A1.pdf.

assume that companies will voluntarily defend the fundamental principles that have made our communications networks great. Particularly where the states have effectively written themselves out of the conversation through deregulation, everyday Americans are relying on federal authorities as their sole defender to protect the reliable, affordable communications access they count on.

5. **Should we analyze interconnection policy differently for best-efforts services and managed services where quality-of-service is a desired feature? If so, what should be the differences in policy between these regimes, and how should communications services be categorized?**

Before we can know what specific interconnection rules may be necessary for various services, we must understand more about the IP interconnection market. This means relevant federal and state agencies must be able to collect data to inform themselves about interconnection agreements and interconnection disputes, and must be able to ensure interconnection continues to encourage universal service and competition.

Viewed in this light, recent peering disputes for Internet traffic simply represent the same sort of problem telecommunications policymakers have always faced regarding whether last mile providers (or some other point in the distribution chain) end up with disproportionate power because of network effects, termination monopoly issues, or other factors that make network economics work differently from general markets.

While we may not arrive at exactly the same interconnection rules for, for example, 9-1-1 voice service and general internet traffic, the same principle underlies both types of traffic. We must ensure carriers can connect to others' networks on reasonable terms for the benefit of customers on all carriers. Where disputes arise or technological complications occur, we need agencies that can step in to protect consumers and promote competition.

6. **Much of the committee's focus in the #CommActUpdate process has been on technology-neutral solutions. Is a technology-neutral solution to interconnection appropriate and effective to ensure the delivery and exchange of traffic?**

Interconnection and competition are technology-neutral enduring values of the network, and therefore should apply to all communications services. However, the specific approach to achieve those values may vary if there are material differences between technologies, now or in the future, that require different approaches. This does *not* mean that any service that uses Internet Protocol is automatically exempt from competition concerns or from consumer protections. Regardless of the technology, a network of networks will always need mechanisms to ensure those networks connect seamlessly.

On this point it is worth noting that the Communications Act in its current form is technology neutral on the point of interconnection. Rather than depending on the technology being used, 47 U.S.C. § 251 applies interconnection obligations based on the type of service a carrier offers. For example, the FCC could, if it chose, classify broadband internet access service as a telecommunications service, and the interconnection duties of § 251 would apply to broadband internet access providers just as they do to other telecommunications service providers now.[12] How those duties apply may look different depending on the technology at issue, but the end goal will remain the same.

Public Knowledge urges policymakers to take a technology-neutral approach to issues like interconnection, ensuring that networks have meaningful duty to connect to each other regardless of the protocols they use.

7. **Wireless and Internet providers have long voluntarily interconnected without regulatory intervention. Is this regime adequate to ensure consumer benefit in an all-IP world?**

8. **Is contract law sufficient to manage interconnection agreements between networks? Is there a less onerous regulatory backstop or regime that could achieve the goals of section 251?**

Complaints and disputes have already arisen in the wireless and IP interconnection markets.[13] We cannot assume that existing practices in wireless and IP interconnection will always be successful, because there is substantial reason to question whether they are always successful now. Even if they were, it is not enough to simply hope that nothing will ever go wrong with something as foundational as the ties that bind our communications networks together.

Even now, ISPs can use their gatekeeping status to manipulate interconnection in ways that create harmful impacts similar or identical to blocking or paid prioritization. This is true even in the presence of otherwise strong open internet rules. For example, when ISPs prevent direct interconnection to their networks (either from an edge provider's private network or from one of its preferred transit providers), edge providers and consumers suffer reduced performance, just as if they were subject to paid prioritization or

---

[12] *See* Comments of Public Knowledge, Benton Foundation, and Access Sonoma Broadband, *Protecting and Promoting the Open Internet*, GN Docket No. 14-28, *Framework for Broadband Internet Service*, GN Docket No. 10-127, *Preserving the Open Internet*, GN Docket No. 09-191, *Broadband Industry Practices*, WC Docket No. 07-52 (July 15, 2014), https://www.publicknowledge.org/assets/uploads/blog/Public_Knowledge_NN_NPRM_comments_2014_FINAL.pdf.

[13] *See Petition of Sprint Spectrum LP for Arbitration Pursuant to Section 252(b) of the Telecommunications Act of 1996 to Establish Interconnection Agreements with Michigan Bell Telephone Company d/b/a AT&T Michigan*, Case No. U-17349, *Joint Submission of Sprint Spectrum LP and Michigan Bell Telephone Company d/b/a AT&T Michigan, for Approval of an Interconnection Agreement*, Case No. U-17569, Order (Michigan Public Service Commission, Mar. 18, 2014); Comments of Netflix, Inc., *Protecting and Promoting the Open Internet*, GN Docket No. 14-28, *Framework for Broadband Internet Services*, GN Docket No. 10-127 (July 15, 2014).

blocking. While the technical explanation for this reduced performance may vary, the user experience of slowed speeds for certain content is identical.

Similarly, ISPs can use their control over interconnection to the last-mile network to appropriate oligopoly rents from edge providers or advantage affiliated content. Large content providers often employ privately operated or third-party CDNs, as well as preferred transit partners, to send traffic to residential ISPs to minimize transmission distance and time. ISPs are able to charge significant tolls to edge providers to the extent that they are able to identify and target certain content providers. Such targeting is particularly viable for companies that serve large amounts of data with high quality requirements. Netflix, for instance, has seen substantial provider-specific drops in service that may be the result of such practices.[14] This ability allows ISPs to shift the costs of network investment to edge providers, who then shift those fees to consumers through higher subscription costs. Similarly, ISPs are able to favor affiliated services via interconnection, not just by placing content closer to the end user but also by ensuring that preferred partners or services face no interconnection fees or tolls.

The evolution of our communications networks can be complicated, and no one party will always know the whole story. But understanding the whole story is crucial to ensuring the networks continue to deliver reliable service of sufficient quality to all networks users across the country. This is why Public Knowledge supports increased transparency and disclosure regarding interconnection agreements. The FCC's recent decision to investigate interconnection issues, in particular by obtaining and analyzing the content of recent paid peering agreements, is a promising first step.[15] However, voluntary disclosure is not a complete solution to addressing interconnection disputes, since companies can mislead policymakers and consumers by selectively releasing information. ISPs may also use language that deliberately obscures the consumers' ability to understand the nature of the agreements made and the implications for their service. Policymakers should therefore ensure mandatory, transparent disclosure in terms understandable to an average ISP consumer so that consumers, companies, and interested parties can be more fully aware of the relevant market issues and dynamics.

Conclusion

Interconnection will always be a key value underlying the success of our communications networks. Interconnection promotes universal service and competition. Interconnection ensures data transmission on the network operates seamlessly and without discrimination. Interconnection is what ties our communications networks

---

[14] *See* USA ISP Speed Index Results Graph, Netflix, http://ispspeedindex.netflix.com/results/usa/graph (last visited July 14, 2014) (particularly the results for Comcast between October 2013 and November 2014).

[15] Tom Wheeler*, Statement by FCC Chairman Tom Wheeler on Broadband Consumers and Internet Congestion*, FCC (June 13, 2014), http://transition.fcc.gov/Daily_Releases/Daily_Business/2014/db0613/DOC-327634A1.pdf.

together. As we move to new technologies, the need for a successful interconnection regime, with federal and state agencies able to ensure the network continues to function as it should, remains just as important as it was when the phone network was first connecting network users decades ago.

CONSULTING

August 6, 2014

Mr. David Redl
Chief Counsel – Communications and Technology
US House Committee on Energy and Commerce
2125 Rayburn HOB
Washington, DC 20515

Dear Mr. Redl:

On December 3, 2013, Chairman Upton and Chairman Walden announced their panels would embark on a year-long initiative to review and ultimately update federal telecommunications policy. On behalf of the Alaska Rural Coalition, CalCom Small Company Committee and Idaho Telecom Alliance (Rural State Association Group – RSAG), GVNW[1] submits the attached comments in response to the fourth white paper network interconnection questions from the Committee on Modernizing the Communication Act.

Alaska Rural Coalition. The companies of the ARC that are participating in this filing[2] serve customers in some of the most extreme regions of the United States. Alaska is a uniquely high cost area within which to provide any telecommunications, whether traditional telephony, mobile or broadband. Much of remote Alaska lacks even the basic infrastructure critical to most telecommunications deployment, such as a road system and an intertied power grid.

CalCom Small Company Committee. The California Communications Association (CalCom) is a statewide non-profit trade association with a rich heritage that dates back to 1917. Its small company members[3] are committed to the effort to build state of the art networks across California.

Idaho Telecom Alliance. The companies in the Idaho Telecom Alliance[4] work collectively to support the advancement of their members and promote services to rural telecommunications subscribers throughout the rugged terrain of Idaho.

---

[1] GVNW is a management consulting firm that provides regulatory and legislative advocacy support for communications carriers in rural America.
[2] The ARC members in this filing include Arctic Slope Telephone Association Cooperative, Inc.; Bettles Telephone, Inc.; Bristol Bay Telephone Cooperative, Inc.; Bush-Tell, Inc.; Copper Valley Telephone Cooperative, Inc.; OTZ Telephone Cooperative, Inc.; Alaska Telephone Company; and North Country Telephone, Inc.
[3] TDS is not participating in this White Paper 4 response.
[4] Fremont Communications is not participating in this White Paper 4 response.

The challenge facing the Committee as it seeks to rewrite federal telecommunications law is to enable broadband to truly be embedded in the national infrastructure while creating a framework of rates so that service and cost are reasonably comparable. This will require universal service provisions to continue to be an important part of any rewrite effort.

This series of questions on network interconnection is germane to the current update analysis. Over the last several years, smaller wireless carriers have been challenged to obtain roaming agreements (voice and data) with the larger national carriers. The FCC was forced to step in and issue mandates. There is no record evidence that leads us to believe that the larger carriers will somehow be more reasonable with respect to IP network transport. And the rural call completion issue is at a level that House Communications Subcommittee Vice Chairman Bob Latta (R, Ohio-05) and Peter Welch (D, Vermont) introduced House Resolution 536 to correct this problem.

In closing, we express our appreciation to Chairman Upton and Chairman Walden for initiating this important review of our federal telecommunications law. The legislative action related to this effort will impact every customer in each of the states.

Please call me on 503-612-4409 or contact me at jsmith@gvnw.com if you have any questions.

Regards,

s/JHS 8/6/14

Jeffry H. Smith
President and CEO

Copy to
Chairman Fred Upton, House Energy and Commerce Committee
Ranking Member Henry Waxman
Chairman Greg Walden, Communications and Technology Subcommittee
Ranking Member Anna G. Eshoo

Mr. Ray Baum

RESPONSE OF THE RURAL STATE ASSOCIATION GROUP (RSAG) TO HOUSE ENERGY AND COMMERCE COMMITTEE

Modernizing the Communications Act
NETWORK INTERCONNECTION QUESTIONS FOR STAKEHOLDER COMMENT

Due Date of August 8, 2014

We offer responses to the questions posed by the Committee by emphasizing four major points, and reference individual questions as needed in each section. Our four sections for this fourth white paper focusing on network interconnection issues are shown below.

THE ROLE FOR THE FCC IS TO REGULATE PURSUANT TO TITLE II ALL TRANSPORT AND TRANSMISSION CAPACITY OFFERED ON UNDERLYING NETWORKS

RURAL CALL COMPLETION ISSUES WILL STILL REQUIRE REGULATORY ATTENTION IN AN IP ENVIRONMENT

CHANGES IN NETWORK INTERCONNECTION SHOULD NOT BE USED AS AN EXCUSE FOR LARGE URBAN CARRIERS TO IMPOSE UNREASONABLE COSTS ON RURAL CARRIERS

EMERGENCY COMMUNICATIONS REQUIRE SPECIAL ATTENTION IN POLICY FORMATION

We appreciate the opportunity to offer input on these network interconnection issues and look forward to the remaining two white papers on media and public safety that the Committee intends to release during 2014.

THE ROLE FOR THE FCC IS TO REGULATE PURSUANT TO TITLE II ALL TRANSPORT AND TRANSMISSION CAPACITY OFFERED ON UNDERLYING NETWORKS

This section of our response addresses White Paper 4 questions[1] one, two, five and six. As a preface to this section, we share the foundational public policy principles[2] of the Rural State Association Group that guide our responses in every House Energy and Commerce white paper filing:

1 – Affordable broadband should be available to all Americans

2 – Federal universal service support should be sufficient and predictable

3 – Policies should promote competition while protecting consumers

4 – Public safety and national security should continue to be a priority

5 – Comparable rates for comparable services

The Committee must determine in its effort to modernize the Communications Act whether the desire to protect consumers in a broadband paradigm is still a relevant public policy concept. We respectfully submit that it remains an important cornerstone for any attempt to enact a forward-looking **national** public policy.

In order to ensure seamless interconnection across the entire network of networks, we support the concept introduced on July 18, 2014 by NTCA – The Rural Broadband

[1] Q1: In light of the changes in technology and the voice traffic market, what role should Congress and the FCC play in the oversight of interconnection? Is there a role for the states? Q2: Voice is rapidly becoming an application that transits a variety of network data platforms. How should intermodal competition factor into interconnection mandates? Does voice still require a separate interconnection regime? Q5: Should we analyze interconnection policy differently for best-efforts services and managed services where quality-of-service is a desired feature? If so, what should be the differences in policy between these regimes, and how should communications services be categorized? Q6: Much of the committee's focus in the #CommActUpdate process has been on technology-neutral solutions. Is a technology-neutral solution to interconnection appropriate and effective to ensure the delivery and exchange of traffic?

[2] We included the first 4 principles in our response to the first white paper and added the fifth in the third White Paper.

Association in the Federal Communications Commission (FCC) Open Internet proceeding (GN Docket No. 14-28). NTCA suggested correctly that one of the roles that the FCC should continue to play is to ensure that with respect to the **transport and transmission capacity on all networks over which data travel,** the networks will 1) be interconnected on just and reasonable terms; 2) that the FCC has the clear and unquestionable ability to step in when those networks do not interconnect seamlessly, and 3) that the important public policy goals of consumer protection, universal service, competition, and public safety are not threatened by the unjust and unreasonable acts or omissions of any given network operator.

NTCA offered a very narrow application of Title II as a means to this important end. So as not to evoke a reflexive negative reaction from any distinguished members of this Committee with the mere mention of Title II, let's first be clear about what is not proposed in this regard. What is not proposed is any overarching Title II authority. Simply stated, Title II would not be appropriately applied to any services that are offered atop of these regulated networks unless those services would qualify for that treatment on their own merits. There is no implied desire to "regulate the Internet", but rather a desire to avoid a problem larger than rural call completion presents today. Instead, as NTCA stated at page 11 of its FCC filing: "*Applying Title II to networks that merely transmit data between points does not hinder innovation in 'the Internet ecosystem' or saddle new services with legacy regulations. Rather, it simply ensures that the networks upon which Internet data travel will be interconnected on reasonable terms, that the Commission can step in when those networks do* **not** *interconnect seamlessly, and that important public policy goals of consumer protection, universal service, competition, and public safety are* *not threatened by the unjust and unreasonable acts or omissions of any given network operator."*

What is proposed by NTCA is a **limited and targeted** application of Title II regulation **specifically and only with respect to** the transport and transmission capacity of all networks and content delivery networks (CDNs) involved in the transmittal of data between points (the network layer as opposed to the service layer). What do we gain by making this assignment? Several notable benefits accrue to such an approach. Sections 201 and 202 provide a solid basis for carefully crafted rules to protect consumers. In Section 201, rules are in place that require service to be provided upon reasonable request and details a carrier's duty to interconnect. Section 202 provides rules that prohibit unjust and unreasonable discrimination. Sections 206, 207, and 208 provide a backdrop for the resolution of complaints and offer enforcement mechanisms. Title II as applied in this context would ensure that consumer connectivity is not lost or impaired due to a disagreement or dispute between underlying network operators.

RURAL CALL COMPLETION ISSUES WILL STILL REQUIRE REGULATORY ATTENTION IN AN IP ENVIRONMENT

Question 4 in this white paper[3] focuses on the troubling issue of rural call completion. At the end of the background section of this 4th White Paper, the Committee offers the observation that "*the federal government has been reluctant to engage in disputes regarding IP interconnection*." Quite frankly, while some discernible progress has indeed occurred, the FCC has been reluctant to bring to an immediate END the problem of rural call completion. It has been over three years since the FCC was made aware of the magnitude of this problem in a series of ex parte[4] letters, and reports from operating companies indicate the problem persists.

The very fact that the problem is still not solved despite FCC attempts to enforce existing rules is instructive for this Committee inquiry. Put simply, the "market" is not producing an equitable solution for rural customers in the nation. A need remains for targeted and reasonable regulation to protect both the rural consumer experience as well as promote universal service and public safety. Even with rules in place, problems persist. Removing basic rules would lead on a path to even larger problems and serve to thwart progress to achieving a truly national broadband platform.

As the Committee seeks to modernize the Communications Act, we respectfully request that public safety and universal service issues not be a secondary thought to how fast large carriers can shed regulation. Such a balance is a key to the public policy solution that will meet a multitude of customer needs.

---

[3] Ensuring rural call completion has always been a challenge because of the traditionally high access charges for terminating calls to high-cost networks. Does IP interconnection alleviate or exacerbate existing rural call completion challenges?

[4] Ex parte letter from NTCA CEO Bloomfield to FCC Chair Genachowski dated September 20, 2011, following up on a June 13, 2011 letter to FCC Enforcement Bureau from NTCA counsel and other parties.

CHANGES IN NETWORK INTERCONNECTION SHOULD NOT BE USED AS AN EXCUSE FOR LARGE URBAN CARRIERS TO IMPOSE UNREASONABLE COSTS ON RURAL CARRIERS

With the eighth question[5] posed in this White Paper 4, the Committee tees up the concept of allowing means other than current FCC rules to oversee the interconnection process.

In determining how best to modify legislative policy, it is important to review recent interconnection location proposals by AT&T and the implications that such a change would have on rural carriers and customers. In a January 24, 2014 ex parte[6] with the FCC, AT&T suggests that the model for both Tier 1 IP voice and peering interconnection is 5 to 8 interconnection points in total for the entire country. Under this proposed scenario, the use of fewer interconnection points covering much larger geographic areas would result in a significant increase in costs on rural ISPs and ultimately rural consumers and business customers. This increase is caused by the smaller providers having the full responsibility for transporting traffic to interconnection points a great distance from their facilities in such a proposed arrangement, in many cases over facilities owned[7] by large carriers such as AT&T.

We respectfully suggest to this Committee that underlying networks are not "free" in an IP-enabled paradigm any more than they are in a TDM world. Small or rural ISPs possess little or most likely NO bargaining power with respect to negotiating

[5] Is contract law sufficient to manage interconnection agreements between networks? Is there a less onerous regulatory backstop or regime that could achieve the goals of Section 251?

[6] AT&T's Director – Federal Regulatory filed an ex parte letter in GN Docket No. 13-5, WC Docket Nos. 13-97 and 10-90 that showed the essence of its proposal at presentation slide 11.

[7] Perhaps this helps explain why AT&T's July 15, 2014 presentation entitled "The Internet Interconnection Ecosystem" stresses at pages 15-18 that the carriage of traffic is not without cost and that there are cost implications of carrying additional traffic. It is comforting to see that our colleagues at AT&T have discovered that bill and keep methodologies may not be universally applicable.

interconnection terms with large national operators. Without something resembling "rules for the IP road" in place, such cost shifts will not be borne equitably across the networks and such an outcome will serve to drastically impact the goal of universal service in our emerging broadband world. One solution is that it may be necessary to distinguish[8] between large and small companies.

EMERGENCY COMMUNICATIONS REQUIRE SPECIAL ATTENTION IN POLICY FORMATION

Question 3 in this fourth white paper[9] states: "How does the evolution of emergency communications beyond the use of traditional voice support impact interconnection mandates?"

In the response to the third white paper, we asserted that competition is at a different point on the continuum if you compare the current state of affairs for customers in either Anaheim, California (population 343,298) or Annapolis, Maryland (population 38,620) versus the number of customers served by Arctic Slope Telephone Association Cooperative in Anaktuvuk Pass, Alaska, with a population of 332. The same may be said for issues related to emergency response. Just like in real estate, location matters. The more remote the location, the more crucial the access to emergency services. In portions

[8] In the RSAG response to the first White Paper, we offered a recommendation that regulation be bifurcated for large companies and small companies. Specifically, we stated on page 12 of our response to the first White Paper that: *Changes to FCC structure would necessarily follow with decisions made to the platform used to regulate carriers. For example, if the decision is made to shift from regulation of services to regulation by size of entity, then the Bureau designations at the FCC might well change to Large Company Oversight Bureau (LCOB) and Small Company Oversight Bureau (SCOB).*

[9] Question 7 is addressed tangentially here as well: "Wireless and Internet providers have long voluntarily interconnected without regulatory intervention. Is this regime adequate to ensure consumer benefit in an all-IP world?"

of Northern California, Siskiyou Telephone offers service to customers in rugged, mountainous terrain. They are the lifeline for these customers.

With respect to the wireless issues posed in Question 7, one of the RSAG members (Copper Valley in Valdez, Alaska) filed an ex parte with the FCC on June 18, 2014 detailing the importance of interconnection for rural wireless carriers in an account of how the Copper Valley Wireless cell site located at Cannon Hill in Chitina, Alaska, saved the life of a pilot and passenger who had crash-landed in a very remote location. As Copper Valley Wireless stated in part in their June ex parte: "*Copper Valley Wireless strives daily to deliver mobile wireless services to a corner of the United States that is both majestically beautiful but exceedingly rugged. Where our urban and suburban counterparts use buried electric lines, roof-top antennas, and service vans, Copper Valley Wireless engineers and technicians use propane generators, helicopters and mountain-top towers. But all these technological marvels come at a great cost – a cost that can only be overcome with sufficient*[10]*, but ultimately supplemental, universal service fund support. A remote cell tower deployed and maintained using universal service support allowed for the quick rescue of Mr. Cyr and the pilot on June 10, 2014.*"

[10] While interconnection is important, so too is universal service support. We will expand on this issue when the House Energy and Commerce Committee releases the upcoming white paper on public safety.