

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

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**

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

This document and any map included herein are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.



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

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

FOREWORD.....	2
MAIN POINTS	3
MAIN POINTS	6
INTRODUCTION AND EXECUTIVE SUMMARY	8
Performance of the Market	8
Structural Evolution of the Market	9
Legal and Policy Frameworks	10
Challenges for the Future.....	11
Challenges for Internet Architecture.....	11
Challenges for Technical Innovation	11
Challenges for the Internet model of traffic exchange.....	12
Challenges for Policy	12
Market Performance Measures	14
CHALLENGES FOR THE FUTURE.....	17
Architecture	17
Challenges for the Internet mode of traffic exchange.....	21
Incumbent perspective	21
Entrant perspective.....	21
Policy maker perspective	22
Existing frameworks and applications of regulation.....	22
ONGOING DEVELOPMENT OF THE MARKET	25
Adaptations of the Internet model.....	25
Regional peering	25
Multilateral peering	25
Structural change in the Internet.....	26
Share-shifting backbone carrier.....	26
Growth of CDNs.....	27
Share-shifting services	27
Investment in access networks	28
NEW MODELS FOR POLICY	30
Establishing a bright line	30
The question of voice.....	30
Reform of TDM interconnection	32
Allowing for disagreement	32
Pricing models to support future investment	32
Traffic exchange and network neutrality	34
REGIONAL SURVEY.....	36
Africa	36
Latin America	37
The Caribbean.....	39
The Middle East.....	39
Europe.....	40

Canada and the United States	42
The Pacific	42
Australia and New Zealand.....	43
South Asia.....	44
East Asia	44
Regional and National Trends	46
APPENDIX 1: NATIONAL INTERNET STATISTICS	49
APPENDIX 2: INTERNET EXCHANGE POINT REGIONAL FIVE-YEAR STATISTICS	54
APPENDIX 3: INTERNET EXCHANGE POINT NATIONAL FIVE-YEAR STATISTICS	55
APPENDIX 4: COUNTRIES AND TERRITORIES STILL LACKING AN INTERNET EXCHANGE POINT	59
ANNEX 1	61
SURVEY OF CHARACTERISTICS OF INTERNET CARRIER.....	61
INTERCONNECTION AGREEMENTS	61
The Survey	61
Informal agreements	62
Symmetric terms	62
Governing law.....	63
National interconnection partners	63
Degree of interconnection.....	65
Unexpected results	66
Further work necessary	67
ANNEX 2 REGIONAL PEERING	68
ANNEX 3 CLOUD COMPUTING.....	71
ANNEX 4 WHO PAYS FOR WHAT?.....	73
ANNEX 5 IPV4 ADDRESSES AND THE FUTURE GROWTH OF IXPS.....	77
ANNEX 6 PRACTICAL IMPLEMENTATION: MECHANISMS AND PRACTICES	78
Internet exchange points	78
Domain name service.....	79
Local content.....	80
Competition and regulation.....	80
International connectivity	81
Voice-over-IP and ENUM	81
IPv4 and IPv6 addresses and transition.....	82
ANNEX 7 WHY HAS THE INTERNET MARKET PERFORMED SO WELL?.....	83
ENDNOTES	90

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 2002¹ and again in 2006². 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 2006 OECD 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 very 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 2011 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 as 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 parameters, 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% utilization 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 "Quavvis," 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 “de-privatisation” 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 neighbor,

with three new IXPs founded in Neuquén, Rosario, and Bahía Blanca in 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. Wholesale 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 Bangladesh 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.

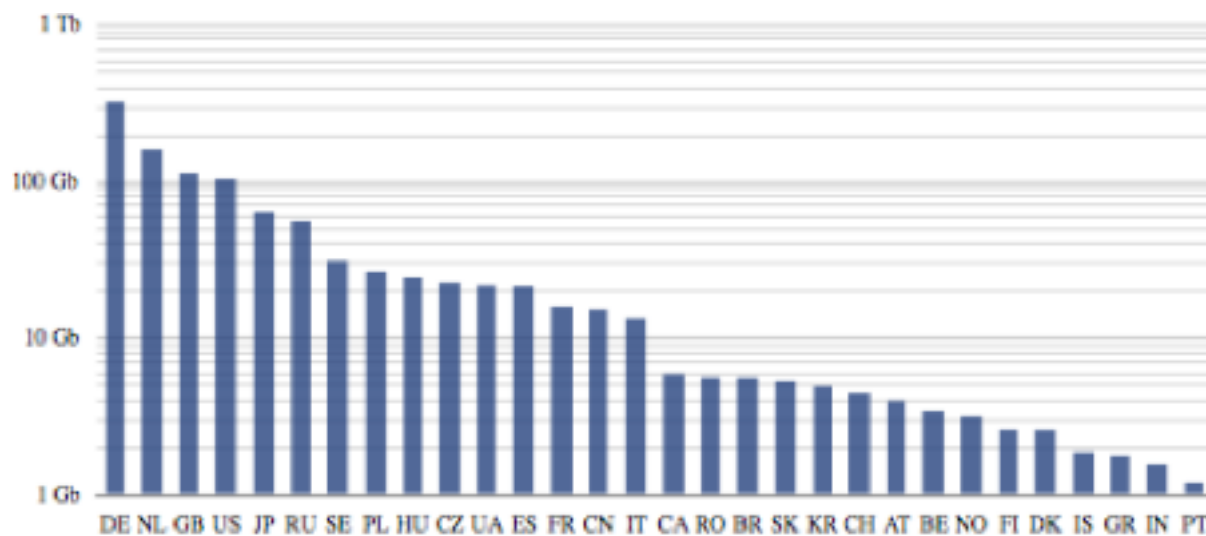
Japan continues to lead the Asian region in most Internet statistics, due to a long head start and a combination of relatively aggressive investment in research and development and a high demand for Internet services in their domestic population. In addition, Japan's high concentration of transpacific 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

Country	IPv4 Addresses				IPv6 Addresses				Autonomous System Numbers				Network Adjacencies			
	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												

DSTI/ICCP/CISP(2011)2/FINAL

	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				Autonomous System Numbers				Network Adjacencies			
Iraq	/16	/15	+16	+66%		/31	+31	new	5	11	+6	+120%	9	24	+15	+167%
Ireland	/10	/10	+14	+7.2%	/27	/27	+30	+19%	107	114	+7	+6.5%	289	333	+44	+15%
Isle of Man	/17	/17	+21	+9.2%	/32	/32	+48	+0%	16	17	+1	+6.3%	119	123	+4	+3.4%
Israel ⁷⁵	/10	/10	+13	+16%	/29	/29	+31	+29%	229	250	+21	+9.2%	738	929	+191	+26%
Italy	/7	/7	+11	+8.2%	/20	/20	+27	+0.55%	565	617	+52	+9.2%	1910	2306	+396	+21%
Jamaica	/15	/15	+17	+19%	/32	/32			9	9			27	38	+11	+41%
Japan	/5	/5	+9	+5.5%	/19	/19	+21	+27%	762	795	+33	+4.3%	2843	3092	+249	+8.8%
Jersey	/17	/17			/32	/32			4	4			22	23	+1	+4.5%
Jordan	/13	/13	+16	+18%	/31	/29	+30	+250%	24	26	+2	+8.3%	64	77	+13	+20%
Kazakhstan	/12	/11	+13	+49%	/32	/31	+32	+100%	56	67	+11	+20%	222	353	+131	+59%
Kenya	/13	/12	+12	+202%	/30	/29	+30	+100%	35	49	+14	+40%	83	171	+88	+106%
Kiribati	/20	/20														
Korea	/6	/5	+7	+33%	/20	/20	+29	+0.12%	862	916	+54	+6.3%	2795	2906	+111	+4%
Kuwait	/12	/12	+14	+25%	/32	/30	+31	+200%	37	41	+4	+11%	137	149	+12	+8.8%
Kyrgyzstan	/15	/15	+19	+6.3%	/32	/32			13	14	+1	+7.7%	33	62	+29	+88%
Laos	/17	/16	+19	+21%		/30	+30	new	6	9	+3	+50%	8	12	+4	+50%
Latvia	/12	/12	+15	+6.9%	/30	/28	+29	+140%	198	219	+21	+11%	494	554	+60	+12%
Lebanon	/14	/13	+15	+33%	/30	/29	+32	+20%	39	45	+6	+15%	108	150	+42	+39%
Lesotho	/18	/18				/32	+32	new	4	4			6	5	-1	-17%
Liberia	/22	/19	+19	+800%					1	2	+1	+100%	1	2	+1	+100%
Libya	/14	/14	+20	+1.4%					1	1			5	3	-2	-40%
Liechtenstein	/16	/16	+21	+3.8%	/31	/31			7	8	+1	+14%	36	44	+8	+22%
Lithuania	/11	/11	+17	+1.2%	/30	/29	+30	+125%	108	118	+10	+9.3%	295	314	+19	+6.4%
Luxembourg	/13	/13	+17	+5.8%	/28	/28	+30	+23%	30	35	+5	+17%	109	142	+33	+30%
Macao	/14	/14	+16	+23%	/31	/31			4	4			9	9		
FYR Macedonia	/13	/13	+20	+0.92%	/30	/30			23	24	+1	+4.3%	100	109	+9	+9%
Madagascar	/17	/16	+18	+42%	/32	/32			3	3			4	6	+2	+50%
Malawi	/18	/16	+17	+319%					4	7	+3	+75%	3	9	+6	+200%
Malaysia	/10	/10	+12	+26%	/27	/27	+28	+48%	89	107	+18	+20%	316	421	+105	+33%
Maldives	/17	/16	+18	+44%		/32	+32	new	3	3			4	6	+2	+50%
Mali	/17	/17	+18	+50%	/32	/32			4	4			5	8	+3	+60%
Malta	/13	/13	+15	+24%	/30	/29	+32	+20%	22	25	+3	+14%	64	94	+30	+47%
Mariana Islands	/18	/18							1	1			5	3	-2	-40%
Marshall Islands	/21	/21				/32	+32	new	1	1			5	2	-3	-60%
Mauritania	/17	/17							1	1			4	2	-2	-50%
Mauritius	/14	/13	+14	+100%	/30	/29	+30	+75%	13	16	+3	+23%	7	11	+4	+57%
Mexico	/8	/7	+10	+23%	/28	/27	+28	+143%	238	253	+15	+6.3%	663	740	+77	+12%
Micronesia	/21	/19	+20	+200%	/32	/32			1	3	+2	+200%	6	5	-1	-17%
Moldova	/13	/13	+15	+39%	/30	/29	+30	+75%	28	37	+9	+32%	106	155	+49	+46%
Monaco	/16	/16			/32	/32			1	1			5	3	-2	-40%
Mongolia	/15	/14	+16	+50%		/30	+30	new	27	34	+7	+26%	56	108	+52	+93%
Montenegro	/15	/15	+17	+45%					9	11	+2	+22%	26	32	+6	+23%
Montserrat	/22	/22														
Morocco	/12	/12	+14	+27%	/31	/31			7	7			20	19	-1	-5%
Mozambique	/15	/15	+19	+5.1%	/32	/32			14	16	+2	+14%	14	75	+61	+436%
Myanmar	/18	/18				/32	+32	new	2	2			3	7	+4	+133%
Namibia	/15	/15	+18	+11%	/48	/48			9	11	+2	+22%	34	44	+10	+29%
Nauru	/19	/19				/32	+32	new		2	+2	new				
Nepal	/15	/14	+15	+121%	/31	/29	+29	+300%	28	32	+4	+14%	73	137	+64	+88%
Netherlands	/8	/7	+12	+5.3%	/23	/23	+26	+9.4%	444	502	+58	+13%	2299	2530	+231	+10%
Netherlands Antilles	/14	/14	+18	+5%	/30	/29	+30	+100%	23	26	+3	+13%	41	47	+6	+15%
New Caledonia	/16	/15	+17	+56%	/32	/30	+30	+400%	4	4			6	6		
New Zealand	/9	/9	+13	+12%	/27	/26	+27	+82%	268	306	+38	+14%	607	1096	+489	+81%
Nicaragua	/14	/14	+18	+9.6%	/32	/32			14	16	+2	+14%	36	23	-13	-36%

DSTI/ICCP/CISP(2011)2/FINAL

	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

Region	Internet Exchange Points				Domestic Bandwidth Production			
	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

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



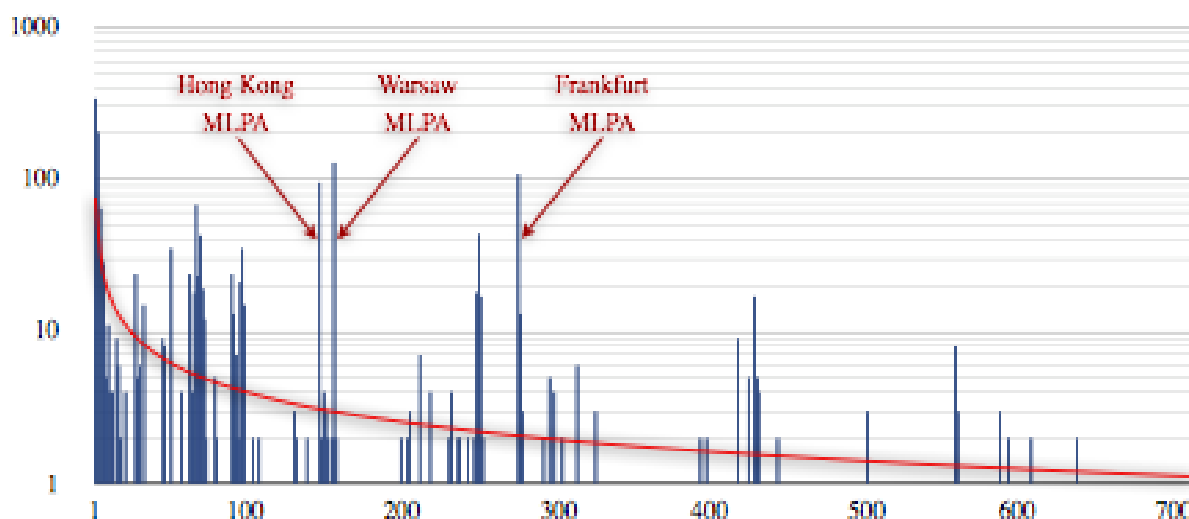
Source: Packet Clearing House.

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.

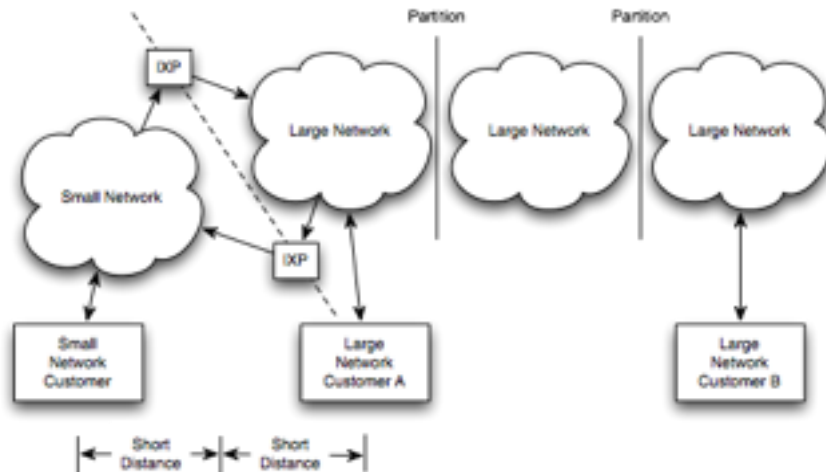
Figure A2.A. Regional Peering



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 over-investment 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 CO₂.⁸⁷ Meanwhile, each Google search has been calculated to generate 5–10 g of CO₂.⁸⁸

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.

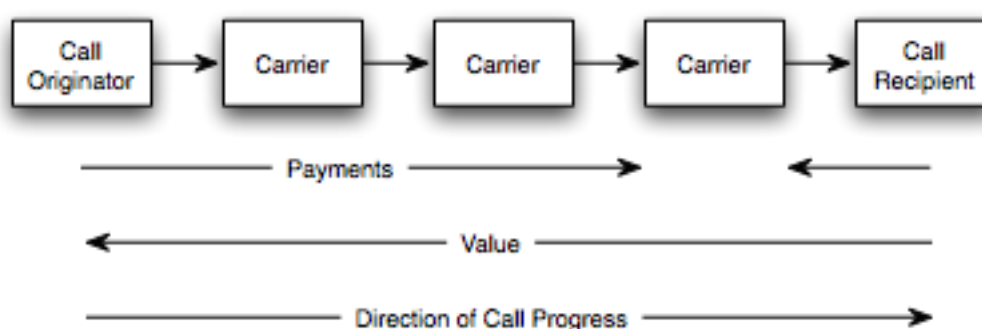
Because it aggregates supply efficiently and provides a high degree of liquidity in the marketplace for computational and storage resources, cloud computing is valuable and is unlikely to be a passing fad. At the same time, it poses substantial security challenges, so it is probably not reasonable to 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 jockeyed 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 “hot potato 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/info/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](http://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 Sarkozy 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](http://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.kcnep.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 States 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.