Organisation de Coopération et de Développement Economiques
Organisation for Economic Co-operation and Development

OLIS : 22-May-1998
Dist. : 26-May-1998

Or. Eng.

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

DSTI/ICCP/TISP(98)1/FINAL
Unclassified

Working Party on Telecommunication and Information Services Policies

INTERNET TRAFFIC EXCHANGE: DEVELOPMENTS AND POLICY

65832

Or. Eng.

TABLE OF CONTENTS

Tables

FOREWORD

In March 1998 this report was presented to the Working Party on Telecommunications and Information Services Policy (TISP) and was recommended to the made public by the Information, Computer and Communications Policy (ICCP) Committee.

The report was prepared by Dr. Sam Paltridge of the OECD's Directorate for Science, Technology and Industry. It is published on the responsibility of the Secretary-General of the OECD.

MAIN POINTS

Discussion of Internet traffic exchange is important because some parts of the communication industry are asking governments to address an increasing number of regulatory issues. Local telecommunication carriers are seeking redress for inequities they believe are arising from the final delivery of Internet traffic for dial-up users in respect to regulation aimed at telephony. Some smaller Internet service providers (ISPs) claim that arrangements for traffic exchange with larger ISPs are not transparent and non-discriminatory. Some international facilities-based ISPs argue that they are not being fairly compensated for carrying traffic for other parties. Finally, some ISPs are calling for professional management of the Internet's global root name servers that provide crucial information for the Domain Name System (DNS). All these issues will require decisions by policy makers, even if those decisions ultimately give governments a minimal, but supportive, role in finding solutions, for example, by providing a neutral forum for industry to reach consensus or by providing a neutral point at which statistics can be aggregated to inform debate.

The Internet is turning traditional point-to-point communication models on their head. For example, in the world of public switched telecommunication networks (PSTN), there was generally a correlation between the amount of traffic exchanged between two countries and the amount of capacity allocated by infrastructure providers for this route. For the Internet, the points of origin and termination of packets of data have very little relationship either to the traffic that may be carried over a given route or to traditional PSTN traffic patterns. For example, PSTN traffic between Singapore and Canada makes up well under one per cent of each country's outgoing international traffic. In December 1997, by way of contrast, nearly half of the international Internet capacity from Singapore's Internet exchange point (STIX) was deployed between Singapore and Vancouver. Similarly, in the PSTN world, telephone numbers were administered on a national basis. For the Internet, more than two-thirds of second-level domain names registered by Canadian, French and Spanish users are administered in the United States.

The Internet has evolved in a fundamentally different way from the world's PSTNs. National PSTNs were established locally and connected internationally many decades later. The reverse is largely true for the Internet. Until relatively recently, an e-mail message between two users located in the same building, but with different ISPs, was likely to travel over inter-continental backbone networks and receive DNS information from a server located on the other side of the globe. In some countries, Internet performance is better on an "international" than on a "national" or "local" basis, and in a number cases, this is directly related to past, and present, regulatory obstacles to infrastructure competition.

Due to the Internet's evolution, as well as the price and availability of national and regional infrastructure, a great deal of traffic between users in one geographical area is being pushed onto international or intercontinental infrastructures, particularly where this improves performance. However, this may negatively affect performance for users in other countries by creating congestion in traffic exchange (including domestic traffic exchange) in those countries. Local infrastructure performance is no longer isolated from local and regional performance elsewhere in the world. Consequently, governments

need to increase the momentum for liberalising infrastructure provision, notably by implementing the world Trade Organization (WTO) agreement, not least because poorer levels of performance in the local infrastructures of other countries affect their own markets.

The private sector, and in particular ISPs, are actively developing infrastructure to "localise" Internet traffic flows by improving the performance of local infrastructure. Their initiatives include the establishment of an increasing number of Internet exchange points (the places where ISPs exchange traffic) and greater global distribution of the infrastructure supporting the DNS. In these cases "localisation" does not indicate a specific geographic area but rather the fact that content, services, and some network functions are being shifted closer to the user to increase network efficiency (e.g. to bring users faster response times). These developments are critical in improving the performance of the Internet for electronic commerce. This document aims to increase awareness of how traffic exchange occurs on the Internet and how infrastructure and financial interconnection arrangements (including peering and transit) to accomplish traffic exchange are developing. The role of government in these developments is to eliminate existing and potential barriers by monitoring and reforming regulation as appropriate.

Because of the international nature of the Internet and the rapid pace at which the services and infrastructure are developing, some of these issues are complex. Moreover, the nature of traffic exchange on the Internet, relative to the PSTN, does not easily lend itself to past commercial settlement models. For example, if the traffic flowing between a country in the Asia-Pacific region and the United States were contained within a single international link and within the geographical borders of both countries, a model for sharing infrastructure costs might be self-evident. However, the Internet does not transport traffic in such a precisely defined or bounded way. Not only might the IP (Internet Protocol) packets travel along different paths and through different countries, so might the DNS information drawn from global root servers.[1]

From the perspective of ISPs in the United States, the greater the capacity foreign ISPs put in place to Internet exchange points based in the United States, the greater their costs in providing domestic infrastructure that is used for international transit. From the perspective of ISPs outside the United States, the existing peering model (akin to "sender keeps all" in the PSTN world) does not fairly compensate them for costs incurred in carrying the traffic of US users. There is little information available for understanding the balance of these costs, or even the patterns of international traffic (including transit traffic and traffic related to international DNS requests), although such information is necessary to inform discussion of the financing of international infrastructure linking OECD countries. At this stage, the best way forward is for industry to initiate discussion on the financing of Internet traffic exchange, for example, via the Asia-Pacific Internet Association’s call for comments and other industry forums. The role of government is to stay abreast of these discussions and support industry-led solutions.

INTRODUCTION

The Internet has evolved in a fundamentally different way from the world's public switched telecommunication networks (PSTN). National PSTNs were established locally and connected internationally many decades later. For these public networks, the first telephone exchanges defined the calling opportunities and the resulting traffic patterns. When these exchanges became connected on a national, and then international, basis traffic patterns started to change. At that point, subscribers to one local exchange could call, for the first time, those of another. The reverse is largely true for the evolution of the global Internet.

The Internet evolved from an international "private" network first used by the US military, and then by academic and research institutions. The Internet's core infrastructure, such as the first network access points (NAPs) and global root-level servers, were mostly established in locations far remote from the bulk of the initial users. This was made possible, of course, because the Internet has been grafted onto the world's PSTN infrastructure via lines leased by Internet service providers (ISPs) or capacity allocated by public telecommunication operators (PTOs) for their own traffic using the Internet protocol (IP). To exchange traffic, and to receive Domain Name System (DNS) information from far distant root-level servers, ISPs connected to the original NAPs using national and international leased lines.

To contrast initial Internet traffic flows with the PSTN consider the example of two users in Boston, Massachusetts, making a local telephone call or exchanging an e-mail at the beginning of 1997. In the case of local telephony, the call would have been switched via an exchange in Boston and the two parties would be connected without the traffic, or signalling information, needing to go outside the local calling area. In the case of an e-mail between the same users (as customers of different ISPs) the traffic would most likely have been routed via NAPs in Washington, DC or New Jersey. These were two of the original four NAPs funded by the US National Science Foundation (NSFNET). A similar IP exchange between two users in London, until 1995, would very probably have traversed one of these same NAPs in the United States[2] At the same time the DNS information, may have been drawn from one of a number of the nine original global root-level servers in the United States or Sweden at the core of the domain name system (DNS).

For a number of reasons, which will be examined below, Internet infrastructure providers are endeavouring to "localise" IP traffic and the DNS functions provided by the global root-level servers. Accordingly, a great many new Internet exchange points are now being established in major US cities to augment the original National Science Foundation (NSF) sponsored sites. By way of example, Boston had its first Internet exchange point established in 1997 so that traffic between Bostonians could be exchanged locally between ISPs operating in Massachusetts. At the same time new Internet exchange points were being established in Manchester (1997), Milan (1996) and Grenoble (1997) to compliment earlier exchange points in London, Rome and Paris. Similarly, in 1997, two global root-level servers were transferred from the United States to Japan and the United Kingdom.

In this way, in contrast to the way PSTNs were built up from local exchanges, the public Internet is gradually seeing exchange points built on more localised basis. This process is critical to efforts to upgrade information infrastructure in ways that will make it more suitable for electronic commerce. Yet, to date, there has been little discussion of its importance and what, if any, might be the implications for communication policies supporting the development of infrastructure for electronic commerce. This source document aims to describe the process that is under way in "localising" Internet traffic exchange and the policy issues that are arising. Before doing this it is necessary to set the context.

The OECD's first consideration of IP traffic flows and infrastructure development was based on a discussion of "webcasting" and some of the policy implications of convergence.[3] That discussion noted that the locality of the most accessed content over the World Wide Web was concentrated in the United States, and more specifically in California. The background report highlighted a number of initiatives that are under way to take this content closer to users and improve network efficiency such as caching, IP multicasting, mirror sites and digital warehousing.

This document does not intend to repeat the analysis of IP traffic flows in relation to webcasting but rather to build on that work. However, several additional factors should be considered owing to the dynamic nature of developments. One factor is that deployment of streaming media technologies used for webcasting is proceeding very rapidly. Based on data collected by RealNetworks, a leading supplier of webcasting tools, there were 103 000 web pages using streaming media in September 1997, with 11 per cent having a video component.[4] By November 1997 they had grown to 178 000 pages with 14 per cent having a video component. Adding to the increase in IP traffic, due to the webcasting of World Wide Web content, is the expected increase in the use of multimedia e-mail.[5]

While accounting for a smaller percentage of IP traffic than the World Wide Web, the use of e-mail is expanding at a prodigious rate. The major reason for this is the increased use of e-mail for social and commercial reasons, as well as the increase in "spam" (the Internet's equivalent of "junk mail" in the postal system). In December 1997, Deja News, a site which archives postings to Usenet newsgroups, said it was collecting about 730 000 messages per day. Deja News noted "...that nearly two-thirds of the content added to Usenet newsgroups daily is spam, or messages sent to cancel spam, creating a major impediment to information access and causing frustration for users".[6]

Including audio and video clips in e-mail communication, even if it remains unopened as in the case of some spam, can be expected to further increase IP traffic. Many of the solutions aimed at taking content closer to users are not designed to deal with such developments. Therefore, the process of localising traffic exchange, where possible, takes on further significance. Until the process of establishing new Internet exchange points in OECD countries is further developed, a great deal of "regular" and "multimedia" IP traffic will continue to traverse continental and intercontinental backbone networks. Infrastructure providers are responding by increasing the amount of capacity available on national and international routes at an unprecedented rate. In December 1997, MCI Communications said it doubled the core circuit capacity of its Internet backbone to dual 622 megabits per second and planned to increase its backbone core circuit capacity even further to 2.5 Gbits per second by year-end 1998. ISPs are increasing the capacity of their international backbones at a similar pace. However, Internet traffic continues to spiral, placing increased traffic demands on existing NAPs and other exchange points.

There is some evidence that some US backbone networks connecting to major NAPs, such as MAE East and MAE West, experienced deteriorating performance during 1997.[7] One study, by Keynote Systems, found the average performance of Internet backbone networks in North America, as measured by page download speeds, decreased by 4.5 per cent during 1997.[8] Significantly, the backbone network with the best performance in the Keynote survey, SAVVIS, has a strategy of bypassing "public NAPs" such as

MAE East by purchasing direct connections to larger backbone networks for IP traffic exchange.[9] The term "public NAPs" or public Internet exchange point is not used to indicate ownership status but rather that any ISP that meets certain requirements may exchange traffic at these points. Private Internet exchange points, such as those being used by SAVVIS, are direct connections with larger backbone networks.

While communication policy makers have a long-standing interest in promoting regulatory frameworks that encourage infrastructure development, this is being given added impetus by the growing requirements for electronic commerce. In terms of the Internet some studies suggest performance problems are generally not located on user's sites (e.g. server problems). For many Web sites used for electronic commerce, Keynote System's measurements indicate that most of the performance problems occur in the Internet's infrastructure somewhere between the Web site and its users, such as at the public exchange points and NAPs where backbone providers interconnect, in one or more routers along the communication path, or in the DNS.[10] How much can be concluded from different efforts to measure backbone network performance is currently a contentious issue. Some critics of Keynote's methodology claim their results reveal more about the performance of different backbone provider's Web sites than their about actual backbone networks.

The amount of traffic generated by the DNS may also be considerable but may depend on where samples are taken. One survey of Internet usage, undertaken in late 1997, found that DNS requests accounted for 24.4 per cent of packet traffic and 9.9 per cent of byte volume on one segment of the Internet backbone -- the largest single category of Internet traffic.[11] Bellcore says this is because, "Every time a Web site is searched or an e-mail is sent the correct IP address must be found and validated -- a process that can take several round trips to different DNS servers anywhere on the Internet."[12] One way to reduce this traffic is to "localise" DNS information by caching the most widely used IP addresses. Accordingly, in December 1997, Sun Microsystems and Bellcore announced the availability of the Soliant Advanced DNS system, which, by enabling ISPs to cache the most accessed IP addresses, reduces IP name searches to one trip. Sun and Bellcore say this will free up bandwidth and quicken response times for users. They believe widespread deployment could enable users to increase the speed with which they access data from a Web site by 10 to 15 per cent. Another study of traffic, undertaken by MCI in 1997, found DNS traffic made up a much smaller proportion of total packets and volume. MCI found that on domestic US and an international (US - UK) link DNS traffic comprised between 3-5 per cent of the packets and 1-2 per cent of bytes but as much as 25 per cent of the flows.[13]

To the extent that the initial "public NAPs" in the United States are points of congestion, the problem is compounded by the fact that many international networks exchange traffic both with the United States and "third country" ISP networks at these points. It has been reported that more than half of intra-European and intra-Asian IP traffic is transported via the United States.[14] While little data are available to provide confirmation of the exact amount of intra-regional traffic passing via the United States it is almost certainly very considerable. Furthermore, the great bulk of intercontinental traffic between the Asia-Pacific region and Europe is transported across the United States. Where data are available they show far more traffic being shipped from the United States to other countries than vice versa. For example, on average, Swisscom's transatlantic Internet link carries six times as much data from the United States as it carries to it.[15]

There are a number of reasons why the United States has developed as a "global hub" for Internet traffic. As will be examined below, these include historical factors (the original core infrastructure was in the United States), as well as the economic and network performance incentives some foreign ISPs have to route international transit traffic via the United States. In some cases, this is because a lack of infrastructure competition has led to prohibitively priced intra-regional capacity in

Europe and Asia. In other cases, it could be because telecommunication carriers are withholding capacity from competitors in adjacent value-added markets such as Internet access.[16] It has also been suggested that because some telecommunication carriers do not exchange traffic at national Internet exchange points, for reasons of commercial strategy or regulatory obstacles, this forces a great deal of domestic traffic to be unnecessarily routed via the United States.[17]

A further factor which might be researched by the Internet community is how much traffic flows back and forth between countries due to requests for DNS information. The results cited by Bellcore and MCI give different indications of the amount of traffic which may be generated between foreign countries and the United States, because the latter is the location for ten of the thirteen global root level servers. A further question is whether greater use of generic Top Level Domains (gTLD such as **.com, .net, .org**) in some countries outside the United States, as opposed to national TLDs (such as **.fr** for France), impacts on the balance of traffic between those countries and other countries where the global root level servers are located. Requests for root servers supplying DNS information for TLDs are a further factor, with around half of the 172 DNS servers being located outside the country concerned. Generally, DNS servers keep track of the round trip times for DNS packets to each of the root servers and will choose to query the server that responds most quickly over time. When traffic traverses international links, it would be the root server that is the closest according to network topology. The traffic impact is essential because once a server queries for a certain domain name, it caches the response and does not need to contact the root again until the cached value expires.

Two other factors are also very important in understanding why the United States is a global Internet hub. First the location of the most accessed "non-local" content is overwhelmingly located in a handful of US States such as California. Second, using an analogy from the PSTN, the greatest number of "calling opportunities" are in the United States. In other words, the greatest number of visible hosts -- the greatest number of connected users and range of content (as corroborated by a similar survey of servers) -- on the public Internet are in the United States. This is significant because the balance of Internet hosts in the United States is significantly ahead of its share of telecommunication mainlines (Table 1). Whereas the United States has just under a quarter of the world's telecommunication mainlines it has just over half of all public Internet hosts.

The US share of Internet hosts is smaller than the share calculated for that country in the past. The reason is that in times past there were very little data on the public record as to how many registrations under gTLDs were made by users outside the United States. Accordingly, the hosts surveyed under gTLDs were assigned to the United States. For network planners this made global surveys of Internet hosts less useful than they might otherwise have been in those countries that make greater use of gTLDs. This is particularly true for Canada, which is the largest user of gTLDs outside the United States. While surveys by Statistics Canada reveal that the number of Canadian households with access to the Internet nearly doubled (843 000 to 1 500 000) between 1996 and 1997, the number of hosts under **.ca** only increased by 62 per cent.[18]

By making an allowance for users of gTLDs in other countries, it is possible to give a more accurate estimate of Internet penetration where official data are not available. This enables policy makers to have a better understanding of how the Internet is developing and to inform discussion of Internet traffic patterns and exchange, including what, if any, impact registration patterns may have on traffic patterns generated by DNS requests. Accordingly, in the concluding sections, this document discusses domain name registration and its impact on Internet host statistics and traffic exchange. This analysis shows that while the United States accounts for half of the hosts accessing the Internet, the distributions of hosts in other countries is significantly higher that it has been previously possible to determine. When this

fact is added to the high concentration of popular content in the United States, it helps to understand the current trend of spiralling international traffic loads. In other words, there are more users outside the United States making use of US content than was previously evident.

This document also discusses national and international traffic exchange and its relationship to Internet exchange points and the international DNS infrastructure. It is very important for policy makers to understand better how traffic exchange occurs on the Internet. One reason is the need to increase the momentum of introducing infrastructure competition. It is ironic that efforts to restrict the opening of national markets has actually pushed increases in IP traffic, sometimes including purely domestic traffic, into being hubbed in foreign countries with the most competitive communication markets.

Policy makers need to ensure that there are no regulatory barriers restricting incumbent telecommunication carriers and private networks from exchanging traffic at public and private Internet exchange points. Particular issues to review are how existing regulation might impact on an incumbent telecommunication carrier's ability to enter into selective peering arrangements as opposed to fulfilling non-discrimination requirements. Similarly, there is a need to ensure that regulation does not have the unintended consequence of designating some private or user networks as facilities-based carriers if they participate at Internet exchange points.

A further initiative some OECD governments might like to consider is the promotion of co-operation between ISPs at the sub-regional and municipal level in the creation of Internet exchange points. In those countries that do not have Internet exchange points, or only have exchange points in the largest city, government authorities at the state or provincial level could stimulate awareness of the benefits of ISPs co-operating to establish local points of traffic exchange. While industry will no doubt accomplish this in the medium term, because the economic incentives are very large, many thousands of ISPs are extremely small and may not be well placed to initiate such efforts. In the case cited above the City of Boston played a role in supporting the establishment of a metropolitan Internet exchange point.

Finally policy makers need to increase their awareness of Internet traffic exchange because of the increasing number of issues that some parts of the communication industry are asking them to address. These include some local telecommunication carriers seeking redress for inequities they believe are arising from the final delivery of Internet traffic for dial-up users with respect to regulation aimed at telephony; some smaller ISPs claim that arrangements for traffic exchange with larger ISPs are not transparent and non-discriminatory; some international facilities-based ISPs arguing that they are not being fairly compensated for carrying traffic for other parties; and some ISPs calling for professional management of the Internet's global root name servers which provide critical information for the DNS. All these issues will require decisions by policy makers, even if those decisions are ultimately for governments to play a minimal role in finding solutions -- such as by providing a neutral forum for industry to reach consensus or by providing a neutral point at which statistics can be aggregated to inform debate.

Table 1. **Internet and PSTN Development**

	Share of world's Internet hosts adjusted for gTLD registrations (per cent)	Share of world's telecommunication mainlines (per cent)	Ratio
United States	51.5	23.8	1.2
Canada	6.1	2.6	1.4
United Kingdom	5.5	4.2	0.3
Germany	5.2	5.8	-0.1
Japan	5.2	8.8	-0.4
Australia	3.8	1.3	1.9
Netherlands	2.1	1.2	0.8
France	2.0	4.7	-0.6
Finland	1.8	0.4	3.4
Sweden	1.8	0.9	1.1
Italy	1.4	3.6	-0.6
Norway	1.1	0.4	2.2
Switzerland	1.0	0.6	0.5
Denmark	0.9	0.5	0.9
Spain	0.9	2.2	-0.6
New Zealand	0.8	0.2	2.4
Korea	0.8	2.7	-0.7
Austria	0.5	0.5	0.0
Belgium	0.5	0.7	-0.2
Czech Republic	0.3	0.3	-0.3
Ireland	0.2	0.2	0.1
Hungary	0.2	0.3	-0.4
Poland	0.2	0.8	-0.7
Mexico	0.2	1.3	-0.8
Turkey	0.2	2.0	-0.9
Portugal	0.1	0.5	-0.8
Greece	0.1	0.7	-0.8
Iceland	0.07	0.02	2.4
Luxembourg	0.03	0.03	-0.2
EC	23.0	26.1	-0.1

1. Host data is from Network Wizard's July 1997 Survey. Mainlines data is from the ITU for 1996. Weighting methodology is described in later sections of this document and registration data is provided by Imperative.

Source: OECD

NATIONAL AND INTERNATIONAL TRAFFIC EXCHANGE

National traffic exchange: peering and transit issues

Traditionally ISPs have exchanged traffic among themselves using a different financial settlement model from PTOs.[19] When domestic traffic is exchanged between PSTNs it is usually accomplished via an interconnection or access payment. The bulk of international PSTN traffic is exchanged via a system of financial settlements, sometimes know as the accounting rate system. By way of contrast ISPs, in the main, continue to exchange traffic via a system known as "peering" although there is increasing use of interconnection payments between ISPs.

An agreement to peer means that two ISPs will exchange necessary routing information so that traffic can be exchanged between their networks at no charge (i.e. similar to the sender-keeps-all system sometimes used for international PSTN traffic).[20] This exchange occurs at public and private Internet exchange points, which are points of traffic exchange and provide access to backbone networks. In times past, there were only a small number of Internet exchange points or NAPs and, because of the origins of the Internet, they were all located in the United States. At present the number of Internet exchange points is growing very quickly, with a much greater geographical dispersion both within the United States and in other OECD countries (see next section).

ISPs will consider several factors in approaching negotiations for peering agreements (Table 2). These include consideration of their prospective peer's customer base (e.g. number and type of customers), as well as the reach and breadth of their network. Peering between ISPs of equal "size and shape" is relatively straightforward at the local and national level. The largest ISPs in the United States, peer with each other on a national basis because they recognise the mutual benefits. The smaller ISPs, usually operating in a limited geographical region, also peer among themselves on a fairly straightforward basis because they too recognise the mutual benefits. This is a major factor in the proliferation of new "local and regional" Internet exchange points (note: these new exchange points are sometimes known by different terms, such as MXP, as explained below).

The different characteristics of ISPs mean that they do not necessarily enter an interconnection negotiation with equal bargaining strength. There are more than 4 000 ISPs of greatly differing characteristics in the OECD countries.[21] Some have customers providing content and services much in demand by the customers of other ISPs. At the same time, some ISPs have a customer base that the customers of other ISPs want to access. However negotiations for peering do not just occur horizontally between ISPs but also vertically between "small local ISPs" and "large national ISPs". In the latter case the "large national ISPs" have a stronger bargaining position because they not only provide access to their customer and content base, but also act as a gateway to the rest of the Internet. As a result the exchange of Internet traffic operates with two parallel systems. The first is peering, whereby traffic is usually exchanged without payment, and the second is a system involving transit payments. One commentator has defined transit along the following lines:

> "Transit comes into play when a provider wants to reach customers of some third party that the first provider doesn't peer with. If the ISP that peers with the first provider also peers with the third party, then that provider is in a position to offer the first provider transit to the third party. Transit will normally cost a flat monthly charge."[22]

To understand the relationship between peering and transit payments it is necessary to recall the "non-commercial" origins of the Internet. In its first guise as a defence network, with a single user (the US military), commercial traffic exchange was explicitly excluded. In the 1980s, as the network was increasingly opened to academic and research institutions, the Internet's commercial utility was recognised but limited by the so called "acceptable use policy" (AUP). In the 1990s, as the AUP was relaxed, the first commercial NAP was established in the United States. In 1991, the Commercial Internet Exchange or CIX, now better known as a leading industry association of ISPs, obtained its name by being the first entity to establish an exchange point for commercial users.[23] The original rationale for a commercial NAP was to enable companies using the Internet to have more flexible routing. Prior to the establishment of the CIX, while one division of a company might have had access to the Internet, other divisions of the same company may not have had the same access rights because of the AUP. In this situation, such a company would therefore have had to route the other divisions' traffic via another network.[24] The establishment of a commercial NAP allowed the original CIX members to route all company traffic over the same network without fear of violating NSFNET or Internet AUPs.

In 1994, the NSF awarded contracts to replace the NSFNET's Internet backbone. These contracts were for backbone transport, the routing arbiter and traffic exchange points (NAPs). The four original NSF-sponsored NAPs were located in San Francisco (operated by Pacific Bell), Chicago (operated by Ameritech), New York (operated by Sprint in New Jersey), and Washington, DC (operated by MFS).[25] In 1995 NSFNET was retired and US backbone traffic was routed via interconnected commercial networks. These networks continue to use the four original NSF-sponsored NAPs, now commercially funded, and the CIX.

In the transition to a commercial Internet, ISPs continued largely to employ the peering model they had inherited from NSFNET. One reason was the fairly small number of ISPs, with relatively uniform characteristics, so that their bargaining power was more or less even. In addition, the relatively small number of exchange points meant that all ISPs wanting to peer had to build or purchase infrastructure (i.e. leased lines from telecommunication carriers) to reach the available NAPs. In other words, the levels of network investment or costs were broadly similar and this infrastructure could, in turn, be used by all ISPs to mutual advantage. However, as the number of ISPs, and the number of Internet exchange points, increased, the economics of providing networks and the relative bargaining strength of the parties began to change.

The strains in certain peering relationships in the United States, which began to show in 1996, became much more visible in 1997. What brought this debate to public attention was a move by some of the larger ISPs to bring to an end certain existing peering arrangements, or a decision not to enter into new peering arrangements, with smaller ISPs. In the United States, the most discussed interconnection debate occurred between UUNET, a subsidiary of WorldCom, and a number of smaller ISPs, including Whole Earth Networks.[26] In this instance UUNET, by some measures the largest Internet service provider in the world, stated that it would no longer accept peering requests from other ISPs whose infrastructures would not allow for the exchange of similar traffic levels.

For UUNET the reasons for ending some peering relationships were clear and compelling.[27] According to UUNET, the company was not only receiving peering requests from smaller ISPs that could not route traffic on an equitable and bilateral basis, but also from companies running "Web server farms"

that were not providing inter-networking services. From UUNET's perspective, this meant they were being asked to provide national and international data transport, as well as connectivity and support services, to companies that could not provide similar services in return. In other words, UUNET said, these companies were seeking to use its network free of charge and not providing it with a return on its growing investment in infrastructure.

Critics acknowledge the validity of many of the points made in support of the position adopted by companies such as UUNET, such as the substantial cost differences in providing different networks. However they are concerned with the transparency of the interconnection process and the potential for anti-competitive behaviour. One critic has stated,

> "In the context of peering, many mid-sized Internet Service Providers have currently or have been willing to build out their networks to exchange traffic with the largest networks in multiple geographically diverse points only to find that these larger networks will neither exchange traffic once these competitive networks have arrived at these points, nor will these large network operators even disclose under what criteria they would exchange traffic over these geographically diverse points. This refusal to make public their criteria for interconnection is at the heart of a very serious threat to the continued growth and openly competitive nature of the Internet."[28]

For its part, UUNET responded to criticism of this type by explicitly stating that it was in favour of interconnection and would not deny network access to any ISP.[29] UUNET further stated its policy was to offer peering with any ISP that operated a national network, with a diversely routed DS-3 (a network operating at T3 speeds) backbone, and which could connect to UUNET at DS-3 or greater speeds in at least four geographically diverse locations in the United States. For ISPs not meeting these criteria, UUNET announced monthly interconnection rates of US$ 2 000 for a T1 connection, and US$ 6 000 for a T3 connection to their network. AGIS, another Internet backbone provider, also discloses its criteria for peering, although the company notes that some companies do not publish these requirements.[30] In the United Kingdom, UUNET has also published its peering policy (Box 1). Clearly stated is the fact that UUNET reserves the right to peer in the United Kingdom with other ISPs, even if they meet all the criteria specified. Similarly in Switzerland, Unisource Business Networks (UBN) reserves the right to decline to peer with other ISPs even if they meet UBN's specified criteria (Box 2). Policies of this type have generated complaints from smaller ISPs in some markets.

This debate raises a number of complex issues for consideration by policy makers. However, before considering these issues it is necessary to note a series of other developments that need to be taken into account. At a time when several large ISPs moved to end peering relationships PSINet, manager of one of the world's largest and most advanced packet switched networks, announced it was prepared to offer free peering to any ISP for its planned OC-48 (a network operating at 2.5 Gbit/s) Internet-optimised backbone. According to PSINet while "free peering" has obvious benefits for smaller ISPs that are denied peering by other larger operators, it would also benefit PSINet's corporate customers by streamlining communication between those customers and PSINet's ISP partners.

In a world where communication carriage has until recently been the preserve of monopoly operators, the offer by PSINet is a somewhat novel example of the market at work. Even when markets have been opened, there is inevitably a considerable period of time necessary for new entrants to roll out alternative infrastructure and test different strategies. Accordingly, in those instances where it has been deemed that incumbent operators of essential facilities were acting in an anti-competitive manner,

regulators have had to act before competitive markets have been able to exert discipline. What is refreshing about the Internet is that market forces are, in a growing number of cases, acting ahead of the need for traditional communication regulation.

The other interesting aspect of the divergent paths adopted by PSINet and UUNET is that the Internet industry is still experimenting with different business models. Both models appear to have attractive features in terms of either offering incentives to build new infrastructure by providing a clear path for financial return (i.e. transit payments) or by opening markets for smaller ISPs (i.e. peering). In contrast to the telecommunication markets, the evolution of the Internet is taking place in a largely unregulated environment which is enabling business models to be tested by the market. To date, as concerns have arisen about questions such as "interconnection transparency" between ISP networks or competition among backbone providers, the market has responded.

While the larger Internet backbone providers appear to be going through a period of consolidation, the number of players appears large relative to other communication infrastructure markets (Table 3). If a backbone company is defined as one that does not have to buy Internet access from any other company, there were nine such entities in the United States in September 1997, although this number may be reduced owing to subsequent announcements of mergers and acquisitions.[31] By including companies that pay transit fees that number may increase to almost 50 backbone providers. Nevertheless, some have questioned whether the proposed merger between WorldCom (including UUNET) and MCI might raise competition concerns.[32] Some analysts put the combined WorldCom-MCI share of Internet backbone traffic at 50-55 per cent.[33] That being said, in the absence of unified Internet traffic statistics, it is not entirely clear how some of these estimates were calculated.

Determining market shares for the Internet using available indicators is extremely difficult.[34] The use of indicators from traditional communication markets, such as measurements of traffic or revenue by market segment, is not generally possible because the aggregated data are not available. While the leading players may know this information for their own activities, they may not have data for the whole Internet market. In the absence of such measurements, some analysts have used the InterNic database of Autonomous System Numbers (ASNs) and the routing table from the MAE West exchange point to approximate the number of routes connecting ISPs to major backbone providers.[35] These analysts then endeavour to assess a backbone provider's position in the market by its share of routes to other ISPs.

The OECD has experimented with another indicator which uses traceroutes from Web sites served by the networks of major backbone providers to 100 of the most accessed Internet sites. This provides an indication of which backbone providers carry this traffic on an end-to-end basis (i.e. back and forth between the origin of the traceroute and the leading Web sites) and which backbone providers exchange traffic with another backbone provider to access some of the most popular content on the Internet. The three originating points for the traceroutes (with their initial US backbone connection in brackets) were Global One (Sprint), Mids/Alexa (Alternet/UUNET), and Beachnet (Cerfnet). The respective results, for each of traceroutes sites to Web21's leading 100 Internet sites in March 1998, are shown in Appendix 1 (Figures 1, 2 and 3). For example, a series of traceroutes from the Global One Web site, where Sprint provides the initial backbone connection, shows that Sprint carries the traffic on an end-to-end basis for 18 of the 100 leading Web sites. For the other 82 Web sites, Sprint passes the traffic to a second backbone provider, the largest of which is the Worldcom group of Internet companies.

By themselves, these three traceroutes should not be taken to measure market share and more extensive measurements would need to be undertaken to provide such an indication. Rather, they represent the perspectives of different backbone providers on the traffic exchange with other companies, which is needed to provide connectivity to popular Web sites for their customers. Accordingly, the traceroute tool

might be further developed to provide an indication of the market power different participants (including ISP to ISP and between some large scale content providers and ISPs) have in relation to negotiations on Internet traffic exchange. The results shown in Figure 1 could be described as the Sprint "perspective" on Internet traffic exchange to leading Internet sites. A series of traceroutes from a Web site served by another backbone provider will produce a different result. Accordingly, a series of traceroutes from Mids/Alexa, which uses Alternet as the initial infrastructure provider, shows that the Worldcom group of Internet companies can carry traffic on an end-to-end basis to a much greater number of these same Web sites than Sprint (Figure 2). The difference in the number of Web sites for which Worldcom passes traffic to Sprint (as opposed to vice versa) is due to some leading content providers multi-homing (i.e. using more than one ISP to provide infrastructure). In other words for many of the sites for which Sprint provided end-to-end carriage, the Worldcom group could also provide direct end-to-end carriage. A third perspective is supplied by undertaking the same series of traceroutes from a Web site served by a smaller backbone provider -- CerfNet (Figure 3).

The competition concern raised by some smaller ISPs is that an operator with a very large share of IP backbone traffic and connectivity might be able to leverage higher rates for "interconnection" with their backbone network. Critics, of this view, point to the high existing level of Internet backbone infrastructure competition and to a number of fairly recent market entrants in the United States which are building extensive networks, such as Level 2 and Qwest. Nevertheless even if the United States does not yet have cause for concern, because the earlier introduction of infrastructure competition has encouraged numerous backbone networks, this may not be true in other OECD countries that are opening their markets to competition in 1998 and beyond. A further concern has been raised by some smaller ISPs that do not receive direct allocations of IP addresses. These ISPs borrow IP addresses from upstream backbone providers and may have to give them up if they change provider. While recognising that alternatives exist in the backbone market, small ISPs say that it is expensive and time-consuming, as well as inconvenient for their customers, to reconfigure their networks if they change backbone provider.[36] Policy makers need to be vigilant in ensuring non-discriminatory and transparent access to essential facilities in those countries where alternative backbone infrastructure is not yet widely available and further investigate the significance of the lack of IP address portability between backbone providers for smaller ISPs (for which there may be sound technical reasons, such as minimising the load on routing tables), in terms of the issues raised for competition policy.[37]

Regulatory designation

The other issue which is arising from current developments is the question of regulatory designation. Most ISPs would prefer to be designated as value-added service suppliers rather than as telecommunication carriers. While ISPs would generally like to be eligible for the privileges granted to telecommunication carriers in terms of co-location, for example, they would not like to have the traditional regulatory burden that has attended common carriage providers. In terms of traffic exchange between ISPs, where the debates clearly have similarities to telecommunication interconnection, the best way forward appears to be not to apply traditional telecommunication regulation to ISPs as long as there is sufficient competition in the market. In other words, ISPs should continue to be free to select the type of traffic exchanges they wish to enter into with other ISPs. At the same time, emerging best practice industry self-regulation involves transparency in terms of publication of peering and transit policies.

Although the regulatory designation of ISPs has received the most attention with respect to Internet traffic exchange, there may also be a need to review communication regulation in relation to how it affects telecommunication carriers and user networks. For example, if telecommunication regulation specified that a common carrier had to offer interconnection on a non-discriminatory basis, then this may

curb a telecommunication carrier's range of options for traffic exchange compared to an ISP. In other words, an ISP might be able to choose to peer with other companies that offer equivalent services while charging transit to others. However, if a telecommunication carrier offering peering to one company had to offer it to all applicants, irrespective of their ability to offer equivalent services, the carrier might choose only to offer transit. To date, it seems unlikely that this has been a restraint on the commercial freedom of carriers. Certainly Swisscom's participation in UBN does not seem to have posed a problem in terms of reserving the right to choose whether it will exchange traffic with other ISPs and how this will be undertaken (see Box 2). In most cases if the telecommunication carrier chooses not to peer it is probably because it does not believe that smaller ISPs can offer equivalent levels of service.

Similar questions are emerging in terms of telecommunication regulation and the participation by user networks wanting to peer at Internet exchange points. In Australia, the establishment of an Internet exchange point in the state of Western Australia (WAIX) raised the question of whether academic and government networks were going to peer at WAIX. Under new legislation, the Telecommunications Act 1997, the question was raised as to whether these organisations might be classified as telecommunication carriers if they were to peer at WAIX.[38] At the time of writing submissions were up before the Australian Communications Authority for exemption or relief. Before continuing the discussion of traffic exchange at the international level it is necessary to better understand the role of Internet exchange points and how they are developing in the OECD area.

The regulatory designation for the exchange of traffic between telecommunication carriers and ISPs are also generating questions at the level of local access networks. In the United States some Regional Bell operating companies (RBOCs) have expressed a desire for calls to ISPs originating on their networks, but terminating on the local networks of others, to be paid for in a different way from local calls.[39] The RBOCs say that with regular telephony patterns of use, traffic exchanged between local networks would tend to be in balance over time. However, they argue that if competitors target ISPs they will be in the position of having a significant financial deficit with other local network providers. Their main point is that regulatory frameworks designed for telephony should be reviewed in the light of the different patterns of use generated by Internet access. At the same time, ISPs argue that incumbent telecommunication carriers sometimes leverage their bottleneck control over the local loop infrastructure in ways that are anti-competitive. Until alternative local infrastructure is available on a sufficient basis to enable ISPs to have competitive access options, policy makers need to ensure that competitive safeguards are in place.

Table 2. **Benefits and drawbacks of peering for ISPs**

Upside	Downside
An ISP can send traffic free to the customers of another ISP.	Another ISP can send traffic free to your customers.
Peering does not require accounting systems in the same way as a settlements system.	Peering may be more difficult to administer than purchased transit.
Some ISPs may benefit from not having to pay their fair share of infrastructure costs associated with transit.	Some ISPs may not receive a fair return on their investment in infrastructure used for transit.
Peering with additional ISPs may reduce the number of "hops" traffic must pass in traversing networks.	This may be at the risk of carrying additional transit traffic.
Large customers have indicated they favour peering among service providers even though there is little to indicate one system is necessarily superior from a performance perspective.	Some ISPs may not manage their networks efficiently.
	For designated telecommunication carriers regulation may mean that if they offer peering to one company they may have to offer it to all applicants irrespective of their ability to provide an equivalent service.

Source: OECD, based on Schwandt (1997).

Table 3. **Selected US Internet backbone providers**

US backbone provider	Status
AGIS (www.agis.net/)	AGIS (Apex Global Internet Services, Inc.) was founded in 1994 to provide Internet backbone services and corporate intranets. AGIS has equipment present at the major Internet peering points around the United States: MAE-East in Washington, DC; the New York NAP at Sprint in Pennsauken, NJ; the AADS NAP in Chicago, IL; the Commercial Internet Exchange (CIX) in Santa Clara, CA; the Pacific Bell NAP in Palo Alto, CA.; and MAE-West in San Jose, CA. AGIS currently uses Worldcom's ATM service at DS-3 (45 Mbps) rates to haul IP data across the United States, and has begun its migration to OC-3 (155 Mbps).
BBN/GTE (www.bbn.com/) (www.gte.com/)	On August 15, 1997, GTE Corporation acquired BBN Corporation, which became a new subsidiary of GTE. GTE Internetworking, the new data unit, includes BBN and the existing GTE Intelligent Network Services organisation.
MCI (http://www.mci.com/)	MCI Internet customers are connected to the Internet through MCI's Internet backbone. Operating at 622 megabits per second (Mbps), it is one of the fastest and largest backbone networks of its kind in the world. Competitive bidding is under way for MCI with offers made by two other companies owning backbone networks, WorldCom, the owner of UUNET, and GTE, the owner of BBN.
Netcom/ICG (www.icgcomm.com/) (www.netcom.com/)	Merger announced in October 1997. ICG has extensive fibre-optic networks and offers local, long distance and enhanced telephony and data services in California, Colorado, the Ohio Valley and parts of the southeastern United States. ICG is a leading national competitive local exchange carrier. The combined company will be served by more than 2 600 employees and will have a network platform interconnecting 330 Internet points-of-presence, over 40 000 dial-in access ports, 18 telephony switches, 15 frame relay switches and nearly 2 900 fibre route miles -- with an additional 1 117 fibre route miles under construction.

Table 3. **Selected US Internet Backbone Providers (continued)**

PSINet (www.psi.net/)	Another Internet pioneer PSINet has 225 points-of-presence in the United States and more than 350 worldwide.
Sprint (www.sprint.com/)	Sprint is an Internet pioneer and says it is the carrier of nearly two-thirds of today's Internet traffic worldwide.
UUNET/WorldCom	UUNET became a subsidiary of WorldCom, Inc., in 1996. In the same year WorldCom acquired MFS. In September 1997, WorldCom acquired AOL and CompuServe's network services company. AOL's company was previously called ANS Communications and CompuServe's company CNS.

1. A list of backbone providers and comparison of performance can be found at: http://www.keynote.com/measures/backbones/backbones.html

Source: OECD.

Box 1: Extract from UUNET (UK) AS-1849 Peering Policy, Source: UUNET at: http://www.uk.uu.net/network/peering/policy/

UUNET (UK) will consider peering with all ISP organisations within the United Kingdom (including the LINX, the London Internet Exchange) provided they meet the following conditions:

- The ISP must be a nation-wide ISP, offering service to its customers throughout the United Kingdom.

- The ISP must have at least 2Mbps bandwidth connecting from their backbone to their router at the interconnect point (for example, the LINX), or to the UUNET (UK) backbone in the case of a direct connection.

- The ISP must announce at least a /15 network block allocation or its equivalent from RIPE or other such registry. This does not include customer class B networks.

- The ISP must aggregate route announcements to UUNET (UK). UUNET (UK) will dampen heavily any networks with /24 mask and filter those with a longer prefix.

- The ISP will receive all UK-based networks connected to the AS1849 backbone, as detailed in the RIPE Routing Registry AS-UUNETPIPEXUK macro. The UUNET International business networks will only be available by separate agreement. Please contact intl-peering@uu.net if you wish to discuss.

- The ISP must have established a full peering with the LINX collector router (If LINX GIX)

Even if these conditions are met, there is no guarantee, implied or otherwise, that UUNET (UK) agree to peer with the ISP or continue to peer at some stage after agreement has been reached. All decisions taken in respect of peering are at the sole discretion of UUNET (UK). UUNET (UK) reserves the right to change any of these conditions at a later date. This policy is intended to serve as guidelines and clarify our peering policy.

Box 2: Extract from Uniplus Internet Peering Policy: http://www.unidata.ch/backbone/policy.htm

To optimise as far as possible the connectivity of the Internet, Unisource Business Networks Switzerland (UBN) has an open policy for "zero-settlement" peering with other Internet service providers (ISPs). A zero-settlement peering is one where both ISPs assume that the traffic is approximately equal in both directions, and both benefit equally from the connectivity. Neither ISP buys a service from the other, and hence neither ISP bills the other.

Rather than measure traffic in order to make sure that a peering is not disproportionally in favour of one party or the other (which is difficult at public LAN exchange points anyway), UBN defines the following criteria for entering into a zero-settlement peering. These criteria are designed to ensure that both UBN and the peering partner will benefit equally from the peering.

- Points of presence (POPs) in the 8 largest Swiss cities
- At least 1 Mbps of dedicated transatlantic Internet trunk capacity
- Direct (leased line) Internet access to customers as a standard offering at all POPs
- RIPE Local Registry
- 24 hour-per-day/7 day-per-week support function

Note that these criteria are only guidelines. UBN reserves the right to decline to enter into a zero-settlement peering even if the above criteria are nominally met by a potential peering partner.

In the case of ISPs who want improved connectivity to the Swiss Uniplus network, but who do not meet the zero-settlement criteria and do not want or need the global connectivity offered to Uniplus Internet customers, there is the option to become a "Local Connectivity" customer. This service offers connectivity to all other customers of the Swiss Uniplus Internet service, and the zero-settlement peers of the Swiss Uniplus network, at a lower price.

Local connectivity customers are normally multi-homed, since they will get global Internet connectivity from another ISP. Therefore they must have their own autonomous system (assigned by RIPE), and the BGP4 protocol must be used on the connection between UBN and the customer. Normally, a local connectivity customer will connect to the Uniplus Internet service via PTT leased line to the nearest Uniplus POP. It is also possible to make the connection at the CIXP, provided that the connection is still via a dedicated serial connection.

Peering, transit and Internet exchange points

The Internet consists of a patchwork of independent networks using the same protocols. When packets need to be exchanged between IP networks they either go via direct interconnection between these networks or via a public Internet exchange point. As recently as 1995, there were only a small number of Internet exchange points which were almost all located in the United States. This meant that content being requested by a user from a server in the same geographical locality might traverse continental and intercontinental networks before being received. In the United States, for example, prior to the establishment of an Internet exchange point in Boston, traffic between users in Boston would mostly be exchanged in New Jersey or Washington, DC.[40] This was the case even if the content provider and the user were physically located in close proximity. The same situation existed internationally. Before the

establishment of the "NAP Roma" in May 1995, data packets to and from Italian users were routed over long paths ranging across Europe and the United States.[41] Indeed, what was true for Boston and Rome applied to virtually the whole of the Internet.

The commercialisation of the Internet, and the rapidly growing traffic it has generated, has provided tremendous incentives for ISPs to increase the number of Internet exchange points. It is an axiom of Internet network management that ISPs are endeavouring to take content closer to customers. This not only provides better response times for applications, such as "surfing the web", but can cut the transit payments smaller ISPs need to make to larger ISPs. In terms of response time, "local Internet exchange points" mean that local traffic is not competing for resources at the larger, and busier, exchange points and NAPs and that local content providers can be directly connected via their ISP. For example, the City municipality of Rome is connected by a local area network to "NAP Roma".

Between 1995 and 1996 Internet exchange points were established in the largest city of most OECD countries outside the United States. In 1997, the trend has been to establish new Internet exchange points in an increasing number of regional cities and centres (Table 4). The country most advanced along this path is the United States which has, in some cases, competing Internet exchange points in the same city (Table 5). In other countries while the first Internet exchange points were often established in capital cities (e.g. Paris, London and Rome), new Internet exchange points have now been established in regional centres such as Grenoble, Manchester, and Milan. The benefits for users are readily apparent. Before the Grenoble exchange point was established, a user on one ISP's network wanting to access local content on another ISP's network had wait while this traffic passed through the only French interconnection point located in Paris (i.e. the SFINX). The short-cut provided by the new exchange point has meant decreases of more than a factor of ten in reported response times.[42] Smaller exchange points than the initial NAPs are sometimes referred to as Metropolitan eXchange Points (MXP). The main difference is that while national service providers exchange traffic at the larger peering points and NAPs, an MXP aims to service local or regional traffic without burdening backbone networks.[43]

The reasons for installing regional exchange points are not just compelling increases in response times for users. The new Internet exchange points also increase the reliability of the Internet. Before MaNAP started operations in Manchester, an estimated 98 per cent of UK Internet traffic was passing through a single building in London and through one set of equipment.[44] A failure of the London LINX exchange, although extremely rare, meant that all UK traffic would have had to be routed to the United States.[45] However the growth of UK Internet traffic means that routing all traffic via the United States is no longer a viable option, not only because the available international links would not have the capacity to carry all intra-UK traffic, but also because US users and ISPs would not want US NAPs to handle this traffic. The process of establishing new exchange points can be expected to continue as some of the first exchange points outside the United States still handle the bulk of national traffic. For example, the Copenhagen Internet exchange (DIX) was founded in May 1994, and still exchanged more than 90 per cent of the Danish inter-network traffic by November 1997.[46]

Internet exchange point policies and pricing

The policies of Internet exchange points and NAPs are drawn up by the founding members mostly in the form of a memorandum of understanding (MoU). In general they are commendable examples of industry participants co-operating to produce workable models for traffic exchange. Significantly, these arrangements have been made without the need for government involvement and

regulation in contrast to interconnection between PSTNs. This is largely because initial participants were of relatively equal strength compared to newly liberalised telecommunication markets where the incumbent is a dominant operator of essential facilities.

While most of these MoUs are not lengthy documents it is not possible to reproduce them here. In most cases, they are published on the relevant Internet site associated with the exchange point. As such, only a selection of the types of policies found in MoUs are described, mainly focusing on traffic exchange (Table 6). In many cases the MoUs have been modelled on the first NAP or Internet exchange in that country (e.g. MaNAP based on LINX) or on international examples. Accordingly, there are a number of common specifications such as the number of independent connections (sometimes international) an ISP must have with other Internet exchanges or the number of other ISPs which must be peered with at the Internet exchange point concerned. One issue that may arise here is whether regulation of common carriers in any way impedes their ability to conform with or take advantage of these such policies.

The pricing of some NAPs appears to vary widely even where it is stated that activities are undertaken on a not-for-profit basis (Table 7). No doubt there are a number of different explanations for these differences, such as the size of the Internet exchange, the services or equipment included in published prices as opposed to those supplied by each ISP, and so forth. As with the policies associated with peering, there appears to be little need for any direct government involvement in these industry-driven pricing arrangements for the Internet exchange points. In contract to traditional telecommunication networks, there are no incumbent operators controlling essential facilities. Moreover, as the policies are agreed by industry participants, they are mostly transparent and non-discriminatory. However, it is also true that pertinent information for prospective members, such as pricing and traffic exchange policies, are not uniformly available on the Web sites of some Internet exchange points. In addition some Internet exchange points provide much better information for end users in terms of network performance. In this respect, one of the best Web sites is Singapore Telecom's site for its Internet exchange (STIX), which provides near real-time performance indicators (http://www.stix.net/).

Table 4. **Selected Internet exchange points in the OECD area**

	Network exchange point	Status	URL
Australia	Internet exchange (WAIX)	Western Australian Internet Association (WAIA)	www.waia.asn.au/Issues/Peering/index.html
Australia	AUIX (Australian Internet Exchange)	Public exchange point for ISPs in the largest Australian cities.	www.auix.net/
Austria	Vienna Internet eXchange (VIX)	Vienna University Computer Center which may be used by Internet Service Providers (ISPs) to exchange traffic at the national or international level. The VIX is a service for commercial ISPs and academic networks operating in the central and eastern European region. Not-for-profit.	www.vix.at/

Table 4. **Selected Internet exchange points in OECD area** (continued**)**

Belgium	BNIX: The Belgian National IP eXchange	The place where ISPs can interconnect in Belgium. It is aimed at the IP traffic exchange between each connected ISP at national or international level. Any ISP with an arrangement for traffic exchange with any of the ISPs already connected to the BNIX can connect to the BNIX.	www.bnix.net/
Canada	Montreal Internet Exchange (MIX) and others. (No web page available for CANIX)	The Montreal Metropolitan Internet eXchange backbone was created in 1993 as a agreement among 5 ISPs to exchange local Internet traffic. A few months later, Toronto MIX and Quebec Cité MIX were created. As of 1995 in Quebec, 47 ISPs and their customers, including large corporations, exchange local traffic over the MIX backbones, thus increasing the performance of Internet communication by avoiding long detours via the United States.	cgat.bch.umontreal.ca:8080/sps3.html
Czech Republic	Neutral Internet eXchange (NIX.CZ)		www.nix.cz/
Denmark	DIX(Danish Internet eXchange point)	UNI-C Network Operations Center, Lyngby	www.uni-c.dk/dix/
Finland	Finnish Commercial Internet Exchange (FICIX)	Consortium of Finnish Internet technology-based data communication providers	www.ficix.fi/.
France	GNI (Grenoble Internet Initiative)	Grenoble's Proximity Exchange Point is a place where Grenoble area ISPs may interconnect their backbones, and exchange local traffic.	www.gni.fr/PEP/
France	SFINX: Service for French Internet Exchange (also GIX)	Paris-based facility enabling ISPs to exchange traffic without passing through transnational networks. The GIX is managed by Renater	www.urec.fr/Renater/Sfinx/French/SFINX.html and www.urec.fr/Renater/gix/gix1000.html
Germany	DE-CIX	Based in Frankfurt.	www.eco.de/
Greece	Athens Internet Exchange (AIX)	Ministry of Development, General Secretariat for Research and Development. GRNET (Greek Research and Technology Network).	www.grnet.gr/index_en.html
Hungary	BIX - Budapest Internet eXchange		goliat.c3.hu/bix/
Iceland	NA	NA	NA
Ireland	INEX (Internet Neutral Exchange)	Facility for Irish Internet Service Providers. Not-for-profit.	www.inex.ie/
Italy	MIX - Milan Internet eXchange	Association of Italian Internet Providers (AIIP)	www.aiip.it/mixit.html
Italy	NapRoma	An Internet exchange point hosted by the CASPUR (Consortium for the Applications of Supercomputation for University and Research) facilities at the University of La Sapienza in Rome and co-operatively operated by its participants.	www.nap.inroma.roma.it/
Japan	JPIX (JaPan Internet eXchange)		www.jpix.co.jp/

Table 4. **Selected Internet exchange points in OECD area** (continued**)**

Japan	NSPIXP (Network Service Provider Internet eXchange Point)		xroads.sfc.wide.ad.jp/NSPIXP/ www.inoc.imnet.ad.jp/noc/nspixp2.html
Korea	NA	NA	NA
Luxembourg	NA	NA	NA
Mexico	NA	NA	NA
Netherlands	Amsterdam Internet Exchange (AMS-IX)	Used by ISPs to exchange traffic at a national or international level.	www.ams-ix.net/
New Zealand	New Zealand Internet Exchange (NZIX)	University of Waikato	www2.waikato.ac.nz/NZIX/
Norway	NA	NA	NA
Poland	NA	NA	NA
Portugal	PIX (Portuguese Internet eXchange point)		www.fccn.pt/PIX/
Spain	ESPANIX (Spanish Neutral Interconnection Point)		www.espanix.net/
Sweden	DGIX - KTHNOC	Royal Institute of Technology (KTH), Stockholm.	www.sunet.se/dgix/
Switzerland	SIX - the Swiss Internet eXchange	ISP national and international traffic exchange. SIX-B - Bern SIX-Z - Zürich SIX-L - Lausanne (planned)	www.six.ch/
Switzerland	CERN Internet eXchange Point (CERN-IXP)	Open to all ISPs having a point of presence in Switzerland and/or France.	wwwcs.cern.ch/wwwcs/public/ip/cernixp.home.html
Turkey	NA	NA	NA
United Kingdom	LINX (London InterNet eXchange)	Not-for-profit industry association of ISPs.	www.linx.net
United Kingdom	MaNAP (Manchester Network Access Point)	Not-for-profit industry association of ISPs.	www.manap.org/rel_22_jul.html
United Kingdom	i-Exchange	Located in Telehouse London, i-Exchange is a neutral peering point allowing UK ISPs to exchange traffic with each other on a regional or national level.	www.i-exchange.co.uk/
United States	Commercial Internet eXchange (CIX) and many others. Refer Table 7	CIX is a not-for-profit Industry association of ISPs.	www.cix.org/

1. In some case information was not available in English or French on the above sites.

Source: OECD and Bill Manning's http://www.isi.edu/div7/naps/

Table 5. **Selected US Internet exchange points**

State	Network exchange point	Status	URL
Arizona	The Tucson Interconnect	The Tucson Interconnect allows taking IP packets that go between one Tucson ISP and another Tucson ISP off the expensive national circuits and onto the inexpensive local circuits. Membership is by invitation, free, and includes automatic peering with all participants. Only Tucson businesses are welcome.	www.tti.aces.net/
Arizona	The Tuscon NAP	Non-NSF Regional Network Access	www.ttn.rtd.net/
California	Digital Internet Exchange	Digital's Internet Exchange in Palo Alto is a data and communications center at which ISPs and their customers can locate equipment for redundant access, reliability and operational stability. In addition, the Commercial Internet Exchange (CIX) has installed their router, which provides free multilateral peering to CIX members at the Exchange.	www.ix.digital.com/
California	MAE-West	MAE West is interconnected with the Ames Internet Exchange, operated by NASA at the Ames Research Center.	http://www.mfsdatanet.com/MAE/
California	MAE-Los Angeles	A WorldCom public exchange point for ISPs.	http://www.mfsdatanet.com/MAE/
California	LAP/MAE	LAP is located in the LA area, at ISI and connects to MFS Datanet ((MAE-LA).	www.isi.edu/div7/lap/
California	Pacific Bell NAP	Pacific Bell's San Francisco NAP is one of the four original NSF-sponsored network access points for the Internet infrastructure.	www.pacbell.com/products/business/fastrak/networking/nap/index.html
Colorado	The MAX: Mountain Area Exchange	The MAX is a public Internet traffic exchange point in Denver.	www.themax.net/
District of Columbia	MAE East	A NAP provided by MFS Datanet in Washington, DC One of the four original NSF-sponsored network access points for the Internet infrastructure.	www.mfsdatanet.com/MAE/

Table 5. **Selected US Internet exchange points** (continued)

Georgia	Atlanta exchange point	The Atlanta-NAP houses fibre systems from Bell South, MFS, and MCI Metro.	www.atlanta-nap.net/
Georgia	Atlanta Internet Exchange (AIX)	Proposed.	www.com/aix/
Illinois	Chicago NAP	Ameritech's Advanced Data Services Network Exchange Point. One of the four original NSF-sponsored network access points for the Internet infrastructure.	nap.aads.net/index.html
Illinois	MAE-Chicago	A WorldCom public exchange point for ISPs.	http://www.mfsdatanet.com/MAE/
Indiana	IndyX, Indianapolis Data Exchange	The exchange functions as a gateway to a new national switched Internet backbone, as well as a local exchange between ISPs.	www.indyx.net/info/
Maryland	Baltimore NAP	The Baltimore NAP provides a common forum for mutual inter-exchange of Internet network traffic among multiple ISPs.	www.baltimore-nap.net/
Massachusetts	Boston MXP	The Boston Metropolitan Exchange Point or MXP is a project undertaken by MAI Network Services with the sponsorship and support of the City of Boston. The functional goal of the exchange is to allow ISPs, businesses, universities, and any other organisations with large IP networks to develop faster connectivity between one another and to rely less on Internet infrastructure located in geographically distant locations.	www.mai.net/bostonMXP/
Michigan	Detroit MXP	Services local or regional traffic	www.mai.net/mxp/MXP.HTML
Missouri	STLOUIX	The St. Louis Open Internet eXchange was created to improve the speed and reliability of Internet traffic for St. Louis and Midwest regional businesses and individuals.	www.stlouix.net/
New Jersey	IPeXchange	IPeXchange is a public/private joint venture of the AV-Network and IPeXchange participants.	www.avnet.org/isg/njipxdes.html

Table 5. **Selected US Internet exchange points** (continued)

New Jersey	Sprint NAP	One of the four original NSF-sponsored network access points for the Internet infrastructure.	www.merit.edu/nsf.architecture/Sprint/.index.html
New Mexico	New Mexico NAP	Under construction	www.nmnap.net/
New York	MAE-New York	A WorldCom public exchange point for ISPs.	http://www.mfsdatanet.com/MAE/
New York	Telehouse NY IIX	Telehouse operate two peering services -- one aimed at serving the needs of global ISPs and the other aimed at bringing together the ISPs in the New York metropolitan area.	www.telehouse.com/Telehouse/InternetOffer.htm
Ohio	FibreNAP	FibreNAP is a layer 2 NAP where ISPs, Network Service Providers, and corporations, can meet to exchange traffic (peer), sell services to their customers (transit), or privately connect two networks.	www.fibrenap.net/
Oregon	Oregon Internet Exchange	The Oregon Internet Exchange (Oregon-IX) provides rich Internet connectivity for ISPs and high-volume networks throughout the Northwest region.	antc.uoregon.edu/OREGON-EXCHANGE/
Pennsylvania	Philadelphia Internet Exchange	The Philadelphia Internet Exchange is a public Internet exchange point located in Philadelphia for ISPs and others looking for better connectivity to others in the region.	www.phlix.net/
Tennessee	Nashville CityNet	Nashville CityNet is the intra-city computer network of Nashville.	nap.nashville.net/
Texas	MAE-Houston	A WorldCom public exchange point for ISPs.	mae.houston.tx.us/
Texas	MAE-Dallas	A WorldCom public exchange point for ISPs.	http://www.mfsdatanet.com/MAE/
Texas	Metro Access Point	Aim to exchange IPv4 packets between service providers and interested parties with a low barrier to entry and a low recurring operation costs. Austin: AMAP (online); Dallas-Fort Worth: DFWMAP (proposed); Houston: HMAP (In progress); San Antonio: SAMAP (In progress).	www.fc.net:80/map/

Table 5. **Selected US Internet exchange points** (continued)

Utah	Utah REP	Utah REP is a regional exchange point.	utah.rep.net/
Vermont	Vermont Internet eXchange - VIX	Proposed.	www.hill.com/trc/vix/index.html
Washington	SNNAP: Seattle Network-to-Network Access Point	Facilitates high performance, reliable connectivity among networks in the Puget Sound region and Washington state	weber.u.washington.edu/~corbato/snnap/
Washington	InterNAP	A Private Network Access Point (P-NAP): providing high bandwidth TCP/IP connectivity between the leading national and global ISPs, InterNAP, and its customers.	www.internap.com/nap4.html
Washington	NIX: Northwest Internet eXchange	This point was designed with the intention that providers in the northwest would get together and exchange traffic with each other, as an alternative to using the heavily used backbone providers and NAPs.	www.structured.net/nix/
Washington	Eastern Washington Internet Exchange	The EWIX is an exchange point in Spokane Washington where members connect via Frame Relay or dedicated 10 Mbps ethernet connections for the exchange of TCP/IP based traffic. There is a significant amount of e-mail and Web traffic that originates and terminates on local ISPs networks. Private exchange points, like the EWIX, provide a more direct path for traffic between ISPs and companies in a particular region.	www.dsource.com/ewix/

1. Shaded NAPs are the four original NSF-sponsored network access points for the Internet infrastructure.

Source: OECD

Table 6. **Selected peering policies and prerequisites of Internet exchange points**

	Peering and Related Policies.
Australia (WAIX)	Each network agrees to exchange traffic with the others at no cost, this not only eliminates traffic costs, but the savings means much faster links can be used. Peering does not replace the need for transit links. Participants still need to maintain an arrangement with a carriage provider to offer a "default route" for full Internet connectivity. ISPs may be able to purchase this transit by means of a direct connection inside the WAIX facility, but outside the WAIX peering fabric by arrangement with another participant.
Australia (AUIX)	A multilateral peering agreement participating ISP is obligated to advertise all its (participating) customers' routes to all the other MLPA (multilateral peering agreement) participating ISPs; obligated to exchange traffic among customers of all MLPA participating ISPs; entitled to select routing paths among the MLPA participating ISP. A multilateral peering agreement ISP is not obligated to provide transit to other MLPA participating ISPs; or obligated to announce the routes obtained from its other bilateral peering agreement partners to the MLPA participating ISPs. No monetary settlements are required by this agreement.
Austria (VIX)	VIX members need to agree on bilateral peering arrangements for traffic exchange. There is no obligation to exchange traffic with all other participants. A VIX member is required to be an ISP with its own international Internet connectivity. This connectivity must not be solely provided by other VIX member(s). A VIX member must provide Internet access to its customers at the IP level. In general, mere content provision does not qualify for VIX membership. Traffic is only permitted between VIX members having an explicit peering agreement. Injection of traffic into routers of non-peers is prohibited. VIX members will document their peering status in the RIPE database and notify VUCC of any changes.
Belgium (BNIX)	Peering with all other ISPs at the BNIX is not mandatory. Separate peering agreements have to be negotiated. ISPs are however asked to have at least one peering active.
Denmark (DIX)	Each party agrees not to charge the other party for interconnection-related matters, including charges based on traffic volume, commonly called "settlements", until mutually agreed by the parties. Transit traffic is traffic that has its origin or destination in a network which is not part of this agreement. Such traffic should not be covered by the agreement. (suggested DIX Peering Agreement).
France (GNI)	NA
Ireland (INEX)	Members must have their own permanent international connection to the Internet. As a rule of thumb, new members must have a route from their network to MAE-East in the United States which does not pass through an existing INEX member. Each member must be licensed by the Department of Communications. End-users who do not sell Internet services may not connect to the INEX. Internet resellers who buy connectivity from existing INEX members, and/or who do not have international capacity independent of members are likewise excluded from membership. Members may not connect more than two wide-area circuits to their router housed in the INEX rack, nor may they directly connect customers via circuits to their router. Each member must publish the contact to whom requests for peering should be sent. Any peering request by a potential new member must be responded to within seven working days of the request. Members will not install "sniffers" to monitor traffic passing through the INEX.
Netherlands (AMS-IX)	After being connected and up and running the ISP needs to arranges its own peerings.
Sweden (D-GIX)	Each ISP manages its own router and decides with which other ISPs it will set up peering sessions for exchanging traffic at the D-GIX. There is no obligation to exchange traffic with all other participants, but each ISP must peer with at least two other ISPs on the D-GIX.

Table 6. **Selected Peering Policies and Prerequisites of Internet Exchange Points** (continued)

Switzerland (SIX)	Peerings occur between all ISPs within the SIX i.e. once an ISP joins the SIX it has established peerings with all other ISPs within the SIX. Agreement to join SIX confers no rights upon either party to use another party's international connections. Operators joining the neutral LAN network are free to set up bilateral network connections independently of the SIX. Routing must be structured so that a network is always identified as being the same AS as "source as" even when it involves more than one peering point. Any ISP (no end user organisations) is eligible to join SIX, plus a couple of exceptions including governments (e.g. Swiss federal government). and academic institutions of higher learning. A SIX organisation must have one other permanent Internet connection.
Switzerland (CERN-IXP)	The CERN-IXP is a neutral Internet exchange point between Internet operators and is not a "service provider" *per se*. CERN operates an IXP in order to maximise its own Internet connectivity. Traffic in transit through the CERN-IXP is not subject to any particular restrictions, but must conform to all applicable laws and guidelines, and to the usage policies of the source and destination ISPs. CERN accepts no liability in any respect for the nature of transit traffic. Connection to the CERN-IXP does not provide any automatic connectivity. In order to get connectivity to other networks (e.g. Ebone, Eunet, Europanet, Transatlantic services, etc) it is the full responsibility of the ISP to make agreements and/or subscribe to such services via other ISPs located on the CERN-IXP or elsewhere.
United Kingdom (LINX)	Applicants must have their own independent, permanent, international connection to the Internet. As a rule of thumb, applicants must have paths from within their UK network to four of the Internet "root" name servers, which do not pass through existing LINX members. Members must have operational peering agreements with at least 20 per cent of existing LINX Members. Members must publish service details, including at least one public service allowing customers to connect to the Internet. Members must respond to a peering request by a potential new member within two working days of the request. Members must publish their contact to whom requests for peering should be sent. Members may not directly connect customers via circuits to their router housed in the LINX rack. Members will not install "sniffers" to monitor traffic passing through the LINX.
United Kingdom (MaNAP)	Members must have operational peering agreements with at least one existing MaNAP member. Members must publish prices, including at least one public service allowing customers to connect to the Internet. Members may not directly connect customers via circuits to their router housed in MaNAP rack. Members will not install "sniffers" to monitor traffic passing through the MaNAP.
United States (Philadelphia Internet Exchange)	PhlIX participants are welcome to set-up bilateral peering arrangements (or not) with any other participants as they see fit. To help reduce the work load in arranging peering, PhlIX management will maintain a list of those wishing to participate in an MLPA. PhlIX management recommends participation in the MLPA.
United States (Ameritech's Chicago NAP)	Attaching customers should intend to form bilateral or multilateral agreements with other NAP-attached networks; upgrade attachment technology and protocols as appropriate, and participate in the Chicago-NAP mailing list. A physical connection to a NAP should not be considered as an "Internet Connection." The NAP is an exchange point and peering arrangements between ISPs should be made before connecting. In addition, once the MLPA is in effect, additional bilateral transit agreements may be made between network providers.

Table 6. **Selected peering policies and prerequisites of Internet exchange points** (continued)

United States (CIX)	Member networks have a fundamental agreement to interconnect with all other CIX members. There is no restriction on the type of traffic that may be routed between member networks. The value of this basic agreement to exchange all legitimate traffic will continue to increase as the number of CIX member networks grow. There are no "settlements" nor any traffic-based charges between CIX member networks. Each member network connects to all other member networks directly or indirectly through the CIX router at no additional cost to member networks.
United States (PAC Bell NAP)	Multilateral peering agreement participating ISP is obligated to advertise all its (participating) customers' routes to all the other MLPA participating ISPs and accept routes from the customer's routes advertised by the ISPs; obligated to exchange traffic among the customers of all the MLPA participating ISPs; entitled to select routing paths among the MLPA participating ISPs; and entitled to make Bilateral peering agreements with non-MLPA participating ISPs. A MLPA participating ISP is not obligated to provide transit to other MLPA participating ISPs; or obligated to announce the routes obtained from its Bilateral peering agreement partners to the MLPA participating ISPs.

1. Wording is taken from sites concerned.

Source: OECD.

Table 7. **Internet Exchange point pricing**

	Pricing Policies
Australia (WAIX)	Any network may link into the WAIA NAP. The cost of connecting is one off set-up fee of US$ 370 and monthly maintenance fee of US$ 111. The peering point is neutrally managed by the Association on a non-profit basis.
Australia (AUIX)	The following charges are proposed by AUIX for the purpose of participating at the exchange points: a one-off setup charge of US$ 370; A monthly maintenance charge of US$ 111; for those participants requiring interNAP carriage, a monthly fee of US$ 148 for a shared 64Kb test circuit. The costs as indicated above go towards the following: 1. Cost of putting in place an Ethernet switch and router at exchange point; 2. Cost of putting in place a rack enclosure at the exchange point; 3. Cost of rental of space for rack enclosure at the exchange point; maintenance of route server will be provided at no charge by auix.net
Austria (VIX)	VIX services are provided on a not-for-profit basis; hence tariffs aim at cost recovery only. VUCC will charge each VIX member an annual tariff of US$ 2 128, for standard equipment housing (up to three height units in a 19" rack) and US$ 709, per height unit for any further rack space needed. The annual charge includes the connection to one 10BaseT switch port, and overheads like floor space, UPS power, A/C and VIX management.
Belgium (BNIX)	Prices not published on Web site. Available on request.
Denmark (DIX)	The network cover all expenses in the establishment of its connection. A low speed connection is presently free of charge. A high speed connection at the FDDI ring is charged a yearly sum of US$ 1 740. Around the clock (7*24) access to the DIX site is charged an annual US$ 1 160 per network. Access during normal working hours is free of charge.

Table 7**. Internet Exchange Point Pricing** (continued)

France (GIX)	Tariffs as from 1st June, 1996 : managed GIX on shared Ethernet : US$ 9 893 per year (not including the leased line) plus cost of the leased line; managed GIX on switched Ethernet : US$ 12 177 per year (not including the leased line) plus cost of the leased line; managed GIX with joined Ebone access either US$ 9893 or US$ 12177 per year (depending on the choice of the managed GIX) plus cost of the leased line; an additional cost of US$ 4110 per year (exploitation and serial link between the GIX and Ebone EBS and exploitation) EBONE costs (without the US$ 8174 for the router) hosted GIX US$ 5023 per year (not including the leased line) plus cost of the leased line.
France (GNI)	The access to the PEP is free during the experimental phases, from June 1st 1997 until December 1st 1997. To get connected to the PEP you must lease a leased line from the PEP (Proximity Exchange Point) to your local Point Of Presence.
Ireland (INEX)	New members, on joining, will pay the agreed annual fee (currently US$ 3106). New members joining more than halfway through the year will pay half of the annual cost.
Netherlands (AMS-IX)	The costs of the connection to the physical infrastructure are as follows: Connection to the switched ethernet 10BaseT: US$ 531 per month (excl. VAT) Connection to the switched ethernet 100BaseTX: US$ 902 per month (excl. VAT)
Sweden (D-GIX)	All ISPs are welcome to connect to the D-GIX. There is a modest fee to cover housing and power expenses. Connecting ISPs are also responsible for providing the local loop to reach the D-GIX. Fee not published on home page.
Switzerland (SIX)	Connection to the SIX is not accompanied by any form for compensation between the parties except for a monthly fee of US$ 37 covering costs housing (for 1 router/modem), electricity (for 1 router), accessing the SIX Ethernet segment.
Switzerland (CERN-IXP)	CERN will charge a standard fee for establishing a new connection and an annual fee for maintaining each connection. This will include physical access to the CERN-IXP (initially Ethernet or FDDI, ATM later), as well as overheads such as floor space, UPS electricity, air conditioning and administration. The standard CERN-IXP connection charges during the 1996-1997 start up period, payable in one settlement, are the following: One time installation charge: US$ 1485. Recurrent monthly charge: US$ 495. First year total (full year): US$ 7426. The prices above only apply to shared Ethernet connections. In case, the ISP requires a dedicated Ethernet segment, a port on an FDDI concentrator, or an ATM connection, extra charges will apply based on full cost recovery basis.
United Kingdom (LINX)	New members joining in the first half of the year will pay the agreed joining fee (currently US$ 14749) for new members for that year, plus the half-annual fee (currently US$ 7375) for existing members for that 6 month period.
United Kingdom (MaNAP)	New members joining in the first half of the year will pay the agreed joining fee (currently US$ 2950, proposed) for new members for that year, plus the half-annual fee (currently US$ 1475, proposed) for existing members for that 6 month period.
United Kingdom (i-Exchange)	There is an initial joining fee, of US$ 2950 (10baseT port) or US$ 7375 (100baseT port), along with a half annual fee of US$ 1475. Additional ports on the i-Exchange switch can be provided for a one off setup charge, currently US$ 2950 (10baseT) or US$ 7375 (100baseT).
United States (Boston MXP)	The Boston MXP is priced on a cost recovery basis. There are two choices of connectivity and the pricing is listed below: US$ 500 per year and US$ 500 install for either 10 megabit or 100 megabit switched access. Peering and Transit services will carry additional charges. MAI will, of course, peer with anyone who connects, but that does not include Transit.

Table 7. **Internet Exchange Point Pricing** (continued)

United States (Baltimore NAP)	Membership in the Baltimore NAP is currently free. Space for a CSU/DSU, and router is provided at no charge by ABSnet for a period to last no longer then June, 1998, after which time ABSnet has the right to re-negotiate rack space rental at reasonable, current rates in line with other ABSnet co-lo space and other NAP co-lo space, if it so chooses. ABSnet understands it is providing the space free of charge as a "seed" donation for the building of a successful NAP, and that charges for space and NAP access, at appropriate rates, are a very real possibility in the future to recover it's costs of service. Membership fees for the NAP will probably be instituted as the NAP incurs any operating costs. NAP members are encouraged to donate needed equipment to the NAP to defray operating costs.
United States (CIX)	Yearly fees for CIX members connecting to the CIX router, in addition to a one time $1 500 installation fee, are : T-1 Port (US$ 10000); SMDS T-1 Port (US$ 5000); SMDS Above T-1 (US$ 7500); Frame Relay T-1 (US$ 5000); FDDI (Co-located at Digital for US$ 15000).
United States (PAC Bell NAP)	Our NAPs are located in the San Francisco Bay Area (Service Area 1) and the Los Angeles area (Service Area 5). Pacific Bell will work with any interexchange carrier of the NAP customer's choice. Rates are: DS3 (45 Mbps) US$ 1500 (installation), US$ 5500 (monthly service); OC3c (155 Mbps) US$ 3000 (installation), US$ 6956 (monthly service)

1. Currency conversion for US$ is based on 1996 purchasing power parities. Wording is taken from sites concerned.

Source: OECD.

International traffic exchange and infrastructure financing

Before examining the different aspects of how international backbone networks are financed it is necessary to note how the current arrangements emerged over private and public networks. This distinction is important because private and public networks have historically had very different models of traffic exchange. The terms "private" and "public" do not indicate ownership status but rather whether the use of such networks were restricted to a defined group of users or open to the public. In most countries this distinction was significant in times past because virtually all telecommunication services offered to the public were reserved for monopoly operators. In other words, a user who leased dedicated capacity across the PSTN could not provide certain services to the public. For example, "reselling" PSTN capacity or providing switched telecommunication services would have placed a private network operator in contravention of government imposed monopolies. The reason users opted for private networks was that dedicated capacity was priced less expensively by telecommunication carriers than transporting the large amounts of the customer's own traffic over public switched networks.

Traffic exchange between private networks, in those instances where it occurred, did not incur a settlement or interconnection fee. If a private network spanned different geographical monopolies, such as in the case of an international leased line, then the user paid both operators for their nominal half of the circuit.

By the time NSFNET was created, the United States had liberalised its telecommunication market, but the Internet was still considered a private network because it was restricted to a defined group of users. Qualified academic or research institutions could join NSFNET and exchange traffic among themselves on a non-commercial basis. The non-commercial nature of NSFNET was the reason for the acceptable use policy. This aimed to preclude commercial traffic being exchanged over "public infrastructure" owned by US telecommunication carriers but leased by NSFNET via funding from the United States government.

Thus, to summarise developments, in its first guise as a US military network, the question of commercial arrangements for Internet traffic exchange between networks owned by different parties was not at issue. In its next guise, as a US academic network, managed under various contracts with telecommunication infrastructure providers, the issue of commercial traffic exchange was similarly not at issue because users were not in the business of selling communication services to each other. As the Internet was opened to academic institutions in other countries the same situation applied. To connect to the NSFNET a foreign academic network needed to purchase an international leased line and pay a telecommunication carrier in its home country and carrier at the US end for both their halves of the circuit.

According to the NSF, network managers outside the United States wanted to connected to the NSF backbone network for two reasons. The first reason was that the United States was the location of "so many of the Internet resources" and initial Internet infrastructure.[47] The second reason was the pricing on international circuits. According to the NSFNET management, the earlier introduction of competition in the United States made that country's half of the international leased lines much less expensive than that of most other countries. This meant that network managers had an incentive to connect to the NSF backbone network rather than directly to another country.

When the acceptable use policy was ended and the Internet opened for ISPs to sell communication carriage to the public the inherited models of traffic exchange were maintained. As previously noted, the word most commonly used to describe these arrangements is peering. While other methods of interconnection payment exist among ISPs, peering still accounts for the bulk of international traffic exchange. This means that individual ISPs generally pay telecommunication carriers the full cost of the circuits connecting their networks to international peering points. In other words an ISP still pays for two half-circuits but, due to peering, carries another ISP's traffic. In those cases where the ISP owns the international infrastructure, such as a telecommunication carrier owning a share of an undersea cable, the ISPs pay the full cost of both circuit halves to reach the international peering point. Accordingly, the Internet model for financing international infrastructure, for those networks that peer, effectively shifts the financial mid-point for traffic exchange from oceans (as in the case of cables) and geostationary orbit (as in the case of satellites) to Internet exchange points.

The above description of the evolution of traffic exchange is not controversial. However, the way these arrangements have evolved is emerging as a contentious issue among some of the owners of backbone IP networks connecting the Asia-Pacific region to the United States. In a competitive telecommunication world, it is natural for end-to-end services to emerge across international borders. The corollary is that companies will build their own end-to-end infrastructure or pay full circuit costs to the providers of this infrastructure. However the question remains open for commercial negotiation as to how ISPs that do not wish to peer want to charge each other for transit.

The chief complaint of the Asia-Pacific carriers such as Telstra of Australia and KDD of Japan is that in shifting the financial mid-point to US Internet exchange points they are paying the full cost of international carriage for both their customers and customers of ISPs based in the United States.[48] Both KDD and Telstra have argued that the costs of the international links should be met by the parties using that infrastructure. KDD says that in Asia many carriers are already sharing carriage costs among themselves. The Asia-Pacific Internet Association (APIA), has also raised the issue of funding the Internet's international backbone networks.[49] In late 1997 APIA requested comments on how future Internet infrastructure growth should be financed and the models that might best guide this development. In addition, Telstra has mounted legal action in the United States in an attempt to get the US Federal Communications Commission to review this issue.[50]

In support of Telstra's plea for the discussion of financing of international infrastructure costs and APIA's explanation of its call for comments, these entities note the shifting balance of traffic. Before NSFNET was retired as the Internet's premier backbone network, in May 1995, data on inbound and outbound traffic for the United States showed a large imbalance. More content was originating from the United States and being shipped to the rest of global Internet than was originating elsewhere and being terminated in the United States. By 1997 Telstra estimated the flow was of the order of 70:30 United States-to-Australia versus Australia-to-the United States.

It is at this point in the debate that the issues involved begin to get more complex. If the traffic flowing between an Asia-Pacific country and the United States was contained within a single international link and the geographical borders of both countries a model for sharing infrastructure costs might be self-evident. However the Internet does not transport traffic in such a precisely defined or bounded way.[51] Not only might the IP packets travel along different paths, traversing different countries, but so might the routing information drawn from global root servers. From the perspective of ISPs in the United States, this means that the greater the capacity foreign ISPs put in place to US-based Internet exchange points, the greater their costs in providing domestic infrastructure used for international transit. While there is little information available to understand the balance of these costs, or even information on patterns of international traffic (including transit traffic and traffic related to DNS requests) necessary to inform an exercise aimed at allocating costs on an international basis, it is possible to indicate how connectivity is evolving across the entire Asia-Pacific region.

Asia-Pacific regional and inter-continental internet connectivity

A 1992 paper by members of the Intercontinental Engineering and Planning Group (IEPG) had a discussion of the considerations for the Asia-Pacific region in connecting to NSFNET (Box 3). This paper noted that the Asia-Pacific region's Internet connectivity was almost exclusively via the United States. This included traffic exchange among entities in the same country and all transit traffic destined for or originating in Europe. In addition to tariff considerations the paper noted "policy considerations" relating to connectivity having been configured in this way. These considerations included routing stability with the authors expressing the opinion that putting in place direct independent routes between the Asia-Pacific region and Europe might "erode stability".

By 1993, there were still virtually no intra-regional links between Asia-Pacific countries except those transiting North America (see Appendix 1 for maps prepared by the Asia-Pacific Networking Group[52]). The exception was a leased line between Korea and Japan. The first independent direct connections to Europe were added in 1994 from Japan, Korea and India and a second intra-regional link between Japan and Chinese Taipei was also added. Over the next two years, many new intra-regional and links to Europe were added. However not only were the number of direct links to the United States increasing but so was the size of these connections. In terms of large-scale connectivity, by the first half of 1996 there were only three connections with 34 Mbps or higher capacities. Two of these were from Japan to the United States and one from Australia to the United States. In the second half of 1996, the first regional link of 34 Mbps or higher was added between Korea and Japan, as well as one addition link between Japan and the United States. In 1997, as the APNG maps indicate, the amount of connectivity across the Pacific spiralled with more than 20 links of 45 Mbps or higher. At the same time, there was only one intra-regional link of this scale (between Korea and Japan) and no direct links from the Asia-Pacific region to Europe of this scale.

The pace at which demand for Internet services drove new capacity allocation across the Pacific is unprecedented. At the start of January 1995 Telstra's international links for IP traffic amounted to 6 Mbps. A year later this had increased to 82 Mbps and by August 1997 had climbed to 140 Mbps. Telstra stated in late 1997 that demand for international Internet capacity was growing at around 10 Mbps per month and accelerating. Much of this demand is being translated into bandwidth allocated across the Pacific because this is where the available capacity IS to be found. Thus, even if North America is not the location for the most Internet users and the most accessed content, it might still be the case that a great deal of transit traffic, including intra-regional traffic and traffic to and from Europe, needs to be passed over the most readily available international capacity. For example, of the undersea cables originating and terminating in Japan, some 70 per cent of the total capacity is between Japan and the United States (Table 8) and only 30 percent to regional or other destinations.

It should be noted that international patterns of traffic for the Internet are radically different from those for international PSTN traffic. Singapore provides one example of these differences. The premier destination for Singapore's international PSTN traffic is Malaysia with just over one-third of all outgoing traffic being terminated in that country. By way of contrast, only 0.4 per cent of the international capacity allocated for global Internet connectivity, by Singapore Telecom's Internet Exchange (STIX) connects directly to Malaysia (Table 9) (Appendix 3, Figure 4). The difference in Internet and PSTN traffic patterns between Singapore and Canada is even more extreme. While less than one per cent of Singapore's outgoing PSTN traffic terminates in Canada, a massive 46.7 per cent of the capacity allocated for Singapore, global Internet connectivity is on a route to Vancouver.

How traffic traverses Asia-Pacific links is a result of routing policies, which depend on peering or transit agreements between partners. Accordingly, a packet may not travel by the most direct or obvious route between two countries. A user in Singapore wanting to access Japanese content might find that traffic goes via the United States. A further consideration is the quality of service within the Asia-Pacific region compared to intercontinental links. It is still the case that for some Asia-Pacific countries packets travel faster across the Pacific than between relatively close neighbours. For example, the average time it takes for packets to make a round trip between Singapore and the rest of the Asia-Pacific region is double the time it takes for a round trip to North America (Table 10) (Appendix 3, Figure 5). Even more remarkable is that a round trip from STIX to Europe can be eight times faster than the regional average. In the latter case, this may indicate that the backbone between STIX and Monaco is relatively under-utilised compared to other links.

It is not easy to discern what overall impact the fact that intra-regional service is often slower than intercontinental has on network planners and users. A current axiom of Internet infrastructure planning is to try to take content and services closer to users. Yet even if a content or service provider creates a mirror site in an effort to provide a more localised service, it may be still quicker for a user to access the same site on another continent rather than in a nearby country. This is because the "local infrastructure" in some countries and regions does not perform as well as intercontinental links. This problem is not confined to the Asia-Pacific region. According to studies of backbone performance undertaken by Keynote Systems:

> "Deploying a mirrored Web site in Europe to serve European users may not increase performance for those users. Our measurements show that a web server geographically close to its users can often deliver worse performance than a more geographically remote server."[53]

Accordingly, users may opt to use popular sites such as search engines at distant rather than at local sites generating increased intercontinental traffic. As for network planners, Keynote Systems say their studies lead them to conclude that:

> "Network engineers at backbone providers tend to focus on optimising traffic flow within their own networks. They tend to de-emphasise or ignore connectivity and end-to-end response time to users on other networks. The Internet, however, is an interconnection of many backbones and private networks, with the result that users rarely access Web sites that are directly connected to the same backbone they are."[54]

It is even more difficult to discern what all this means for discussions and negotiations over the financing on international infrastructure. Telstra and KDD's point about shifting the financial mid-point from within international links to Internet exchange points is a good one, but it is not clear how this should be interpreted in terms of different parties paying a fair share of infrastructure costs. Indeed before this question could be addressed in any meaningful way at great deal of information would need to be gathered. For example, the issue of Internet traffic transiting via a third country would need to be studied in greater detail than has been done to date as would the costs of providing infrastructure for this traffic. Moreover, while some have cast this as a debate between certain Asia-Pacific and US-based ISPs, it clearly has implications for ISPs in all countries.

Before proceeding to discuss further the issue of financing international infrastructure, it is interesting to consider how ISPs, in countries that have largely not been party to this debate, might be affected. For example, how might ISPs in Canada or Europe view this discussion, particularly given the remarkably large amount of capacity between STIX and Canada compared to that between STIX and Europe. The first thing to note is that it is self-evident that the capacity allocated between STIX and Vancouver is far in excess of what would be required for transporting traffic solely between Asia-Pacific users and Canadian users. This capacity is, of course, being used to transit traffic to other destinations and is related to hubbing strategies. This means it would be extremely complex to try to determine some type of cost allocation among all beneficiaries.

The example of connectivity between STIX and Vancouver demonstrates the incompatibility of traditional cost allocation in financing international infrastructure, such as sharing of half-circuit costs, with current Internet developments. Nevertheless the question of cost allocation is a valid one if the current arrangements are not equitable in terms of use made of infrastructure relative to financial contribution. Here, it is interesting to contrast the situation in Europe with that in the Asia-Pacific region. In the traditional PSTN model, European and Asia-Pacific ISPs would have paid half-circuit costs. In the Internet model, given the small amount of direct connectivity, it seems as if the financial mid-point has been shifted to North America. For European ISPs, the current arrangements may be a less expensive option to connect to the Asia-Pacific regional than, for example, sharing direct half-circuit costs.

While listed half-circuit prices for leased lines are not what users pay, they provide one starting point for analysis of the cost to an ISP of purchasing intercontinental links. This analysis shows that if an ISP in Belgium, Norway and the UK purchased a 2 Mbps leased line to the United States and paid for both half-circuits it might be less expensive than sharing the average of a half-circuit cost for the same link to Australia and Japan (Table 11). At the same time because the discounts are likely to be greater on routes to the United States than to Asia-Pacific countries this situation probably applies to ISPs in many other European countries. Recent analysis shows that it is often less expensive for European ISPs to purchase trans-Atlantic capacity, to traffic exchange points in the United States, than to purchase equivalent trans-European capacity to European Internet exchange points.[55] Moreover, the countries with less expensive trans-Atlantic prices may be the points of departure for trans-European networks to aggregate international traffic. For example, UUNET's largest connections between Europe and the United States are from the United Kingdom.

A further factor in the financing of international infrastructure is how ISPs, with their own intercontinental infrastructure, view these issues either because they own it directly or via alliances. One reason why some European telecommunication carriers may not view the issue in the same way as some Asia-Pacific carriers is that they have been more active in investing in the US market. For example, France Telecom and Deutsche Telekom's investments in Sprint, which has one of the largest Internet backbone networks, mean that international IP traffic exchange could be handled by partners within the Global One alliance. Similarly, Cable and Wireless has a leading IP backbone network (CWIX) in the United States which it can connect to its other global networks. At the same time US-based ISPs have been much more aggressive in entering European markets to offer services to business and consumers than in the Asia-Pacific region. Companies such as WorldCom/UUNET own and manage international networks providing end-to-end services for their business and dial-up customers. [56]

Companies such as UUNET also sell transit across the United States and around the world for other ISPs. This transit is priced at both flat and measured rates.[57] The latter service allows users to have usage-sensitive prices for "burstable connections" enabling large amounts of data to be transported in bursts, during periods of high demand, as required. The growth in scale of UUNET's network capacity across the Atlantic quadrupled to 45 Mbps in August 1996 from the beginning of that year, and the doubled to 90 Mbps by October 1996. That same month, UUNET's parent company, WorldCom, announced that a US$ 500 million, 10 Gbps transatlantic link, was to be constructed and that it is expected to commence service in 1998.[58]

Liberalisation of international infrastructure markets has underpinned the transfer of the financial mid-point from transmission links to public and private Internet exchange points. The bargaining power each party has at a Internet exchange point no longer rests on regulation, as it did when financial mid-points resulted from the connection of national monopoly networks. In the new environment, bargaining strength relies on several other factors. One factor is the network reach which parties can offer each other via their own networks or those of their partners. As the leading global Internet hub, ISPs in the United States are in the strongest bargaining position. The fact that the most accessed Internet content resides in the United States adds to this strength because the customers of ISPs in other parts of the world want direct connections to this content. Foreign ISPs want to reduce the number of hops between networks experienced by their customers to minimise response times. Direct peering and transit connections at public and private US NAPs and other exchange points are one way to minimise the number of hops. This is driving the creation of high-speed international connections to the United States in advance of some intra-regional connections elsewhere in the world.

In the newly liberalised environment for international infrastructure provision regulators would, no doubt, be loath to intervene in the commercial negotiations between ISPs. Moreover, it is not easy to discern what would be the consequences of such intervention. If US policy makers mandated that US-based ISPs had to pay half-circuit prices it would push financial mid-points away from Internet exchange points and NAPs and back to transmission links. For European and Asia-Pacific ISPs this would increase the financial appeal of linking to each other via the United States. This might increase congestion at North American Internet exchange points for users and generate new investment demands on ISPs based in the United States. From the perspective of US-based ISPs they are already paying for a global hubbing function for the Internet. On the other, hand the existing system means that European and Asia-Pacific traffic will continue to add to the burdens of the US-based Internet exchange points because this is the least expensive option for Europe, and the Asia-Pacific ISPs have a more limited range of options.

In the face of such challenges, Asia-Pacific ISPs do have some strategies available. First, the creation of attractive content at the national or ISP level would place ISPs in a stronger position to negotiate joint sharing of infrastructure costs. Second, there is a growing incentive to use caching

technologies at national or ISP levels to reduce the amount of international and national bandwidth required and therefore produce a lower cost. Third is the creation of Internet exchange points (and numerous MXPs) in those countries and cities that are not well served. Fourth is for ISPs, which are not presently PTOs, to consider construction of their own infrastructure. According to one example given in *CommunicationsWeek International*, owning cable infrastructure is considerably less expensive than leasing capacity.[59] According to one commentator, leasing a 2-Mbps circuit between the US mainland and Australia from two PTOs would typically cost US$ 98 000 a month. A 2-Mbps link, known as a minimum investment unit (MIU), for the same route would cost only $12 638 per month -- an 87 per cent discount.[60] Fifth could be the deployment of IP multicasting on an international basis for the most accessed audio and video content, in an effort to reduce the increase in traffic due to webcasting. Sixth, could be a strategy by PTOs, outside the United States, to offer discounts to ISPs in an effort to become an attractive hub along the lines STIX.

Box 3: Extract from "Connectivity within the Internet - A Commentary", by Geoff Huston (AARNET), Elise Gerich (MERIT), and Bernhard Stockman (SUNET/NORDUnet), 1992. http://www.iepg.org/docs/IEPG-connect.html

From the perspective of the global Internet the picture is not complete without adding the connectivity issues of the Asia / Pacific region. Here policy objectives and link tariff constraints dictate that there is little international regional infrastructure within the region (unless you consider that Hawaii and the United States itself are an integral part of the Asia / Pacific region!). The overall structure of regional connectivity here is that each national entity (and in some cases more than one entity within a nation) implements its primary connectivity through a link into the US infrastructure, as being an implementation of their primary policy objective. There are no regional *IX, *BONE or *EX structures in place, as a consequence (in part) of these policy objectives and tariff constraints.

The current picture of connectivity within the region is there are direct connections into Hawaii, the US west coast and even the US east coast, and a large proportion of inter-regional traffic transits portions of the US infrastructure as a consequence. At this point in time there are no direct infrastructural Pacific regional connections to the European infrastructure. While from a routing management perspective this may be considered to be a reasonable position, when one takes into account the fact that there are at present no open transit paths within the United States itself, the end result is that inter-Pacific and Pacific / European connectivity is constrained by the policy provisions of the various US entities who are in a position to undertake a transit role for such traffic. Unless this situation is altered in the near future a natural consequence of the growth in connectivity requirements will be a number of Pacific regional entities making direct connections to the European infrastructure, adding even further to the routing management complexities and further eroding the levels of stability of the overall Internet connectivity structure

Table 8. **International undersea cable capacity to and from Japan (July 1997)**

Cable Name	Japan - United States. Number of 64-kbit/s circuits	Cable name	Japan - Regional and other, Number of 64-kbit/s circuits
TPC-3, Chikura - Hawaii	3 780	H-J-K, Chikura - Hong Kong	3 780
TPC-4, Chikura - Point Arena	7 560	H-J-K, Chikura - Korea	3 780
TPC-5CN, Miyazaki - Hawaii	120 960	APC Miura- HongKong	7 560
TPC-5CN, Bandon - Ninomiya	120 960	APC Miyazaki-Toucheng APC Miyazaki-Singapore C-J FOSC Miyazaki-Nanhui R-J-K, Naoetsu- Russian Fed. R-J-K, Naoetsu-Korea APCN, Miyazaki - Asia	7 560 7 560 7 560 7 560 7 560 60 480
Total	253 260		113 400
Per cent of total	69.1		30.9

1. This table does not include KDD capacity on cables that do not originate or terminate in Japan, such as across the Atlantic.

Source: KDD Annual Report.

Table 9. **Singapore's Internet and traditional PSTN traffic patterns**

	Outgoing MiTT as per cent of total outgoing PSTN traffic (1996)	Internet capacity as a per cent of total capacity allocated for Internet on international links from Singapore (STIX) to global Internet (November 1997)
Malaysia	34.5	0.4
Indonesia	8.5	1.9
Hong Kong	7.4	0.4
United States	4.8	24.9
Japan	4.8	6.8
Australia	4.8	1.6
China	4.8	0.4
Thailand	3.7	0.0
United Kingdom	3.2	0.0
India	3.2	0.2
Philippines	3.2	0.4
Korea	1.4	0.4
Brunei	0.9	6.2
France/Monaco	0.9	6.2
Canada	0.8	46.7
Total of above	86.9	96.5

1. MiTT (Minutes of Telecommunication Traffic) is for 1996. The MiTT share for Korea and France are based on 1995 data. The MiTT for Canada is an upper bound estimate.
2. STIX stands for Singapore Telecom Internet exchange. Refer to http://www.stix.net/.

Source: OECD based on STIX, Telegeography.

Table 10. **STIX Regional capacity and performance**

STIX to:	Capacity (kbit/s)	Per cent of total capacity in operation	Average round trip time (ms) (1)
Vancouver	15360	46.7	331
Los Angeles	4096	12.5	390
San Francisco	4096	12.5	424
Brunei	2048	6.2	270
Monaco	2048	6.2	97
Tokyo	1984	6.0	248
Tokyo	1920	0.0	Not in operation during survey
Surabaya	512	1.6	792
Sydney	512	1.6	508
Tokyo	256	0.8	289
Jakarta	192	0.6	1331
Beijing	128	0.4	602
Hong Kong	128	0.4	435
Johor Bahru	128	0.4	763
Jakarta	128	0.4	487
Lahore	128	0.4	1475
Manila	128	0.4	516
Noumea	128	0.4	1208
Seoul	128	0.4	272
Taipei	128	0.4	281
Bombay	64	0.2	291
Colombo	64	0.2	1235
Dhaka	64	0.2	1578
Dhaka	64	0.2	908
Dhaka	64	0.2	1244
Karachi	64	0.2	1408
Karachi	64	0.0	Not in operation during survey
Kathmandu	64	0.2	1677
Lahore	64	0.2	1837
Noumea	64	0.2	1208
Phnom Pehn	64	0.2	1096
Total (2)	32896	100.0	808
Total to Asia-Pacific Region (3)	7296	22.2	885
Total to North America	23552	71.6	382
Total to Europe	2048	6.2	97

1. These are averages taken from the STIX site in late November early December 1997 over several weeks at different times of day. Data from STIX: refer to http://www.stix.net/. Destinations mentioned two or more times have multiple links.
2. Excluding capacity not in operation. Simple rather than weighted average.
3. Excluding North America and capacity not in operation.

Source: OECD

Table 11. **International leased line prices (2Mbps, $US)**

From/To	Australia	Japan	United States	Ratio (%) of average price from Europe to Australia/Japan compared to price to United States.
Norway	58445	58445	20920	279
United Kingdom	68215	68215	29990	227
Belgium	61064	61064	27926	219
France	70030	41787	30046	186
Switzerland	N/A	43292	23252	186
Ireland	51760	51760	28468	182
Netherlands	30444	30444	16710	182
Denmark	61736	61736	36525	169
Germany	49944	44910	28659	165
Poland	34166	62634	34166	142
Finland	52104	52104	37778	138
Portugal	60656	60656	47146	129
Sweden	35034	35034	27604	127
Italy(2)	54534	54534	44223	123
Spain	45354	45354	45354	100
Average	52392	51465	31918	163

1. Monthly half-circuit charge. Price data from Tarifica.

2. The prices in these tables are subject to rapid change. Between the time of writing and publication the prices in Italy were reduced to US$ 50 050 for the route to Australia and Japan and US$ 40 600 for the route to the United States (using the May 1998 exchange rate of Lit 1775/US$).

Source: OECD.

MAPPING THE INTERNET BY DOMAINS

All indications are that the amount of traffic being transported over the Internet is increasing very rapidly. As a result infrastructure providers are having to escalate the amount of capacity they are making available on existing national and international routes. Few data are available to enable analysis of national and international IP traffic flows on the Internet, but all indications are that it is fundamentally different from PSTN traffic patterns. The OECD report "Webcasting and Convergence: Policy Implications" showed that by far the most accessed content on the Internet was located in the United States, and more particularly in California. To complement that analysis of content location, this section attempts to give a better indication of levels of Internet host penetration and users. It makes allowances for generic TLD registrations, which account for the bulk of Internet domain name addresses (Box 4). The other reason for examining gTLD registrations is to better understand how traffic may be generated between some countries because of the relative use of gTLDs as opposed to TLDs.

Distributing domain name registrations

By September 1997, there were just over 2 000 000 domain names registered in the world (Table 12). The largest number of registrations occur under the gTLDs which account for just under 75 per cent of all domain addresses. The greatest number of domain registrations in the gTLD category is under **.com**. Registrations in national registries of TLDs (such as **.be** for Belgium) account for just over 25 per cent of all domain addresses.

A small number of OECD countries utilise gTLDs far more than other countries. By far the largest user of gTLDs is the United States followed by Canada. By mid 1997 some 98 per cent of all second level domains registered in the United States were under gTLDs and only 2 per cent registered under the **.us** domain name (Appendix 3, Figure 6). The next largest user of gTLDs is Canada. In fact Canadians have opted to use gTLDs over the **.ca** TLD by a ratio of nearly five to one. The only other OECD countries where the number of gTLD registrations at InterNIC outnumber TLD registrations in national registries are France, Korea and Spain. By way of contrast, the balance of domain registrations in other OECD countries favours TLDs rather than gTLDs. In the Czech Republic, New Zealand and Poland, more than 90 per cent of all registrations are made under **.pl**, **.nz** and **.cz** respectively.

Within the United States the largest concentration of domain registrations is in California (Table 13). In September 1997, California had just under 250 000 active gTLD registrations, just under 22 per cent of all gTLD registration in the United States. At the same time, this represented just over 13 per cent of the world's combined total for gTLDs and TLDs. In fact California, in its own right, is the largest market for domain names in the world, with twice as many registrations as the next largest markets in Canada and the United Kingdom. Significantly, US states in their own right make up 19 of the top 30 domain name markets.

Registries located in OECD countries for gTLD and TLD registrations were responsible for 93 per cent of the world's domain addresses by mid 1997. Some 96 per cent of gTLDs and around 85 per cent of all TLDs are registered in OECD countries. That being said users in non-OECD countries, on average, tend to favour registration under gTLDs over TLDs much more than users in most OECD countries (Table 14). Outside OECD countries, the largest markets for gTLDs are Hong Kong, China and the Arab Emirates. Brazil has the greatest total number of gTLDs and TLDs outside the OECD area. The locations outside the OECD area with the largest number of gTLDs per 1 000 inhabitants are the British Virgin Islands, Liechtenstein and the Cayman Islands (Table 15).

Reasons why users prefer registering domains under various gTLDs and TLDs are examined in the OECD document "Domain Names: Allocation Policies".[61] The benefits cited for registering under a gTLD, such as **.com**, are that they are viewed as not being identified with a particular geographical location. This is a particularly attractive quality for a company with international operations. It is also true that as domain name users have opted for **.com**, particularly in North America, this has built momentum owing to the belief that Internet users will use the DNS as a directory service. In other words, users wanting to find a company or subject will try the company name or subject area followed by **.com** by way of instructing their browser.

Other reasons for registering under a gTLD, rather than a TLD, include the relative efficiency of registries. Even in 1997 some TLD registries are still staffed on a voluntary basis, although most have recently moved to a more commercial footing. This has generally meant that users opting for a gTLD have received a more efficient service at InterNIC than at their national registry. Moreover, the price for TLD registration, where the service is charged for, is usually more expensive than InterNIC's gTLD registration. In addition national TLDs often have restrictive policies in terms of who can register and how many registrations a user is allowed. In some countries individuals are not permitted to register under TLDs and must have a presence in that country. Moreover, in some countries multiple registrations are not permitted under TLDs. A further consideration is that some TLD registries only allocate third-level domain names to some users. Accordingly, some users may prefer the simplicity of a second-level domain name under a gTLD or may not want to be tied to a particular second-level TLD category (e.g. a particular state or province). The other reason some users prefer gTLDs is that they do not want to be directly identified with a particular country because of regulatory or legal considerations. For example, the most accessed Web sites conducting online gambling all use **.com** domain addresses rather than the TLDs for their operational locations.

Weighting Internet host distribution

The most common indicator used to measure Internet development is the survey of Internet hosts undertaken by Network Wizards and RIPE (Réseaux IP européens). The Network Wizards survey includes all gTLD and TLD registrations and is undertaken every six months. The RIPE survey is undertaken monthly but is limited to European TLD registrations. While both surveys are much appreciated by the Internet community the results need to be qualified and have several limitations. The first qualification that needs to be made is that host data do not indicate the total number of users who can access the Internet. At best, they may be interpreted as the minimum size of the public Internet.

The second factor that has made it difficult to undertake comparative analysis between countries has been that there was no way to distribute Internet hosts under gTLD registrations on a national basis. In other words, the reachable hosts of a user in France registering under a gTLD would appear under domains such as **.com** or **.net** rather than **.fr**.

The availability of gTLD registrations by country gives the first possibility to consider redistributing Internet hosts under domain names such as **.com** to individual countries. The simplest option is to weight the number of hosts under gTLDs according to the number of gTLD registrations from a particular country. In other words if 5 per cent of the total gTLD registrations are from a particular country, then 5 per cent of the total number of hosts surveyed under gTLDs are reallocated to that country. This methodology would, no doubt, be subject to a number of caveats. Nevertheless, it seems reasonable to assume that this approach gives a more accurate distribution of Internet hosts in OECD countries than allocating all hosts under gTLD registrations to the United States.

The results of the weighted methodology are most striking in the case of Canada where there is a 72 per cent increase in the number of hosts over the number of hosts surveyed solely under **.ca** (Table 16) (Appendix 3, Figure 7). Other countries to record significant increases, albeit from smaller base numbers of hosts, were France, Luxembourg, Spain and Turkey. All these countries recorded a relatively large increase in the number of hosts relative to the average OECD increase of 21 per cent. The countries for which this made very little difference are those where users mainly rely on national TLD registrations such as Iceland, the Czech Republic, Finland, New Zealand and Poland.

In terms of the ranking countries by the number of Internet hosts per 1000 inhabitants, compared to the previous methodology, Canada rises three places and Luxembourg climbs two places. The United States falls five places, as would be expected, and Hungary declines two places. Most other countries either hold their ranking (12 countries) or rise or fall one place (13 countries).

Box 4: The Internet Domain Name System

The DNS essentially maps Internet addresses. It works for any Internet service that requires domain names: e-mail, WWW, FTP and so on. To function as part of the Internet a host needs a domain name that has an associated Internet Protocol (IP) address record. This includes any computer system connected to the Internet via full or part-time, direct or dial-up connections. A top-level domain name (TLD) is either an ISO country code (for example, .be stands for Belgium) or one of the generic top level domains (a so called gTLD such as .com, .org, .net). Internet domain names consist of a number of domains joined together by a dot (".") following a form similar to the following example: www.oecd.org. This example has three separate domains (www, oecd, and org). There can be four or more domains within the domain name, but it is often impractical to go much beyond four. The domains follow a hierarchy where the left-most domain is the lowest level, and the right-most domain (known as the top level domain) is the broadest coverage. Left of the TLD is the is the second level domain (i.e. oecd), then a third level domain if applicable and so on. The OECD's domain name is registered under a generic top level domain (i.e. .org) with InterNIC. This name provides a user friendly address which overlays a numeric address (i.e. 204.180.228.0). The OECD has also registered a second-level domain name under a TLD in France (i.e. oecd.fr), although this address is not currently used as part of the Organisation's universal resource locator (URL) on the World Wide Web or for the e-mail address.

In 1997 two registration processes operated for gTLD and TLD addresses. In 1995 the US National Science Foundation contracted Network Solutions, Inc., to manage certain gTLDs and authorised the company to charge users for registration from September 1995. In most OECD countries, the registrars responsible for TLDs have followed the reforms instituted at InterNIC. Most commonly, registrars, once housed in universities, have been spun off into private companies or associations run by members. Typically, the members of organisation-based registrars are ISPs, and in a growing number of countries, the marketing of TLDs is done through their offices. TLD registrars have also followed InterNIC's lead and introduced fees for registration. As a result the incentives for registrars of TLDs, and resellers, to market these services have radically changed between 1995 and 1997. Whereas university-based registrars did not market TLDs in a commercial manner this was changing by 1997. A commercial customer of a registrar, who may once have been referred to InterNIC, will now have a TLD recommended.

Table 12. **Domain registration in OECD countries**

	gTLD registrations (September 1997)	TLD domains (July/Aug-97)	Total domain registrations	gTLD as % of total registrations	TLD as a % of total registrations
United States	1141455	21590	1163045	98.1	1.9
Canada	101540	21753	123293	82.4	17.6
Spain	12486	5829	18315	68.2	31.8
France	19733	10769	30502	64.7	35.3
Korea	6346	5686	12032	52.7	47.3
Turkey	3067	3664	6731	45.6	54.4
Sweden	14281	19745	34026	42.0	58.0
Italy	11407	17377	28784	39.6	60.4
Netherlands	12028	18948	30976	38.8	61.2
Belgium	3302	6220	9522	34.7	65.3
Switzerland	8006	15061	23067	34.7	65.3
Germany	29542	56059	85601	34.5	65.5
United Kingdom	38615	76916	115531	33.4	66.6
Ireland	1360	2899	4259	31.9	68.1
Luxembourg	317	765	1082	29.3	70.7
Japan	11492	31143	42635	27.0	73.0
Greece	528	1531	2059	25.6	74.4
Finland	1890	5687	7577	24.9	75.1
Austria	2721	9013	11734	23.2	76.8
Portugal	684	2325	3009	22.7	77.3
Australia	7646	28523	36169	21.1	78.9
Mexico	1400	5827	7227	19.4	80.6
Norway	2340	10124	12464	18.8	81.2
Denmark	6609	30873	37482	17.6	82.4
Hungary	349	1961	2310	15.1	84.9
Iceland	50	573	623	8.0	92.0
Czech Republic	273	4010	4283	6.4	93.6
New Zealand	874	12929	13803	6.3	93.7
Poland	259	5544	5803	4.5	95.5
OECD	1440600	433344	1873944	36.1	67.3
Non-OECD	63340	74603	137943	57.7	42.3
World	1503940	507947	2011887	74.8	25.2

1. Iceland's gTLD registration is an OECD estimate.
2. OECD average is a simple average rather than a weighted average. The weighted average is 76.9 for gTLDs and 23.1 for TLDs. The average for non-OECD is a simple average. The weighted average for non-OECD is 45.9 for gTLDs and 54.1 for TLDs.
3. gTLD registration data supplied by Imperative (http://www.imperative.com/ and http://www.internet.org).

Source: OECD.

Table 13. **US Domain name markets**

	gTLD registration by US State (1)	US State share of gTLDs	State share of global gTLD/TLDs	Rank	Top 30 Domain Markets
California	248 467	21.8	13.3	1.	California
New York	92 322	8.1	4.9	2.	Canada
Florida	76 196	6.7	4.1	3.	United Kingdom
Texas	68 633	6.0	3.7	4.	New York
Illinois	44 292	3.9	2.4	5.	Germany
Massachusetts	41 624	3.7	2.2	6.	Florida
New Jersey	41 133	3.6	2.2	7.	Texas
Pennsylvania	37333	3.3	2.0	8.	Illinois
Washington	31 960	2.8	1.7	9.	Japan
Ohio	31 094	2.7	1.7	10.	Massachusetts
Virginia	29 551	2.6	1.6	11.	New Jersey
Georgia	25 367	2.2	1.4	12.	Denmark
Colorado	24 942	2.2	1.3	13.	Pennsylvania
Michigan	24 691	2.2	1.3	14.	Australia
Arizona	22 912	2.0	1.2	15.	Sweden
Maryland	22 439	2.0	1.2	16.	Washington
Minnesota	20 666	1.8	1.1	17.	Ohio
Oregon	19 676	1.7	1.0	18.	Netherlands
North Carolina	18 644	1.6	1.0	19.	France
Connecticut	17 524	1.5	0.9	20.	Virginia
Missouri	15 771	1.4	0.8	21.	Italy
Indiana	15 007	1.3	0.8	22.	Georgia
Wisconsin	14 465	1.3	0.8	23.	Colorado
Tennessee	12 954	1.1	0.7	24.	Michigan
Nevada	11 197	1.0	0.6	25.	Switzerland
Utah	10 896	1.0	0.6	26.	Arizona
Kansas	10 515	0.9	0.6	27.	Maryland
Louisiana	8 344	0.7	0.4	28.	Minnesota
District of Columbia	8 242	0.7	0.4	29.	Oregon
Other	92 618	8.1	4.9	30.	North Carolina

1. Includes state registrations under .com, .net. .org, and .edu. Excludes state registrations under the .us domain.
2. gTLD registration data supplied by Imperative (http://www.imperative.com/).

Source: OECD.

Table 14. **Domain registration in selected non-OECD countries**

	gTLD registrations (September 1997)	TLD domains (July/Aug-97)	Total domain registrations	gTLD as % of total	TLD as a % of total
Cayman Islands	146	0	146	100.0	0.0
Virgin Islands (British)	726	6	732	99.2	0.8
Arab Emirates	5 342	92	5434	98.3	1.7
Saudi Arabia	721	19	740	97.4	2.6
Virgin Islands (US)	316	9	325	97.2	2.8
Bahamas	824	28	852	96.7	3.3
India	4 218	281	4499	93.8	6.2
Panama	493	33	526	93.7	6.3
Barbados	253	22	275	92.0	8.0
Netherlands Antilles	198	21	219	90.4	9.6
Hong Kong	10 316	1616	11932	86.5	13.5
Pakistan	537	190	727	73.9	26.1
Jordan	280	119	399	70.2	29.8
Dominican Republic	228	112	340	67.1	32.9
Venezuela	1 268	650	1918	66.1	33.9
Thailand	1 621	859	2480	65.4	34.6
Liechtenstein	324	184	508	63.8	36.2
Indonesia	1 674	1019	2693	62.2	37.8
Trinidad & Tobago	169	109	278	60.8	39.2
Philippines	887	605	1492	59.5	40.5
China	6 125	4534	10659	57.5	42.5
Costa Rica	483	364	847	57.0	43.0
Ecuador	234	178	412	56.8	43.2
Chinese Taipei	1 915	1784	3699	51.8	48.2
Bermuda	179	187	366	48.9	51.1
Colombia	750	899	1649	45.5	54.5
Guatemala	115	152	267	43.1	56.9
Malaysia	1 218	1670	2888	42.2	57.8
Lebanon	190	374	564	33.7	66.3
Israel	1 955	3902	5857	33.4	66.6
Egypt	126	253	379	33.2	66.8
Brazil	4 824	9885	14709	32.8	67.2
Singapore	2 126	4525	6651	32.0	68.0
Argentina	1 809	8370	10179	17.8	82.2
South Africa	1 489	9055	10544	14.1	85.9
Chile	253	1787	2040	12.4	87.6
Slovenia	153	1358	1511	10.1	89.9
Russian Federation	578	5915	6493	8.9	91.1
Slovakia	63	1856	1919	3.3	96.7
Other countries	8 214	11581	19795	41.5	58.5
Total non-OECD	63 340	74603	137943	57.7	42.3

1. Non-OECD average is a simple rather than a weighted average. The weighted average is 45.9 for 54.1
2. gTLD Registration data supplied by Imperative (http://www.imperative.com/).

Source: OECD

Table 15. **Domain registrations per 1 000 inhabitants**

OECD	Active gTLD registrations per 1 000 inhabitants	Total gTLD and TLDs per 1 000 inhabitants	Non-OECD	Active gTLD registrations per 1 000 inhabitants	Total gTLD and TLDs per 1 000 inhabitants
Denmark	1.26	7.17	Virgin Islands (British)	55.02	55.48
United States	4.34	4.42	Liechtenstein	10.38	16.27
Canada	3.43	4.16	Bermuda	2.88	5.89
New Zealand	0.24	3.86	Cayman Islands	4.21	4.21
Sweden	1.62	3.85	Bahamas	2.94	3.04
Switzerland	1.13	3.26	Virgin Islands (US)	2.87	2.95
Norway	0.54	2.86	Arab Emirates	2.24	2.28
Luxembourg	0.77	2.62	Singapore	0.71	2.22
Iceland	0.19	2.33	Hong Kong, China	1.67	1.93
Netherlands	0.78	2.00	Netherlands Antilles	0.99	1.10
Australia	0.42	2.00	Barbados	0.97	1.06
United Kingdom	0.66	1.97	Israel	0.35	1.04
Finland	0.37	1.48	Slovenia	0.08	0.76
Austria	0.34	1.46	Colombia	0.19	0.42
Ireland	0.38	1.19	Slovakia	0.01	0.36
Germany	0.36	1.05	Argentina	0.05	0.29
Belgium	0.33	0.94	South Africa	0.04	0.25
France	0.34	0.52	Costa Rica	0.14	0.25
Italy	0.20	0.50	Trinidad & Tobago	0.13	0.21
Spain	0.32	0.47	Panama	0.19	0.20
Czech Republic	0.03	0.41	Chinese Taipei	0.09	0.17
Japan	0.09	0.34	Malaysia	0.06	0.14
Portugal	0.07	0.30	Chile	0.02	0.14
Korea	0.14	0.27	Lebanon	0.05	0.14
Hungary	0.03	0.23	Jordan	0.06	0.09
Greece	0.05	0.20	Brazil	0.03	0.09
Poland	0.01	0.15	Venezuela	0.06	0.09
Turkey	0.05	0.11	Russian Federation	0.004	0.04
Mexico	0.02	0.08	Dominican Republic	0.03	0.04

1. Data are for mid-1997. gTLD registration data supplied by Imperative (http://www.imperative.com/).

Source: OECD

Table 16. **Internet host penetration weighted by gTLD registration**

	Internet hosts as reported by Network Wizards (July/Aug 1997)	Estimated additional hosts based on active gTLD registrations (September 1997).	Total hosts	Change in number of hosts (%)	Traditional presentation of Internet host penetration per 1 000 inhabitants (mid-1997)	Internet Host penetration per 1 000 inhabitants weighted by gTLD registration (Mid-1997)
Finland	335956	9316	345272	2.8	65.8	67.6
Iceland	14153	246	14399	1.7	53.0	53.9
Norway	209034	11534	220568	5.5	47.9	50.6
New Zealand	155678	4308	159986	2.8	43.5	44.7
Australia	707611	37686	745297	5.3	39.2	41.3
Canada	690316	500476	1190792	72.5	23.3	40.2
Sweden	284478	70389	354867	24.7	32.2	40.2
United States	11829141	-1773488	10354699	-15.0	45.0	38.2
Denmark	137008	32575	169583	23.8	26.2	32.4
Switzerland	148028	39460	187488	26.7	20.9	26.5
Netherlands	341560	59284	400844	17.4	22.1	25.9
United Kingdom	878215	190328	1068543	21.7	15.0	18.2
Luxembourg	3854	1562	5416	40.5	9.3	13.1
Austria	87408	13411	100819	15.3	10.9	12.5
Germany	875631	145608	1021239	16.6	10.7	12.5
Ireland	33031	6703	39734	20.3	9.2	11.1
Belgium	86117	16275	102392	18.9	8.5	10.1
Japan	955688	56642	1012330	5.9	7.6	8.1
France	292096	97261	389357	33.3	5.0	6.7
Czech Rep.	49104	1346	50450	2.7	4.8	4.9
Italy	211966	56223	268189	26.5	3.7	4.7
Spain	121823	61542	183365	50.5	3.1	4.7
Korea	132370	31279	163649	23.6	3.0	3.7
Hungary	33818	1720	35538	5.1	3.3	3.5
Portugal	18147	3371	21518	18.6	1.8	2.2
Greece	19711	2602	22313	13.2	1.9	2.1
Poland	43384	1277	44661	2.9	1.1	1.2
Turkey	22963	15117	38080	65.8	0.4	0.6
Mexico	35328	6900	42228	19.5	0.4	0.5

1. gTLD registration data supplied by Imperative (http://www.imperative.com/).

Source: OECD

Global root-level servers

DNS servers perform the necessary function of translating back and forth between names and numbers and are therefore a critical element in traffic exchange between ISP networks. These servers contain databases of IP addresses and corresponding domain names and are interrogated when a user wants to send an e-mail or request data over the World Wide Web.[62] For example, if a government user in Ottawa wanted to send an e-mail to a colleague in the Japanese Ministry for Posts and Telecommunications (e.g. **person@mpt.go.jp**), and copy that message to a colleague in Industry Canada (e.g. **person@ic.gc.ca**) their mail programme would initiate a request to the DNS server of that person's Internet service provider.

Owing to the fact that the Canadian colleague's machine is hosted on the same government network, his/her address would be located on the same DNS server. By way of contrast, because the Canadian government DNS server would not contain a record for the domain name **.go.jp** it would initiate a request to a root-level server. Root-level servers contain databases with information about which DNS servers on the Internet act for which domain names. The root-level server would, in this case, point the Canadian DNS server to a counterpart that knows the IP address for the **.go.jp** domain name. Following this, a request is made to the DNS server hosting **.go.jp** and the DNS server then returns the IP address which receives mail for **person@mpt.go.jp**. The e-mail can then be sent and received.

There are 13 global root-level servers in support of the IANA recognised DNS (Table 17). Four contain information for the root alone and these "root only" servers were launched to test the feasibility of maintaining the root information separate from the information for the top level domains.[63] InterNIC/Network Solutions, Inc., operate two of the global root-level name servers. In May 1997, a third global root-level name servers, formerly operated by Network Solutions (InterNIC), was shifted to the United Kingdom. In addition a further global root server was established at Keio in Japan. There are now two such name servers located in Europe, one in Japan and ten in the United States. Under the current addressing system it has been suggested that 13 global root servers if the optimal number (Box 5).[64]

Discussion of location and management of these global root level servers takes place in the Internet Engineering and Planning Group (IEPG). The IEPG is an Internet operational group intended to assist ISPs to interoperate within the global Internet. The goals and activity domains of the IEPG are described in the IEPG Charter, and are summarised in RFC1690. In June 1996 the IEPG discussed,

> "... the logical location of root nameservers in respect to provider and exchange topology. It was noted that while the placing of a root server on an exchange was a topological neutral location both in terms of traffic flow and in terms of relation to potential inter-provider settlement structures, the location did hamper effective management of the root name service. Placing a root server within a provider network offers improved management capability and places the onus on the provider to provide high quality connectivity to the server, and was generally considered to be a more stable deployment structure."[65]

A further consideration of location has been the funding of root servers. At present a number of the root-level servers are ultimately funded by governments, with US taxpayers making the largest contribution. As IEPG discussions bear witness, some of these root servers have relied on volunteer hosting. This has raised concerns as the Internet moves to being a fully commercially driven network and CIX has recommended:

"... that the US Government study the desirability and feasibility of full time professional maintenance of the DNS root servers. The root servers are the equivalent of the 'air traffic controllers' for the Internet. Few of us would like to imagine commercial aviation in which the air traffic control towers are managed by volunteers, or solely by the carriers themselves. The purpose of this study would be to ensure that these vital Internet components are -- and will continue to be -- fully secure but also are fully supported to minimise network disruption."[66]

Table 17. **Operators of the Internet's root-level name servers**

Operator	Status	Location
Network Solutions, Inc. (A)	Private company	United States http://www.netsol.com/
Army Research Laboratory (B)	Military	United States
Performance Systems International Inc. (C-NYSER)	Private company	United States http://www.psi.com/
University of Maryland (UMD-TERP) Computer Science Center	University	United States http://www.umd.edu/
NASA Ames Research Center, E	Government	United States
Internet Software Consortium (ISC),F	Non-profit organisation.	United States http://www.isc.org/isc/
GSI (DIIS-NS)	Military	United States http://www.nic.ddn.mil/
University of Southern California (ISI2) Information Sciences Institute	University	United States http://www.usc.edu/
US Army, H	Military	United States
Network Solutions, Inc. (J)	Private company	United States http://www.netsol.com/
M-Wide Keio	University: "Wide Project"	Japan
NORDUnet	Private company (1)	Sweden http://www.nordu.net/
RIPE NCC. (K)	European Network Co-ordination Centre http://www.ripe.net/	LINX (United Kingdom) http://www.linx.net

1. In May 1997 K.root-servers.net was shifted from Network Solutions (InterNIC) to be housed within LINX and managed by RIPE NCC. LINX is a London-based exchange point for Internet traffic in the United Kingdom and externally. Two other global root name servers exist at ISI at the University of Southern California, one of which was moved to Keio, Japan, in August 1997.

Source: OECD based on http://nic.mil/DNS/root-server.html and ftp://rs.internic.net/domain/named.ca.

Box 5: IEPG Advisory Note - Root Name Servers (Source: Bill Manning, October 1996. http://www.iepg.org/docs/IEPG-ADV-9610.html)

The question has been raised in the IEPG regarding reports of less than satisfactory responses to root nameservers. These reports were substantiated and it was suggested that a couple of actions be taken:

Separate the machines serving the root zone from those supporting top level domains.

Move some root nameservers to be "closer" to end users.

Add more machines as root nameservers.

It was noted that given current constraints, that all root servers need to be able to fit into a single UDP packet, which works out to a total of 13 machines, given the current naming scheme. There are currently 9 in service. Several discussions on placement of the remaining four nameservers ensued. The IEPG discussed direct attachment to exchanges, placement just off an exchange or within a service provider. It was suggested that perhaps the best approach, to meet the proposed guidelines in the manning-dnssrv draft, would be to place them, under contract from the IANA, one hop off exchange points and in control of an ISP. It was suggested that perhaps the best way to maintain the desired level of topological insensitivity would be to inject 13 host routes into the global routing system from a well known prefix. This would ensure that regardless of where these servers were placed over time, there would always be a route to them. Such a prefix has been identified and has been prepared for this purpose. The prefix in question is: 192.0.0.0/23 We will run some experiments using addresses within this prefix before deploying new root nameservers or migrating existing servers. We then discussed potential candidates and found no volunteers in the Asia-Pacific region, none in Africa and only one in Europe. The biggest concern was the level of global traffic that would attempt to resolve root queries. So, at this time, only Keith Mitchell of LINX has offered, with some qualifications, to host a root name server in Europe.

Intellectual Infrastructure Fund

An important element of the core infrastructure supporting the public Internet are the DNS databases. In September 1995, following NSF authorisation, Network Solutions began charging for gTLD registrations. The co-operative agreement issued by the NSF to Network Solutions, to provide InterNIC registration services, provides for 30 per cent of the funds collected for registration and renewal of domain names to be:

> "...placed into an interest-bearing account which will be used for the preservation and enhancement of the "Intellectual Infrastructure" of the Internet in general conformance with approved Program Plans. [Network Solutions] will develop and implement mechanisms to insure the involvement of the Internet communities in determining and overseeing disbursements from this account. [Network Solutions] will also establish and maintain publicly available records of all deposits to and disbursements from the account."[67]

Between September 1995 and August 1997, US$ 34.2 million was deposited into the "Intellectual Infrastructure" account by Network Solutions. From publicly available data it is not possible to determine how much of this money was contributed by users in different countries. Since the introduction of charges applicants have been charged US$ 100 per initial registration of a second level domain name with US$ 30 of that money being put into the "Intellectual Infrastructure" fund. The initial fee covered a period of two years. However, by using the number of active domains in each country, as at

September 1997, it is possible to give a good indication of likely proportional contribution by country for one year (Table 18). This shows that users in the United States have contributed around 76 per cent of money for the "Intellectual Infrastructure" fund and users in other countries 24 per cent. Outside the United States the largest contribution was made by Canadian users and the largest regional contribution made by European users. Reflecting their relatively low use of gTLDs OECD countries from the APNIC area (Australia, Japan, Korea and New Zealand) would have contributed less than 2 per cent.

In keeping with the original goal of the Intellectual Infrastructure, Fund the US Congress has proposed devoting a major part toward funding "Internet 2", a high-speed network connecting US research and university campuses. With respect to the contribution from users outside the United States, some have raised the question of whether this money should made available for similar worthy projects in appropriate countries or regions. A good case can be made for this point of view. On the other hand, this needs to be balanced against the ongoing funding by US taxpayers of certain parts of the DNS operation and management. In addition, many regard this fee as a form of "tax". They might point out that US users pay value added taxes for registering second level domains in many national TLD registries around the world. Given the higher pricing of these TLDs the tax often costs more than US$ 15 per annum per name. Those countries without charges for registering or without VAT could no doubt counter this point.

Table 18. **Indicative contributions to the Intellectual Infrastructure Fund**

	Active gTLDs (September 1997)	Indicative Annual Contribution to Intellectual Infrastructure fund, US$ (1)	Per cent of World Total
United States	1 141 455	17 121 825	75.90
Canada	101 540	1 523 100	6.75
United Kingdom	38 615	579 225	2.57
Germany	29 542	443 130	1.96
France	19 733	295 995	1.31
Sweden	14 281	214 215	0.95
Spain	12 486	187 290	0.83
Netherlands	12 028	180 420	0.80
Japan	11 492	172 380	0.76
Italy	11 407	171 105	0.76
Switzerland	8 006	120 090	0.53
Australia	7 646	114 690	0.51
Denmark	6 609	99 135	0.44
Korea	6 346	95 190	0.42
Belgium	3 302	49 530	0.22
Turkey	3 067	46 005	0.20
Austria	2 721	40 815	0.18
Norway	2 340	35 100	0.16
Finland	1 890	28 350	0.13
Mexico	1 400	21 000	0.09
Ireland	1360	20 400	0.09
New Zealand	874	13 110	0.06
Portugal	684	10 260	0.05
Greece	528	7 920	0.04
Hungary	349	5 235	0.02
Luxembourg	317	4 755	0.02
Czech Republic	273	4 095	0.02
Poland	259	3 885	0.02
Iceland	50	750	0.003
Total OECD	1 440 600	21 609 000	95.79
Total Non-US	362 485	5 437 375	24.10
EU Area	155 503	2 332 545	10.34
Europe (RIPE area)(2)	169 847	2547705	11.29
Asia (APNIC area)(2)	26 358	395370	1.75
Non-OECD	63340	950100	4.21
World	1503940	22559100	100.00

1. This represents the “Intellectual Infrastructure” contribution for one year of US$ 15 per second level domain registration under a gTLD. Some second level domain name owners would have paid this on an going basis over several years.
2. OECD countries only.

Source: OECD.

ANNEX 1

Figure 1: Peer/transit/internal backbone access to top 100 Internet sites via Global One

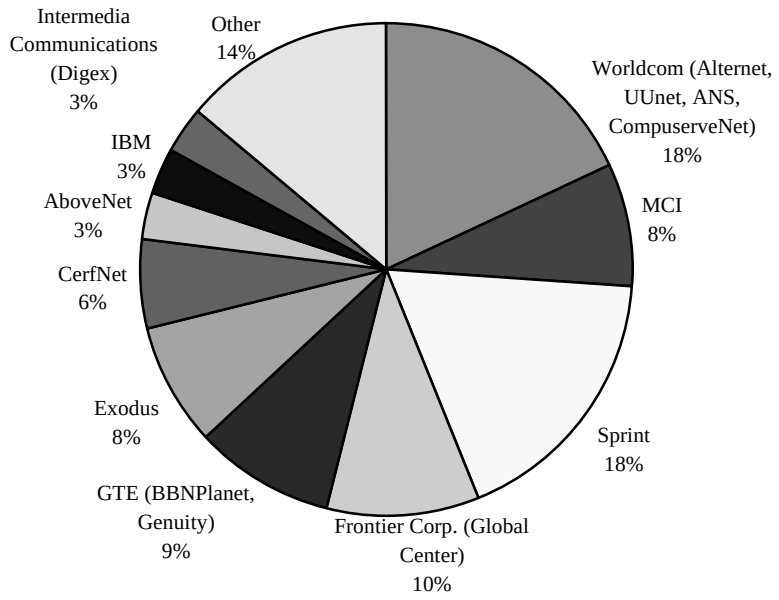

Hot100 from Web21/Trace from Global One

Source: OECD.

Figure 2: Peer/transit/internal backbone access to top 100 Internet sites via Worldcom

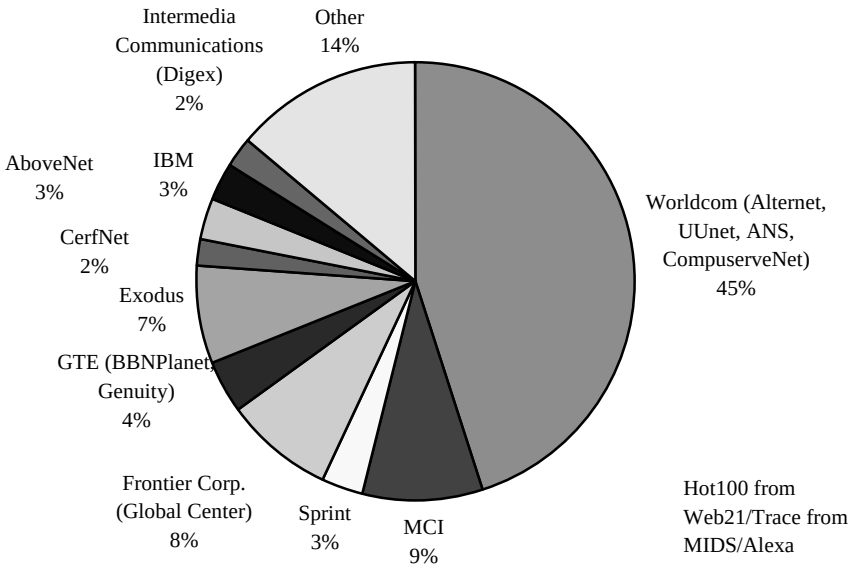

Source: OECD.

Figure 3: Peer/transit/internal backbone access to top 100 Internet sites via CERFNET

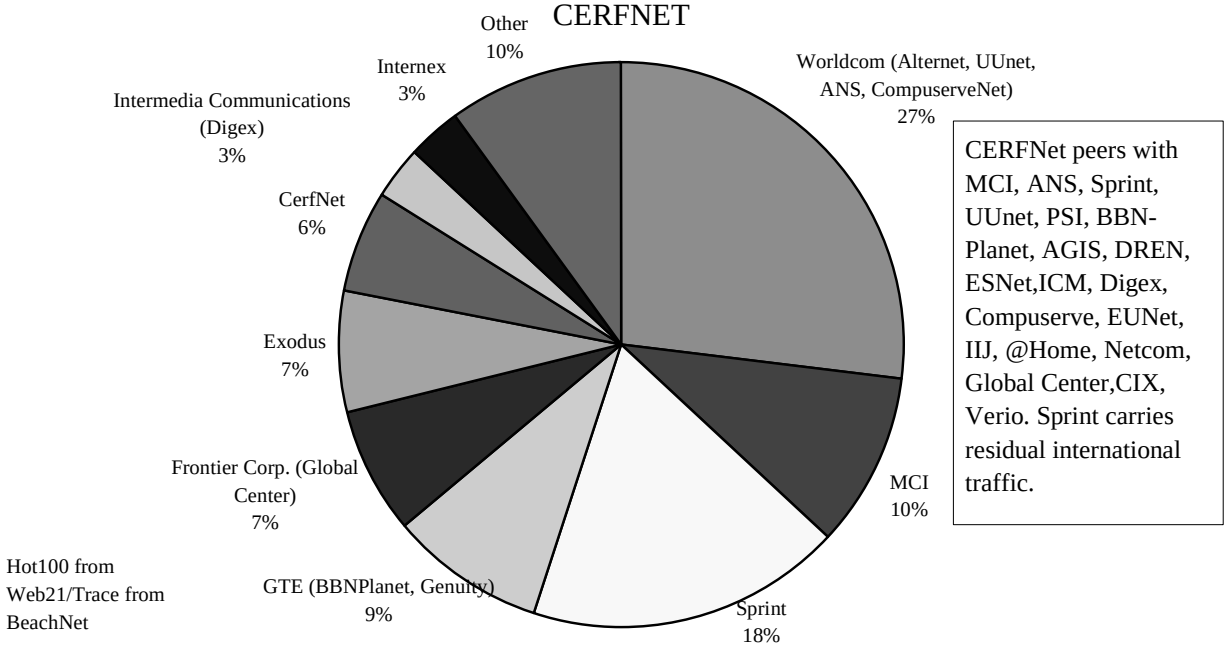

Source: OECD.

ANNEX 2

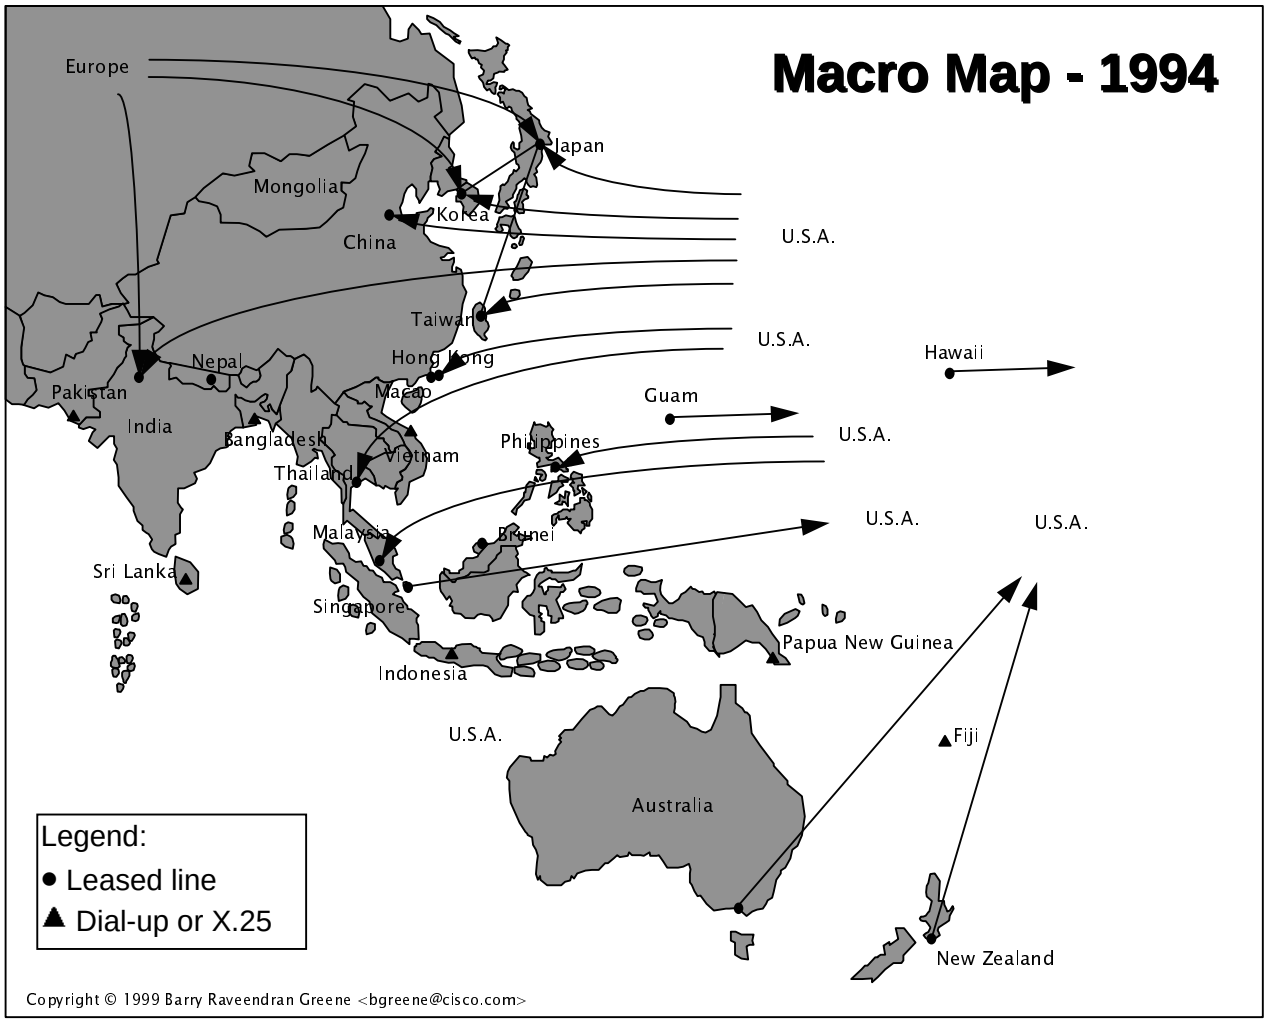
Macro Map - 1994
Europe
Japan
Mongolia
Korea
China
U.S.A.
Taiwan
U.S.A.
Hawaii
Nepal
Hong Kong
Pakistan
Macao
Guam
India
Bangladesh
Vietnam
Philippines
U.S.A.
Thailand
U.S.A.
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Papua New Guinea
Indonesia
U.S.A.
Fiji
Australia
Legend:
Leased line
Dial-up or X.25
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

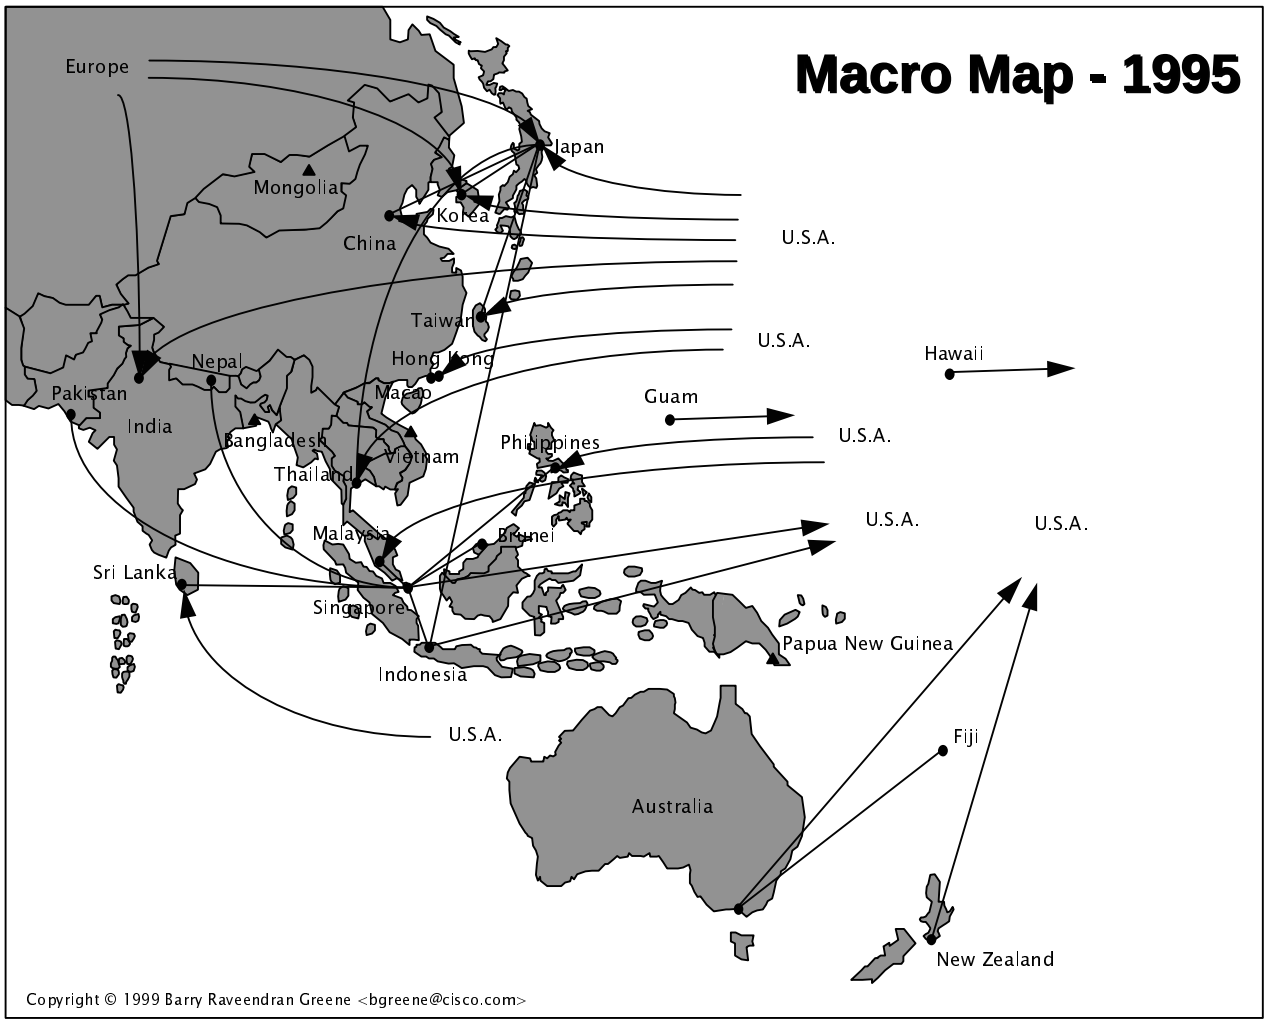
Macro Map - 1995
Europe
Japan
Mongolia
Korea
China
U.S.A.
Taiwan
U.S.A.
Hawaii
Nepal
Hong Kong
Pakistan
Macao
Guam
India
Bangladesh
Vietnam
Philippines
U.S.A.
Thailand
U.S.A.
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Papua New Guinea
Indonesia
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

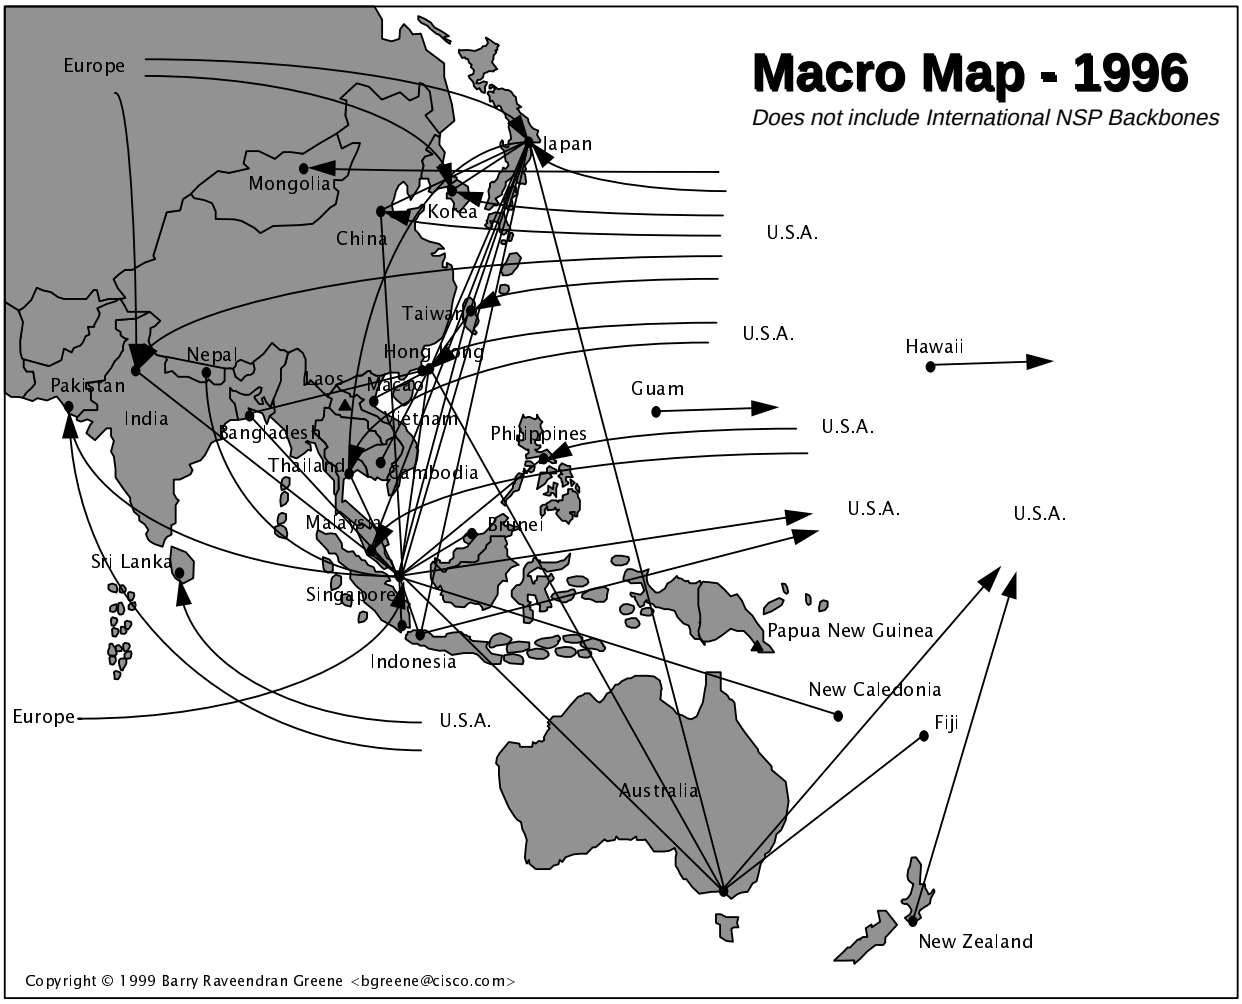
Macro Map - 1996
Does not include International NSP Backbones
Europe
Japan
Mongolia
Korea
China
U.S.A.
Taiwan
U.S.A.
Hawaii
Nepal
Hong Kong
Pakistan
Laos
Macao
Guam
India
Vietnam
U.S.A.
Bangladesh
Philippines
Thailand
Cambodia
U.S.A.
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Papua New Guinea
Indonesia
New Caledonia
Europe
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

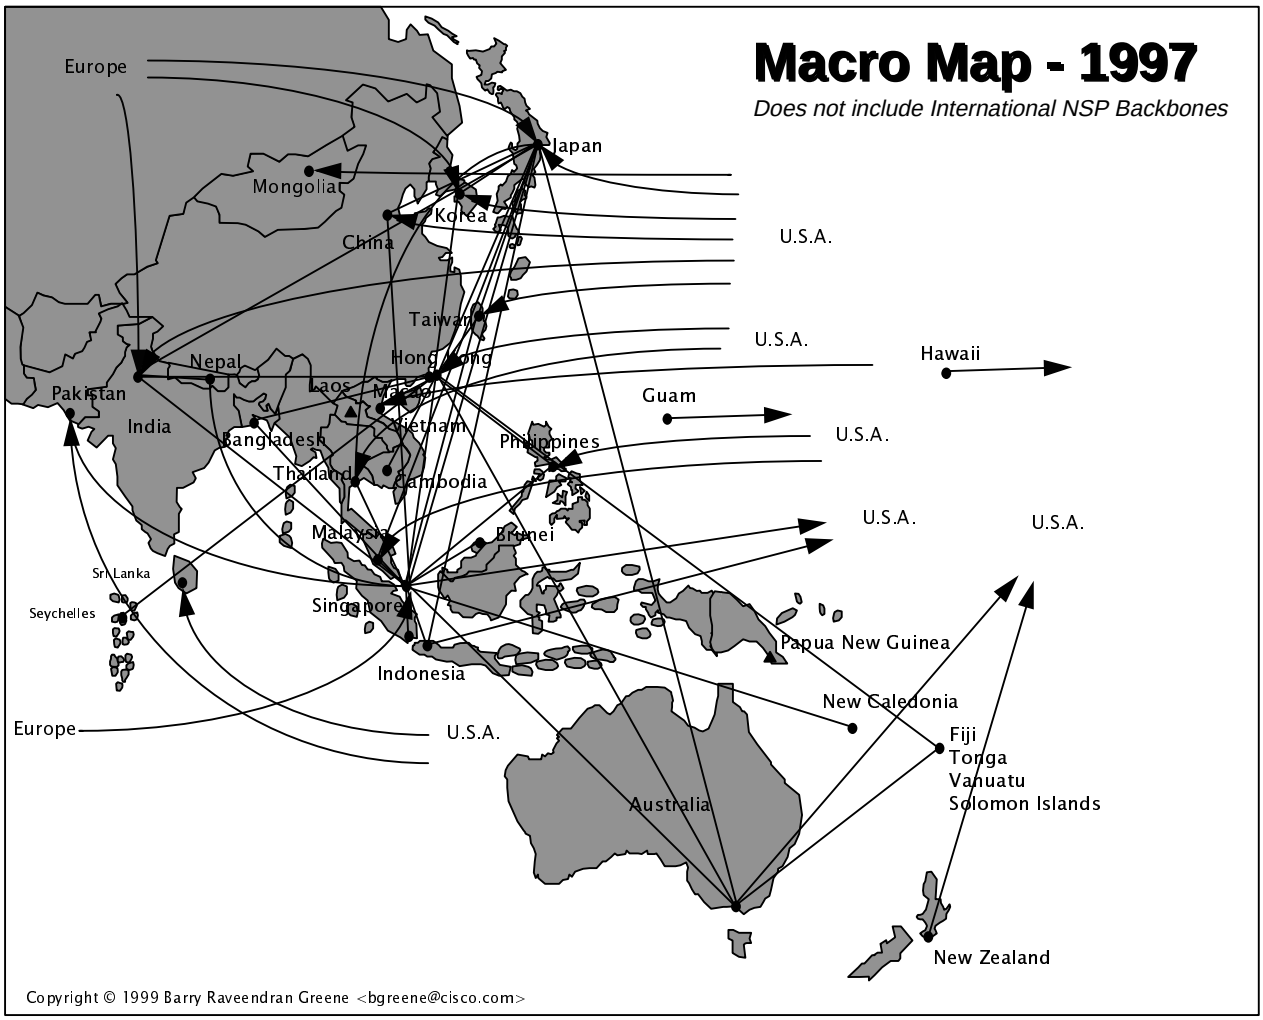
Macro Map - 1997
Does not include International NSP Backbones
Europe
Japan
Mongolia
Korea
China
U.S.A.
Taiwan
U.S.A.
Hawaii
Nepal
Hong Kong
Pakistan
Laos
Macau
Guam
India
Vietnam
U.S.A.
Bangladesh
Philippines
Thailand
Cambodia
U.S.A.
U.S.A.
Malaysia
Brunei
Sri Lanka
Seychelles
Singapore
Papua New Guinea
Indonesia
New Caledonia
Europe
U.S.A.
Fiji
Tonga
Vanuatu
Solomon Islands
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

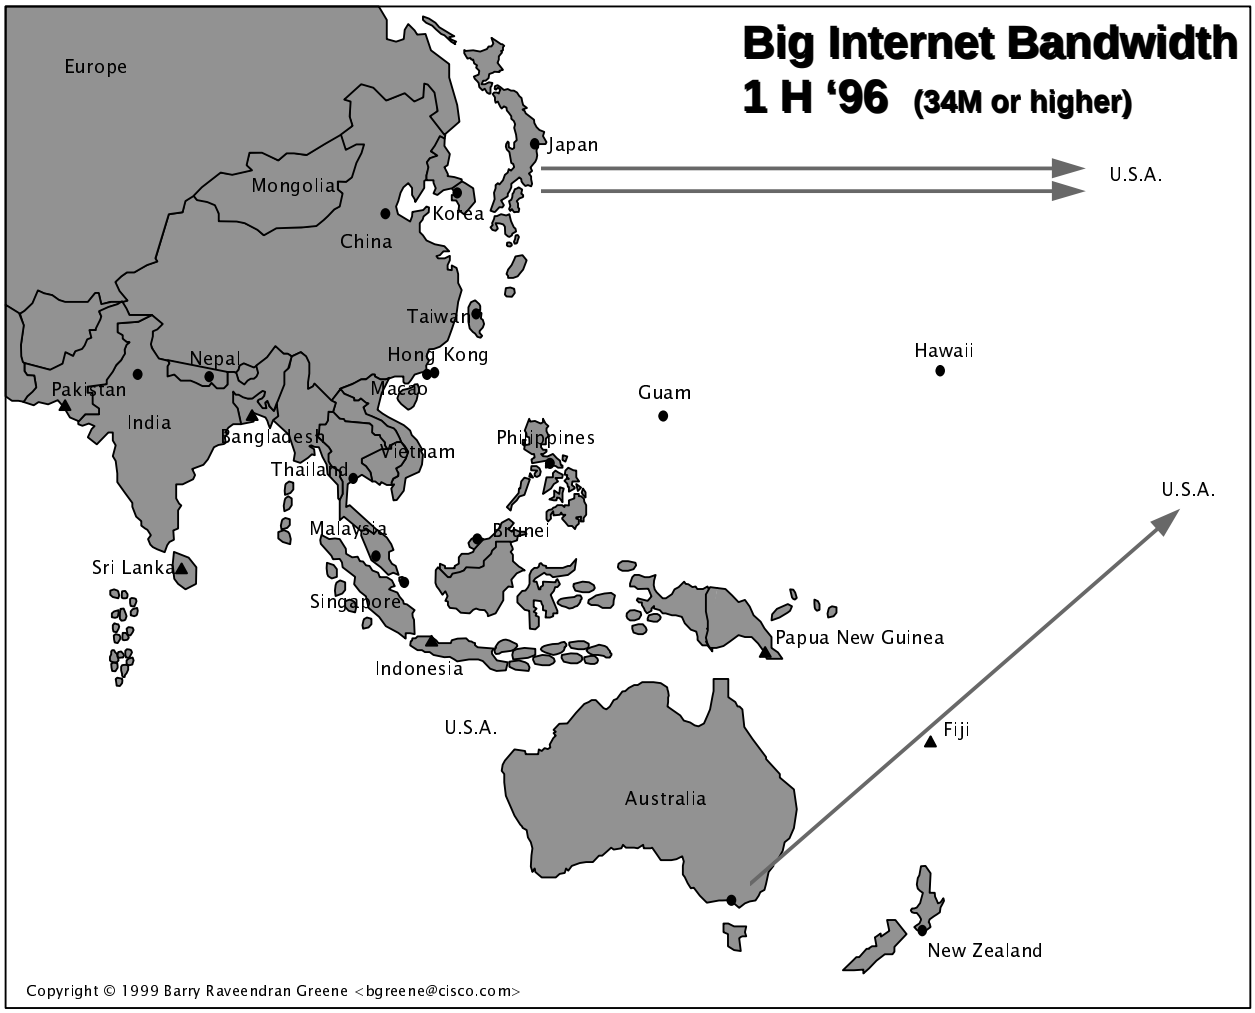
Big Internet Bandwidth
1 H '96 (34M or higher)
Europe
Japan
U.S.A.
Mongolia
Korea
China
Taiwan
Hong Kong
Nepal
Pakistan
India
Macao
Bangladesh
Vietnam
Thailand
Philippines
Guam
Hawaii
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Indonesia
Papua New Guinea
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

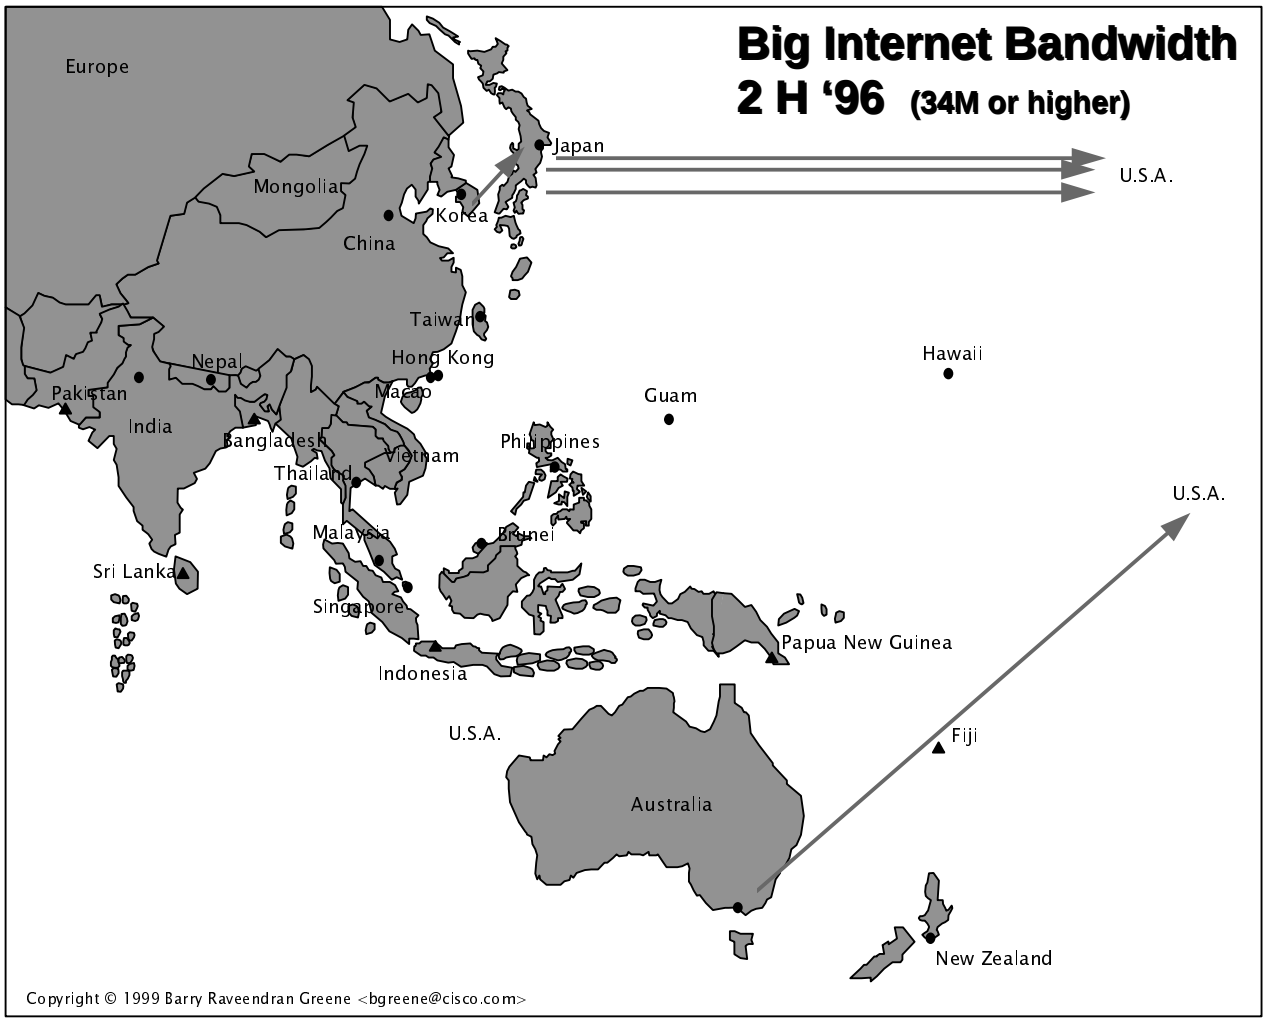
Big Internet Bandwidth
2 H '96 (34M or higher)
Europe
Mongolia
China
Korea
Japan
U.S.A.
Taiwan
Hong Kong
Macao
Nepal
Pakistan
India
Bangladesh
Vietnam
Thailand
Philippines
Guam
Hawaii
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Indonesia
Papua New Guinea
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

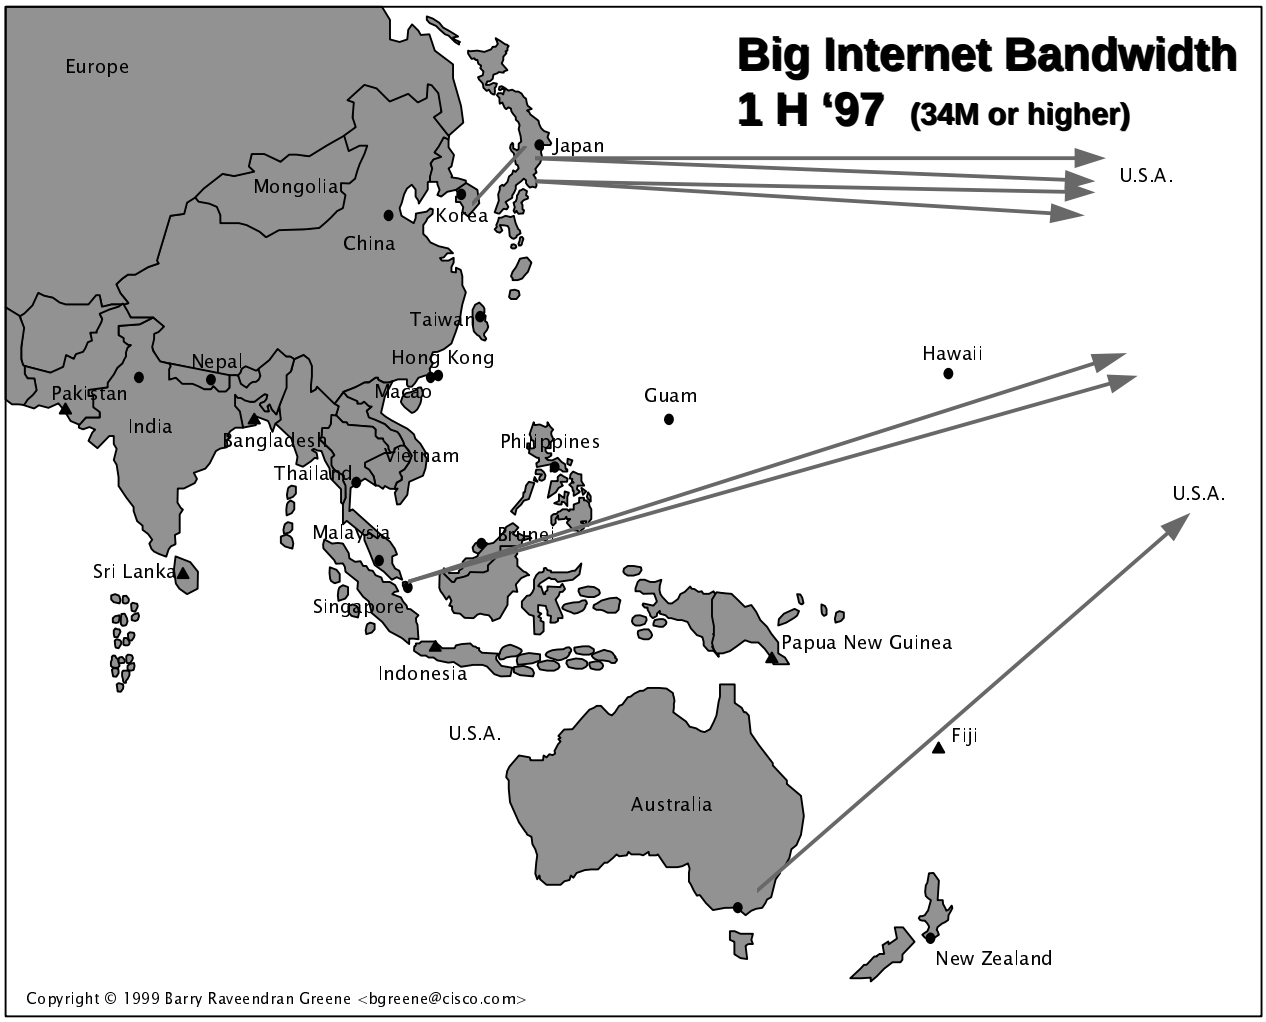
Big Internet Bandwidth
1 H '97 (34M or higher)
Europe
Mongolia
China
Korea
Japan
U.S.A.
Taiwan
Hong Kong
Macao
Nepal
Pakistan
India
Bangladesh
Vietnam
Thailand
Philippines
Guam
Hawaii
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Indonesia
Papua New Guinea
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

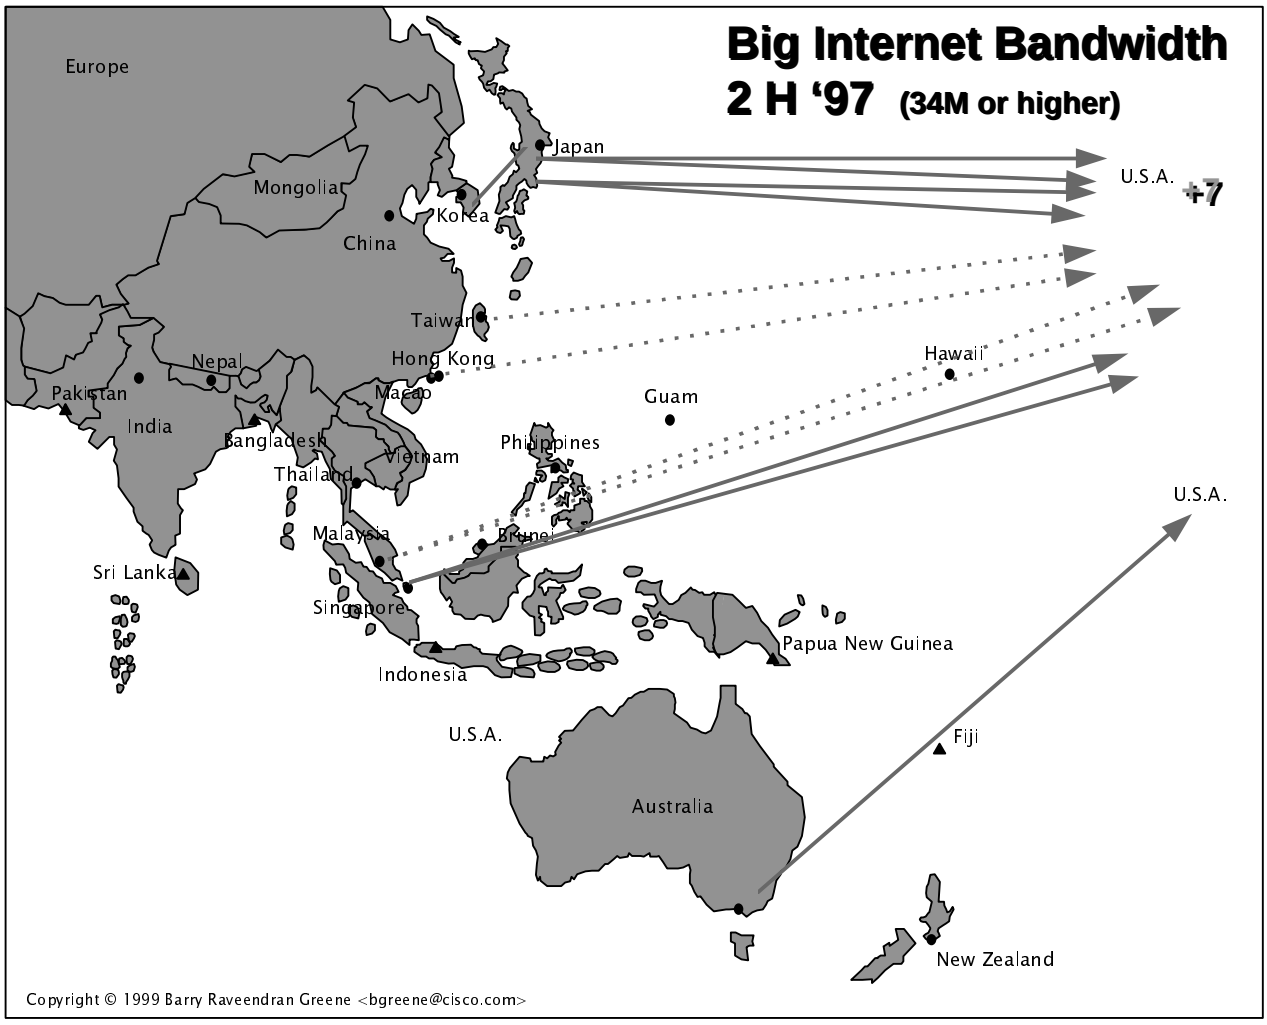
Big Internet Bandwidth
2 H '97 (34M or higher)
Europe
Mongolia
China
Korea
Japan
U.S.A.
+7
Taiwan
Hong Kong
Macao
Nepal
Pakistan
India
Bangladesh
Vietnam
Thailand
Philippines
Guam
Hawaii
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Indonesia
Papua New Guinea
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

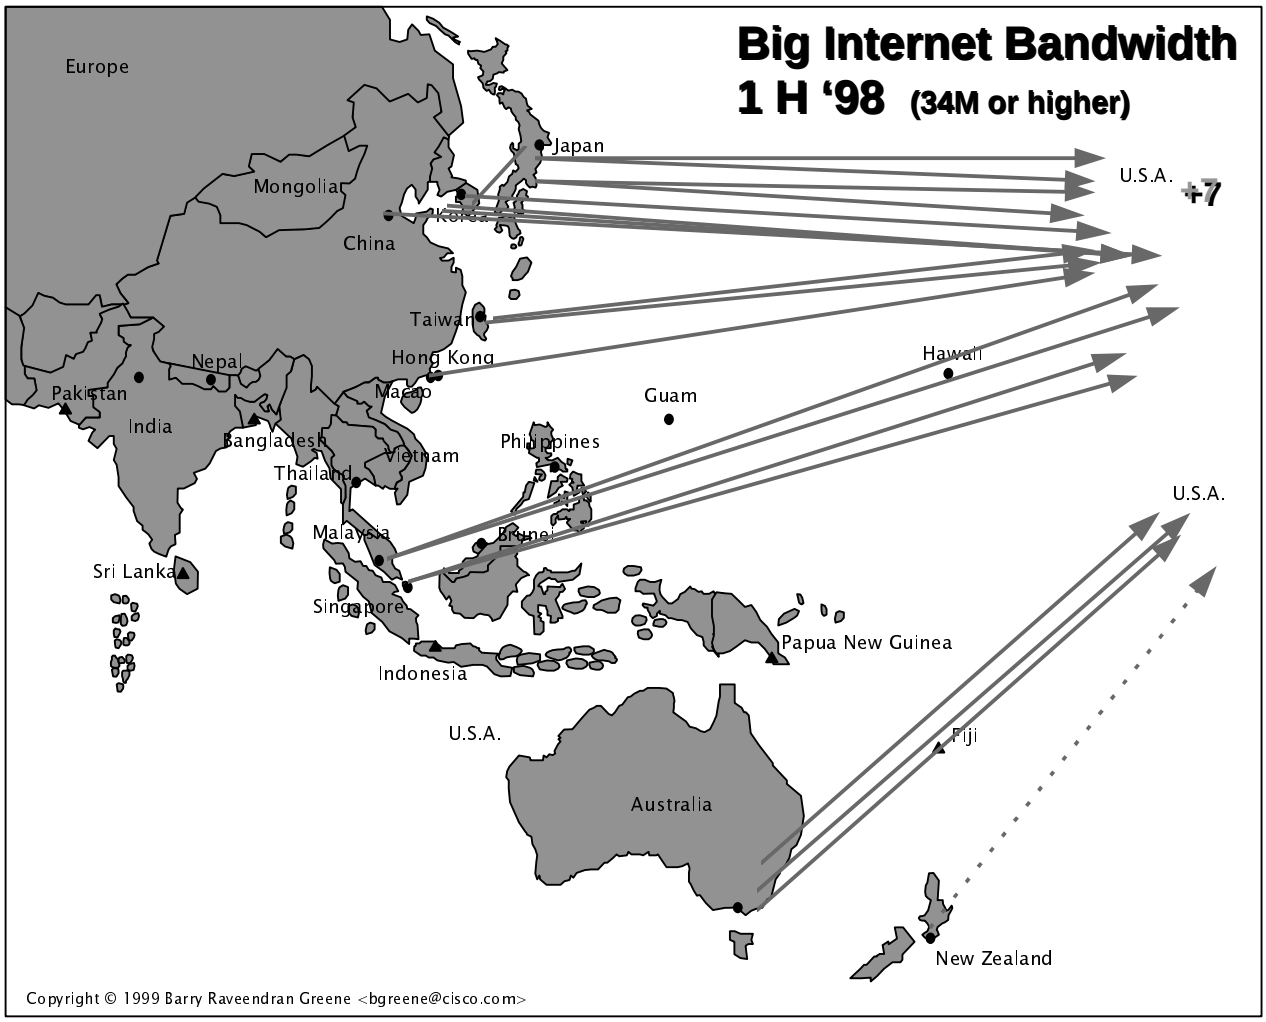
Big Internet Bandwidth
1 H '98 (34M or higher)
Europe
Mongolia
China
Korea
Japan
U.S.A.
+7
Taiwan
Hong Kong
Macao
Nepal
Pakistan
India
Bangladesh
Vietnam
Thailand
Philippines
Guam
Hawaii
U.S.A.
Malaysia
Brunei
Sri Lanka
Singapore
Indonesia
Papua New Guinea
U.S.A.
Fiji
Australia
New Zealand
Copyright © 1999 Barry Raveendran Greene <bgreene@cisco.com>

ANNEX 3

Figure 4: Singapore's PSTN Traffic and Internet Connectivity

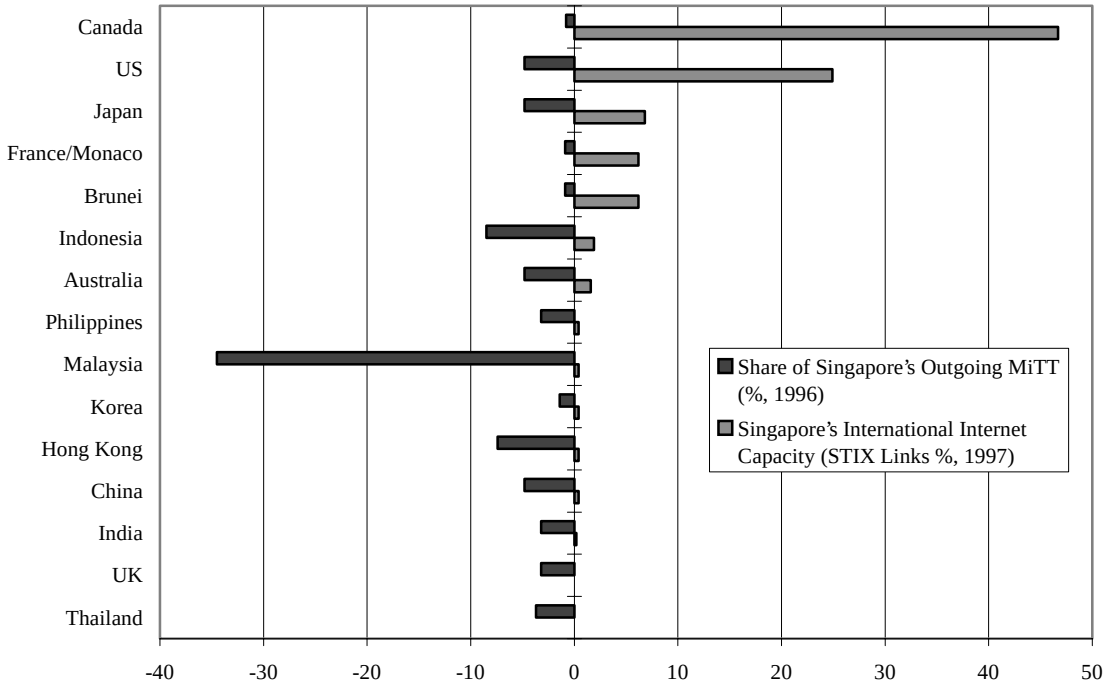

Figure 5: Average round trip time to and from STIX

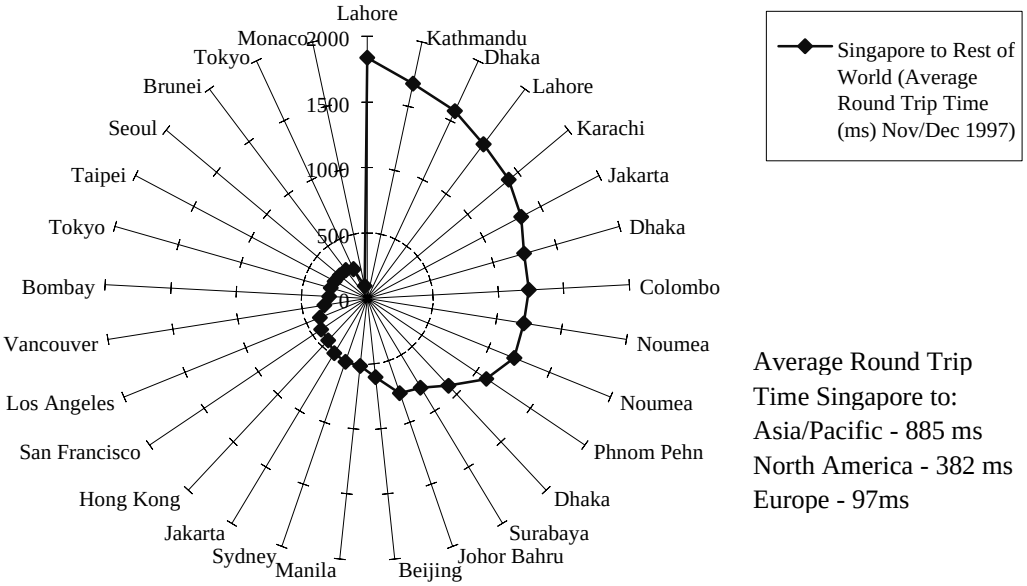

Figure 6: Balance between second level domain registration under TLDs and gTLDs

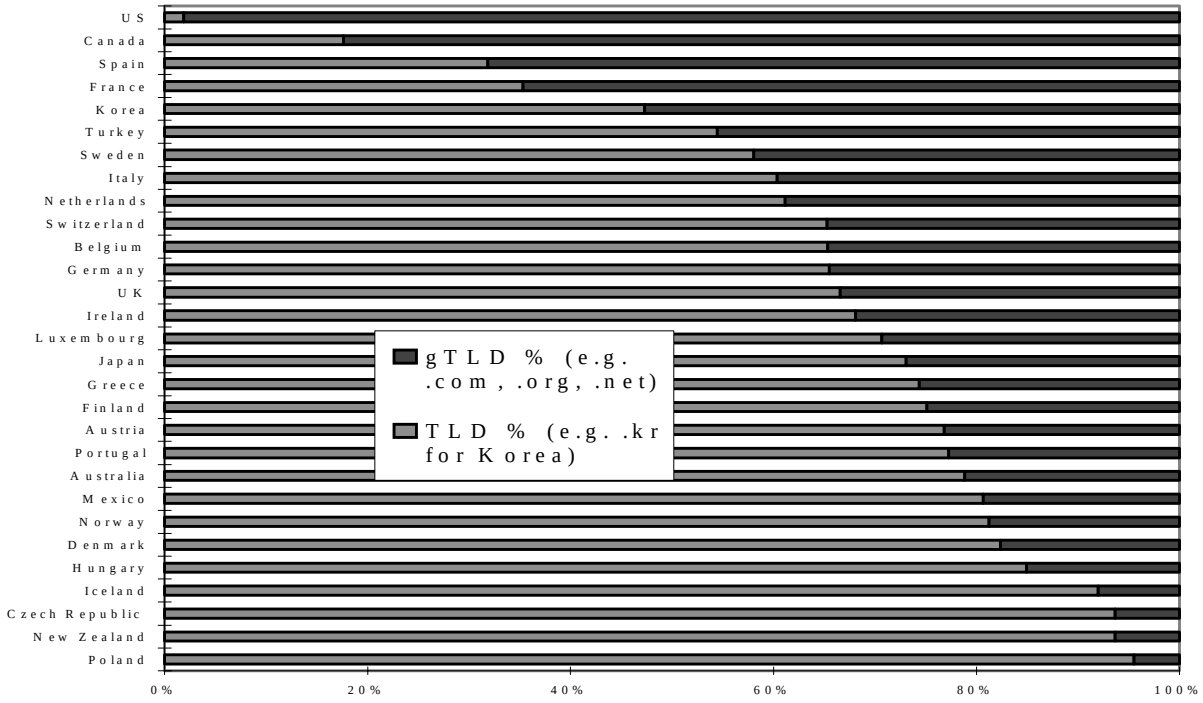

Figure 7: Internet Hosts per 1000 Inhabitants adjusted for gTLD registrations

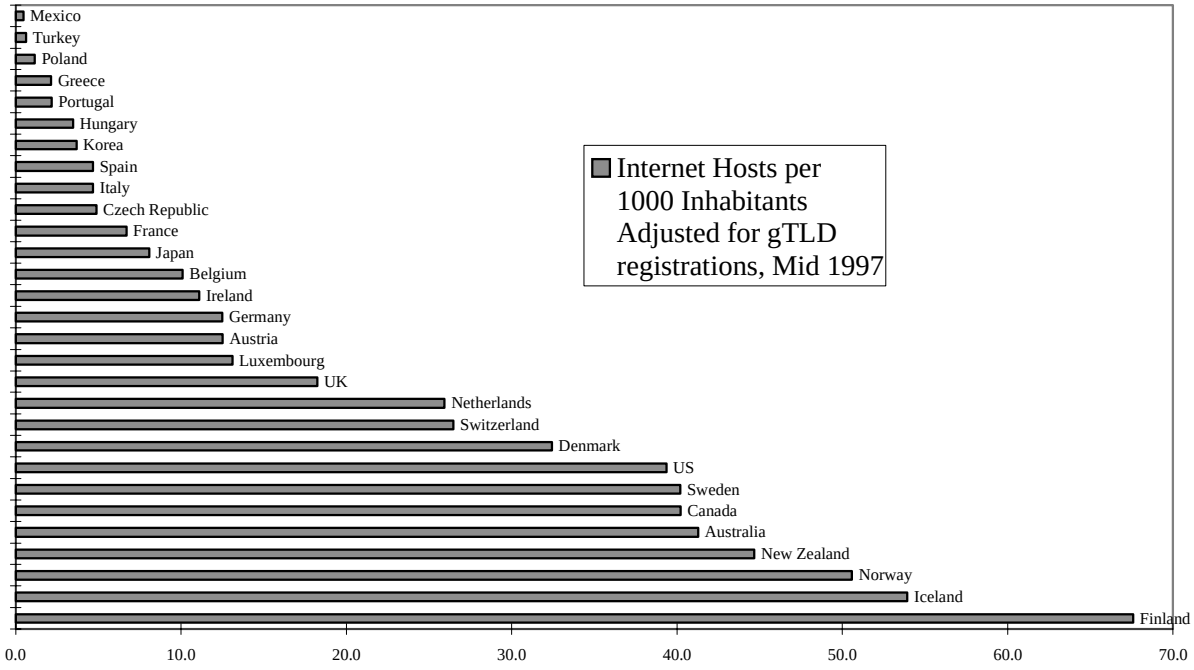

NOTES

1. The DNS is a mapping service which translates domain names and IP addresses back and forth. It is separate from the routing system (except that the routing system is needed to enable kDNS traffic to flow). DNS root servers contain no routing information. Routing information is distributed and there is no root. Each ISP announces its routes to its peers via the Border Gateway Protocol (at exchanges) or statically. There are some route servers located at exchanges, but they are the exception rather than the rule and are not essential to the operation of the routing system.

2. The first Internet exchange point in the United Kingdom was established in London by a small group of IAPs in October 1994.

3. OECD, "Webcasting and Convergence: Policy Implications", OCDE/GD(97)221, 1997.

4. Bruce Haring, "Video on the Net snaps into focus", *USA Today*, 25 November 1997. http://www.usatoday.com/life/cyber/tech/ctb705.htm

5. Steve Rosenbush, "E-mail takes on a video dimension", *USA Today*, 28 November 1997. http://www.usatoday.com/life/cyber/tech/ctb710.htm

6. Deja News, "Press Release: Deja News Launches Major Assault On Spam Offers Internet Users "Spam-Free" Access to Internet Discussion Forums, including Usenet Newsgroups", 8 December 1997. http://emarket.dejanews.com/emarket/about/pr/dnpr_971208.shtml

7. Gene Koprowski, Backbones Wheel and Deal to Keep Net Moving", *Wired*, http://www.wired.com/news/news/technology/story/8642.html

8. See http://www.keynote.com/measures/backbones/backbones.html

9. See http://www.savvis.com/11-11-97-savvis-is-tops.html

10. Keynote Systems, "Top 10 Discoveries about the Internet", 1997. http://www.keynote.com/measures/top10.html

11. "Sun Microsystems, Bellcore Announce Plans to Deliver First Jointly Marketed Product; New Solution, Advanced Domain Name Server, Will Reduce Internet Traffic Bottlenecks", Business Wire, 10 December 1997. See also http://www.soliant.com

12. Ibid.

13. Kevin Thompson, Gregory j. Miller and Rick Wilder, "Wide-Area Internet Traffic Patterns and Characteristics, IEEE Network, November/December 1997 or See http://www.vbns.net/presentations/papers/

14. Andreas Evagora, “World Wide Weight”, *tele.com*, September 1997. http://www.teledot.com/0997/features/tdc0997globe.html

15. Data for monthly average in the late 1997. Refer to monthly averages at: http://www.unidata.ch/info/public/usa-tot.html

16. Ibid.

17. Ibid.

18. Reuters, “Internet Use Doubles in Canada”, 28 November 1997. http://www.yahoo.com/headlines/971128/tech/stories/canada_1.html

19. Public Telecommunication Operators are here defined as entities offering public switched telecommunication services. The term is not used here to indicate ownership status.

20. For a glossary of Internet terms, some of whose definitions are drawn on in this document, see Gordon Cook’s Glossary at: http://www.cookreport.com/

21. See the “The List” at http://thelist.internet.com/

22. Gordon Cook, “Glossary of Internet Terms”, http://www.cookreport.com/

23. Commercial Internet eXchange (CIX) Association, Inc. formed by General Atomic (CERFnet), Performance Systems International, Inc. (PSInet), and UUNET Technologies, Inc. (AlterNet), after NSF lifts restrictions on the commercial use of the Net. See http://info.isoc.org/guest/zakon/Internet/History/HIT.html

24. Originally the organisation, under award from NSF, which provided routing information at each NAP was the Routing Arbiter (RA). The RA provided customised routing information at each NAP that will reflected all bilateral agreements between that NAP's clients. http://www.busn.ucok.edu/desci/binning/nsp_rnp.htm

25. See http://www.pacbell.com/products/business/fastrak/networking/nap/features.html#howitworks

26. Jeff Pelline, “Whole Earth dumps president”, NewsCom, 30 April 1997. http://www.news.com:80/News/Item/0,4,10247,00.html

27. UUNET, “UUNET Details Peering Strategy”, Media Release, 12th May 1997. http://www.us.uu.net/press/peering.html

28. David Holub, “Peering/Interconnection on the Internet as a Telecom Carrier”, Internet Press Release to Various NewsGroups, 5th May 1997.

29. UUNET, op.cit.

30. Jeff Ubois, “Peer Pressure: An Interview with Philip Lawlor”, August 1996. http://www.internetworld.com/1996/08/peer.html

31. Janet Kornblum, “Will WorldCom own the backbone business?”, *NewsCom*, 11 September 1997. http://www.news.com/News/Item/0,4,14171,00.html

32. Kenneth Cukier, “MI-WorldCom faces Internet probe...”, *CommunicationsWeek International*, Issue 195, 24 November 1997. p 1

33. Ibid.

34. Comments by CAIDA Concerning the Federal Communications Commission Review of the Acquisition of MCI Communications Corp. by Worldcom, Inc. April 27, 1998. http://www.caida.org/Caida/fcc-98.html

35. Bell Atlantic, Filing to US Federal Communications Commission CC Docket No 97-211, 5 January 1998. (Appendix A) http://www.fcc.gov/Bureaus/Common_Carrier/Comments/worldcom/

36. See Simply Internet, Filing to US Federal Communications Commission CC Docket No 97-211, 1998. http://www.fcc.gov/Bureaus/Common_Carrier/Comments/worldcom/

37. A plot of the Internet routing table is available at http://www.employees.org:80/~tbates/cidr.plot.html

38. See http://www.waia.asn.au/Issues/Peering/index.html

39. Jermey Scott-Joynt, “Pressure Builds on FCC to Reconsider Internet Access Call Status”, *Totaltele*, http://www.totaltele.com, 14th November 1997.

40. See http://www.mai.net/bostonMXP/general.htm

41. See http://www.nap.inroma.roma.it/

42. See www.gni.fr/PEP/

43. See http://www.mai.net/mxp/MXP.HTML

44. See MaNAP press release at: http://www.manap.net/rel_22_jul.html

45. Ibid.

46. DIX is located at Lyngby just north of Copenhagen . See http://www.uni-c.dk/dix/

47. Interview with Steve Goldstein, Cook Report, January 1995. http://cookreport.com/icm.html.

48. Telstra, Filing with the Federal Communications Commission “In the Matter of International Settlement Rates) IB Docket No. 96-261”, op.cit.

49. APIA’s CFC is at http://www.apia.org/call4.htm

50. Telstra (the Petitioner) initiated action versus the FCC and the United States (Respondents) in the United States Court of Appeals for the District of Columbia Circuit on 1 October 1997.

51. For a discussion of what connectivity is, see Geoff Huston, Elise Gerich and Bernhard Stockman, “Connectivity within the Internet - A commentary”, 1992, http://www.iepg.org/docs/IEPG-connect.html

52. The APNG maps are made available courtesy of Barry Raveendran Greene.

53. Keynote Systems, “Top 10 Discoveries about the Internet”, 1997. http://www.keynote.com/measures/top10.html

54. Ibid.

55. See EuroISPA analysis of costs at http://euro.ispa.org.uk/papers/telecoms1.html.

56. For UUNET's major backbone routes, see http://www.uk.uu.net/network/connectivity/

57. See http://www.us.uu.net/products/ustrans/pricing.shtml

58. See http://www.uk.uu.net/international/faq/

59. Camille Mendler, "Subsea cables make waves: New entrants climb on board..." CWI News Listing for Issue 186, Monday, 2 June 1997.

60. Ibid., This would assume an upfront payment of US $1.3 million spread over the cable's expected 15-year lifespan - US$ 7 222 per month over 180 months.

61. OECD, "Domain Names: Allocation Policies", OCDE/GD(97)207, 1997.

62. See Lanminds "How Does the Domain System Work" http://www.lanminds.com/dns/tld.html

63. See http://rs.internic.net/newsletter/apr97/dnsgraphic.html

64. See http://www.iepg.org/docs/IEPG-ADV-9610.html

65. See IEPG Meeting Notes, June 1996. http://www.iepg.org/mjun96.html

66. Barbara Dooley, Executive Director of CIX, Testimony before the House Subcommittee on Basic Research House Science Committee, Washington, 30 September 1997. http://www.cix.org/congtest.html

67. See http://rs.internic.net/announcements/iif-update.html