

Historic Role Of The Commercial Internet eXchange Router And Its Impact On The Development Of Internet eXchange Points [IXs]¹

Farooq Hussain

Introduction

The world's first commercial Internet exchange point was set up more than a decade ago between March and October 1991. The Commercial Internet eXchange or CIX was established on the principle of open competition on a *level playing field* between large and small operators. The CIX was also a proponent of self-regulation of the Internet industry. Since that time Internet eXchange Points (IXPs) have grown to more than one hundred and fifty worldwide with a variety of interconnection policies and technologies employed. Perhaps now is a good time to reflect on the lessons of the past and to look to the future.

In the early 1990's all of the countries of the Americas were required to connect to an Internet backbone service provider in the United States in order to have their traffic carried to users on other networks. Sometimes these other networks were in the same country. Absurdly inefficient routing carried traffic thousands of kilometers across the oceans to sometimes deliver it a few kilometers from its point of origination! The reason for this was that competing networks were not willing to inter-connect locally. This was the same for other regions of the world such as Europe, Asia, the Middle East and Africa. New IXPs emerged as the rationale for exchanging Internet traffic locally (within a country) or regionally became economically compelling for ISPs that otherwise perceived each other as competitors. Equally important is that these circumstances created policy pressures for changes in the way Internet traffic might have otherwise been regulated in these countries. Though there are specific regional characteristics, local and country specific issues that have generated special conditions for the development of IXPs in the Americas, their overall growth has been generally comparable to that of other regions.

There are already several IXPs in the Americas. In Argentina (CABASE), Brazil (ANSP, Laboratório Nacional de Computação Científica LNCC, Embratel, GT-ER-PIR), Chile (NAP CHILE), Colombia (NAP of Columbia), Panama (NAP Panamericano SA - Panama, InterRED – Panama) and Peru (NAP Peru). Most other countries in the Americas are planning or actively considering establishing IXPs. These Americas IXPs today implement a variety of interconnection policies including multilateral and bilateral peering. There is also an evident interest regulation mandating interconnection between service providers. Whether one of these approaches will emerge as the dominant interconnection policy is still to be determined. It may well be that different approaches will be applicable at local, national, regional and international levels.

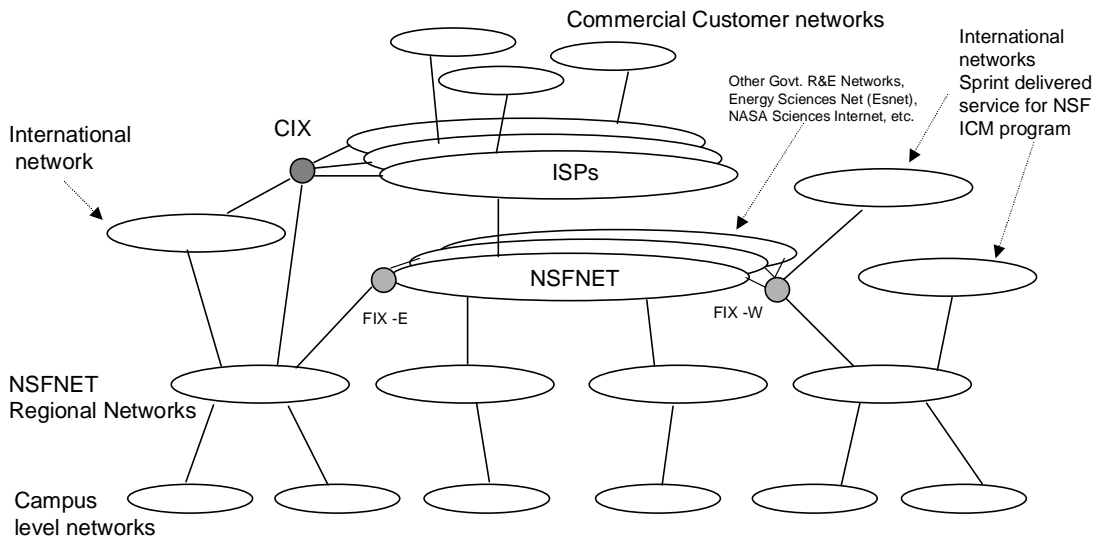
Perhaps, we are now at a watershed in the economic development of the Internet. The impact of this year's collapse of the telecommunications sector in the stock market has forced consideration of significant changes in the outlook for traditional carrier based services. In the future these services will include Internet Protocol as core network infrastructure thus raising new challenges for policy, consideration of regulation and the impact of new technology. One of impacts of these changes may be to redefine the future operation and regulation of IXPs in the global Internet infrastructure.

Background

In early 1991 PSInet, UUNET and CERFNETⁱⁱ formed the basis of what became the Commercial Internet eXchange [CIX], the world's first commercial Internet eXchange Point. Up to then Internet traffic was still largely associated with research and education (R&E) traffic.

At this time these three networks together with ANS were the only national level commercial ISPs in the United States. PSInet hosted the CIX router in its Bay Area POP and soon the group was joined by Sprint, the first of the carriers to deploy a national commercial Internet service. At that time the US government funded National Science Foundation (NSF) Network, NSFNET was the Internet backbone and the traffic that could be carried on it was restricted to non-commercial purposes under an Acceptable Use Policy [AUP]. The NSFNET was operated by a consortium called Merit, created under the auspices of the University of Michigan together with IBM and MCI. Merit, IBM and MCI subsequently created Advanced Network Services [ANS] to operate and manage the NSFNET for the US Government under the terms of the National Science Foundation grant.

Figure 1 Pre NSFNET Transition



FIX Federal Internet exchange (East and West)

NSFNET Operated by Advanced Network Services (ANS) for Merit, Inc. a consortium of IBM, MCI, University of Michigan. Merit was awarded the NSF contract for the development of the NSFNET

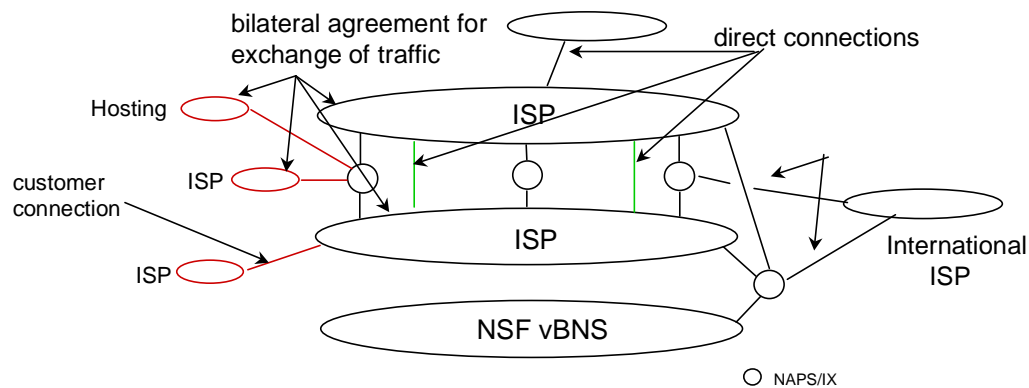
ICM International Connections Manager. NSF program to connect other countries to the NSFNET

CIX Commercial Internet eXchange

The rationale for the CIX was to provide the basis for the open competitive development of commercial Internet in the United States. At that time ANS was seen as having a de facto monopoly because it both operated the NSFNET under an AUP while at the same time using the same network infrastructure to operate a commercial Internet service for profit. ANS faced mounting criticism that it was using government funds to operate a commercial service and using its position as the operator of the NSFNET to restrict the emergence of new Internet Service Providers that could compete directly against it. ANS had only reluctantly agreed to peer with PSInet, UUNET, CERFNET and Sprint and was widely accused of operating restrictive practices against smaller ISPs through its interpretation of the requirements of the NSFNET AUP.

While debate raged on the appropriateness of the NSFNET contractual arrangements with ANS and the US Congress considered a transition of the NSFNET to commercial service, the CIX grew rapidly in strength and popularity with start-up ISPs. It championed and shepherded the emergence of a burgeoning ISP industry. The IXP that formed the rationale for its existence provided a critical service that enabled new entrants who would very likely have otherwise been unable to effectively offer commercial Internet service.

Figure 2 Post NSFNET Transition



- NSF VBNS is policy restricted and is evolving within the context of the Internet II and NGL initiatives. Access will be broadened to a wider set of US and international institutions.
- Commercial ISPs exchange traffic at Network Access Points/Interexchange points or by direct connections. ISPs enter into bi-lateral agreements
- ISPs that do not have backbone network infrastructure required for obtaining an agreement with a Level 1 ISP may obtain service by purchasing dedicated access from one or more Level/Tier 1 ISPs

The Pursuit of Peering Agreements: Multi-lateral and Bi-lateral Peering

All members of the CIX big and small agreed to interconnect to each other for the purposes of exchanging routes, and to carry member's traffic whose destination was on their network. The arrangement was defined as a multi-lateral peering agreement - all members of the CIX agreed to exchange traffic with each other. In doing so the group also was able to overcome the restrictions that ANS was selectively applying by refusing to peer with new networks. Since ANS peered with some networks that were CIX members it ended up receiving the traffic from all of them.

As the time approached to shut down the NSFNET and transition to fully commercial Internet services in the United States new concepts of Interconnection were being explored. The CIX was joined in 1992/3 by a private peering arrangement between PSI, UUNET, Sprint, and ANS in Washington DC. This used a Metropolitan Area Ethernet [MAE] to provide the interconnection between these networks. The MAE was a commercial service offered by Metropolitan Fiber Systems (MFS). The purpose in establishing the MAE was in part to explore new architectures to replace and provide alternatives to the Federal Internet eXchanges that operated out of College

Park, Maryland [FIX-East] and NASA's Ames Center in the San Francisco Bay Area [FIX-West]. Following the shutdown of the NSFNET it was envisaged that interconnection between the new commercial ISPs would be managed through Network Access Points or NAPs. MAE-East as it became known [since a comparable MAE was also envisioned for the west coast], operated on bi-lateral agreements between the networks connected to it.

As the shut down of the NSFNET and the transition to commercial service drew close in 1994 the NSF had mandated that new commercial Internet service providers interconnect at five designated NSF Network Access Points [NAPs] across the country. These NAPs required ISPs to enter into Bi-Lateral interconnection [peering] agreements with each other. The interconnecting networks defined the structure of the bi-lateral agreements. They followed the some basic principles which have defined bi-lateral peering agreements ever since. First, interconnecting ISPs agreed not to charge each other for the exchange of traffic because at the time the interconnecting ISPs believed that the traffic to be exchanged was in approximate balance and that and that the costs to each party for carrying the traffic were approximately equivalent. This became known as the no settlement clause. Secondly, each ISP agreed not to permit its network to be used for transit, that is traffic exchanged between bi-lateral peers terminated on their networks *not* on terminations of a third party network. For the purposes of routing efficiency and cross country capacity costs interconnecting ISPs agreed that traffic to be exchanged would be routed via the shortest exit path. This has often been called Hot Potato routing. Lastly, interconnecting networks agreed to maintain the highest level of privacy and security for their customers by not permitting any form of monitoring at the point of interconnection. This was familiarly known as the no-wire tapping clause.

How were Internet Peering Agreements First Negotiated?

In the year leading up to and following the shut down of the NSFNET, ISPs actively negotiated bi-lateral interconnection arrangements based on these four principles: no settlement for the exchange of traffic; no third party transit; shortest exit routing; and no wire-tapping. Many peering relationships were made on the fly between operations groups and were also undocumented. However, the main national level backbone networks did maintain a requirement for documented and clearly defined agreements thus setting a de facto standard in terms of perceived requirements for peering between networks. In the United States at the this time those requirements were formulated as follows:

A peer network needed to have:

- a minimum capacity of 45Mbps in a full mesh cross-country infrastructure
- to be present at a minimum of three NAPs
- to present a 24x7 network operations capability

The above criteria made sense for the national backbone networks but smaller start-up ISPs regarded the criteria as unfair since they could not afford presence at three or more NAPs and did not necessarily have a full 45Mbps network capacity. The seeds of future disputes between the national backbone backbone networks that were to replace the NSFNET and the smaller regional and local networks were sown here. The NSF had instructed national level networks to negotiate peering agreements with each other making it clear that NAPs would not be implementing interconnection policies rather they would only provide the interconnection environment. The MAE-East operated in this way, any network could purchase MAE service from MFS but they needed to independently and directly negotiate peering agreements with other networks attached.

The primary motivation for peering was then and still is for cost saving. At that time the NSFNET structure was comprised of a hierarchical architecture with the NSFNET itself being the national level backbone. Directly connected to the NSFNET were large regional networks. These regional networks were partly subsidized by the NSF to purchase access to commercial backbone networks following the NSFNET's shutdown. The institutions and enterprises that formed the majority of the connections for the regional networks were also required to purchase commercial

connectivity and could in theory select from a broad range of potential providers. That is they could stay with their existing provider or potentially move to one of the new commercial national level backbones.

Essentially, this situation placed the large regional networks in competition with the commercial national level backbones. Thus BBN having acquired NEARNET one of the large regionals and a customer network of MCI soon acquired SURANET and BARRNET and moved from being a paying customer of MCI to a peer network. Other start up networks were equally anxious to reduce their costs of transit by gaining peering relationships with national backbones. These conditions naturally created tremendous pressure on the national backbone networks from the new start-ups for peering relationships. For the national level networks the initial rationale of peering relationships between a few competing national backbones that would replace the NSFNET was instead viewed as an environment which would soon develop to more than twenty national level backbone networks that more or less complied with the basic peering criteria described above.

Although no more than five national backbones carried more than two thirds of the traffic it was not uncommon for ISPs to end up with between twenty to fifty peering agreements. Some of these included variations on the basic structure whereby some of the costs were distributed asymmetrically in favor of the larger network with the greater amount of traffic.

By 1995 when WorldCom bid to acquire MCI there were more than twenty national level backbones in the US listed in the Boardwatch Directory. There were huge strains between the national level backbones resisting peering agreements with new entrants and the perception of new entrants that such resistance was restrictive. Additionally, some national level backbones perceived that the combination of UUNET with internetMCI would create a dominant player in terms of global Internet traffic carried and would thus be anti-competitive. This view was accepted by the US Department of Justice and the European Commission and WorldCom was required to divest internetMCI, selling it to Cable & Wireless as a condition of the planned acquisition.

Amongst the many complex issues that arose during the Department of Justice and European Community inquiries into the WorldCom acquisition of MCI was the importance of the policy neutrality of the IXPs. Though the CIX was formed as a policy based IXP and the Federal Internet exchanges had policy regulated by the US government agency networks the MAE-East environment and the formulation of the NSF NAPs were policy neutral. The NAP operator played no policy role; rather they emphasized a policy neutrality and presented their facilities as carrier or service provider neutral for the purposes of the physical location and means of the interconnection between ISPS. At the time of the WorldCom acquisition of MCI the issue of WorldCom's ownership of the MAEs {East and West by this time} raised concerns regarding the perceived maintenance of neutrality and whether regulation might be required to maintain the neutrality of IXPs. Within two years of the NSFNET shutdown and prior to the WorldCom MCI merger, the NSF had decided that the transition to commercial service had been successful and ended its formal support of NAPs focusing its efforts exclusively on the very High Speed Backbone Network Service [vBNS] and other projects for next generation applications including assistance to Internet 2.

What were the Critical Factors in the Creation of the NSF NAPs and why did the NSF withdraw from continued support for them?

In formulating the requirement for NAPs in the planning for the NSFNET shutdown the principle of policy neutrality and geography were formalized though they had previously existed. The need for the interconnection facilities to be present on either coast of the US and in the center was following the pattern of the FIXes East and West and Washington DC also had the MAE. Additionally, The NSF never specified that the NAPs be geographically specific to any City for example the California NAP and the New York NAPs referred to regional geography rather than cities. Following the NSFNET shutdown the main concentration of IXPs based on aggregate

volume of traffic has stayed in Washington DC, the San Francisco Bay Area, and Chicago. Other facilities have been established in Seattle, Dallas, Los Angeles, Miami and elsewhere over the years. In the US nearly all the IXPs have operated as carrier neutral facilities supporting bi-lateral interconnection agreements between service providers. While the WorldCom MCI merger raised the question of government regulation of the IXPs, regulation has generally been resisted both by the government agencies and by the smaller service providers who have been amongst the strongest critics of the peering agreement behavior of the large service providers.

The NSF awarded the NAPs to Sprint, PAC Bell, Ameritech and MFSⁱⁱⁱ. But these have subsequently lost their position as key interconnection locations to the MAEs (East and West) and the PAIX. Additionally, the major carrier based ISPs in the US have in recent years established direct interconnection between each other reducing or eliminating their need for the use of 'public' peering points. Public peering refers here to the use of a 'carrier neutral' switching fabric provided by the IXP. Certain IXPs today such as PAIX, Equinix and Xchage Point (in Europe), are amongst those that provide both the means for carrier private interconnection and the use of an IXP provided switching fabric. The largest ISPs that have pursued direct interconnection between each other have mainly done so to improve the overall quality of service for their respective customers. The shared environment of an IXP provided switching fabric was too frequently saturated resulting in traffic congestion at these points. But the move towards direct interconnections was conducted in parallel with other changes in peering policy that has made the requirements for peering significantly more restrictive for the smaller ISPs and any new entrant. Amongst the most notable of these has been the requirement for a peer to present traffic in approximate balance (within ten to fifteen per cent of parity). While this requirement was not new in principle it strictly impacted all but the largest service providers. This is the culmination perhaps of the various peering disagreements between service providers.

The decommissioning of the NSFNET in the US had substantial domestic and international implications. The domestic issues arose out of the desire of smaller regional networks to compete with much larger providers with national level infrastructures and therefore required interconnection [peering] agreements with them. This first became apparent with the consolidation of three former NSF regional networks SURANET, BARRNET and NEARNET by BBN. The consolidation by BBN was notable in that it created a very substantial competitor to the then 'incumbent' national level ISPs. But for some time prior to and certainly long after this, start up networks began often as customers and evolved to the point where they were requesting interconnection/peering relationships with their 'upstream' ISPs. The relationship between ISPs and hosting and collocation providers created similar desires for interconnection agreements between the hosting companies and their ISPs. The first notable case was perhaps between Exodus and BBN [by then acquired by GTE] in summer of 1998.

Peering Wars

Thus the present situation in the United States where the largest providers are attempting to induce service providers to purchase transit connectivity rather than pursue peering relationships is the outcome of a trail of peering disputes over the years. The first major peering dispute arose around the decision of UUNET in 1994-5 to de-peer a large number of ISPs. From UUNET's perspective the rationale was that a large number of its peers did not meet the basic peering criteria that was in common use by other service providers. Though UUNET attempted to provide alternative arrangements to those networks to continue to carry and deliver traffic they now required payment for carrying the traffic. The dispute though acrimonious in its early stages established the basis for some form of payment for peering where there are traffic asymmetries or other factors to exclude a no-settlement arrangement.

Other peering disputes included WorldCom and Level 3 where WorldCom denied Level 3's requests for peering. Level 3 managed to work around this by acquiring a small ISP, Geonet, which maintained peering agreements with large ISPs established at the time of the NSFNET

shutdown. Level 3 became one of many larger players entering the US Internet backbone business that followed the acquisition path to obtaining peering.

GTE/BBN and the former hosting company Exodus (the company recently applied for Chapter 11 protection) engaged in a substantial peering dispute in 1997, which highlighted the difference between providers of content and providers of network infrastructure. The situation between Exodus and GTE was blurred by the fact that Exodus also operated significant network infrastructure. Ultimately, GTE's position in common with the views of other large service providers prevailed. GTE wished to be paid for carrying Exodus's traffic and denied realizing any direct commercial benefit for the direct access that its users gained to Exodus hosted content.

The peering disputes were by no means exclusively domestic. Telstra was perhaps the most vocal critic of the policy of US ISPs to charge international half circuit access costs and Internet access port charges, and steadfastly decline to move toward a peering relationship. None of the international ISPs were prepared to meet the domestic US criteria for peering. Telstra in particular began a long campaign to seek a settlement-based exchange of traffic broadly supported by a number of Asia Pacific countries. It carried the issue to the ITU and to APEC's Telecommunications group where it continues to be an area of reporting and investigation.

Bi-lateral Interconnection [Peering] Agreements between the national level network service providers having been established during the period immediately preceding and just following the NSFNET transition. But for emerging commercial service providers the CIX was sometimes the only way to route traffic to terminations off the major networks in circumstance where the emerging service provider did not want to pay for the connection to that network. It should be noted that many networks successfully operated by maintaining connections to upstream ISPs that they purchased as a customer. Such networks tended to be local or at most regional and as they grew to operate at a national level they needed to have interconnection agreements to improve operating costs and diminish the prospect that they might have supply dependencies from a competitor network. To summarize, a start up service provider or hosting company was able to purchase connections from upstream providers but as its business grew the cost of purchasing bandwidth outgrew the costs of operating under interconnection agreements in which providers pick up their own costs of presence at IXPs and the associated bandwidth costs.

In Europe the cost of inter-European bandwidth was comparable to bandwidth costs between Europe and the United States. But these Europe - US costs were still too expensive for the majority of start-ups who sought to obtain connection to the Internet infrastructure through an upstream provider or through purchasing consortium arrangements such as the E-Bone and transit and interconnection associations such as the LINX. In the Americas, Asia, the Middle East and Africa comparable cost of bandwidth issues existed and the traffic from these regions was in the majority, flowing between the region and the United States rather than as inter-regional traffic via the United States. Regional initiatives have been made in the past to build high capacity connectivity 'rings' interconnecting the countries of the region as well as the emergence of inter-country if not true inter-regional level interconnection. Such initiatives are perhaps worth revisiting in the context of the next generation of network infrastructure and types of services - to include voice and multimedia for example that would be supported.

Peering and Transit as International Interconnection and Settlement Mechanisms

At the start of international Internet connectivity to the NSFNET the research institutions of countries trying to connect to the NSFNET needed to fund their costs of connectivity. In practice this meant that the US Government paid for the United States half-circuit and the half circuit cost for the connecting country needed to be paid for by their government or some other source. Simply the governments of the United States and the connecting country put up the costs of the circuit other costs were relatively small by comparison and were generally met by the connecting institutions. In some cases the US government sponsored the whole circuit where there were opportunities for the NSF to fund the circuit to support development of the Internet in countries

would not otherwise be able to support the costs of the connection. When the NSFNET shut down commercial service providers in the US had no rationale by which they could underwrite the cost of international capacity to a large number of countries thus requiring countries to meet the full international circuit cost and Internet access port charge to reach the US Internet infrastructure.

The circumstances of the US environment have led to the dominance of the Internet transit capacity market by fewer than ten networks. That is that more than ninety per cent of all Internet traffic is carried by these networks. Additionally, the majority of the global Internet traffic has characteristic pattern of flowing into the US from other countries with the US being a hub for other countries to reach each other. Although there have been many efforts to argue that this pattern is changing there is little proof that it is doing so in any significant way. Being packet-based IP networks have never been technically suited to reciprocal settlement models particularly since it is not feasible to account for the source and destination of traffic for charging purposes, as is the case for telephony.

Since the mid-1990's there has been a substantial change in the business of international transit service – that is Internet connectivity between the rest of the world and the United States. The majority of service providers from these countries were international carriers who were also operating ISPs. They negotiated with US service providers for connectivity based on the combined cost of the international half-circuit and the Internet access port. This simple business relationship became stressed as the demand for international Internet capacity from some of these countries exploded. US providers for the most part could not reconcile peering relationships with the main carriers of international Internet traffic based on settlement models of the balance of traffic. However, the international carriers themselves were able to take advantage of alternatives such as wholly owned international capacity to the US and the establishment of interconnection arrangement within the US through IXPs. Of course, in notable instances there have been mergers and acquisitions such as Cable & Wireless's acquisition of internet MCI, NTT's acquisition of Verio, and the Qwest KPN Merger. There are a number of service providers operating international trans-oceanic capacity between Asia Pacific and the US and Europe and the US. Additionally, perhaps most importantly there is available capacity to meet demand and the pricing is likely to remain highly competitive. Whereas in the past US service providers were able to sell transit services to the carriers of other countries, the transit providers today are operating globally and have reduced capacity and operating costs. However, the market for international transit generated regionally in Asia, Europe and elsewhere still requires the purchase of international transit and these services are now more broadly available through purchasing consortiums and operating efficiencies have been increased by the introduction of IXPs for the exchange of local traffic separately from international.

The developments here have generated more pressures for interconnection/peering agreements between carriers. However, competitive advantages are held by existing service providers whose networks have become large 'sinks' for traffic. Thus these networks are also important peers for other service providers. Some of the existing service providers are seen to maintain standards of interconnection that severely limit the number of potential peers. These standards include for example the requirement to maintain 'traffic balances' and the capacity and number of nodes and interconnection points of the network. While these issues are definitely of concern they are perhaps passing problems that will be overcome by an open competitive market and rapidly advancing technology. The extent to which government regulation or intervention may be required to maintain a fair competitive market is still open to question.

Overview of Internet eXchange Points Today

Though Internet eXchange points formed a crucial role in the early Internet and at its transition to commercial service in the United States they are now only one of the means by which service providers may establish interconnection. The original IXPs have operated under broad and sometimes very different criteria. The CIX for example, was formed with a framework that

required policy compliance from its members. By contrast the NAPs though requiring attached networks to maintain interconnection policies were themselves neutral regarding the policies providing only an interconnection environment for the service providers.

None of the original NAPs were particularly successful as independent stand-alone businesses. In part they were never able to successfully accommodate the needs of large and small service providers. The recent consolidation in the US Internet service provider industry essentially creating an environment of less than twenty very large service providers may both define a new future for IXPs but also constrain it to those physical locations where a substantial number of the larger providers are present together. It has always been the case that the successful IXP would be the one where the larger transit service providers were willing to be present and establish peering agreements. This is likely to hold true for the future. The criteria that will draw larger service providers together at common IXP locations in the future is unlikely to be driven as it was at the shutdown of the NSFNET by government policy. Rather it is likely to be driven by collective interests of the larger providers to exchange traffic with each other and to find new customers for transit services. These customers may of course include each other where route specific capacity allows improved cost performance for one provider over another. The global distribution of Internet eXchange Points has not changed very much in the past decade in that the IXPs with highest density of traffic and the routes with the highest volume of traffic have not changed. While it has been important to have new exchange points emerge these have established local and regional efficiencies which though important in themselves have not much altered the global picture.

There will continue to be an interest in finding new charging arrangements and to explore the means by which voice and multimedia services may be handled as the traditional telephony carriers increasingly deploy network infrastructures based on the Internet Protocol. It is reasonable to anticipate that such changes will significantly impact traditional models of settlement for the exchange of traffic between service/providers. To the extent that in some countries at least these service providers may only be at an early stage of privatization or de-regulation may be a source of further complication for how local, regional and international Interconnection policies may be implemented.

What is the CIX's Impact on the Development of IX's ?

The CIX was established by a small group of service providers primarily to challenge the specific perceived anti-competitive circumstances that existed in the United States in 1991 as described above. The multi-lateral interconnection arrangement of the CIX continues to be of value in circumstances where smaller providers see benefit in aggregating outbound traffic prior to presenting it to another 'peer'. For example, regional ISPs may wish to establish an IXP to exchange local traffic between each other and present an aggregate of their traffic to all other destinations to another ISP[s] and to collectively seek a peering relationship with that ISP[s]. But for the most part multi-lateral interconnection arrangements have been superseded by new technology and quality of interconnection environments that permit direct interconnection between pairs of networks even permitting for the cost and operational management of asymmetric traffic volumes between them.

In broader policy terms the CIX as an organization sought to support all interconnection frameworks that facilitate interconnection between Internet Service Providers on the basis of a 'level playing field' providing for fair competition between all providers. However, over time, the membership of the CIX very largely outgrew its dependence on the organization for the provision of services associated with interconnection. This was also influenced by the emergence of commercial interconnection points as well as the preference for bi-lateral and direct interconnection/peering arrangements between providers.

In December 2001 the CIX Router was shut down. It was fitting that the last peering session that was taken down was that of UUNET one of the three networks that helped get it started. The CIX

Router that had been located at the Palo Alto Internet exchange (PAIX) for a decade was shipped to be stored for conservation – for it is a piece of the Internet's history, and an important one. Later, in the Spring of 2002 the CIX itself closed.

Rationale for the development of IX's Internationally

Since the shutdown of the NSFNET, the rationale for a regional IX in the Americas, Europe, Asia Pacific and other regions of the world has been motivated by a complex mix of economic, political and technical issues. In the first instance, connectivity to the Internet was achieved through government-subsidized payment for an international private line between the institution of the connecting country and the NSFNET. Following the transition to commercial service, operators in countries outside the United States found it difficult to accept that they needed to absorb the cost of their own half-circuit and purchase the US half-circuit of the international private line, in addition to paying the costs associated with the Internet port. Many entrepreneurs in the region felt that these 'international transit' costs of reaching the Internet unfairly hampered their ability to develop Internet services within their respective countries and even the region with respect to the United States. From the perspective of US service providers the issue was seen from a different perspective. The business models for US ISPs were based on 'free' local dial access with highly competitive Internet access charges. The cost of deploying national level backbone infrastructure in the US was also expensive. No specific means was available to determine the specific financial value for 'international access' for users in the US. Internet could not be divided into local, long distance, and international services and the business models that were attractive at the time for connectionless data network services found nothing to borrow or adopt from traditional telephony.

One of the consequences of bandwidth pricing between countries in the Americas and between these countries and the United States is that it has been significantly cheaper to allow inter Americas traffic to be exchanged at IXPs in the United States. Of course, this characteristic is common to the circumstances of other regions such as Europe, Asia Pacific, the Middle East and Africa. Recently, there has been renewed interest in the need for Internet traffic in Americas to be exchanged within the region and not via the USA. At present some x% of Americas traffic is exchanged in Miami, New York. For Asia Pacific some 70% of the international traffic converges on Palo Alto in CA at the PAIX. Additionally, most of the traffic originating from Asian countries is for terminations in the United States. While there has been some significant success in establishing IXPs within some Asian countries and also in inter-country exchange of traffic these efforts have been very small by comparison to the continued growth of Asia – US traffic volumes.

Regional IX's in Asia

Future initiatives to establish regional IXPs should take these factors and the lessons of past initiatives to establish inter-regional traffic exchanges in Asia into consideration. The economics of trans-oceanic bandwidth capacity and the differentials between the Asia – US compared with the inter-Asian pricing will continue to be a challenge. National carriers have to address not only issues of inter-Asian bandwidth pricing but also national practices so as to avoid a bottle neck in the local loop, often arising from a reluctance to activate latent bandwidth based on narrow business considerations. The improved monitoring of network traffic flows and observation of routing practices might also help in driving potential inter-regional connectivity towards better cost-efficiencies and performance for network operators and users alike. The introduction of national and regional level initiatives for the coordination of the planning and management of shared network infrastructure would help sustain service provider industry lead efforts to establish improved inter-regional Internet traffic exchanges. The potential resistance of some of the network service providers to interconnect can be better managed through communication and interaction via such initiatives. Although many may see the importance of inter regional connectivity initiatives the implications of current charging practices for peering and transit of Internet traffic should not be underestimated. If these practices continue to progress along the lines of making it harder for the new and smaller service providers to enter into competitive

commercial service offerings then the industry, as a whole will have badly served itself and its users.

What is the future for IXPs?

The technology on which IXPs have been based have generally constrained the potential opportunity for flexible interconnection policies. These constraints may be lifted as new routing technology becomes viable^{iv} IXPs began as direct connections to shared routers using Ethernet and at the time of the NSFNET shutdown much emphasis was directed at ATM based switching fabric which never really came through. Private direct cross connections, gigabit Ethernet, SONET, and most recently the potential for optical BGP routing all present interesting paths forward. None of these of course would specifically require an IXP environment to be implemented though such an environment may be a convenient way for service providers to pursue such technologies of interconnection. However it is with source-destination routing models such as may become feasible with the approaches being explored in OBGIP that truly revolutionary flexibility in interconnection may become possible. Such flexibility may help realize the vision of a level playing field between large and small service providers in an open competitive and global market place.

It appears that the role of Interconnection Points continues but they serve different purposes for service providers depending on their stage of development. Start-up service providers see significant attraction in potentially meeting and exchanging traffic with peer networks at interconnection points. The networks that have become the larger sinks of traffic however, dislike the costs associated with the continued increase in bandwidth requirements that occurs with their presence at IXPs. Much criticism was directed by smaller providers against these networks for 'throttling' bandwidth at IX's such as MAE-East resulting in poor performance for the users of smaller service providers.

The larger service providers saw incentives to directly interconnect their networks by-passing the need to increase capacity at IXP's with resulting improvements in performance for their users. Smaller service providers faced difficulties in obtaining direct interconnection arrangements with larger providers based on relative volume of traffic exchanged and other criteria that made it difficult for them to comply.

In many circumstances IXPs have been planned and launched in order to provide smaller service providers with a better competitive positioning against their much larger competitors. Looking to the future, IXPs are only one of a number of options that smaller networks have in terms of positioning to compete effectively against established large incumbents. In this regard routing practices and future development of routing technology are amongst the most important potential areas to observe. The future relevance of IXPs will depend on national and regional geographic and economic factors which will make them more or less important depending on the country and region of the world in which they are deployed. In the US it is interesting to note the overwhelming concentration of interconnection at MAE-East has given way to a more even distribution across the country and via direct interconnection between networks. There has not been any significant growth in regional US IXPs and the NAP of the Americas in Miami is perhaps one of the rare new IXPs to emerge with any significance. The continued perceived importance of a neutral interconnection location is demonstrated in the number of networks locating at the PAIX and those that are doing so at Equinix locations also. Overall though, there appears to be a slight decline and consolidation of IXPs in the US. Although the growth and contraction of IXPs is perhaps a characteristic of the development the number of networks, it will also be impacted by the consolidation of the recent past resulting in a few very large networks.

The means of interconnection and the possible efficiencies of operation including routing issues have many ways to be successfully managed between providers and IXPs are only one means to this end. Thus it is perhaps better to separate the best principles and practice of interconnection and traffic exchange between providers as national, regional, and international codes of conduct

than to have groups of service providers establish these practices in isolation from the global perspective. While there have been significant areas of difficulty in managing interconnection relationships for the most part this has been very successfully self-regulated in the United States. There are many new challenges ahead both for industry policy and the technical means of implementing interconnection between service providers as traditional carriers deploy core network infrastructure for voice, multimedia and data services based on the Internet Protocol.

About the Author

Farooq Hussain is a Principal and Analyst with Network Conceptions LLC a Washington DC based telecommunications and Internet consulting services firm. Hussain is also publisher and co-editor of Peering Business News that covers interconnection developments and provides consulting services internationally. Previously, he was President of mediagate Inc., from April 1999 until April 2001. mediagate Inc. pioneered digital signal processing (DSP) technology for Internet-based universal communication applications for next-generation Internet service providers. Prior to joining mediagate, he was executive vice president, Global Multimedia and Data Services for AGIS where he was responsible for the development and management of global multimedia products and services such as voice, video and universal messaging. Hussain previously served as director, Internet Strategy and Planning for MCI. He played a key role in defining and directing the implementation of the company's Internet strategy, including new services and policies, and developing strategic business alliances. Prior to joining MCI, Hussain held several positions at Sprint. He was a National Science Foundation's (NSF) Principal Investigator for the NSFNET transition, managing the design and deployment of the New York NAP (Network Access Point). Hussain supported the development of network products associated with high-performance computing and communications at Sprint and also had product planning responsibilities for the company's commercial Internet service, SprintLink.

He served on the Board of the Federation of American Research Networks from 1996-98 and from 1997-2001 as Treasurer and member of the Board of the Commercial Internet eXchange [CIX] a US based organization representing the ISP industry. He is a consultant to several Internet technology and service provider companies and holds a Doctor of Philosophy degree in science and international affairs from King's College, University of London.

Notes

ⁱ Earlier versions of this paper were presented at the First Conference of Latin American NAPs, Office of the Colombian Minister of Communications, Colombian Chamber of Information Technology and Telecommunications (CCIT). Cartagena de Indias, Colombia October 22, 23, and 24, 2001 and to the INFOCOMMS Forum @ Communicasia 2001, Singapore, June 19-22, 2001.

ⁱⁱ CERFNET [no relationship to Vint Cerf] was a not for profit ISP founded and owned by General Atomic who managed the University of California, San Diego Super Computer Center. CERFNET was the first not for profit to transition to commercial service offering. It was subsequently acquired by TCG [a cable operator] which was itself

ⁱⁱⁱ Sprint still operates the NAP, PAC Bell and Ameritech are now both owned by SBC and WorldCom is the owner operator of the MAEs.

^{iv} The notable work in this area is being undertaken by CANARIE in Canada