



Toward Efficiencies in Canadian Internet Traffic Exchange

Bill Woodcock, Research Director, Packet Clearing House
Benjamin Edelman, Associate Professor, Harvard Business School
September 12, 2012

1. Summary

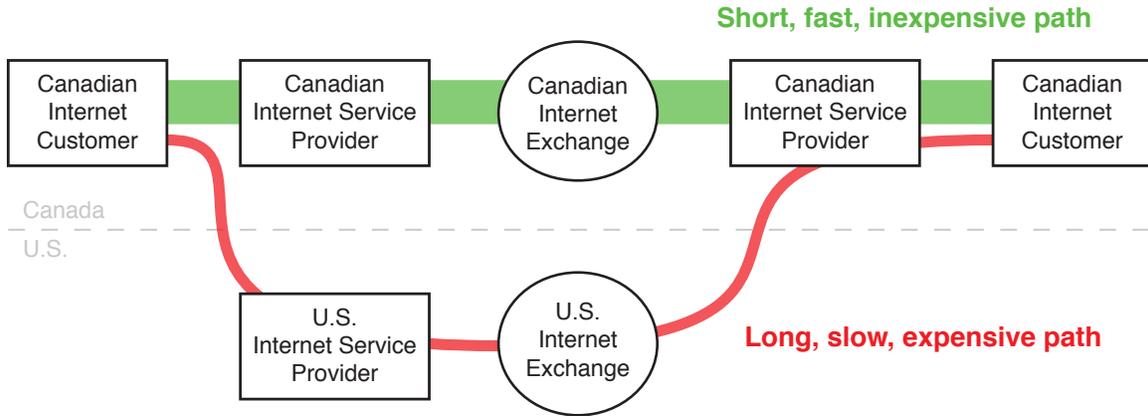
Canadian Internet access is heavily and unnecessarily dependent upon foreign infrastructure, especially U.S. infrastructure. This dependence imposes significant burdens upon Canadian Internet users:

- Service prices are higher than would be the case if Canadian networks more densely interconnected domestically.
- Network speed is slower than would be the case if Canadian networks more densely interconnected domestically.
- When data en route from one Canadian network to another passes through other countries, the data is subject to examination by companies and government authorities in those countries. Canadian data-protection laws are understood not to protect data as it passes through other countries.

Despite these challenges, experience in other countries shows a clear way forward. By establishing more Internet exchange points (IXPs) within Canada, Canada can reduce the portion of network traffic that travels from one point in Canada, through the United States or other nations, and back to another point in Canada. The key benefits:

- By reducing reliance on costly international data transit, additional IXPs will reduce networks' ongoing operational costs. These cost savings will flow to Canadian Internet users, and unnecessary export of capital will be reduced.
- By providing high-speed domestic links, additional IXPs will increase the amount of bandwidth available to Canadian users, mitigating networks' bandwidth shortages and removing networks' incentives to impose bandwidth throttling and usage caps.
- By favoring shorter and more direct routes, additional IXPs will reduce network latency, improving the performance of new services like video and cloud-based applications.
- By allowing Canadian data to remain in Canada as much as possible and as often as possible, additional IXPs will reduce the risk of Canadian data becoming subject to foreign laws and practices.

- By increasing the richness and density of connections between Canadian networks, additional IXPs will increase the reliability of Internet access in Canada and its resilience to disaster and attack.



The figure above shows alternative paths to connect two Canadian Internet customers. The top option shows a short, efficient, direct path via a Canadian Internet exchange point, whereas the bottom path detours via a US ISP and Internet exchange point.

Relative to comparable international peers, Canada is well behind in its provision of IXPs. Canada currently has two operational IXPs (in Toronto and Ottawa) and three in planning – approximately one operational IXP per 17 million people, three to thirty times fewer than other similarly developed nations. A void of Canadian Internet strategy and policy has resulted in Canadian Internet users inheriting U.S. policy, costs, and flaws rather than enjoying an environment deliberately crafted to Canadian benefit.

The provision of IXPs is not automatic: a group of networks must step up and take the lead in addressing the physical, managerial, and technical requirements. Moreover, the mere presence of an IXP is not sufficient to improve conditions; ISPs must actually make the effort to use them. Indeed, as many Canadian networks peer in London, UK, as in the existing exchange in Ottawa. Fortunately, IXPs typically cost less than \$40,000 to establish, and return on investment can be seen in as little as a few days.

This document proceeds as follows: In section 2, we present the relevant technical underpinnings and resulting incentives. In section 3, we explore the benefits of increasing the number of IXPs in Canada. In section 4, we offer recommendations for the number, location, and structure of IXPs in Canada. In section 5, we flag possible challenges and offer specific recommendations.

2. The technologies and incentives of Internet exchange points

The Internet is a “network of networks.” It is not a centralized, organized system. Rather, many independently-operated businesses carry data from one point to another and exchange data at the borders between their networks. Occasionally, data exchanges take place via simple bilateral pairings. Far more commonly, data exchanges take place at rich meeting points where many networks intercommunicate. These meeting points are IXPs.

A small network typically begins operation by paying another network to transport its data. This relationship, “transit,” gives the small network immediate connectivity, typically to the entire Internet. Yet transit carries a high cost: the small network typically pays a service provider network for every byte of data sent and received – often paying in terms of total transit in bytes, peak usage in bytes per second, or otherwise in proportion to usage.

More mature networks prefer to interconnect at IXPs.¹ An IXP is a location where participating networks exchange traffic without payment, an interconnection relationship known as “peering.” To the extent that a network has data to exchange with another network, and both these networks are connected to one or more IXPs near their locations, the networks can typically send and receive such data in virtually unlimited quantities without significant incremental cost. More than 300 IXPs exist across the world.² In the regions where IXPs are densest and most common, networks enjoy the highest levels of growth and profitability, and consumers’ Internet access tends to be particularly fast and inexpensive. Medium-sized networks may peer in several locations, typically connecting with dozens of other networks. The largest networks peer in dozens of locations, often connecting with thousands of other networks.

For most networks, transit is a major expense. By increasing the use of peering and reducing dependence on transit, a network can retain greater profit, offer its customers lower prices, or reinvest in faster growth – and often it can partake of all these benefits.

An IXP ordinarily offers the greatest benefit to a network if the IXP and network are proximate. A network seeking to connect to a local IXP – potentially an IXP in the same building where the network already has equipment – can often connect to that IXP with little or no recurring cost, and at a speed limited only by available equipment. In contrast, if a network must connect to a distant IXP, the connection is more costly. That cost often prompts the network to choose a lower-speed link, but the slower link undermines much of the benefit of connecting to that IXP. Additionally, a nearby IXP offers a particularly efficient path for sending data to and from other local networks. In contrast, if a network can connect only to a distant IXP, then communications with a local network require sending data from the first network to the remote IXP, then from that IXP back to the second network – an unnecessarily lengthy and expensive route that consumes capacity on both networks’ long-haul links. To connect many networks in many locations efficiently, numerous IXPs in numerous locations are necessary.

IXPs offer the greatest benefit when they are widely used. In particular, an IXP with many connected networks allows each network to exchange data with all others – sharply reducing those networks’ reliance on costly transit service or long-distance links to a remote IXP. Indeed, the number of possible connections increases at a greater rate than the number of participating networks. For example, an IXP with four connected networks facilitates six different bidirectional data flows; an IXP with six connected networks facilitates fifteen bidirectional data flows; with ten, 45.

3. The benefits of additional Internet exchange points in Canada

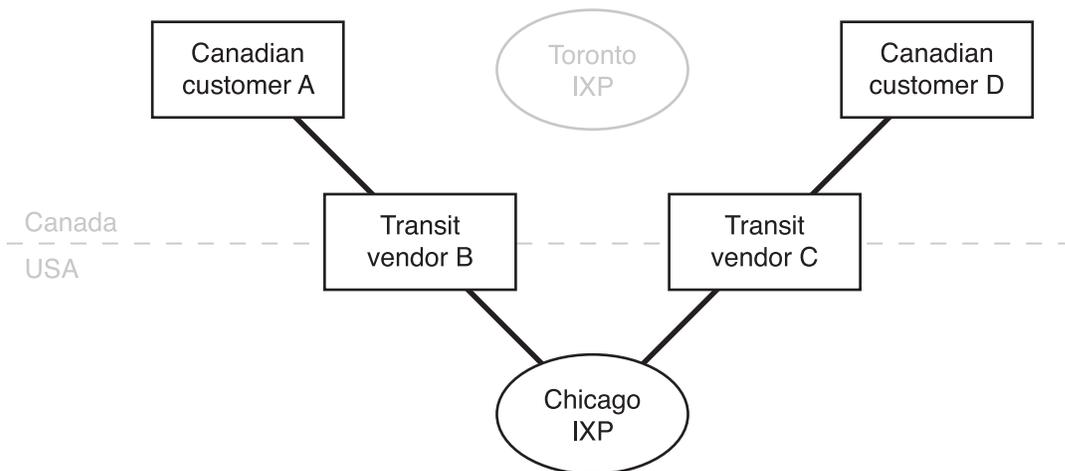
We now turn to the specific benefits likely to result from establishing additional IXPs in Canada.

Cost savings for Canadian networks and Canadian Internet users

By reducing networks' reliance on costly international data transit, additional IXPs reduce networks' ongoing operational costs. From the perspective of an individual network, the savings result from the substitution of low-cost peering for high-cost transit – letting Canadian networks exchange data with their domestic peers without paying high transit costs, and reducing networks' average per-bit delivery cost (APBDC). From the perspective of the system as a whole, the savings come from the increased efficiency of local IXPs – reducing the need for the indirect routings imposed by transit and distant IXPs.

The provision of additional IXPs in Canada would address a long-standing failure of coordination among Canadian networks. By all indications, Canada's dearth of IXPs results in large part from Canada's proximity to the United States, where IXPs are widespread. An individual Canadian network may find it easier to connect to an IXP in the United States, or simply to buy transit, than to coordinate with its competitors to form additional IXPs in Canada. Yet, in aggregate, such decisions impose unnecessary costs on Canadian networks and on all who use the Internet in Canada: anyone using the Internet in Canada is, indirectly, paying for multiple high-speed links to and from the United States, at Canadian expense and to the benefit of U.S. interests. Notably, this cost applies even to data that originates at one point in Canada with the recipient at another point in Canada.

The scarcity of Canadian IXPs poses particularly large costs to Canadian networks because of the limited competition in Canadian long-haul and international data communications. If two Canadian networks cannot peer at a local IXP because such an IXP does not exist, the networks typically either buy transit from one of only a small number of Canadian transit vendors or alternatively turn to their U.S. competitors. Moreover, the Canadian transit vendors have a natural incentive to structure network communications so that Canadian networks incur additional costs. For example, transit vendor B might announce that anyone wishing to send data to its transit customer A should deliver that data to B in Chicago. Transit vendor B can then bill customer A for the cost of transporting that data from Chicago to A's location in Canada. This approach is particularly harmful when traffic originates with another Canadian customer. Customer D must mirror A's unwarranted expense and send the data to Chicago via its transit vendor C in order for the data to reach A.



We have repeatedly observed data routed from one point in Canada to another point in Canada via an IXP in the United States. This occurrence is often known as “tromboning” because, like wind through the slide of a trombone, data flows to a distant intermediary, only to return to a destination near its origin. For an example, see the traceroute in Exhibit 3. Numerous, widespread Canadian IXPs could sharply reduce the prevalence of this wasteful practice.

To confirm the scope of Canadian data routed via other countries, we analyzed traceroutes collected by the University of Toronto IXmaps Project and the University of Washington iPlane project from 2010 through the present.³⁴ We identified 24,705 traceroutes between pairs of end users both located in Canada. Of these, 60% stayed entirely within Canada, whereas the remaining 40% were routed through one or more other countries. Most common by far was routing via the United States: fully 38% of traceroutes between two points in Canada were routed through the United States (occasionally in addition to at least one other country). It is difficult to confirm the degree to which the University of Toronto dataset is representative of Canadian Internet traffic as a whole, but these data nonetheless confirm that it is not unusual for Canadian traffic to be routed via the United States.

Avoiding unnecessary round trips to and from other countries would help Canadian networks enjoy lower costs. If network operators are able to operate at lower costs, a portion of the savings will be retained by these networks as increased profits, and a portion will flow to consumers through lower prices and higher service quality. The allocation of savings depends on the relative elasticity of supply of Internet access and demand for Internet access. Where consumer Internet access is competitive and margins are low, most of the benefits flow through to consumers.

Avoiding unnecessary international detours would also greatly increase the performance of Canadian domestic traffic. The median round-trip time for the paths in our dataset that remained within Canada was 33.37ms, while that of the paths that passed outside of Canada was 48.21ms. Thus we can see that even with only two IXPs in Canada, a performance increase of more than 30% has already been achieved, and that figure can only increase with the construction of more IXPs and shorter paths.

Increased bandwidth to Canadian users

By providing high-speed domestic links, additional IXPs will increase the amount of bandwidth available to Canadian users, mitigating networks’ bandwidth shortages and thus reducing networks’ incentives to impose bandwidth throttling and usage caps.

Data transmission within IXPs can be strikingly large – flows between 10Gb/second and 100Gb/second predominate between networks, via standard equipment at modest cost. Routing data through fast, inexpensive IXPs frees up capacity on slower, more expensive transit links, allowing for additional traffic growth on those transit links without upgrades or associated costs.

Canadian bandwidth limits result in part from the high transit costs Canadian networks face in transporting users’ data to and from desired destinations. IXPs serve to lower transit costs via dramatic efficiencies for data that Canadian networks can deliver through peering at local IXPs. By lowering networks’ average data transmission costs, additional IXPs will dull networks’ concerns about bandwidth usage and allow networks to offer users increased capacity. Meanwhile, increased use of local IXPs will yield cost savings that increase the funds available for investment in last-mile improvements.

While IXPs reduce unit cost, they also increase the amount of bandwidth available to networks to provide to customers. This in turn allows ISPs to offer bandwidth at a lower price, which increases the pool of customers who can afford to purchase the service and expands the market as a whole. Thus, IXPs support the virtuous cycle that drives the Internet’s growth and expansion. Conversely, a lack of IXPs interrupts that cycle and stunts the growth of a region’s network relative to its neighbors and the Internet as a whole.

Shortening routes to reduce network latency and improve performance

By creating shorter and more direct routes, additional IXPs reduce network latency, improving the performance of new services like video and cloud-based applications.

Several forms of delay result when Internet traffic takes a longer and more complex a route, such as Canadian data that is routed from Canada through another country and back to Canada. These delays include retransmission delay (dropped packets), systematic delay (increased latency due to distance and processing), unpredictable delay (jitter), and differential delay (out-of-order delivery). These delays particularly disrupt voice, video, and other types of streaming and time-sensitive communications.

The newest interactive network applications are particularly sensitive to network latency. For example, VoIP, videoconferencing, and video games greatly benefit from low-latency links. Similarly, remote computing applications – be they full-on Citrix-style remote operating systems or “cloud-based” web applications like Google Apps – require that every mouse click and every screen display entail a sequential series of round-trip transmissions between a user and a remote server. Lowering network latency improves the performance of these new applications – facilitating transition to systems widely understood to be the future of computing.

Retaining the benefits of Canadian privacy and other law

By allowing Canadian data to remain in Canada as much as possible and as often as possible, additional IXPs reduce the risk of Canadian data and personal information becoming subject to U.S. and other foreign laws and practices.

Once data is routed outside Canada, foreign companies may track, analyze, and even store the data pursuant to their respective privacy policies. Foreign governments also track, analyze, and store the data. These concerns are more than speculative; in 2006, a whistleblower revealed large-scale U.S. government collection of data at an AT&T network switching center in San Francisco, and subsequent investigation confirmed that such inspection was occurring elsewhere.⁵ More recently, it was revealed that the U.S. National Security Agency is building a vast datacenter to permanently archive and analyze such data.⁶ If Canadian data is kept in Canada, it need not be subject to this tracking or analysis.

Conversely, the use of U.S. or other foreign transit vendors can impede Canadian law enforcement operations. A transit vendor doing business in Canada may be subject to Canadian law, but may not be accustomed to Canadian procedures or requirements, or may impose additional fees or impediments due to international data flows. Canadian law enforcement officers have conveyed to us that they face additional impediments when Canadian data is routed via the U.S. or other countries.

Improving reliability

By increasing connections between Canadian networks, additional IXPs will increase the reliability and resilience of Internet access in Canada.

Without IXPs, networks typically find it cost-effective to rely on just one or two links to key transit providers. In this sparsely interconnected network architecture, if one of a network’s links fails, the network often suffers severely degraded performance.

In contrast, a robust web of interconnected networks, with substantial local peering at IXPs, gives networks many ways to reach each other. Thus, if one link fails – whether through operator error, equipment malfunction, maintenance, sabotage, or worse – communications continue across the many remaining links unimpeded, and performance is not degraded. IXPs make it easy and cost-effective for a network to have substantial excess capacity – capacity that is easily and automatically invoked as needed.

Recent cyber-attacks confirm the benefits of IXPs in improving network reliability. During 2007 cyber-attacks, Estonia benefited from two domestic IXPs as well as transit and peering in diverse allied nations. During similar attacks in 2008, Georgia lacked those benefits. Estonia weathered the attack easily, whereas the entire government of Georgia was forced to migrate onto Google's free hosted services until the attack subsided, months later.

4. Recommended approaches

Number and general location of IXPs needed

Canada currently has two operational IXPs, in Toronto and Ottawa – approximately one operational IXP per 17 million people. By comparison, the United States has eighty-five (one per 4 million), England twelve, Australia eleven, and New Zealand seven (one per 5 million, 2 million, and 0.5 million, respectively). See Exhibit 1. This comparison reveals Canada's laggard status in the provision of domestic IXPs.

Based on experience in other countries, we believe that IXPs should be viable in most Canadian urban areas. An IXP is more likely to thrive where the population has sufficient income to support substantial Internet use. An IXP is more likely to thrive the further it is located from existing large IXPs (which serve as partial, albeit imperfect, substitutes). An IXP should be viable anywhere where three or more networks are willing to peer with each other.

Exhibit 2 summarizes our preliminary review of prospects for IXPs in Canada, including recommended short-term and medium-term priorities for IXPs in key cities and regions. IXP locations were evaluated based on factors including population density, per capita income, Internet usage, and political centrality. Because networks' savings are typically proportional to the amount of data that can be exchanged at a given IXP, the benefit will be greatest at IXPs in the largest Canadian cities. With Ottawa and Toronto already enjoying IXPs and Winnipeg and Calgary IXPs in planning as of fall 2012, it is natural to proceed to Edmonton, Halifax, Montréal, Québec City, Vancouver, and Windsor – the proposed IXP locations detailed in the center panel of Exhibit 2. However, IXPs in all the listed cities and regions – both short- and medium-term – are likely to be cost-effective, so if local interest quickly arises in an IXP for a smaller city, the local support justifies prioritizing effort there.

We believe it is realistic for Canada to seek to establish seven to ten IXPs in total (five to eight new IXPs) in the short run (within the next 18 months). Within four years, it is realistic for Canada to host twenty-five IXPs, which would yield one IXP per 1.3 million people – between the rates now seen in Australia and New Zealand.

Specific location of IXPs

The location where an IXP is housed is crucially important, both for financial efficiency and to maximize networks' willingness to participate.

Some IXPs have been established in commercial datacenters, generally in their own dedicated rooms, or occasionally in space controlled by a participating network. It is preferable for an IXP to operate within its own dedicated space to accommodate future expansion, to secure access, to avoid the perception of non-neutrality, and to avoid the privacy and security complications resulting from sharing a participant's space.

Because an IXP requires the cooperation of members who are otherwise often competitors, it is important that each IXP be located somewhere that is and, equally importantly, is widely perceived to be neutral. An independent commercial datacenter or a centrally located office building typically satisfies this concern. A location owned or controlled by any of the participating members typically aggravates this concern, no matter how well intentioned the member is in offering to host the IXP. A location controlled by a dominant network arouses fears of abuse of market power, and one controlled by a weak network may enter bankruptcy or be acquired by a

dominant network. In either case, networks hesitate to invest in connecting to a location in which they have limited confidence.

Cost structure and technical requirements of IXPs in Canada

To date, Toronto and Ottawa are the only Canadian cities to enjoy IXPs. Established in 1998, Toronto's TorIX currently serves 156 participating networks with 79 gigabits per second of traffic. TorIX began as an unincorporated organization but after approximately five years incorporated as a corporation without share capital in the Jurisdiction of Ontario. TorIX's board is selected by annual vote of participating networks, one vote per network. TorIX is among the minority of IXPs that charge fees for participation. Initially TorIX offered Fast Ethernet access without charge, Gigabit Ethernet at \$1,000 one-time, and 10G Ethernet at \$1,000/month. Today TorIX no longer offers Fast Ethernet, and offers Gigabit Ethernet at \$1,200/year and 10G Ethernet at \$12,000/year, as well as a sub-rate 10G Ethernet capped at three gigabits for \$2,400/year. Like most IXPs, Torix receives dedicated space at no cost from the owner of the building in which it resides, for the presence of Torix vastly increases the value of the building to other tenants.

Established in 2001 and incorporated soon after, Ottawa's OttiX currently serves thirteen networks with 256 megabits per second of traffic. OttiX's board is selected by vote of participating networks, one vote per network. OttiX charges no fees. Keeping an IXP cost-neutral in this manner is a widespread best-practice, and greatly simplifies its governance while simultaneously increasing its attractiveness to participants.

Because a key goal of an IXP is to reduce its member networks' costs, it is important to keep the cost of participating in an IXP low. A network considering participating in an IXP faces multiple cost components. Most obvious is any fee the IXP may charge each participating network. IXPs typically avoid incurring expenses, so fees are never necessary, and create many perverse incentives. Space, power, cabling, and equipment all have inherent costs, but often, as in the case of TorIX, the owner/operator of the host facility (where the IXP is to be installed) donates at least space and power – recognizing that the presence of the IXP makes the building a more desirable location for networks to connect to. Equipment vendors happily donate the necessary switches. Such donations make fees unnecessary and make governance substantially easier, alleviating objections and accelerating launch.

Generally, successful IXPs limit the shared expenses incurred by the IXP as a whole. Typically, a successful IXP provides core network infrastructure, but each participating network must provide its own link to/from the IXP and its own equipment to be installed at the IXP. In this way, an IXP can avoid imposing charges on members – letting the IXP launch and grow significantly more quickly. Interested members can still add specialized equipment or additional connections if they so choose. But when each network makes these investments independently, the IXP avoids increasing costs for networks that do not require the additional services. Meanwhile, retaining networks' independence simplifies IXP governance by reducing the scope of issues requiring coordination and agreement.

In particular, substantially all IXPs have adopted a Layer 2 switched-Ethernet architecture. Networks provide their own backhaul link and a router to plug into the IXP in order to deliver and receive traffic. Participating networks own and/or operate their infrastructure at the exchange. In this way, the cost of participating is minimized, and each network has the greatest possible ability to choose cost-effective equipment to meet its needs.

Like any computer network, an IXP can be designed to be very complicated, but such complexity is unnecessary, particularly for a new IXP. In most regions, new IXPs can begin operation within three to six months using less than \$40,000 of capital. Occasionally, IXPs have been proposed that would be significantly more costly and slower to build (\$10 million and three years in recent plans from the ITU).⁷ But if such complex IXPs are needed anywhere, they are not necessary for new installations, and it is difficult to imagine how they would overcome the handicap of initial over-investment to reach profitability. If an IXP operates as an association of networks doing

business in a given region, the networks usually realize that the greatest efficiencies and benefits are gained by keeping costs as low as possible and maximizing the return on investment.

A network's biggest expense in participating at an IXP is the telecommunications infrastructure connecting the IXP and the network's users. A suitable choice of IXP location can reduce this cost. Usually, the best location turns out to be in the center of a dense urban environment – where space is more expensive, but where the IXP is closer to the users who depend upon it (reducing the average distance and cost to connect those users to the IXP). It is also desirable for an IXP to be served by competitive telecommunications providers because a competitive telecommunications marketplace will reduce the largest expense of IXP members.

Governance structure of IXPs in Canada

An IXP's participants generally share a written understanding of the business and technical terms under which they come together, often calling this an "IXP policy document" or similar. That document is generally as light-handed and non-prescriptive as possible.⁸ In light of both their neutrality commitments and the business interests of their participants, successful IXPs welcome to their premises all potential participants without discrimination. However, each individual participating network may determine which other networks it will exchange traffic with, and which of its routes it will advertise as reachable at the exchange.

An IXP may be a formally incorporated entity or an informal project. The latter usually suffices initially. IXPs tend to incorporate only when needed, e.g., to hold an insurance policy or to sign to a lease.

Because the formation and operation of an IXP require participation of multiple networks, including networks that are otherwise competitors, the successful formation of an IXP requires establishing trust and cooperation. Transparent governance models and financial management are crucial. Reviewing successful IXPs in other countries, we see four key commonalities:

- Successful IXPs are consortia of their network-operator participants. Usually, each participating network receives one vote in matters of governance.
- Successful IXPs operate at little or no expense, imposing little or no cost burden on their participants. Successful IXPs leave the profits to be realized directly and efficiently by each of their participants.
- Successful IXPs are, and are perceived to be, neutral. If an IXP falls disproportionately under the control of any one of its participants or any identifiable subset of its participants, it will not be trusted by the remainder, and it is likely to fail.
- Successful IXPs do not compete with their participants in any way. In particular, successful IXPs do not operate business ventures or carry traffic beyond the point where participants meet.

Most IXPs include a voting mechanism to let participating networks decide key questions of management and policy. Votes are typically limited to networks connecting to the IXP. Usually, each network gets a single vote. (Larger networks do not receive extra votes.) Votes are typically used to select board members or fill other roles. Votes may also be held on individual policy decisions or hardware adoption.

Fewer than ten of the world's roughly 350 IXPs are operated by research and education networks, which tend to call them "GigaPOPs" to distinguish them from things that are more clearly IXPs. Canada has two such, the QIX in Montreal and VANTX in Vancouver. Exchanges operated by R&E networks generally adhere less to the principles of neutrality and open and transparent governance than the more common not-for-profit consortia of network operators.

1. Possible challenges and suggested responses

Dominant networks refuse to participate in an IXP

In some regions, dominant networks refuse to participate in an IXP that is being formed or refuse to participate in an operational IXP. For example, historically Mexico's Telmex (94% market share as of 2011) has refused to participate in IXPs within Mexico, instead peering in other countries and forcing its domestic competitors to reach those exchange points over layer 1/2 circuits, over which Telmex is similarly dominant. As a result, Mexico is the only OECD country without an IXP, and Mexican Internet access is far more expensive than in other similarly developed economies.⁹

Similarly, Telstra historically refused to connect to IXPs in Australia, forcing Australian networks to peer with them in Los Angeles or Hong Kong, or not at all. To peer with Telstra in Los Angeles or Hong Kong, Australian networks would need to buy costly international data links to get data to/from those locations, and as the primary shareholder in the submarine fiber systems connecting those locations, Telstra would be a likely vendor of those links, so Telstra knew that peering in Los Angeles or Hong Kong would entail costs sufficiently high that a network might instead buy transit from Telstra. Australia's attempted solution was regulation, but this solution has brought limited success due to over-prescriptive regulation that failed to follow established best practice, and resulted in the formation of the "gang of four" oligopoly.

Likewise, Bell Canada does not peer in Canada, preferring to force its Canadian counterparts to meet it in Seattle, New York, or San Jose. (See Exhibit 4, tabulating peering locations of Canadian transit networks.)

Whether or not a dominant network participates, other networks can enjoy the significant benefits of using IXPs. In particular, the benefits detailed in preceding sections do not anticipate that a network will deliver all of its traffic via IXPs. Because the benefit of an IXP increases when a network can deliver more traffic via the IXP, a network can enjoy some benefits even if IXP participation is limited. Moreover, the networks that connect to an IXP enjoy a cost advantage for their communications with each other – generally allowing them to provide higher-quality service to their customers at lower prices, allowing them to expand market share and thereby increase the proportion of traffic deliverable via their IXP. Thus, if a dominant network refuses to participate in an operational IXP, other networks' best response is to embrace the IXP and the competitive advantage it confers.

Dominant networks try to dominate the process

Sometimes dominant networks also try to dominate the process of forming and running an IXP. In Malaysia, dominant networks Telekom Malaysia and Jahring successfully opposed the first attempt to form an IXP. When it later became clear that networks would form an IXP anyway, the dominant networks offered to host the IXP in their datacenters, but they proposed to charge such a high price for this service (above the price of transit) that the IXP would not have been cost-effective, and they thereby successfully blocked the second attempt. In a third iteration, dominant networks proposed a three-location IXP, using the facilities of Telekom Malaysia and Jahring as well as at a third location to be chosen by the smaller networks. All of the other networks chose to connect to that third location. Initially, the terms imposed by Telekom Malaysia and Jahring required the costs of circuits between the three locations to be divided equally between all members, but paid to Telekom Malaysia and Jahring, as the providers of those circuits, regardless of whether other networks wanted them. Now that most networks connect at a single location, the networks are well positioned to cease paying for the links between that location and the premises of Telekom Malaysia and Jahring.

Networks compromise to agree on IXP location

One of the most important decisions in the process of forming an IXP is the selection of the location for the exchange. Networks sometimes struggle to come to agreement in this matter. For

example, each network typically prefers that an IXP be placed at a major service point for its own operations. Attempts to find suitable locations through open discussion are sometimes stymied by the awkwardness of the give-and-take – often, a combination of networks wanting to appear generous yet also wanting to serve their self-interest.

PCH has found that an expedient approach is voting through prioritized lists each network prepares separately. Each network is instructed to prepare an ordered list of the locations where it would be willing to connect to an IXP, from most preferred to least preferred. Usually, each network’s single most-preferred location is a location of its own, of limited or no interest to others. But many networks often share several highly ranked locations. The stylized example below illustrates typical preferences. PCH has successfully used this method to guide the selection of IXP locations more than one hundred times. Usually PCH styles the prioritized list as a “homework assignment” to be prepared by network staff separate and apart from other networks, a process that seems to yield thoughtful non-self-serving responses.

	ISP A's choices	ISP B's choices	ISP C's choices	ISP D's choices
First choice, weighting multiplier: 4	ISP A's headquarters	ISP B's headquarters	ISP C's point of presence	ISP D's datacenter
Second choice, weighting multiplier: 3	ISP D's datacenter	1 Main Street	50 Central Avenue	1 Main Street
Third choice, weighting multiplier: 2	10 Broad Street	50 Central Avenue	1 Main Street	50 Central Avenue
Fourth choice, weighting multiplier: 1	50 Central Avenue	10 Broad Street	10 Broad Street	10 Broad Street

When the weighted votes are tabulated, we get the following results:

	Weighted Votes
1 Main Street	8
50 Central Avenue	8
ISP D's datacenter	7
10 Broad Street	5
ISP A's headquarters	4
ISP B's headquarters	4
ISP C's point of presence	4

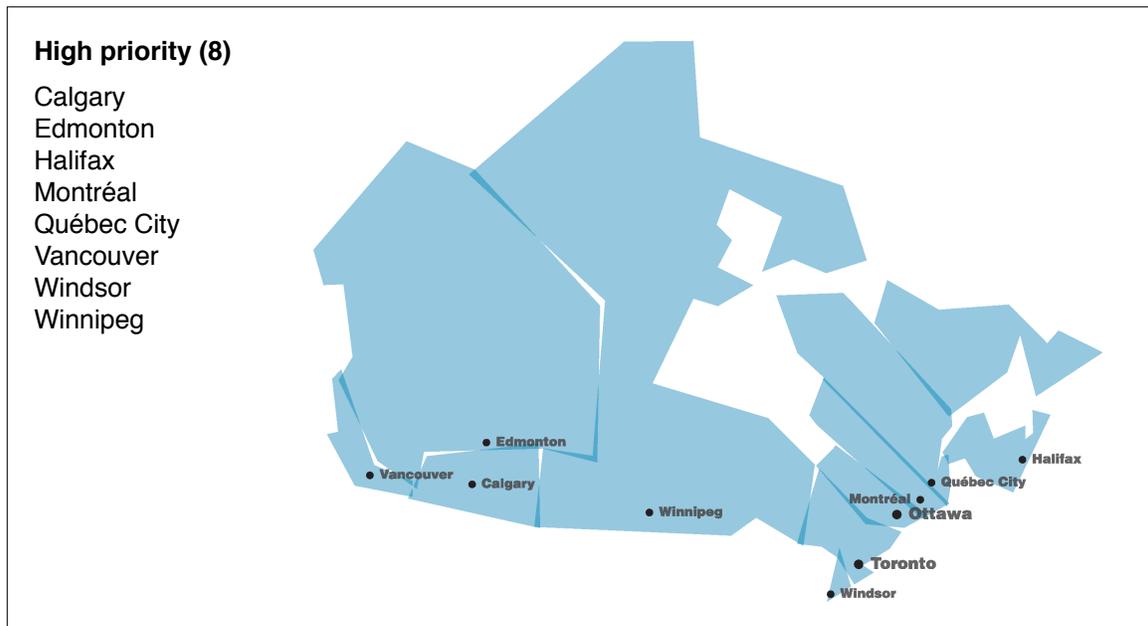
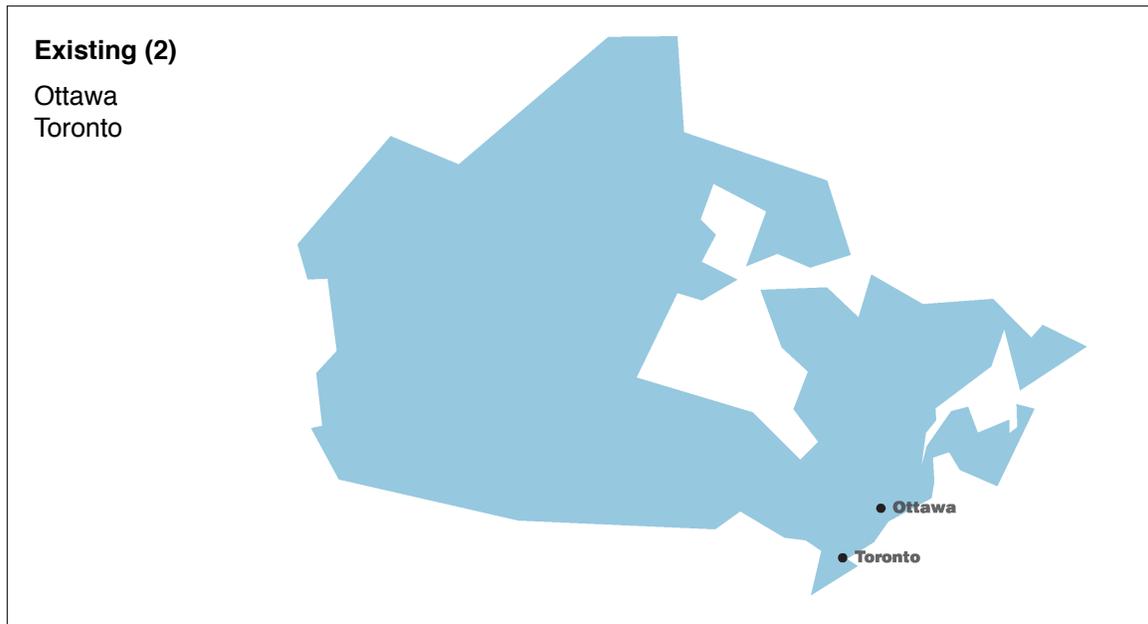
In this case, 1 Main Street and 50 Central Avenue are broadly acceptable, even though these locations were not the first choice of any of the participating networks. ISP D's datacenter is a close runner-up because it is also used by ISP A. Following PCH's usual procedure, the working-group of participating networks would send representatives to the owners of the buildings at 1 Main Street and 50 Central Avenue. In discussions with those building-owners, the working-group representatives would explain the nature and requirements of an IXP, seeking proposals for space, power, and other resources the buildings would offer to obtain the IXP's tenancy. With two or more proposals in hand, the working-group would choose a location balancing the preferences of participating networks with the costs and other terms from building owners.

2. Exhibits

Exhibit 1: IXPs by country (selected countries)

Country	GDP (M CAD)	Population	Area (km ²)	IXPs	Bandwidth (Mb/sec)	Broadband per 100	Pop/IXP (1,000s)
Argentina	366,566	40,091,359	2,780,400	1	1,420		40,091
Australia	1,223,184	22,705,020	7,692,024	11	3,500	24.1	2,064
Austria	373,073	8,404,252	83,871	2	45,600	23.89	4,202
Belgium	461,019	10,918,405	30,528	2	31,500	30.85	5,459
Brazil	2,069,411	190,732,694	8,514,877	19	73,300		10,039
Canada	1,558,310	34,581,000	9,984,670	2	55,600	30.72	17,291
China	5,819,474	1,210,193,422	9,640,011	3	153,000		403,398
France	2,556,702	65,821,885	640,294	15	38,100	33.66	4,388
Germany	3,282,487	81,751,602	357,114	14	1,380,000	31.93	5,839
India	1,522,586	1,210,193,422	3,287,263	7	8,330		172,885
Indonesia	699,668	237,556,363	1,910,931	7	3,410		33,937
Italy	2,034,563	60,626,442	301,336	7	84,000	22.07	8,661
Japan	5,404,283	127,950,000	377,930	16	577,000	26.74	7,997
Kenya	32,093	38,610,097	580,367	2	936		19,305
Mexico	1,028,730	112,336,538	1,964,375	0	0	10.45	n/a
Nepal	14,957	28,584,975	147,181	1	19		28,585
Netherlands	775,460	16,695,800	37,354	5	1,330,000	38.09	3,339
New Zealand	126,680	4,315,800	270,467	7	540	24.93	617
Norway	410,317	4,968,200	323,782	2	33,000	34.65	2,484
Poland	463,854	38,186,860	312,685	5	191,000	14.2	7,637
Russia	1,450,428	142,914,136	17,098,242	14	759,000		10,208
Saudi Arabia	439,254	27,136,977	2,149,690	0	0		n/a
South Korea	997,013	48,988,833	99,828	4	437,000	34.03	12,247
Spain	1,395,847	46,125,154	505,992	6	256,000	23.36	7,688
Sweden	451,290	9,440,588	450,295	12	154,000	31.85	787
Switzerland	518,534	7,866,500	41,277	3	25,300	38.07	2,622
Taiwan	426,274	23,188,078	36,188	4	4,730		5,797
Uganda	16,950	31,800,000	241,550	1	52		31,800
United Kingdom	2,224,980	62,435,709	242,900	12	848,000	31.94	5,203
United States	14,511,222	312,200,000	9,629,091	85	832,000	27.74	3,673

Exhibit 2: Prospective IXP locations in Canada



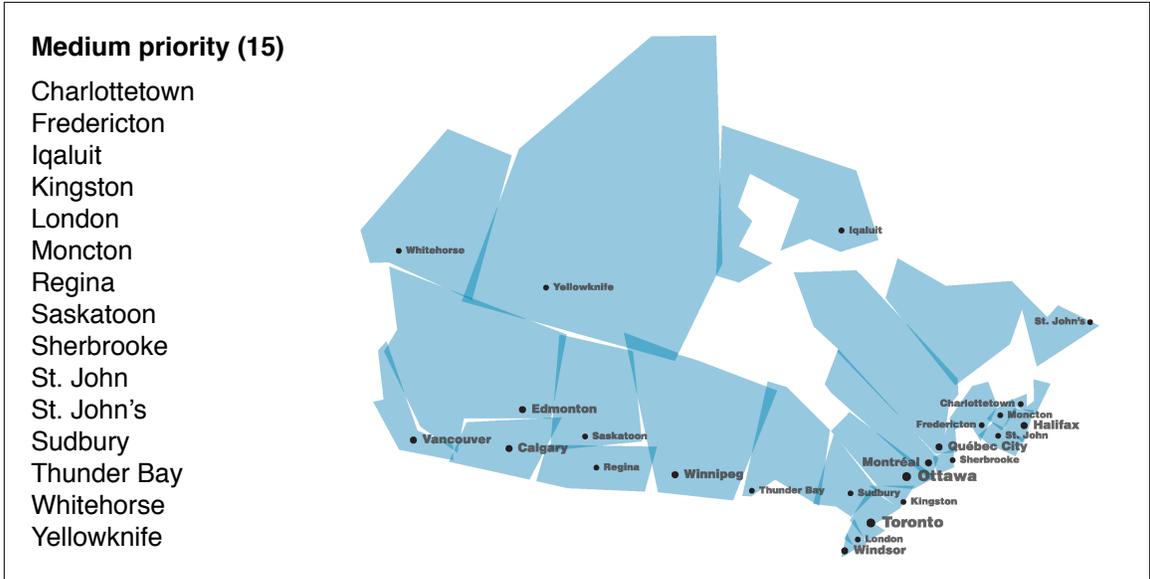


Exhibit 3: Traceroutes from One Point in Canada to Another Point in Canada, passing through the US

The following typical “tracroutes” illustrate paths taken by traffic routed from one point in Canada to another point in Canada via the United States. In this first example, the first router on the path was 206.248.154.0, an IP address registered to Teksavvy Solutions (of Chatham, Ontario), from which the traffic was routed to Washington, DC, and New York, NY, en route to Athabasca University in Athabasca, Alberta. This 2009 example is drawn from the IXmaps Project corpus of traceroute archives, as discussed on page 9.

Host Name	IPv4 Address	Company	Country
(no host name provided)	206.248.154.0	Teksavvy	Canada
2110.ae0.bdr02.tor.packetflow.ca	69.196.136.34	Packetflow	Canada
ge-7-2-6.was12.ip4.tinet.net	77.67.68.125	Inteliquent	United States
so-5-1-0.nyc22.ip4.tinet.net	89.149.187.53	Inteliquent	United States
telus-gw.ip4.tinet.net	77.67.68.42	Inteliquent	United States
edtnabxmgr01.bb.telus.com	154.11.10.141	Telus	Canada
edtnabkddr02.bb.telus.com	154.11.5.36	Telus	Canada
atbcab03-athu01.ab.tac.net	216.123.198.147	Telus	Canada
urania-a.cs.athabascau.ca	131.232.193.100	Athabasca University	Canada
ren.pc.athabascau.ca	131.232.31.232	Athabasca University	Canada

In this second example, utilizing IPv6, the first router in the path belongs to PriorityColo, of Toronto. The traffic is sent to Hurricane Electric, a US network, in Toronto, which backhauls it to New York, where Bell Canada is willing to receive it, before returning it to Canada.

Host Name	Company	Country
gi-2-16.dist01.tor1.ip6.prioritycolo.ca	PriorityColo	Canada
te-1-2.core01.tor1.ip6.prioritycolo.ca	PriorityColo	Canada
2001:504:1a::34:112	Hurricane Electric	Canada
10gigabitethernet4-1.core1.nyc4.he.net	Hurricane Electric	United States
paix-ny.bell.ca	Bell Canada	United States
2001:4958:5:1::2	Bell Canada	Canada
2001:4958:5:8::2	Bell Canada	Canada
2001:4958:5:9::2	Bell Canada	Canada
2001:4958:5:3::1	Bell Canada	Canada

Exhibit 4: Peering Locations of Canadian Transit Networks (October 2011)

Canadian ISP	ASN(s)	Toronto	Ottawa	Seattle	New York	London	San Jose	Washington	Amsterdam	Chicago	Miami	Los Angeles	San Diego	Hong Kong	Oslo	Paris	Singapore	Stockholm	Total
Accelerated Connections	21570	●																	1
Advanced Knowledge Networks	14453	●																	1
Allstream	6453 15290	●		●	●	●	●	●	●		●	●	●	●	●	●	●	●	15
Atria	26230		●																1
Bastionhost	36119						●												1
Beanfield	21949	●			●														2
Bell Aliant	855	●				●													2
Bell Canada	6539			●	●		●												3
Cipherkey	25668			●															1
E-Gate	13657	●																	1
Galaxybroadband	14500	●		●															2
Hydro One	19752	●																	1
Hypnovista	10533		●																1
Internet Light and Power	12059	●																	1
Mohawk	14537					●													1
MTO Telecom	21548	●																	1
Neutral Data Centers	33554	●																	1
Nexicom	11666	●		●	●	●		●											5
Primus	6407 7788	●	●	●	●				●										5
Priority Colo	30176	●																	1
Rogers	812	●			●														2
Shaw	6327	●		●	●	●	●	●	●	●									8
SmarttNet	30295			●															1
Storm	13319	●	●																2
TekSavvy	5645	●	●																2
TeraGo	20161 25976	●		●															2
Trends	10678	●																	1
Videotron	5769	●			●			●		●									4
Total		21	5	9	8	5	4	4	3	2	1	1	1	1	1	1	1	1	

- ¹ For an in-depth look at the interconnection practices of Internet networks, see Bill Woodcock and Dennis Weller, *Internet Traffic Exchange: Market Developments and Policy Challenges*. Organization for Economic Cooperation and Development; Directorate for Science, Technology and Industry; Committee for Information, Computer and Communication Policy; Working Party on Communication Infrastructures and Services Policy, DSTI-ICCP-CISP(2011)2, June 2011.
- ² Packet Clearing House, “Internet Exchange Directory,” available at <https://pch.net/ixpdir>.
- ³ David J. Phillips, et al., The IXmaps Project, available at <http://dev.ixmaps.ischool.utoronto.ca>.
- ⁴ Ethan Katz-Bassett, et al., iPlane: An Information Plane for Distributed Services, available at <http://iplane.cs.washington.edu>.
- ⁵ Ryan Singel, “Whistle Blower Outs NSA Spy Room,” *Wired Magazine*, April 7, 2006, available at <http://www.wired.com/science/discoveries/news/2006/04/70619>.
- ⁶ James Bamford, “The NSA Is Building the Country’s Biggest Spy Center,” *Wired Magazine*, March 2012, http://www.wired.com/threatlevel/2012/03/ff_nsadatacenter/all/1.
- ⁷ See, for example, the International Telecommunications Union’s February 2012 proposal “Internet Exchange Points (IXPs): Keeping Local Internet Traffic Local.”
- ⁸ We document IXP Policy Document best practices at <https://wiki.pch.net/pch:public:ixp-policy-document>.
- ⁹ Christian O’Flaherty and Bill Woodcock, “La Interconexión de Redes en Internet,” January 8, 2011. *Política Digital en Línea*. <http://www.politicadigital.com.mx/?P=leernoticia&Article=21023&c=9>.



© 2012 by Packet Clearing House on behalf of and commissioned by the Canadian Internet Registry Association. Published under the Creative Commons Attribution-ShareAlike (CC BY-SA) License, which is included by reference. <https://creativecommons.org/licenses/by-sa/4.0/legalcode>