

Internet Service Providers and Peering

William B. Norton <wbn@equinix.com>

Abstract

Internet Service Provider (ISP) peering has emerged as one of the most important and effective ways for ISPs to improve the efficiency of operation. Peering is defined as “an interconnection business relationship whereby ISPs provide connectivity to each others’ transit customers.” ISPs seek peering relationships primarily for two reasons. First, peering decreases the cost and reliance on purchased Internet transit. As the single greatest operating expense, ISPs seek to minimize these telecommunications costs. Second, peering lowers inter-Autonomous System (AS) traffic latency. By avoiding a transit provider hop in between ISPs traffic between peering ISPs has lower latency. So how is peering done?

This paper details the ISP peering decision-making process. Interviews with Internet Service Providers¹ have highlighted three distinct decision phases of the peering process : Identification (Traffic Engineering Data Collection and Analysis), Contact & Qualification (Initial Peering Negotiation), and Implementation Discussion (Peering Methodology). The first phases identifies the who and the why, while the last phase focuses on the how.

The appendix includes the description of a Peering Simulation Game that has been used in workshops to play out peering negotiations.

Introduction and Definitions

Internet Service Providers (ISPs) connect end-users and businesses to the public Internet. They compete with each other on price, performance, reliability, etc. but they also must cooperate with each other to provide global connectivity to all other attachments on the Internet. The cooperation is explicitly stated at demarcation points where they interconnect, called imprecisely “peering points”. This is an imperfect name since interconnection typically takes one of two forms: a peering relationship or a transit relationship, and both use the Border Gateway Protocol (BGP) for routing announcement exchange. Problems arise when the term “peering” is used interchangeably with a “transit” relationship. To lay the groundwork for this paper, we introduce the following working definitions.

Definition: Peering is the business relationship whereby ISPs reciprocally provide to each other connectivity to each others’ transit customers.

To illustrate peering, consider figure 1 below showing a much simplified Internet: the Internet with only three ISPs: WestNet, USNet, and EastNet. WestNet has customers shown as green circles. USNet has customers of its own (beige circles) and EastNet has its customers shown as yellow circles.




Figure 1 - Peering and Transit relationships

In this example, WestNet has a peering relationship with USNet in which USNet announces reachability of its beige customers to WestNet, and WestNet announces reachability to its green customers to USNet. This is the essence of the peering relationship; each ISP reciprocally provides access to each others customers. EastNet also peers with USNet, announcing its yellow customers to USNet while USNet announces its blue customers to EastNet.

It is important to note that WestNet and EastNet can not access each others customers in this configuration. (The boxes below the ISPs show their respective routing tables.) WestNet only knows how to get to blue and green customers, and EastNet knows how to reach only blue and yellow customers. The fact that they both peer with USNet is inconsequential; **peering is a non-transitive relationship.**

Since peering is a reciprocal non-transitive relationship, EastNet and WestNet must peer with every other ISP or find another way of accessing every other ISP.

¹ Interviews with 100 ISPs over the course of two years along with presentations of the findings to ISPs at NANOG, RIPE and IEPG substantially validate the findings.

Definition: *Transit* is the business relationship whereby one ISP provides (usually sells) access to **all destinations** in its routing table².

Consider a simplistic Internet model below in figure 2. In this picture, EastNet **purchases transit** from USNet, and in return gets connectivity to **all** networks in USNet’s routing table. As a customer of USNet, USNet announces EastNet routes across all of its peering and transit interconnections.

By EastNet purchasing transit, EastNet is announced by USNet to USNet Peering and Transit interconnections alike.




Figure 2 - Transit Relationship - selling access to entire routing table

To better illustrate transit, note that there are over 8800 ISPs in the US³ alone; imagine the complexity and cost of trying to peer with all of them! (In reality you wouldn’t have to peer with all ISPs, but you would need to peer with their upstream (transit) providers to avoid transit entirely.) A more effective demonstration of transit is shown below, showing the opaque conglomeration of International networks accessible via a transit relationship with a large transit provider. Compared against the relatively small number of routes received from a single peering relationship, one can see that transit is a valuable and different service from peering.




Figure 3 - Transit Providers with large network

As a side note, some service providers⁴ prefer a transit (customer) relationship with ISPs for business reasons, arguing that the threat of lost revenue is greater than the threat of terminating a peering arrangement if performance of the interconnection agreement is inadequate.

Now that we have introduced the notion of peering and transit relationships and the difference between them, we will examine the role of the Peering Coordinator.

I. Phase 1: Identification of Potential Peer: Traffic Engineering Data Collection and Analysis

We spoke with over 100 peering coordinators in this study to document how peering works, and how they approach peering from a practical perspective. Peering Coordinators are typically charged with establishing and managing the interconnections between their network and others. This multidiscipline job crosses the boundaries of network architecture, technical (routing logistics), business, and legal. This job therefore requires a mix of skill sets to be executed effectively.

We’ll first examine the peering coordinator motivations for peering and selection of peers.

² Note that increasingly in Europe ISPs are offering and obtaining hybrids. For example, they may purchase “Regional Transit” from global players in a region without adequate coverage. In a few rare cases ISPs have arranged “paid peering” to eliminate the cost of peering to one or the other. INSNNet and GXNet for example.

³ See the Boardwatch list: <http://www.thelist.com> for a list of ISPs.

⁴ Conversation with Allan Leinwand, Founder of Digital Island. Allan indicated that the financial “teeth” are much stronger with transit ISPs than with “peers”, and the threat of lost revenue provides better quality and reliability in a transit relationship.

Motivations: Why Peer?

Lower Transit Costs. Choices made by Internet Service Providers (ISP) are often dominated by telecommunications cost issues. Highest among these costs is Internet transit service that provides the ISP with connectivity to the global Internet. Transit Prices for DS-3 transit for example can be as high as \$50,000/month⁵, and OC-3 transit can cost up to \$150,000/month⁶. To reduce these costs, ISPs seek peering (zero or reduced cost) relationships with other ISPs that provide more direct traffic exchange and reduce the load on these expensive transit services (as shown below).




Figure 4 – Migrating traffic from Transit to Peering Interconnection

Lower Latency. As a side effect of interconnecting directly with peers, ISP customers experience lower latency to the other ISP's customers. Traffic destined for a local competitor's customers may need to traverse a couple of transit providers and potentially across great distances (with high latency) before reaching the other customer. The worst example highlighted traffic between the United Arab Emirates and Saudi Arabia traversing an overloaded exchange point in Washington DC⁷. Through direct interconnections (using direct circuits or regional exchange points) ISP customers realize better performance.

Usage-based traffic billing. Some ISPs charge customers based upon metered traffic. Since packet loss and latency slows traffic consumption, they benefit from a lower latency, lower packet loss Internet. It is in their best interest therefore to assure that customers use as much bandwidth as possible by minimizing loss and latency through effective traffic engineering⁸.

Conversations with European ISPs showed increasing adoption of the usage-based billing model. This has motivated them to compete for traffic and therefore revenue.

Why not peer?

On the surface peering appears to be a good idea from a financial and technical perspective. However, the topic has generated more heat than light due to the following conflicts of interest between ISPs.

- **Traffic Asymmetry and Investment asymmetry** means that one party bears more of the cost as a result of peering. For example, consider the figure below where Exodus peers with GTEI. Web traffic (the dominant traffic flow on the Internet) is inherently asymmetric. Exodus is a net source of content, therefore more GTEI resources (bandwidth) than Exodus bandwidth are consumed as a result of peering. GTEI could say that Exodus was “dumping traffic” onto GTEI's backbone and GTEI was forced to “carry” the traffic across the great distances. Meanwhile, Exodus and Exodus customers are able to sell advertising on their web pages and yield great revenues off of GTEI's customers. In some cases ISPs will peer without settlement up to a certain traffic ratio (for example 4:1 traffic out to traffic in) and then on a usage basis beyond that on a Mbps basis⁹. In the Exodus-GTEI negotiations, rumor has it that the solution was to peer at more locations and to engineer cold-potato¹⁰ routing to reduce the distance the traffic had to spend on the

⁵ Varies by transit provider, backhaul costs vary by circuit miles, carrier competition, etc.

⁶ Dave Rand interview with the Cook Report, and Author interview with Pat Binford-Walsh (UUNet) in 1998. Note that the bandwidth and transit costs have dropped.

⁷ Consulting work with the United Arab Emirates PTT.

⁸ Interview with Avi Freedman, AboveNet.

⁹ Conversations with Frontier Global Center engineers at RIPE 35 in Amsterdam.

¹⁰ Cold-potato routing is a routing discipline whereby one ISP carries traffic as far as possible before handing it off to another ISP.

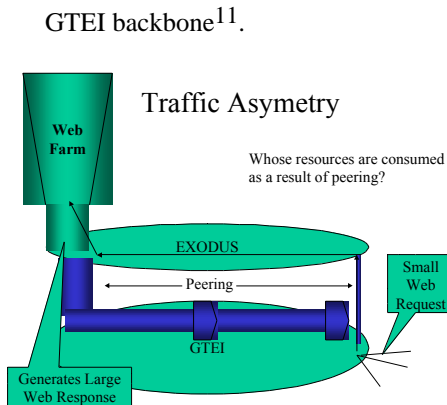


Figure 5 - ISP Traffic Asymmetry

- There may be the potential for transit sales if they don't peer¹². Some ISPs will not peer if there is **any** existing or pending customer-provider relationship between the parties, even if the sale is completely unrelated to interconnection (i.e. fiber sale or colo sale)¹³.
- Peering consumes resources (router interface slots, circuits, staff time, etc.) that could otherwise be applied to revenue generation. Router slots, cards, interconnection costs of circuits or Internet Exchange environments, staff install time are incremental expenditures¹⁴. Further, operating costs, particularly for peering sessions with ISPs without the necessary on-call engineering talent, can require increased processing power for filters and absorb time better suited to paying customers.
- Motivation not to commoditize IP transit.

Tier 1 ISPs¹⁵ compete on the basis of better performance. They accomplish better performance due to the large customer base of direct attachments and high-speed interconnections with other Tier 1 ISPs. Since peering with other ISPs improves the performance of the "peer" it effectively makes them a more powerful competitor. Therefore, there is a strong disincentive to peer and increase the number of top tier competitors.

- As a "peer" there are no Service Level Agreements (SLAs) to *guarantee* rapid repair of problems. Both parties benefit from the reparation of outages, and may even have clauses in the peering agreements to work diligently to repair the problem. However, a customer relationship (with or without SLAs) generally has more contractual teeth (financial repercussion for failure to perform).

Traffic Ratio-based Peering. As a result of these forces, a traffic ratio-based **paid peering model** is emerging. In this approach, peering is free until traffic asymmetry reaches a certain ratio (4:1 is common). At this point, the net source of traffic will pay the net sink of traffic a fee based upon traffic flow above this ratio.

With Whom to Peer?

If peering makes sense from a technical and financial perspective, the next question is, "With whom should we peer?" To identify potential peers, ISPs use a variety of criteria.

Quantities of traffic distributed between networks often sets the pace of the negotiation; to quantify this, ISPs may systematically sample inbound and outbound traffic flows. Flows then are mapped to originating AS, and calculations are made to determine where peering (direct interconnections) would most reduce the load on the expensive transit paths. There is substantial work involved here, as this traffic sampling results in a large number of data. Alternative measurement methods include measuring port statistics¹⁶.

¹¹ "GTEI, Exodus Make Peace On Peering", Randy Barrett, Inter@ctive Week, September 16, 1998 and anonymous interviews with parties involved with the negotiations.

¹² A case recited to the author: Level 3 refused to peer with GST Networks in the Washington DC area since GST was a Level 3 transit customer in New England. While technically the request could be accommodated, Level 3 rationally preferred the transit revenue.

¹³ Apricot 2000 presentation by Lauren Nowlin, Onyx, March 2, 2000 in Seoul, Korea on "Peering and Interconnection Panel" with the author.

¹⁵ Tier 1 ISPs are defined as ISPs with global coverage and a full Internet routing table without acquiring transit from any ISP.

¹⁶ Avi Freedman, AboveNet citing ATM and other switch measurement methods in use.

Many peering coordinators indicated that peering selection is accomplished by intuition¹⁷. Their sense was that they knew where traffic was and would be headed.

In either case, the end result of this first phase is list of the top 10 ISP candidates for peering. Interviews with Peering Coordinators highlighted a few other considerations.

Broader business arrangements between ISPs may circumvent the peering negotiation phase and expedite discussions directly to Phase III, the peering methodology negotiation phase.

Peering policies range across a wide spectrum from “open peering policy” meaning “we will peer with anyone”, to “if you have to ask, we won’t peer with you.”¹⁸ Peering policies are often exposed only under Non-disclosure agreements, and these policies reduce the number and type of ISPs that are peering candidates.

In many cases peering requires interconnections at multiple peering points, explicit specifications for routing, migration from public (shared switch) peering to private (non-shared switch) peering after a certain traffic volume is reached, etc. It is beyond the scope of this document to fully explore the technical and political motivation for peering policies; it is sufficient to be aware that these discussions can be cumbersome and require a combination of technical and financial negotiation.

The greatly simplified peer qualification decision tree looks something like this:

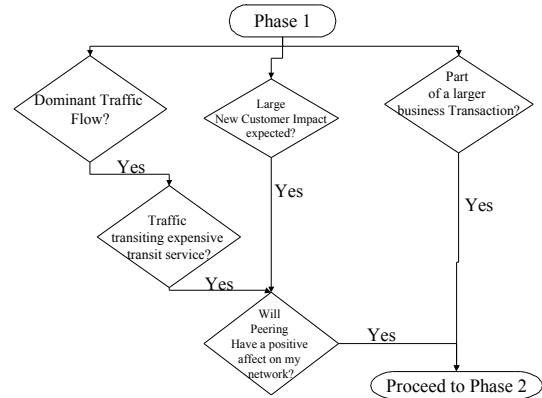


Figure 6- Phase 1 of Peering Selection Decision Tree

Once the measurements have been made and analyzed, and it appears to be beneficial to peer, the ISP enters into Phase 2, Contact & Qualification, Initial Peering Negotiation.

Emerging Migration Path from Transit to Peering. Interviews with tier 2 ISPs highlighted an emerging peering transition strategy:

- 1) Access the Internet via transit from a global provider,
- 2) Pursue peering arrangements on public switches at exchange points to reduce load on transit links and improve performance
- 3) Migrate high traffic public peering interconnections to private interconnections (via fiber or direct circuits).
- 4) Ultimately migrate traffic away from transit purchase and negotiate (free or for-fee) peering with former transit provider.

To illustrate this path, consider Telia, a global ISP based in Sweden. Telia analyzed their transit costs and recognized that approximately 85% of their traffic at MAE-East was to their transit provider and the remaining 15% was through peering relationships. By focusing on establishing peering relationships with the top 25 destination ASes they shifted the mix to 70% through private peering at an exchange with the remaining 30% of traffic heading toward their transit provider¹⁹.

¹⁷ NANOG 17, Montreal, Peering Birds Of a Feather (BOF) meeting held by author when about two-thirds of the audience indicated that they use ad-hoc, predictive, or intuition for selecting peering candidates.

¹⁸ Sentiment articulated by Sean Doran, shortly after leaving SprintLink. Peering policies are a politically sensitive subject, and peering policies are often not explicitly articulated.

¹⁹ Interview with Anne Gibbens (Telia)

The result was increased traffic efficiency and a reduction in the cost of transit²⁰.

It should be stated that phase four of the migration strategy listed above may be overly optimistic and/or challenging for several reasons. First, transit providers prefer paying customers to peers. Second, transit providers typically have much more ubiquitous network infrastructure than their customers, and therefore will not see their customers as equal contributors. Finally, the transit providers have an incentive to reduce the number of their own competitors.

To illustrate this migration path, Raza Rizvi (REDNET) said “We had to leave our upstream provider for 16 months with alternative access to their route before they considered us not as a customer lost but as a potential peering partner.”

After the top 10 potential peers are identified, peering coordinators proceed to Phase 2: Contact & Qualification, Initial Peering Negotiation.

II. Phase 2: Contact & Qualification, Initial Peering Negotiation

Internet Service Providers typically have a person or group specifically tasked with peering and traffic engineering issues. For example, UUNet has a “Peering Steering Committee” to evaluate peering requests²¹. Some variations of the following steps lead to the parties either leaving the negotiation or proceeding to peering methodology discussions.

Interviews have highlighted a key challenge for ISPs. **Finding the right person to speak with at the target ISP** is a difficult and time intensive process. Peering Coordinators change jobs and there is no standard way to find out who handles this task. Mergers and acquisitions cloud lines of communication. Even if the name is known, Peering Coordinators are often traveling, way behind in e-mail, and prioritizing e-mail based on the subject or the sender. This is where “people networking” helps a great deal, and hiring expertise for their contacts speeds this initial contact process up quite a bit. In some cases, peering is expedited between ISPs simply because the decision makers have a previous

relationship²². This was the dominant mode of operation in the early days of the Internet.

In any case, peering contacts are initiated in one of the following ways:

- a) via electronic mail, using the pseudo standard peering@ispdomain.net or a personal contact,
- b) from contacts listed on an exchange point participant list,
- c) with tech-c or admin-c from DNS or ASN registries,
- d) informal meeting in an engineering forum like NANOG, IETF, RIPE, etc.,
- e) at trade shows from introductions among speakers, or with booth staff,
- f) from the target ISP sales force,
- g) from the target ISP NOC,
- h) as part of a larger business transaction.

Second, mutual non-disclosures agreements (NDAs) may be negotiated and signed, and a discussion of peering policy and prerequisites follow. Note that NDAs are an optional step, and many ISPs do not require signed NDAs prior to discussions²³. Traffic engineering discussions and data disclosure may be used to justify the peering relationship. Each ISP typically has a set of requirements for peering that include peering at some number of geographically distributed locations, sometimes at public exchange points.

Traffic volume is usually a key determining factor. The decision rule hinges upon whether or not there is sufficient savings from peering to justify spending capital on a port on a router and/or a portion of the interconnection costs or augmenting existing capacity into an exchange point. A *Bilateral Peering Agreement*²⁴(BLPA) is the legal form that details each parties understanding of acceptable behavior, and defines the arms length interactions that each would agreed to.

²⁰ As compared with growing the transit connection.

²¹ Point made by Paul McNulty at the 1999 Apricot Session titled “Next Generation Internet Infrastructure”.

²² Discussion with Vab Goel, former VP of Engineering at Qwest.

²³ NANOG Peering BOF, NANOG 17 in Montreal, about 70 Peering Coordinators of 125 indicated they do not require NDAs.

²⁴ See <http://www.linx.net/joininfo/peering-template/agreement-v4.html> for sample BLPA

Another motivation for peering to factor in includes lower latency and/or more regional distribution of traffic than existing connections allow.

This process is diagrammed below.

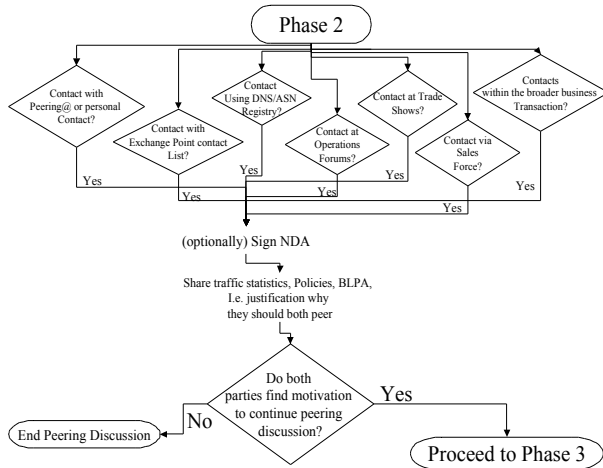


Figure 7 – Phase 2: Contact & Qualification Decision Tree

After this initial discussion, either party may decide to walk away from the peering discussions until certain criteria are met²⁵. If both parties agree that their requirements are sufficiently met to discuss methodology (they both benefit from the peering relationship), they move onto Phase 3: Implementation Discussions.

III. Phase 3: Implementation Discussions: Peering Methodology

Since peering is seen as being of mutual benefit, both parties now explore the interconnection method(s) that will most effectively exchange traffic. The primary goal is to establish point(s) of interconnection, and secondarily detail optimal traffic exchange behavior (using Multi-Exit Discriminators (MEDs) or other traffic weighting techniques).

To interconnect, ISPs face two distinct options: Direct Circuit Interconnection or Exchange-Based Interconnection (or some global combination thereof).

The “Interconnection Strategies for ISPs” white paper²⁶ quantifies the economics and technical tradeoffs between the first two options. To summarize this report, the preferred methodology depends on the number of peers participating in the region and bandwidth required for its regional interconnections. ISPs that expect to interconnect at high or rapidly increasing bandwidth within the region, or expect interconnections with more than five parties in the region prefer the exchange-based solution. Those that do not anticipate a large number of regional interconnects prefer direct-circuits and typically decide to split the costs of interconnection with the peer by region. On occasion the costs are covered in whole by one peer²⁷.




Figure 8 – ISP Physical Interconnection Methods

For direct-circuit interconnects, key issues center upon interconnection location(s) and who pays for and manages the interconnection. This becomes a material cost issue as traffic grows and circuits increase in size and cost.

In either case, ISPs generally have the following goals for establishing peering:

1. get peering set up as soon as possible,
2. minimize the cost of the interconnection and transit costs,
3. maximize the benefits of a systematic approach to peering,

²⁵ According to participants at the NANOG17 Peering BOF led by the author, government agencies in Israel and Australia forced ISP peering!

²⁶ Interconnection Strategies for ISPs, W. B. Norton, June99, presentation: <http://www.nanog.org/mtg-9905/norton.html>. A copy of this report can be requested via e-mail to wbn@umich.edu.

²⁷ Interviews found a pattern in which PSINet would peer with ISPs provided that peer covered all interconnection costs.

4. execute the regional operations plan as strategy dictates (may be architecture/network development group goal), and
5. fulfill obligations of larger business agreement.

Exchange Environment Selection Criteria

This section details the selection criteria an ISP typically uses when selecting an exchange. Note that these issues are listed in no particular order. These issues are shown graphically as flowcharts and discussed in detail in the paragraphs below.

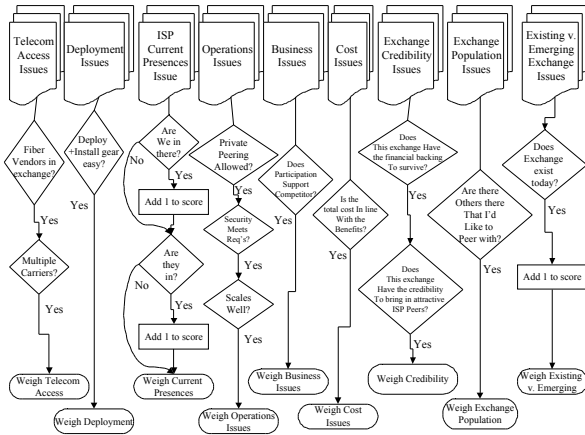


Figure 9 - Exchange Environment Selection

Telecommunications Access Issues

These issues have to do with getting telecommunications services into the exchange. How fast can circuits be brought into the interconnection environment? How many carriers compete for business for circuits back to my local Point of Presence (POP)? For facilities-based ISPs, what is the cost of trenching into the exchange (how far away and what obstacles present themselves)? Are there nearby fiber providers that lease strands? These questions will help answer the most important question to ISPs: How fast can my peer and I get connectivity into the exchange? Multiple carriers lead to speed and cost efficiencies. Some ISPs have volume deals with certain carriers or otherwise prefer carriers and therefore prefer exchanges where these carriers can quickly provision circuits. These answers strongly impact the desirability of the exchange environment.

Deployment Issues

These issues have to do with getting equipment into the exchange. How do I get my equipment into the exchange (assuming it supports collocation)? Do I ship

equipment in or do I have to bring it with me as I fly in? Will someone act as remote hands and eyes to get the equipment into the racks or do I do the installation myself? Comparing exchange environments in this context, what are the costs associated with deployment (travel, staff time, etc.) into this exchange? Does the exchange have sufficient space, power, air conditioning, etc. The answers to these questions impact the deployment schedule for the ISP(s) engineers and the costs of the interconnection method.

ISP Current Presences Issues

This issue is based on the following observation by the peering coordinators: *The most inexpensive and expedient peering arrangements are the ones made between ISPs that are already located in the same exchange.* There is a hidden assumption here that there is sufficient capacity to interconnect at the exchange. Cross-connects or switching fabrics can easily establish peering within a few hours or at most days. ISPs will prefer to interact where one or both ISP already has a presence.

Operations Issues

These issues focus on the *ongoing* operations activities allowed within the exchange after initial installation. Does the exchange allow private network interconnections? Are there requirements to connect to a central switch? How is access and security handled at the facility²⁸? Is there sufficient power, HVAC, capacity at the switch, space for additional racks, real time staff support²⁹? Is it easy to upgrade my presence over time? Upgrading in this context means the ability to increase the speed of circuits into the exchange, the ability to purchase dark fiber, the ability to increase the number of racks and cross connects in the exchange, the ease of increasing the speed of interconnection. ISPs will prefer bandwidth-rich, ISP-friendly exchanges over those with restrictions over future operations.

²⁸ For example, the NSPIXP is a major exchange in Japan yet has no staff on-site so engineers need to be called in for support. Escorted access could take hours to be coordinated.

²⁹ MAE-East has been widely criticized for being a major interconnection point in the US without sufficient infrastructure (power, A/C) to support expansion.

Business Issues

“Bandwidth, strategic partner alliances, and corporate ties often override the technical justification³⁰.” – Lauren Nowlin, Peering Coordinator for Onyx Networks.

Perhaps the most far-reaching issue is strategic: do we want to support this exchange operator, and do their interests enhance or conflict with ours?

Will using this exchange support a competitor (contribute to their net income, their credibility, their positioning)? A neutrally operated exchange (defined as one that is not owned or aligned with any carrier, fiber provider, or ISP) provides an open distortion-free marketplace for carrier and ISP services.

Market distortions often result when an exchange is owned by one of its participants. This often manifests itself in requirements (required use of their carrier or ISP services) that constrain the market for services within the exchange³¹. Since it is difficult and disruptive to move equipment out of an exchange, ISPs will prefer a neutrally operated exchange environment that will not suffer from market distortions and limitations due to business conflicts of interest.

Cost Issues

This broad issue crosses all other issues. What is the cost of using this exchange? What are the rack fees, cross connect fees, port fees, installation fees? What are the future operating fees going to be? What are the motivations and parameters surrounding these fees? Cost issues shadow most of the other issues listed in this paper. All else being equal, ISPs will seek to minimize the costs, particularly upfront costs, associated with the interconnection for peering.

Credibility Issue

The credibility issue is twofold.

First, credibility goes to the financial support of the exchange. Does the exchange exist today and will it exist tomorrow? During the early stages of the exchange, ISPs are asked to make a leap of faith when committing, and therefore prefer an exchange with strong backing and the credibility to survive.

³⁰ Discussions at NANOG 17 in Montreal.

³¹ MAE-East is owned and operated by MCI Worldcom and requires use of MCI circuits to access MAE-East services. Exodus requires use of its network and at one point restricted direct access between ISPs and Carriers.

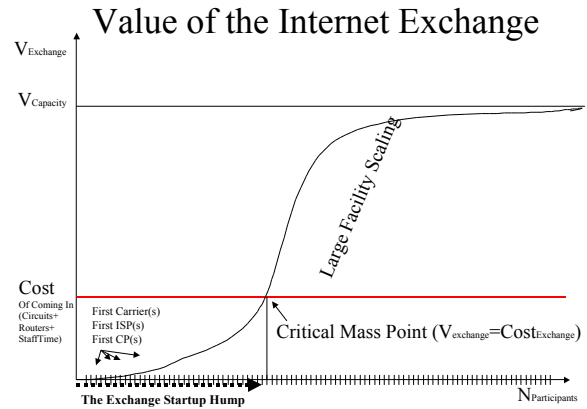


Figure 10 - Value of exchange varies over population

Second, does the exchange operator have the backing and credibility to attract the more valuable peering candidates? Since the value of the exchange (shown in the graph below) is proportional to the number and type of participants. Does this exchange have the backing to attract my peers? Who is managing the exchange and what technology is in use? These answers signal the credibility and survivability of the exchange. ISPs will prefer an exchange with credibility – one that is financially and technically well backed and likely to attract the most desirably peering candidates.

Exchange Population Issues

These issues focus on the side benefits to using this exchange. Are there other ISPs at this exchange that are peering candidates? Are there transit sales possible at the exchange? In the context of the credibility issue discussed above, who will likely be at the exchange in the future, and when will the cost of participation equal the value of the interconnection (also known as the Critical Mass Point)? ISPs will prefer an established and well-populated exchange, particularly one with potential customers that can generate revenue.

Existing Exchange vs. New Exchange?

There are many operational exchange points in each region of the U.S. There are also *emerging* (soon to exist) exchanges that may be considered as peering points. However, given the pace of ISP expansion, it is unlikely that *emerging* exchange offerings are differentiated or compelling enough to be preferred over existing exchanges. Chronic traffic congestion can influence the decision to plan to peer in an existing malfunctioning exchange or wait until a better exchange opens. Customers with heavy flows of

regional traffic can also influence the decision. Long term benefits (scalability) may lead to preferring a next generation exchange. *However, all else considered equal, ISPs generally prefer an existing exchange to an emerging one.*

One Final Note on Exchange Criteria:

Weighting

The ISPs we spoke with shared with us varied weightings of the importance of each of these issues. To some, the most important issues were the business issues, and others weighted more heavily the operations issues. Each ISP places higher or lower importance on different issues and not surprisingly select their operations environment based on their specific criteria.

IV. Summary

This paper provides a rough description of the decision processes ISPs follow to identify and establish peering relationships. It explores the implementation phase and the criteria for exchange point selection.

The results of the interviews with ISP Peering Coordinators can be summarized with the following observations:

- 1) ISPs seek peering primarily to reduce transit costs and improve performance (lower latency).
- 2) Peering goals for ISPs include a) get peering set up as soon as possible, b) minimize the cost of the interconnection and their transit costs, c) maximize the benefits of a systematic approach to peering, d) execute the regional operations plan as strategy dictates (may be architecture/network development group goal), and e) fulfill obligations of larger business agreement.
- 3) The selection of an exchange environment is made relatively late in the peering process.
- 4) ISPs highlighted 9 selection criteria for selection of exchange environment: telecommunications access issues, deployment issues, current presences, operations issues, exchange population, cost issues, and credibility of the exchange environment operator. ISPs weight these issues differently and will prefer and exchange environment that best suit these needs.
- 5) One major challenge facing Peering Coordinators is the identification of potential peers and initiating discussions.

Acknowledgements

This paper was the result of interviews with key Internet peering coordinators, interactions with ISP support staff, and side conversations at various meetings and forums. Special thanks to a few folks who contributed their time and ideas to help create this early draft report: Ren Nowlin (Onyx networks), Joe Payne (IXC), Dave Diaz (Netrail), Jake Khuon and Alan Hannan (Global Crossing), Dan Gisi and Jeff Rizzo (Equinix), Patricia Taylor-Dolan (Level 3 Communications), Sean Donelan (AT&T Labs), Avi Freedman (AboveNet), Patrick Gilmore (Onyx), Geoff Huston (Telstra), Steve Meuse (GTE), Aaron Dudek (Sprint), Raza Rizvi (REDNET), Wouter van Hulten (InterXion). Special thanks to John Odenwelder (Equinix) for editing the document.

References


- [1] "Maturation in a Free Market: The Changing Dynamics of Peering in the ISP Market" by Jennifer DePalma
- [2] "NAPs, Exchange Points, and Interconnection of Internet Service Providers: Recent Trends, Part I: 1997 Survey of Worldwide NAPs and Exchange Points", Mark Knopper
- [3] Peering Contact Database, Restricted Access; for Peering Coordinators only. To receive send e-mail to wbn@umich.edu with AS#, company name, peering e-mail address, mailing address, phone numbers, and contact names.
- [4] List of ISPs and NOC contact info: <http://puck.nether.net/netops/nocs.cgi>
- [5] ISP Survival Guide, Geoff Huston, Wiley Publishers
- [6] Data Communications "Old Boys' Network", Robin Gareiss <http://www.data.com/issue/991007/peering.html>
- [7] Australian's take on Interconnection competition issues, <http://www.accc.gov.au/>
- [8] IWS2000 Presentations on Internet Exchange futures: http://hsn-crl.koganei.wide.ad.jp/ipsj-hqi/prev_workshops/5.html

Appendix A – Calculating the Financial Benefits of Peering

Since the cost of transit is substantially higher³² than the typically zero cost of peering, ISPs try to reduce this cost with peering relationships. Before we discuss the tactics used to establish peering, we will take a brief diversion to quantify the financial value of the peering.

The Financial Value of Peering

Peering is seen as so valuable that companies are acquired because of pre-existing peering relationships³³. But how does one assess the value peering relationships? A rough approximation can be³⁴ made by measuring the traffic flow across the peering connections and comparing the cost of that traffic if sent across a transit interconnection. One can then use a perpetual annuity function to roughly value the peering sessions. To be complete one would have to factor in (at least modest) growth in traffic traversing that peering session, the chances of peering termination, improved performance and effect on customer retention, and a variety of other factors.



example: In a simple example, assume that ISP A and ISP B both pay for transit at a transit fee of

\$100,000³⁵ per month for an OC-3 (155 Mbps) worth of transit. Assume further that one-tenth (15.5 Mbps) of that is traffic between ISP A and ISP B. If ISP A and ISP B agree to peer over a large point-to-point (OC-3) at \$11,400/month³⁶, each pays half of that cost³⁷, and each pays for an interface card to support private peering over that circuit, both ISP A and ISP B reduce their transit load by 10%³⁸.

Assume a peering coordinator establishes one peering session this year. One can approximate the financial value of a peering coordinator by the cost savings to the ISP:

$$\begin{aligned}
 &10\% \text{ savings on the transit OC-3: } \$10,000 \\
 &\text{-monthly cost of peering: } \$11,400/2 = \$ 5,570 \\
 &= \$5,430/\text{month} * 12\text{months} = \$66,840/\text{year}
 \end{aligned}$$

Note 1: These cost figures are highly variable and are based on early 1999 quotes. Since then the prices have dropped dramatically, and ISPs have started offering tiered pricing. Fundamentally this doesn't change the dynamic; peering over point to point telco circuits (or less expensively over private cross connects within an exchange) has a profound effect on the cost of telecommunications for ISPs.

Note 2: Since only 10% of the interconnect is used between these two ISPs a smaller lower cost circuit could suffice. This approach allows the interconnection to scale and helps ensure that this interconnection is not a congestion point for some time.

Note 3: This is a financial benefit and ignores the decreased latency and perhaps decreased packet loss that leads to increased traffic. This benefit ultimately results in more bandwidth usage and therefore more revenue for usage-based ISPs. This additional revenue effect is ignored in this calculation.

³² We'll show some actual transit price quotes from a previous research study. You can also look at the Boardwatch annual survey for this data also.

³³ Anonymous source claims GeoNet was acquired by Level 3 primarily because of its rich peering relationships, and NextLink purchased Concentric for the same reason.

³⁴ Conversation with Nigel Titley (Level3) where they calculated the peering cost savings to Level 3 in the millions.

³⁵ Price quotes from interview with Dave Rand early 1999.

³⁶ The circuit prices we used were taken directly from a carrier in the Ashburn, VA area. We were quoted an OC-3 for \$11,400/month and OC-12 for \$23,000/month.

³⁷ See the Interconnection Strategies for ISPs white paper describing

³⁸ Lots of assumptions here: each ISP will save 10% if traffic volume is symmetric.

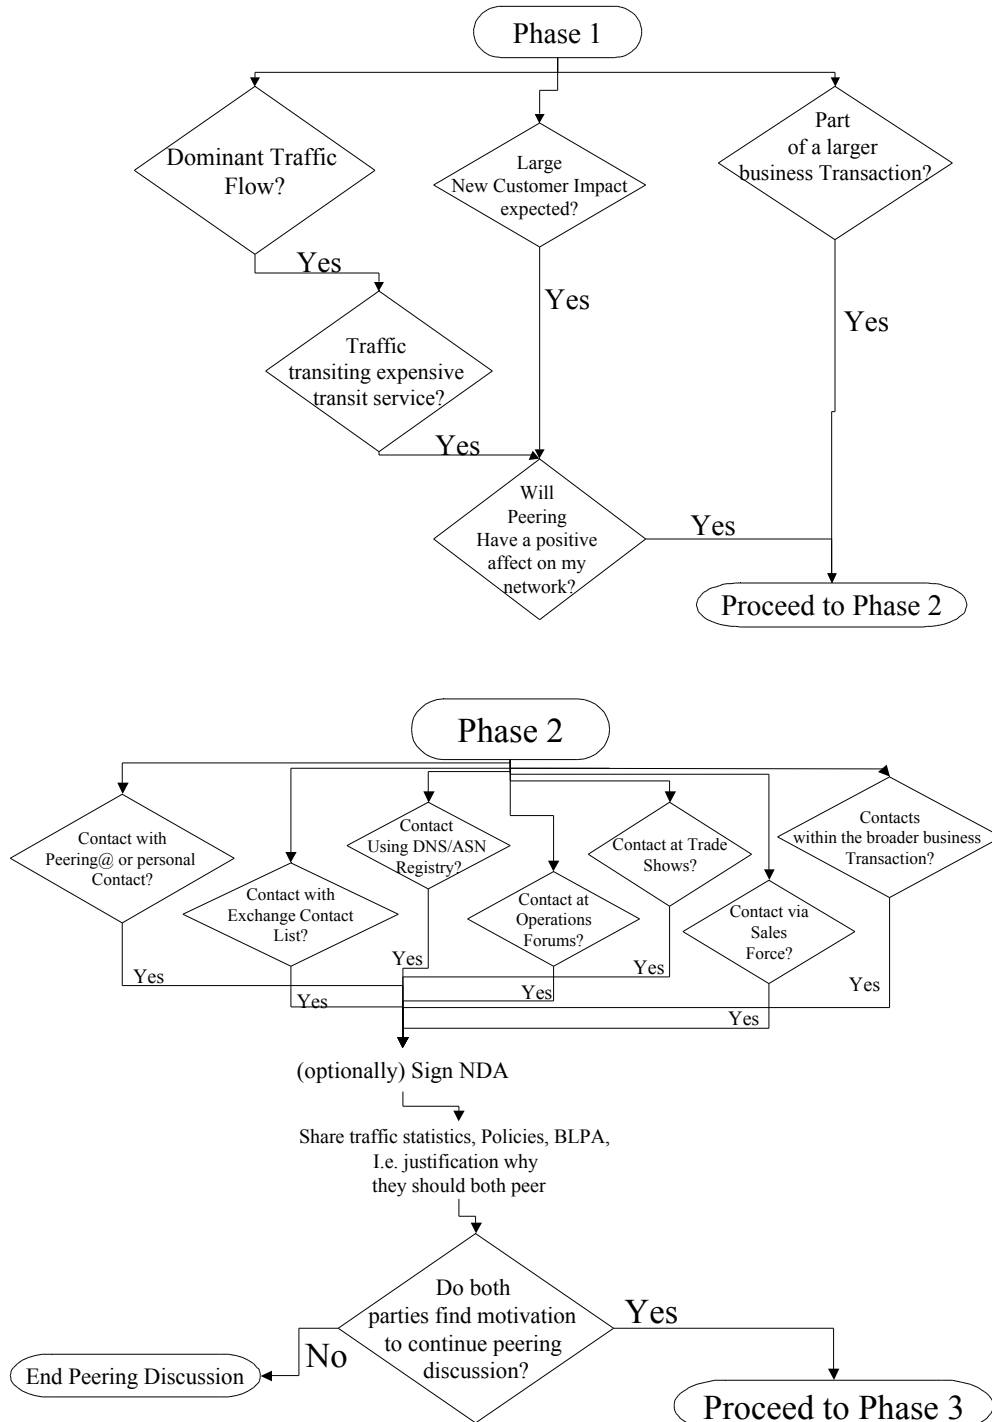
Appendix B – European Peering Differences

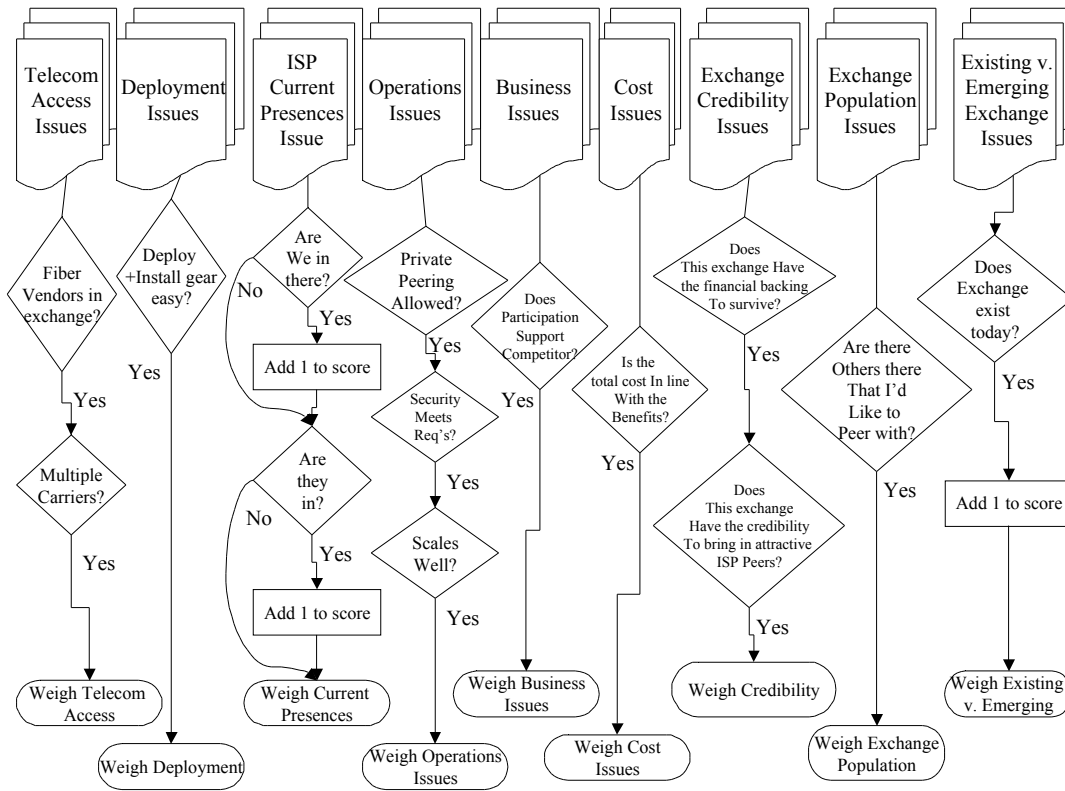
During research for this paper we found several distinct differences between the US and Europe in terms of peering in practice. For example, rather than only peering or transit, the Europeans more commonly bought or sold partial transit.

Interestingly enough, the peering wasn't based upon Europeans gaining partial transit access to the US, but rather European ISPs getting partial transit access to other European cities! It was seen as far cheaper to buy transit to European cities from a single provider already expanded into those areas than to build across international boundaries themselves. Long haul circuits and fiber to the US and establishing peering on the East coast with other US players was seen as relatively easy.

Peering	Paid Peering	Partial Transit	Full Transit
\$0	\$	\$\$	\$\$\$
Customer Routes Only	Customer Routes Only	Customer & Regional Routes	All Routes

Appendix B: Peering Decision Tree





Appendix C – Peering Simulation Game

Setting

In order to illustrate the strategic and financial role peering plays in an ISPs strategy, we created a peering simulation game³⁹. In this game, four ISPs (A, B, C, D) seek to maximize their revenues and minimize their costs. The **revenues** are determined by the number of regions or “squares” they occupy representing their market coverage and a quantum of content traffic (revenue) that market generates. The **costs** are determined by the number of squares that others occupy, representing the transit expense to access the rest of the Internet. The game board is shown below.

The Board




Figure 11 - The Peering Simulation Game

Play. Each ISP rolls the die and selects the number of squares indicated by the die, building into the exchange points desired. For each square occupied,

write your name in the square and collect \$2,000 transit revenue. The ISP must then pay its upstream transit provider (shown as Transit Provider X or Transit Provider Y around the border of the board) \$1,000 for each square the **other** ISPs own. The ISP fills in the score card and (if at an exchange point) can proceed to the peering negotiation stage.

Peering Negotiation. ISPs can reduce their *transit* costs by building into an exchange point and *peering* with the other ISPs there. If both ISPs agree to *peer*, the transit costs to the other peer’s squares are eliminated. (Both ISPs’ transit costs are reduced by the number of squares the other ISP occupies). However, the ISPs must collectively cover the cost of peering (\$2,000 per round and two lost turns), split however they see fit. (This is the peering negotiation.)

Objective

Generate as much revenue as possible by growing your network and establish peering or transit relationships to reduce network costs. Play is ended when any player can no longer execute a play (when the board is filled up). The winner is the player with the most money at the end of the game.

Variations

Transit Sales Negotiation (Optional). ISPs can buy/sell transit to each other at a reduced rate of \$500/square. In the transit sale, the transit provider gets the \$500/square transit revenue and the transit purchaser saves \$500/square (compared with buying transit at \$1000/square). The cost of transit (\$2,000 per round and two lost turns) is identical to the cost of peering and is split however the ISPs negotiate.

Merger and Acquisition (Optional). ISPs can agree to pool their interests and merge into a single ISP. There is functionally no difference in play except money can flow between the players and the new

0 0 0 ROUND	D C B A PLAYER	Roll (Copy down)	# Squares Owned	Revenue (Squares * \$2000)	#OthersSquares	Transit Cost (\$1000)	Peering Costs	Net	Running
	A	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
	B	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
	C	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
	D	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
copy	A	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
copy	B	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
copy	C	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
copy	D	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	\$0
1 A	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	(\$1,000)	
1 B	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	(\$1,000)	
1 C	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	(\$1,000)	
1 D	FALSE	1	\$2,000	3	(\$3,000)	\$0	(\$1,000)	(\$1,000)	

XpeerY (clear contents)	PLAYER	Pay for Transit to A?	Pay for Transit to B?	Pay for Transit to C?	Pay for Transit to D?	Sum of Transit \$\$\$ paid to X	Sum of Transit \$\$\$ Paid to Y
A						\$0	
B						\$0	
C						\$0	
D						\$0	
A						\$0	
B						\$0	
C						\$0	
D						\$0	
A						\$ (3,000)	
B						\$ (6,000)	
C						\$ (3,000)	
D						\$ (6,000)	

³⁹ First played at the Interconnect Billing and Accounting

merged company gets two turns. Transit fees must be paid until the two ISPs peer.

How this simulation is different from peering reality:

1. The board is veiled allowing for gaming and bluffing during peering negotiations
2. ISPs move serially in the game, while in the real world action is parallel.
3. The meaning of the board squares is severely overloaded to mean regional coverage and corresponding revenue, a quantum of traffic generated, and a quantum of traffic transitted to all others. All customers are not equal in revenue, traffic.
4. Customer transit revenue gained does not cause any additional financial load for the ISP in the game.
5. Traffic quantum is a vague notion that ignores the asymmetric nature of traffic today.
6. Shared squares should cause revenue and costs to be divided
7. Everyone starts with the same number of squares.
8. Everyone is financially backed to support infinite periods of financial loss. Well, that may reflect reality for some period of Internet time.
9. If ISPs fail to peer they must pay transit to get access to these squares. In reality, content multi-homes allowing alternative paths to the same content.
10. Business motivations to sell transit instead of peer are an ignored dynamic in the game.

Summary

The basic peering game does a good job of highlighting the issues ISPs face when peering. Several comments from ISPs offering enhancements add reality to the game at the cost of complexity. For example, ISPs capture market share in order to be an attractive acquisition target. Adding merger rules adds a real complexity, somewhat tangentially related to peering and transit. Adding rules for ISPs to buy/sell transit to each other similarly adds complexity but adds a negotiating dynamic that ISPs face today. Balancing the desire to explain and explore against the desire for the simulation to match reality has proven to be a challenge.