

14.4 Traffic Characterization in Light of Shortest-Exit Routing

Whether the peering connection is implemented at a shared interconnection (public peering) point or through direct interconnection (private peering), the prevailing pattern in the Internet today is to perform routing between peers on the basis of *shortest exit*, sometimes called *hot-potato routing*.

Shortest exit is, from a routing technology point of view, fairly simple. If a pair of ISPs or NSPs peer in more than one location, the *sender* determines the interface over which to send the data. The sender will generally choose to send the data at the earliest possible opportunity, in order to minimize cost to itself (while maximizing cost to the provider with which it peers).

This resulting system is, overall, fair to both parties under most circumstances; however, the economic implications of shortest-exit routing may not be immediately obvious to the casual observer. For that matter, the trade press has been exceptionally confused when it comes to shortest exit.

Suppose, for example, that a Massachusetts GTE customer chooses to access a Web page maintained by a Los Angeles customer of, say, Sprint. What happens? As we work through this example, pay close attention to Figure 14-4. The GTE customer selects the URL that he or she wants, and presses Enter. The customer's PC system will determine the IP address of the destination and will then start to open a TCP connection by sending what is

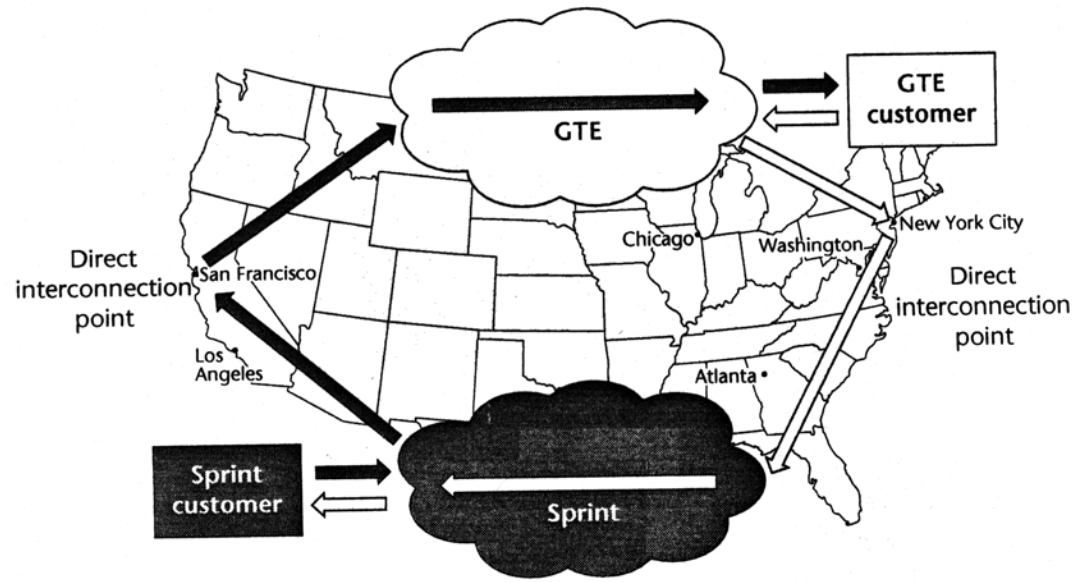


Figure 14-4: An example of shortest-exit routing.

called a SYN packet to the destination host. Routers within the user's company will recognize that this IP address is external to the company and will route it to their Internet provider, GTE (following the light gray lines). GTE's routers will, in turn, recognize that the IP address was assigned by Sprint and will route the datagram to the nearest interconnection point with Sprint. Let's suppose that the interconnection point is located near New York City.

At this point, Sprint has the obligation of carrying the datagram across the country using its own infrastructure. When it arrives at a Sprint router in Los Angeles, the datagram is sent to the customer's router, which in turn routes it to the Sprint customer's Web server.

The Web server generates a TCP acknowledgment of the SYN packet. Routers within the customer's internal network send this datagram to Sprint (the dark gray line). Sprint routers recognize that this datagram is destined for a GTE customer, so Sprint routes it to the nearest interconnect point with GTE, which, for our example, might be at Palo Alto, California. The traffic is handed off to GTE at that point and then carried on GTE's network back to Massachusetts, where it is finally handed off to the GTE customer who originated the request.

This somewhat arcane system has some subtle implications.

- Traffic flows are generally asymmetric.
- The system tends to be fair as long as data volumes are roughly balanced between transmit and receive, with no systematic bias toward or away from any particular geographic location, and as long as both providers are hauling data comparable distances.
- It is indeed more blessed to give than to receive: better to transmit data than to receive it. The recipient of traffic is burdened with the cost of hauling it for long distances. Shortest exit thus tends to favor providers with Web hosting traffic (asymmetric flows favoring transmission) and to put dialup providers (asymmetric flows favoring reception) at an economic disadvantage.
- There are innumerable ways to "game" the system—to so structure your business as to artificially reduce your costs, at the expense of other providers.

In light of shortest exit, large backbone ISPs tend to offer peering privileges at no cost only to other backbone ISPs that can peer in locations on both coasts and that can provide sufficient bandwidth between the coasts, subject to various other technical and business considerations. These represent necessary, but not sufficient, conditions for shortest-exit peering to be

reasonably equitable. And it is for these same reasons that backbone ISPs either do not offer peering or else insist on charging for peering for small ISPs who can peer in only one location—the small ISP cannot meaningfully reciprocate the service that the large backbone ISP provides, because the distances over which it operates are not comparable.

14.5 International Internet Traffic Flows

Not surprisingly, the NSF's view of the evolution of the Internet was focused primarily on its evolution within the United States. What about the rest of the world?

For historical reasons, the Internet evolved in a U.S.-centric way. It started in the United States, which today still represents the vast majority of the traffic on the Internet, although the rate of growth in other parts of the world may be even higher than that in the United States.

The prevailing tendency until quite recently has been for each foreign ISP to maintain a connection to a U.S.-based backbone ISP, at the foreign provider's expense. Thus, connectivity to the rest of the world was provided by the United States—a system sometimes called “hub and spokes.” This is depicted in Figure 14-5. From the perspective of foreign providers, the system could be quite irritating; it implied that traffic to any other provider's customers would generally go to and from the United States, even if both customers in question are located in, say, Europe.

Moreover, the foreign provider would pay for the circuit to connect to the U.S. NSP. Those circuits are expensive! A T-1 or an E-1 circuit across the Atlantic can cost almost as much as a T-3 across the entire continental United States. It is not surprising that foreign ISPs have been unhappy with this system.

It has been suggested that U.S. backbone ISPs should pay for half of the cost of circuits to foreign providers; this, however, is clearly a wrongheaded notion. The distribution of market forces does not support this distribution of cost; were the situation otherwise, it would already be in effect.

NAP-like public interconnects have appeared in a number of European locations and in Kobe, Japan. These interconnects can provide local concentration of traffic. During 1998, traffic interchange in Europe increased markedly, and it became increasingly common for traffic between European providers to be exchanged within Europe.

Today, the vast majority of Internet content is based in the United States. This reflects the reality that, in many parts of the world, circuits between two adjacent countries may cost nearly as much as a circuit from either country

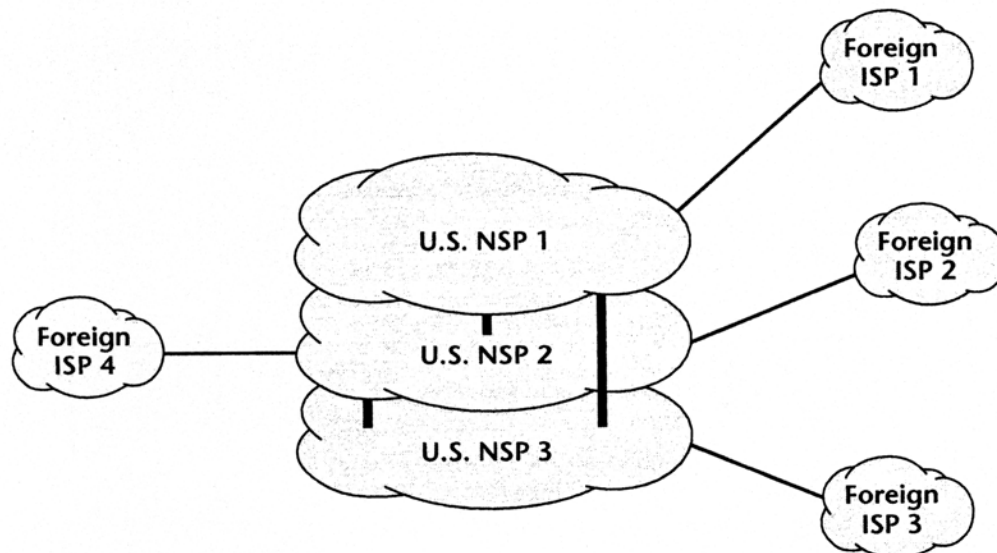


Figure 14-5: The "hub and spokes" system whereby foreign ISPs connect to a U.S. NSP.

to the United States. As a result, if a foreign firm wishes to place its content in a single location that will provide optimal international visibility at minimum costs, that location will usually be in the United States, *even if the target audience is primarily foreign*. Circuit prices in Europe have begun to plummet even faster than transoceanic circuit prices over the past year as a result of deregulation. As a result, these economics have begun to reverse, and the system is visibly starting to right itself.

In 1999, backbone providers began to offer global Internet access to foreign ISPs from POPs in major overseas markets. The price of these wholesale services is considerably less than that of a comparable transoceanic circuit. These services are likely to lessen the cost disparity between U.S. and foreign ISPs.

14.6 Traffic Statistics

It is natural to want to model traffic flows through the Internet as a whole, particularly when we find that we are not getting the performance that we would like. Unfortunately, at present, very little global data is being captured, and the quality of what exists is uncertain.

The Router Arbiter project developed a number of statistics about the Internet as a whole. Of particular interest were measurements based on a tool called NetNow, which attempted to characterize delay and packet loss

through the backbones of various backbone ISPs. The definition of the associated metrics has since been taken up by the *IP performance metrics (IPPM)* activity of the IETF.

A number of other statistics-gathering initiatives are in the pipe. It is not clear that *any* of the existing or emerging studies will generate statistical data that is both valid and useful. Partly, the technical problems are daunting; partly, the decentralized nature of the Internet, as well as the legitimate desire of ISPs to protect their proprietary data, make it difficult to capture useful and meaningful overall statistics.

For the foreseeable future, we will all be working largely in the blind as regards the characteristics of the Internet as a whole. Individual backbone ISPs will, in some instances, have good data about their own networks, and may in some cases be able to draw valid inferences about the Internet as a whole. Nobody, however, will have good, comprehensive data about the system as a whole.

14.7 Internet Access

The typical access to the Internet happens in one of two ways: either dialup access, including ISDN access, as is commonly used by residential consumers, or dedicated access by means of leased lines or Frame Relay. Dialup access is implemented by using terminal servers. For dedicated access, a few additional considerations are worth noting.

The most common dedicated access is implemented over a leased line, as shown in Figure 14-6, using a CSU/DSU and a router or, equivalently, a router with an integrated CSU/DSU.

As a customer, you should carefully consider the placement of this Internet connection. In general, you will want to minimize the cost of access (for instance, a leased line from the LEC) to the provider's facilities. But if you are operating a large Web hosting operation, for example, it is likely to be important to situate the Web servers at a point where your provider has good connectivity to other providers, bearing in mind that most of the traffic is probably destined for customers of other ISPs, since no single ISP has a majority of the market as a whole.

In fact, it is for this reason that you may wish to consider *Web hosting* services, where the provider houses the Web server at the provider's own premises. This enables the Web server to gain very direct access to the ISP's infrastructure, which can be highly advantageous; however, it may also imply more distance from the server to the rest of the customer's facilities. Whether

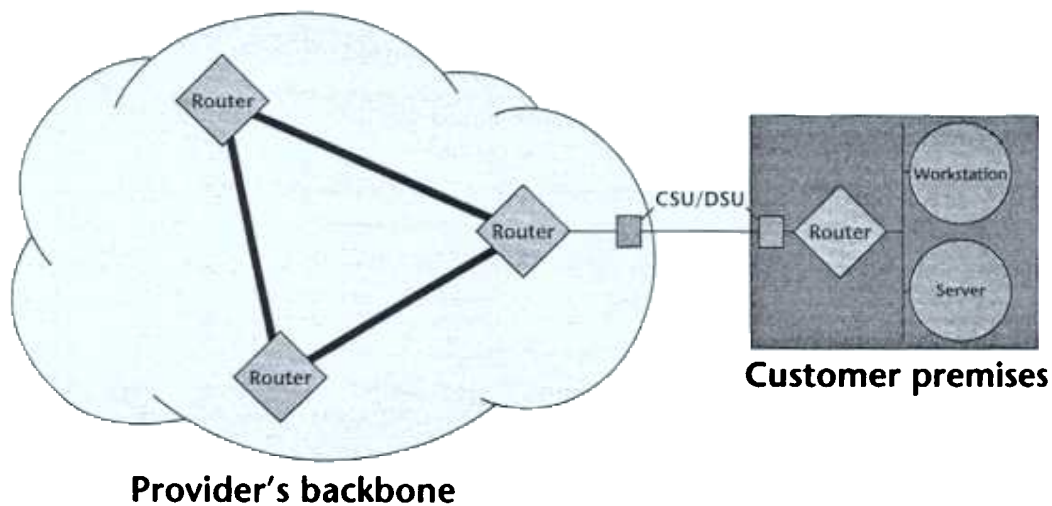


Figure 14-6: A typical dedicated Internet connection.

this is a good thing or not needs to be evaluated carefully in light of the design of the customer's underlying applications.

A single line from the ISP's router to your router can represent a single point of failure. Depending on the nature of your applications, this may or may not be acceptable. There are various alternatives for connecting to your ISP over multiple links or at multiple locations. These may differ from ISP to ISP. Some of these provide good throughput but limited failover capability; others provide good failover but do not provide maximum effective capacity. This is a choice that you would need to make in careful consultation with your ISP.

You may also wish to consider maintaining connections to multiple ISPs. In general, doing so will necessitate that you run BGP-4 exterior routing protocols, a major increase in complexity.

BGP-4 is a very powerful and flexible protocol, but you are nonetheless likely to find that some things become more difficult as a result of using it. Load sharing across multiple links, in particular, can require significant fiddling to get right. Relatively few people in the industry have experience with advanced applications of BGP-4, and many of them already work for ISPs. In general, then, you are likely to want connections to multiple ISPs only after careful consideration of the pros and cons, and only if your organization has a strong skill base in IP routing.

For residential consumers and *Small Office—Home Office (SOHO)* use, dialup Internet access is usually much more cost-effective than leased line

access. Many dialup providers offer service at a fixed monthly rate, irrespective of the amount of traffic you generate or the number of hours for which you are connected. There has been a trend over the past few years for some providers to charge for hours beyond some fairly large fixed threshold, largely as a means of discouraging customers from staying continuously connected for the entire month; for those users who never reach the threshold, these services are still effectively flat rate. In the United States, it is often possible to find an ISP that provides service that is within your local calling radius and that is, therefore, free of per minute charges from the LEC under current regulatory policy.

In all cases, when you select a dialup ISP, you should consider how many hours a month you are likely to use the service and whether it provides local access with no per minute charges from the LEC. You would also want to consider the quality of the services provided; however, this will be largely subjective or anecdotal, since there are, in my opinion, no sound, objective publicly available comparative measures of dialup ISP quality. The various surveys of various ISPs are at best suggestive, not definitive.

You should also consider high-speed access alternatives available in the consumer space, particularly if you are a "power user," someone who makes heavy use of the Internet. These high-speed services include ISDN, ADSL, and IP over cable. Make inquiries to determine the actual throughput capacity of each service. For the service to provide several megabits per second of access speed to your home means nothing to you if the service provider's access to the Internet constrains you to a few kilobits per second of actual throughput, as is often the case. Considerations for ISDN are very similar to those for dialup: Local calling radius is important. In all three cases, you would have to contact your cable provider, LEC, or ISP to determine whether service is available to your home or office; even if ADSL coverage is available in your area, your particular line might not be suitable. If the service is available, with ISDN or ADSL, you will generally be free to choose from among several ISPs; with IP over cable, however, you may, under current regulatory policy and cable industry business practice, typically be constrained to use the cable provider's "captive" ISP (RoadRunner for Time Warner, @Home for TCI). Several municipalities and state governments have challenged the cable providers' right to impose these exclusive arrangements, but there is as yet no consumer protection at the national level to ensure freedom of choice of ISP for cable IP customers.