14 The Public Internet: Unique Design Considerations

Much of our discussion up to this point has focused on design for private intranets and enterprise internetworks. How do things change when we design for the global, public Internet? How might the design of the public Internet interact with design considerations for enterprise internetworks?

The first six sections of this chapter deal with the overall design of the Internet, explaining how that system has evolved over time. They will be of particular interest to anyone designing for an *Internet service provider*, particularly for anyone designing a network for a *backbone ISP*, but they have other applications as well. For instance, understanding the structure of the Internet helps a business to order the most appropriate Internet access service and to position its Web servers and other sources of high traffic volume. Thus, the final section of this chapter contains helpful hints for ordering Internet access from an ISP.

14.1 Evolution of Today's Internet

Today, there is no overall central planning function for the Internet. Changes set in motion a few years ago have led to the present system. In order to understand how the system works, and how it was intended to work, it is helpful to return to the year 1993—the last point in time at which a grand vision existed for the evolution of the Internet as a whole.

Until 1995, the U.S. National Science Foundation provided the core of the Internet in the form of the NSFnet, a publicly funded backbone network. The NSFnet was restricted to carrying only traffic associated with research or education. By 1993, the NSF felt that the Internet had sufficient commercial viability to stand on its own two feet. Moreover, the NSF understood that as long as the government provided the backbone of the Internet at no cost, there could never be a broad private Internet industry. Private industry could not compete with a service provided at zero cost. In order to permit an industry to spring into being, the National Science Foundation would have to step aside.

The NSF initiated this transition through a public solicitation document, Solicitation 93-52, in which the NSF described a future Internet structure comprising

- *National service providers (NSPs),* providers of Internet access operating on a national scale
- Internet service providers (ISPs), operating on a smaller scale than NSPs
- Public shared interconnection points called *network access points,* or *NAPs*
- A transit service to interconnect the regional ISPs formerly funded by the NSF
- A *very high speed backbone network service (vBNS)* to interconnect a handful of research laboratories and to serve as a testbed for new high-speed Internet services
- A *router arbiter* service, both to rationalize routing in the global Internet and to provide a central anchor point for routing research and statistics

Figure 14-1 shows a notional view of what the NSF had in mind.

On balance, the NSF was very successful in taking itself out of the innards of the global Internet without massive disruption. However, some aspects of its plan were more successful than others. It is instructive to look at the results of each element of the plan.

14.1.1 National Service Providers and Internet Service Providers

A number of NSPs and a far larger number of ISPs were already in operation at the time of Solicitation 93-52. Today, there are at least five major backbone ISPs and somewhere between six and perhaps thirty smaller backbone ISPs that could be described as NSPs. We will discuss the distinguishing hallmarks of a backbone ISP later in this chapter.

By all accounts, there are several thousand ISPs today. *Boardwatch Magazine* currently lists some 4,855 ISPs in North America.[1]

The NSF did not attempt to create NSPs and ISPs; rather, in exiting the public Internet space itself, it created the market forces that would foster the formation of NSPs and ISPs.

1. *Boardwatch Magazine,* http://boardwatch.internet.com/isp/spring99/introduction.html, June 1999.

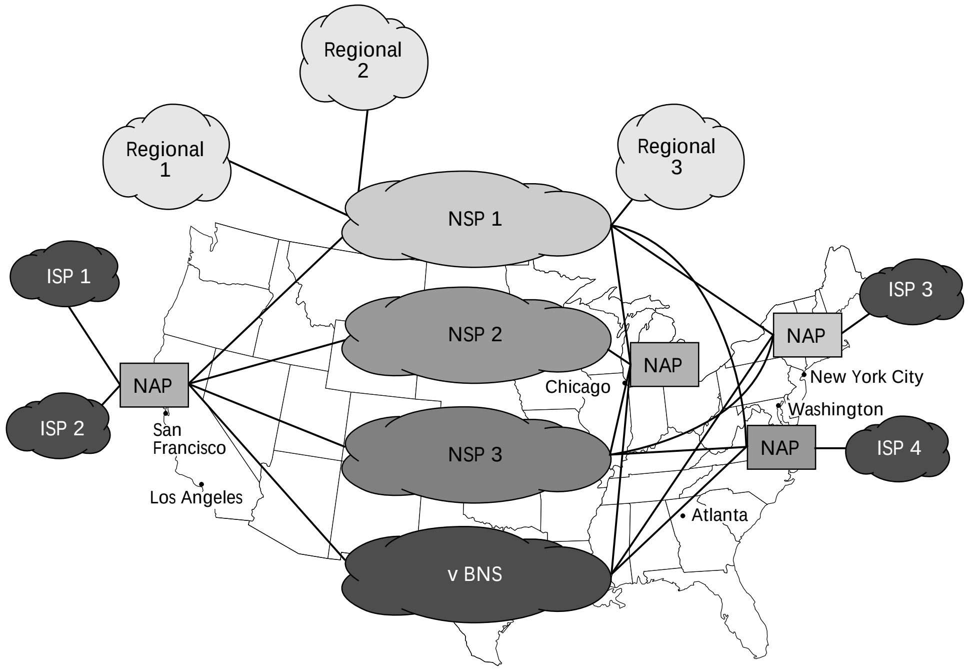

Figure 14-1: The architecture of the Internet, according to NSF Solicitation 93-52.

14.1.2 Network Access Points

Solicitation 93-52 sought to create at least three network access points, or NAPs, located in the San Francisco Bay area, the Chicago area, and the New York City area, as well as additional NAPs to the extent that funding might permit. The NAPs were expected to operate at the data link layer of the OSI Reference Model; they might be implemented as a shared high-speed LAN or as an SMDS or ATM service providing shared interconnection among NSPs/ISPs.

Each NAP would provide interconnection among NSPs and ISPs for purposes of traffic interchange—known in the trade as *shared interconnection,* and sometimes referred to as *public peering.* In an *interconnection,* or *peering,* relationship, Internet providers agree to accept traffic destined for one another's respective customers but not necessarily for other third parties. In a *customer,* or *transit,* relationship between Internet providers, by contrast, the transit provider typically agrees to accept traffic destined for any point in the global Internet.

The solicitation resulted in the NSF's supporting four NAPs. Those NAPs, along with their sponsoring firms, are shown in Table 14-1.

Table 14-1: Network Access Points and Sponsoring Firms

NAP	Sponsoring Firms
Chicago	Ameritech/BellCore
New York[1]	Sprint/San Diego Supercomputer Center
San Francisco	PacTel/BellCore
Washington, D.C.	Metropolitan Fiber Systems (MFS), in a facility called MAE East

1. New York is something of a misnomer: Sprint responded to the NSF's request with a proposal to locate its NAP in a facility situated some 90 miles southwest of New York, across the Delaware River from Philadelphia.

In the flurry of acquisitions that followed the Telecommunications Act of 1996, the identity of a number of these organizations is changing or has changed. BellCore has been acquired by SAIC; PacTel has been acquired by SBC; Ameritech is in the process of being acquired by SBC; and MFS was acquired by WorldCom, which subsequently became MCI WorldCom as a result of its acquisition of MCI.

The NSF's program was basically successful, even though things did not work out exactly as initially envisioned. The Washington NAP, better known as MAE East, and the Sprint NAP evolved into full-fledged national and global traffic interchange points; the Chicago and San Francisco NAPs, however, have been problematic since their inception. Initially, they suffered from technical problems associated with attempting to drive large volumes of IP traffic over still immature ATM switching products; subsequently, these NAPs have lacked a critical mass of large national backbone ISPs. As a result, they function, in practice, in the role of regional traffic concentration points. Quite a few additional regional public interchange points have come into existence in the United States in recent years, although they do not appear to play a strong role in the overall traffic flows of the Internet.

Meanwhile, the lack of a stable public peering point on the West Coast of the United States was, briefly, problematic for the NSP/ISP community. In practice, the large national backbone ISPs soon converged on a facility called MAE West, operated jointly by MFS and NASA.

As a result, the work originally envisioned for the NAPs can be viewed, for most purposes, as being performed by MAE East, MAE West, and the Sprint NAP. All three of these were based in the recent past on Fiber Distributed Data Interface (FDDI) LAN technology, and all three augmented the FDDI with high-speed LAN switches (gigaswitches). The gigaswitches support full-duplex FDDI, thus offering, in theory, 100Mbps of input and output simultaneously between each pair of routers.

By 1996, all of these facilities were suffering from inadequate capacity. In 1998, MCI WorldCom upgraded its MAE facilities in Washington (MAE East), San Jose (MAE West), and Dallas to offer modern ATM switches as a high-capacity alternative to the FDDI/gigaswitch architecture.

As we shall soon see, a largely unanticipated consequence of the evolution of the Internet away from the pure NSF 93-52 model has been a migration of interconnection traffic away from these three locations and into private arrangements among the NSPs. These *direct interconnections* are sometimes referred to as *private peering.*

14.1.3 Transit Service among Regional ISPs

Eight large regional ISPs went through a competitive bidding process, which resulted in a series of contracts for *transit* service—where transit represents carrying of data to other ISPs—being awarded to MCI. MCI connected the NSF-sponsored regional ISPs to one another and carried their traffic to other NSPs and to the NAPs. The arrangement worked well, in general, but it has largely been phased out today. The original regional ISPs have outgrown their need for transit services; however, most backbone ISPs now offer transit services, which continue to play a huge role in the public Internet today.

14.1.4 The Very High Speed Backbone Network Service (vBNS)

The NSF also awarded the vBNS to MCI. MCI runs routers over an infrastructure of ATM switches to interconnect a number of NSF-sponsored research institutions, including the supercomputer centers at Cornell University, Pittsburgh, San Diego, NCAR, and NCSA. The vBNS was initially operated at OC-3 speeds (155Mbps) and was subsequently upgraded to OC-12 (622Mbps).

The program has generally worked as intended; nonetheless, there are issues. First, it serves only a portion of the institutions to which connectivity is offered. (Thus, it is not available to all of Cornell University). Second, the original intent of using the vBNS as a testbed for research into high-speed internetworking may be inappropriate for a production network used to accomplish real work.

More recently, some of the functions of the vBNS have been subsumed by two newer initiatives: the *Next-Generation Internet (NGI),* and the *Abilene* project of *Internet 2.* The NGI is a U.S. government–sponsored initiative to provide a high-speed private internetwork among a number of major U.S.

agencies. Abilene is a high-speed private internetwork interconnecting members of UCAID, a consortium of research universities and private industry. Abilene has benefited from significant "in-kind" contributions of circuits and equipment from Qwest, Cisco, and Nortel.

Meanwhile, the original vBNS contract will soon expire. MCI WorldCom intends to replace it with a privatized next-generation vBNS network based on packet-over-SONET at OC-48 speeds.

14.1.5 The Router Arbiter (RA) Project

The Router Arbiter project was intended to conduct research into routing and to provide a database of Internet topology and policies, in order to enhance the stability, robustness, and manageability of the Internet. The Router Arbiter project also produced statistics about the Internet as a whole (see Section 14.6).

14.2 Structure of the Internet Today

Traffic interchange among backbone ISPs is fundamental to the operation of this system. As previously noted, in an *interconnection* relationship, Internet providers agree to accept traffic destined for one another's respective customers but not necessarily for other third parties. This is different from a *customer*, or *transit*, relationship between Internet providers, where the transit provider typically agrees to accept traffic destined for any point in the global Internet.

Historically, interconnection was called *peering*, in order to imply that traffic was interchanged among providers that were similar in size and capability. Over time, it came to be recognized that peers need not be similar in size; rather, what was important was that there be comparable value in the *traffic interchanged*.

I find it convenient to think of today's Internet as comprising two kinds of ISP: backbone ISPs and all others. The backbone ISPs interconnect with all other significant ISPs by means of a full set of interconnection relationships. Other ISPs may have some interconnection relationships, or they may not, but they have a significant dependency on a customer or transit relationship to one or more backbone ISPs. (Thus, the ability to reach all Internet destinations without the need for a transit relationship—sometimes called, somewhat inexactly, *default-free* status—is a strong indicator that an ISP should be viewed as a backbone ISP.) This yields the complex, Medusa-like structure shown in Figure 14-2.

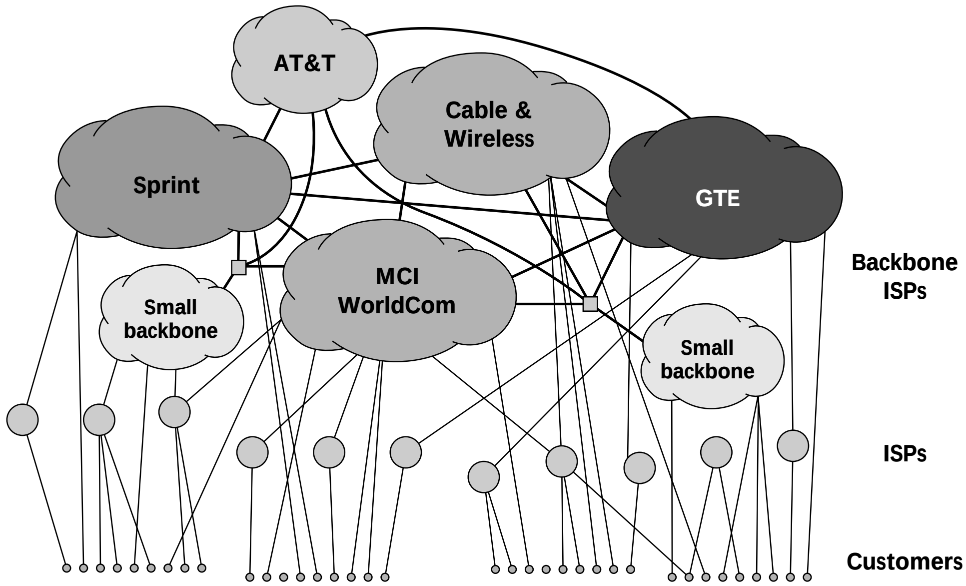

Figure 14-2: Present-day Internet structure: Backbone ISPs and other ISPs.

By these criteria, Cable & Wireless (formerly internetMCI), MCI WorldCom (including UUnet, ANS, and other previously independent ISPs), Sprint, AT&T, and GTE (including the former BBN and Genuity) should clearly be viewed today as backbone ISPs. Somewhere between six and perhaps thirty other ISPs could also be viewed as backbone ISPs. The vast remainder are dependent on the backbone ISPs for global interconnectivity.

14.3 Direct and Shared Interconnections (Public and Private Peering)

Over the past few years, there has been a marked trend away from the use of the shared interconnection points on the part of the large backbone ISPs. In absolute terms, there is still significant traffic growth at MAE East, MAE West, and the Sprint NAP; however, the growth is not commensurate with the growth in overall traffic in the Internet. Thus, these facilities are losing “market share” in terms of the number of bits that flow through them.

The trade press has been perplexed by this trend and has occasionally presented it as if it were a predatory tactic on the part of the large backbone ISPs. In fact, simple economic and technical considerations drive the move to direct interconnections (private peering). For that matter, the migration to private peering is, for the most part, neutral in its impact on smaller ISPs—which indeed have severe problems with the affordability of interconnections but not because of the upsurge in direct interconnections among large backbone ISPs.

Even our way of thinking about shared interconnection has evolved. Historically, we referred to the NAPs and MAEs as *public peering points*. This was something of a misnomer, as, in almost all instances, peering was not public! Interconnection (peering) at a shared interconnection point would be established as a bilateral business relationship between two backbone ISPs, often at no direct cost to either party, based on the shared perception that both parties benefited from that interconnection.

In 1993, it might have seemed that the NAP-based architecture shown in Figure 14-1 could be expanded indefinitely. As each component was outgrown, it would simply be upgraded. When the NAP shoe started to chafe, the NAP provider might simply buy a shoe of larger size.

Things have not played out quite the way that they were expected to. First, LAN technology did not keep pace with the growth in Internet traffic levels. As traffic at the three major interconnects grew, the shared-medium FDDI (with 100Mbps of bandwidth to *all* ISPs present) was upgraded several times but never seemed to be able to keep pace with traffic. More recently, the largest MAE facilities have migrated to ATM; nonetheless, the main reason that the Internet as a whole has not long since collapsed under its own weight is that the major backbone ISPs have diverted most of their traffic away from the shared interconnection points.

The second main factor in this migration is that the economics of interconnection work differently as the bandwidth ramps up. Under the original concept of a NAP, a single T-3 connection from each NSP would carry all of the peering traffic for that NSP for a major portion of the United States. The NAPs made economic sense because they allowed multiple peers to share resources. They also provided economies of scale, because operating a single shared T-3 circuit was far more cost-effective than operating one or more T-1 circuits to each of several peers.

Today, if a given backbone ISP has enough traffic to another backbone ISP to fill a T-3 circuit, there is no incentive to use a shared interconnect for that traffic. Indeed, in light of current Internet traffic levels, use of a public interconnect

- Would introduce the risk of overloading the shared interconnect, to the detriment of both backbone ISPs and possibly also of third parties

- Offers no additional economic incentives in the form of economies of scale
- Potentially complicates future upgrades, in that three parties are involved (two backbone ISPs and the shared interconnect facility manager) instead of two (just the two backbone ISPs)

For a large backbone ISP, the natural tendency is to use

- Private interconnects to those backbone ISPs with which one has a lot of traffic to interchange
- Shared interconnects to interchange traffic with smaller backbone ISPs and, possibly, with small ISPs

This yields a system that looks like Figure 14-3. The majority of interconnections among providers may continue to take place at shared interconnection points; however, the preponderance of traffic already flows across private peering interfaces, and this tendency is sure to accelerate in the coming years.

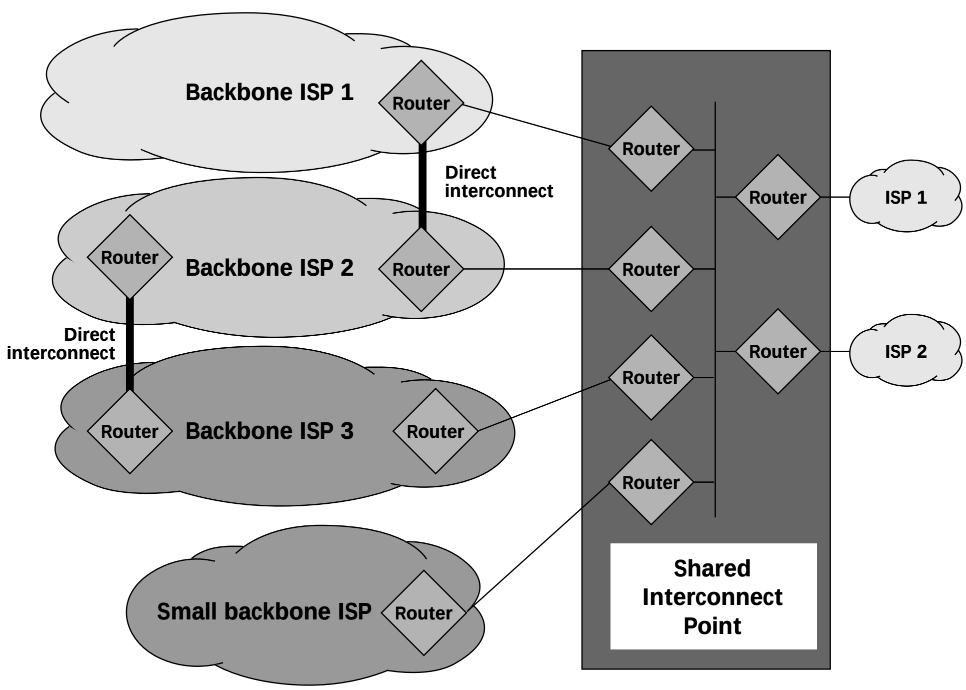

Figure 14-3: Direct and shared interconnections (private and public peering).

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

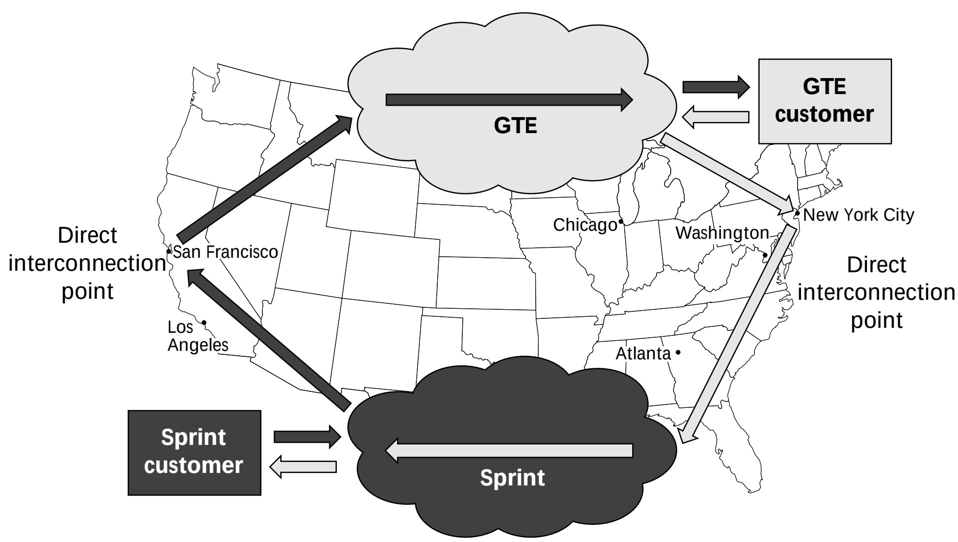

Figure 14-4: An example of shortest-exit routing.

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

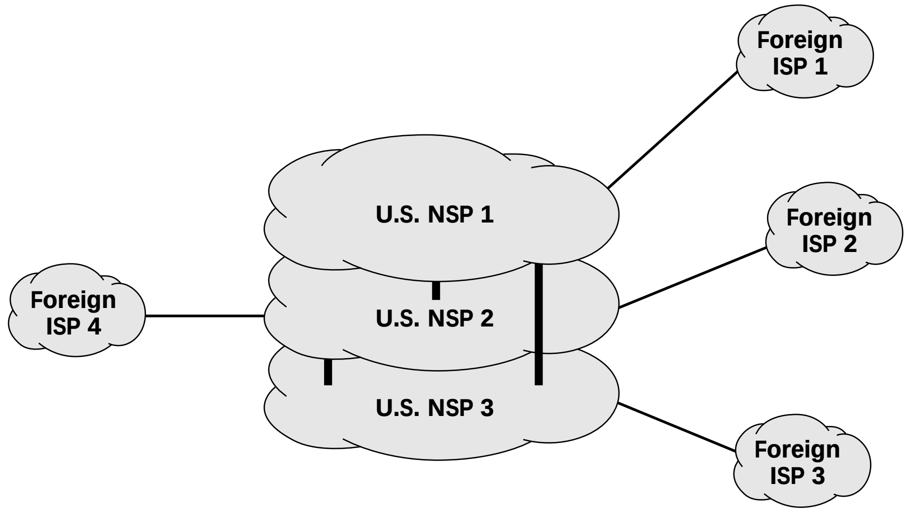

Figure 14-5: The "hub and spokes" system whereby foreign ISPs connect to a U.S. NSP.

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

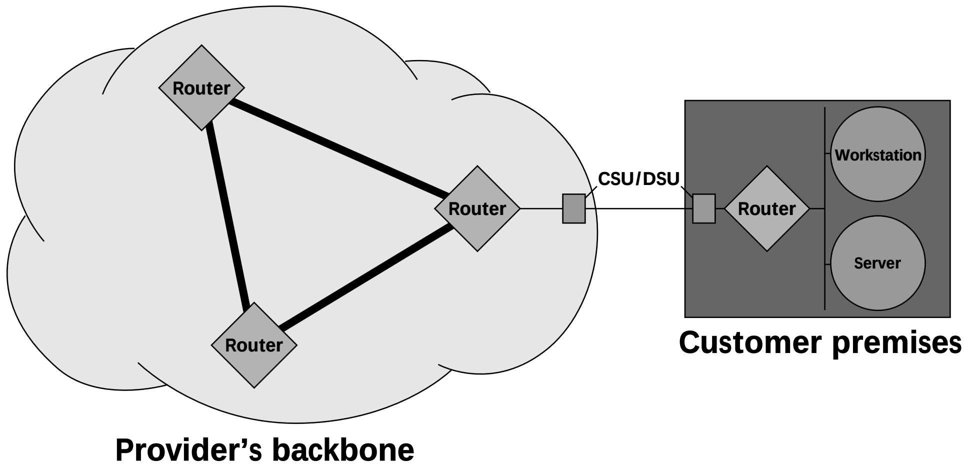

Figure 14-6: A typical dedicated Internet connection.

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.