Establishing a Presence on the Internet

Getting Connected

The Internet at 56K and Up

O'REILLY®

Kevin Dowd

Getting Connected: The Internet at 56K and Up
by Kevin Dowd

Published by O'Reilly & Associates, Inc., 101 Morris Street, Sebastopol, CA 95472.

Editor: Mike Loukides

Production Editor: Jane Ellin

Printing History:

June 1996: First Edition

This book is printed on acid-free paper with 85% recycled content, 15% post-consumer waste. O'Reilly & Associates is committed to using paper with the highest recycled content available consistent with high quality.

ISBN: 1-56592-154-2 (domestic)
1-56592-203-4 (international)

[5/99]

Preface

I have had the privilege of witnessing many people's first Internet experiences. It is as if they're trapped in a darkened basement, clinging to a candle. There are cobwebs in the rafters, and plumbing going every which way. And all the while their guide is gushing, "Isn't this the greatest thing you've ever seen?!"

Often, the guide is a book. Whole shelves are dedicated to exploration, browsers, and do's and don'ts; they talk about the Internet as a playground. This book is different—it's about the plumbing. We'll cover the gory details—how the network works, how to maintain a full-time connection, Internet security, even wide area networking protocols.

If you're the person who will be bringing Internet connectivity to a business, school, or club, then you have probably already realized that it is going to be a multidisciplinary challenge. You will have to be part businessperson, part electrician, and you will have to crawl over networking software like a spider eating its mate. Then, if your Internet connection gets breached by crackers, there's a chance that your company will lop your head off. And lastly, there's the dreaded end-user problem; you will be answering silly questions forever. In the short term, you are going to have to figure out how to shop for a connection, prepare a server, and secure your site. Longer term, you will need to set up information services, glue your electronic mail systems together, and extend Internet connectivity to the desktop.

Still, it's great fun. One rarely gets to undertake projects where the success rate is high, the visibility is high, and the rewards are good. If you have been nominated for the job, then good for you. If not, you might consider championing Internet connectivity. Your company is going to be connected sooner or later. You may as well be the person who gets the credit.

Be forewarned, however: there are prerequisites. It will take perseverance and talent to get the job done. I'm not so interested in your Internet experience; Internet connectivity may be completely new to you. However, I *am* going to assume that you have a natural affinity for things technical, even though your strengths may be in areas other than Internet connectivity—PC LANs or software development, for instance.

Most technical Internet books build from the bits upwards. They start with a general discussion of packet switching, then carefully build up through IP addressing, gateways, DNS, and specific Internet services. Several excellent books that take this approach are available from O'Reilly & Associates, including *TCP/IP Network Administration*, by Craig Hunt, *DNS and Bind*, by Paul Albitz and Cricket Liu, and *Managing Internet Information Services*, by Liu, Jerry Peek, and friends. We're going to do something different, however. We are going to start with function, and work both sideways and downward.

You will find that the book has several personalities. Some sections—particularly near the front of the book—are lighter reading, approachable by all audiences. Depending on your background, the material may already be familiar to you. On the other hand, you may find welcome insight into aspects of the Internet that you assumed you understood thoroughly. Later chapters are fairly technical. In a few spots, we pursue related topics just because they are interesting.

I know that priorities change: what's really important to you today may not be so important tomorrow. Accordingly, I have tried to keep all chapters modular.

Chapter 1 provides a framework for bringing in an Internet connection, and answers some common questions about the Internet's funding, maintenance, and structure. In Chapter 2, we explore service options and attempt to understand how "big" a connection you will need.

Chapter 3 helps you critically analyze and compare Internet Service Providers (ISPs), and helps you understand contract options and fees. In Chapter 4, we investigate circuit and equipment needs.

Chapter 5 discusses the task of programming a router and provides a few examples. Chapter 6 provides some explanations of ISO and IP networking as background information for programming the router's serial interface.

Chapters 7 through 12 provide technical background and practical configuration advice for a variety of physical and link layer network technologies. These chapters are important because your Internet service will be provided using a combination of these technologies—Frame Relay, ATM, SMDS, X.25, PPP, SLIP, or HDLC. Chapter 13 briefly discusses routing. Chapter 14 explains physical network technologies—DDS (56K), ISDN, T1, and more. Chapter 15 talks about circuit turn-up and debugging.

Chapter 16 addresses security, system audits, and Internet firewalls in detail. Everyone connecting to the Internet will want to become familiar with this material. For those who are comfortable with UNIX and C, the chapter goes on to discuss inhouse firewall construction.

Chapters 17 and 18 cover key infrastructure pieces: domain name service and electronic mail. We will look at baseline configurations, plus configurations for operation behind a firewall. We will also talk about marrying Internet services to inhouse mail systems, such as Lotus cc:mail and Microsoft Mail.

Chapter 19 discusses IP connectivity for Macs and PCs. We will learn how to configure IP stacks for a variety of configurations, and explore dynamic address assignment.

My deepest hope is that you will say: "Ah! This is the book I have been looking for." It's the result of years of work, and years spent connecting organizations to the Internet. It's also the book that *I* was always looking for. I feel lucky to be the person who got to write it.

Conventions

Italics	used for emphasis, as well as to denote file, directory, domain, and command names. Also used for *ftp* and Web sites.
Bold	used in code examples and in discussion to show user-typed input.
`Constant Width`	used to set off indented examples of computer output, code, or, in discussion, to refer back to code. Character strings such as "a" are set in quotation marks.

Acknowledgments

Unbounded thanks to my wife, Paula, for putting up with the lost weekends and non-vacations that went into this book. And congratulations to her for bearing a son, Ian, during the time the book was being written. Of course, Ian too, *in utero*, missed a few weekends and a vacation ... Also thanks to my parents, family, and in-laws.

Great thanks to Mike Loukides for his extraordinary abilities as an editor. Thanks too to Tim O'Reilly for his faith and patience. I'd also like to thank my co-worker Eric (Rick) Romkey for reviewing the material as it was extruded. My thanks to reviewers, including Bob Tinkelman, Eric Raymond, Howard Burke, Ed Dorso, John Curran, and Eric Pearce, plus other facilitators, including John Crouse and

Darryl Pope. And of course, thanks to the crew at O'Reilly & Associates, including Chris Reilley, who did some incredible illustrations (again); Edie Freedman for the cover; Jane Ellin for the production work; Lenny Muellner for Tools support; Clairemarie Fisher O'Leary for doing quality control, organizing freelancers, and helping with Appendix G; Nancy Priest and Eden Reiner for their work on the design; Kismet McDonough-Chan for quality control; Seth Maislin for writing the index; David Futato, Michael Deutsch, and Barbara Coughlin for their work on Appendix G; Nicole Gipson Arigo for stepping in at the last minute to help finish the production work; and Sheryl Avruch for making sure everything ran smoothly.

Table of Contents

In this chapter:

- ***An Internet Connection***
- ***Timeline***
- ***How Much Is This Going to Cost Me?***
- ***Where Do We Go From Here?***

1

Getting Connected to the Internet

Say that over lunch I shake my tuna sandwich at you and remark, "Fred, our companies work together often. We should join networks so that we can share information." That afternoon you lean out the window and toss the tail end of your network backbone across to me, I pull it through the bathroom window, and join it to our network through a device called a *router* (or, alternatively, a *gateway*). Traffic from one building, bound for the other, passes through the router. The rest remains on its respective sides. By joining our networks in this way, we have formed a small *internet*.

Other organizations could connect up too, adding routers at convenient spots. Say, for instance, that Joe in the building across the way also "gateways" to you—this time through the cafeteria window. Now, in order for me to reach Joe's network, I will have to transit your network (passing first through my bathroom, and then through your kitchen). There may even be several paths between my network and Joe's, some of which don't involve you at all. It depends on how the networks join together.

Like the wires dangling from the story above, the *Internet* is a network of networks—just a whole lot bigger. It spans the globe, reaching hundreds of thousands of individual networks, and several million computers. All of them speak *Internet Protocol* (IP), developed for DARPA (Defense Advanced Research Projects Administration) in the late 1970s. IP is very flexible; the internetwork topology can be anything you like—a straight line, a tree, a god-awful tangle. IP routing protocols help networks "discover" one another, without help from humans. The same qualities also make IP networks robust; if a backhoe cuts through one portion, the rest of the network automatically updates its routing tables to avoid the dead sections.

Ownership and management of the Internet—like the Internet itself—is noncentralized; no single organization is in charge. Observers and participants sometimes like to tease themselves that there's a recreational chaos to the whole thing, but there really isn't. The Internet runs pretty well, especially considering its size and complexity. It works because there is agreement about how it should work, because there are organizations that watch out for its well-being, and because there is money pouring into it all the time.

An Internet Connection

I guess I don't hear the question quite as often these days, but people used to ask: "Isn't the Internet free?" The misconception came from the fact that governments around the world have historically funded large tracts of the network with tax dollars; for some people, access *was* (or is) free because they were shielded from the costs. Furthermore, there is a tremendous amount of information available for the taking—again, free. However, the Internet is not free; it is largely a commercial operation. *Somebody* pays for it. Typically, it is you, the person who orders Internet service.

Who do you pay? Your provider might be a national or international *Internet Service Provider* (ISP), such as PSI, UUNET, BBN, Sprint, MCI, etc.* Each of these companies operates a wide area network, of which you become a part when you sign up for service. You can reach everyone else on the Internet because these companies trade traffic with their competition via routers, or high speed networks, tucked into wiring closets and sprinkled about the Internet.

You might, on the other hand, connect through a local Internet service provider. Like the national firms, a local provider maintains a wide area network, which grows to include you when you sign up for service. Technically, local providers are once removed from the "centers" of the Internet; they purchase bandwidth from one of the national providers, and resell it. On the surface, this would appear to put them at a disadvantage because their bandwidth is twice warmed-over. However, good service is a function of more than the speed of an ISP's backbone. Network loading, the quality of the company, and network operations support all play a part.

In addition to an ISP, you will also need to hire an exchange carrier—a telephone company. They will operate the physical (and possibly virtual) network that brings routed Internet service to your organization. If your exchange carrier is an ordinary local phone company, then your physical network will most likely be an ordinary leased line of 56 Kbit/S or higher capacity. It may also be something

* See Appendix B, *Calculating Your Bandwidth Requirements.*

Internetworking the Commercial Internet

Some of the main trunks of the Internet—both in the United States and abroad—are funded by national governments in support of specific kinds of activities, such as research and education. These networks (MILNET, NSI, ESnet, and NORDUnet, to name a few) operate under "acceptable use policies," which restrict commercial activities. Other portions of the Internet bear no commercial restrictions. These networks are private ventures, paid for by their subscribers, often for the express purpose of conducting business. Both sections—publicly funded and privately funded—are interconnected. The trick has been to try to keep commercial traffic off the public sections of the Internet.

In 1991, a handful of network carriers joined together to form the Commercial Internet Exchange Association (CIX), a nonprofit corporation devoted to commercial Internet cooperation. At the time, the National Science Foundation (NSF) backbone was the main Internet trunk. The idea was that by expressly routing commercial traffic through the CIX exchange point(s) (there are two at present, in Santa Clara, CA, and in Washington, D.C.), commercial organizations could have visibility in publicly funded networks without passing an undue amount of commercial traffic through them. CIX now includes 177 member networks.

Other North American (non-CIX) exchange points exist as well, including MAE-West (San Jose), MAE-East (Washington, D.C.), and some upcoming MAE locations—all operated by MFS Datanet. Additionally, there is an exchange point called SWAB, operated by Bell Atlantic (Washington, D.C.), and Federal Exchange Points (FIX-West and FIX-East), located at NASA Ames, and College Park, Maryland, respectively.

The NSF backbone is being reworked as well. In 1994, MCI received a National Science Foundation contract to operate an OC-3 (155 Mbit/S) backbone, connecting supercomputer centers and four regional *Network Access Points* (NAPs), located in New York, Chicago, San Francisco, and Washington, D.C. The NAPs are managed by Sprint, Ameritech and Bellcore, Pacific Bell and Bellcore, and MFS Datanet, respectively.

European regional network exchange points include Ebone (continent), LINX (UK), and several *Distributed Global Information Exchanges* (D-GIXes), located in Stockholm and Paris. NSF, Sprint, and NORDUnet have sponsored transatlantic traffic, tying regional European networks and exchange points together with North America.

more exotic—ISDN, or perhaps optical fiber. In many cases, you won't hire the exchange carrier directly; Internet service providers often pass the charges through on your bill. Sometimes (as in the case of Sprint or MCI, for instance), the ISP and exchange carrier are the same company.

Responsibility for the rest of the pieces—equipment, day-to-day support, security, Internet applications—can fall on either side of the fence; they can be your problems, or they can belong to the provider. Depending on who you ultimately decide to go with, you may be able to pick from an assortment of options, ranging from complete outsourcing of installation, security, and maintenance, to a roll-your-own kind of connection. Usually, it comes down to a question of which parts you want to pay for, which parts you are willing to manage on your own, and how much you are willing to pay.

Timeline

There's a lot to do, but as with any large project, a little bit of planning will make an overwhelming list of tasks feel manageable. I can't tell you *exactly* how much time will pass from the moment when you decide to go ahead with an Internet connection to the moment when you puff a fat cigar and pull up the first Web page. But I am going to guess it will be eight weeks, maybe less.

Week 0

At the outset, you need to identify the Internet applications people will use, estimate bandwidth requirements, understand the equipment needed, evaluate Internet service providers, and pick service options. Of course, you will want to choose a domain name too (e.g., *atlantic.com*).

It will also be a good idea to think about your security needs before proceeding; uncertainty about Internet security has torpedoed more than a few otherwise enthusiastic Internet connectivity plans. Will you need a firewall? If so, will you construct it, buy it, or hire somebody to secure your site for you? Most organizations can't proceed until these questions are answered.

Week 2

Once you have your plans, you will need to make a contract with an Internet service provider, and possibly hire an exchange carrier for the physical connection. If you are in charge of purchasing your own data communications equipment, you should get that under way too; some brands of routers have long manufacturer delivery times.

It will take a while for the circuit to come in—probably from four to six weeks. In the meantime, you should probably get started on your security plans, prepare a server (or servers) (for DNS, SMTP mail, the Web, FTP, and so on), and plan for desktop access to the Internet (provided you haven't outsourced these pieces). This is also a good time to think about electronic mail integration; how will you glue Internet mail into your current email system?

Week 6

Tempus fugit. By the sixth week, you should be sitting on your hands and swinging your feet, bored. At long last, you will be visited by the phone company, come to install the physical circuit. They will pull a line into your wiring closet, screw a network termination unit into the wall, and disappear. At this point, you may wish to connect the data communications equipment, and give your Internet service provider a call. Once the physical circuit is in, the service provider should be able to give you a definite date for circuit "turn-up," probably within a few days.

Week 7

During the seventh week, you and the Internet service provider will rendezvous on the telephone, to bring up the Internet connection. Your data communications equipment should be powered up and functional. Your router should be set to go: addresses, routing, and endpoint identifiers programmed in advance. Likewise, your server hosting DNS, SMTP mail, etc., should be configured and ready for all comers.

Once the connection is "up," the clock starts ticking; your Internet service provider will begin charging for the service.

Week 8

A few details usually linger another week or so. If you are getting a news feed, for instance, it may be several days before activation. This, too, may require a phone call with the service provider. Likewise, DNS usually needs some tweaking after circuit turn-up. Particularly, you may need to register reverse (address-to-name translation) zones for your network, or you may need to migrate responsibility for the primary nameserver for your domain away from your ISP, and over to a machine on your network.

How Much Is This Going to Cost Me?

In the chapters that follow, we will expand considerably on the scenarios you might want to investigate—do-it-yourself connectivity, long distance connections, connectivity with the right to redistribute bandwidth, firewalls. In the meantime, I am going to assume that you are interested in a simple full-time connection with nearby access, a secured router, and a modest server. I'll also assume that you plan to bring the connection up yourself (hence, this book). I'm putting my head in the sand a little; there are true costs to labor, of course. As a benchmark, let me say that if you hired me, I would probably charge you between $6,000 and $10,000 for the labor covered in this book. Your time has similar value.

There are some one-time and some recurring components associated with your connection. The one-time components include hardware: a CSU/DSU or multiplexor, a router, and a server. There will also be one-time connection set-up costs, charged by both your Internet service provider and your exchange carrier. The breakdowns can be a little confusing; some service providers charge you individually for every little portion of the connection—the cost of the circuit, rental of a port on their router (called a port charge), cost for the routed Internet traffic, a PVC charge (for a virtual circuit). I'll try to assemble these into the big picture.

Representative Costs

Let's assume that you will connect to an access provider with a *Point of Presence* (POP) nearby. Accordingly, the physical circuit will involve your local telephone company alone. Long-haul connections, by comparison, can involve both your local phone company and an interexchange carrier; there can be several additional components added to the bill (we discuss circuit types and charges in more detail in Chapter 4, *The Circuit*).

Table 1-1: Representative Connection Costs

	56 Kbit/S FR	T1 Dedicated	10 Mbit/S SMDS	T3 ATM NAP
CSU/DSU or Mux	$ 250	$ 1,000	$ 3,000	$ 5,000
Router	$ 1,800	$ 1,800	$ 2,500	$ 6,000
Server	$ 2,500	$ 2,500	$ 3,500	$ 10,000
Transport One-time	$ 800	$ 1,200	$ 3,000	$ 5,000
Transport Recurring (annual)	$ 1,500	$ 4,800	$36,000	$ 25,000
Internet Service One-time	$ 1,200	$ 2,500	$ 3,500	$ 5,000
Internet Recurring (annual)	$ 4,800	$18,000	$30,000	$ 60,000
Year One	$12,850	$31,800	$81,500	$116,000

Table 1-1 suggests some possible first-year costs for a variety of connection types.* The 56 Kbit/S Frame Relay charges, shown in the first column, are typically fixed, without regard to distance. T1 (1.54 Mbit/S) dedicated access (column 2), on the other hand, is usually mileage-sensitive; the farther away you are from the POP, the more it will cost. T1-speed connectivity may be available to you in Frame Relay too, though T1 Frame Relay circuits can cost more than dedicated connections for nearby destinations.

The last two columns represent connections that aren't terribly common, though you *can* buy them now, and you will probably see more of them in the next few years. SMDS, shown in the third column, is a high-speed, metropolitan area network technology. As with Frame Relay, the pricing assumes that you are within a geographical area served by the network, and will pay a flat-rate tariff. I show a 10 Mbit/S price estimate; however, speeds for SMDS networking can go as high as 45 Mbit/S and beyond. The connection we are talking about here is pretty fast; compare with a vanilla Ethernet network—also rated at 10 Mbit/S.

DS3 (T3) connectivity is shown in the last column. Surprisingly, a dedicated T3 (45 Mbit/S) or OC-1 circuit doesn't cost proportionally more than a dedicated T1, given the great difference in performance. The estimate shown for T3 Internet service presumes service through an NSF NAP. Other varieties of T3 access are available too, particularly through a few of the national ISPs.

You will note that I included the cost of an Internet server in the table. This would be a machine that could act as a mail hub, news server, DNS server, or even a firewall. Naturally, the server horsepower should match the kind of traffic you expect. You might find it surprising, however, that for everything through T1 connectivity, a high-end PC running UNIX makes a fully capable server.† You will also note that I bumped up the server cost estimates for both the SMDS and ATM scenarios. This is because I imagine that you will want more capable hardware servicing the end of a very high-speed connection. In fact, the notion of a single server is fairly ludicrous; if you are buying a DS3 circuit, you probably have some grand plans for the bandwidth—perhaps involving a whole computer center.

There can be fees beyond those in the table, including charges for domain name service or a news feed, for instance. Likewise, there may be fees for managing data communications equipment, or providing site administration. Other costs involve the desktop environment and a dial-up strategy for members of your organization when they are out on the road. You will also want to choose a set of

* Note: these are 1996 guesstimates. Over time, they can be expected to drop, with hardware costs receding most quickly, and service/support costs lagging behind.

† Consider a high-end PC running *Linux*. Linux is a very capable UNIX rewrite, available at no charge. Everything you might need in a server has already been ported and tested under Linux. We use it all the time.