Chapter 2

Key Technology Trends

"Hope is like a country road. At first there was no path. But after many people walked on it, it was created."

-- Lu Xun

"Everything that's going to happen in 25 years, it's already happening, somewhere."

-- Bob Metcalfe, Ethernet inventor

This report is a foundation for a discussion of management issues and industry impacts, this chapter reviews key trends in Internet technology that are most relevant to retail electronic financial services.

To understand the technical capabilities of the Internet as a platform for electronic finance, it is helpful to recall that -- Silicon billionaires notwithstanding – it started out as one of the world's most successful *not-for-profit* experiments. It had its roots, not in visionary risk-taking by entrepreneurs or strategic initiatives by the likes of Microsoft, IBM, or Netscape, but in a series of highly innovative experiments that were undertaken largely for their own sake by researchers, academics, and government bureaucrats.

Today, key features of the Internet's architecture – from its highly-distributed architecture to its bias toward open standards -- still derive from these non-commercial roots. Much of the critical software that runs the Internet is "shareware," and its dominant applications -- e-mail, chat, browsing, and newsgroups -- were all originally designed for non-business use. A key factor responsible for its rapid takeoff, flat-rate pricing, is also rooted in this non-profit service model.

However, as we saw in Chapter I, all this is now changing rapidly, as the Internet evolves from a free range to a vast expanse of valuable commercial real estate. Some leading Internet experts have voiced doubts about whether its basic architecture is really up to the task of providing reliable commercial service for millions of users. Of course financial service companies also have a long history of preferring to build their own proprietary networks and dial-up platforms, in part because they are also especially concerned about reliable, secure service.

In the last two years there have also been some rather dramatic service interruptions on the Internet, and many other individual cases where subscribers have experienced congestion when trying to log in. In the most recent episode, in late October 1997, thousands of customers for brokerage services like e: Trade and e: Schwab were unable to access their accounts at a critical time. Such episodes, while brief, have reminded many users that the Internet is certainly not yet as reliable as the nation's telephone network – which has also been known to have service interruptions.

Finally, in the last decade there have also been several widely-reported examples of Internet security breaches by outside intruders, and several well-publicized warnings by government agencies about "cyber-terrorism." While the actual incidence of such behavior appears to be very low, these reports have also helped to raise public concerns about the Internet's vulnerability to snoopers, swindlers, and vandals."

Together, these concerns make it sensible for those of us who want to understand the potential of Internet-based financial services to start out by reassuring ourselves that the Internet is indeed technically capable of becoming a reliable services platform. Along the way, we will also glance at some of the technology "upsides" that may also make it a high-value, highly competitive services platform.

Overall, on the basis of interviews with more than fifty leading experts in Internet technology and services, we have concluded that while the Internet does indeed have many unsolved problems, most of them are well-understood, and there do not appear to be any fundamental technical barriers to the widespread deployment of robust Internet financial services in the next 5-7 years.

In fact, if anything, a veritable tidal wave of new technologies is likely to arrive in this timeframe. This will make it possible to offer much *better* service -- security, reliability, decision support, accuracy, timeliness, interactivity, and ease of use at lower cost -- than other existing financial channels do today.

On the other hand, making sense of all this new technology and applying it is likely to be a serious management challenge for many financial institutions, and the technology itself is likely to have several important impacts on the overall industry. Chapter III will examine these technology management challenges and industry impacts more closely.

I. Overview – Key Technology Issues

As noted in Chapter I, the Internet's extraordinary growth -- more than 100 percent a year so far this decade -- has occurred despite the fact that much of its fundamental technology is still under construction. To understand how its platform is likely to bear up under continued growth, it is helpful to organize our discussion around the "stack" approach summarized in **Figure 2.1**. While there are many technology trends like improved CPU speeds and new programming languages that lift productivity across the board, the stack disaggregates technology changes into three basic groups – those that primarily affect network services and connectivity, those that affect access devices and interfaces that users actually see, and those that mainly affect our ability to combine networks and endpoints into useful solutions.

Figure 2.1 State of the Art Internet Technology "Stack"

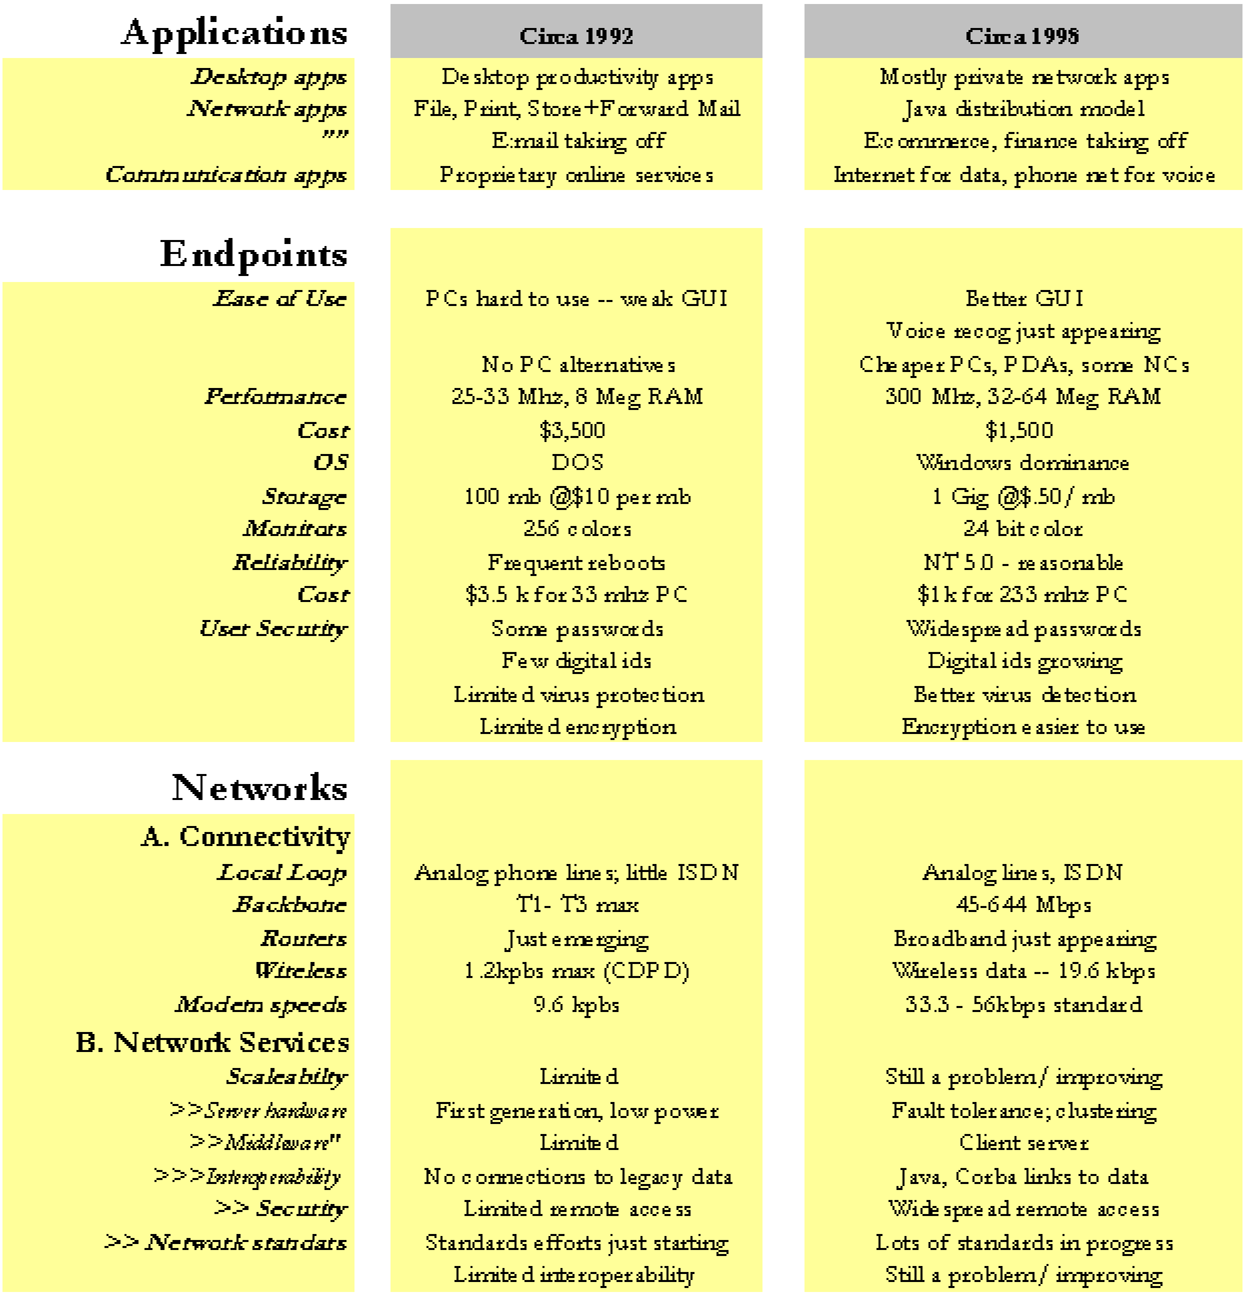

Applications	Circa 1992	Circa 1998
Desktop apps	Desktop productivity apps	Mostly private network apps
Network apps	File, Print, Store+Forward Mail	Java distribution model
""	E:mail taking off	Ecommerce, finance taking off
Communication apps	Proprietary online services	Internet for data, phone net for voice
Endpoints		
Ease of Use	PCs hard to use -- weak GUI	Better GUI
		Voice recog just appearing
	No PC alternatives	Cheaper PCs, PDAs, some NCs
Performance	25-33 Mhz, 8 Meg RAM	300 Mhz, 32-64 Meg RAM
Cost	$3,500	$1,500
OS	DOS	Windows dominance
Storage	100 mb @$10 per mb	1 Gig @$.50/ mb
Monitors	256 colors	24 bit color
Reliability	Frequent reboots	NT 5.0 - reasonable
Cost	$3.5 k for 33 mhz PC	$1k for 233 mhz PC
User Security	Some passwords	Widespread passwords
	Few digital ids	Digital ids growing
	Limited virus protection	Better virus detection
	Limited encryption	Encryption easier to use
Networks		
A. Connectivity		
Local Loop	Analog phone lines; little ISDN	Analog lines, ISDN
Backbone	T1- T3 max	45-644 Mbps
Routers	Just emerging	Broadband just appearing
Wireless	1.2kpbs max (CDPD)	Wireless data -- 19.6 kbps
Modem speeds	9.6 kpbs	33.3 - 56kbps standard
B. Network Services		
Scaleabilty	Limited	Still a problem/ improving
>>Server hardware	First generation, low power	Fault tolerance; clustering
>>Middleware"	Limited	Client server
>>>Interoperability	No connections to legacy data	Java, Corba links to data
>> Security	Limited remote access	Widespread remote access
>> Network standats	Standards efforts just starting	Lots of standards in progress
	Limited interoperability	Still a problem/ improving

It is helpful to remember, as **Figure 2.1** indicates, just how far most of the key elements in this technology stack have come in the *last* five years. As of 1993, for example, PCs were relatively costly, low-powered devices with limited storage, low-resolution monitors and weak built-in networking. Graphical user interfaces (GUIs) and built-in networking were just beginning to appear on PCs other than Apples. PCs had trouble networking even locally – by far the hottest growth segment in the software industry consisted of adding basic local area networking features like printing, file sharing, and e-mail to stand-alone PCs. For purposes of reaching "wide area" networks like the Internet, PCs were almost unbearable to use. Their communications software was complex, the best modems of the day were limited to 9600 baud connections, and endpoint devices were too slow to handle more than one application while on-line. Finally, all this cost at least 15 to 150 times as much as simpler, more reliable communication appliances, phones and fax machines.

At such low connection speeds, even when one did connect there was not much to do "online." It took forever just to send files and simple text messages. The leading online services of the day -- Prodigy, CompuServe, and tiny AOL – charged by the hour for clunky services that didn't inter-operate. They had a few hundred thousand subscribers at best. While the Internet had existed in a primitive state since at least late 1969, the easy-to-use search, browsing, and e-mail applications that we all now take for granted did not exist. Mosaic, the first Internet browser with a GUI, and the inspiration for Netscape's browser, was only introduced on Unix machines in February 1993. Internet service providers as such also did not exist. Internet access was available only to universities, government agencies, think-tanks, and a few network-centric businesses. The Internet's fiber backbone consisted leased T1 lines. No self-respecting telephone, cable, or electric utility showed any interest in upgrading their millions of wires and fiber connections to the home to deliver what was then quaintly referred to as "data communications services."

Meanwhile, most other large enterprises were having a hard time just installing and maintaining their own internal PC networks and integrating them with their mainframe data centers. No one was thinking seriously about using the Internet to connect to thousands of outside customers. As of 1993 there were less than 200,000 people in the world who even had Internet access. Accordingly, the kind of "public network services" – security, registration, authentication, directory, billing, and customer care – that phone companies take for granted were scarce even among the leading online services.

Fortunately, for a variety of reasons that would take us too far afield to explore here, society was decisively deflected toward a higher-growth path for Internet services. In the intervening half-decade, significant progress has been made toward relaxing many of the key technical constraints on the Internet's growth. Cheaper, more reliable endpoint devices are becoming available, network infrastructure is becoming much more robust, and ordinary retail customers are beginning to have more options for getting on-line – no longer just by way of analog phone lines and PCs, but also through cable TV set-top boxes, so-called "network computers," Web phones, and Web pagers.

This is not meant to be some kind of Media Lab sermon on the glories to come of the digital age. But, as discussed below, a closer look at each of layer of this technology stack indicates that on the horizon, even greater progress is likely to be made in the next 5-7 years. Indeed, the consensus of the experts that we polled is that we may be on the verge of a veritable "technology glut," with service providers, service designers, and customers all struggling to keep up with the changes and opportunities. In particular,

- Endpoint devices, user interfaces, intelligent agents, and network services like connectivity, security, and "middle-ware" are all likely to show dramatic improvements in this time frame, eliminating some of the most important remaining obstacles to the widespread deployment of Internet services.
- The technologies required for the mass deployment of new retail electronic payment systems, especially electronic billing, presentment, and electronic cash, will all mature to the point where they will be poised for dominance.
- Concerns about Internet security will be substantially alleviated. In fact, as discussed below,

the Internet is already well-positioned to be recognized as a *much more secure* place to do commerce and finance than traditional financial channels that are based on paper checks, 800-number operators, unencrypted credit-card authorizations, and bricks-and-mortar, and trusted restaurant waiters.

- The one remaining area for very serious concern is scalability. Hyper-exponential growth in network traffic is likely to put severe stresses on network software and infrastructure in the next few years. For financial institutions that plan to offer large-scale Internet services, this is an important, often-overlooked issue. Even here, though, significant technical progress is likely to be made.

The following sections take a closer look at each of these key trends.

II. The Customer Side – Key Trends

To begin with, on the customer side, the consensus of the experts that we interviewed is that in the next five years all the sharp improvements in the PC's costs, processing power, local storage, monitors, and multimedia capabilities are likely to continue, although at a lower rate. There are also several trends that will enable the deployment of new endpoint devices and applications. These trends, summarized in **Figure 2.2** and discussed below, will be especially important to retail financial service customers.

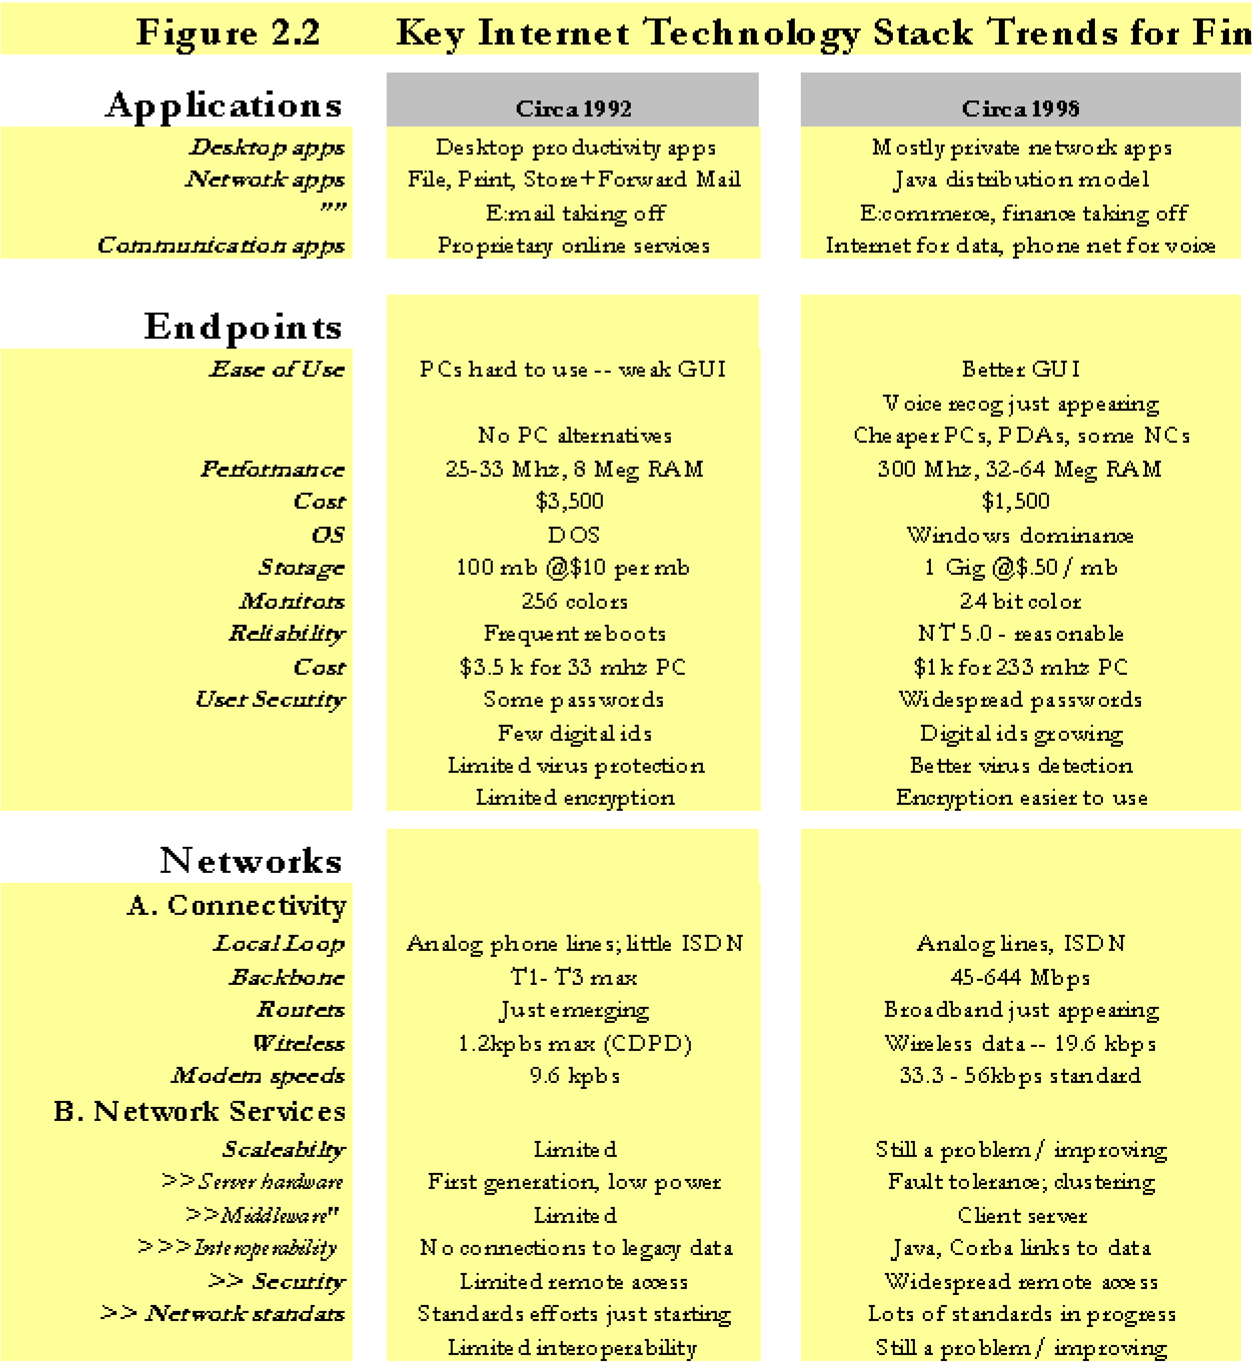

Figure 2.2 Key Internet Technology Stack Trends for Fin

Applications	**Circa 1992**	**Circa 1998**
Desktop apps	Desktop productivity apps	Mostly private network apps
Network apps	File, Print, Store+Forward Mail	Java distribution model
""	E-mail taking off	E-commerce, finance taking off
Communication apps	Proprietary online services	Internet for data, phone net for voice
Endpoints		
Ease of Use	PCs hard to use -- weak GUI	Better GUI
		Voice recog just appearing
	No PC alternatives	Cheaper PCs, PDAs, some NCs
Performance	25-33 Mhz, 8 Meg RAM	300 Mhz, 32-64 Meg RAM
Cost	$3,500	$1,500
OS	DOS	Windows dominance
Storage	100 mb @$10 per mb	1 Gig @$.50 / mb
Monitors	256 colors	24 bit color
Reliability	Frequent reboots	NT 5.0 - reasonable
Cost	$3.5 k for 33 mhz PC	$1k for 233 mhz PC
User Security	Some passwords	Widespread passwords
	Few digital ids	Digital ids growing
	Limited virus protection	Better virus detection
	Limited encryption	Encryption easier to use
Networks		
A. Connectivity		
Local Loop	Analog phone lines; little ISDN	Analog lines, ISDN
Backbone	T1- T3 max	45-644 Mbps
Routers	Just emerging	Broadband just appearing
Wireless	1.2kpbs max (CDPD)	Wireless data -- 19.6 kbps
Modem speeds	9.6 kpbs	33.3 - 56kbps standard
B. Network Services		
Scaleabilty	Limited	Still a problem / improving
>>Server hardware	First generation, low power	Fault tolerance; clustering
>>Middleware"	Limited	Client server
>>>Interoperability	No connections to legacy data	Java, Corba links to data
>> Security	Limited remote access	Widespread remote access
>> Network standars	Standards efforts just starting	Lots of standards in progress
	Limited interoperability	Still a problem / improving

New Endpoints

The first such trend is toward the proliferation of new low-cost, Internet-based access devices. An early step in this direction was the flurry of "personal digital assistants" (PDAs) that appeared in the last year, using "Windows CE™" software, a light-weight version of Windows, as an operating system.

But PDAs are essentially shrunken PCs. A more important development is a whole new breed of devices called *network appliances*. These appliances may not run conventional PC operating systems at all. Typically they have their own "embedded" operating systems that have been specifically designed to run on small ROM chipsets. They depend for their utility and low cost on the fact that --- like TV sets, radios, and telephones before them – they rely on an external network for applications and services.

It appears that the market for such appliances could be very large – at least the computer industry thinks so. In the last two years, the world's leading semiconductor manufacturers have added substantially to their processor chip capacity, preparing the way for a "processor glut," in addition to the "bandwidth glut" described later. Recent projections are that in the next 3-4 years, the number of Internet access devices – including such appliances as well as PCs -- will grow to more than 300 million, compared to just 32 million in 1996

All this bodes very well for consumers and providers of Internet services. Among the leading candidates for new endpoint devices are the following:

- **Internet TV**

One new family of devices may eventually make TV sets a pervasive outlet for Internet services. By adding digital modems and local memory to set-top boxes, cable operators are able to offer high-speed Internet access to their subscribers. There are already about 100,000 external "cable modems" in service at two dozen US locations. In the next five years, leading cable companies like TCI, Time Warner, US West Media, and Comcast will roll out several million more. Digital set-top boxes take the next step, integrating cable modem and digital TV functions into one device. They also leverage last year's mandate by the FCC for all US TV sets and TV broadcasting infrastructure to become "digital" by the year 2005. This will improve TV picture quality to PC resolution levels, providing much better reception of Internet content. A primitive version, introduced last year by WebTV, already provides Internet access over an ordinary phone line, using TV monitors for display.

Digital set-top boxes will eventually deliver Internet services at speeds up to 10 Kbps, about 40 times as fast as today's best analog modems, and 15 times faster than the most widely- used high speed telephone services, ISDN. Depending on the demand for such services, and what share of cable networks are upgraded to deliver them, they might eventually reach a significant portion of the seventy percent of US households that now subscribe to cable. Recent estimates are that in the next five years, cable modems and digital TV set-tops boxes may yield at least 3 to 7 million new Internet subscribers in the US, most of them from top-income groups.

- **Network information appliances** Another new class of endpoints focuses on making computers simpler, less costly, and more reliable. Last year, "network computers" were introduced by leading manufactures like IBM and Sun. Their aim was to significantly reduce the "total ownership cost" of PCs by eliminating local disk storage and relying on network servers or service providers to store files, data, and applications. This has special importance to large enterprises, where network hardware and software maintenance can easily cost $10,000 per user per year or more – up to 70 percent of total IT (information technology) operating costs.

But the concept of narrow-function devices is also relevant to the retail services market, where many computer users just want simpler appliances that let them get on line quickly and send e: mail or faxes, browse, make Internet voice calls, and check their bank balances and stock portfolios. From a service provider's standpoint, such devices also reduce the costs of customer service, attract new customer segments, and enable the more rapid deployment of new "thin-client" software applications and upgrades over the Internet, direct from the providers' servers.

As early as 1993, this desire for simpler endpoints and greater reliance on network services prompted Citibank, Phillips, and AT&T to experiment with special "screen-phone" terminals for home banking. These experiments turned out to be premature, because the Internet and several critical technology components simply weren't ready. By now, however, all the pieces that were missing have come together, and are taking the form of several new device categories:

- **Web screen-phones.** This is the closest thing to a reincarnation of the Phillips/Citibank screen-phone, updated to take advantage of better Internet access, more processing power, and better displays. A typical one – like the prototypes recently offered by Intel, Cidco, Mitsui, Intelidata, and Uniden – has its own IP address, a touch-screen menu, basic Internet e-mail, browser, and "push" applications, built-in modem, an online rolodex and calendar, and a smart-card reader. Depending on the device, they may also feature "instant-on" capability, some local memory to store numbers and some information downloads, flat-panel displays, and wireless keyboards. While initial models have been priced in the $500 range, the expectation is that as the market develops, unit prices can decline to less than $250 to the consumer. Recent forecasts by industry analysts suggest a US market for such IP phones on the order of 4-5 million units in the next five years.

- **Web kiosks.** This new device category represents the fusion of the Internet appliance with the public payphone. Designed for public spaces (e.g., airports, bank lobbies, convenience stores, supermarkets, malls, hotels, libraries) , it provides users with smart-card-id.-protected access to e-mail, bank balances, weather reports, and other Web information -- and perhaps e-cash downloads to smart cards as well. The deployment of these kiosks by banks would be consistent with the recent trend toward "off-premise" ATM machines, which many banks have used to reduce branch costs, encourage customer self-service, reach particular customer segments more effectively, and provide "7 x 24" support. But non-financial enterprises have taking the lead in deploying them – as they did with off-premise ATM machines. These non-bank sponsors, indeed, may be among the first adopters of "generic" smart cards. (See below.)

- **Real-time Web video.** In the next five years real-time Internet video conferencing is also likely to take off, because of a combination of increased bandwidth, increased local processing power, reductions in storage and digital camera costs, and improvements in the protocols for "multicasting" over the Internet. This means that another important new class of Internet endpoints will be PCs, TVs, or network appliances that include video cameras and microphones for real-time video conferencing, "remote monitoring, " and application sharing. Early versions of Web-based cameras and video phones are already in the market, some of which operate through conventional TV sets. Significant improvements have also recently been made in the software interfaces available for multi-point video conferencing, permitting shared white-boarding and a kind of "Hollywood squares" format for displaying multiple connections on one screen.

- **Wireless/ portable Web devices.** Yet another new form factor leverages recent improvements in wireless communications, to deliver email, browsing, paging, and transactions to mobile users. Three segments are likely to be the most important – Internet car terminals, now in development by most major auto manufacturers; digital cell-phones with screens for pager-like messages and small keyboards, like the one recently introduced by Nokia, and narrow-band two-way pagers with their own Internet addresses for simple messaging.

- **Generic Web-based "smart cards."** Smart cards – wallet-sized plastic cards with embedded computer microprocessors -- have been around since the mid-1960s. At least 32 different proprietary systems are now available, mainly in Europe, where they are already widely used as stored-value and identity cards. As of 1996 there were about 65 million bank-issued "smart cards" (e.g., cards with their own microprocessors, not just memory chips) in circulation worldwide, including 27 million "cartes bancaires" in France, 9 million ChipKnip cards in the Netherlands, 2.5 million Banksys cash cards in Switzerland, 400,000 cards in Sweden, and 165,000 "Mondex" cards in Hong Kong, the UK and New York City. Leading card system suppliers like Banksys Proton, Mondex, and Visa predict that this number will swell to more than 200 million by 1999.

  For our purposes, the news here is the trend toward integrating smart cards with the Internet. This can help to provide secure on-line identification, a storage medium for e-cash, and many other services. Several computer manufacturers have already announced plans to add smart card readers to their PCs or network appliances, allowing the cards to be used as "pass keys." This will permit users to authenticate themselves and gain access to all their Web services from multiple endpoints.

  There is also a trend toward developing smart cards and readers that – like PCs – can run multiple applications written in standard languages. For example, the Java Card, recently announced by Sun's JavaSoft, can download and execute any Java applets, permitting new applications to be added without issuing new cards. The key applications for smart cards include digital identification, digital signatures, digital cash, credit and debit cards, "custom/discount" credit services, and loyalty group programs (e.g., frequent flyer awards) that provide benefits automatically tied to a customer's identify and card activity. Such multi-function cards would be easier to manage than a whole wallet full of incompatible credit, debit, and cash cards from different banks, assuming that they are widely accepted by merchants and banks. As discussed later, they would also help to reduce fraud and theft losses.

  The hope is that the addition of multi-function capability and standardized application interfaces will help smart cards take off outside Europe. However, in the US, the use of smart cards for financial services also depends on e-cash standards (see the discussion below), on upgrading existing ATM, credit and debit card reader infrastructure, and on replacing some 405 million existing credit cards and 220 million debit cards. The unit costs of the cards, at $8 -$18 (depending on the type), also needs to decline to encourage widespread adoption. All this will take time, and non-financial applications like identify and health benefit cards may actually lead the way.

Summary – New Endpoint Trends

Overall, the trends toward the proliferation of new Internet endpoints could go a long way toward replacing our view of computers as expensive, intricate, finicky devices for *aficionados* that are "on line" only part-time, with the notion that they can be friendly, easy-to-use "helpers" that simply run in the background, always on line to send messages, deliver news, analyze problems, and take orders at the touch of a button or a vocal command. For financial service providers, this will be especially important for several reasons:

1. First, it could substantially broaden the Internet's reach. In the US, for example, even today, 55% of households still don't own PCs and 85% lack Internet access. In most other countries, except Nordic ones, the penetration is even lower. The value of many financial services, like payment systems, are proportional to the number of users on the network, and this trend could put many more customers within reach.
2. The proliferation of new appliances also means that the number of transactions and the degree of customer control associated with any given customer base will increase, since customers have more ways of staying in touch and asserting control. For some lines of business – especially those that have depended on "dead balances" -- this may mean more competition and lower profitability; for others, like advisory services, it may mean far more customers in far more places at much lower service distribution costs. In the limiting case, distribution costs are zero; the truly scarce resource is expertise.
3. As discussed below in the section on scalability, the proliferation of new Internet access devices makes it even more critical for designers of Internet-based services to use common "backend" platforms that can easily be extended to multiple devices. This common platform approach differs from the "vertical stovepipe"/ product-centric approach to network design used by many financial service companies.
4. The existence of all these new endpoints will also enable new combinations of services, tailored to specific endpoints. For example, a Web-based "family finance page" might be distributed by way of digital TV, with different modules for each family member, and a combination of special live video "news flashes" and Internet content. A personal credit analyzer might be delivered through Web kiosks at shopping malls, to show balances and payment status for customers about to make purchases. Mobile devices might be closely integrated into information services, as Barclays Bank has done in the UK, providing customers with digital cell-phones and two-way pagers to access their balances and do transfers.

Other Application-Enabling Technologies

In addition to new endpoints, in the next five years customers will also benefit from several interface technologies that will help to improve the immediate experiences they have with computers and Internet services.

- **Java and VRML**

Java is a programming language that was especially designed in the early 1990s to support the kind of networked, highly-distributed applications that the Internet calls for. It is now supported by every major software company, including IBM, Oracle, Computer Associates, Motorola, Netscape, Novell, Apple – and at least officially, Microsoft itself. The first thing that is very appealing to all these players about Java is that it permits large-scale software applications to be decomposed into small programs called "applets" that can be easily downloaded over the Internet, relatively securely and reliably. This delivery model solves several important problems:

- **Upgrading.** Java's network distribution model is a natural for mass-based Internet applications, because it sharply reduces the expense and bother of distributing and upgrading applications to millions of people. It permits older versions of applets to be replaced as frequently as desired; in the extreme case, with network appliances, new applets can be transparently downloaded at every sign-on.

- **Customization.** Java also permits specific versions of applets to be downloaded to different customers, depending on who they are or where they are connecting from. This enables a high degree of customization, with, for example, each private banking client receives a different version of a portfolio analyzer tuned to his portfolio, special interests, and degree of sophistication.

- **Roaming.** For users switching from one portable device to another – say, from a PC to a mobile phone to a Web kiosk – the Java distribution model also lets the system "tune" the service to the device – for example, by varying graphical displays for device screen type.

- **Rapid Application Development.** Applets can also be used to build certified components that can sent securely over networks to users, then joined together to form large applications. Combined with the ability to update the components frequently, this building-block approach is extremely valuable for rapid application development and deployment, especially in fast-changing arenas like financial services.

- **Support for Internet Appliances.** Java also supports a "thin client" application model that is required by Internet appliances. Unlike PC programs, which permanently reside on a local computer's system disk, Java applets – generally much leaner than PC applications -- can be downloaded securely from an Internet or network server, executed on the fly, and then deleted after each execution.

- **Platform Independence.** Java is also especially well-suited for Internet appliances because it has been designed to be platform-independent. Java code runs in a "virtual machine" that is specific to the operating system of the actual hardware it is using. So the very same program can run on Unix systems, Windows, NT, the MacOS, and embedded operating systems. Whether or not Microsoft ends up supporting Java, this feature alone is likely to insure that Java thrives.

- **Special Network Security.** To allow for secure distribution of applets over public networks, Java has been specially designed to insure that applets are safe from viruses. Applets can be digitally signed by their authors or by certificate authorities, letting users know that they actually came from the organization that claims to have written them. This also helps to prevent the introduction of illegal code or viruses during downloads, and the counterfeiting of applets. Even signed Java applets have to run in a secure "sandbox" on user machines, which restricts their access to critical system files and communication channels.

It is important to emphasize that all this is still very much work in progress, and that the next two years will see a veritable outburst of new Java features and applications. In this regard, one very interesting development will be the appearance of specialized applet vendors that supply Java components specially designed for Web-based services like e-commerce and finance. As noted, another keen area of interest is embedded systems, where Java will be used by software developers to do "write- once" applications that can run on multiple devices, including set-tops, network appliances and smart cards.

Finally, for end users, another important new interface technology that is consistent with the Java model is the "Virtual Reality Modeling Language" (VRML), a new light-weight language specifically designed to deliver high-quality 3D images over the Internet. This would permit, for example, the deployment of a "virtual ATM or branch" 3D application environment on customer screens, providing them a familiar metaphor for conducting financial services online.

- **Electronic Payments**

Electronic payments – electronic cash, bill presentment, and bill payment -- are yet another important arena on the Internet services frontier. At the moment they are even more of a work in progress than Java, but they are also poised for a take-off.

As of 1998, for example, there are at least a half dozen incompatible systems available for electronic cash, with no common standards for security, storage, transmission, or hardware. However, if standards can be developed, we believe that both electronic cash and bill payment could become major factors in Internet-based financial services in the next five to seven years. As discussed in Chapter III, depending on how this transition to new payments mechanisms is managed, banks could either be tremendously advantaged, or they could begin to lose their whole central position in the payments system.

- **Electronic Cash**. At the moment electronic cash, in particular, suffers from having more variables than constants:

- Some versions of electronic cash, like Mondex or Visa's stored-value card (SVC), reside on a smart card and can be used in lieu of currency, as well as over the Internet. Others, like DigiCash's Ecash™, are limited to on-line use, because the cash only resides on a PC.

- Some schemes permit electronic cash to be downloaded over the Internet from a bank account to a PC or smart-card, while others require an outside system, like a bank ATM, to do the downloads.

- In some schemes a person who pays with electronic cash is anonymous; in other schemes both payers and payees can be tracked by financial intermediaries or others.

- Non-bank issuers of electronic cash, like Mondex and Digicash, are not yet subject to regulation or reserve requirements by central banks, unlike bank issuers. This disparity is a source of concern for some regulators, like the US Federal Reserve.

- Some schemes permit direct transfers of electronic cash among card holders without third-party intermediaries; others require intermediaries to be in the middle.

Underlying all these differences, the fundamentals of electronic cash are pretty simple. When a customer wants to withdraw cash in electronic form, his bank essentially creates a block of electronic digits that represents the cash, and sends them to the customer's smart card or PC . The bank "mints" the cash by encrypting the block of digits with its own private key, using procedures that are a standard part of public key encryption technology. This means that the digits that represent the money cannot be altered. However, anyone with the bank's public key could decrypt the message and verify that the electronic cash, indeed, came from the bank.

In the same way, customers can use their own private keys and digital signatures to request electronic cash transfers securely, and banks can use the customers' public keys to unlock the requests and verify whom they came from. Electronic cash might also then be passed from one customer to another, if the card readers permit transfers from one smart card to another.

All this basic technology, especially encryption and the encoding and protection of smart card chips and readers, are relatively mature. As discussed below in the security section, public key encryption is a well-tested method for exchanging private information over public networks. The smart card experts that we interviewed agreed that while cards are not totally impermeable to tampering, this can be made so difficult that the profit from breaking any particular card is small. The risk of "double-spending," the duplication of digital money already been issued, can also be minimized.

The key obstacles to the widespread deployment of digital cash, therefore, are not technical, but market-related:

- **Fuzzy Customer Benefits**. While in theory stored value cards offer many potential benefits to customers and financial institutions, in practice these benefits have been slow to materialize, at least outside Europe:

- First, in principle, electronic cash might appear to make sense for payments made over the Internet, especially small ones ("micro-billing") where transaction processing costs are much smaller than for credit cards. Originally it was expected that this Internet micro-billing opportunity would be a large one, because the Internet was supposed to create a huge market for "digital rights" -- electronic magazines, proprietary data, music, and so forth. However, this digital rights market has failed to develop, because of legal issues and the fact that digital rights management technology is still cumbersome. So almost all the information distributed on the Web today is either provided voluntarily or supported by advertising.

- In the US, credit card use is very high, compared with Europe or Canada. Whether on the Internet or off it, therefore, electronic cash starts out with a more serious competitor than in these other markets – and more than two-thirds of the world's Internet users and merchants are still located in the US. Even for purchases of non-digital items over the Internet, therefore, direct electronic cash transactions with merchants is far behind credit cards as an

Internet medium of exchange.

- Off the Internet, in theory electronic cash might be a great substitute for currency and coins, because it reduces the nuisance costs of cash management and the risk of theft. Assuming that downloads are permitted over the Internet, it also provides the equivalent of a virtual ATM machines in the privacy of one's own home. It might also be a time-saver in "wireless" situations – for example, by permitting commuters to pay without stopping as they drive by toll points.

  > However, in practice, electronic cash is much more complex for many users to handle than simple currency. Some versions are also not as anonymous – a major consideration in the demand for larger-denominations bills.
  >
  > There are also trust issues involved in sending cash payments directly over the Internet, just as in sending cash through snail-mail. For some reasons it just *seems* safer to call many people to an 800-number operator and give one's credit card number over the phone for a purchase – or for that matter, to send a check through the mail, or give a credit card to a waiter, despite the fact that, as we will see below, all these *off-line* payment mechanisms are subject to far *greater* risk.
  >
  > As for the "ATM in the home" concept, this feature is not available on all electronic cash schemes, and users may in any case doubt trusting their bank accounts to the instabilities of modem connections and Windows 95™.

- On the supply side, electronic cash also promises significant theoretical benefits for financial institutions and merchants. Depending on the scheme, for financial institutions these include being able to invest "float" that would not be available with currency; reduced handling costs; reduced consumer and merchant fraud over credit cards (because no user numbers are provided to merchants in the case of a cash card); and the opportunity to gather information on aggregate customer spending patterns. For merchants, in principle the benefits include faster credits for cash receipts, lower currency handling costs, reduced risk of theft, and the ability to service new customers who use e-cash.

- In practice, many of these benefits also remain unproven, since they depend on widespread acceptance of the cards. As one New York merchant who was participating in a recent Mondex card trial put it, "For me it's *more* expensive, because now I have another card reader to maintain, in addition to credit and debit cards, and I still have to handle currency too. Maybe I'm missing something here?"

- **High Initial Costs/ Poor Performance**. As noted above, smart cards, card readers, and the overall card systems are relatively expensive, partly because the market is so small that unit costs have enjoyed the benefit of experience curves, and partly because upgrade costs for existing credit and debit card infrastructure in the US are very high. Performance, system costs, and ease of use also remain an issue for the proposed Secure Electronic Transactions (SET) standard for secure card payments.

- **No Standards.** As noted above, the smart cards used by most of the technologies for store value have been incompatible with each other, capable of running only stored value and other applications that were written by their sponsoring institutions. In Europe, the larger market is divided among several strong local and foreign card vendors; in the US, where the market is just emerging, no clear leader has emerged to drive a *defacto* standard.

Yet standardization may be a precondition for anyone to make money with such new payments systems. These system are subject to the same kind of network externalities as telephone networks -- their value is proportional to their number of users. So if rival systems compete vigorously, no one system may be able to reach the critical mass that is required for any one sponsor to make money.

In the case of e-cash, this absence of card standards is partly just a matter of technology -- the very fact that new hardware is required means that there is more of an opportunity for established players to fight over rival proprietary systems. As we will see below, in the case of electronic bill payment and presentment, the fact that the technology consists merely of software and standard Internet transport and security has helped it grow more rapidly.

Overall, therefore, e-cash appears to be stuck in what economists call a low-level equilibrium trap. On the Internet, as noted, the volume of electronic cash transactions is miniscule. While some stored-value card experiments have succeeded – for example, the 1997 introduction of Metrocards for the New York City Transit System -- and several large-scale trials are under way in the US, Canada, and the UK, most observers agree that we are several years away from widespread deployment.

There are several factors that might turn this situation around. The growth of a more lively Internet digital rights market, facilitated by higher bandwidth and increased availability of audio and video content, could create a need for micro-payments. As smart cards become more standard and multi-purpose, their costs will come down. Their readers will also be integrated into standard network appliances and combined with those for debit and credit cards. Stored-value applications will then be easier to promote as just one of several optional "wallet card" applications that are available to merchants, end users, and banks on standard computer network gear.

The proliferation of digital IDs as a much more secure way of handling electronic mail and e-commerce will also help. Finally, if consumers become interested in selling services or merchandise to each other over the Internet, by way of exchanging anonymous electronic cash – sort of an online equivalent of the untaxed yard sale – this might also help to kick-start this market. But there is little consumer-to-consumer selling over the Internet at this time.

- **Bill Payment/Presentment.** Meanwhile, electronic bill payment and presentment over the Internet have started to take off, at least in the US. Although they still account for only a small fraction of all check payments, their use is now more than doubling each year. This growth is highest in the consumer-to-business segment that accounts for about forty percent of all US paper checks. The penetration rate differs strikingly across industries – interestingly, US banks as a whole are well behind insurance companies, electric utilities, and phone companies in the use and acceptance of electronic payments, as well as in the encryption of their electronic communications. By the year 2005, according to one conservative estimate, these electronic bill payments are expected to account for at least ten percent of all US transactions payments.

As noted earlier, this rapid take-off has occurred in part because these bill payment services are technically much simpler than electronic cash. They don't require extensive hardware upgrades or new hardware standards for their deployment. Their cause has also been helped by recent trends toward more powerful endpoints, the proliferation of home-finance applications like Quicken™ and Microsoft Money™, improved software for electronic forms and remittance processing, and middleware.

All the main retail bill payment services follow the same basic flows; the main differences are mainly with respect to service fees and institutional roles rather than underlying technology. Electronic home banking, the simplest, consists of having customers with checking software or online bank accounts authorize payments by their banks or third parties to particular billers. In many cases this "half-paper" method still results in a paper check for the biller.

The full-blown version of Internet bill presentment lets customers receive their bills in electronic form from billers by way of the billers' Web sites, through aggregators like CheckFree, or through banks that have signed up billers and deliver their bills for them. In the simplest case where no aggregator or bank is involved, the biller prepares an electronic bill and makes it available to the customer at the biller's Web site, secured by the customer's password or digital certificate. The customer visits the Web site, accesses the bill, and uses one of several alternative on-line payment methods. Or he may decide to send a paper check. When an aggregator is involved, the only significant difference is that the bill is sent to the aggregator, and the consumer visits the aggregators' Web sites to retrieve bills and make payments. More advanced versions of bill presentment could employ new "push" and "publish and subscribe" technologies to regularly round up bills from multiple biller and aggregator Web sites and deliver them directly to the customer's machine.

For retail businesses that send out lots of direct bills, the completely-electronic version of this payment/presentment system has many advantages – assuming that enough of their customers use it. First, it eliminates much of the paper handling and cumbersome record-keeping associated with paper checks, and the even worse "check and list" system imposed on billers by first-generation PC banking. Second, they often get paid earlier, because there are fewer delivery and processing delays. Third, depending on whether billers use aggregators or deliver their bills directly, they can use billing as an opportunity for other marketing communications with their customers. Fourth, there is much less check fraud.

Finally, as more and more businesses establish electronic payment for their customers, they also are more inclined to use it for business-to-business and business-to-consumer transactions, so these also become easier. Depending on how competitors and customers react, the combination of all these effects may permit billers to reduce their transactions costs, increase margins, improve terms, or lower prices. All told, it is no accident that a majority of the nation's largest billers are now moving rapidly toward electronic systems.

Many of these same benefits also apply to retail customers, assuming that enough billers and banks start using the system to simplify rather than complicate bill management. In general, for both billers and customers, electronic payments reduces their transactions costs and reduces their desired equilibrium level of cash balances.

For financial services institutions, especially banks, the consequences of electronic payment are more mixed. On the upside, some banks may be able to realize significant processing cost savings. According to one recent analysis, payment processing now accounts for as much as a fifth of all non-interest expenses at some large US banks, and electronic processing might save at least $.75 per check. Some banks may also be able to use bill aggregation to strengthen their ties to other billers and retail customers.

The problem arises, however, because many banks – especially the largest ones -- reap a surprising share of their non-interest income from their *privileged* roles in the current *domestic and international* payments systems. Electronic payment technology, especially that which might readily be made available over a secure, pervasive, global Internet, exposes the banking sector as a whole – including reserve banks like the Federal Reserve-- to a fundamental risk of disintermediation by other payment system providers. Left to its own devices, therefore, the whole payments system might soon follow the way of bankcard drafts, three-quarters of which are now handled electronically by non-banks.

Chapter III will examine the potential industry impacts of these new payments mechanisms in more detail. The main point to take away here is that all these questions of fundamental roles and interests have been begged by the unstoppable momentum of these essentially Internet-enabled, software-defined, real-time, globally available new electronic payment technologies.

- **Intelligent Software Agents**

The concept of an "intelligent agent" refers to the notion of software that resides on a network and performs services for customers in the background, even while they are not connected.. The basic notion is that "the truth is out there" model of the Internet is not enough. It only becomes compelling if the network can reach out to the customer and inform him that, say, "Your CDs are maturing, there are several alternative investments that look appealing, and here’s what we need to consider in making this decision."

The underlying technology that provides such capability is a combination of distributed object languages like Java (for agent applets that are downloaded on the fly), new "search engines" and "publish and subscribe" technologies that can filter and distribute data much more efficiently, and decision-analysis tools that have been revived from the good days of artificial intelligence.

Armed with these tools, there are actually several different kinds of agents that can be turned loose. One simple kind, a "search-bot," periodically searches for updates on topics of special interest to the customer -- say, the latest Indonesian financial crisis. A "decision assistant " might accepts inputs over the Web, prowl around, and return with an analysis of alternative mortgage or college loan financing options. Simple versions of decision assistants are already deployed in the on-line mortgage, car loan, and credit card markets. A "virtual portfolio manager" might monitor a customer’s portfolio and propose alternative investments on the basis of market conditions, the client’s risk preferences, and his tax status.

As the technology matures and customers become familiar with it, there will be an opportunity for financial service providers to develop more sophisticated versions of their own branded Internet agents and decision assistants. On the other hand, agent technology also provides an opportunity for third-party advisors to offer "provider-neutral" analyses of terms and features – an approach already taken on Web sites like InsWeb. As this technology progresses, therefore, it may begin to have similar effects on industry competition as independent insurance agents or mortgage brokers.

- **Other Internet Applications – Groupware and Telephony**

Until recently, point-to-point e-mail and file transfer were the predominant ways of communicating over the Internet, but this is changing rapidly. Internet chat-rooms, popularized by AOL, already permit groups with similar interests to type messages back and forth in real-time. In the next five years, this simple, synchronous, text-based messaging will be supplemented by many more group communication options – which may be very helpful to financial services companies that wish to stay close to their customers.

Groupware was first introduced in the late l980s, most popularly by Lotus Notes™, one of the first PC software products that addressed the opportunity to provide low-cost workgroup communication over computer networks. The essential Notes™ insights were twofold. First, a great deal of daily work gets done by virtual teams whose members span organizational boundaries and are scattered all over the planet; second, "synchronous" – e.g., simultaneous real-time – communication among team members is often unnecessary, and indeed, downright counterproductive.

In other words, it is often sufficient for team members to be able to connect to a shared virtual data base at their own convenience, upload the latest team work, and download their own contributions. For Lotus, this "asynchronous" model of workgroup communications was an adroit discovery, because it was all the limited bandwidth of the late 1980s could handle. But there were still many technical obstacles to implementing the asynchronous model in that pre-Internet period.

In the last decade, the Internet's maturation, the inroads that Microsoft Windows™ has made against other operating systems, and increasing network bandwidth have made all this much easier. Real-time technologies like Internet voice and videoconferencing have also started to emerge from the labs. The next five years, therefore, will see a much richer variety of alternatives become available for conducting group communications over the Internet. These include the following:

- Customers and account representatives who use different Internet service providers will be able to join chat-rooms and 3D "virtual offices," three-dimensional representations that allow documents, analysis, and transactions to be shared in real time.

- Notes™-like shared discussions over the Internet will be available on the fly, permitting secure discussions among customers, account officers, and experts from different locations.

- Improvements in IP telephony and multicasting technology will permit real-time video conferencing, voice conferencing, program sharing, IP faxing, paging, and white boarding to be conducted over the Internet among multiple parties.

From the standpoint of financial services, as noted, the importance of these group communications tools is that they can facilitate much closer contacts with customers and much more timely services. The instant availability of Internet data stimulates the need for much more frequent contact that the usual monthly client newsletter. A broker might offer his clients daily face-to-face Internet video or telephony briefings with leading analysts or fund managers around the globe, accessible from any digital TV, at much lower cost than non-Internet video conferencing. Interest group discussions might be organized to put clients in touch with experts in taxation, finance or technology. Internet telephony might be integrated into customer support so that a client just clicks a button on his Web phone and gets a call back from the help-desk – whether he is connecting *from* Sao Paulo or Kuala Lumpur *to* Sao Paulo or Kuala Lumpur.

III. The Network Side – Bandwidth, Scalability, Security, and "Middleware"

So far we have focused mainly on the "client" side of technology trends, where change will be most visible. On the network side, there are also several trends under way that will extend the reach and power of Internet-based services in the next few years. The most important have to do with improvements in network bandwidth, scalability, security, and so-called middleware.

Bandwidth

As summarized in **Figure 2.3**, Internet access speeds are determined by the processing power, bandwidth, services, and the kinds of network management available at three different levels of a network's architecture -- the *local area network,* the *local loop* that connects local area networks to service providers, and the *backbone* that connects service providers.

Figure 2.3 Basic Internet Connectivity - Key Elements

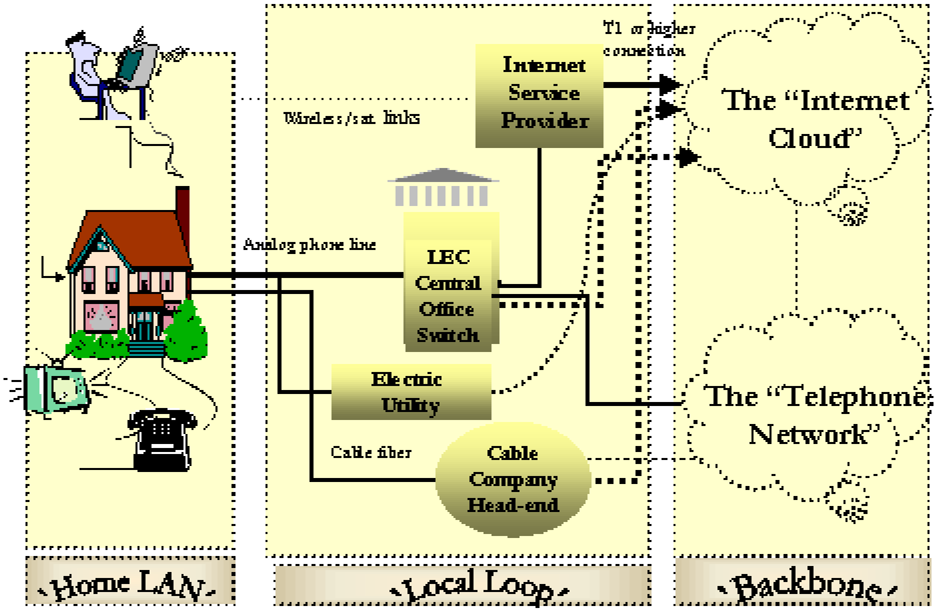

Raw bandwidth, the capacity to send a given volume of digital bits at a particular speed, is only one determinant of the Internet's performance. If a Web site's or ISP's servers and routers are short of processing capacity to handle requests quickly enough, no amount of extra bandwidth will fix this problem. In general, computing power, router capacity, and storage capacity are a substitute for bandwidth, up to a point. For example, local caching servers – which store copies of frequently-accessed Web sites on nearby servers, and only reach out to the actual Web for updates – can significantly reduce the amount of long-distance bandwidth required to display those pages quickly.

Despite this role of infrastructure, we have grown used to thinking of bandwidth as *the fundamental* bottleneck. This is partly because computer processing speeds have been historically improved at a much faster rate than access speed. Following "Moore's Law," CPU and router processing power have increased about ten times in the last five years, while modem speeds have only increased by a factor of four. This is only partly due to technology. It is also because bandwidth improvements have until recently been under the control of phone companies, cable companies, and electric utilities, a group that, let us say charitably, has a rather mixed track record with respect to customer service, competitive pricing, new technologies, and support for the Internet's growth.

The good news is that this bandwidth bottleneck is about to disappear – for wireless and wired technologies alike, and at the local loop and backbone alike.

- **Home/ office LAN connectivity**

We usually associate local area networks with large enterprises that have thousands of users and their own IT departments. Indeed, in the early l990s, enterprises with more than 100 users did account for the great majority of networking investments, and they have also been the earliest adopters of Internet services, remote access, and high-speed connectivity. However, several trends point to the emergence of a "home/office" LAN market that may, in the aggregate, be quite large, and will demand many of the same higher-bandwidth technologies that have been so far available only to enterprises.

- First, a growing share of homes and small businesses now have multiple PCs that are capable of being networked. At last count, at least 10 percent of US households homes have more than one home PC, and among households with incomes greater than $100,000, the share increases to about a quarter.. The proliferation of Internet appliances noted above will only add to the number of devices that could be coordinated by a home LAN.
- Second, an increasing share of the work force needs to share applications and services between home, road, and office. Recent estimates are that 12 percent of household heads already regularly telecommute to work, and another 13 percent have their primary place of business in the home.

- Third, there is the concept of the "silicon furnace," a low-cost server for the home that keeps track of e-mail, voice mail, Internet faxes, calendars, home entertainment, and family Internet access accounts, provides local caching of frequently-accessed Internet services and remote control of household features like temperature and home security, and keeps track of calendars, electronic bills and payments. Within the home or small office LAN, it is quite easy to network multiple PCs together using conventional 10/100 Mbps Ethernet

technology, hubs, and routers, or wireless LAN technology. This is now beginning to receive serious attention from several PC server vendors, notably IBM and HP. Home servers could include set-top boxes as nodes on the network, integrating messages that are received over cable services.
- Another interesting variant on this theme is the "digital apartment/ office complex" recently pioneered by commercial real estate developers in San Jose and Manhattan. These feature built-in high-speed connections from each apartment or office to the Internet, with the complex essentially acting as an Internet service provider.

From the standpoint of retail financial services, early adopters of home/ office LAN technology may also be leading-edge customers for "advanced" retail services like electronic payments, integrated banking-and-brokerage, small business planning, and real-time video advisories.

- **"Last mile" technology**

Most of us have long since given up on the mid-1990s view that telephony, data, and multimedia networks are "converging," in the sense that millions of telephone customers would someday soon able to download movies, Internet services, and video conferencing over fiber links to their doorsteps. This vision proved vastly more costly to implement than it was to talk about, mainly because of the exorbitant costs of providing fiber connections to all the residential customer on the local loop – the so-called "last mile" problem.

Still, there is now much more technical rivalry than ever before with respect to last-mile technologies. Each one has its advantages, and they are all being refined, so it is impossible to say which ones will dominate. It doesn’t much matter – the point is that we may soon have more local bandwidth than we know what to do with.

- At the moment, more than seventy percent of all Internet users in the US are connected to the Internet over ordinary phone lines at average speeds of 33 Kbps or less. The latest generation of analog modems provides 56 Kbps connections, and many observers believe this is an upper bound for modem links over single copper phone lines.
- The next step up is to a digital switched service called ISDN, a technology that has been available for more than a decade. It offers up to 128 Kbps in both directions over a single phone line. But ISDN has been has been plagued by standards issues, is hard to configure, and usually requires an installation visit from the phone company. In the US it has also has taken years for phone companies to upgrade their central offices to handle it. By now there are only about 500,000 ISDN subscribers in the US, including 186,000 residential customers. It has done much better in Germany, where Deutsche Telecom has installed more than 1 million ISDN lines.
- Satellite companies like Hughes Electronics recently introduced a "DirectPC" service that permits download speeds of up to 400 Kbps to computers connected to a $350 wireless dish, with a regular phone line for uploads at 56K. However, reception problems in dense urban areas, plus the extra costs of the service, have limited its deployment to less than 50,000 subscribers so far.

- As noted above in our Internet TV discussion, another higher-bandwidth alternative is the cable modem, which provides Internet access over the cable operators’ hybrid fiber-coax

(HFC) links to the home at average speeds of 1-2 Mbps. To date, about 100,000 cable modems have been installed in the US. Compared with alternatives like ADSL, cable modems are more expensive to install, because cable operators have to upgrade their networks to permit two-way service. Of the 65 million US households that now subscribe to cable, only about 9 million are on cable networks capable of handling two-way signals. Cable-based Internet services are also shared rather than switched – unlike phone service. So they are inherently less secure and subject to more congestion.

On the other hand, once cable networks have been upgraded, they may also be able to handle much higher bandwidth – up to 5-8 Mbps or more in two directions. Cable-based Internet services also have inherently lower networking costs than phone-based alternatives. Internet access over cable is "always on;" with no dial-up is required to access the network. Quality of service also tends to be more uniform. Cable’s security problems can also be solved pretty easily with the help of encryption and digital ids. Finally, as discussed below, cable modems may also present fewer conflicts of interest for cable operators than high-bandwidth alternatives do for local phone companies.

All this, plus the "mega-deals" for digital set-top boxes signed by the largest US cable operators with Sun, Microsoft, and General Instruments in December 1997, makes it likely that cable-based Internet access will finally take off in the next two-three years.

- Another high-bandwidth alternative is provided by a group of connection technologies called "xDSL," especially ADSL (Asymmetric Digital Subscriber Line). This technology permits telephone companies or electric utilities to offer Internet access over ordinary copper cable at download rates of 1.5 Mbps to 9 Mbps, so long as users have special modem-like terminators and service providers have installed the appropriate Ethernet concentrators and routers in their central offices. So, unlike cable modems or ISDN, ADSL requires no rewiring. A user just connects a phone line to a $350 terminator and plugs in his PC. The line can still be used for voice calls as well as for Internet access, and the connections are inherently private, unlike cable. In recent trials, such systems have delivered reliable service at download speeds of 1.5 Mbps, with 64 to 350 Mbps on the return channels.

  So why are there only about 4,000 ADSL users in the US right now . This is mainly because local phone companies (and electric utilities) have been very slow to embrace the Internet as a new channel. This may be an instructive parable for financial service companies. As noted earlier, at first many local phone companies were obsessed with the notion of building their own proprietary switched fiber networks, to provide "video on demand" (VOD). That not only proved very costly, but even if the phone companies had succeeded in capturing the entire annual consumer budget for video rentals and pay-per-view services, they would not have earned a decent return on their VOD investments.

  At the same time, burned by their poor experiences with ISDN, the phone companies were generally slow to provide ordinary Internet service. At first they shunted the business to the roughly 4000 independent ISPs that appeared in the US market between 1994 and 1996. Then, as that business took off, they started to pick fights with the ISPs over flat-rate pricing and access charges, which did not endear them to their own customers. As for ADSL, which had been invented by Bellcore, their own research lab, in 1989, the phone companies were at first chiefly concerned that it might cannibalize their profitable leased line businesses.

So until recently, ADSL has had to rely on a more progressive phone companies, like Canada's Sasktel, Finland Telecom and Helsinki Telecom, a few electric utilities that wanted to use their copper networks to provide Internet services, and a few ISPs that aggregated demand from apartment complexes, provided residents with ADLS modems, and resold T1 capacity they had purchased from local phone companies.

In the last year, however, local phone companies in the US have finally realized that VOD is dead and that cable modems and digital set-top boxes are about to eat their lunch. ADSL, on the other hand, may provide them a low-cost way of providing higher speed access, and a backdoor way to push ISPs out of the way and supply their own links between their central office switches and the Internet. So more than a dozen ADSL trials have recently been started in Boston, San Francisco, Atlanta, and several other cities. In January 1998, Intel, Compaq, Microsoft, and four out of the five leading local telephone companies also announced an alliance that would role out 1.5 Kbps ADSL service to the home on a large scale by 1999. By making built-in ADSL modems a standard option in new PCs, the alliance hopes to avoid the nasty configuration barriers that held up ISDN, and cut terminator costs. Depending on how all this goes, the industry experts that we interviewed are hoping for as many as 10 million installations in the next five years.

- Yet another category of new access technologies would bypass all these entrenched cable and telco interests and offer very high bandwidth (e.g., 27 Mbps or more) in at least one direction over fixed wireless (LMDS) connections. LMDS-based services are already being tested in Brazil and the US, where the FCC will hold a first round of LMDS license auctions in February 1998. ISPs and other service operators are expected to be important participants. LMDS is a new wireless technology that offers such high speed connections over distances up to five kilometers -- quite adequate for Internet subscribers in flat urban centers like Brasilia, Iowa City, Tulsa, or New Orleans. Several other versions of fixed wireless technology are also under development. Of course cable modems and ADSL are more robust to weather and building interference than these wireless alternatives, but they will have trouble matching all this short-distance bandwidth.

- Meanwhile, in continental Europe, where leased lines and ISDN have long been more dominant, alternatives like cable modems and ADSL face an up-hill battle. In Scandinavia, the "convergence" vision of fiber to the doorstep still holds. Telia, the leading Swedish telco, is deploying a "Asynchronous Transfer Mode" (ATM) fiber network to the home. In the US, a coalition of universities, government agencies, and the N.S.F. are also proceeding with a high-fiber project, "Internet2." But this is not open to commercial users. Plans by US phone companies to build fiber to the door have been shelved.

Whoever gets there first, all this competitive activity probably guarantees that faster "last mile" access will eventually be available in the US and most other markets. For financial services companies, this trend has several implications.

- First, if we assume that average Internet access speeds will jump to 1-2 Mbps or more in the next five years without a significant increase in cost, this means the Internet is sure to become an even more valuable general-purpose medium for a majority of potential retail customers than it is already.

- Second, higher-speed connectivity will also permit higher-performance applications – for example, the ability to download video advisories on the fly, provide 3-D data analysis, and push news, information, and new applications to clients continuously.
- Third, as noted, all the increased access bandwidth will put more strain on the Web backbone and server infrastructure, and make it even more important for financial service providers to deploy highly-scalable service platforms that can keep up with the fatter pipes. (See below.)
- Fourth, financial institutions should realize that as, in effect, some of the world's largest Internet service providers, they have a profound stake in Internet regulations that pertain to local loop competition. For the most part, however, they have so far stayed on the sidelines when it comes to debates among phone companies, ISPs, cable companies, and the FCC over data service regulation and pricing policy.

- **"Backbone" technologies**

At the backbone level, there is also very good news regarding the outlook for improvements in Internet bandwidth. Indeed, many observers believe that we may be on the verge of a bandwidth glut. We think that this view may be a little premature, because there are still serious hurdles to overcome with respect to scalability. However, there is no question that there will be huge increases in backbone capacity in the next three-to-five years.

To begin with, planned investments in new global Internet backbone capacity by leading backbone suppliers in this period already add up to a several-fold increase in capacity. The leaders include long-distance carriers like MCI/ Worldcom, Sprint, and AT&T, and several new players that have specialized in backbone routes, including Quest, IXC, and Williams. Sprint, the traditional leader, has just upgraded its entire national backbone from 155 Mbps to 622 Mbps. Quest, a company that was only founded in the mid-1990s, is building 16,000 miles of fiber links along railroad rights of way, and already has several links in place with terabit (1000 gigabit) capacity --- enough fiber to carry *five* times the annual volume of all US telephone voice traffic.

Several other backbone providers also plan large capacity additions in the next two years. Accordingly, all these providers are looking for downstream demand, signing up "local loop" partners like ISPs, cable companies, and phone companies who can deliver customers to use all this capacity. A spot market in IP backbone capacity is also beginning to develop.

In addition to all this raw fiber, there are also several technologies that will increase backbone throughput and quality dramatically in the next few years:

- **Wave division multiplexing (WDMA),** a new approach to combining multiple digital signals over fiber, was developed by the hottest IPO of 1997, Cienna Corp. The technology is already in pilot at AT&T, Quest, Sprint, and several other backbone providers. Its supporters claim that it can boost the productivity of existing fiber backbones by a factor of 8 to 10, bringing "Moore's Law" to bear on backbone productivity.
- **Gigabit Ethernet** is the most recent improvement on the networking technology that was

pioneered by Bob Metcalfe at Xerox in the early l970s for use in local area networks. Available this year from leading suppliers like 3Com, Bay Networks, and Cisco, it will provide bandwidth in the gigabit (1000 Mbps) range. It will first be employed on private corporate network backbones, where it is backwards-compatible with the 10- and 100-Mbps Ethernet that most companies already use. But it may also soon be adopted by some Internet backbone providers.

- **Asynchronous transfer mode (ATM)** technology**,** the "other ATM, " is also finally getting off the ground. This data switching technology has been around since the early 1990s, and it has long been predicted to displace the Internet's aging TCP/IP protocol and the router technology that goes with it. ATM is actually a whole new set of standards and protocols for data communications, plus the underlying hardware and software. It was originally developed by a research arm of AT&T and several other phone companies. Indeed, one of its key problems has been that its design was left in the hands of these slow-moving entities – until it was taken over by the ATM Forum, which has not been much faster. Meanwhile, the simpler IP protocol was vigorously implemented by networking leaders like Cisco, Bay, and 3Com, to the point where the entire Internet now runs on TCP/IP. As noted earlier, ATM is finally now being deployed on a mass basis by Sweden's Telia and a few other carriers.

  Despite its slow takeoff, ATM does have some technical advantages over IP, especially its ability to handle high-bandwidth video connections with guaranteed quality and timing, something that the "connectionless" IP protocol has trouble doing. In the last two years several products have also emerged from companies like Ipsilon and Cisco to run the IP protocol over ATM hardware. These allow a smoother transition between IP and ATM. But the worldwide volume ATM hardware sales only just recently passed $1 billion, and most observers have long since stopped holding their breath.

Figure 2.4. Internet Bandwidth Performance Trends

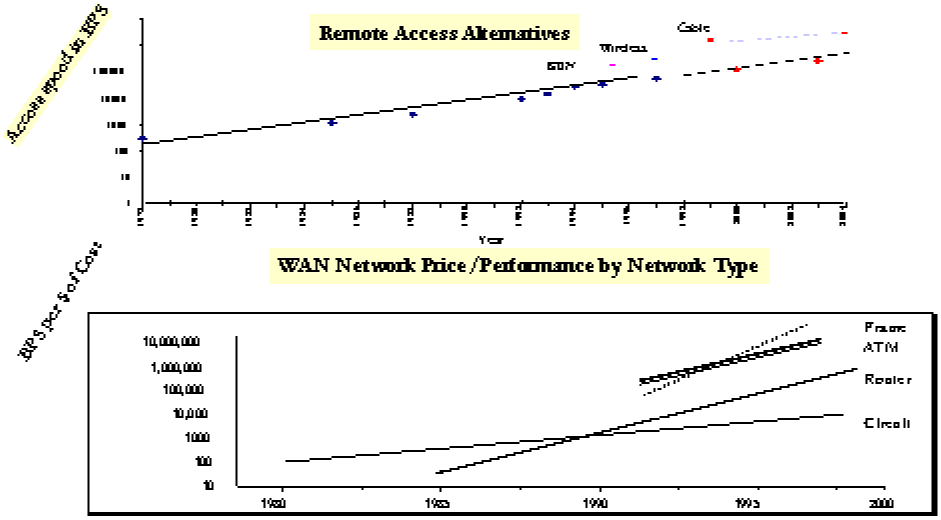
Remote Access Alternatives

WAN Network Price/Performance by Network Type

Source: SHG Interviews

Together, all these trends toward enhanced backbone speeds strengthen the case for the Internet as a powerful new two-way medium that is here to stay. If the backbone enthusiasts that we interviewed are correct, contrary to the doomsday forecasts noted at the beginning of this chapter, congestion on the Internet's backbone may be an occasional inconvenience, but gridlock and systemic collapse are unlikely.

Among other things, this means that the complex congestion pricing schemes for Internet services that have been proposed by some well-known economists are likely to be still-born, just as schemes for the marginal-cost pricing of mainframe time-sharing services were thirty years ago. If the bandwidth optimists are right, the dam is breaking and the flood will have to pass before we can begin to for water. (See **Figure 2.4**.)

However, before we celebrate, there is one other serious concern. This is the question of whether or not the Internet's infrastructure will be able to handle all the traffic generated by all this new bandwidth. This issue of scalability is a crucial one for any potential provider of large-scale Internet services. As we will see, while there is reason to be hopeful, the issue is by no means resolved.

Scalability

Until recently many Internet service providers have assumed somewhat blithely that the Internet's infrastructure is inherently scalable. After all, its servers, routers, and network software have so far been able to support the Internet's rapid growth without many breakdowns. Contrary to the doomsayers, to date there is no evidence that the Internet's *average* performance has deteriorated with growth. Indeed, just the opposite – it is handling more users at higher access speeds and larger volumes of content with shorter delays, on average, ever before.

However, one implication of all the predicted improvements in Internet bandwidth and local endpoint devices -- combined with continued growth in the sheer number of Internet subscribers -- is increased stress on this infrastructure. This means that the Internet's scalability may only really be tested in the next few years, as all these new endpoints and bandwidth come together.

Scalability means, first, that the Internet's servers, routers, and software systems should be capable of expanding smoothly from a small number of initial users. Second, the Internet should also be capable of absorbing new users reliably and economically at any scale, with average costs of service declining as numbers increases. Third, the infrastructure should also permit new services to be added easily without duplicating initial investments.

There are several reasons for large scale service providers to be concerned about scalability in all three senses.

- To begin with, much of the Internet's network infrastructure has been duplicated, lock, stock, and barrel, from the "client/server" model of distributed computing that was originally developed for private networks. This was never really intended to provide secure, reliable services to millions of users on a daily basis, nor was it intended to provide economies of scale and scope.

- Most of the network software now being used by leading service providers like AOL or E-Trade for authentication, e-mail, directories, firewall security, and billing was custom-built from the ground up. There is a shortage of standard, modular "retail" software that has been field-tested under varied conditions by many different customers.

- Third, on the hardware side, there are indeed some very knotty unsolved scalability issues that have to be dealt with in order to facilitate the Internet's continued high growth, especially with respect to the current generation of servers and routers.

- On the server side, there is a need to bring features like fault-tolerance, clustering, and shared backup down from the heights of "specialty" servers (like those made by Tandem and Stratus) and incorporate them in the Windows NT™/ Intel servers that many ISPs are struggling to employ as Web servers.

- For routers, backbone providers like Quest and Sprint have already discovered there are simply no routers available that can interconnect more than a few high-capacity lines of 155 Mbps or larger. Even that requires whole room full of racked Cisco routers. Quest, in particular, has several terabit-capacity fiber lines in place that it can't light, because no terminating equipment exists that can handle terabit data flows.

This problem may get worse before it gets better. If the number of Internet users grows *exponentially* – it has recently doubled every 9 months or so –the *volume ofinterconnections* and the resulting total load on network interconnections may grow *hyper-exponentially*. So network traffic would grow even faster than chip speeds and router performance, both of which are (only) subject to Moore's (exponential) Law.

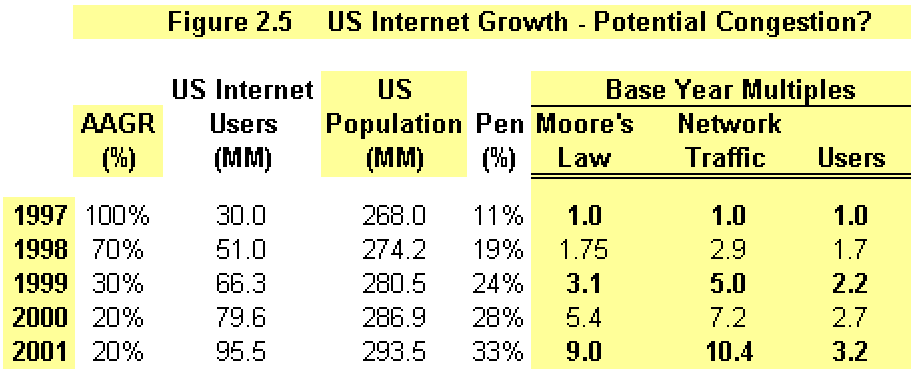

Figure 2.5 US Internet Growth - Potential Congestion?

		US Internet	US		Base Year Multiples		
	AAGR	Users	Population	Pen	Moore's	Network	
	(%)	(MM)	(MM)	(%)	Law	Traffic	Users
1997	100%	30.0	268.0	11%	**1.0**	**1.0**	**1.0**
1998	70%	51.0	274.2	19%	1.75	2.9	1.7
1999	30%	66.3	280.5	24%	**3.1**	**5.0**	**2.2**
2000	20%	79.6	286.9	28%	5.4	7.2	2.7
2001	20%	95.5	293.5	33%	**9.0**	**10.4**	**3.2**

As shown in **Figure 2.5**, for reasonable assumptions about Internet use and processing power, if we assume that network traffic is "increasing" in the number of interconnections on the network, and that these are proportional to users, then congestion – measured by the ratio of traffic increases to gains in processing power – may indeed grow sharply in the next five years, *despite* bandwidth improvements.

However, as we discovered in our interviews, all these scalability problems are now being tackled by all the leaders in the networking industry. On the server side, a new generation of standard "network operating systems for ISPs" is now emerging, with built-in security, billing, directory, e-mail, and caching modules that can be scaled from a few hundred users to several million. New "publish and subscribe" software will also help to improve scalability by reducing the amount of traffic generated by individual users. And the increased use of application proxies and caches for leading Internet applications by service providers and corporate networks will also help to reduce the load on the Internet's backbone infrastructure.

On the hardware side, more powerful transactions and applications servers, with greatly-improved clustering, shared storage, and failover features are also appearing. Improvements in endpoint processing power are allowing the use of better compression algorithms and the local execution of Java-based programs, further reducing network traffic. Finally, the terabit routing problem is also being tackled by several new companies that are trying to break the constraints of Moore's Law by applying massive parallel processing architectures, extending today's 155 Mbps routers to 2.4 Gbps or more.

Still, scalability is likely to remain a hot spot for Internet service providers throughout the next decade, as the entire industry comes to grips with what a huge job it is to provide reliable electronic services to millions of customers.

Indeed, our interviews with senior managers at leading Web-based financial services confirmed that scalability is already perceived as one of the most important "sleeping" issue that they face. In the words of one technologist at a leading British bank, " This has already been very expensive for us -- we had to rebuild our entire Internet banking service from the ground up, because the network architecture that we started with just wouldn't stretch."

Security

Another basic network ingredient in Web-based financial services is of course security. According to the standard definition, secure transactions consist of being able to transact in secrecy ("*privacy*"), know whom we are dealing with ("*authentication*"), make sure that our communications have not been tampered with *en route* or after-the-fact ("*integrity*"), and also make sure that at the end of any given transaction, the other parties can't disavow whatever they agreed to ("*non-repudiation*").

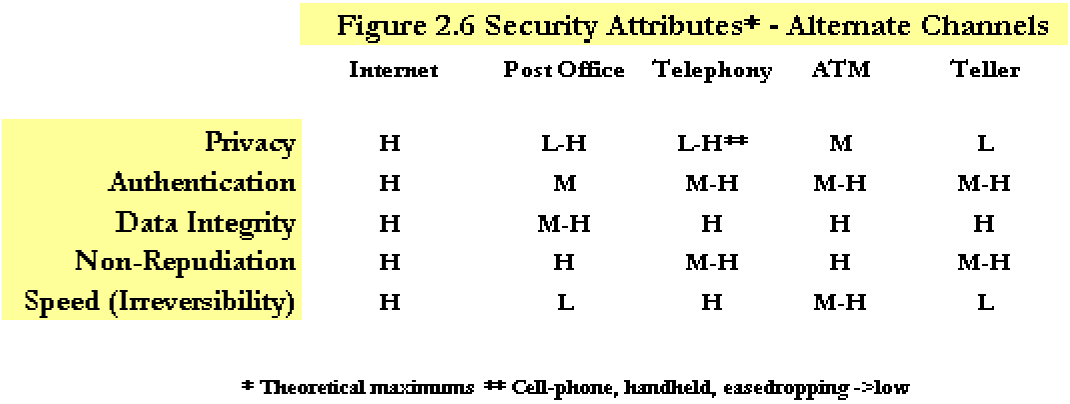

Figure 2.6 Security Attributes* - Alternate Channels

	Internet	Post Office	Telephony	ATM	Teller
Privacy	H	L-H	L-H**	M	L
Authentication	H	M	M-H	M-H	M-H
Data Integrity	H	M-H	H	H	H
Non-Repudiation	H	H	M-H	H	M-H
Speed (Irreversibility)	H	L	H	M-H	L

* Theoretical maximums ** Cell-phone, handheld, easedropping ->low

As shown in **Figure 2.6**, all these core security attributes can be supplied in varying degrees by conventional financial service channels -- sending a check through the mail, checking one's bank balance over the phone, paying a bill with a credit card at a restaurant, transferring money with the help of a teller or an ATM. In addition, these channels differ in the *speed and convenience* with which they supply these attributes. They also differ in the degree to which customers are exposed to attacks that have nothing to do with financial services *per se*, but only with the technologies they use (e.g., computer viruses spread over the Web; live viruses spread by way of "snail-mail.")

Concerns about the Internet's security, especially its vulnerability to outside "hackers," are indeed among the public's most important reasons for their reluctance to use it for electronic finance or commerce. In fact it turns out, however, that security is already one of the Internet's strongest features.

- First, all the basic technologies needed to deliver the basic security attributes over the Internet -- including public key encryption, digital certificates, secure operating systems and firewalls, and anti-virus tools – are already relatively mature, as we'll see below. And it is becoming much easier to implement these technologies because of recent improvements in computer processing speeds, encryption algorithms, standard "application programming interfaces' for security software, and the development of certificate authorities.

- Second, the Internet's basic approach to security, which relies on encryption to protect the security of messages rather trying to protect the *physical* security of any particular communications channel, turns out to be much more robust than the "physical guard-dog' approach. In contrast to the Internet, according to a recent study, only about forty percent of US banks use any kind of data encryption whatsoever in their internal communications, while credit card authorizations and even ATM network communications are often not encrypted at all.

- Third, the available empirical evidence also favors the view that encryption-based security is far tighter than the security of other retail channels. For example, the introduction of smart cards in France has sharply reduced the volumes of credit card fraud there. And recent estimates for the losses incurred by banks due to online fraud are miniscule – a recent ABA estimate is that it amounts to only $5 million a year.

Of course this is likely to grow. But it will have to grow a lot to match the security losses recorded in other conventional channels. For example, check fraud in the US now approaches $14 billion a year, greater than all losses due to armed robbery, credit card fraud, and securities theft combined. For its part, credit card fraud now exceeds $850 million a year.

Figure 2.6. Security: "The enemy is...*us*."

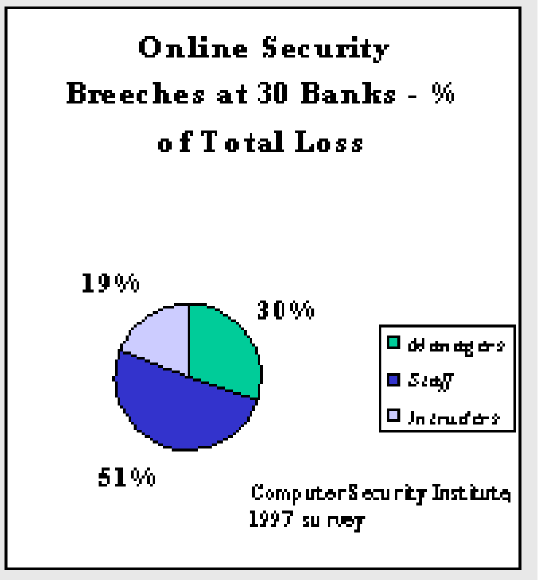

Furthermore, as noted by a recent study by one leading professional organization for computer security experts, the vast majority of computer security problems are not due to clever, malicious outside intruders, but to security breaches by internal staff. (See **Figure 2.7.)** As one security expert that we interviewed put it, "This is not new, you find the same thing in garment district stores or Macy's – the security guards are always the first ones you check. As the CIA was reminded in the

Aldrich Ames case – the real enemy is often…*us*."

Of course the digital security methods preferred by the Internet will not protect against sheer stupidity – the sort uncovered, for example, by one security expert in Chicago, who recently found that more than ten percent of one large company's employees were using the word "Bulls" as their passwords. Ordinarily, however, unless one is up against a malicious attack from a well-funded competitor or a government agency, financial institutions and their customers who follow best practices should have little to worry about on the Internet. In fact if they are really serious about reducing fraud losses, these institutions and customers should really switch to Internet-based transactions, Internet security methods, and smart cards as quickly as possible.

- **Encryption**

To understand why public key encryption is so important to financial institutions, we need to understand the two types of encryption that are used to keep data exchanges secret -- symmetric key encryption and public key encryption.

Symmetric key encryption is the old "shared secret code' technology, where the same code is used to encrypt and decrypt messages by all parties to a transaction. Since security depends on the secrecy of a single code, those who want to exchange something in secret start out with a problem – how do they all get the secret code? Somehow they have to exchange a secret code in advance, before any encrypted messages have been sent or decrypted. This is not much of a problem locally, but it creates more serious problems if we are talking about strangers, or people who are far apart. Anyone intercepting the secret key might also be able to decrypt messages and then masquerade as one of the communicating parties.

Public key encryption, invented in 1976 by Whitfield Diffie and Martin Hellman, eliminated all these difficulties, and has deservedly become the foundation for Internet security. The technology uses two keys instead of one -- a public key that everyone has access to, and a private key that only its owner has. The real breakthrough is that the public and private keys have a unique one-to-one correspondence, through the mathematics of factors and modular arithmetic. Without going into the gory details, the power of such encryption methods is that in order to break the code – e.g., determine the private key from the public key, so that plain text can be derived from code text – a prohibitively expensive factoring exercise involving an immense amount of computer power has to be undertaken, even for a relatively short key length.

So anything encoded with the public key can only be decoded with the private key, and only something encoded with the private key can be decoded with the public key. This means that, for example, anyone who wants to send a secret message to John can encrypt it using John's public key, and only John will then be able to decrypt it, using his private key. To enable secret communications between two distant strangers, they just exchange public keys (by e:mail, for example). So long as their private keys remain private, their secrets are safe.

The exchange of private, unalterable, remote messages and transactions among complete strangers is just one reason that public key encryption is so important to on-line financial services. Another is for the authentication. If, for example, John receives a message that can be *decrypted* by Sharon's public key, then John knows that only Sharon had to be the one to encrypt it, using her private key.

In effect, then, Sharon has "signed" the message – he knows it is from her. That is all we mean by a "digital signature."

The latest trends in public key encryption are towards longer key lengths, faster algorithms for public key encryption, and more frequent "re-keying." The emergence of cryptography on smart cards is also very significant. Smart card chips are becoming powerful enough to generate public-private key pairs and perform encryption/ decryption on the card, without the private key ever leaving the card. This means that private keys never have to be loaded onto a user's PC, where they might be snooped by viruses or be vulnerable to other kinds of break-ins. In fact, the owner of a key never has to know it.

As processors get more and more powerful, the frequency with which key pairs can be changed will increase. In the not too distant future a smart card might generate a new key pair for every transmission, sending a new public key to other parties for each communication session and "forgetting" the key after the session completes. Even if the key were somehow broken, it would only be useful for the messages sent in that one session.

What does all this mean for financial transactions over the Internet? First, sending credit card numbers, other financial data, or a digitally-signed check over the Internet can easily be made much more secure than transmitting such data by, say, voice over a wireless phone, by punching numbers into a public payphone, or through snail-mail.

Second, encryption technology can also reduce insider fraud as well. For example, a programmer might be given the ability to *encode* data for a system without having the ability to *decode* it.

Third, protocols like SSL and SET, which incorporate basic public key encryption into standard procedures for exchanging digital signatures and conducting card transactions, are beginning to make the use of all these tools much easier. New developments in API software , like the GSS API for smart cards, as well as the high level of security built into languages like Java, will make all of these mechanisms more convenient to use. This is important, because the real weak link turns out to be – rather like condoms -- quite simply whether people bother to use the protection that is available.

- **Digital Certificates/Ids**

Public key cryptography by itself provides privacy and integrity. Digital certificates and digital IDs build on this foundation to deliver secure authentication and non-repudiation.

A digital certificate is simply a secure ,unique piece of digital data that certifies that the holder of that data is a specific real person (or organization.). At present, an individual's public key is almost always used as the unique piece of data. The digital certificate says the equivalent of " the holder of public key [long number] is John Jones, social security number 123-45-6789 , who lives at 123 Front Street, was born on April 1, 1975 and has red hair and blue eyes. " The Certificate Authority that issues the certificate has the job of verifying that this information is true before it issues the certificate.

So what can John do with his digital certificate? Suppose we want to make sure that we are in fact communicating with John Jones, and vice versa. We exchange digital certificates that we decrypt with the Certificate Authority's (CA's) public key. Reading the certificate, we see a certification by the CA that John's public key is [long number], and John also sees ours. Without certificates, we cannot be sure that that public key we got really does belong to John Jones. So the CA basically insures against fraud.

Another important benefit of digital certificates is that they implement non-repudiation. For example, if a message can be decrypted by John's certified public key, it came from John. Nor can he claim that the contents were changed. Unlike symmetric key encrypted messages, a receiving party can't decrypt a message, alter something, and re-encrypt it again.

Over time, the role of the Certificate Authority is likely to expand beyond simply certifying identities. They might also certify financial status, software applications, legal documents, time-stamps, and the validity of Web merchants, as a sort of on-line "better business bureau."

Playing this role might be interesting for major banks or other "trusted" financial services institutions. After all, the business is not very far removed from other authentication services that banks and credit card companies already provide. Eventually it might also become a large source of fee revenue, at least as large as the $1 billion a year that US banks now collect from ATM transfers. Depending on offshore regulations, it might also provide a global market that a few well-known institutions may dominate -- with every merchant server, every individual who engages in e-commerce, and perhaps every Java applet needing its very own certificate. There could be significant first mover advantage here, given the scale economies.

The interesting question, then, is, will the global certificate leaders be software companies, insurance companies, encryption companies telephone companies, postal authorities, package delivery companies, currency printing houses, or major banks? We will return to this question in Chapter III.

- **Viruses/Illegal Code**

A third key ingredient in Web security is virus protection. We tend to associate viruses with malicious programs that are sent from user to user on disks in a matter of days, and run one's files. In fact they have now matured to become programs that can spread over the Internet in minutes and try to pierce one's security for financial gain. Fueled by the growth of network applications and the potential gains from "snooping" financial data, they are now more powerful and dangerous than ever. .

There are several important trends that are helping to contain this situation, however. First, Java is a very effective anti-virus mechanism. The corruption of servers and desktops by viruses from downloads can now be virtually eliminated if only digitally signed, "sandboxed" Java applets are downloaded from the network.

Another important trend in fighting these viruses is "capability-based" security systems. This extends the notion of the "sandbox, " providing each software object a set of digital certificates that it owns, defining its authority perform specific functions. Without such authorization, the object can't do anything. This means that capability-based systems makes it possible to severely restrict the functions that any imported piece of software code can perform. Some capability-based systems are already being developed for wide-scale deployment on the Internet in 1998. JavaSoft has recently announced refinements to its "trusted applets" model, along the lines of capability-based limits

Another key trend with respect to virus protection is digitally-signed code. Browsers now routinely ask permission before downloading unsigned code from an un-trusted source. They also ask before downloading or opening macros or other code where viruses can effectively hide. Combined with increasingly sophisticated virus detection programs for detecting viruses in email, if proper download procedures are taken, security against viruses can be very strong.

The final trend here is toward more secure operating systems and firewalls, not only to protect networks from users but also users from networks. This not just a matter of virus protection, though viruses do have greater access to certain operating systems than others. It is also a matter of preventing outsiders from exploiting the notorious security bugs of operating systems like Unix or Windows that permit, say, outsiders to ransack a user's PC anytime he is connected to the Internet without a firewall. Of course this has special significance for financial services companies, because while the theft of an individual credit card number over the Internet might be costly, the theft of 50,000 passwords or account numbers from a Web authentication server would be devastating. However, a great deal of effort has already been applied to filling such holes, and standard precautions exist for all but the most exotic defects.

In short, as in the case of other Internet security issues, most of the methods required to protect against malicious viruses – like scan programs, digital signature requirements, capability tools, secure operating systems, and firewalls – are "reasonably well-known," as one security expert said. The major problem – and loophole --- is that users and network providers often don't always make use of the tools and best practices that are available.

- **"Middleware"**

Middleware sits between the application layer of a software program and the operating system, providing some service to the application. All the functions provided can usually be written into the application layer, but by using middleware, code can be often developed much more quickly.

For example, standard middleware for legacy system integration is now available that makes it straightforward to "join" old mainframe-based services and data bases with new Internet-based systems. Using messaging middle-ware and transaction monitors, existing back-office systems can communicate with and connect with Internet based systems. This "coupling" of old and new systems eliminates a large barrier to deploying on-line services.

Summary – "Frequently Asked Questions," Key Technology Trends

- **What are the overall implications of these technology trends for FIs?**

The financial services industry will have an ample supply of raw materials available in the next few years to construct much more powerful Internet-based financial services, at least in First World countries. While there are plenty of unsolved problems, the Internet is within reach of becoming a relatively safe, reliable, scalable, place to do robust e-commerce and finance – and eventually, the dominant channel for such services.

- **Which kinds of new applications will these technology trends enable?** On the upside, as we have seen, there are many opportunities to use technology to develop stronger, more continuous ties to customers, as well as new delivery vehicles for bundled services.

New Internet appliances will make it possible to provide endpoints to retail customers that are simpler and cheaper than today's PCs. They will get access to information and transactions from a wide range of wired and wireless devices, and from many more locations – the ATM and the branch will, in effect, become virtual. New application interfaces like Java and VRML will make also help to make the interfaces to such virtual services much more accessible to the ordinary user. Intelligent software agents will be on-call to track accounts , analyze financing and investment alternatives, and shop for the best yields. On line. Internet-based telephony applications will make it possible to provide integrated voice support to customers over the Internet, and lower-cost, computer-based alternatives to existing calls centers and voice-response units – with some video as well. Internet groupware will make it easier to arrange private discussion groups, news and analysis sessions with clients over the Web.

On the other hand, some other new Internet applications, especially in the areas of electronic payments, advisory services, and certificate authorities, will pose serious threats of disintermediation to the industry. We will explore the potential impact of re-intermediation in Chapter III.

- **Is the Internet safe enough for retail financial service applications?** While there are always risks to guard against, the fact is that all the basic security tools needed to enable secure e-commerce and finance on a large-scale basis are now in place. Indeed, we have argued that Internet-based services can generally be made even more private and secure at lower cost than financial services delivered through non-Internet channels. Assuming that certificate authorities, in particular, become more widely adopted, we see no fundamental obstacles to the growth of the Internet as the world's largest, most secure platform for retail finance and commerce.

- **Will there be adequate bandwidth?** Indeed, in a five year time frame there is likely to be a surge of bandwidth available at affordable prices at both the backbone and the local loop levels. Making using of all this new bandwidth will be a challenge for application developers.

- **Is the Internet likely to scale to handle services for millions more users?** While there is no evidence as yet that the Internet is suffering from congestion or performance degradation, continued growth in users, Internet-based applications and devices, and the likely surge in bandwidth (ironically) will put severe strain on the software platforms, routers, and other network infrastructure. While the whole networking industry is devoting a great deal of attention to this issue, this is clearly a potential trouble spot that designers of new services should watch carefully – especially when they chose service platforms.

- **Is there any evidence yet that Internet-based services are subject to economies of scale or product scope?** We will have more to say about this key issue in Chapter IV. However, the key points to note from a technology trends standpoint is the overall trend in Internet-based services toward lower-cost, highly-distributed, "open" systems. Here, network economies – the ability to provide millions of users with common interfaces that support multiple applications whose value increases in the number of users on the network -- are very powerful. So in general, there's now the equivalent of a public transportation system that is available for any potential supplier to reach any customer – those suppliers who, in an earlier era, invested heavily in their own private canals may now no longer gain as much of a competitive advantage on the basis of infrastructure ownership alone.

On the other hand, there may be significant economies of learning, advantages that accrue from entering a complex market early and making useful errors that increase one's options down the road.

- **What are the implications for management?** Nearer term, there is likely to be something of an indigestion problem, as financial institutions find that they have their hands full just staying on top of this "technology glut." This highlights the importance of *technology management* and *partnering skills*, as opposed to proprietary systems and sheer technical know-how.

In the next chapter we will take a closer look at some "lessons learned" from other recent efforts to launch new retail financial services, and consider some of the overall impacts of these services on the industry.