Evolving From Telephony To IP Networking

N. F. Maxemchuk

AT&T Labs - Research

ABSTRACT

The separation that the IP layer provides allows the Internet to change services and transmission facilities more easily than the telephone network. As telephony becomes less profitable, and other communications services more lucrative, we have an incentive to adopt an IP architecture in order to introduce new services more quickly. Changing the current telephone network into an IP network, while retaining complete communication connectivity and service quality, is not trivial.

In this note I present an evolutionary strategy. The strategy at the network layer is based upon the island and tunneling techniques that are used to introduce new technology into the Internet. These techniques encourage trying alternatives rather than making irrevocable commitments. They also encourage continuous changes in the network rather than embracing a technology until it becomes obsolete. The strategy at the services layer has both supported and experimental components and results in continuously changing IP platforms (plural).

1. Introduction

The IP layer in the Internet is between the services that are used by the end users and the transmission facilities. The layer was designed to isolate the services and transmission facilities so that one could change without changing the other. Independence has resulted in both services and transmission evolving quickly.

The current digital telephone network is completely integrated from the end user's terminal, through the services that are provided, to the transmission formats that are used on the various media. For instance, changing the current 64 kbps voice coders to a lower rate would leave very few pieces of the telephone network unchanged. When the digital phone network was designed in the 1970's all available information, about all of the layers, had to be used in order to construct an economically competitive network.

The IP strategy also evolved in the 1970's, but the charter of the government organizations that funded the work made them consider precompetitive technologies. Now that the Internet's technology is competitive, we would like to apply it to the AT&T network to obtain the advantages that is has provided to the Internet.

Evolving into an IP network is particularly difficult because IP is a moving target. In addition to the continuous evolution of IP, that is orchestrated by the IETF, the government is funding a precompetitive technology that may result in a disruptive change in IP. It is important that we understand this change because it may have a significant impact on the IP isolation strategy and the relationship between services and technology. In section 2 I will describe the reasoning that resulted in IP isolation and why a change in IP is necessary.

In section 3 I describe a strategy for evolving from the current telephone network to an IP network. There is concern that IP networking implies a reduction in voice quality. If IP is applied to our network properly we can preserve voice quality. Our objective should be to use this technology to provide better quality services.

The Internet is the domain of the entrepreneur and start up companys. In order to compete in this environment, we must establish platforms that take advantage of our size and persistence. We are moving to an IP network to take part in the rapid changes that are occurring in networks. Our IP platforms must continuously change to reflect the network. In section 4 I explain the need for several independent platforms, and an evolutionary process for the services layer.

2. The IP Isolation Strategy

When the ARPAnet was started in the late 60's communications software was implemented as a single routine that was part of an application program. By the early 70's network implementors realized that many of the communications functions are independent and could be changed without affecting the others. For instance, the procedure for checking the data for errors and recovering lost messages can be the same whether or not a message is subdivided into packets for transmission. Separating functions resulted in layered architectures, such as IBM's 5-layer SNA and the ISO/OSI 7-layer architecture. The 3-layer IP architecture, that is used on the Internet, won, and is the last healthy layered architecture.

The three layers in the IP architecture separate the functions that are performed at the end user, such as error recovery, the functions that are performed in the network nodes, such as routing, and the transmission technology. A significant difference between the IP architecture and the other layered architectures is that the layers in IP correspond to the physical devices in the network, while the other architectures drew lines in the software that was often on a single machine.

IP isolates, in addition to separating, services and transmission facilities. There is no channel that is intended to pass information between a service and the transmission facility. The decision to isolate is intentional and separate from layering. For instance, the architects could have left a field in the IP header that the service set to request a specific channel on a transmission facility.

Isolation permits services and transmission facilities to evolve independently. If a service knows enough about the transmission facility to request a specific part of that facility, and the facility is changed, then the service must also change. The rapid evolution of the Internet proves that the decision to isolate was the proper decision --- at the time it was made.

One reason that the IP isolation strategy worked as well as it has is that it matched the technology of the 70's. In the 1970's dedicated logic controlled transmission equipment and the network oriented operations were relatively simple. Most channels were just fixed rate bit pipes. The advantages that can be obtained by communicating with a dumb bit pipe can be obtained by other means. Therefore, IP gave up very little.

An important change since the 70's is the increased, economical use of processing. Transmission technologies are becoming more intelligent, but IP cannot take full advantage of the intelligence. For example, when we send IP over ATM, IP cannot take full advantage of the quality of service options. ATM is only the first in future generations of intelligent transmission facilities. It is likely that all future transmission technologies will have embedded processors.

IP must change to take advantage of the intelligence in future transmission facilities. IP may change by interpreting a packet's control bits, as described earlier, but this results in a tight coupling between services and transmission facilities. The tight coupling can be reduced by using the active networking technology that is currently being funded by DARPA and the NSF.

In an active network instructions are passed across the IP layer, from the applications to the transmission facilities, along with the data. The instructions are much more general than the control bits in current communications protocols. Instead of an application using a few bits to access a useful characteristic of a specific transmission facility, it uses more general commands to describe the service it requires. A transmission facility does the best it can to provide the service, but the instructions can be interpreted differently as technology changes. Active networking is the juncture of communications protocols and general purpose computing languages. The more general language makes it more likely that we can change applications and services independently and retain a separation.

Active networking is a precompetitive technology, in that there are no economically viable active networks. In addition, there are technical problems that must be solved before active networks are deployed. Java has demonstrated that external users can be given safe access to processors, providing the effect of the program is constrained in a "sandbox." A similar, network sandbox, must be defined before applications are given

access to transmission equipment. While active networks are currently precompetitive, it is likely that they will become competitive more quickly than the quarter of a century that IP required.

3. An Evolutionary Strategy

The entire telephone network cannot be converted to an IP network in a single instant. We must develop a strategy that makes it possible for pieces of the telephone network to evolve into an IP network while remaining connected with the rest of the network. The strategy must also preserve the quality that voice users expect.

IP is not a stagnant protocol. It changes continuously. There is a tunneling strategy that permits new versions of IP to be deployed on islands in the network, without disrupting the network. The current version of IP is IP version 4, IP version 5 failed because it did not solve the important problems, and IP version 6 is currently being tested.

The strategy that is used to convert the telephone network to an IP network should be similar to the strategy of gateways, islands, tunnels and encapsulation that is used to convert between versions of IP. This strategy does not require a large, initial investment in infrastructure, nor an irreversible commitment to a specific technological approach. Instead, we start with several testbeds, that may use different technologies, in several geographical areas. The successful testbeds grow to become the network.

Once adopted, the strategy of islands and tunnels permits a network to evolve gracefully by laying the ground work for its own destruction. When a better technology exists, it is introduced into the network as islands, and, if it is truly better, it eventually replaces the current technology. More than two generations of technology may operate at the same time. For instance, we can still have legacy telephone networks in operation when we move from the generation of IP that we are moving toward, to the generation that will replace it.

3.1 IP Encapsulation vrs. Protocol Conversion

The IP evolutionary strategy is based upon packet encapsulation and relaxed constraints on the end-to-end operation of new versions of IP. In this section I will explain the difference between the IP strategy and a general protocol conversion.

In the Internet a new version of IP is introduced by grouping routers that interpret the new IP protocol into islands. Between the Islands and the main body of the Internet there are gateways. The gateways make certain that the packets can be interpreted by the routers in the region that they are entering. They do not change the packets that are transmitted by the application. Instead, they either encapsulate an IP packet inside of a packet that can be interpreted by the routers, or they remove the encapsulation to recover an appropriate packet.

When a packet that uses a new version of IP travels between islands, the gateway encapsulates the new IP packet in an IP packet that is familiar to the main network. The encapsulated packet is destined for a gateway to another island. When the packet arrives at the next island the new IP packet is removed from the capsule. An old IP packet traverses an island by being encapsulated in a new IP packet.

There are important differences between IP encapsulation and a more general protocol conversion:

1. The IP strategy obtains the best possible service in each part of the network, rather than finding a service that can be provided end-to-end;
2. New versions of IP are backward compatible; and,

3. IP is a stateless protocol.

The first difference gives the end users some advantages before significant parts of the network are changed. The last two differences make IP evolution simpler that a general protocol conversion.

Most protocol converters enable communications over paths with multiple protocols by using a common subset of the functions that are supported by the protocols. The Internet is connectionless. When a packet enters the network it does not know the path or the protocols it will encounter. When a packet traverses segments that use different versions of IP, the strategy is to use the capabilities that a service requests, where ever they are supported. As a result, the services in an IP network are softer than in other networks, but the network can be improved gradually.

For instance, if IPv6 provides reservations and IPv4 does not, then the IP encapsulation strategy makes reservations in the parts of the network that use IPv6, even though it cannot make end-to-end reservations. The new IP service cannot guarantee end-to-end reservations, however, as the packet spends more of its trip in parts of the network that support IPv6, the service improves. In contrast, a standard protocol converter does not offer any improvement until reservations can be made for the entire path.

New versions of IP are backward compatible. We do not remove capabilities as we move to new generations of IP. Backward compatibility makes protocol conversion simpler:

— When a packet enters a region of the network that uses a newer version of IP, all of the functions required by the application's version of IP can be performed.

— When a packet enters a region of the network that uses an older version of IP, a proper subset of the functions can be performed.

The IP protocol carries all of the information that is needed to process a packet in the packet header. The stateless nature of IP makes it possible for a gateway to operate on individual packets. When a packet enters an IP region that is different from the source, the gateway determines all of the functions that must be included in a new header from the the header that is encapsulated. When an encapsulated packet is disassembled, the encapsulated header contains all of the functions specified by the source.

3.2 Using IP Evolution Techniques In The Telephone Network

The initial change from a telephone network to an IP network is a general protocol conversion, not an IP evolution. Before we can use an encapsulation and tunneling strategy across the two networks we must resolve some basic differences between IP and telephone networks.

1. The telephone network transmits a continuous signal, while the Internet transmits packets as needed.
2. The IP network is a stateless network, while the telephone network stores state at every switch on the path.
3. The set of functions performed inside the telephone network and the Internet are not only different, but neither is a proper subset of the other.
4. The services carried by the telephone network have quality requirements, which the Internet does not currently provide.

Solving the first three problems is straightforward, if not trivial. The fourth problem is more complex, not because there is no solution, but because there are several, incompatible solutions. It is also very likely that the best solution to the fourth problem today will be different than the best solution in a year or two. We will deal with the fourth problem in the next section.

The first problem is solved by organizing the continuous signal on the telephone network into packets at the first gateway that users on the telephone network encounter, and converting the data in the packets into a continuous signal at the last gateway to users on the telephone network. These "edge" gateways are more complex than gateways inside the network. As the user equipment becomes IP capable these gateways are

only needed to reach legacy equipment or users on old style, circuit switched networks.

Note that there is no need to convert between the packet format and a continuous signal if IP islands are connected together with circuits. The circuit switched network can carry packets as a digital bit stream. In the next section we will explain why it is very likely that we will connect IP islands together with circuits, at least in the short term.

Obviously, when we convert from a continuous signal to packets, and back, we have all of the standard problems of picking coders, compressing silence and other less useful information, restoring timing and deciding whether to treat analog data modems and fax machines as data or continuous signals. All of these problems have been investigated for decades and we have many answers, if not a universally accepted answer. These problems are not addressed here.

An IP packet's header has all of the information that a router needs to process the packet. The information is explicit. There is no implicit information that has been deposited by other packets. By contrast, the telephone network uses a circuit set-up procedure that deposits all of the information that a switch needs to process the data on the circuit at each switch on the path. It is not necessary for each transmitted data element to carry information for the switches. ATM networks straddle the two approaches. An ATM cell contains some information that a switch uses to determine how to treat the cell, but the cell header is small relative to an IP packet header, and other information is deposited at the switch during a circuit set-up procedure.

There are at least two approaches for dealing with data that must traverse packet, circuit and cell switched networks that maintain different amounts of state. The first, and most straightforward approach, is to use a circuit set-up procedure that selects the gateways between the networks, deposits state information at the gateways, and sets up circuits in the circuit switched network. This approach can provide almost the same quality of service as the current telephone network, and is suggested in the next section for the initial phases of the the migration. The advantage of this approach is that it requires no changes in the current IP or telephone networks. We must create a circuit set-up protocol that traverses both the circuit and packet networks, but this is a straightforward task.

The second approach uses quality of service routing that uses circuits to bypass congested regions of the Internet. This approach requires a change in the operation of some of the routers. The modified routers are connected to both the packet and the circuit switched networks. When modified routers exchange packets over the Internet they measure the current loss and delay. When a packet arrives at a modified router, the router decides whether or not the Internet delay to the next modified router is likely to satisfy the quality of service constraints of the packet. If the Internet path is congested, the packet is sent over a circuit. If the modified router runs out of circuits, it uses the circuit switched network to add capacity.

Many of the differences in the functions performed in the telephone network and the Internet are caused by the different capabilities of the user devices. The computers that are connected to the Internet can perform more functions than telephones. As a result, the switches in the telephone network perform both the network and application layer functions in the Internet. The network layer functions in both networks are concerned with directing the data. Although the techniques are different, packets that traverse both networks can use the network functions in each. The application layer functions, that are performed in switches but cannot be performed in routers, can be performed in application layer servers on the Internet.

For example, in a conference call on the telephone network signals from the users are added together in a switch. A router in an IP network does not process the data in a packet and cannot perform this function. A similar function can be performed by routing voice packets to an application layer network server, processing the data in the arriving packets, and transmitting new packets.

3.3 Quality in an IP Telephone Network

The price of telephone communications is decreasing to the point where price alone will not entice many voice users to move to a lower quality IP network. Voice communications over a new IP voice network must be as good or better than the current telephone network. The advances in voice coding that have occurred since the adoption of the current 64 kbps telephone standard make it reasonable to expect much higher fidelity communications in the same bandwidth.

The quality of voice communications on an IP network is impacted by the delay and losses that occur because the network is shared with unpredictable users. The extent of the problem is different in different parts of the network. The solutions differ depending upon the fraction of the users who have moved to IP, and are changing as technological advances occur.

The access link from a user to the ISP, or the first router or LAN switch in a private network, is not shared. (Older Ethernets share the access link, but there are other ways, which will not be discussed here, to guarantee quality on those networks.) A user who owns the entire link can perform access control and prioritize his own traffic so that all of the services that gain access to the link can achieve their required quality of service. The quality may be preserved beyond the access link depending upon how well the network is provisioned, the priority policies enforced by the routers, and which technological advances are implemented.

At present, end-to-end quality cannot be guaranteed for every connection on the Internet. The strategy for providing high quality, end-to-end communications is to create islands in which the quality can be preserved and tunnel between the islands. The means for tunneling between the islands is likely to change with time.

Initially, a tunnel may be a single telephone circuit per user. IP voice packets will enter the network on an island, and remain on the Internet until a congested region is encountered. At that point, the network provider dials up a telephone circuit and sends all of the packets from the user on that circuit. If the first router dials up a circuit to the destination telephone or final router for each of its voice connections, then the service and quality is very similar to the current telephone network, except for the packet assembly delay. The advantage of using IP in this limited environment is that multiple sources can be multiplexed onto a wide band access pipe and non-voice sources can be placed directly onto the Internet.

Using a single circuit per user strategy makes sense when a small fraction of the population uses IP. In order to have complete connectivity, most of the IP traffic must pass through a gateway to the circuit switched network. It makes sense for the traffic to stay on the circuit switched network in its final form once the IP network becomes inadequate. The final form on the telephone network is a single circuit per user.

As the IP population and the number of IP islands increases, circuit aggregation between islands is preferable to single circuits. Consider a network with IP islands in San Francisco, New York, and Chicago. In a circuit aggregation network there are wide band pipes between the three islands. When traffic, that requires a higher quality of service than the Internet can provide, enters on one of the island and is destined for another, capacity is reserved and all future packets bypass the Internet. The advantages of shared bandwidth over dedicated circuits are that

— the shared facility can take advantage of traffic characteristics, such as the silent intervals in voice, to reduce the average bandwidth, and

— new classes and types of services, that are not related to 64 kbps, can be provided between the islands.

The Internet community is currently investigating reservation techniques to provide quality of service guarantees. Good reservation techniques share bandwidth as on the aggregated circuits, but the bandwidth comes from a pool that is shared with other services, rather than on separate facilities. The advantage of

reservations over separate facilities is that the fraction of the facilities that is assigned to each service can change. Movable boundary approaches to multimedia systems were studied in the mid 80's and their advantages over dedicated networks are well understood. When a reservation technique that provides an acceptable quality of service is implemented, it will replace aggregated circuits.

Active networking is a wild card in the evolution of quality of service guarantees. The instructions that are included in an active packet enable communications between users and transmission facilities, in spite of the IP layer. In a current implementation of IP over ATM an individual user cannot obtain the QoS available in an ATM switch. In an active network, in which ATM switches interpret the instructions in a packet, different grades of services can be provided to each packet.

The quality of service islands that we have described differ from the islands in a current IP network. In the current IP network an island is fixed size and is usually determined by the capabilities of the routers. For instance, as multicast is introduced into the Internet the multicast islands are the routers that can and will multicast. By contrast, the size of a quality of service island may be different for each service and may change with the time of day. For instance, when data and voice packets arrive on an access link from a user the data packets may consider the island to be the entire Internet while the voice packets may consider the island to be a small body on which it can obtain much higher quality. Even for real time services the island size may differ. For instance, entertainment radio can tolerate much larger delays than interactive voice and may consider the island to be larger. In addition, during high use periods of the day islands may shrink, making it necessary for users who had been able to communicate directly to use a tunnel. Locating gateways between islands and the main network is much more challenging in the QoS regime.

4. The IP Platform

An IP platform is a set of tools, services, experiences, etc. that we use to compete more effectively. When we start on the platform we can reach an objective more easily than a competitor that starts at ground zero. The problem with defining a single platform for AT&T is that AT&T is not a one dimensional company. Each of the businesses has different competitors and different objectives.

To build a platform, we should look at our customer bases one at a time, identify our competitors, and determine how we can invest to obtain an advantage. The AT&T platform is the union of the individual platforms. Elements that we are not in the union are not of use to any of our businesses and do not represent a sound investment.

The number of AT&T's businesses should result in a considerable overlap between the individual platforms. The expense of developing platforms in areas that overlap is shared by more than one business and provides an advantage over companies that are in only one of the businesses.

The applications on the Internet are continuously changing. Each platform should consist of stable components and experimental components in order to evolve with the applications. The stable components are the elements with proven value, that we guarantee will continue to be supported, and the experimental components are the elements that we offer to more adventurous, forward looking, users, on a trial basis.

From the invention of the telephone until today AT&T has invented and advanced the technologies that we use. Our engineering can provide us with a source for the experimental components of the platforms, as well as a proprietary advantage over competitors. Platforms that are built solely with publically available parts can be duplicated.

In order to demonstrate what I mean I will look at three sets of customers: service providers, end-users, and the platform itself. Each set of customers has its own platform. For each platform there is an objective and competitors. Based upon the objective and competitors there is a set of standard components, that should be obvious, a set of experimental components, and a set of technology differentiators. AT&T has many

more sets of customers, but these three are adequate for demonstration.

4.1 Services Network

An important set of customers for the Internet are service and information providers. Our competitors are other network providers. Our IP platform should make it easier for service providers to do business if their servers are based on our network. The standard part of our platform should include network servers that perform customer identification, a network notary, an interface to credit card companies, ...

The experimental component could include research servers and protocols, such as CRMP and the reliable broadcast protocol. The CRMP servers improve the quality of multicast video and audio without changing the operation of the sources or receivers. The reliable broadcast protocol trades delay and loss to meet the requirements of customers, such as the New York Stock Exchange.

The technology differentiators could include our most recent attempts at intelligent and active networks. There is a considerable effort to use API's to make the intelligence in routers available to service providers. While the past ten years is littered with failed experiments in intelligent telephone networks, recent advances such as Java, network sandboxes, and active networks may make intelligent networks possible. The first network provider that can safely and economically implement an intelligent network will have a major advantage.

4.2 Consumer Network

The largest set of customers on the Internet are individual users. Our competitors in the Internet are the other ISP's. Our competitors in telephony include the other network owners, the regional operating companies, the resellers, the dial around companies, ... The objective of this platform is to make it easier and safer for the consumers to use the network from our access points.

At present America On Line has an Internet platform. The AOL consumer platform makes it easier for its customers to find information on the Internet and communicate with one another, directly or in groups. Most other ISP's just provide access to the Internet and conventional email.

The conventional part of our platform should include:

— tools that make it easier for users to search for information,

— servers that maintain the large set of passwords users are being required to remember,

— servers that authenticate the service providers that are asking the customers for information,

— ... .

Initially, the conventional part of our platform should include the components that AOL has demonstrated to be useful. The experimental part of the platform will eventually distinguish our platform from AOL, and can include things like proxy processors that protect our customers from viruses. We should also introduce technology differentiators, such as our proprietary information separation techniques that maintain our customer's privacy.

4.3 The Composite Platform

The platforms are the customers of our engineering organization. The objective of engineering is to introduce or improve the services in the platform. The ability to construct our own services and network servers provides us with a proprietary advantage over network providers who do not have in house engineering. If we use our own services without regard for cost or quality, our network will not be

competitive. Therefore, our platform implementation must have its own platform in order to compete with outside providers.

In the Internet the main group of competitors for software services are startup companies. Therefore, our objective is to implement new services more economically and more quickly than a startup company. On a level playing field we would have a very difficult time.

One advantage that we have over a startup is that our implementation is a persistent process, rather than a one time effort. Most internet services have similar components, including WEB interfaces, databases of user profiles, connections with related Internet information sources, ... To the extent that we can reuse software from previous services, we can obtain an implementation advantage over startups. The standard part of our Internet platform is the software components that have been successful and that we reuse. Creating this portion of the platform has all of the problems that we have encountered in the past, such as common interfaces for the software and cataloging the components.

The experimental component of this platform naturally contains any new software or tools that we need to implement a new service. The technical differentiator should take advantage of engineering advances that are not available to our competitors. For instance, most network servers do not take advantage of the fact that they are connected to the Internet. The WEB is a vast store of poorly organized information. We have a considerable investment in the technologies that mine information on the WEB. If we can tie information mining into a service, we can distinguish our service from other services. For instance, instead of setting up ecommerce sites that just offer products for sale we can organize and present information about the products from other WEB sites, in order to help consumers make up their minds.

5. Conclusion

The Internet uses an evolutionary strategy that continuously introduces new technologies and services. The telephone network has optimized the use of technology for a class of services. The telephone strategy results in a more economical solution for a particular set of services and technologies. The Internet is less likely to reach the point where it is obsolete and must undergo a major change. In the current environment, where services and technologies are changing rapidly, the Internet approach is preferred, and is probably more economical in the long run.

There are major differences between the telephone network and the Internet that make it difficult to use the Internet strategy to evolve the telephone network into an IP network. I’ve explained the differences and showed how to modify the Internet’s "island and tunneling" strategy so that it can be applied to the transition from a circuit switched network to an IP-based, packet network.

The basic strategy is straightforward. Everything in the network is packets, except the initial and final connections with current end-user telephone equipment. The telephone network tunnels through the IP network by putting continuous signals in packets and IP islands are connected together by sending packets on circuits. The details, such as maintaining quality of service of voice telephony and performing the application layer services of the telephone network in the Internet, lead to an interesting set of technical problems. Potential solutions have been outlined for many of the technical problems.

The evolutionary strategy is non-committal. Several techniques may be tried on different islands, and unsuccessful techniques purged. The strategy does not require as large an initial investment as strategies that start at the core of the network and invest in a backbone before connecting users.

In our move to an IP network we are introducing a new layer into the simple IP architecture. Our IP platform. The platform gives us a competitive advantage for providing and implementing customer services. One reason for moving to an IP network is to be able to introduce new services quickly. The platform itself must be part of our evolutionary strategy. New versions of the platform can be introduced as

islands in the network. An experimental component guarantees that the platform evolves continuously. The relationship between the experimental and conventional components of the platform, and between the experimental component and technology differentiators, have been described for several sets of customers.