Realtime publishers
"Leading the Conversation"

The Definitive Guide™ To

Successful Deployment of VoIP and IP Telephony

This eBook proudly brought to you by

PROGNOSIS®

Jim Cavanagh

Copyright Statement

Chapter 2: The IP Telephony Life Cycle

Every system has a rhythm, a natural life cycle. If you can understand a system's rhythm and "go with the flow," your job of managing that system will be made much easier. The traditional telephony systems that we are now replacing have a natural rhythm, as do the IP Telephony (IPT) systems with which they are being replaced. This chapter is about understanding the life cycle of your enterprise IPT system and harnessing that understanding to increase the success of your implementation as well as your financial rewards—be those rewards direct tactical benefits such as cost savings or indirect, and often elusive, strategic benefits. This chapter is also about squeezing as much use as you possibly can from your existing telephony systems, a move that can make a lot of business sense regardless of your natural desire to discard the old and to start using the shiny new thing.

The Bigger Business Picture

From the 1960s to the early to mid 1990s, a technology project often took 12 to 18 months or more with little variation in schedule from large to small organizations. Any technology project was a major undertaking and included an exhaustive process of a Request for Proposals, presentations, site visits, and analysis, often by an internal staff and a cadre of outside experts. In many cases today, however, it is only the "Fortunate 500" who still have sufficient resources to be able to take a careful, project-oriented approach and to think and rethink the various outcomes of possible choices or approaches—and often even those organizations will not apply sufficient resources to assure success. These days, smaller organizations are often forced to react rather than plan and to put out fires rather than building a fire-proof system in the first place.

The following section explores the facts surrounding a project that was carefully thought through, argued out, and managed in a thorough and professional manner as an example. The project described was the development of a global Voice over Packet (VoP) blueprint and subsequent implementation for a Fortune 100 client. Although the client company is a global energy giant operating in 162 countries, there are lessons in this story that can benefit organizations of any size, industry, or geographic reach.

After being briefed by the client on the intended scope of the project—a global migration to VoIP—my first step was to define terms and to redraw the scope of the project through a collaborative, interactive, and often heated and animated debate. The initial thinking on the client's part was to create a new world of telephone communication that would connect all corners of their vast global organization relying largely on their existing IP network coupled with a new application called VoIP. Their primary motivation was to retire their old telephony gear and in fact, all their traditional telephony headcount and associated costs. Their drivers were as much a desire to get the benefits of "free voice" as anything else as well as to lower costs while maintaining the same level of voice quality and connectivity.

A key part of the success of the project redefinition effort was the rewriting of the project objective. The objective was originally "to implement a VoIP system and reduce or eliminate costs for internal voice communication" and was rewritten as "provide a global migration path toward a multi-media network that will leverage existing technologies and systems, provide a baseline of voice, data, and video services and allow implementation of important new technologies when and where needed, leveraging existing assets, and to positively contribute to organizational success." Even though the updated objectives sound a lot like a really bad PhD thesis title, the new objectives incorporated a variety of factors that must be included for the long-term success of any project: a consideration of how the project contributes to the long-range success of the organization and best utilizes existing assets.

This project occurred in late 2003, and one interesting lesson learned during this early phase was how much things have changed since the 1990s. For instance, an effort to identify organization-specific strategic and tactical benefits in order to establish milestones and success criteria for the move to a VoP system did not get very far. The Global Network Manager explained that in their world, tactical benefits were called "hard benefits" and hard benefits could be measured in dollars and cents. He further explained that strategic benefits were called "soft benefits," and that there was no tangible way to measure strategic benefits. The project must, therefore, be measured strictly in terms of hard, tactical benefits. As a further historical footnote, he added that strategic benefits had not been considered since the late 1990s.

Packet Telephony Is Inevitable

Packet-based telephony is, in my strong view, inevitable, but VoIP, per se, is not. The first step was to redefine the project as the Voice over Packet (VoP) project, as opposed to a Voice over Internet Protocol (VoIP) project. The main difference is that VoIP originated as a way of allowing two people to communicate by voice between two computers while VoP includes a comprehensive toolkit that goes beyond VoIP to include other services required by traditional telephony, such as signaling, accounting, security, and enhanced 9-1-1 emergency calling. In addition, VoP extends the scope of communications to include a variety of services and capabilities envisioned in initiatives such as the Session Initiation Protocol (SIP) and the Internet Multi-Media Subsystem (IMS).

VoP

The broad term VoIP has been accepted by the industry, and consumers, to refer to a new, inexpensive way of making telephone calls using the Internet. Although "insiders" understand that the IP network is not always the Internet, the common definition is convenient and easy. What VoIP also means is that the voice information is encoded and sent, historically, using proprietary methods or the H.323 protocol—and, increasingly, via the Session Initiation Protocol (SIP). The term VoIP is exclusive and limits the range of choices, however, the term VoP is inclusive and is a more general term meaning the coding of voice into binary and transmitting the voice over a broader range of options—including Voice over ATM, Voice over Frame Relay, and Voice over DSL. By considering a wider range of VoP options, not just VoIP, an organization can often maximize use of its current networks, get around regulatory issues in certain countries, and extend the life of existing assets.

Session Initiation Protocol (SIP)

Early versions of VoIP used proprietary "homegrown" protocols to convert voice waves into bits and package those bits into discrete packets for transmission to a distant PC where they were depacketized and replayed over the speakers, and eventually the headsets of the distant computer. Early efforts to standardize the process led to the adoption of the H.323 protocol, really a family of multi-media call setup and media negotiation protocols, from the International Telecommunications Union (ITU), a global standards body under the United Nations. Eventually, a much more powerful multi-media control protocol, the Session Initiation Protocol (SIP), which can define not only voice but other types of multi-media sessions such as Instant Messaging and mobility, emerged from the Internet Engineering Task Force (IETF), which is the de facto standards body for the Internet. SIP is generally accepted as the primary standard protocol for VoP communications today and for the foreseeable future.

Internet Multi-Media Subsystem

The Internet Multi-Media Subsystem (IMS) is an initiative of the next-generation wireless network architects intended to create a single operational structure that incorporates wireless and wired systems designed for multi-media communications. Although there are other, competing, approaches, IMS is a prime example of a vision for a single architecture that incorporates static, nomadic, and mobile multi-media systems.

Figure 2.1 is very similar to the diagram used to explain the benefits of taking a VoP approach. Analog telephony is the base of the family tree. The importance of starting with analog established that analog still exists in the network, and always will, because the endpoints that are communicating—humans—communicate by modulating analog waves and, at least for the foreseeable future, will continue to. The discussion of analog communication planted many seeds that would bear fruit later. For the traditional "Bellheads," the members of the team who run the old phone network, the discussion afforded a nice walk down memory lane, jogging memories and knowledge all but forgotten, or at least not appreciated. For the "Netheads," the discussion was the first of many reminders that there was much to learn about the world of voice if it were to be successfully grafted onto their IP network.

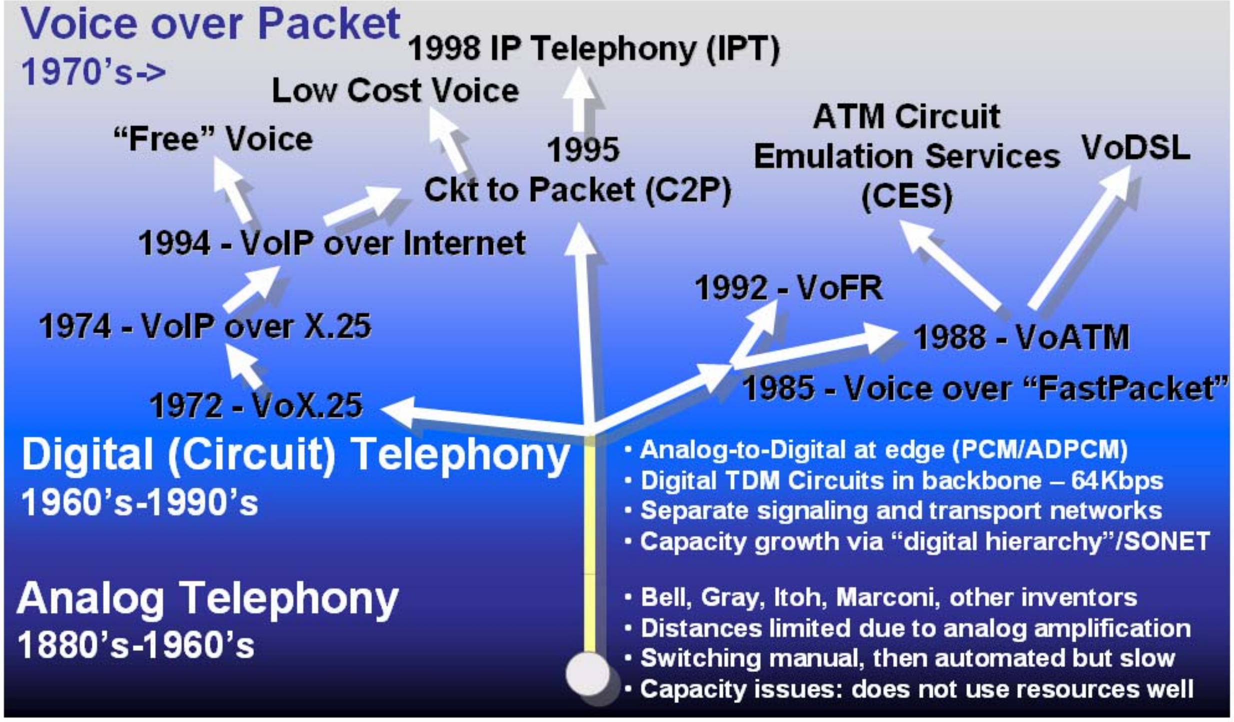

Figure 2.1: The voice technology family tree.

Key points about analog telephony are that it works by translating continuous analog acoustic waves into continuous analog electrical waves and sending the electrical waves over a distance. Transmission distance is limited because the electrical signal becomes "tired" (the correct term is *attenuated*) over a distance and must be amplified to travel farther. The amplification process is not ideal because it amplifies the desirable signal (speech) as well as undesirable background noise. After a number of amplifications, the desirable signal is all but indistinguishable. Analog switching is also slow and problematic and analog systems have capacity limitations. Even so, analog is still used in today's networks, even if just from a desktop or residential telephone to a Private Branch eXchange (PBX) or telephone company switch.

One additional benefit of a review of traditional telephony was to provide a basis for understanding voice quality. Regardless of the many theories, metrics, and measurements, the true measure of voice quality has to do with how closely the analog wave arriving at the destination, after its trip through the phone network, resembles the wave that was sent. This basic fact is true regardless of whether the transmission system is analog or digital circuit or a digital packet system. In fact, an early objective called "toll quality" was based on the quality of the signal at the destination as judged by a panel of human listeners on a 5-point scale called the Mean Opinion Score (MOS) that has been imitated but never duplicated.

After a thorough discussion of the analog era and the knowledge from that era that could be applied to the current project, we moved to the discussion of the digital circuit era. The comprehensive understanding of the digital circuit era that comes from operating a global digital voice network was fresh in the minds of the Bellheads because this is the network that they were operating every day and was the network from which they were attempting to migrate. The Netheads, however, got another reminder of how much they did not know about voice communications and how much there was to know.

In this phase, we discussed the translation of analog waves, at or near the edge of the network, to digital bits using Pulse Code Modulation (PCM) or Adaptive Differential Pulse Code Modulation (ADPCM). We discussed the fact that, unlike voice coding techniques developed for VoIP/VoP systems that are optimized for human speech, PCM and ADPCM are analog-to-digital schemes that work for all sounds. In addition, although the VoIP/VoP systems would give a wide range of clever, bandwidth efficient choices, we can still use PCM and ADPCM in the VoIP/VoP architectures to maintain a higher voice quality. We also reviewed the global telephony network that existed within the organization and found that Private Branch eXchanges (PBXs), literally scaled-down telephone switches designed as the cornerstone of organizational voice networks, connected 64Kbps voice circuits to larger switches, which connected to high-capacity fiber-based connections formatted according to the rules of a standard called Synchronous Optical NETworks (SONET) in the US and Canada and a very similar system call Synchronous Digital Hierarchy (SDH) elsewhere. We also discussed the fact that regardless of the capacity of the SONET or SDH optical systems, they were nothing more than bigger and bigger bundles of the 64Kbps voice connections, at least in the older format in which they existed with this network. In the current network, voice was completely separate from data, video, and telemetry, and the multi-media bits never, ever, commingled within the system.

Among other considerations was the fact that critical signaling information needed to establish and take-down the voice calls traveled in packets in a parallel network completely separate from the voice connections. On the plus side, the parallel signaling networks, governed by rules titled System Signaling 7 (SS7) in the US and Canada and Common Channel Signaling 7 (CCS7) elsewhere, were designed to be secure and virtually indestructible—without them, calls didn’t get connected and phone company revenue did not flow. On the minus side, were the costs of maintaining two complete networks, a circuit network for voice and a packet network for the signaling, just to set up voice connections.

As we moved from our historical inventory to get a glimpse at the future, we had some area of common ground after which the IP folks began to feel more comfortable. The common ground was in the area of Voice over Fast Packet and its various branches. Unbeknownst to many, the global telephony standards organizations had begun work on a packet-based infrastructure for voice as early as the late 1970s with standards such as Broadband-ISDN beginning to appear in workable form in the early 1980s.

The right side of the voice technology family tree (see Figure 2.1) is, in fact, representative of the work of the telephony traditionalists and includes a range of technologies from Voice over ATM to Voice over DSL, including Voice over Frame Relay and Circuit Emulation Services (see Figure 2.2). The left side of the family tree shows initiatives, as early as 1972, to put voice into discrete packets and send them over a shared network with data. These initiatives begin with little-known Voice over X.25 and VoIP over X.25 and include the much more recent, and more memorable, VoIP over the Internet projects in 1994 that were the true ancestors of VoIP as we know it today.

Until 1995, however, progress in packet-based voice was interesting but not of commercial interest. It was, in fact, very much in the hobbyist realm, as we might think of ham radio, for instance, until the invention of the packet-circuit gateway in 1995 made it possible to call from a PC to a traditional telephone. It was not long after that calls from traditional telephones to PCs were demonstrated and then, the service that proved most disruptive to the established traditional telephony companies, calling from a traditional telephone to another traditional telephone completely transparently over an IP-based bypass network. This is where the fun began and this is, precisely, the objective of the VoP migration project: to gain the benefits of a combined voice, data, and video network and bypass traditional telephone carriers for all internal and some external calls while maintaining connectivity to the traditional telephone network to be able to reach persons, be they on stationary traditional telephones or the growing population of mobile phone users, as and when needed.

In addition to the raw technical arguments for or against a specific packet voice option, experience and analysis showed two more important dimensions: legal/ regulatory and operational. From a legal/regulatory standpoint, it is interesting that packet voice is not legal in many of the jurisdictions in which this global organization operates and, even though there has been a broad relaxation of anti-VoIP statutes, there are still nations in the world where VoIP, or at least non-circuit voice communications, breaks certain statutes, violates specific tariffs or contracts, or is out-right illegal. In these cases, the business and technical arguments have no weight against the legal ones and legal and regulatory compliance is still a check-off for any change to the telephony status quo.

As far as the operational dimension of VoP systems, the technology choices were weighed against voice quality, delay, and delay variation issues, price, multi-plexing efficiency and system complexity and some interesting relationships emerged (see Figure 2.2). Overall, voice quality was deemed to be highest in the systems in which all aspects of voice transmission could be most strictly controlled—that is, in circuit systems. The less control and the greater the influence of other traffic, the lower the quality. Considering delay and delay variation, the further you go away from circuits, the more delay and delay variation. Considering these items only, it would appear that circuits are ideal and packet-based systems inadequate—that is, until price and multi-plexing efficiency are factored in. VoIP systems that use bandwidth-saving silence suppression and overhead compression schemes to provide the best multi-plexing, or sharing, with data and video have the lowest overall costs though they are somewhat more complex to operate than circuit-based systems.

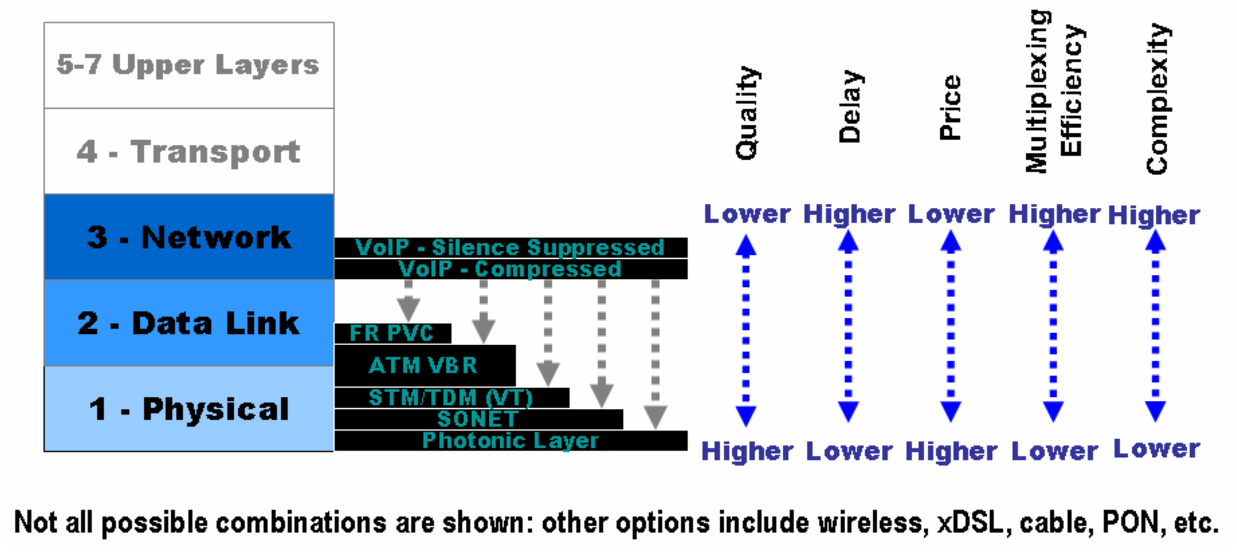

Figure 2.2: Analyzing operational dimensions of VoP systems.

Multi-Protocol Label Switching (MPLS)

Multi-Protocol Label Switching (MPLS) is an approach to networking characterized by this author as "the best of both worlds." MPLS is a way of delivering Quality of Service (QoS) in Virtual Private Networks (VPNs), the secure package in which many organizations are purchasing network services, by allowing a dynamic choice to be made whether to route or switch information traversing a network. Routing, the traditional Internet way of sending information across a network, has the benefit of not having an initial delay to establish a path from end-to-end across the network, but each packet must be analyzed and forwarded individually. Switching, in contrast, the traditional way that telephone networks work, does have an initial delay but whisks each subsequent bit quickly to its destination. The real answer to which approach is best cannot be answered until we know what information is being sent and, more specifically, if it is a long-lived or short-lived flow.

For a short-lived flow, something like a small Web page, for instance, routing is best because each packet is forwarded on its own: there is no need to go to all the trouble and delay of setting up a switched path for such a small amount of information. A long-lived flow, however, such as a longer VoIP conversation or a video over IP program, will benefit greatly from having a pre-established path and not having to go through the requisite packet analysis and forwarding of each packet.

MPLS is "the best of both worlds" because MPLS is capable of routing all traffic until it has determined whether a flow is long-lived or short-lived and setting up a switched path for those flows that turn out to be long-lived. Early implementations of MPLS would always route over the available IP routing mechanism and switch over Frame Relay, but newer variants are capable of switching over everything from ATM to Wave Division Multi-Plexing to SONET/SDH to dedicated fibers, including Frame Relay and the IP RSVP bandwidth reservation protocol.

In the case of the global energy giant, the impact of this analysis was felt more on the prioritization of systems to migrate and a heightened comfort level with leaving some non-IP systems in place for a longer period of time during migration, but had no appreciable impact on the ultimate objective, which was still to eventually have an all-IP multi-media architecture. This analysis did, however, support their decision to migrate much of their global IP traffic, both voice and non-voice, over to a managed MPLS-based VPN to get the combined benefits of IP-based applications, QoS, and security offered by such VPNs.

In other instances, however, a broader look at VoP caused several large organizations to rethink their VoIP objectives. In many cases, including a North America-wide system for a major manufacturer of industrial tools, the organization realized, prior to embarking on a traditional VoIP approach, or early enough in the cycle to retreat and regroup, that they did not really need the power of multi-media or even magic of SIP: what they truly desired was less-expensive voice communications at a level of quality comparable to, or just less good than, their traditional dial network. These organizations often implemented Voice over Frame Relay or, less frequently, VoDSL or Voice over ATM.

Impact of VoP

Having considered VoP almost in isolation, the next step was to consider the impact of VoP on the existing network. The existing network was a predominantly data network with a rapidly growing set of video services. The data, however, includes the traditional menu of services—browsing, email, file transfer—and the movement of gargantuan volumes of near-real-time telemetry data that most network operators do not have to deal with.

The initial "VoIP is just another IP application" mantra was quickly thrown out in favor of the "VoP is a unique application with its own special requirements" mantra. One reason for this change of mantra is that early, especially residential and consumer flavors of VoIP, used very low-bandwidth compression algorithms which, while bandwidth efficient, provided a lower quality, or at least a different quality of voice. Many businesses have opted to continue to use the traditional Pulse Code Modulation (PCM) voice coding scheme even after their migration to packet voice. Keeping PCM end-to-end, regardless of whether the bits are traveling in circuits or packets, assures better compatibility between different systems and reduces the number of translations between different coding schemes, which also improves voice quality. What this means in operation is that instead of 6 to 8Kbps per voice connection, each connection consumes 80 to 110Kbps of bandwidth, though the actual utilization may be considerably less if silence detection systems are used to not send packets containing silence.

Another impact consideration is that voice systems can provide less delay, less delay variation, and less susceptibility to packet loss by putting fewer voice samples in each packet. Although this approach does have its benefits, as stated earlier, it also has the drawback of taxing routers and transmission systems by requiring a lot more work for a lot fewer bits being transmitted—and all this is going to have to be considered when assessing the true impact of adding telephony to the network.

VoIP and IPT Return on Effort

The financial measurement of Return on Investment (RoI) is a calculation of the financial benefit achieved for a specific investment of capital. An additional metric that takes into account intangible—and, yes, admittedly, soft/strategic benefits—is Return on Effort. RoE had no value to the global energy company, but many organizations find a consideration of RoE a good way to take into account aspects of a move to a VoP system that cannot be captured any other way.

For instance, in one of the case studies reviewed in Chapter 1, Heritage Bank, the financial RoI calculation was sufficiently terrible as to scream to anyone who would listen "YOU SHOULD NOT HAVE DONE THIS!" but the RoI calculation neglected the fact that the system being replaced by the new VoIP system was no longer available from the manufacturer, spare parts were no longer available, and that the bank had the choice of a forklift upgrade or ceasing to add staff, branches, and new clients. For the bank, as for many companies adopting VoIP, it is more expensive than the current system and an upgrade does not make financial sense, but there is no other option. RoE is also a way of identifying other intangibles, such as the school district that was able to provide in-room phones for teachers where none had previously existed or the pharmaceutical lab that was able to display employee and contractor photos on phones for enhanced security.

Enterprise Telephony Migration

At this point, it has been determined that the organization will migrate to a new packet-based telephony system. What lies ahead is the inevitable budget process, the migration itself, and, once across the bridge, the IPT life cycle of planning and assessment, pre-deployment testing, ongoing operations, and optimization.

The Budget Process

For the global energy company, as for many organizations giddy with excitement at the prospect of finally being able to retire traditional voice assets and headcount and realize the benefits of "free voice," the excitement disappears quickly as the budget process begins.

The first realization was that the skills of the traditional telephony people are still in demand and, in fact, are becoming more unique and valuable than their IP counterparts due to the natural attrition of the traditional skill sets. In the case of the global energy company, and many other companies, the following type of plan was anticipated—as soon as the first packet carrying voice flowed from one IP phone to another, the un-needed old telephony regime could be shown out of the building and their cost saved. What actually happened was that the best of the traditional voice and IP team were kept on staff and the less useful, less-skilled individuals from both groups were terminated, though nowhere near 50 percent of the staff was cut. For this situation, it turned out that it took about 70 percent of combined prior staff to run the new, consolidated network, though it was felt that cross-training over time to create a combined data/voice/video multi-media support tech could provide greater than 30 percent savings on staffing costs.

The second realization was that packet voice requires special hardware. The new hardware would take the form of either new, specialized phone devices—at roughly the cost of earlier digital and PBX phones—or gateway systems. The gateway systems would be needed to adapt the traditional telephone sets to the new IP connection. In most cases, some of each new system would be required, at least through some migratory period. This is an additional, and often unexpected, financial burden on the organization, but a surprise that can be avoided with good planning. Another related surprise that can be avoided is the fact that new phone sets must be provided with AC power or powered via Power over Ethernet (PoE), both of which require additional planning and costs.

> In some cases, organizations spend as much, or more, than they spend on the new voice over packet system to provide power to the phone systems because, unlike their analog predecessors, the new systems rely on "wall power."

The third realization is that, like their predecessors, IP-based multi-media systems are complex systems that require management, measurement, troubleshooting, and problem resolution—in effect, life cycle management at each step through their operational cycle. In fact, the management required by traditional telephony systems is often hidden "behind the scenes" because dial tone has often been purchased as a service rather than a bunch of piece parts to be assembled by the end user organization, such is the shift in thinking from telephony as a service to telephony as a managed application. But do not be fooled, there are important management functions that must be completed routinely at each step of the life cycle as will be described in more detail later.

Reuse of Existing Assets

In the ideal situation, existing assets can be reused. This is true of human assets that can be retrained or cross-trained in the needed disciplines. Hardware assets can often be reused as well, with the addition of a translation system or gateway. One benefit of gateways is that they can provide access to voice connections on existing networks as well as special features provided via the Advanced Intelligent Networks (AINs) that traditional carriers have built over the years. The physical positioning of the gateway also impacts the cost and ability to reuse existing assets as well as the ease of support of the new system. Consider Figure 2.3. It might be possible, for instance, to assess the needs of the user community and divide users into three groups: those users who need new IP capabilities now, those who will need IP capabilities at some time in the future, and those who simply need traditional dial tone.

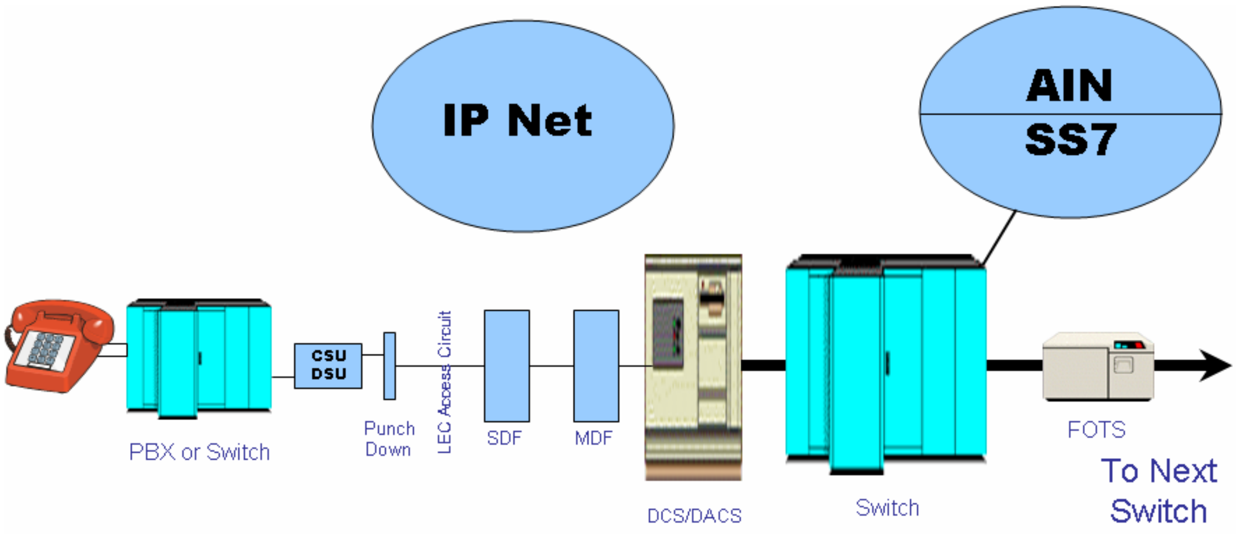

Figure 2.3: Determining which assets can be reused.

The first group would be provided with new IP phones and directly connected to the IP network. Because this group is a subset of the entire user population, however, the cost of the new phones is a fraction of what it would be to provide new phones for everyone, and the resource impact on the IP network is also a fraction of what it would be if all users were migrated. Some scheduling algorithm could be applied to the second group to determine at what point they would be migrated to the new IP system, and the third group could simply be provided with "personal gateways," often called Analog Terminal Adapters, which would allow their traditional telephones to be connected to the IP network. They would still use bandwidth on the IP network, but the hardware cost and complexity of putting this set of users on the network would be substantially lower than moving all users to new IP phones. At this point, the overriding consideration might well become the calculation of the cost of maintaining the old traditional network just to support those users who had not yet needed to migrate, and it may be advisable to complete the migration sooner, rather than later, and retire the old network.

However, if a larger gateway was positioned on the network side of the PBX or premises switch it is possible to reuse all the premises phone systems and wiring and attach all existing telephone users to the IP network via the gateway. This may be an interim step until a local shared IP softswitch is implemented or until a remote IP-attached softswitch is implemented. In any case, it is an approach that can be used to extend the life span of the premises voice solution.

New Assets Needed

In any case, new assets will be needed in order to extend the life of the old system and to bring new features and capabilities to the users. New assets will include, at minimum, hard or soft phones, gateways, softswitches, MPLS-capable routers and MPLS-aware end systems, Virtual Private Networks (VPNs) built with those tools, and additional bandwidth. New assets will invariably include new management, monitoring, and troubleshooting tools as well as tools to ascertain compliance of services with Service Level Agreements (SLAs).

Assessing Urgency and Prioritizing Migration

Except in the smallest, single office situations, it is not advisable to flash cut all telephony users to a new system. For this reason, it is necessary to assess the urgency and prioritize which areas of the organization will be migrated in which order. As difficult as it might seem, political pressure is the last thing that should be considered when determining the order of priority for migration. In fact, the most visible or boisterous departments should be migrated last to assure a minimum negative impact on the overall migration. Migration order should be based on more reasonable considerations:

- Which office or region has outgrown, or is closest to outgrowing, their existing system?
- Which region has bandwidth or IP network infrastructure good enough to carry voice, and other resources to accommodate the new telephony upgrades?
- Which region has the most technically savvy support staff?
- Which region was involved with early testing and/or staging and is, therefore, most familiar with the new system?
- Which office or region is closest, because the earliest migrations are always learning experiences and should be expected to go slower and be more problematic?

Migration

Migration can be divided into three categories: avoiding migration, partial migration, and full migration. Within the partial migration and full migration categories, the roll-out may be phased based upon a variety of factors, as will be discussed later.

Avoiding Migration

Avoiding migration is usually "avoiding migration *for now,*" which means that the office, division, or region under consideration either has a recently acquired system that has not yet been depreciated or that the group will be sold or otherwise liquidated before it is required to join the new VoP network.

One strategy adopted by the global energy company that not only eased migration but also stretched the use of their assets, thus reducing their overall budget, was to establish regional refurbishing and staging centers for equipment that had been taken out of service as a part of the upgrade. As an example, let's say that an analog key system in Tokyo no longer had expansion capability because the upgrade components were no longer manufactured and that this is the only reason why Tokyo was being upgraded to a new VoP system. Sydney had exactly the same issue and the same key system. In this case, Tokyo might be upgraded to the new system, thereby freeing components that could be used to upgrade Sydney and avoid migration to the new VoP system, at least for now, or vice versa. If there were no immediate need for the components removed from Tokyo, they could be refurbished and placed into an upgrade inventory for the time when they would be needed to upgrade a system and avoid migration.

Partial Migration

Partial migration is accomplished as previously discussed by using strategically placed gateways to allow existing assets to be kept in place, thereby extending their useful life, especially where only a limited subset of the new VoP capabilities were needed by the office or region being partially migrated.

Full Migration

Full migration means that all non-VoP components are removed and replaced with their VoP counterparts.

Delaying VoIP/IPT

It might seem a bit late in the process to ask this question, but can a move to VoIP/IPT be delayed? One reason why it is not too late in the process is that up to this point all the planning exercises are what military planners call “tabletop exercises”—no changes should have yet been made, no systems implemented. Everything up to this point should be on paper, but not yet realized in the real world.

Thus, what are the pros and cons of delaying a move to VoIP and IPT? If there are no competitive or operational pressures to make the move, the organization should consider delaying the move. Each day that goes by, prices for components and services get lower, standards solidify further, and the various aspects of implementation and operations are better understood. For each day that passes, traditional telephony components are further depreciated, and for each day that passes, there is one more day without the potentially unnecessary disruption to operations that virtually any change, especially the change to a system as fundamental as your phone system, can bring. It is only prudent to stop at various times in the planning process to revalidate your objectives and to consider any change in business climate that would mandate either slowing down or speeding up the migration process.

Systems Integrator vs. Do It Yourself

Once the decision has been made to move to VoP and it has been determined the extent to which it will be rolled out in the organization, it is necessary to confront the issue of phasing and project management. Some organizations choose to migrate entire departments or offices to VoP, while other organizations will phase-in VoP by job function—for example, phasing in support staff first to allow them to gain experience with the technology. Once the phasing has been determined, a project management schedule can be developed. It is at this juncture that many organizations realize that they do not have the resources to do the task. Smart organizations realize that if they are resource-constrained before a VoP rollout project, they must line up additional resources before they can begin their VoIP rollout.

Some organizations attempt to be clever and hire and train college interns or other “temporary” workers to assist during the rollout. Although the concept is right—workers will be in high demand during the rollout but will not be necessary after migration—the approach is wrong. A key consideration in this scenario is that knowledge gained during the rollout should be captured in some way because it is valuable to the support process. For this reason and others, many organizations consider engaging a systems integrator to assist in the rollout.

When an organization weighs the systems integrator vs. DIY approach, the organization should remember that the VoP migration is not simply a matter of plugging VoIP phones into existing Ethernet ports and making phone calls. Instead, it is a highly complex, and highly visible, project within the organization. The systems integrator should be evaluated based upon demonstrable success, site visits to the systems integrator, and interviews with project managers and other key persons. In addition, the systems integrator should be involved as early as possible in the planning phase.

It is also important to consider the systems integrator vs. DIY debate as early in the migration planning as possible so that the systems integrator may be considered in the budgeting process. Although it is not possible to predict prices without more detail about an individual project, it is not uncommon for an excellent systems integrator, who can assure project success, to cost between twenty and thirty percent of the overall price of a project.

Managed Service vs. Do It Yourself

For small and midsized organizations, voice has largely been a "managed service" that has been provided by a telephone company with the occasional small to medium enterprise venturing into the "do-it-yourself" voice business. Large to very large organizations have been quite the opposite and many have, since the late 1970s, provided their own dial tone via PBXs or technical arrangements with similar names. As we move into the multi-media era, it seems as though many organizations are viewing VoP as a new application for their IP intranets and Internet networks and, therefore, increasingly as a do-it-yourself project. But, this does not need to be the case, and, in fact, the migration to a packet-based voice system may very well be the opportunity to rethink the way voice is delivered within the organization and between the organization and the outside world.

If an organization does decide to choose a managed service approach, they can ignore many of the technical details of the inner workings of the system and instead "take the temperature" from the outside and ask the doctor how the patient is doing. This can be accomplished using service provider-offered tools that provide a snapshot of the network or even embedded measurement systems in the customer's own equipment that might give a more accurate, or at least more timely, report on the health of the overall system and its individual connections. In any case, these tools provide nothing more than a second opinion that might—or might not—correlate with the health ratings offered up by the service provider.

If an organization does decide to take the do-it-yourself approach, the implementation of VoP services will require the use of measurement and monitoring tools that reveal far more "personal" results—results that not only indicate how the system is performing but also what the organization needs to do to repair any problems or, hopefully, the level of pride they may take in their own good results.

In either case, the measurement and analysis tools are useless without baseline measurements and a set of rules that typically take the form of a document called the SLA—and even though the SLA for years has been thought of as a way of assuring compliance of the service provider, and it still is, the definition of service provider probably needs some updating. In the managed service scenario, the service provider is still who you think it is—"them". But, in the do-it-yourself scenario, the service provider has become "us," because you are now an internal service provider and the internal departments, divisions, regions, and employees are the consumers of this service—they are the customers.

A further consideration is the interplay of multiple SLAs from multiple parties. Take, for example, a scenario in which a managed service provider is providing VoIP service on a MPLS VPN provided by a separate provider. In this case, the two SLAs are inextricably interwoven. Can you penalize the VoIP service provider for not connecting calls on an underlying IP network that cannot connect across the MPLS network, or if the MPLS network does not meet specific QoS metrics? And what if there are multiple MPLS networks provided by multiple providers? This is an issue often seen with larger, multinational organizations. These items must be considered to achieve a reliable and fully functional VoIP implementation. It is also possible that these items have already been considered, and resolved, in some other networking context, so it is worth researching before diving in.

Although the network is evolving into a complex creature delivering real-time, near real-time, and non real-time voice, data, and video, the SLA must still be written in terms of metrics that are meaningful to the specific services being delivered by the network to the user/consumer. Let's take a closer look at this critical management tool—the SLA.

SLAs

One of the key items that will require revisiting, and refreshing, is the SLA. The SLA is a wonderful tool the full potential of which has never been realized. In its most basic form, the SLA is a contractual vehicle which, at the very least, documents the minimum level of service that a customer is willing to accept before they are due some penalty or refund on their service, but it can, and should, be so much more. It can, and should, also provide a description of the actions taken by a carrier or service provider should the user not keep up their part of the arrangements, such as using a service class for which they have not contracted.

One possible model that might emerge is that transport becomes a commodity over which differentiated services flow. In this model, the prior performance of various transports, representing some combination of different Class of Service (CoS) offerings from one or more transport providers, is known and the closest CoS to that needed by a given service at a moment in time can be selected, possibly on a "pay as you go" basis. An alternative to that would be that bids would be requested, in real time, from available transport services to provide a certain CoS at a moment in time. How would this work in practice? Simply recall the example in Chapter 1 of the call from a Milan coffee shop.

Penalties

Most SLAs are lopsided, especially when it comes to penalties. In most cases, the SLA is written in such a manner that any shortcomings on the part of the service will result in penalties from the service provider. However, the method of proving that penalties are owed is notoriously difficult and only the most exacting and picky end user organizations routinely see any penalties. Both of these problems must be addressed.

SLAs should impose a penalty on the service provider (be they external providers or an in-house group) if they do not perform; in addition, SLAs should impose penalties on the user organization if they don't perform, such as if they continually ask for bandwidth beyond the contracted rate or impose a special availability requirement for a certain location. Finally, SLA compliance must be automatic. At the present time, too few carriers and service providers have embraced the SLA as a strategic customer relationship management tool and require the customer to jump through hoops to get the penalties they are due. SLAs and the compliance information supporting them should be available to the customer when and as needed. We will see a lot more of this in the multi-media future, and sophisticated monitoring and management tools will be required to keep customers apprised as to the QoS delivered by a network versus what was promised and contracted. In fact, the educational process that goes along with this effort should be initiated as early in the process as possible because the process of educating the user on SLA compliance is also the process of educating the user on how to select multi-media services without relying strictly on price.

True Impact of SLA on Business and Operations

A fully integrated and automated SLA-based network monitoring and management system can provide both long-term and "canary in the mine" types of visibility to overall and service specific operations. In an environment of enlightened collaboration between user and service provider, the SLA is a key tool to successful operation and a path to continuous improvement. SLA metrics can, and should, be modified over time to reflect actual network objectives and realities.

The IPT Life Cycle

The four phases of the IPT life cycle are planning and assessment, pre-deployment testing and implementation, ongoing operations, and optimization. These four phases are discussed in this section.

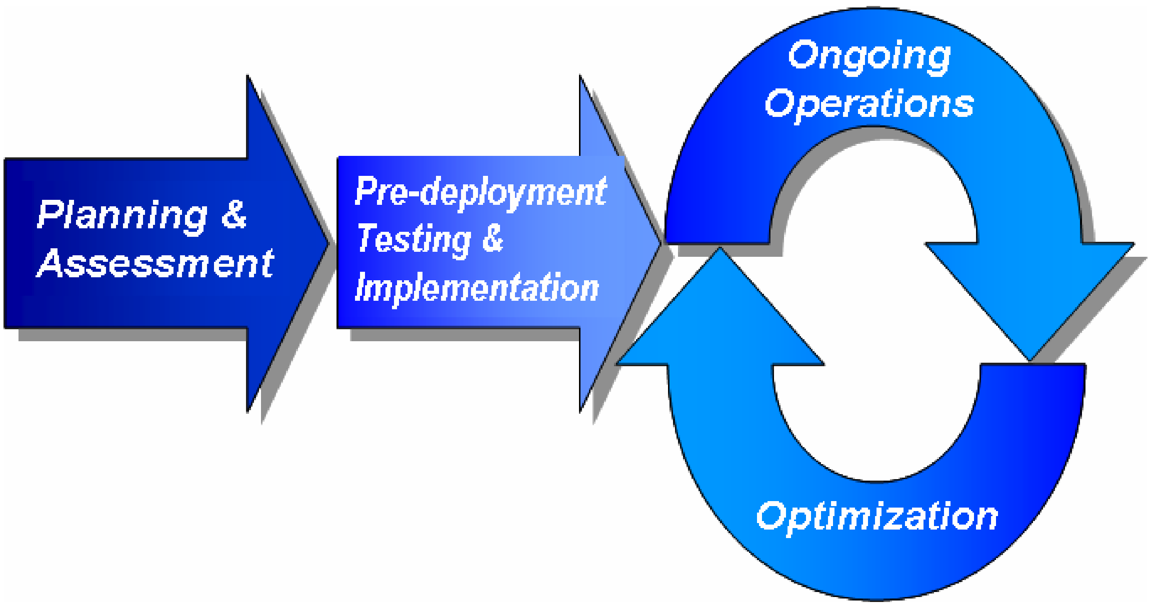

Figure 2.4: The IPT life cycle.

Planning and Assessment

The objectives of planning and assessment are to develop a replacement system that initially delivers replacement functionality and eventually adds important strategic functions. The migration process is characterized by a new, fresh commitment to traditional telephony characteristics such as reliability, security and fraud control, and customized services deployment as well as making desirable new aspects, such as a multi-media focus, available organization-wide to all users who need them. All this must be taken into account in the planning and assessment phase. It can also not be stressed enough that management expectations must be set so that a sufficient amount of time and resources can be invested in the planning and assessment phase to make the phases that follow it successful. VoIP is not just another application that can be launched by plugging in a few IP phones. A proper planning and assessment phase will take all of the needed factors into account.

To be successful in your next-generation network venture, expectations should be set as follows—only in so doing will the results you achieve have any hope of being judged a success when compared with the expectations you set:

- Users—Different is not bad. The new system will initially provide all the important functions of our current system and then, after the baseline is established, and the migration is complete, will provide important new capabilities. The system will sound different but this is not necessarily a bad thing. Our system voice quality will be ____________(not as good as, nearly the same as, better than—fill in the blank based on your objectives) and the new system will be more flexible and give you more freedom while allowing you to stay connected.
- Management—The cost of our new network, when taken as a whole, may be the same or higher than what we are paying now, but we must move forward to assure the continuity of our voice communications and eventually acquire advanced capabilities of strategic benefit, though our unique combination of geographic coverage and scale may allow us to realize overall raw cost savings. There will be four specific phases, each with its own specific objectives. Investing more time, effort, and attention in earlier phases will lower the overall costs and increase the value to our organization. The phases are planning and assessment, pre-deployment testing and implementation, ongoing operations, and optimization. Contrary to popular belief, VoIP is not "free voice" and VoIP is not "just another application."

Organizations adopting multi-media will not, in most cases, actually 'move' at all. In most cases, they will already be providing or using traditional and/or VoIP-based systems and will simply transition to multi-media as non-disruptively as possible. The first step is to ascertain the current services and features that are being utilized to assure their accurate reproduction in the new system. The next step is to move as quickly as possible through the transition phase, though many organizations would be well advised at this point to skip the VoIP step and move directly to the ultimate converged IP-centric multi-media IPT phase. This step avoids the risk of being in the transition period and reduces negative user impact.

A critical aspect of planning and assessment is ensuring that the new system performs at minimum the key functions of the system being replaced and that transitioning to the new system does not present any unexpected and unpleasant surprises. One way to do this is to do a functions audit on the old system in the planning phase and even prior to acquisition of a new system as a precursor to defining mandatory functionality during procurement.

The obvious approach would be to ask the systems administrator or manager of the old system which functions must be supported, but the result of this request usually produces an operations manual or sales brochure listing all possible functions. It is not desirable, for numerous reasons, to attempt to replicate all functions of the old system. A reasonable middle ground is to print colored cards with all the key functions from the operations manual and a blank for "write your function here" and distribute the cards to the users. When a telephony user performs a specific function, they write their phone extension number or other identification and check off the function performed. This would allow an audit of used phone functions to be performed. An alternative to this would be if your PBX, phone switch, or phone service had a reporting function that identified, at minimum, which features were programmed and, at best, how often the function was used and by which users. After the audit, it might not be possible to guarantee that the new VoP system will perform all the functions of the traditional system, but it will be possible, in advance of migration, to identify potential problem areas and workarounds.

Pre-Deployment Testing and Implementation

The objectives of pre-deployment testing and implementation are to assure the outcomes from the planning and assessment phase and to determine the impact of adding the new applications to the existing network—and then to rollout the new system into the organization, assuring that the systems are operating in compliance with the SLAs. The use of sophisticated, specialized modeling and planning tools in the first two phases will allow you to assure the maximum benefits in the shortest amount of time with the least possible disruption to your operations.

Ongoing Operations and Optimization

Onging operations and optimization are two phases with opposite objectives. It is the objective of ongoing operations to assure consistent functionality across the network, regardless of where an individual is in the organization geographically. Ongoing operations is constantly measuring the network's performance against established benchmarks to assure unwavering compliance and absolute consistency of network services with SLAs. Optimization, in contrast, has as its primary objective improving operational aspects of the network services and adjusting the operational baselines to reflect new needs and requirements. Examples might be to improve voice quality and update SLA levels, which are operational improvements, or to decommission PBXs or IP gateways, which will provide a cost benefit.

Although ongoing operations is about consistency, the optimization phase is about a regular, constant, ongoing program of positive change in network operations. Optimization takes the measurement and SLA compliance statistics gathered during ongoing operations and seeks ways to improve those statistics with specific target outcomes. Desirable objectives will include financial or operational goals. The process will involve conceptualizing improvements and validating assumptions using modeling or simulation tools and then testing them in a laboratory or isolated network segment before implementing the changes in the actual network. After network performance has been assessed and any needed fine tuning done to the changes, the network operational benchmarks are adjusted and new optimization goals are established.

Managing the IPT Life Cycle

Now that you have a better understanding of what the IPT life cycle encompasses, you must now develop a better sense of how to manage the life cycle. This guide will provide much more insight into the management process in later chapters but here will take a look at the tools that must be in your toolkit in order to meet the challenge.

Enterprise-Accessible Tools

If you have opted for the "do-it-yourself" approach, you will have access to sophisticated tools embedded in the routers, switches, gateways, end systems, and other components of your networks. Even if you haven't, there are still a rich set of tools that can be used to provide "the big picture" as well as pinpoint areas that require a more granular focus.

MSP-Provided Customer Network Management

Managed service providers (MSPs) allow end-user organizations to view thin slices of their network: the slices that are important to the transmission of the customer's information, masking information about other customers and providing a true VPN view. This capability is called Customer Network Management (CNM)—allowing the customer to see their part of the shared virtual network—and has been around in some shape for well over two decades. Early CNM systems were periodic summary snapshots showing not-so-very-near-real-time (NSVNRT) statistics that bore little resemblance to the real world or the SLAs, while more current systems, especially those based on MPLS VPNs, can provide darn-near-real-time (DNRT) views down to a connection level and do, in many cases, allow customers to adjust certain characteristics, such as CoS or available bandwidth, in real time. Although not ideal, CNM is a valuable tool for end-user organizations that must rely on third parties to provide and manage their service.

Manufacturer-Provided Systems

Manufactures provide the most accurate, most granular, and seemingly most desirable systems embedded within their hardware. The tight coupling of hardware and operating software provides the finest possibly granularity of analysis but, in fact, tends to be limiting and confounding unless the entire network, from the physical layer transmission systems to the application servers themselves, are provided by a single manufacturer and, of course, that is rarely the case. The proprietary, manufacturer provided, tools do have their place, when optimizing a network, for instance, or at higher levels of escalation for system-specific trouble shooting, but lack the broader, more standardized and less platform-specific capabilities of systems provided by third parties. It is also true that third-party systems lack any vendor bias and can be counted on to be more impartial when a system is comprised of components from multiple manufacturers or as a good "second opinion" in those rare cases that a system is constructed completely from the components of a single manufacturer.

Third-Party Tools

Third-party tools for modeling, monitoring, measuring, and managing are invaluable in that they provide broad, cross-platform visibility into aspects of network operation that far exceed what humans alone can do without those tools. The consistent and proper application of these tools at various stages in the life cycle of the multi-media network spells the difference between success and failure. The selection of these tools is the subject of the third-party tools selection report card described in the next chapter and their application to the life cycle of the network are described in further detail in subsequent chapters.

Importance of Modeling, Monitoring, Measuring, and Managing

As important as the right tools to dimension and manage emerging multi-media networks are, they seem to go against conventional wisdom about how a multi-media IP network is implemented and managed. Time after time, the biggest and most visible example of IP networks today, the Internet, has been described as "plumbing," or a common conduit for any type of information, and the term plug-and-play has been used over and over to the point that it has become ingrained in the subconscious, and almost into the genes, of 'Net users and managers who must make budgetary decisions. "Free Voice" is another obstacle that must be overcome in order to be granted the funds needed to operate a network that is, consistent with another metaphor, "the nervous system of an organization."

What management must understand, and subsequently follow with appropriate funding, is that the multi-media IP network, be it Internet, intranet, or extranet, requires special tools to build and maintain its health. Experience has shown that an initial cost of 5 to 7 percent for procurement and training, with an ongoing 10 to 12 percent of operating budget dedicated to staffing, ongoing training, support, and network optimization will pay substantial dividends in network reliability and cost reductions. Proper funding of personnel and systems, will ensure that organizations reap all the rewards of their move to a new voice platform.

Put in its proper perspective, the shift from traditional telephony to its packet-based replacement is like the differences between a 1967 Chevy and a 2007 Lexus. Both have four tires and a steering wheel and are designed to get you from one place to another. But, what is "under the hood" varies dramatically. Your '67 Chevy was so simple that Floyd down at the service station could pretty much provide any service you needed. In fact, the term "keep it humming," relative to automobiles, came from the fact that much of the diagnostic process involved listening to how the car sounds. The VoP system is much more like your 2007 Lexus. It requires specialized service at a Lexus dealer or service point that is capable of utilizing the diagnostic interface, and Floyd at the service station wouldn't have a hope of understanding the sophisticated, computerized guts of this vehicle. If you have an accident, a signal is sent to a customer care center that calls you and will dispatch 9-1-1, an ambulance, and service vehicle within seconds, even if you are incapacitated. And what about locking your keys in the car? Well, the car can be unlocked from the satellite, as well. As much as the Lexus may resemble the Chevy in form and function, it is a fundamentally a different creature with different capabilities, different management requirements, and different maintenance needs. So, too, is the case with VoP and traditional telephony.

Summary

This chapter has begun the practical job of looking at the key elements that further set the stage established in Chapter 1. It has also begun looking at the key phases of the life cycle of the packet-based telephony system and how to harness your knowledge of the natural rhythms of that system. The next chapter will continue to dive deeper into the life cycle.

As promised, the tone and intent has shifted from one of establishing a baseline understanding of some broad, abstract concepts, to providing specific, actionable items that you can perform for your own situation. The next chapter will continue this approach.