Realtime publishers

"Leading the Conversation"

The Definitive Guide™ To

Successful Deployment of VoIP and IP Telephony

This eBook proudly brought to you by

Jim Cavanagh

Copyright Statement

Chapter 6: Ongoing Operations

Most, if not all, successful start-up companies are founded by creative, energetic, entrepreneurial individuals whose 16-hour days, eating at the desk or on the run, decisions on-the-fly and out-of-the-box thinking that could not possibly be sustained over the long run, are replaced by a new team. After the "start up" phase, top management positions are occupied by less creative, more down-to-earth, and more operations-focused personnel whose job it is to standardize operations and run the company going forward. Companies with foresight often move the founders into key roles of business development or product improvement. Those companies that do not do so often stagnate without the enthusiasm and commitment of the original team.

This transition happens as often in Voice over IP (VoIP) projects as it does in entrepreneurial start-ups and for many of the same reasons. The job of creating a project, selling it to management and users, and making it happen—in any size organization—is different to the job of running the system day-to-day. It is the rare individual who can fill both roles. If your job is to take the reins, create harmony from chaos, and run the new system day-to-day, this chapter is for you.

Network Operation

The first step as you emerge from the implementation phase and enter the operations phase is to go back to the beginning and review the original plans, goals, and objectives and to ask several questions:

- Are the original goals and objectives still valid?
- Have technologies changed significantly?
- Are standards more mature, and does it matter?
- Have laws or regulations in any of the areas in which you operate changed in any important ways?
- Have your operating procedures or business environments evolved in a way that impacts your new telephony system?

It may have been several months or longer since the VoIP project was first conceived and planned. Things may have changed: the business direction, the services your users need, and costs, technologies, and products. This review step should also include input from management and users to ensure that all current views are considered.

Once you have validated the original goals, or modified them to reflect any changes, it is time to review the Service Level Agreement (SLA), or agreements, to ensure that they reflect the original goals and objectives and to make adjustments to the SLAs, if needed, in a similar manner to what was done to the project objectives. Once this administrative step is completed, it is time to schedule a review cycle to periodically assess the validity of the project goals and set new goals based upon the network optimization process that will be ongoing in parallel with the operations cycle.

The important question is, "How often should operations be reviewed?" Some organizations are fairly static and have, in fact, been doing much the same job in much the same way for many years. Those organizations might very well be able to have a comprehensive review scheduled annually if one is not triggered sooner by the optimization process. Other organizations that are more dynamic and are undergoing more rapid change—or growth, moving into new markets, evolution of processes and procedures, or similar changes—may be at the other end of the review spectrum and might, at least for the first year or so, want to conduct a formal review monthly. The timing is really up to the organization and their needs, but the key is to get it on the calendar and move on to operations management. Using the SLA as a guideline, take a closer look at network capacity, reliability, security, voice quality and call quality, growth management, and how to measure and monitor the network.

Capacity, Performance, and the Broadband Flip

Anyone who has been in networking before the start of the low-cost bandwidth era is very focused on capacity issues. The industry veteran most likely has a genetic predisposition to focusing on optimizing bandwidth for purposes of capacity and cost savings. The reason for this is very simple: prior to the "broadband era" bandwidth was outrageously expensive and bandwidth demand was calculated and optimized for capacity in much the same way conveyor belts in a factory are sized: getting the most done with the least resources by filling the conveyor system up as full as possible. But, for several reasons, this is no longer the case and you must now do the "broadband flip"

For lack of a better label, I will refer to what came before the broadband era as the narrowband era. The narrowband era was characterized by individual circuits, or channels, that were dedicated to the use of a single communication, be it a stream of voice, data, or video bits or packets.

Bandwidth Demand Projection: Sample

growth: 10% *per month*
circuit order cycle: 3 *months*

Month	Bandwidth	Units	
January	5.20	kbps	(Measured)
February	5.72		(Projected)
March	6.29		(Projected)
April	6.92		(Projected)
May	7.61		(Projected)
June	8.37		(Projected)
July	9.21		(Projected)
August	10.13		(Projected)

Figure 6.1: Narrowband bandwidth calculation.

In narrowband, a general rule of thumb, the two-thirds rule, or a similar approach is often applied. When utilization grows to about two-thirds of the capacity of the circuit, or channel, it is time to request a larger circuit. If a network manager budgets properly, tracks the rate of growth, and knows the installation lead time for a new or upgraded circuit, the management of this process can be fairly straightforward. For instance, as Figure 6.1 shows, if a remote site is:

- Connected to the central data center by a 9.6Kpbs circuit that was averaging a utilization of 5.2 kbps in January.
- Utilization was growing at a rate of 10% per month.
- It takes 3 months to order, install, and test a new 19.2Kbps circuit.

In this case, simple math shows that the new circuit should be ordered no later than April 1st, and probably sooner, to be up and running by August 1st, before demand exceeds the capacity of the 9.6 kbps circuit. Lacking unforeseen growth or upgrades or changes unreported by the user groups, the system worked. Ahhh, life was so much simpler back then…but not necessarily better.

To get from "back then" to "right now," you must do the broadband flip. Stop thinking about bandwidth for capacity and begin thinking about bandwidth for performance. Consider a traditional, channelized, narrowband T1: it has 24 individual channels each capable of carrying exactly 64,000 bits every second. In fact, the type of sharing, or multiplexing, as the engineers prefer to call it, that the narrowband T1 uses—Time Division Multiplexing (TDM)—gives each of up to 24 simultaneous users the impression that they each have a connection that is dedicated to their own use and, in fact, they really do.

Figure 6.2 represents the narrowband T1. Let's take a closer look and ponder some possible scenarios and how they might play out in the channelized narrowband world. In the hypothetical T1 that Figure 6.2 shows, the darkest of the 24 channels, such as the topmost channel, are both configured/reserved and are in use at the current moment. The lighter colored channels, such as the bottom-most channel, are configured/reserved but are not currently being used. The white channels are neither configured, reserved nor in use.

What if the top channel needed more than 64 kbps of capacity? Could it share the channel below if that was already configured/reserved? No, it could not. Could it use any of the idle channels? No, not dynamically. To do so would require a rather lengthy reconfiguration process. Is it possible to use more or less than 64 kbps for a single connection? Yes, but not dynamically: a lengthy reconfiguration and provisioning process is required. The point is, narrowband connections are "nailed up" and not capable of changing dynamically.

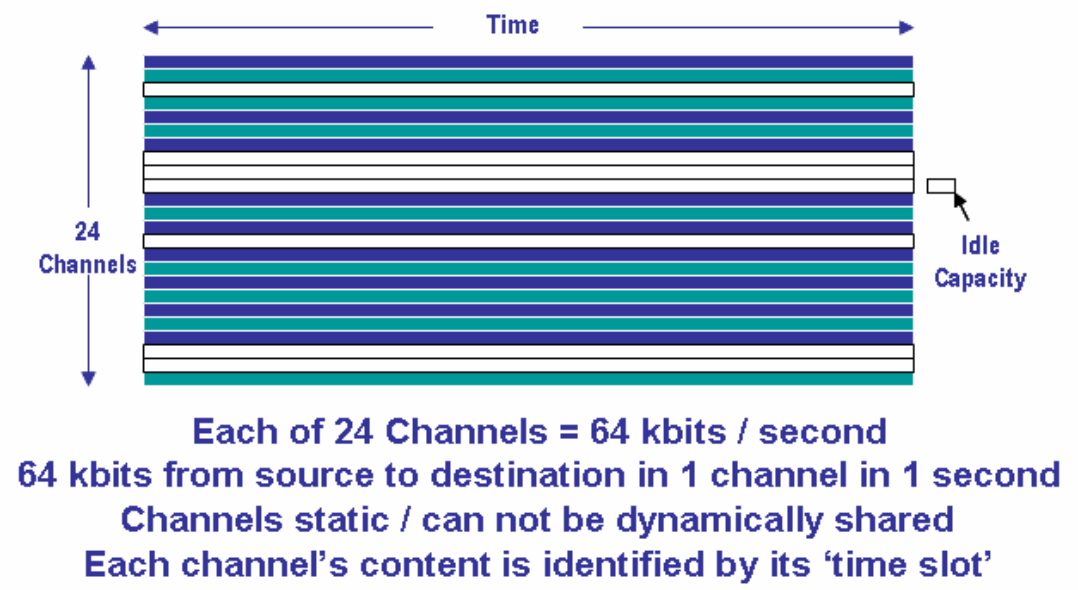

Figure 6.2: Channelized Narrowband T1.

And, what about capacity? How is capacity adjusted? Capacity is increased in groups of the basic 64 kbps channels, called DS0s, up to large capacities, as shown in Table 6.1. When the demand exceeds 24 individual channels in a single group, called a T1, the T1 can be replaced with a T2, which can accommodate up to 96 simultaneous 64 kbps connections. It is rare, however, to encounter a T2, as it is more customary to make the jump in most cases directly to a T3, which is capable of supporting as many as 672 simultaneous connections, or a fractional T3 comprising some number of T1s or T2s. The next step before the Synchronous Optical Network (SONET) is a metallic T4, which can accommodate up to 2016 simultaneous connections. The system was built to accommodate progressively larger numbers of the same size, 64 kbps, connections, not to allow for greater bandwidth, as broadband does.

SONET, as defined by GR 253-CORE from Telcordia, is a method of communicating digital information using lasers or LEDs over optical fiber

Level	x Prior Level	DS0s	BW
DS0	-------	1	64K
T1	24	24	1.544M
T2	4	96	6.312M
T3	7	672	44.736M
T4	3	2,016	139.264M

Table 6.1: North American (Narrowband) digital hierarchy.

Figure 6.3 shows the narrowband T1 from Figure 6.2 reconfigured for broadband use. Instead of being carved into 24 identical small channels the non-channelized broadband T1 is a single large channel that can be shared by multiple connections, possibly more than 24, but the key point is that the full single channel is dedicated to a single connection at a moment in time. How is this accomplished? In the narrowband T1, we could associate specific bits with specific channels by their *positions*. The first group of eight bits read from the wire belongs to the first connection; the second group of eight bits belongs to the second connection, and so on up to 24. Any idle or non configured channels have filler bits to maintain the positioning and this process is repeated 8000 times per second to deliver 64 thousand (8 x 8000) bits per second per channel.

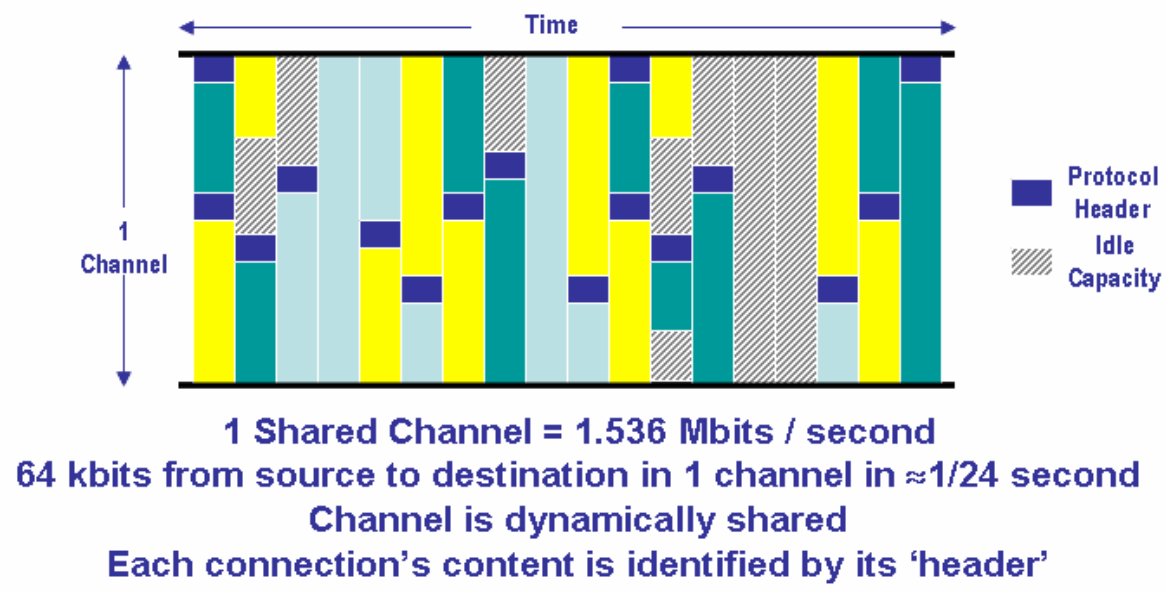

Figure 6.3: Non-channelized Broadband T1.

In the case of the broadband T1 shown in Figure 6.3, bits are not associated with their connections by their position but rather by additional protocol bits that are used to create frames, packets, or cells, which are three different types of electronic containers for transporting broadband bits. Broadband is actually not as efficient as narrowband for transporting multiple identical channels of information, but that is rarely the objective anymore. Broadband is good at dynamically allocating bandwidth resources to the rapidly varying demands of multimedia such as voice, data, and video.

In Figure 6.3, it is also possible to observe that idle capacity is not trapped in individual connections that cannot be dynamically shared, as it is in narrowband. In broadband, all capacity can be used dynamically and idle bandwidth only occurs when all information has been transmitted for all connections.

And what about capacity? Broadband capacity is enhanced by increasing the size of the "pipe," not by increasing the number of same-size pipes in some sort of bundle as is the case with narrowband. It is possible to configure narrowband connections—such as T1, T2, T3, and T4—for broadband access, but it is also possible to use Digital Subscriber Line (DSL), cable modem connections, metro and wide area Ethernet, and other broadband services to allow you to get as close as possible to the right size "pipe" for both capacity and performance needs. This is where you must make the big leap from considering bandwidth for capacity—though capacity is still a consideration—to considering bandwidth for real performance and doing the broadband flip.

As Figure 6.4 shows, increasing the capacity of a narrowband circuit increases the aggregate capacity but does nothing for transmission performance. Regardless of the number of 64 kbps connections bundled together, the individual channel capacity does not change. Figure 6.5 shows that increasing the size of the actual "pipe" and utilizing it in a broadband fashion have an important impact on performance per connection.

Transmission Time for 64K bits
(Narrowband Example)

1 DS0 Channel	= 1 second
1 DS0 Channel in a T1	= 1 second
1 DS0 Channel in a T2	= 1 second
1 DS0 Channel in a T3	= 1 second
1 DS0 Channel in an OC-3	= 1 second
1 DS0 Channel in an OC-12	= 1 second
1 DS0 Channel in an OC-48	= 1 second

Figure 6.4: Narrowband transmission times.

Transmission Time for 64K bits
(Broadband Example)

Broadband T1	= 1/24 second
Broadband T2	= 1/24/4 second
Broadband T3	= 1/24/4/7 second
OC-3c	= 1/24/4/7/3 second
OC-12c	= 1/24/4/7/3/4 second
OC-48c	= 1/24/4/7/3/4/4 second

Figure 6.5: Broadband transmission times.

Table 6.2 and Table 6.3 show the wide range of possible bandwidth needs, per connection, for a variety of combinations of Voice over Packet and transmission protocols. The difference between Table 6.2 and Table 6.3 is the packetization interval. Table 6.2 shows the impact of creating and sending packets every 10 ms, resulting in more packets but a smoother flow of voice that is less susceptible to packet loss; Table 6.3 uses a packetization interval of 30 ms. With a 30 ms packetization interval, bandwidth is consumed less quickly because more voice samples are placed in each packet, thereby reducing the number of overhead bits per sample; however, the connection is impacted more heavily by packet loss as the loss of a single packet will result in the loss of three times as many voice samples.

				10 ms Packetization, VAD off			10 ms Packetization, VAD on		
				G.711	G.729	G.723	G.711	G.729	G.723
Voice Type	over Link Type	Full/Half Duplex	Link Speed (kbps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)
Digital	TDM	Full	n/a	64,000	8,000	6,300	64,000	8,000	6,300
VoIP	802.3	Half	10,000	252,800	140,800	n/a	176,960	98,560	n/a
VoIP	802.3	Half	100,000	252,800	140,800	n/a	176,960	98,560	n/a
VoIP	802.3	Full	100,000	126,400	70,400	n/a	88,480	49,280	n/a
VoIP	Frame Relay	Full	any	100,800	44,800	n/a	70,560	31,360	n/a
VoIP	PPP	Full	any	102,400	46,400	n/a	71,680	32,480	n/a
VoIP	ATM (AAL-5)	Full	any	127,200	84,800	n/a	89,040	59,360	n/a
VoIP/hc	Frame Relay	Full	any	72,800	16,800	n/a	50,960	11,760	n/a
VoIP/hc	PPP	Full	any	74,400	18,400	n/a	52,080	12,880	n/a
VoFR	Frame Relay	Full	any	71,200	15,200	n/a	49,840	10,640	n/a
VoATM	ATM (AAL-5)	Full	any	84,800	42,400	n/a	59,360	29,680	n/a

Table 6.2: VoIP bandwidth estimate with 10 ms packetization interval.

Source: Matthew Michels, Nortel Networks, used with permission.

				30 ms Packetization, VAD off			30 ms Packetization, VAD on		
				G.711	G.729	G.723	G.711	G.729	G.723
Voice Type	over Link Type	Full/Half Duplex	Link Speed (kbps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)	Bandwidth demand (bps)
Digital	TDM	Full	n/a	64,000	8,000	6,300	64,000	8,000	6,300
VoIP	802.3	Half	10,000	169,600	57,600	54,400	118,720	40,320	38,080
VoIP	802.3	Half	100,000	169,600	57,600	54,400	118,720	40,320	38,080
VoIP	802.3	Full	100,000	84,800	28,800	27,200	59,360	20,160	19,040
VoIP	Frame Relay	Full	any	76,267	20,267	18,667	53,387	14,187	13,067
VoIP	PPP	Full	any	76,800	20,800	19,200	53,760	14,560	13,440
VoIP	ATM (AAL-5)	Full	any	84,800	28,267	28,267	59,360	19,787	19,787
VoIP/hc	Frame Relay	Full	any	66,933	10,933	9,333	46,853	7,653	6,533
VoIP/hc	PPP	Full	any	67,467	11,467	9,867	47,227	8,027	6,907
VoFR	Frame Relay	Full	any	66,400	10,400	8,800	46,480	7,280	6,160
VoATM	ATM (AAL-5)	Full	any	84,800	14,133	14,133	59,360	9,893	9,893

Table 6.3: VoIP bandwidth estimate with 30 ms packetization interval.

Source: Matthew Michels, Nortel Networks, used with permission.

These types of bandwidth utilization models can also be used to determine the impact of voice activity detection (VAD), which eliminate packets that are carrying silence. By eliminating silence packets, it is possible to get additional benefits with a negligible impact on quality assuming the VAD system is properly configured and tuned. Proper configuration involves the treatment of the absence of silence packets at the destination and how silence packets are chosen for non-sending at the source. If in the absence of voice packets, no sound is played into the listeners' ear, as in the earliest implementations, VAD can be very disruptive to the call and perceived call quality, causing the talker, upon hearing nothing from the listener, to pause and ask "are you still there?" Insertion of comfort noise, also known as Comfort Noise Generation (CNG), is vital. At the source, if packets are not sent that are *predominantly* silence, chopping or clipping of syllables can occur. Not setting stringent enough criteria for packet non-sending will lessen the value of VAD, which can be as much as 4:1, meaning that an actual call using VAD will require 25% of the bandwidth required by the same call not using VAD.

Results of this kind can be used to play "what-if games" and to develop realistic bandwidth estimates for different types of connections and, thereby, be able to establish baselines for the operation of your specific network. It is likely that you will end up, with at the very least, two profiles: inside calls and outside calls. But you may very well have different bandwidth requirements for a much more specific set of user requirements, much along the lines of the user profiles developed in Chapter 4. It might be possible to establish as few as four user profiles that are common to all departments, though it is more likely that there will be far more profiles for any given situation. Let's look at four basic profiles, just to give an idea of the type of feature sets they might need.

Profile	Possible Job Functions	Telephony Features
Basic Telephony User	Office worker, inside sales person, general staff, supervisor, low level manager, help desk worker	▪Dial Tone ▪In-house (extension) dialing ▪Local dialing ▪Long Distance dialing ▪Basic Voice Mail ▪Call Waiting ▪Caller ID
Intermediate Telephony User	Secretary, senior secretary, telemarketer, senior help desk worker	***Basic Telephony User PLUS:*** ▪3-way Conference calls ▪Skills-Based Call Routing ▪Call Forwarding
Advanced Telephony User	Administrative assistant, group admin, senior telemarketer	***Intermediate User PLUS:*** ▪Enhanced Voice Mail (add FAX storage/retrieval and advanced messaging and retrieval features) ▪Enhanced Call Forwarding ▪Multi-party conference calls
Power Telephony User	Senior sales person, senior manager	***Advanced User PLUS:*** ▪Multimedia conferences (web plus audio) ▪Enhanced Call Forwarding with Remote Access
Multimedia User	Senior manager, technician, developer, product marketing manager, senior sales manager	▪Telephony ▪Call Routing Management and Follow-me Features ▪Voice Messaging ▪Instant Messaging ▪Multimedia Conferencing ▪Collaboration
Road Warrior	Sales executive, account manager, senior sales manager, sales VP, sales director	▪Remote Telephony ▪Remote Data ▪Remote Video ▪Advanced Call Routing and Follow-Me Features

Table 6.4: Sample user profiles and telephony features.

It is clear that many organizations will have additional requirements, such as various skill levels of Call Takers, Call Taker Supervisors, operators, executives, administrative assistants, and so on, but this basic mapping of user profiles to features will show a way of simplifying both the testing and overall implementation issues.

So what should the operations personnel *do* with all this information? Using these considerations, baselines should be established, growth should be observed, and system capacity tracked. Additionally alternative routes should be established in case of failure, possibly with multiple carriers and other back-up plans put into place, as described in the following section on reliability.

There are many modeling, simulation, and tracking tools available to help with this task, as even the smallest networks are too complex to be tracked manually. Tools to help with this function are described at the end of this chapter.

Reliability

Establishment of "reliability" baselines and system monitoring to ensure that all components are operating within those parameters is very important to the proper ongoing operation of the network. Reliability must be carefully considered from the standpoint of the impact on the end user. Reliability issues may be considered, and prioritized, in terms of "service impacting" and "non-service impacting" outages. To implement a reasonable approach to reliability, one must consider the percentage of time a system is needed, that the system is actually required; versus the total amount of time the system is available.

Beyond the service impacting and non-service impacting categories, reliability must be monitored and addressed in terms of "root cause analysis," and this must be a key part of training and standard operating procedures. Root cause analysis, in the terminology of the medical profession, allows you to treat the illness, not the symptom. If we stay for a moment with the medical analogy, we would describe the visit of a patient to the doctor. The patient complains about a temperature. Does the doctor tell the patient to lie down in a bath of ice water to lower their temperature or does the doctor ask further questions to ascertain why the patient has a temperature? Remarkably, many call centers and help desks are designed to get the "patient" to fill a bath with cold water. These help desks are usually characterized by setting goals for the call takers such as "closed tickets per hour" or a similar metric. The call centers and help desks that will go to the next step are usually setting goals for their call takers based upon user satisfaction and solving user problems. They will ask the next diagnostic questions and consider the systemic issues of any problem—and that is what must be done.

Reliability is related to the "end-to-end view"—the closer a problem is to the end user, the fewer users that problem will impact. The other side of that example is that the closer a problem is to the core of the network, the more users it will impact. For example, a problem with an individual user's SIP phone will only impact that user's ability to make and receive calls. Their voicemail will still be working, but they will not be able to engage in interactive conversations. Although this does represent a problem, the scope of the problem and the set of users a problem impacts are limited. However, if a problem affects the VoIP call server, it will impact all users of that server. These considerations should drive the investment of time, money, and effort to prevent problems and your investment in mitigation when problems do occur.

Security

An important part of ongoing operations is the security of the entire system generally and specifically; in this case, the Voice over Packet services. There are three distinct areas of concern relating to security of voice, specific security threats and vulnerabilities within each area. In hybrid systems, combining traditional systems and Voice over Packet, the number of distinct areas of concern increases substantially. As Figure 6.6 shows, the three areas are the phone itself, the network, and the server. Within the network, there are also three areas of vulnerability: the phone access to the network, the network backbone, and the server access to the network.

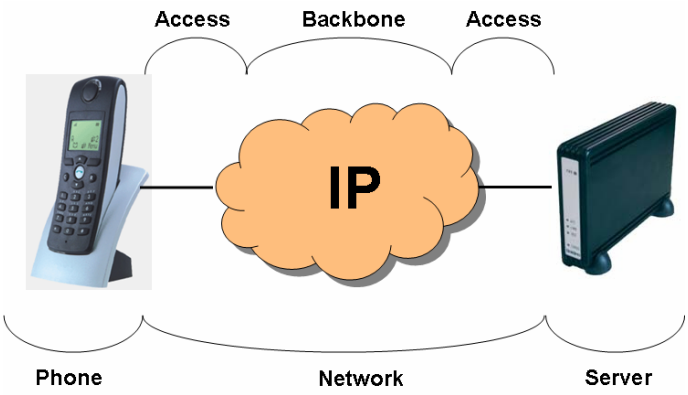

Figure 6.6: Areas of vulnerability for IP-based voice systems.

Let's start a discussion of security from the phone end of the connection. What many people forget is that Voice over Packet-capable phones, regardless of connection type, are actually computers running specific applications. Very often, the phones are programmable in some generally available programming or mark-up language, such as the eXtended Markup Language (XML), and can be reprogrammed as easily as they can be programmed. The types of vulnerabilities, therefore, are very similar to vulnerabilities in the server with the exception that security issues with individual phones may impact one or a small subset of users while server breaches impact all users of that server. Phone and servers are susceptible to three types of attacks: malicious software, or malware-enabled attacks, as they have come to be known; content compromise; and Denial of Service (DoS). An additional consideration, even though most organizations consider VoIP and its related services to be "free," (organizational telephony services are about as free today as 800 service was free on traditional telephone networks)—is that it's not. Someone has to pay for the bandwidth and, maybe more importantly, the time of the person using the service. Many organizations are shocked when they do an audit of the protocols running on their network and find a large percentage of their bandwidth is used by personal Skype, unauthorized NetMeeting, and other non-approved system uses. These are often the silent destroyers of response times and performance that chokes off organizational productivity.

Malicious software, or malware, is software that is surreptitiously downloaded to a phone or server that does things of which the user is unaware and would probably not approve if they were aware. Malware can be as simple as compromising dialed phone numbers or calling directories by sending them to an unauthorized third party or disrupting normal operation of the phone, such as not ringing or selectively not ringing, or allowing unauthorized use of the phone. Content compromise can allow access to the voice communication by unauthorized third parties and can possibly overcome encryption methods by allowing the unauthorized third parties access to encryption keys. DoS attacks are the simplest type of attack and simply block the user from establishing all calls or specific calls or from using certain system features. DoS attacks are the simplest but the most common.

Within the "network" portion of the call, the vulnerabilities are typically of the eavesdropping or DoS varieties. There is greater vulnerability, typically, in the user/phone access than there is in the server access and backbone. The backbone is usually the least vulnerable to ad hoc attacks, especially if the backbone is a Virtual Private Network (VPN) provided by a trusted third party.

Another aspect of security that must be managed operationally is systems that are put into place to enhance security—such as firewalls, stateful inspection engines, and Network Address Translation (NAT)—but that often collide with the Voice over Packet systems and prevent their proper operation. With packet voice services playing an increasingly important role in organizational communications, these are considerations that must be made prior to implementation of the security systems or upgrades and enhancements to the capabilities of the existing systems.

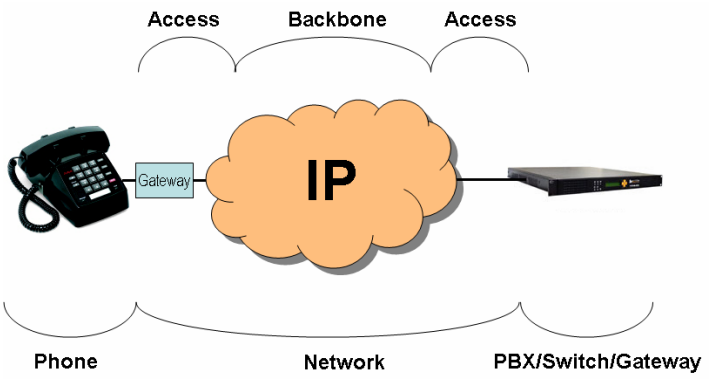

Figure 6.7: Voice over Packet system implemented via a gateway.

A second consideration is packet voice systems that are implemented over traditional telephony systems. They share the basic issues with the systems implemented via IP-enabled phones except that the target for malicious software will be the gateway device that connects the traditional telephone devices to the IP network. This arrangement has not proven to be more or less vulnerable to hacking than the situation in which the gateway functionality is embedded in the telephone device itself.

There is yet a third situation, which Figure 6.8 shows that must be considered from a security standpoint: the situation whereby a traditional telco network, either public or private, and an IP voice network are interconnected. This is a very common situation that, the call at some point in its duration will be subjected to the security issues discussed earlier relative to IP-based voice and to all of the security issues associated with traditional telephony. Although traditional telephony security issues do overlap with some IP-related issues, the traditional telephony part of the call has some unique security vulnerabilities as well. The issues that overlap have to do with unauthorized use of services, and the ones that are unique are related to these issues. The cost per packet voice call is usually very low compared with the cost of traditional calls, so, therefore, as mentioned earlier, the largest component of the true cost is usually productivity and time costs, as opposed to per minute charges, referred to in old telephone network parlance as “toll charges”—the same term that gives us the word “toll fraud” for the theft of services with those charges associated.

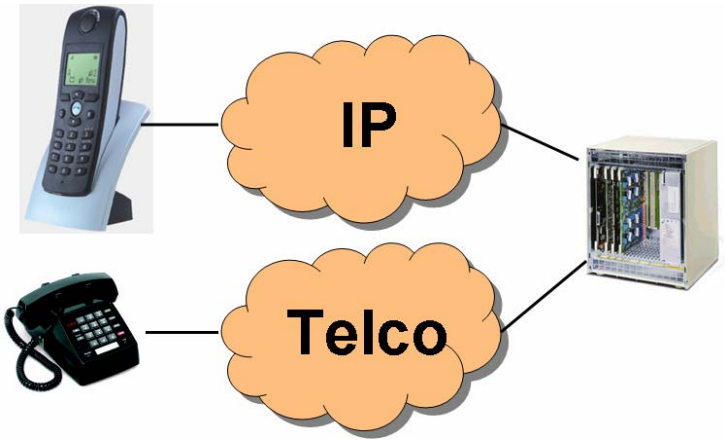

Figure 6.8: Voice over Packet and traditional telephony hybrid.

The toll fraud issue in a hybrid network is actually a double-edged sword: the IP network can provide new, often unsecured access to traditional telephone networks that have toll charges associated, and the traditional telephony networks can provide unchecked access to the IP networks. In addition to newly minted exploits, the old ones still work and are still an issue: Mailboxes that are set up on corporate systems and resold as part of a service package or used by terrorists or for untraceable or difficult-to-trace drug transactions. Sale of long distance, often international, calls via corporate PBXs; and other similar loopholes and backdoors cost organizations billions of dollars a year. This will continue as long as the IP and traditional networks are interconnected, which in many cases will be for some number of decades.

What can an organization do? The best bet is to work with your security team, whether they be an inside team or a part of your service provider, to monitor activity and to thoroughly check out any activity that appears to be abnormal—or an unusually high volume of normal activity. Eccentricities and anomalies that should be monitored are unusual protocols, unusually high or low call volumes, appearance or disappearance of encryption, time of day shifts in usage, unusually high or low volumes for a certain phone number or user, and similar patterns that can mean abuse. Restoring software to IP phones and servers, PBXs, and gateways periodically is also a good idea as is using secure versions of the protocols and systems used to install, maintain, and operate them. Using Secure Shell (SSH), for instance, to maintain gateways, as opposed to the common Telnet protocol that is available to every hacker, is a good idea as is using secure versions of SNMP, OSPF routing protocol, and other fundamental workhorse protocols used to run the networks that carry the increasingly important Voice over Packet traffic. These protocols have many benefits when compared with their insecure variants in that they employ encryption, tunneling, and other security capabilities to ensure greater security of the applications.

Call Quality and Voice Quality

We have previously, and vigorously, defined call quality and voice quality and described the difference between the two. Nowhere is this distinction more important than in the operations phase of the Voice over Packet network's life cycle. Call quality is directly related to QoS and its underlying methods and measurements. Call quality and QoS measure and attempt to optimize, in some cases dynamically, packet loss, delay, delay variation, and system availability. Voice quality, however, is directly related to Quality of [User] Experience (QoE) and is very tightly coupled with the opinion of the humans using the system. QoE is only loosely coupled with QoS even though it is an automated approximation of human QoE opinions derived from the statistics of QoS upon which Service Level Agreements (SLAs) and network-based alarms are based. The reason for this is that QoE scores are based upon the opinion or averaged opinions of one or more humans. Those humans are subject to a variety of factors that are impossible to model and reproduce mathematically, but those humans are also not available to judge voice quality on a given call, so it is important to use automated systems. The benefit of automated systems is that, although they do not exactly reproduce the human scores, they are consistent and always available. The drawback is that they only roughly approximate the humans' opinion of the voice quality.

The two types of testing are non-intrusive and intrusive. Non-intrusive testing does not require a special test call and is based on actual conversations or is calculated from metrics gathered during an actual call. Intrusive requires a special test call and is based on comparison of source signal with the signal after transmission. Both are valuable operationally but the non-intrusive testing will raise a realistic red flag sooner as it relates directly to the quality of the voice as experienced by the user because it is a set of measurements made on real calls. Intrusive testing not only is artificial in that it is not based on actual calls but also can place an additional burden on a network for the test calls themselves that may distort the test results and even have a negative impact on real-live calls in progress. This is especially true if intrusive testing is performed at the most desirable time to do testing: when the live call load on the network is particularly high and nearing upper limits.

Non-intrusive tests include the Mean Opinion Score (MOS), derived MOS, the E-Model (ITU-T G.107) R value, and, more recently, ITU-T Calculated Planning Impairment Factor (ICPIF) Scores (ITU-T G.113). MOS testing has been adopted from traditional telephony, and historically has been, at least for North American telco operations, the actual opinion of voice quality from a panel of human judges from central Illinois. Even though one might question whether it is statistically significant, the traditional panel has been composed of 16 people: 8 men and 8 women. The panel of judges is sequestered in a sound-controlled laboratory environment and judges the quality of voice on a variety of systems on a scale from 1 (lowest) to 5 (highest). MOS is useful for human user QoE analysis but is virtually useless for SLAs and network monitoring due to the difficulty in reproducing the results reliably and under a wide variety of circumstances.

The E-Model (ITU-T G.107) is a much newer innovation and is intended as a design tool that predicts average voice call quality using a mathematical model that takes into account the estimated impact of delay, delay variation (also known as jitter), packet loss, and performance of the codec (the voice coder that translates human speech into bits at the source and decodes the bits back into analog waves at the destination). The result of the E-Model calculation is the R factor or R value. The R value is an estimated voice quality rating with a range from 0 to 100. 0 indicates lowest call quality and 100 indicates a perfect quality score. From the E-Model's output, the R value and derived MOS can be calculated, as Figure 6.9 shows.

Figure 6.9: Correlation of E-Model R values to MOS scores for derived MOS.

In actual operations environments, the author recommends a target of 100, which matches to a MOS of 5.0—a perfect and unachievable score. But, like archers, Voice over Packet operations' personnel need to aim above the mark to actually hit the mark and, in this case, aiming for a perfect score will probably put the voice quality in your network consistently in the R value range of 94 or so. This matches a derived MOS of 4.4, which is in the same 4.3 to 4.5 range generally achieved by most traditional voice networks based upon Pulse Code Modulation (PCM) and TDM transport.

The real value of the E-Model lies beyond the design work for which it was originally created and is realized more in the operations area, which is the subject of this chapter. Operationally, E-Model scores are based upon measurable parameters found in the network and which can, if needed, be reproduced in the lab. In actual operation, E-Model evaluates Real-Time Protocol (RTP) Streams based upon source and destination addresses, port numbers, and sequence numbers and creates a jitter profile and predicts the impact of the Internet Protocol (IP) bearer on MOS with an 80 to 90 percent correlation to actual human scores.

An additional non-intrusive evaluation technique is the ITU-T ICPIF score (ITU-T G.113). Some manufacturers are beginning to adopt ICPIF for voice quality calculations due to the additional levels of sophistication being required in assessment of voice quality. In fact, from the perspective of VoIP service providers, this author has always maintained that when all prices are equal—and as low as they can realistically go without bankrupting the service provider—other differentiation factors will emerge that will be the basis for purchase decisions and that the human interpretation of voice quality will be one of those factors. We are already beginning to see this in services offered to the public. From an organizational point of view, this has never ceased to be true for two important reasons. The first reason is that in the organizational/enterprise/corporate/agency environment, cost savings are not seen and felt directly by users in most cases: they only see an improvement, degradation, or no change in the quality of voice calls and the related telephony and multimedia services available to them. The second factor is that although most people's opinion of what call quality is possible, and reasonable, has been strongly, and downwardly, influenced by the use of cell phones, many in the business environment still have a traditional, reliable, high-quality telephone with which to compare new service. Therefore, a lot of work is being done in the area of voice quality, much of it in anticipation of the shift of importance to voice quality and the ICPIF score, available since February of 1996, is a good example of this.

ICPIF takes into account the factors that the E-Model does and additionally gives weight in the evaluation to signal attenuation distortion and loss, circuit noise, codec distortion, group delay distortion, one-way transmission time, talker echo, and some additional parameters of special importance. Figure 6.10 shows the rough correlation of ICPIF scores to MOS. It is, therefore, possible to derive MOS values from ICPIF scores just as it is possible to do so using R values.

ICPIF Range	MOS	Quality Category
0 - 3	5	Best
4 - 13	4	High
14 - 23	3	Medium
24 - 33	2	Low
34 - 43	1	Poor

Figure 6.10: Correlation of ICPIF to MOS.

Intrusive evaluations, while not as valuable operationally as non-intrusive tests, still have operational value. Intrusive tests are often used to perform automated testing during busy hours or non-busy hours, and several companies have systems that are automated and operate free of human intervention. The family of intrusive tests includes Perceptual Analysis Measurement System (PAMS), Perceptual Speech Quality Measure (PSQM), Perceptual Evaluation of Speech Quality (PESQ), and PESQ+. All of these evaluation approaches are derived from PAMS, which was developed by British Telecom as a replacement for human MOS values. Although PAMS and its close relatives that follow do not replace MOS, only approximating it within +/- 10 to 20 percent (a window of up to 40%), they are still excellent for benchmarks and comparisons. This group of tests work by comparing the original analog voice wave with reproduced/transmitted speech using a complex weighting method intended to take into account characteristics important to the human ear. The scale is from 0 to 6.5 with 0 being "perfect" (that is, no difference between waves). PSQM, ITU-T Recommendation P.861, and PESQ, ITU-T Recommendation P.862 and PESQ+ operate very similarly to PAMs but come close to approximating MOS values.

Measurement and Monitoring

So what is the value of this in operation? The SLAs that the operations group inherited from the design and procurement process should be based, to some degree or other, on MOS. Because it is not practical to line up humans to listen to all calls, an automated method is needed and this, or some other variants—some proprietary and some standard—can be established as the network-wide standard. Non-intrusive tests can be used as the basis for evaluating human user complaints and intrusive testing can be used to anticipate issues and take a proactive approach to fixing them before they reach a level sufficiently bad to generate user complaints and the resulting trouble tickets. Baselines need to be established early in the operational life cycle, consistent with the SLA, and any variation above or below the baseline by more than your tolerance—+/- 5% gives a 10% window that is reasonable in most circumstances—is reported, a root cause analysis is done, and whatever steps are needed are performed to get the observed scores back into line with the baseline values for the network.

NetFlow/IPFix

An important consideration for measurement and monitoring is the exact methodology that will be user to implement it. In many cases, organizations opt for bulk statistics or exception reporting through SNMP or Remote MONitoring (RMON), but a methodology that is being standardized and receiving a lot of attention recently is NetFlow, also known, in the standards area, as IPFix. NetFlow is an open but proprietary network protocol developed by Cisco Systems to run on Cisco IOS-enabled equipment for collecting IP traffic information. Cisco routers that have the NetFlow feature enabled generate NetFlow records; these are exported from the router in User Datagram Protocol (UDP) or Stream Control Transmission Protocol (SCTP) packets and collected using a NetFlow collector. Juniper Networks provides a similar feature for its routers called cflowd, which is basically NetFlow 5. Huawei Technology routers also support the same technology, but call it NetStream.

Historically, network flows have been defined in many ways. In the case of NetFlow, Cisco uses the common 5-tuple definition, where a flow is defined as a unidirectional sequence of packets sharing all the following five values:

- Source IP address
- Destination IP address
- Source TCP port
- Destination TCP port
- IP

The router will output a flow record when it determines that the flow is finished. It does so by flow aging: when the router sees new traffic for an existing flow, it resets the aging counter. Also, TCP session termination in a TCP flow causes the router to expire the flow. Routers can also be configured to output a flow record at a fixed interval even if the flow is still ongoing. In Flexible NetFlow (FNF), an administrator could actually define flow properties on the router.

Beyond the Cisco realm, the Internet Engineering Task Force (IETF) is standardizing NetFlow under the title IPFix—IP Flow Information Export—and it is being made available on a wide variety of network products. The gathering and analysis of NetFlow/IPFix data is often best performed on separate management platforms using third-party software of which a wide range is available.

Growth Management

Growth of a Voice over Packet network is a part of the bigger project of managing growth of a multimedia IP network. As such, it has much in common with the bigger project but also has issues of its own. Growth must be managed in terms of fixed assets and equipment budgets, access bandwidth, usage of the shared bandwidth, and shared resources such as switches, routers, telephony servers and gateways, and numbers and licenses.

Fixed Assets and Equipment Budgets

The first job will be keeping track of physical assets, which have both hardware and software elements. The Voice over Packet network will include hardware and possibly software telephony clients for the end users. In many cases, this will include the reallocation of traditional telephony devices such as desk phones to the Voice over Packet inventory. Fixed assets will also include telephony-only systems, such as SIP servers and gateways, and might include an allocation for the part of shared systems, such as routers and switches that are used by voice traffic. Although this is often a secondary consideration on the list of many seemingly more critical items for operations staff, this can be a time-consuming and possibly career-threatening item if ignored. Proper tracking and forecasting of growth includes awareness of pending management decisions regarding mergers and acquisitions, layoffs, new product lines, or anything that can substantially impact the number, and profile, of the user community. The recordkeeping associated with this task will be addressed further in subsequent sections.

Network Bandwidth: Access and Backbone

The second job will be managing the bandwidth pool. Bandwidth utilization issues have been addressed in a prior section and at this point in the life cycle of the Voice over Packet network, calculated values should have been compared with actual values and you should have a clear picture of the bandwidth that is used by each type, or profile, of user. You should also be aware of the number of each type of users and the likelihood of that sub-population to substantially grow, shrink, or be reallocated to a different profile.

Recall that an element of the user profile was the degree of mobility, which can impact user access as well as the amount of shared bandwidth they use in the network backbone. All of these factors need to be managed; human resources and management hiring and acquisition of other organizations needs to be factored in and a budget developed. One of the biggest factors that can influence a finely tuned and well-performing network is a shift in users or among and between user profiles.

How can this be managed operationally? Most systems, especially VPN systems and others that require user sign-ons, allow resources to be tracked to some level by individual user and, at the very least, by user group. The key is to anticipate the need and to establish appropriate profiles in advance and to use the profiles to establish user groupings. Another benefit of such profiles is that, very often, user assignment to specific profiles can be done using a standard template that substantially reduces the chances of errors and makes adding new users much faster and easier.

Shared Resources: Switches, Routers, Call Servers, IP PBXs, and Gateways

Telephony users are typically multimedia users, not telephony-only, though in either case, it is possible to determine the amount of shared resources that a user consumes. This is a very important growth planning and management element and should be tracked, reported, and reviewed regularly. Specific alarm thresholds being set to send up some sort of electronic signal flare should utilization approach or exceed thresholds between reviews. Shared resources to consider include the switches and routers that are responsible for transmitting and receiving all traffic and the call servers, IP-PBXs, and gateways responsible for handling voice traffic.

Numbers and Licenses

Resources that are often not considered, but must be, are numbers and licenses. The two main numbers that are required are phone numbers and IP addresses. The licenses that need to be managed are any software licenses that are charged per user, per seat, per port, or on any other incremental basis.

Phone numbers, most often, must be purchased from an outside telephony organization so that they may be registered in the proper places and so that appropriate routing may be established in the Public Switched Telephone Network (PSTN). Organizations are strongly advised not to use telephone numbers of their own creation, even inside what seems to be a private network, due to the fact that almost invariably there will be a need to connect to the PSTN and number duplication is inevitable. It is worth the extra cost. This is a bit less true with IP addresses because NAT can be used to map internal, potentially conflicting, IP addresses to external, registered IP addresses, but this point should be very carefully considered. It is possible to create and manage an inventory of available phone numbers and IP addresses, to assign them as needed, and return them to the inventory for reassignment when they are not in use.

Licenses, likewise, should be inventoried and returned to the inventory for reassignment when not in use. When a user leaves or is terminated, the company almost always remembers to get that user's keys and delete their passwords. A part of this procedure should also be to return any licenses to the common inventory. Licenses cost money and are needed to operate most Voice over Packet systems.

Recordkeeping and Documentation

This is the boring but inevitable part of any operations environment: Keeping track of the assets. Although boring, and inevitable, this is still a very important function and one that must be kept up to date as information may be requested at any time.

Fixed Asset Accounting

Fixed assets include all PCs, laptops, palm devices, routers, switches, Power over Ethernet racks, and just about anything that one can touch. These are the assets of the organization that have been entrusted to the care of the network or IT department for the benefit of the company. These assets must be managed.

In many cases, the finance department will insist that you use their fixed asset system or, at the very least, provide your data in a manner that can be input to theirs electronically. The problem with using a generic system is that it often does not capture specific characteristics that are desirable to have to manage your network properly. Information such as release levels, software and firmware revisions, and other key characteristics are often lost. It is also true that many of the automated tracking and inventory systems for networked components are not directly compatible with generic organizational fixed asset systems and must be put through an intermediate step, such as a Comma Separated Value (CSV) or Tab Separated Value (TSV) files, Excel spreadsheet, and/or XML output.

If you have not inherited a fixed asset system nor have one already in place, you must procure one as part of the project. The considerations for a fixed asset accounting system closely mirror the considerations for any new software application. What are your basic system requirements? Will the system run on servers with existing applications or on a different platform? What hardware does your MIS or IT department recommend? What operating systems can be supported? How many users will access the system and what are the security requirements? Will shared access over the network be allowed? What additional user rights and permissions must be configured and enabled to support this?

Will the new fixed asset accounting system be a separate, standalone system or an add-on module to your current accounting system? If you require a standalone system, it is important to understand the system's interface capabilities. Will it interface with general ledger accounting modules and tax applications to avoid duplicate entry of data? Does the system allow data import and export? If so, what import and export formats does it support? If you require an add-on module to your current accounting system, you must make sure that the add-on module will interface with your current accounting system.

What do you want in terms of reporting capabilities? Standard reports should include depreciation, such as projected depreciation and depreciation method comparisons, to aid planning. What kind of tracking does it do for financial reporting? Also be sure to investigate the ad hoc reporting capabilities. Is there a limit to the number of user-definable reports you can create? Is it difficult to create ad hoc reports?

What are your data entry and depreciation needs? Many packages offer features that significantly reduce fixed asset entry and calculation complexities. Is the user interface easy to use? Does it offer standard templates to ease data entry? Does it offer automatic data validation? Does it offer automatic calculations on a periodic and daily basis? Does it allow you to customize the information you want to enter? How many depreciation methods are supported?

What are your fixed asset tracking and management requirements? Whether you are a large or small company, it is important to understand your asset tracking and management needs. Does it keep physical track of where assets are located? Does it track multiple locations? Does it manage multiple companies? Does it offer bar code capabilities? How does it track assets of mobile users such as cell phones, PDAs, notebook computers, and similar devices that may be on the road as much as their human users? All of these concerns, and more, must be addressed in close collaboration with the accounting department.

Shipping, Warehousing, and Tracking

Shipping, warehousing, and tracking of assets are critical. In many cases, devices and systems are ordered and shipped directly to the end user, very often being charged to the user's personal or company credit card and charged back to the company through an expense reporting and reimbursement procedure. Are these assets tracked? What is the dollar value threshold for your organization that must be tracked? How much money is effectively lost each year by not doing proper tracking? Or maybe assets such as SIP phones or VoIP over WiFi devices are centrally purchased and shipped to the user from a warehouse. How are these assets tagged and tracked?

Repair/Refurbishment/Return to Service/End-of-Life

Another area of consideration is repair, refurbishment, and return to service. How is this tracked? Are users provided with a loan or replacement system? How often are systems refurbished, either via software refresh and upgrade or cosmetically? What happens to assets at the end of their accounting life? Are they sold, given to schools or charities, or recycled?

Administration

In most organizations, operations consist of some level of operations management, often, in larger organizations, headed by a Vice President of Operations or similar function. Below that rectangle on the organizational chart, the line usually splits into at least two directions. One direction is the operations group of engineers, technicians, and often Help desk people who manage the technical operations and keep the hardware, software, and network running. The other line goes to a group of administrative people who do the paperwork, manage the budgets and inventories, and, basically, keep the group running that keeps the services running. There are many different functions of that second group, but we will consider some of the primary functions here that may be somewhat different or have some unique characteristics for Voice over Packet networks.

End-User Cost Allocation

I well remember the first time I heard the term "funny money." Besides being a two-word poem it caused me to ponder what was so funny about it. Upon further inquiry I learned that computer usage—at the time access to a mainframe computer over 9.6 kbps leased lines—was used to allocate costs to departments but that paper checks were never written nor did real money change hands. The transactions were purely bookkeeping transactions but, to the department heads, were real nonetheless. If they had a surplus of money in their budget, they might not get the full allocation next year; if they ran out of money, they might not be able to access the computer system, and so on—normal budgeting issues. Many organizations still track usage of telephony systems, both the variety of usage that results in monthly charges and usage that is part of a fixed cost system such as a private IP network or flat-rate VPN.

The most logical, and common, metric to track is MOU or Minutes of Use. This is the easiest to track in most cases but might not provide a clear picture of the assets, or parts of shared assets being consumed, such as bandwidth. Two other facets of the MOU tracking can be the quality, or Class of Service (CoS), of the call and the bandwidth being consumed as these other factors should have a "cost" associated with them. An additional reason to use MOU as the basic unit of measure, at least while the telephony network is still a circuit/packet hybrid using both the VoIP and PSTN networks, is that MOU will provide a common tracking and cost allocation metric that is used across both telephony platforms.

Detailed Usage and Billing

There is an inclination when making the move to IPT, and VoIP specifically, to abandon many prior financial controls and recordkeeping. Detailed usage and billing, however, is still required in many organizations for needs that go far beyond "funny money" accounting needs. In the case of professional services, detailed usage information supports client billing: not for the phone call but rather for the professional services delivered over that phone call. Detailed usage is also used to track sales progress or the efficiency of phone support, not to mention use as a management tool to review how an employee spends their time and to question non-business use—an item that is ever smaller in terms of network services costs but ever increasing in the cost of the human time wasted.

Vendor/MSP/Carrier Billing

An important area of any operation is ensuring the accuracy of bills, and nowhere is this more important than in the area of recurring charges—those charges that come up each and every month and must be paid over and over again. A problem with the bill rarely gets fixed and just becomes a part of a bigger problem over time. It is critical, for budgetary purposes, to review all vendor, managed service provider, and carrier bills each and every month against the orders for the services that appear on those invoices. The same holds true for periodic equipment invoices, though one-time invoices for purchases are more likely to be audited at the point of receipt of the product.

In many instances, organizations use the move to a Voice over Packet system as an opportunity to clean up their billing that has slowly become inaccurate over time. They consider the new packet voice system as a clean piece of paper and meticulously account for each and every new charge and track it. When employed by a new operations manager who is looking for a way to sort out the mess they have inherited from a predecessor, this approach has a reasonable chance of succeeding, assuming the new operations manager has the discipline to keep the new system updated. When used by an existing manager who created the mess in the first place, this is often a formula for assured failure: if this manager could not manage the old system what changes are being made to ensure that they can manage the new system?

Cost Reduction and Network Utilization Maximization

The job of looking for ways to reduce costs and maximize network utilization is the task of network operations. The job of putting those ways into production and modifying the underlying operations procedures is the job of network optimization and is covered in the next chapter.

Provisioning

Traditional telephony operates on a system of OAP&M—Operations, Administration, Provisioning, and Maintenance. In effect, these procedural steps are outlined in this operations chapter. The “P” part of the task—provisioning—includes setting up, configuring, installing, testing, and “turning up” any new service or feature a user may need. This includes the initial installation; any subsequent moves, adds, and changes; and the refresh of the technology and audit of the installed system.

Initial installation includes proper configuration, verification, and testing and should also include a security assessment to ensure that installation of the system is proper from a security standpoint and does not introduce any security vulnerabilities into the system. Initial installation should also include user training and execution of any asset tracking documents or other financial recordkeeping.

Maintenance

Once the system and its components have been placed into service, the maintenance cycle begins. Maintenance of hardware, software, and QoS and user QoE levels is a key function of ongoing operations. Basic maintenance functions include managing moves, adds, and changes and handling technology refresh and periodic audits.

Moves, Adds, and Changes

Moves, adds, and changes (MAC), including de-installation and de-commissioning of equipment, are the things that absorb a large part—in some cases, the majority—of the operations budget of any organization. VoIP systems, using Dynamic Host Configuration Protocol (DHCP) and other dynamic configuration options, were supposed to solve the MAC problem and allow users to easily relocate to any location where Internet access was available—inside or outside the company—and to simply connect. Although users of public VoIP services have, in fact, achieved this level of dynamic connection, many within the enterprise and agency networks have failed to do so. Inside the walls of the organization, many times, are internal defensive mechanisms—firewalls and similar systems—deployed in a defense-in-depth arrangement that often prevents easy, user-initiated MAC. An additional tool, but one that also adds complexity, is a VPN, such as those based on wide area Ethernet services and MPLS, and Virtual LANs (VLANs) to carry multimedia traffic. All of these areas must be considered in the light of your specific implementation and security needs.

Technology Refresh

Periodically, a technology refresh will be required. This may be as simple as ensuring that software and firmware are up to the most current release level or may require hardware upgrades or replacement as well. The rate of technology refresh should not be based on the rate of change of the technology or availability of new features but, rather, should be based on the need for the new technologies or features. For example, if it has been determined that a specific profile of VoIP user can get all the functionality that they need from plugging their old desk phones into Analog Terminal Adapters (ATAs) but new, low-cost SIP phones become available, the consideration to replace the ATAs should be relative to the basic operations needed by that profile and not based upon the additional features, however desirable they appear to be. A refresh is just that—refreshing capabilities that exist and have been previously approved not an upgrade or adding additional functionality.

Audits

Periodic audits must occur to verify both functions and assets in place. Audits should include functionality, security, hardware, software, services, and circuits. A comprehensive periodic audit, aided by automated auditing tools, should be performed independently of the requirements of the finance department. Audits of Voice over Packet systems should be sure to include all components and end systems, including those possessed by mobile users and users in their small office/home office environment.

Operations Tools

Decisions must be made about the tools needed, the amount of graphical reporting required and the testing methodology. Whether this is active and non-intrusive during the call, or intrusive and separate from live traffic as well as tracking relative to the SLA and/or other baselines, and the amount of integration that is desirable.

Integration

Generally speaking, the complete operations picture will require several tools—both on the technical operations and financial/budgeting sides. Some degree of integration should be expected between the different tools and may require some custom development work in order to get a full, comprehensive solution. This should be anticipated and worked into any plan that is developed.

Operation

Ongoing operation of network monitoring and management tools will require proper budgeting for training and planning to ensure that qualified personnel are available to use the systems and will most likely be integrated with ongoing network operations. The special training and skills needed to get the most out of the system will be the biggest concern. Besides considering employees, this is also a task that can be outsourced to a third party and managed by the organization.

Summary

Chances are that if you are responsible for the ongoing operation of the Voice over Packet network and its associated services, you have inherited something that is a bit chaotic and not quite fully formed. That is the nature of new things. If the job has been done correctly, many of the important operational considerations—baselines, SLAs, accounting, and asset tracking systems—will have been made and some initial steps taken to equip you to do your job. If not, the scope of your job is bigger.

The job of the operations group is less glamorous in many ways than the job of designing and procuring a new system, wrestling it into place, and making it work—but is no less important. It is the job of the operations group to ensure that the system performs to specifications consistently—to keep humming along and doing its job day and night, weekdays, weekends, and holidays without glitches or errors, through upgrades and surges in user demand.

The next chapter will consider the task of network optimization: improving the network operation in some way and providing new baselines and operational guidelines for the operations group to maintain in, hopefully, a never-ending loop of improvement—operations, improvement, operations…on and on and on.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.