

What You Will Learn

After reading this chapter, you should be able to

- ✓ Describe how VoIP can compete with the reliability of traditional PBX systems.
- ✓ Explain how to replace PBX-to-PBX connections with a VoIP network.
- ✓ Identify router interfaces used for connecting various analog devices (for example, phones and fax machines).
- ✓ Describe how voice-enabled routers connect to digital circuits (for example, T1 and E1 circuits).
- ✓ Explain how dial peers allow voice-enabled routers to forward calls to the appropriate destination.
- ✓ Discuss how VoIP can be used in the home, as opposed to traditional telephone service.



CHAPTER 3

Paving the Pathway to a Voice over IP Network

The transition from a traditional private branch exchange (PBX)-based telephony system to a *Voice over IP (VoIP)* system is not usually an overnight (or over-the-weekend for that matter) process. Instead, we usually take “baby steps,” as they said in *What About Bob?* A first step might be to replace the trunk line that interconnects PBXs at remote sites with an IP *wide area network (WAN)* connection. A next step could be to connect existing analog phones, fax machines, and speaker phones to voice-enabled routers.

The end result of these baby steps is a telephony network, without a PBX, where voice traffic is transmitted over an IP network. In this chapter, we’ll explore how to begin this migration and pave the pathway to a VoIP network.

Competing with the Reliability of Existing Phone Systems

The perception of many in business today is that VoIP simply isn’t reliable enough to support the telecommunication demands of a corporate environment. After all, corporate PBX systems are considered highly reliable, but how many times in a month do you hear users say, “My e-mail isn’t working,” “The Internet is down,” or “I can’t print to the network printer”? Because of such past frustrations with data network applications, this perception of unreliability has unfortunately carried forward to any new application running on the data network, such as VoIP.



Many PBX administrators like to boast, even though they are not always correct, that their PBX has the “five nines” of availability. By the “five nines” they mean that their PBX is available (that is, up and running) 99.999 percent of the time, and this availability isn’t just during regular business hours—it’s 24 hours a day, 365 days a year. If we were to do the math, we would see that if a network is up and available 99.999 percent of the time, then it would only be unavailable for five minutes a year. Consider Table 3-1, which illustrates the yearly downtime associated with various availability levels.

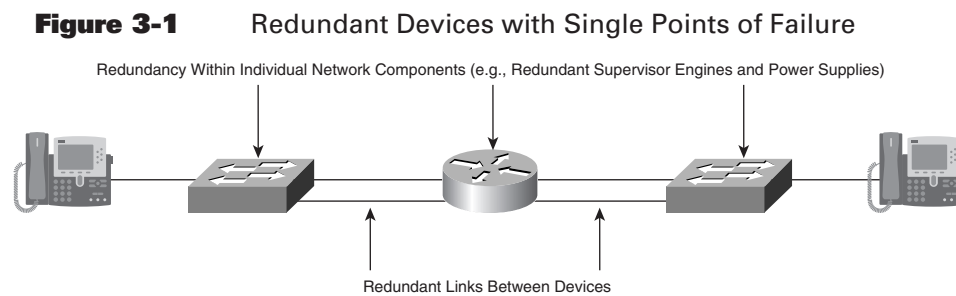
Table 3-1 Availability and Downtime

Availability	Maximum Yearly Downtime
99.000 percent (two nines)	3 days, 15 hours, and 36 minutes
99.900 percent (three nines)	8 hours, and 46 minutes
99.990 percent (four nines)	53 minutes
99.999 percent (five nines)	5 minutes
99.9999 percent (six nines)	30 seconds

Before discussing how a VoIP network can be designed to be more available, we need to distinguish between *reliability* and *availability*. A *reliable* network, as an example, does not drop many packets, whereas an *available* network is up and functioning. Availability is a function of the *mean time to repair (MTTR)* and the *mean time between failures (MTBF)*.

As the names suggest, the MTTR is the average time it takes to repair a failed network component, and the MTBF is the average time between the failures of a network component. A network’s availability can be improved by reducing the MTTR and increasing the MTBF. When purchasing network hardware (for example, an Ethernet switch), many manufacturers, such as Cisco, provide MTBF information; and you can determine the MTTR as part of your network design. For example, you might have spare parts on-site to quickly swap out failed equipment, or you might have redundant components within a chassis (for example, redundant supervisor engines in a Cisco Catalyst switch). These network components can also be interconnected in a redundant fashion (for example, having multiple connections between multiple devices). Let’s consider some of these design approaches in a bit more detail.

One approach is to have fault tolerance built into the network components. So, even though the perception of VoIP reliability is still growing, we can actually design VoIP networks that are just as reliable as legacy PBX systems. Notice in Figure 3-1 that there are dual physical connections between all network components.



As an example, we might have a Cisco Catalyst 6500 multilayer switch with redundant features built into the chassis itself, including:

- Two supervisor engines (the “brains” of the switch)
- Two power supplies
- Two switch fabric modules (which increase the throughput of the switch)


Not only do these modules help minimize the MTTR, but they are also hot-swappable. For example, if one supervisor engine were to fail, the other supervisor engine could step in and take over that responsibility, and the failed supervisor engine could be removed from the chassis and replaced without powering down the chassis.

Instead of having redundancy built into the router or switch itself, another design approach is to have redundancy between devices, as shown in Figure 3-2. Notice that in this topology, any single network link or network infrastructure device (for example, switch or router) can fail (with the exception of the wiring closet switch where the IP phone attaches); and a path will still exist from the host to the server.



Redundant design approaches such as these benefit not only voice networks but also data networks. In fact, many network redundancy features were available well before the introduction of VoIP. However, the critical nature of voice traffic is causing many network designers to beef up the redundancy in their existing data networks.

Figure 3-2 No Single Points of Failure



End systems not running a routing protocol point to a default gateway. The default gateway is traditionally the IP address of a router on the local subnet. However, if the default gateway router fails, the end systems are unable to leave their subnet. Two approaches to Layer 3 redundancy include *Hot Standby Router Protocol (HSRP)* and *Virtual Router Redundancy Protocol (VRRP)*. With both of these technologies, the *Media Access Control (MAC) address* and the IP address of the default gateway can be serviced by more than one router. Therefore, if a default gateway router goes down, then another router can take over, still servicing the same MAC and IP addresses:

- HSRP is a Cisco-proprietary approach to Layer 3 redundancy.
- VRRP is a standards-based approach to Layer 3 redundancy.

Layer 3 redundancy is also achieved by having multiple links between devices and selecting a routing protocol that load balances over the links. EtherChannel is another way to load balance across multiple links. With EtherChannel, you can define up to eight physical links that are logically bundled together, such that the bundle appears as a single link to the route processor.

Although having multiple links between switches is great for redundancy, these links can cause loops in the Layer 2 (that is, switching) network. These Layer 2 loops can cause *broadcast storms*, where broadcast packets circle the network forever, consuming bandwidth and switch processor resources. The IEEE 802.1D standard is the legacy approach for Layer 2 loop avoidance. IEEE 802.1D is better known as the *Spanning Tree Protocol (STP)*. By default, in the event of a link failure, STP takes 50 seconds to recover and start forwarding traffic over a backup link (that is, to *converge*). Cisco added proprietary enhancements to speed up the convergence time. These Cisco-proprietary STP enhancements include:

- **PortFast**—Used on ports connecting to end stations
- **UplinkFast**—Used on building access switches
- **BackboneFast**—Used on all switches in the topology

Each *virtual LAN (VLAN)* can run its own instance of STP. This Per-VLAN STP approach allows different VLANs (that is, subnets) to have different *root bridges* (that is, switches in the Layer 2 network that serve as the points to which other switches forward traffic). However, with the Cisco Per-VLAN STP, every VLAN must run its own instance of STP, which might place unnecessary overhead on the switches.

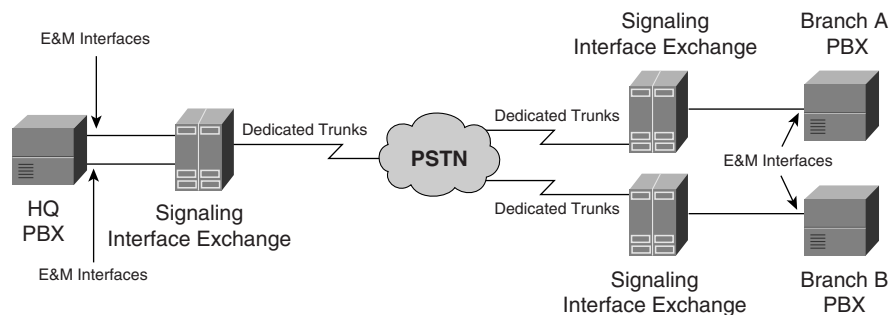
The best of both worlds is achieved with the new IEEE 802.1w and 802.1s protocols. IEEE 802.1w (that is, Rapid Spanning Tree Protocol) dramatically reduces convergence times in the event of a failure. IEEE 802.1s (that is, Multiple Spanning Tree) allows you to create a set of STP instances. Then VLANs might be assigned to appropriate STP instances. This eliminates the Per-VLAN STP requirement that each VLAN run its own instance of STP.




Replacing PBX Trunks: Out with the Old, In with the New

PBX systems often span a company's geographic locations, as shown in Figure 3-3. The connections used to tie the PBXs together are *trunks*, and a company pays monthly recurring charges to its telephone carrier for these trunk connections. However, more than calls are carried over these trunks. PBXs have *signaling protocols* used to communicate call setup information between them. For example, dialed digits and on-hook/off-hook conditions are sent across these trunk connections, too.

Figure 3-3 PBXs Before Convergence



These trunks are likely candidates for replacement as a company begins its migration toward a VoIP network. This initial migration step doesn't necessitate throwing out the PBXs. The PBXs can continue to service the telephony needs of the company. However, the difference, as shown in Figure 3-4, is that the PBXs are interconnected by the IP WAN instead of the trunk lines. By eliminating these trunk lines, the company can eliminate the recurring cost of the dedicated trunk lines.

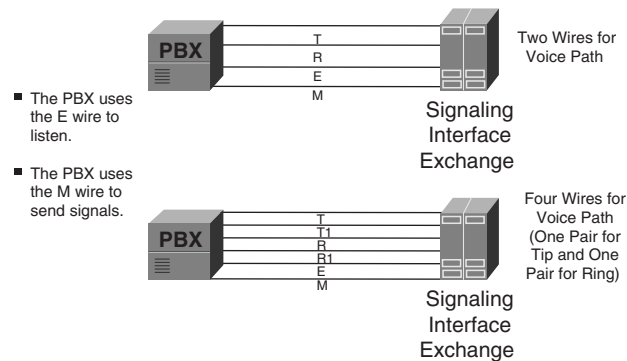
Figure 3-4 PBXs After Convergence

The connection from a PBX to a local router can take one of several approaches. In Figure 3-4, notice that the PBX connects to the router using an *Ear and Mouth* (*E&M*) connection. An E&M interface is an analog interface present in many of today's PBX systems. Many say that the E stands for Ear, and the M stands for Mouth. Other authorities say that the E stands for Earth, and the M stands for Magneto. You might also read that the E is the E in rEceive, and the M is the M in transMit. Personally, I use the Ear and Mouth definition because it gives me a great visual image of the E lead being used for the receive function, and the M lead being used for the transmit function.

Note that the actual voice path doesn't use the E or M leads (that is, wires). The E&M leads are used for call signaling (that is, setting up and tearing down a call), but an E&M connection still uses *tip and ring* wires to transmit the actual voice. Instead of always having a single tip wire and a single ring wire, in some instances, two wires are used for "tip," and two wires are used for "ring," as shown in Figure 3-5.



Figure 3-5 E&M Connections



If a company already has an E&M interface in its PBX being used to form a trunk with a remote PBX, we can help preserve the company’s original investment in its PBX E&M interface by connecting the PBX’s E&M interface to an E&M interface in a router. The router then connects into the IP WAN using a traditional WAN interface, such as a serial interface. We don’t have to use an E&M interface, however, to interconnect a PBX and a router. Other analog and digital interfaces are also available, as we’ll discuss in the next couple of sections.

Connecting a Router to a Phone Line

As voice travels from a LAN to a PBX or to the *public switched telephone network (PSTN)*, it needs a “translator” to convert back and forth between those two environments. Voice on a LAN-based VoIP network takes the form of packets, whereas voice traveling to the PBX or PSTN might be analog waves or digital signals. The job of this translator can be performed by a voice-enabled router. When a router acts in this capacity, it is called a *gateway*.

A gateway typically has at least one interface that connects to the LAN (for example, an Ethernet or Fast Ethernet interface) and at least one interface that connects to the PBX/PSTN environment. These PBX/PSTN interfaces might be either analog or digital, as shown here:


Analog Interfaces:

- FXS (Foreign Exchange Station)
- FXO (Foreign Exchange Office)
- E&M

Digital Interfaces:

- T1
- E1
- Basic rate interface (BRI) /primary rate interface (PRI) (which use *Integrated Services Digital Network [ISDN]* technology)

First, let's consider what connects to an FXS port. A FXS port connects to a *station*, such as an analog phone, fax machine, or speaker phone, as shown in Figure 3-6. Consider the analog phone you have in your home. Just as you can connect that analog phone into your RJ-11 wall jack (which goes back to the telephone company), you can also connect that phone into an FXS port. The FXS port can provide the attached device with -48 VDC to power the phone. Ringing voltage can be sent from the FXS port to the device, and the FXS port can recognize digits dialed by the attached device.

Figure 3-6 FXS Connections



A Cisco router is configured using the Cisco Internetwork Operating System (IOS). Using the IOS, we can configure the characteristics of an FXS port, including the following parameters:

- **Signal type**—An FXS port on a Cisco router defaults to *loop start signaling*. However, for some applications, such as connecting a PBX trunk port into an FXS port, we might prefer to use *ground start signaling*.
- **Call progress tones**—A call progress tone gives the caller an idea of how the call is progressing. For example, if you call your friend, before your friend answers the phone, you hear *ring back* in your ear, indicating that your friend's phone is indeed ringing. If your friend is already on the phone when you call, you might instead hear a *busy signal*. Both ring back and a busy signal are examples of call progress tones. However, these call progress tones might vary from country to country. A Cisco router defaults to call progress tones heard in the United States. However, we can alternately configure the FXS to ports to use call progress tones common to other countries.
- **Ring pattern**—If you live in the United States, chances are, when your home phone rings, the ringing lasts for two seconds, followed by a four-second pause, followed by two seconds of ringing, and so on. This ring pattern (sometimes called the *ring cadence*) might vary in different countries. Fortunately, the Cisco IOS allows us to configure a predefined ringing pattern to be sent out of an FXS port, or we can define a unique ringing pattern.

If you are configuring several phones in an office (for example, in a cubicle environment), it might be wise to configure different ringing patterns for different phones. Then, when all of the employees are gathered around the water cooler, and a phone rings, an employee will be able to know that it is his phone ringing due to the distinctive ringing pattern you configured.


- **Ring frequency**—When I was five years old, my family's home phone was on a *party line*. (No, not one of those 900 number party lines.) The party line allowed more than one home to share the tip and ring wires going back to the *central office (CO)*. As a result, only one home could use its phone at any one time (unless the homes were talking with each other). But the question is, "If we have more than one phone on the same tip and ring circuit,

how can we make only one phone (the phone that was called) on that party line ring?” Back in those days, phones belonged to the phone company. We could not just go down to the local Wal-Mart and buy one like we can today. Because the telephone company controlled who got which phone, it could give a phone with one ring frequency to one party line member and a phone with another ring frequency to another party line member. These phones had a mechanical ringer, and these ringers were tuned to only ring at a specific frequency. If any of these phones are still being used in your VoIP environment, you might need to adjust the ringing frequency used by an FXS port to make the phone ring. However, this is not a concern for most modern phones, which use piezoelectric speakers. These ringers sound the same, regardless of the ringing frequency we specify.

- **Caller-ID information**—A popular feature on many home telephones today is *caller-ID*, which allows a called party to see who is calling them. On an FXS port, we can configure the caller-ID information that the router transmits over the VoIP network to the destination phone.

An FXO port connects to an *office* (that is, a phone switch such as a PBX or a switch in the local CO). For example, you could connect a router’s FXO port to the RJ-11 wall jack in your home (which goes back to the telephone company). Or, you could connect an FXO port into the station side of a PBX, as shown in Figure 3-7. Therefore, we could say that an FXO port acts like a phone. It can place calls, receive calls, and dial digits (using either *dual tone multifrequency [DTMF]* or pulse dialing).

Figure 3-7 FXO Connections





Using the IOS on a Cisco router, we can configure the characteristics of an FXO port including:

- **Signal type**—Just as we can select between loop start and ground start signaling on an FXS port, we can also select the signal type of an FXO port. Like the FXS port, the FXO port defaults to loop start signaling.
- **Ring number**—When I was first learning about VoIP, I connected a router's FXO port to my home telephone line, and set the router up so that a phone connected to an FXS port could call out to the PSTN. What I didn't consider, however, was that an FXO port can answer a call, and, by default, an FXO port answers a call after only one ring. So, when I had that FXO port connected into the phone wall jack in my home, if someone called my home, the FXO port on the router answered. As soon as the router answered, the caller would hear the dial tone, which was understandably confusing for the caller. The reason a caller would hear the dial tone was the FXO's *ring number* (that is, the number of rings received on an FXO port before the port answers the call) was at the default value of one. By default, an FXO port plays the dial tone when it answers a call, allowing the caller to call another number known to the router. However, the FXO port supports other options. For example, an FXO port can be configured to forward a call to a predetermined number after it answers, or the FXO port could look at the dialed digits and forward the incoming call based on those dialed digits.
- **Dial type**—On your home phone, you probably have an option of selecting either pulse dialing or DTMF dialing. Most locations in the United States now have COs that support DTMF dialing. However, in some parts of the world, we might need to use pulse dialing (that is, the type of dialing used by older rotary phones); and we have the option of changing the dial type on our Cisco FXO port from the default of DTMF dialing to pulse dialing.

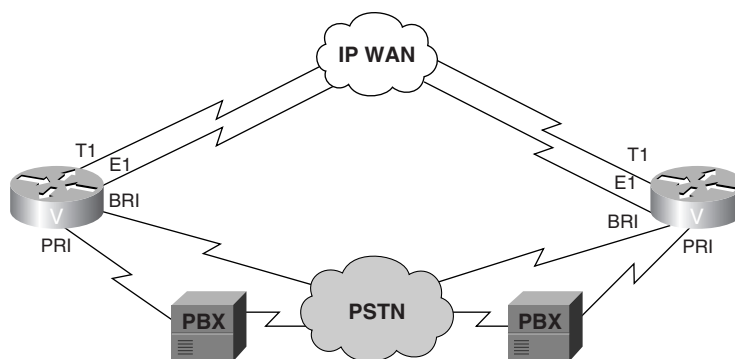
Connecting a Router to a Digital Circuit

Although analog connections are great for connecting a router to a phone, a PBX, or even to the PSTN, as the need for more connections grows, so does the expense

associated with adding more and more interfaces to the router. Digital connections are often a more cost-effective solution when we have more than approximately eight connections, because a single digital connection can carry multiple conversations over a single circuit.

Whereas analog interfaces send and receive analog waveforms that continually vary, digital interfaces send binary 1s and 0s, which are represented on the wire as the presence or absence of voltage. Examples of digital circuits include T1, E1, and ISDN circuits, as shown in Figure 3-8.

Figure 3-8 Digital Connections




The question is, “How can multiple conversations be sent across a single connection?” Just like we learned to do as children, the multiple conversations share and take turns. Specifically, they share the bandwidth by taking turns sending data on the wire.

Consider a T1 circuit, which has 24 separate channels. With a T1, we can do *time-division multiplexing (TDM)*. With TDM, a T1 circuit can send an 8-bit sample from its first channel, followed by an 8-bit sample from its second channel, followed by an 8-bit sample from its third channel, and so on. Each channel, which means each conversation, gets its own time slice in which it can transmit its voice, represented as binary 1s and 0s. In fact, we could say that with TDM, each voice



conversation has, to borrow a line from Whitney Houston, *one moment in time*, as shown in Figure 3-9.

Figure 3-9 Time-Division Multiplexing (TDM)



For years I heard that a T1 had 1.544 Mbps of bandwidth, and I knew a T1 had 24 channels, where each channel had 64 kbps of bandwidth. Then, one day I did the math. I multiplied 24 and 64,000, but to my surprise, I did not get 1,544,000 as a result. Instead, the result was 1,536,000. That really confused me. What happened to the extra 8000 bits?

What I did not consider were the *framing bits*. A framing bit is a single bit that indicates the end of the frame, and a frame contains an 8-bit sample from each of a T1's 24 channels. Once I accounted for the framing bit, the math worked out beautifully.

Each frame is 193 bits in size:

$$24 \text{ channels} * 8 \text{ bits per channel} + 1 \text{ framing bit} = 193 \text{ bit frames}$$

The Nyquist Theorem requires that we send 8000 samples per second:

$$\begin{aligned} \text{Samples per second} &= 2 * \text{the highest frequency being sampled} \\ &= 2 * 4000 \\ &= 8000 \end{aligned}$$

The total bandwidth on a T1 is 1.544 Mbps:

$$193 \text{ bit frames} * 8000 \text{ samples per second} = 1.544 \text{ Mbps}$$

However, in a T1 environment, we don't typically send just one frame at a time. Instead, we connect multiple frames together and send them all at once. Two popular approaches to grouping these frames together are:

- **SF**—Combines 12 standard 193-bit frames into a *Super Frame*
- **ESF**—Combines 24 standard 193-bit frames into an *ESF*

When configuring a T1 interface (also known as a *T1 controller* on a Cisco router), the T1 interface defaults to SF as the framing type. The good news is that we do not have to be concerned with selecting a particular framing type. Because our T1 connects to a service provider, the service provider tells us what framing type to use, and we simply configure our router to match the service provider's parameters.


Another piece of T1 configuration information given to us by our service provider is the *line coding*. A T1 circuit's line coding is the set of rules that dictates how binary 1s and 0s are represented over the wire.

We normally think of binary 1s being the presence of voltage and binary 0s being the absence of voltage. Although that is true, the goal on a T1 line is to keep the average voltage on the line 0 volts, which means when we send a binary one using a positive voltage, the next binary 1 uses a negative voltage. Therefore, on average, the voltage on the wire is 0. Personally, if I put one hand in a bucket of boiling water and the other hand in a bucket of freezing water, on average, I'm not going to be comfortable, but I concede this approach does work for digital circuits.

If two consecutive voltages have the same polarity, an error, called a *bipolar violation*, occurs. The approach of representing binary 1s as alternating voltages is called *alternate mark inversion (AMI)*, as shown in Figure 3-10.

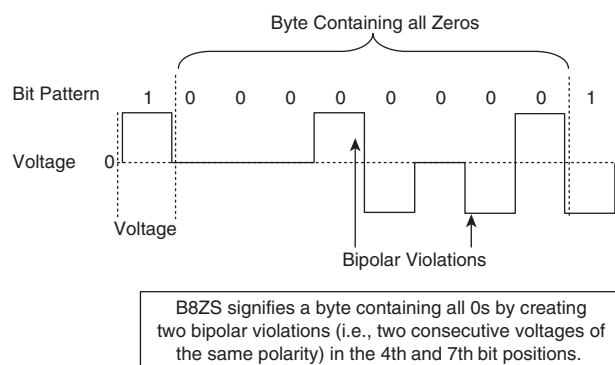


Figure 3-10 Alternate Mark Inversion (AMI)



Although AMI does meet the goal of maintaining an average of 0 volts on the circuit, it has a major challenge. AMI has issues when it attempts to send a byte containing all 0s (that is, eight binary 0s in a row). Although there are various workarounds that address this issue, errors can occur when sending eight 0s in a row over a T1 circuit using generic AMI line coding.

Due to AMI's limitation, another type of line coding was developed. *Bipolar 8-zero substitution (B8ZS)* can represent a byte containing all 0s by creating a couple of bipolar violations. If a T1 circuit using B8ZS line coding experiences two bipolar violations at very specific bit positions, as shown in Figure 3-11, the equipment the T1 connects to (for example, a router) knows that a byte containing eight 0s is being transmitted. Therefore, in the case of B8ZS, two wrongs really do make a right. While T1 circuits commonly use B8ZS, you might see something called High Density Binary 3 (HDB3) used on E1 circuits. Like B8ZS, HDB3 overcomes the limitations of AMI.

Figure 3-11 Bipolar 8-Zero Substitution (B8ZS)

Just as an FXS port needs some type of signaling (for example, loop start or ground start) to determine when a phone is on-hook or off-hook, a T1 circuit also needs a signaling mechanism. Two approaches to sending signaling across a T1 circuit include:

- **Common Channel Signaling (CCS)**—With CCS, one or more channels are dedicated to sending a signaling protocol, while each of the other channels carry, for example, a voice conversation.
- **Channel Associated Signaling (CAS)**—With CAS, framing bits are “robbed” from the Super Frame or Extended Super Frame and used for signaling bits. This approach is sometimes referred to as *robbed-bit signaling*. Because none of the 24 channels are dedicated to just sending signaling information, unlike CCS, all 24 channels can be used.

Let us consider each of these approaches in a bit more detail. The simplest approach to understand is CCS. As the name suggests, all of the channels used for sending voice, video, or data use the same channel (that is, a “common channel”) to send signaling information. A signaling protocol is sent over this dedicated channel.

A popular technology that leverages CCS is ISDN. An ISDN circuit is made up of *B-channels* and a *D-channel*. A B-channel is a “bearer” channel, which carries the voice, data, or video. These bearer channels typically carry information at a rate of



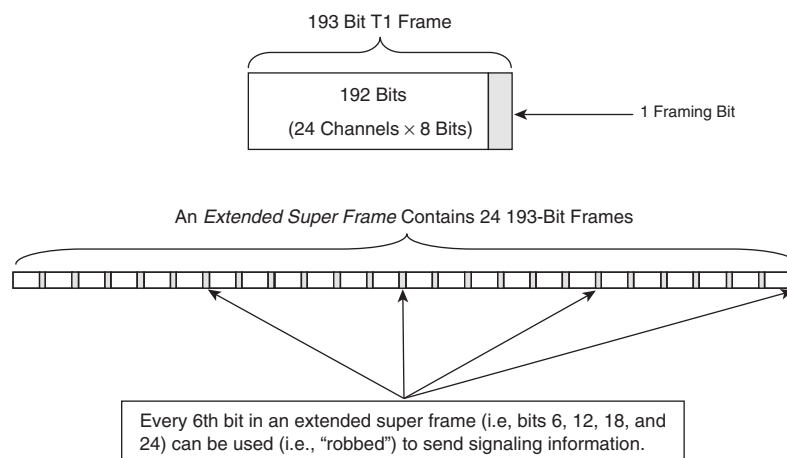
64 kbps. The D-channel acts as the “signaling” channel, meaning that the D-channel carries the data necessary to set up and tear down calls on the B-channels. Depending on your bandwidth needs, you might select either the BRI or the PRI flavor of ISDN.

- **BRI**—BRI ISDN connections contain two 64-kbps B-channels and one 16 kbps D-channel, for a total usable bandwidth of 128 kbps.
- **PRI**—A PRI ISDN connection can use the channels on either a T1 or an E1 circuit. If the PRI is based on a T1 circuit, 23 of the T1’s 24 channels are used as B-channels, and the remaining channel serves as the D-channel, for a total usable bandwidth of 1.472 Mbps. However, if the PRI is based on an E1 circuit, 30 of the E1’s 32 channels are used as B-channels. One of the 32 channels carries framing and synchronization information, while the remaining channel acts as the D-channel, carrying the signaling information for the 30 B-channels.

The D-channel in each of these instances uses *Q.931* as its signaling protocol. PRI ISDN connections are often used to connect a company’s PBX to the PSTN. However, we might see BRI ISDN used in a small office/home office (SOHO) environment.

ISDN was developed during the 1980s and is, therefore, a very mature protocol. When I was first introduced to ISDN, back in 1988, web browsers were not available yet, and the thought of having 128 kbps of bandwidth in a home seemed to be overkill. In fact, we used to say that the acronym ISDN stood for “I Still Don’t Need it.”

Next, consider how CAS carries signaling information for a T1. Recall that a T1 doesn’t send individual frames. Rather, a T1 sends a Super Frame (containing 12 standard frames) or an Extended Super Frame (containing 24 standard frames). Therefore, an Extended Super Frame contains 24 framing bits, one bit from each standard frame it contains. The Extended Super Frame does not need all 24 of these framing bits. So, some of those bits can be used to send signaling information. Specifically, every sixth bit in a Super Frame or an Extended Super Frame can be used as a signaling bit, as shown in Figure 3-12.

Figure 3-12 “Robbed-Bit” Signaling

Because the CAS approach takes these unneeded framing bits and uses them for signaling, this approach is often referred to as “robbed-bit signaling.” With CAS, all 24 of a T1’s channels can be used for voice, data, or video because none of the channels are dedicated solely to signaling.

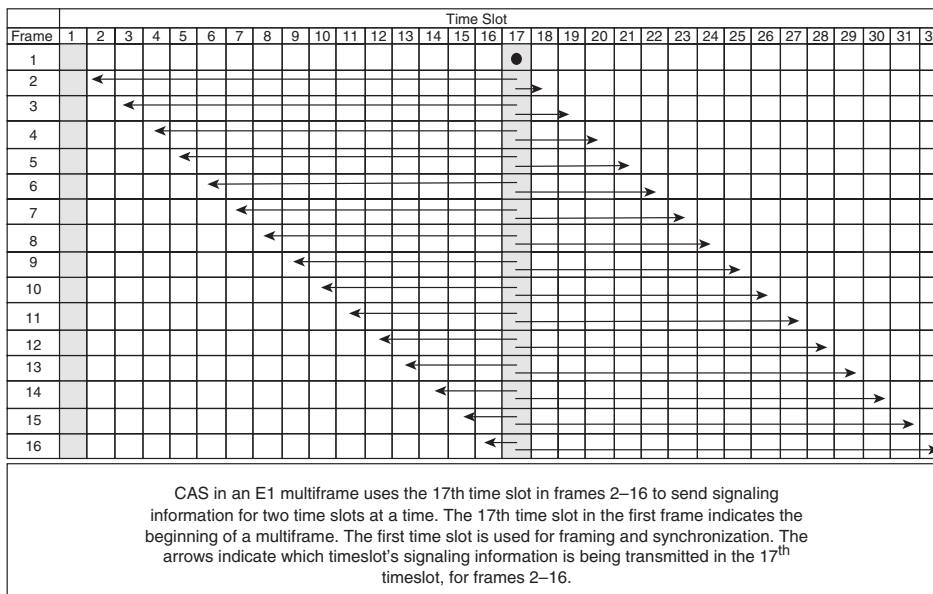
Just as T1 circuits are popular in North America, E1 circuits are commonplace in Europe. An E1 circuit has 32 channels, as opposed to the 24 channels available in a T1. The first of those 32 channels is dedicated to framing and synchronization, while the seventeenth channel is dedicated to signaling. Coming off our discussion of how a T1 can free up its signaling channel using CAS, it might be tempting to think we could do the same with an E1 circuit, giving us 31 usable channels to send our voice, video, and data. However, an E1 circuit approaches CAS very differently than a T1.

On a standard E1 circuit, the seventeenth channel is always used for signaling, regardless of whether we are doing CAS or CCS. The good news at this point is that you don’t have to relearn how CCS is performed, because like a T1, a signaling protocol (for example, Q.931) is sent over an E1’s signaling channel.



We should spend some time, however, delving into how an E1 CAS functions. To begin with, you need to understand that an E1 doesn't use the Super Frames or Extended Super Frames you saw in the T1 world. Rather, an E1 combines 16 frames together in a *multiframe*. If we examine the first frame in a multiframe and look at its seventeenth channel, we discover that the seventeenth channel indicates the beginning of this multiframe. But then if we take a close look at the second frame in a multiframe, we see that its seventeenth channel is used to send signaling information. Specifically, 4 bits of signaling information for channel number 2 and 4 bits of signaling information for channel number 18 are carried in the seventeenth channel of the second frame in an E1 multiframe. Similarly, the seventeenth channel of the third frame in a multiframe carries 4 bits of signaling information for channel 3 and 4 bits of signaling information for channel 19, as shown in Figure 3-13. This process continues for each of the remaining frames in the multiframe, such that the multiframe sends signaling information for 30 channels, which is exactly the number of channels we use in an E1 to send voice, video, and data.

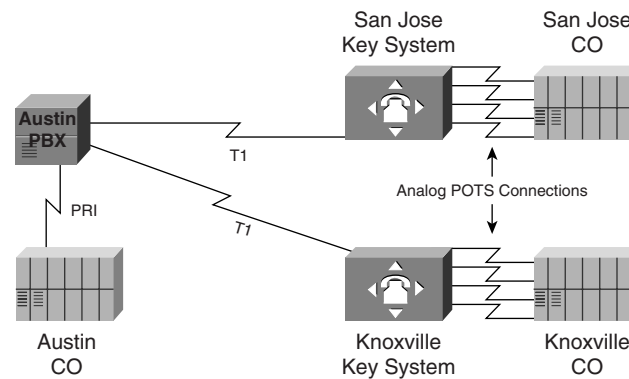
Figure 3-13 E1 Multiframe



Thus far in this chapter, we examined how an IP WAN can replace a PBX-to-PBX trunk connection and how our Cisco router can connect to various analog and digital ports. Let's put all the pieces together by considering a sample VoIP migration scenario.

In this scenario, our company currently has a main office in Austin, TX and two branch offices, in San Jose, CA and Knoxville, TN. The Austin location has a PBX system, and each branch office has a key system. The key systems each have a dedicated T1 trunk connection back to the PBX in Austin. To support the Austin office's relatively high call volume, an ISDN PRI connection connects the Austin PBX to the local telephone company's CO. The branch offices each have four Plain Old Telephone Service (POTS) telephone lines connecting to their local COs to support local calls, as show in Figure 3-14.

Figure 3-14 Scenario Topology – Before Migration



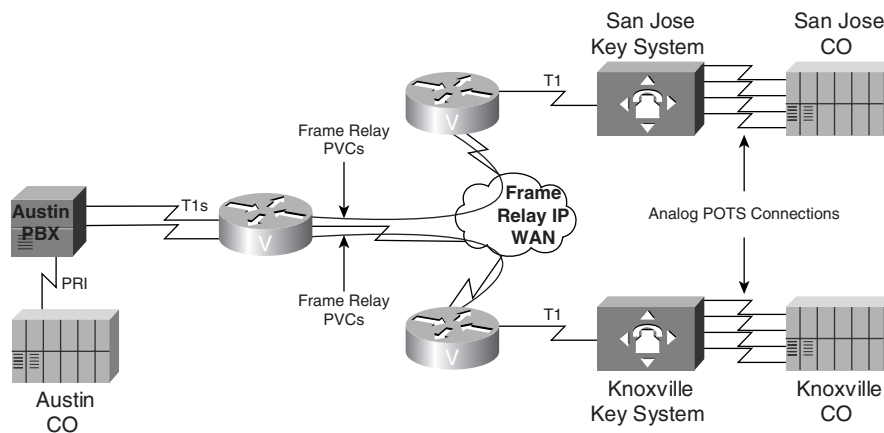
Our goal in this scenario is to replace the key system-to-PBX trunk connections with VoIP connections and, in preparation for removing the PBX, to have the Austin, San Jose, and Knoxville CO connections terminate on a router, as opposed to a PBX or a key system.

As a first step, we can replace the existing trunk connections from the branch offices to the main office with VoIP connections over the IP WAN. Because the PBX and key systems already have T1 interfaces, we can leverage the company's

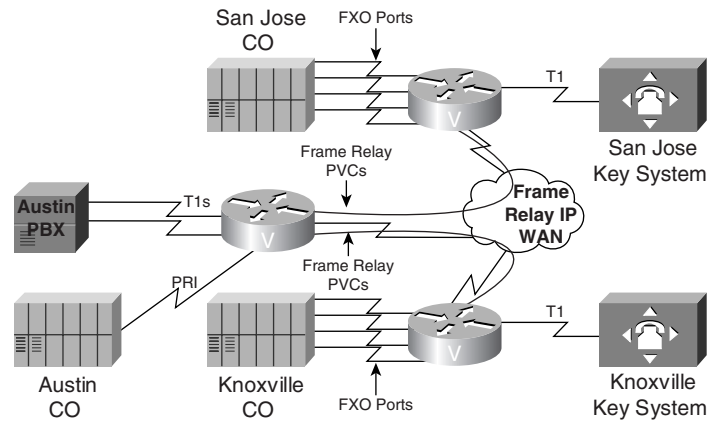


existing investment in these interfaces and purchase T1 interfaces for our Cisco routers. The PBX at the main office (that is, the Austin office) can then connect to a router located at the Austin location via a T1 connection. Similarly, the key systems at the San Jose and Knoxville locations can connect to local routers using T1 connections. The routers at these locations all connect into a service provider's IP WAN. For this scenario, assume the routers connect into a Frame Relay network using a hub-and-spoke topology, where each of the branch offices has a Frame Relay *permanent virtual circuit (PVC)* connecting back to the main office, as shown in Figure 3-15.

Figure 3-15 Scenario Topology – Migration Step 1



Our VoIP migration already eliminated the recurring cost of the dedicated PBX-to-key system trunk connections. However, another requirement was to take the CO connections at each location and terminate those connections on a router. At the main site, the PBX currently connects to the local CO using an ISDN PRI circuit. Therefore, we can install a T1 interface in our Cisco router, and configure that interface to function as an ISDN PRI interface. Then we can move the PRI connection from the PBX to the router. In a similar fashion, we can move the existing POTS telephone lines from the key systems at the branch office locations and terminate those lines on FXO ports in the local routers, as shown in Figure 3-16.

Figure 3-16 Scenario Topology – Migration Step 2

Even though this scenario did not involve converting any of the company's phones to IP phones or connecting any phones directly to a router, we did eliminate recurring costs for trunk lines and simultaneously laid the foundation for an IP telephony network that can, in the future, replace the PBX at the headquarters with a Cisco CallManager cluster; and replace the key systems at the branch offices with, perhaps, a *CallManager Express (CME)* router. Also, this future IP telephony network can replace existing analog and digital phones, currently connecting to the PBX and key systems, with IP phones.

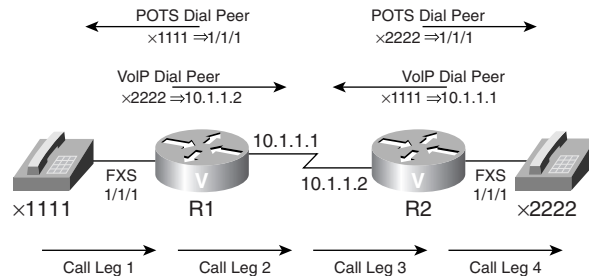
**note**

The Cisco CallManager cluster and the Cisco CME router mentioned in this section are discussed in the next chapter.

At this point, we have seen how voice ports are used on our voice-enabled routers. However, the routers are not yet trained to reach specific destinations. In order to give our routers call routing intelligence, we create *dial peers* that inform our routers how to reach specific phone numbers. Consider the topology in Figure 3-17.



Figure 3-17 Dial Peers and Call Legs



Routers R1 and R2 each have a POTS dial peer that points to their locally attached phone, and a VoIP dial peer that points to the IP address of the remote router.

Therefore, when extension 1111 dials extension 2222, router R1 searches for a dial peer that matches a destination pattern of 2222. In this case, R1 has a VoIP dial peer that points to R2's IP address of 1.1.1.2. R1 then forwards the call to R2. R2 then receives the incoming call destined for extension 2222. R2 searches for a dial peer that matches a destination of 2222, and it finds a POTS dial peer that specifies FXS port 1/1/1. The FXS port then sends ringing voltage out port 1/1/1. Extension 2222 rings and goes off-hook, and the end-to-end connection is complete.

Notice that there are a total of four dial peers, which allow a call in the opposite direction. Also, notice that four stages of the call (that is, *call legs*) are defined—two call legs from the perspective of each router:


- Call Leg 1: The call comes into R1 on FXS port 1/1/1
- Call Leg 2: The call is sent from R1 to IP address 1.1.1.2
- Call Leg 3: R2 receives an incoming call destined for x2222
- Call Leg 4: R2 forwards the call out FXS port 1/1/1

Voice over IP in the Home

Now that we have examined how we can begin to migrate our corporate telephony systems to VoIP, let's bring the conversation home, literally. After all, many homes today have broadband connections, such as digital subscriber line (DSL) or cable modem connections, and many service providers are beginning to offer telephone service over these broadband connections.

You can purchase a small router for your home that has one or more analog telephone jacks and sign up for VoIP service with a carrier, such as *Vonage* (<http://www.vonage.com>). Cisco makes a router that can be used with the Vonage service. This router also has an Ethernet connection, allowing you to connect it into your existing home network, or directly into your DSL router or cable modem, as shown in Figure 3-18. When you pick up your phone to place a call, the router supplies your phone with dial tone, and it can interpret the phone's dialed digits.

Figure 3-18 VoIP in the Home





The router in your home forwards the dialed digits over the Internet to the carrier's equipment located in a telephone company's CO. Often times, the carrier leases space for its equipment in the CO. This type of leasing arrangement is called a *co-lo* (that is, "co-location"). The carrier's equipment in the CO connects into the traditional telephone company's network, using *Signaling System 7 (SS7)* as a signaling protocol.

Once an end-to-end call is set up, the router in your home converts your voice (that is, analog waveforms) into packets. These packets are transmitted over the Internet to the carrier's equipment in the CO, which sends your voice into the PSTN.

You can, in some cases, keep your current phone number if you convert your existing telephone service to a VoIP-based telephone service. However, because these services are not installed in every local CO, if your home is in a more rural location, the nearest CO you can connect to over the Internet might be in another city, which might mean that you cannot keep your existing phone number. Your friends and family might also need to pay a long-distance charge to reach your new number in the other city. These are just a couple of caveats to watch out for when subscribing to one of these residential VoIP services.

Another critical consideration when signing up for VoIP in your home is 911 emergency service. Please check with your VoIP carrier for the specifics of how its 911 service functions. You might, for example, need to activate 911 service for your line, and you might need to specify the physical location associated with your VoIP phone. Not all residential VoIP services offer Enhanced 911 (E-911), which can automatically send the caller's location to the 911 operator. Therefore, a caller might need to clearly state his location to the 911 operator. Because 911 service can literally mean life or death, and because 911 services with VoIP carriers vary, be sure you understand the specifics of how 911 services are provided by any VoIP carrier that you consider signing up with. Also, if communication between your home phone and the PSTN flows through a VoIP router, consider what would happen if you experienced a power outage. Without some sort of power backup, you would not be able to place any calls, because your VoIP router (and any other broadband router/switch equipment) would be unpowered.

Case Study: Your Turn to Put the Pieces of the Puzzle Together

You designed a telephony network for the XYZ Company in the case study in Chapter 1, “Touring the History Museum of Telephony.” This design was based on traditional PBX and key system technologies. Then, in Chapter 2, “Making Waves: Turning Your Voice into Zeros and Ones,” you calculated the bandwidth required to interconnect the XYZ Company’s headquarters with two of its remote offices.

Based on your previous recommendations, the XYZ Company has decided to interconnect its PBXs and key system over the IP WAN, as opposed to using the T1 and fractional T1 connections specified in Chapter 1’s case study. Therefore, your goal in this case study is to create a design that places a router at each XYZ Company location and to specify the router interfaces used to connect remote routers. Also, each XYZ Company location needs connections to its local CO as follows:

- Headquarters—48 connections to local CO
- Remote Office 1—24 connections to local CO
- Remote Office 2—5 connections to local CO

In addition to the required voice bandwidth you calculated in Chapter 2’s case study, the links in this converged network solution also need to transport data. The data bandwidth requirements are as follows:

- Link from HQ to Remote Office 1: 768 kbps for data
- Link from HQ to Remote Office 2: 128 kbps for data

Demonstrate in your design how the router at each XYZ Company location connects into its local CO. This VoIP design lays the foundation for eventually replacing the PBXs and key system units with IP telephony components such as the Cisco CallManager, as discussed in the next chapter.

Use the following area to sketch your design.



Design for XYZ Company:

Suggested Solution

Although multiple solutions exist for the design scenario presented, the following is a suggested solution that meets the design criteria. The router located at the XYZ Company HQ needs WAN connectivity with each remote office. Based on the bandwidth requirements calculated in Chapter 2's case study, we have the following voice bandwidth requirements for each of these WAN links:

- Link from HQ to Remote Office 1: 2026.08 kbps for voice
- Link from HQ to Remote Office 2: 337.68 kbps for voice

In addition to voice, these intersite links need to carry data. The data requirements specified in the case study were:

- Link from HQ to Remote Office 1: 768 kbps for data
- Link from HQ to Remote Office 2: 128 kbps for data

To accommodate both voice and data, we sum the required bandwidth (that is, the voice bandwidth plus the data bandwidth) for these links to calculate each link's total required bandwidth:

- Link from HQ to Remote Office 1: 2794.08 kbps total
- Link from HQ to Remote Office 2: 465.68 kbps total

Although we could select from various WAN technologies to interconnect these sites (for example, Frame Relay, Point-to-Point Protocol [PPP], or Asynchronous Transfer Mode [ATM]), this suggested solution uses PPP links. The headquarters connects to Remote Office 1 using two T1s. Two T1s actually provide more bandwidth than the required 2.794 Mbps. However, these two T1s allow room for future growth. In this design, the two T1s are combined into a single logical link, using *Multilink PPP (MLPPP)*.

The headquarters connects to Remote Office 2 using a fractional T1, running at 512 kbps. Again, the allocated bandwidth is a little more than the required 465.68 kbps, thus allowing for future growth.

The routers at each site should also connect to their local CO. As a best practice, once we exceed eight connections, we should choose digital interfaces instead of

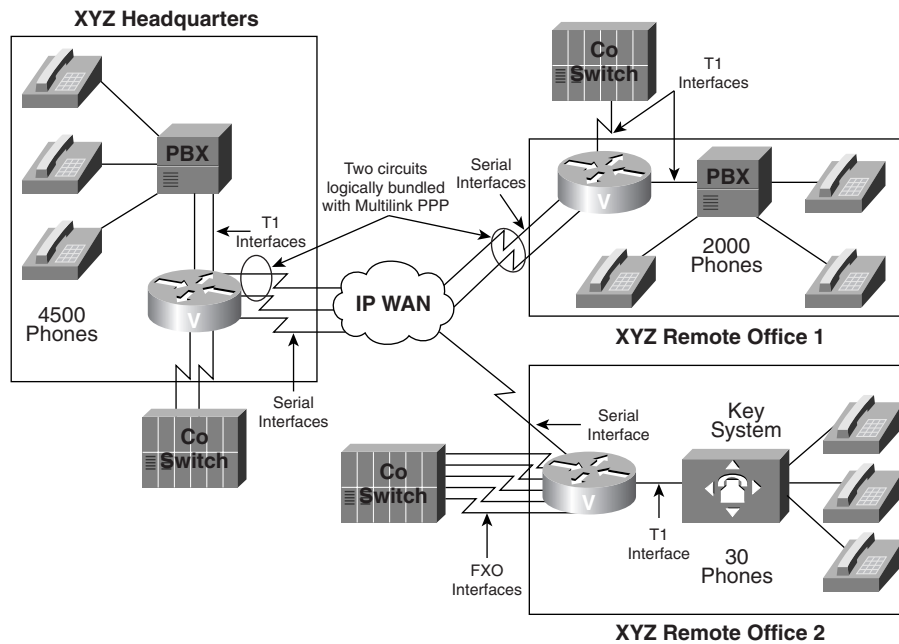


analog interfaces, due to cost considerations. Therefore, this suggested design selected the following CO connections for the XYZ Company routers:

- Headquarters Router: Two T1 connections
- Remote Office 1 Router: One T1 connection
- Remote Office 2 Router: Five FXO connections

These suggested design solutions result in the design shown in Figure 3-19.

Figure 3-19 XYZ Company's Suggested Solution



Chapter Summary

This chapter presented a major design challenge for VoIP: achieving high availability. Specifically, our VoIP availability design goal is the “five nines,” 99.999 percent of availability, which equates to approximately five minutes of downtime per year.

Most companies perform a phased migration to VoIP, as opposed to a *forklift upgrade*. Interconnecting existing PBX/key system units over the IP WAN is often the first step in a phased migration, and it offers *toll bypass* cost savings.

We then considered various analog and digital interfaces available for our routers and what could connect to those interfaces. Analog phones, for example, can connect to FXS interfaces. The telephone wall jack in our home can connect to an FXO interface. Also, PBXs in our company might connect to remote PBXs using E&M interfaces. We can leverage our company's existing investment in those interfaces by connecting the PBX's E&M interfaces to E&M interfaces on our routers.

Whereas FXS, FXO, and E&M are examples of analog interfaces, digital interfaces include T1, E1, and ISDN. Recall that a T1 circuit has 24 channels, and we can use all 24 channels to send voice traffic if we use CAS, which is sometimes called "robbed-bit" signaling. However, the CCS option uses one of the 24 channels just for signaling.

An E1 interface has 32 channels. However, we only use 30 of those channels for voice. The first channel is used for framing and synchronization. The seventeenth channel is used for signaling, and it's interesting that the seventeenth channel is used for signaling in both the CAS and CCS modes.

ISDN comes in two flavors, BRI and PRI. The BRI flavor (which includes two 64-kbps B-channels) might be appropriate for a SOHO environment, whereas PRI (which includes 23 64-kbps B-channels on a T1-based PRI interface) would be more appropriate for a larger environment.

This chapter included an overview of VoIP in the home. Several service providers have equipment in telephone COs scattered across the country. If we have a broadband (for example, DSL or cable modem) connection in our home along with a router that allows us to connect phones, our phones can connect to the service provider's equipment over the Internet (that is, over our broadband connection) and from there connect to the PSTN. VoIP in the home has the potential to offer comparable features to our existing phone service. However, because VoIP in the home is an emerging technology, subscribers should understand exactly how their service provider handles 911 calls.



Chapter Review Questions

1. Approximately how many minutes of downtime per year does a network experience if it has the “five nines” of availability?
 - a. 53 minutes
 - b. 46 minutes
 - c. 16 minutes
 - d. 5 minutes
2. Which of the following offer Layer 3 redundancy in a network? (Select the two best answers.)
 - a. 802.1w
 - b. VRRP
 - c. HSRP
 - d. RSTP
3. An analog phone can connect to which of the following router interfaces?
 - a. BRI
 - b. FXO
 - c. E&M
 - d. FXS
4. The wall jack in your home, where you plug in your analog telephone, can connect to which of the following router interfaces?
 - a. BRI
 - b. FXO
 - c. E&M
 - d. FXS
5. Identify two valid framing types for a T1 circuit.
 - a. AMI
 - b. SF
 - c. B8ZS
 - d. ESF

6. Select the type of line coding that replaces a byte containing eight zeros with two bipolar violations.
 - a. AMI
 - b. SF
 - c. B8ZS
 - d. ESF
7. How many channels does a T1 circuit have?
 - a. 16
 - b. 24
 - c. 30
 - d. 32
8. What channel on an ISDN circuit carries signaling information?
 - a. A
 - b. B
 - c. C
 - d. D
9. Which of the following parameters are configurable on an FXS interface? (Select the two best answers.)
 - a. Signal type
 - b. Ring number
 - c. Ring frequency
 - d. Dial type
10. An ISDN BRI circuit has how much usable bandwidth (that is, not including the D channel)?
 - a. 56 kbps
 - b. 64 kbps
 - c. 128 kbps
 - d. 256 kbps