CHAPTER 2

Audio Basics

This chapter covers the following topics:

Basic Understanding of Sound: This topic will introduce a basic understanding of how sound behaves, including wave propagation, technical properties of sound waves, and noise.

Analog vs. Digital Signals: This topic will compare noise on an analog signal compared to noise on a digital signal, observe the principle of the Nyquist-Shannon theorem, and identify how bandwidth conservation can be achieved through data compression.

ITU Audio Encoding Formats: This topic will identify the most commonly used audio codecs in a Cisco solution and analyze the various aspects of each codec.

The preceding chapter provided a backward look into the evolution of communications and a high-level overview of the technologies that exist today that allow us to collaborate at many levels. This chapter will begin to dive into the intricate details that encompass the vast world of audio communication. As you read through this chapter, you will find that the chasm that embodies audio communication is only a foot wide but a mile deep. Although this chapter will not cover every aspect of audio communication, it will provide a solid foundation to build your knowledge on. Topics discussed in this chapter include the following:

- Basic Understanding of Sound
  - Wave Propagation
  - Technical Properties of a Sound Wave
  - Understanding Noise
- Analog vs. Digital Signals
  - Nyquist-Shannon Sampling Theorem
  - Data Compression Equals Bandwidth Conservation
- ITU Audio Encoding Formats

This chapter covers the following objective from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

2.2 Identify the appropriate collaboration codecs for a given scenario

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 2-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 2-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Basic Understanding of Sound	1–4
Analog vs. Digital	5–8
ITU Audio Encoding Formats	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Through which of the following mediums will sound travel the fastest?
   - **a.** Air
   - **b.** Water
   - **c.** Concrete
   - **d.** Gypsum board
2. The measure of the magnitude of change in each oscillation of a sound wave is a description of which term?
   - **a.** Frequency
   - **b.** Amplitude
   - **c.** Wavelength
   - **d.** Pressure
3. Which of the following is used to calculate the average amplitude of a wave over time?
   - **a.** Newton's second law of motion
   - **b.** Sine waves
   - **c.** Pascals
   - **d.** Root mean squared
4. What is the primary issue with using repeaters for analog signals?
   - **a.** Noise riding on the original audio signal is also amplified by the repeaters.
   - **b.** Distances may be too great to host enough repeaters.
   - **c.** Noise from the repeaters can be interjected into the audio signals.
   - **d.** There are no issues with using a repeater for analog signals.

5. Which of the following is a benefit of digital audio over analog audio?
   a. Digital signals are more fluid than analog signals.
   b. Digital signals have a more natural sound than analog signals.
   c. Digital signals are immune to ambient noise.
   d. Digital signals are not based on timing like an analog signal is.
6. What is the maximum bit rate the human ear can hear at?
   a. 24 bit
   b. 21 bit
   c. 144 bit
   d. 124 bit
7. Which of the following supports 24 circuit channels at 64 Kbps each?
   a. DS0
   b. DS1
   c. T3
   d. E1
8. Which of the following describes lossless compression?
   a. This algorithm searches content for statistically redundant information that can be represented in a compressed state.
   b. This algorithm searches for nonessential content that can be deleted to conserve storage space.
   c. Compression formats suffer from generation loss.
   d. Repeatedly compressing and decompressing the file will cause it to progressively lose quality.
9. Which of the following codecs supports both compressed and uncompressed data?
   a. G.711
   b. G.729
   c. G.722
   d. G.723.1
10. Which of the following codecs supports lossless compression?
    a. G.729
    b. G.722
    c. iLBC
    d. iSAC

Foundation Topics

Basic Understanding of Sound

Sound is the traveling of vibrations through a medium, which in most cases is air, generated by a source as waves of changing pressure. These waves of changing pressure in the air, in turn, vibrate our eardrums, allowing our brains to perceive them as sound. How often these pulses of pressure change in a given time period can be measured. We can also measure how

intense each pressure change is and how long the space is between each wave. Two factors that relate to sound must be considered: the source and the medium.

Anything that can create these vibrations can be the source of sound. These sources could include human vocal cords or animal larynx, such as birds chirping, lions roaring, or dogs barking. Sound could also be sourced from stringed instruments, such as a guitar, violin, or cello. Percussions cause sound, such as drum skins or even the diaphragm of a loudspeaker. All these things in some way vibrate particles to make sound. A tree falling in the forest will make sound, whether you hear it or not.

The medium sound travels through can affect how that sound is perceived. Have you ever dipped your head underwater in a swimming pool and tried to talk with someone else? Although you can hear the other person, understanding what that person is saying becomes difficult, even when you are both in close proximity. It's important to understand that the speed of sound is constant within a given medium at a given temperature. Generally, the term *speed of sound* refers to a clear, dry day at about 20° C (68° F), which would have the speed of sound at 343.2 meters per second (1,126 ft/s). This equates to 1,236 kilometers per hour (768 mph)—about 1 kilometer in 3 seconds or approximately 1 mile in 5 seconds. The denser the medium, the faster sound will travel. Because air is thicker at lower elevations, and temperatures are usually higher, sound will travel faster but will degrade faster as well. However, in a vacuum, sound will travel much slower because there are no particles to propagate the sound waves, but sound can be heard much clearer at farther distances because they will degrade much slower. Table 2-2 illustrates the speed of sound through four different mediums.

Key Topic

Table 2-2 Speed of Sound Through Four Common Mediums

Medium	Speed (Meters per Second)	Speed (Feet per Second)	Speed Factor
Air	344	1130	1
Water	1480	4854	4.3
Concrete	3400	11,152	9.8
Gypsum Board	6600	22,305	19.6

Wave Propagation

Stand next to a still body of water, pick up a pebble, and toss it to the middle of the pond. As the pebble breaks the surface, it displaces the water it comes in immediate contact with. Those water particles displace other water particles as they are pushed out, so you get a rippling effect that grows from the point where the pebble originally made contact out to the edges of the pond. Once the ripples meet the outer edges of the pond, the rings will begin to move back in, and the ripples will continue until the energy created by tossing the pebble into the water runs out.

Key Topic

This principle is the same one on which sound propagates. When sound is created at the source, the vibrations radiate out in waves, like the ripples in a pond, except they travel in every direction. These sound waves bounce off particles of different mediums in order to travel outward. As the sound waves bounce off particles, some of the energy is absorbed so the vibration weakens. The density of a medium plays a part in how fast a sound wave can propagate or move through it. The more tightly that particles of matter are packed together,

the faster a sound wave will move through it, but the sound wave will degrade faster as well. So sound travels faster through a solid than a liquid and faster through a liquid than a gas, such as air. As the sound expands across a larger area, the power that was present at the source to create the original compression is dissipated. Thus, the sound becomes less intense as it progresses farther from the source. The sound wave will eventually stop when all the energy present in the wave has been used up moving particles along its path.

Observe the analogy of the pebble in the pond. When the ripples first begin, they will be taller and closer together. As the energy dissipates, the ripples will become shorter and farther apart, until eventually they cease to demonstrate any effect on the water's surface. This effect can be observed with sound as well. Stand in a canyon and shout out toward the cliffs. The sound you create will travel out and reverberate off the cliffs and return to you. However, what you hear will sound degraded. You may even hear yourself multiple times in an echo, but each time the sound will be weaker than the time before until you can't hear anything at all. If someone shouted out to you under a body of water, by the time the sound reached your ears, it would have degraded so much from bouncing off all the many particles in the water that you would not be able to articulate what was said.

Technical Properties of Sound

All forms of energy can be measured. Light and sound sources in nature are forms of energy. Acoustical energy consists of fluctuating waves of pressure, generally in the air. One complete cycle of that wave consists of a high- and low-pressure half-cycle. This means that half of a sound wave is made up of the compression of the medium, and the other half is the decompression, or rarefaction, of the medium. Imagine compressing a spring and then letting it go. Now imagine that the spring represents air molecules and your hand is the acoustic energy. Figure 2-1 illustrates the technical properties of a sound wave.

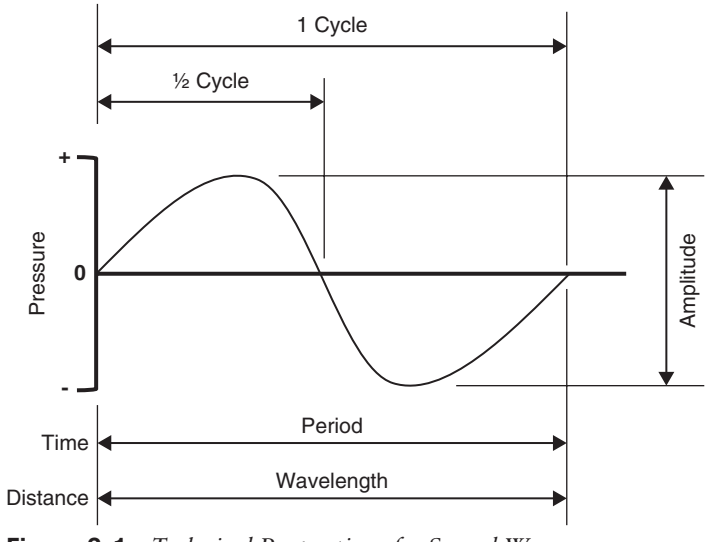

Figure 2-1 *Technical Properties of a Sound Wave*

To better understand how sound waves are measured, we must define some terms. *Frequency* is the rate of the air pressure fluctuation produced by the acoustic energy wave. To be heard by human ears, wave frequencies must be between 20 and 20,000 cycles per

second. Although there is more to it than just frequency, in general, the frequency of a sound wave corresponds with the perceived pitch of the sound. The higher the pitch, the higher the frequency. When we are talking about sounds and music, there is not just one single frequency, but many different frequencies overlapping with each other. This means many different frequencies are being represented within a single signal. The bandwidth of a particular signal is the difference between the highest and the lowest frequencies within the signal.

Amplitude is the measure of the magnitude of change in each oscillation of a sound wave. Most often this measurement is peak-to-peak, such as the change between the highest peak amplitude value and the lowest trough amplitude value, which can be negative. Bear in mind the relationship between amplitude and frequency. Two waveforms can have identical frequency with differing amplitude and identical amplitude with differing frequency. For sound waves, the wave oscillations are representative of air molecules moving due to atmospheric pressure changes, so the amplitude of a sound wave is a measure of sound pressure. Sound pressure is measured in pascals or millibars:

1 pascal = 1 newton per square meter

1 millibar = 100 pascals

A newton is the standard international unit for force. It is equal to the amount of net force required to accelerate a mass of one kilogram at a rate of one meter per second squared. Newton's second law of motion states F = ma, multiplying m (kg) by a (m/s^2); the newton is, therefore, N = $kg(m/s^2)$.

A frequency spectrum is the complete range of frequencies that can be heard, just as the visible spectrum is the range of light frequencies that can be seen. Devices and equipment can also have frequency spectrums, or ranges.

Basic sine waves cannot have an "average" per se. Rather, it would equal zero because the wave peaks and valleys are symmetrical above and below the reference of zero. What is much more helpful when discussing the average amplitude of a wave over time is called root mean squared (RMS). RMS is often used to calculate a comparable measure of power efficiency, such as in audio amplifiers. RMS measures mean output power to mean input power. RMS is just the squaring of every point of a wave and then finding the average. Squaring a negative value always results in a positive value. RMS gives us a useful average value for discussing audio equipment and comparing amplitude measurements. Basically, RMS is much more similar to the way we hear sound, as opposed to measuring just the peaks of a wave, because we don't hear just the peaks.

A sound wave emanating from a source is like an expanding bubble. The power of the sound would be the energy on the surface of the bubble. As the surface of the bubble expands, that energy is used up in order to move, or vibrate, the air ever farther outward. That means eventually the power that the bubble started with at its source would be expended. The power of sound, or sound's acoustical power, is a measure of amplitude over time. The sound had a particular measurement at the moment it was emitted from the source, but as it travels, over time the power decreases as the energy present in the sound wave is expended transmitting itself through the medium. The transmission of acoustical energy through a medium is measured in watts. A watt is just the rate of transfer of energy in one second or one joule per second. A joule is equal to the energy expended in applying a force of one newton through a distance of one meter. The following equation is used to calculate sound's acoustical power:

$J=kg*m^2/s^2=N*4444m=Pa*m^3=W*s$

where

- N is the newton.
- m is the meter.
- kg is the kilogram.
- s is the second.
- Pa is the pascal.
- W is the watt.

Understanding Attenuation and Noise

When Alexander Graham Bell placed his first audio call to his assistant, Watson, it was only to the next room. Bell's journal records that the audio was loud but muffled. There are two reasons why the sound was hard to hear: attenuation and noise.

Attenuation, as has previously been discussed, is the degrading of the sound wave as it loses energy over time and space. When the telephone was first made available to the public market, telephone companies had to set up repeaters to account for this loss in signal strength. These repeaters could boost the intensity of the analog wave at different points along the path between two telephones, or nodes. This allowed telephone calls to traverse much longer distances. However, telephone repeaters solved only half the problem.

In audio terms, noise doesn't necessarily refer to something you hear. *Noise* can also refer to inaccuracies or imperfections in a signal transmission that was not intended to be present. Noise comes in different forms, like acoustical, digital, or electrical. Noise can be introduced to a signal in many ways, like faulty connections in wiring or external signal interference. In telephony, as audio signals are transmitted across a wire using electrical signals, those same wires can pick up energy from external sources and transmit them along the same current carrying the audio signal. If you are old enough to remember using analog phones, you might also remember hearing a crackling sound in the background during a phone call. That background noise was caused by interference.

When an analog signal must be transmitted a long distance, the signal level is amplified to strengthen it; however, this process also amplifies any noise present in the signal. This amplification could cause the original analog signal to become too distorted to be heard properly at the destination. Therefore, it is very important with analog transmissions to make sure the original audio signal is much stronger than the noise. Electronic filters can also be introduced, which helps remove unwanted frequency from the analog signal. The frequency response may be tailored to eliminate unwanted frequency components from an input signal or to limit an amplifier to signals within a particular band of frequencies. Figure 2-2 illustrates how noise riding on an analog signal can be amplified at repeaters.

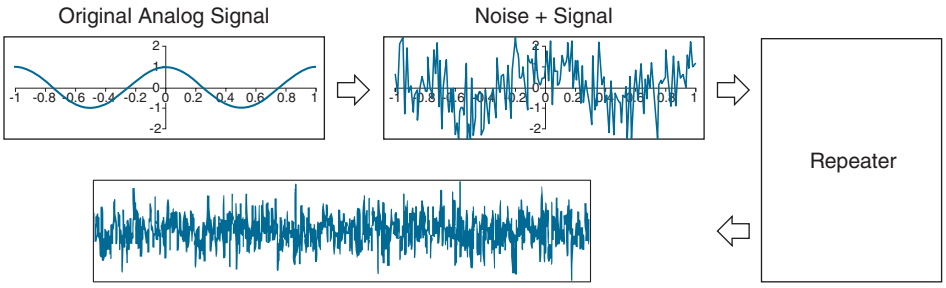

Figure 2-2 *Amplification of Noise on an Analog Signal*

Analog vs. Digital Signals

Key Topic

An analog signal is a continuous signal that contains a time variable representative of some other time-varying quality, such as the voltage of the signal may vary with the pressure of the sound waves. In contrast, a digital signal is a continuous quantity that is a representation of a sequence of discrete values which can take on only one of a finite number of values. There are only two states to a digital bit, zero or one, which can also be perceived as on or off. These are referred to as *binary digits*. Another way to compare analog signals with digital signals could be to think of them as light switches. A digital signal operates like a regular light switch; it is either on or off. An analog signal is more like a dimmer switch that is fluid and allows varying levels of light.

Another good way to mentally picture the difference between analog and digital is to think of two types of clocks. An analog clock has hands that point at numbers, by slowly rotating around in a circle, and digital clocks have decimal number displays. The analog clock has no physical limit to how finely it can display the time, as its hands move in a smooth, pauseless fashion. If you look at the clock at just the right moments, you can see the hand pointing between two minutes or even two seconds, such as it could read 3:05 and a half. The digital clock cannot display anything between 3:05 and 3:06. It is either one or the other; there is no variance between the two times.

When it comes to sound quality, there is much debate as to whether analog or digital sounds better. Vinyl records, or LP (long play) records, are an example of raw analog audio that has been recorded. CDs or MP3s are examples of digital audio that has been recorded. The science behind audio quality definitively determines that analog has a better sound quality than digital because analog is the most natural and raw form of sound waves, and analog is the only form of audio sound waves our ears can articulate. Digital recordings take an analog signal and convert it into a digital format. When that digital recording is played back, it must be converted back to analog before the audio is played over speakers. Because digital conversion cannot make a perfect copy of the original fluid audio sound wave, some of that sound quality is lost when the digital signal is converted back into analog.

This issue of conversion raises a question: if analog audio is better than digital audio, why use digital at all? In conjunction with the preceding examples, digital copies allow a higher quantity of audio to be stored. If you go to a music store or download music from the Internet, you can store more songs in a smaller container, such as an MP3 player or an iPod. Bringing the subject back to communications, an analog signal cannot travel very far without suffering from attenuation. Digital signals can travel much farther than analog signals and

don't degrade in the same manner. Digital repeaters will amplify the signal, but they may also retime, resynchronize, and reshape the pulses. Additionally, analog signals are very sensitive to ambient noise, which can quickly degrade the sound quality. Digital signals use binary digits of zero or one, so they are immune to ambient noise.

It is much easier to identify and remove unwanted noise from a digitally sampled signal compared to an analog signal. The intended state of the signal is much easier to recognize. When discussing potential errors in a transmitted signal, error checking and correction are made much easier by the very nature of digital representation. As a digital signal enters a digital repeater, an algorithm will first check to see if the information that is supposed to be there still exists or if it's missing. Assuming the correct information is there, next it will check the state of each digital bit in question: on or off (a one or zero). Finally, it will check that the position of each bit is correct: is the bit on when it's supposed to be off or vice versa?

When an analog signal must be transmitted a long distance, the signal level is amplified to strengthen it; however, this also amplifies any noise present in the signal. Digital signals are not amplified. Instead, for long-distance transmission, the signal is regenerated at specific intervals. In this manner, noise is not passed along the transmission chain.

Nyquist-Shannon Sampling Theorem

Always keep in mind that digital signals are just long strings of binary numbers. Analog signals must be somehow converted into strings of numbers to be transmitted in the digital realm. When a continuous analog signal is converted to a digital signal, measurements of the analog signal must be taken at precise points. These measurements are referred to as *samples*. The more frequent the sample intervals, the more accurate the digital representation of the original analog signal. The act of sampling an analog signal for the purpose of reducing it to a smaller set of manageable digital values is referred to as *quantizing* or *quantization*. The difference between the resulting digital representation of the original analog signal and the actual value of the original analog signal is referred to as *quantization error*. Figure 2-3 illustrates how quantizing an analog signal would appear.

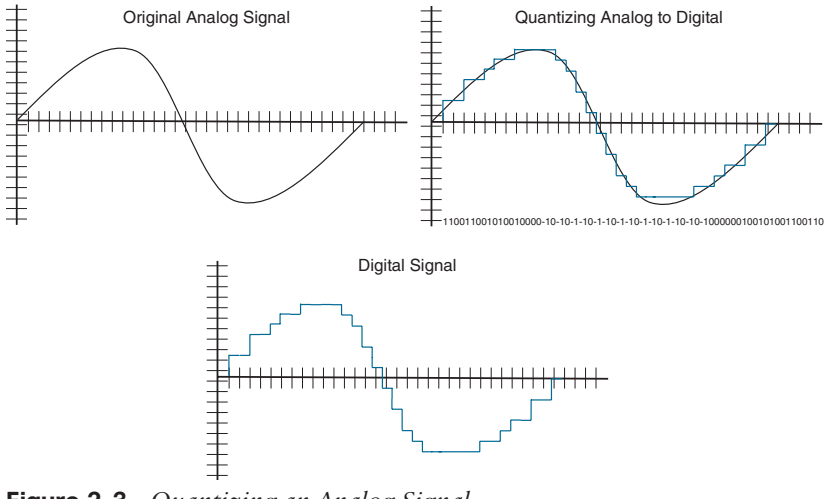

Figure 2-3 *Quantizing an Analog Signal*

In Figure 2-3, each movement upward on the Y-axis is signified by a one-binary digit. Each movement horizontally on the X-axis is signified by a zero-binary digit. Each movement downward on the Y-axis is signified by a negative one-binary digit. In this manner an analog signal can be quantized into a digital signal. However, notice that the binary signal is not exactly the same as the original analog signal. Therefore, when the digital signal is converted back to analog, the sound quality will not be the same. A closer match can be obtained with more samples. The size of each sample is measured in bits per sample and is generally referred to as *bit depth*. However, it will never be exact because a digital signal cannot ever match the fluid flow of an analog signal. The question becomes, can a digital signal come close enough to an analog signal that it is indistinguishable to the human ear?

When an audio signal has two channels, one for left speakers and one for right speakers, this is referred to as *stereo*. When an audio signal has only one channel used for both left and right speakers, this is referred to as *mono*. Using the preceding quantizing information, along with whether an audio signal is stereo or mono, it is relatively simple to calculate the bit rate of any audio source. The bit rate describes the amount of data, or bits, transmitted per second. A standard audio CD is said to have a data rate of 44.1 kHz/16, meaning that the audio data was sampled 44,100 times per second, with a bit depth of 16. CD tracks are usually stereo, using a left and right track, so the amount of audio data per second is double that of mono, where only a single track is used. The bit rate is then 44,100 samples/second × 16 bits/sample × 2 tracks = 1,411,200 bps or 1.4 Mbps.

When the sampling bit depth is increased, quantization noise is reduced so that the signal-to-noise ratio, or SNR, is improved. The relationship between bit depth and SNR is for each 1-bit increase in bit depth, the SNR will increase by 6 dB. Thus, 24-bit digital audio has a theoretical maximum SNR of 144 dB, compared to 96 dB for 16-bit; however, as of 2007 digital audio converter technology is limited to an SNR of about 126 dB (21-bit) because of real-world limitations in integrated circuit design. Still, this approximately matches the performance of the human ear. After about 132 dB (22-bit), you would have exceeded the capabilities of human hearing.

Key Topic

To determine the proper sampling rate that must occur for analog-to-digital conversion to match the audio patterns closely, Harry Nyquist and Claude Shannon came up with a theorem that would accurately calculate what the sampling rate should be. Many audio codecs used today follow the Nyquist-Shannon theorem. Any analog signal consists of components at various frequencies. The simplest case is the sine wave, in which all the signal energy is concentrated at one frequency. In practice, analog signals usually have complex waveforms, with components at many frequencies. The highest frequency component in an analog signal determines the bandwidth of that signal. The higher the frequency, the greater the bandwidth, if all other factors are held constant. The Nyquist-Shannon theorem states that to allow an analog signal to be completely represented in digital form, the sampling rate must be at least twice the maximum cycles per second, based on frequency or bandwidth, of the original signal. This maximum bandwidth is called the *Nyquist frequency*. If the sampling rate is not at least twice the maximum cycles per second, then when such a digital signal is converted back to analog form by a digital-to-analog converter, false frequency components appear that were not in the original analog signal. This undesirable condition is a form of distortion called *aliasing*. Sampling an analog input signal at a rate much higher than the minimum frequency required by the Nyquist-Shannon theorem is called *over-sampling*. Over-sampling improves the quality or the digital representation of the original analog input signal. Under-sampling occurs when the sampling rate is lower than the analog input frequency.

The Nyquist-Shannon theorem led to the Digital Signal 0 (DS0) rate. DS0 was introduced to carry a single digitized voice call. For a typical phone call, the audio sound is digitized at an 8 kHz sample rate using 8-bit pulse-code modulation for each of the 8000 samples per second. This resulted in a data rate of 64 kbps.

Because of its fundamental role in carrying a single phone call, the DS0 rate forms the basis for the digital multiplex transmission hierarchy in telecommunications systems used in North America. To limit the number of wires required between two destinations that need to host multiple calls simultaneously, a system was built in which multiple DS0s are multiplexed together on higher-capacity circuits. In this system, 24 DS0s are multiplexed into a DS1 signal. Twenty-eight DS1s are multiplexed into a DS3. When carried over copper wire, this is the well-known T-carrier system, with T1 and T3 corresponding to DS1 and DS3, respectively.

Outside of North America, other ISDN carriers use a similar Primary Rate Interface (PRI). Japan uses a J1, which essentially uses the same 24 channels as a T1. Most of the world uses an E1, which uses 32 channels at 64 kbps each. ISDN will be covered in more depth in Chapter 5, "Communication Protocols." However, it is important to note that the same sampling rate used with DS0 is also used in VoIP communications across the Internet. This led to a new issue that in turn opened up a new wave of advancement in the audio communication industry: How could digital audio signals be sent over low-bandwidth networks?

Data Compression Equals Bandwidth Conversion

The answer to the question regarding digital audio signals being sent over low-bandwidth networks came with the development of data compression. Bandwidth is the rate that data bits can successfully travel across a communication path. Transmitting audio can consume a high amount of bandwidth from a finite total available. The higher the sampled frequency of a signal, the larger the amount of data to transmit, resulting in more bandwidth being required to transmit that signal. Data compression can reduce the amount of bandwidth consumed from the total transmission capacity. Compression involves utilizing encoding algorithms to reduce the size of digital data. Compressed data must be decompressed to be used. This extra processing imposes computational or other costs into the transmission hardware.

Lossless and *lossy* are descriptive terms used to describe whether or not the data in question can be recovered exactly bit-for-bit when the file is uncompressed or whether the data will be reduced by permanently altering or eliminating certain bits, especially redundant bits. Lossless compression searches content for statistically redundant information that can be represented in a compressed state without losing any original information. By contrast, lossy compression searches for nonessential content that can be deleted to conserve storage space. A good example of lossy would be a scenic photo of a house with a blue sky overhead. The actual color of the sky isn't just one color of blue; there are variances throughout. Lossy compression could replace all the blue variances of the sky with one color of blue, thus reducing the amount of color information needed to replicate the blue sky after decompression. This is, of course, an unsubtle oversimplification of the process, but the concept is the same.

It is worth noting that lossy compression formats suffer from generation loss; repeatedly compressing and decompressing the file will cause it to progressively lose quality. This is in contrast with lossless data compression, where data will not be lost even if compression is

repeated numerous times. Lossless compression therefore has a lower limit. A certain amount of data must be maintained for proper replication after decompression. Lossy compression assumes that there is a trade-off between quality and the size of the data after compression. The amount of compression is limited only by the perceptible loss of quality that is deemed acceptable.

ITU Audio Encoding Formats

Codec stands for coding and decoding. A codec is a device or software program that is designed to code and decode digital data, as well as compress and decompress that data. Therefore, an audio codec is a device or software that is designed to process incoming analog audio, convert it to digital, and compress the data, if necessary, before sending that data to a specific destination. The codec can also process incoming data, decompress that data if necessary, and convert the digital data back to analog audio. Many proprietary codecs have been created over the years. However, in an effort to unify communications, the International Telecommunications Union (ITU) created a standardized set of audio codecs. The ITU is a specific agency of the United Nations whose chief responsibilities include the coordination of the shared global use of the usable RF spectrum, and the establishment of standards to which manufacturers and software designers comply in order to ensure compatibility. Table 2-3 outlines some of the more common audio codecs used in a Cisco collaboration solution. Each codec that starts with a "G" is an ITU codec.

Key Topic

Table 2-3 Audio Codecs Commonly Used by Cisco

Codec and Bit Rate (Kbps)	Codec Sample Size (Bytes)	Codec Sample Interval (ms)	Mean Opinion Score (MOS)	Voice Payload Size (Bytes)	Bandwidth MP or FRF.12 (Kbps)	Bandwidth Ethernet (Kbps)
G.711 (64 Kbps)	80	10	4.1	160	82.8	87.2
G.729 (8 Kbps)	10	10	3.92	20	26.8	31.2
G.723.1 (6.3 Kbps)	24	30	3.9	24	18.9	21.9
G.723.1 (5.3 Kbps)	20	30	3.8	20	17.9	20.8
G.726 (32 Kbps)	20	5	3.85	80	50.8	55.2
G.726 (24 Kbps)	15	5		60	42.8	47.2
G.728 (16 Kbps)	10	5	3.61	60	28.5	31.5
G.722_64k (64 Kbps)	80	10	4.13	160	82.8	87.2
iLBC_Mode_20 (15.2 Kbps)	38	20	4.14	38	34.0	38.4
iLBC_Mode_30 (13.33 Kbps)	50	30	4.14	50	25.867	28.8

Table 2-3 is not an exhaustive list of codecs available today, and it is important to note that ITU codecs are used with SIP communication as well. Table 2-3 is divided into seven columns. The first column identifies the codec and the number of bits per second needed to transmit the payload in each packet for a voice call. The codec bit rate = codec sample size / codec sample interval. The next column is the codec sampling size based on bytes. This is the number of bytes captured by the codec at each codec sample interval. Column three is the codec sampling interval. This is the sample interval at which the codec operates. For example, the G.729 codec operates on sample intervals of 10 ms, corresponding to 10 bytes (80 bits) per sample at a bit rate of 8 kbps. The next column is the Mean Opinion Score (MOS), which is a system of grading the voice quality of telephone connections. With MOS, a wide range of listeners judge the quality of a voice sample on a scale of one (bad) to five (excellent). The scores are averaged to provide the MOS for the codec. The last two columns show the bandwidth needed to transmit the audio with overhead added in. The bandwidth MP or FRF.12 shows the Layer 2 header values added to the original payload, and the bandwidth ethernet shows the Layer 3 header values added on top of the Layer 2 headers.

Based on the Nyquist-Shannon theorem, the first audio codec created by the ITU is G.711. Originally released in 1972, G.711 is also referred to as pulse-code modulation (PCM) and was introduced for use in telephony. It is the required minimum standard in both H.320 for circuit-switched telephony and H.323 for packet-switched telephony. It is considered a *narrow-band* codec since it processes only frequencies between 300 and 3400 Hz, although an annex has been added to extend the frequency range. It uses a sampling frequency of 8 kHz and has a bit rate of 64 kbps. Two algorithm types are associated with G.711: G.711 mu-law and G.711 a-law. G.711 mu-law is a *companding algorithm*, which reduces the dynamic range of an audio signal. G.711 mu-law is used throughout North America and Japan. The algorithm used with G.711 a-law is common throughout the rest of the world where E1 circuits are used, and it is the inverse of the mu-law algorithm. All G.711 codecs are uncompressed.

G.729 is a lossy compressed audio codec and is the most common compression algorithm used in low-bandwidth environments. It has been used in videoconferencing applications for some time and attained its ITU recommendation in 1996. It can operate at a low rate of 4000 Hz. There are several annexes to the G.729 codec, the ones most commonly used with Cisco being G.729r8 and G.729br8. Both codecs operate at 8 Kbps, but G.729br8 contains built-in VAD that cannot be disabled.

Now compare the two previously mentioned codecs with G.722. Released in 1988, this codec can operate as a lossy compressed codec or an uncompressed codec, depending on the annex being run. G.722 addressed some of the speech quality issues presented by the limited bandwidth of the G.711 codec. In contrast to G.711, G.722 is considered a *wide-band* codec and processes frequencies between 50 and 7000 Hz. It also samples audio at 16 kHz and operates at 48, 56, or 64 kbps. When G.722 audio is used over an H.323 call, the packets are framed using the 802.3 standard and are sent at 60-millisecond intervals.

Two other commonly used codecs with Cisco collaboration solutions are iLBC and iSAC. These are not ITU codecs, but they are open standards that can be used by any organization. Internet Low Bitrate Codec (iLBC) was originally drafted for the WebRTC project. It was adopted by the IETF in 2002 and ratified into SIP communications in 2003. Internet Speech

Audio Codec (iSAC) was originally created by Global IP Solutions. iSAC was acquired by Google in 2011, incorporated into the open-source WebRTC project, and later ratified into the SIP protocol by the IETF. The iSAC codec is an adaptive wideband speech and audio codec that operates with short delay, making it suitable for high-quality real-time communication. It is specially designed to deliver wideband speech quality in both low and medium bit-rate applications. The iSAC codec compresses speech frames of 16 kHz, 16-bit sampled input speech, each frame containing 30 or 60 ms of speech. The codec runs in one of two different modes called channel-adaptive mode and channel-independent mode. In both modes iSAC is aiming at a target bit rate, which is neither the average nor the maximum bit rate that will be reached by iSAC, but it corresponds to the average bit rate during peaks in speech activity.

In channel-adaptive mode, the target bit rate is adapted to give a bit rate corresponding to the available bandwidth on the channel. The available bandwidth is constantly estimated at the receiving iSAC and signaled in-band in the iSAC bit stream. Even at dial-up modem data rates, iSAC delivers high quality by automatically adjusting transmission rates to give the best possible listening experience over the available bandwidth. The default initial target bit rate is 20,000 bits per second in channel-adaptive mode.

In channel-independent mode, a target bit rate has to be provided to iSAC prior to encoding. After encoding the speech signal, the iSAC codec uses lossless coding to further reduce the size of each packet and hence the total bit rate used. The adaptation and the lossless coding described here both result in a variation of packet size, depending both on the nature of speech and the available bandwidth. Therefore, the iSAC codec operates at transmission rates from about 10 kbps to about 32 kbps.

The best quality audio codec available to date is Advanced Audio Codec-Low Delay (AAC-LD). It is also referred to as Low-overhead MPEG-4 Audio Transport Multiplex (LATM), and this codec is most commonly used with SIP communication. Since 1997 this codec has been used to offer premium stereo audio by over-sampling analog audio signaling. Sampling rates range between 48 Kbps and 128 Kbps, with a frequency range of 20 kHz. Although this codec does offer excellent audio quality during calls, the trade-off is the high cost in bandwidth when this codec is used.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, “Final Preparation,” and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 2-4 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 2-4 Key Topics for Chapter 2

Key Topic Element	Description	Page Number
Table 2-2	Speed of Sound Through Four Common Mediums	19
Paragraph	Behavior of Sound Waves as They Propagate	19
Section	Technical Properties of Sound	20
Paragraph	Define Attenuation and Noise	22
Paragraph	Define Analog and Digital Audio Signals	23
Paragraph	Samples and Quantization Explained	24
Paragraph	Nyquist-Shannon Theorem Explained	25
Paragraph	Lossless and Lossy Compression Explained	26
Table 2-3	Audio Codecs Commonly Used by Cisco	27
Paragraph	G.711, G.729, and G.722 Explained	28

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Acoustical Power, Aliasing, Amplitude, Analog Signal, Bandwidth, Bit Depth, Data Compression, Digital Signal, DS0, DS1, DS3, Electronic Filters, Frequency, Frequency Spectrum, Lossless, Lossy, Millibar, Mono, Newton, Over-sampling, Pascal, PCM, Quantization, Quantization Error, RMS, Samples, Sine Wave, Sound Pressure, Stereo, Undersampling, Watt, Wavelength

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. Name the technical properties of a sound wave.
2. What are the three digital signal rate forms, and how do they translate into T-carrier rate forms?
3. List the top three ITU codecs used for audio and the three SIP audio codecs mentioned in this book.

This page intentionally left blank

CHAPTER 3

Video Basics

This chapter covers the following topics:

Basic Understanding of Light: Th`is topic will provide a basic understanding of light behavior, break down aspects of light into chrominance and luminance, and discuss light temperature.

Capturing and Cameras: This topic will discuss frame rates, explain what pixels are and how they impact resolution, and overview common encoding techniques for video.

Standard Video Codecs: This topic will examine video compression and discuss varying video quality based on video codecs.

Video Container Formats and Codecs: This topic will compare and contrast the H.264 video codec with the latest H.265 HEVC video codec. This topic will also discuss how content can be shared in a video stream during video communication.

As mentioned in Chapter 1, "Introduction to Collaboration," video communication is a relatively young industry. As such, most of the development within this industry has occurred within the last 30 years. Microsoft released one of the first consumer products that allowed video communication from a desktop application with the release of Windows 95, but it didn't resonate with people for several reasons. Networks couldn't support the bandwidth required, so quality was an issue, and consumers couldn't envision the need for video communication. Most of the early use cases for video communication occurred within the business sector, which in turn influenced the standards that soon followed. In today's market several consumer video communication products are available, such as Skype and Apple's FaceTime. However, it is still the private and public sectors that continue to drive the industry forward. To fully understand how video communication works, we need a more in-depth examination of light behavior. Topics discussed in this chapter include the following:

- Basic Understanding of Light
  - Chrominance and Luminance
  - Color Temperature
- Capturing and Cameras
  - Frame Rates
  - Understanding Pixels and Resolution
  - Common Encoding Techniques
- Standard Video Codecs
  - Video Compression
  - Video Quality

- Video Container Formats and Codecs
  - H.264 Compared to H.265 HEVC
  - Content Channels

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

2.2 Identify the appropriate collaboration codecs for a given scenario

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 3-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 3-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Basic Understanding of Light	1–4
Capturing and Cameras	5–9
Standard Video Codecs	10–11
Video Container Formats and Codecs	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What are the three primary colors?
   - **a.** Red, white, blue
   - **b.** White, black, red
   - **c.** Red, green, blue
   - **d.** Red, yellow, blue
2. Which of the following terms is used to identify colors in light?
   - **a.** Chrominance
   - **b.** Luminance
   - **c.** Composite
   - **d.** Component

3. In the formula G = (Y – 0.299R – 0.114B) / 0.587, what does Y stand for?
   a. Luminance
   b. Chrominance
   c. Composite
   d. Component
4. Which of the following colors has the warmest temperature?
   a. Blue
   b. Red
   c. Green
   d. Yellow
5. Which of the following light filtration techniques lets about one-third of each color in and maps the saturation levels of each color per pixel in a mosaic effect?
   a. Foveon X3
   b. 3CCD
   c. Bayer
   d. 3 Chip
6. What is the lowest frame rate that still allows the human brain to perceive motion?
   a. 15 fps
   b. 24 fps
   c. 25 fps
   d. 30 fps
7. For what reason do NTSC and PAL use different frame rates?
   a. Different standards warrant different frame rates.
   b. Different resolutions warrant different frame rates.
   c. NTSC uses radio waves for broadcasting and PAL uses microwaves.
   d. The hertz rate on different power grids warrants different frame rates.
8. When given the resolution 1280 × 720p30, what does the 1280 represent?
   a. Pixels
   b. Lines
   c. Frame rate
   d. Scanning
9. When given the resolution 1280 × 720p30, what does the **p** represent?
   a. Pixels
   b. Progressive
   c. Picture
   d. Photons

10. What is the standard segmentation for macroblocks?
    a. 2×2
    b. 4×4
    c. 8×8
    d. 16×16
11. Which of the following video codecs was based off the MPEG-4 codec?
    a. H.261
    b. H.263
    c. H.264
    d. H.265
12. What is the maximum bit rate reduction that can be expected using the H.265 HEVC codec over H.264?
    a. 25%
    b. 50%
    c. 68%
    d. 75%

Foundation Topics

Basic Understanding of Light

There has been much debate among physicists over the last 400 years as to whether light is a particle or a wave, and there are conclusive studies that definitively prove both theories. Light was originally believed to be a particle based on a study by Sir Isaac Newton; however, several contemporary physicists concluded that light is a wave, and so the debate began. In the mid-1800s, a Scottish physicist named James Clerk Maxwell, who was studying electromagnetic waves, proved beyond anyone's doubts that light is not just a wave, but an electromagnetic wave. Much went into the explanation of Maxwell's theory, but in the end, Maxwell showed that the speed of an electromagnetic wave is the same as the speed of light. Maxwell ended the debate, proving light is a wave until it was examined by the famous physicist Albert Einstein as he studied quantum mechanics. Einstein's theory is that light is a particle, specifically a photon, and that the flow of a photon is a wave. Based on Einstein's theory, the energy of light is in direct correlation to its oscillation frequency, but the intensity of light correlates with the number of photons. Thanks to Einstein's theory, we now understand that light is both a particle and a wave.

Building on Maxwell's discovery that electromagnetic waves are light waves, you can begin to understand that spectrums of light extend beyond what you see. Just a small portion of the electromagnetic radiation (EMR) spectrum is what we refer to as *visible light*. Some EMR is visible to the human eye and is a form of energy emitted and absorbed by charged particles, called *photons*. EMR has both electric and magnetic field components and exhibits wave-like behavior as it travels through space. Figure 3-1 illustrates where visible light exists compared to other spectrums of light.

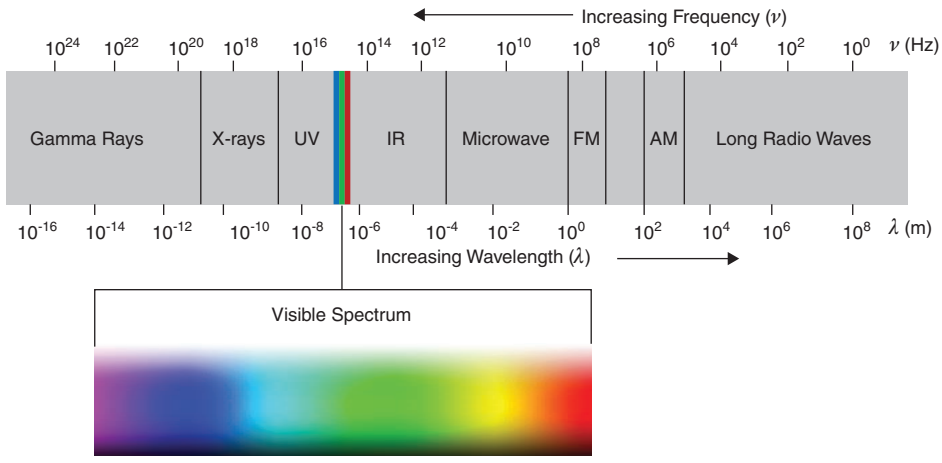

Figure 3-1 *Light Spectrums*

In common with all types of EMR, visible light is emitted and absorbed in tiny packets called photons and exhibits properties of both waves and particles. For example, a lens might refract a single photon or exhibit wave interference with itself but also act as a particle, giving a definite result when its position is measured. In a vacuum, EMR propagates at a characteristic speed, the speed of light. The speed of light is constant in a vacuum at exactly 299,792,458 meters per second. Note that sound waves require a medium and therefore cannot propagate in a vacuum.

Key Topic

The behavior of EMR depends on its wavelength. Higher frequencies have shorter wavelengths, and lower frequencies have longer wavelengths. When EMR interacts with single atoms and molecules, its behavior depends on the amount of energy per photon it carries. Photons are at the lower end of the energies that are capable of causing electronic excitation within molecules. Excitation would lead to changes in the bonding or chemistry of said molecule. At the lower end of the visible light spectrum, such as infrared, EMR becomes invisible to humans because its photons no longer have enough individual energy to cause a lasting molecular change in the molecules within a human retina. It is these photons exciting molecules with the human eye that trigger molecular changes, which cause the sensation of vision. Above the range of visible light, ultraviolet light becomes invisible to humans mostly because it is absorbed by the tissues of the eye—in particular, the lens. EMR is categorized by the frequency of its wave. The electromagnetic spectrum, in order of increasing frequency and decreasing wavelength (refer back to Figure 3-1), consists of radio waves, microwaves, infrared radiation, visible light, ultraviolet radiation, X-rays, and gamma rays.

Chrominance and Luminance

When we see a color in an object, such as a red rose, the rose is not actually red. The surface of those petals absorbs all colors in the light spectrum except red, which are reflected back out. So, when those red light waves enter our eye, we perceive the rose as being red. If an object were to reflect all colors in a light spectrum, then we would perceive that color as white. If an object were to absorb all light in a spectrum, then we would perceive that color as black. The visible color spectrum can be divided into three primary colors—red, blue, and green. Mixing various aspects of these three primary colors is what creates the whole spectrum of color humans can see. Now, this may be confusing to some of you who were taught

since grade-school that the three primary colors are red, blue and yellow. So, let me explain why green is referred to as the third primary color instead of yellow.

Humans see color not just by using the eye, but also through processing the light wavelengths in the brain. The previous section identified differing spectrums of light, which are based on frequency and wavelength. Frequency and wavelength also determine the different colors of light we can see within this visible spectrum. Red light waves have the lowest frequency but the longest wavelength. Blue light waves have the highest frequency but the shortest wavelength. Green light waves, not yellow, are exactly in between the frequency and wavelength of red and blue light waves. As light enters the eye, it is focused through the lens to the back of the eyeball called the retina. The retina is covered in millions of tiny cells called cones and rods, based on their shape. These cells are considered part of the brain because their purpose is to process the light into impulses and pass that information to the cortex of the brain.

Key Topic

Cones are concentrated around the center of the retina near the optic nerve. There are six million cones in each human eye, and they are divided into three types. Each type of cone is sensitive to a different primary color in the visible light spectrum, so as light reflecting off an object enters the eye, the amount and brightness of each light wave provide enough information to the brain to interpret and associate the color being observed. In the video technology world, this is referred to as the *chrominance* of light.

Where cones are concentrated near the center of the retina, rods are concentrated around the edges of the retina. There are over 120 million rods in each human eye, but they mostly process black and white information. This information is used to help our brain interpret depth perception through varying levels of brightness. It is the rods that help your eyes adjust when you turn out the light in a room. In the video technology world, this is referred to as the *luminance* of light. Figure 3-2 illustrates the variances between wavelengths within the visible spectrum of light.

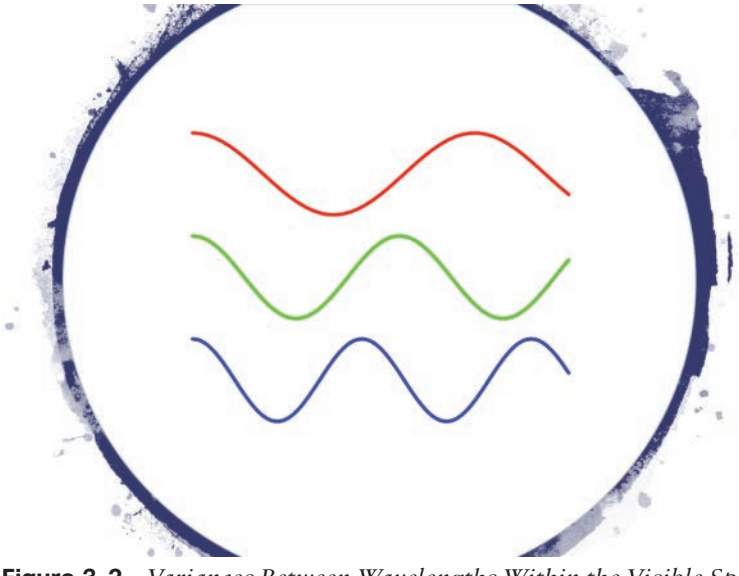

Figure 3-2 *Variances Between Wavelengths Within the Visible Spectrum of Light*

The saying "Not everyone sees things the same way" may not only refer to someone's perspective. All people have varying visual frequency response curves. Changes in how the brain processes light perceived by cones can cause what is known as color blindness. It is very rare for anybody to have total color blindness, where no color can be seen. More commonly, when someone is diagnosed with color blindness, it is usually with one or two colors. Red and green dichromatism is a diagnosis where red and green are indistinguishable from one another. Changes in how the brain processes light perceived by rods can cause issues with depth perception. Since rods deal with light levels, issues with depth perception are more common at night, or in the dark.

Brightness changes can be properly referred to as luminance (luma), and color changes can be properly referred to as chrominance (chroma). As scientists and engineers understand more about the human eye and how people see, technology has evolved to mimic this function through cameras and displays. For many years two organizations shaped how these video technologies were developed. The National Television System Committee (NTSC) was formed in North America and controlled how broadcast television operated in the USA and Canada. Phase Altering Line (PAL) was created in Europe and controlled how broadcast television operated for the UK, Europe, South America, and most of the rest of the world. The NTSC identified that the perceived brightness, or luminance, of an image can be found from the following formula:

Y = *0.299R + 0.587G + 0.114B*

Y = *luminance (perceived brightness of light)*

R = *strength of red light*

G = *strength of green light*

B = *strength of blue light*

This formula allows only three chrominance signals to be sent without the luminance signal, because the luminance signal can be derived from the chrominance signals at the destination. In YUV signaling, which was developed by PAL, the luminance information is sent as one signal, **Y**, whereas the chrominance information, or color, is sent as two signals: **U=Y – Blue** and **V=Y – Red**. We are still sending three separate signals, which allow us to reproduce any color, but the signals have different individual meanings. Therefore, if we transmit the luminance signal (**Y**) along with a red (**R**) and blue (**B**) signal, we can reproduce the green (**G**) information at the receiving end. G = (Y – 0.299R – 0.114B) / 0.587

Other formulas use differing standards, and each can be grouped into two categories:

- In **composite video**, all the video information is combined into a single line level.
- In **component video**, a video signal is split into two or more component channels.

Historically, the term *YUV* is used for a specific *analog encoding* of color information in television systems. YUV is a component video process, where Y represents the luminance component, and UV represents the chrominance components. The YPbPr color model (used in analog component video) and its digital version YCbCr are more or less derived from YUV. The YUV formula was backward compatible with black-and-white TVs and offered the ability to keep luminance separate from chrominance. The chrominance signals could be compressed more, as that information is less important, and emphasis could be put on the

luminance information. Other forms of component video processing could include RBG and HDMI. Examples of composite video could include RCA, VGA, and S-Video. S-Video does break out luminance from chrominance, but because all the chrominance components are sent as one signal, it is still considered composite video.

Color Temperature

All light has a specific temperature and that temperature will affect the lighting arrangement in a video communication environment. If you have a mismatch of light temperatures in the image that you are trying to capture, that mismatch will result in a color variation over the image. Such problems are well known to professional and amateur photographers alike. For example, it is still common today to be able to purchase film designed for either daylight or artificial light. This distinction is particularly important when flash bulbs are required. Often, film sees daylight as having a bluish tinge, whereas artificial light has a yellowish tinge. Film producers refer to light as being either warm (red shift) or cold (blue shift). Theaters and filmmakers place a lot of emphasis on adjusting light toward the lower end of the scale to create a reddish tinge.

Office environments are purposefully designed with a mix of functional lighting that offers a degree of comfort. More importantly, when a room is used for video communication, a range of light known as cool white light should be used. This effect can be achieved by using lighting in the region of 4500 Kelvin. Light temperatures affect room design in more ways than most people generally realize. Every color in a room gives off a temperature. When these temperatures are mixed, some startling effects can occur. This problem is exacerbated by a range of contributors, such as clothing, skin tones, make-up, natural light from windows or skylights, and artificial lighting.

Most light fixtures emit a range of light frequencies, not just a single frequency, so measuring the "color" of a light source is not just measuring the frequency. When objects get very hot, they start to glow and emit light. The hotter the object gets, the higher the frequencies of light it produces. The concept of black-body radiation is used to provide a scale, measured in temperature (Kelvin) by which we can assess the light being emitted by a light fixture.

Capturing and Cameras

Now that the groundwork has been laid explaining how light affects vision, it's time to turn the conversation toward the components used for video communication. Important aspects to understand include how cameras capture images, how moving pictures operate, and how captured images can be re-created on a display at a remote location. It is also important to understand some common encoding techniques.

Digital cameras work similarly to solar panels in that they use photosensitive panels to change light energy into electrical energy. Two main types of digital cameras are available on the market today:

- Charge-coupled device (CCD) image sensors
- Complementary metal-oxide-semiconductor (CMOS) image sensors

An image sensor detects variable attenuation of light waves and converts them into electrical current. The small cameras in laptop computers, smartphones, and tablets usually

use CMOS because they are less expensive and have a lower power consumption. CCD sensors are more commonly used in high-end video cameras.

Cameras detect changes only in light levels, not color, so the light must be split or filtered in some way to get the color values out. Three common techniques are used to separate color in cameras:

Key Topic

- **Foveon X3 sensors** use a method similar to how color film for photography works. An array of layered pixel sensors separates light via the inherent wavelength-dependent absorption property of silicon, such that every location senses all three color channels.
- **3CCD** (also known as three-chip cameras) colors are determined by sending the light through a prism, which separates the light into the RGB spectrum frequencies. Each color is measured on an individual light sensitive chip.
- The most common pattern for the filters is a mosaic called **Bayer**, which lets about one-third of each color in and maps the saturation levels of each color per pixel in a mosaic effect.

Regardless what method is used to filter color from light, the effect is the same. Varying saturations of color are mapped to electrical impulses so that the data can be compressed and sent across a network to a destination. Figure 3-3 illustrates how the Foveon X3 sensor operates compared to the Bayer sensor.

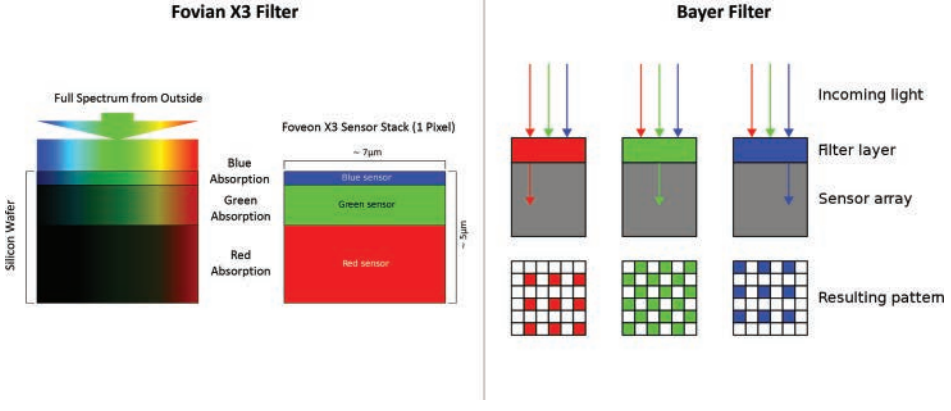

Figure 3-3 *Foveon X3 Sensor vs. Bayer Sensor*

Frame Rates

Now that you understand how cameras capture images, the next topic to understand about video communication is frame rates. Frame rates impact both video capture and video display. When movies first came out, they were referred to as *moving (or motion) pictures*. This is a more accurate description because what we see in video is actually a series of still images, called a reel, that give the illusion of motion. Each still image is called a frame, and the number of frames that are shown per second (frames per second, or fps) is known as frame rate. So how high of a frame rate must be achieved to fool the brain into seeing fluid motion?

Key Topic

The human brain can perceive motion up to about 50 fps. However, even when frames move at slow rates of change, our brains will try to fill in the missing information. Imagine a dog running behind a white picket fence. You know that every few inches a plank of wood obstructs some of your view of the dog, but your brain fills in the missing information unconsciously; there is never any question in your mind that it is anything but a dog running behind a fence. Strobe lights work in a similar fashion. Adjusting the flicker frequency of the light can make images appear as though they are moving in slow motion as our brains stitch the images together. At some point, the frame rate is so slow that our brain will no longer perceive motion. In film circles, the phenomenon of afterimage is referred to as *persistence of vision*. It is generally understood that motion shown at less than 15 fps will be noticeably distracting due to flicker. For images shown above 15 fps, our brains do not realize instantly that a change in an image has occurred. The higher the frame rate, the greater the sense of fluid motion. Film for cinema generally runs at 24 fps, but the projector shutters were designed to flash twice per frame of film, so the screen actually flickers 48 times per second, or 48 Hz, which is less noticeable to the human eye.

Key Topic

In countries where NTSC is used, the frame rate is usually 30 or 60 fps for TV or video communication. In countries where Phase Alternating Line (PAL) is used, the frame rate is usually 25 or 50 fps. The reason for this difference has to do with the power supply to lighting. Did you know that the lights in your home and office flicker? Unless the ballast is broken in fluorescent lights, you probably have not noticed. The reason is that the flicker is so fast that your brain processes it as continuous light, much like with moving pictures. Most power supplied to buildings across the globe uses an alternating current (AC). AC power is supplied in pulses rather than one continuous flow, such as with direct current (DC) power. Low-powered electric fences operate in a similar manner. If you are ever brave enough to touch an electric fence (though, speaking from experience, I do not recommend it), you will feel a shock every second or so. The AC on these types of fences pulse at a much slower rate than the power in a home or office. Throughout North America and in Japan, the AC power operates at 60 Hz, meaning there are 60 pulses per second. Throughout most of the rest of the world, the AC power operates at 50 Hz. This is where it gets interesting. The frame rate must match a multiple of the hertz rate of power; otherwise, the video is affected. Have you ever watched a video or TV show and noticed lines running up the screen? This is one of the effects caused by the frame rate and hertz rate being out of sync. Therefore, in countries that use NTSC standards, the frame rate must be 30 or 60 fps to match the 60 Hz power grid. In countries that use PAL standards, the frame rate must be 25 or 50 fps to match the 50 Hz AC power grid. A few years ago, some televisions that support 120 fps were introduced. In my opinion, purchasing one of these would be a waste of money because the human eye cannot process video beyond about 50 fps. Any frame rates beyond 60 fps have no added benefit to video quality. In contrast, an increase in pixel saturation has much to do with video quality. It is debated heavily that the human eye can detect much faster speeds than what is claimed in this book. However, there is evidence to support both sides of the argument. Therefore, you must decide for yourself if a faster fps display is worth the money. As for me and my house, we will watch TV at 60 fps.

Understanding Pixels and Resolution

All computer monitors, televisions, and displays of any sort that are used in technology today use pixels to create the images people see. *Pixel* is a contraction of the words *picture* and *element*, and this term generally is used to describe the smallest component of a digital

image. Pixels are tiny colored dots that, when combined, create a larger image, similar to a mosaic. The image in Figure 3-4 was blown up three times to reveal the pixels that make up the picture.

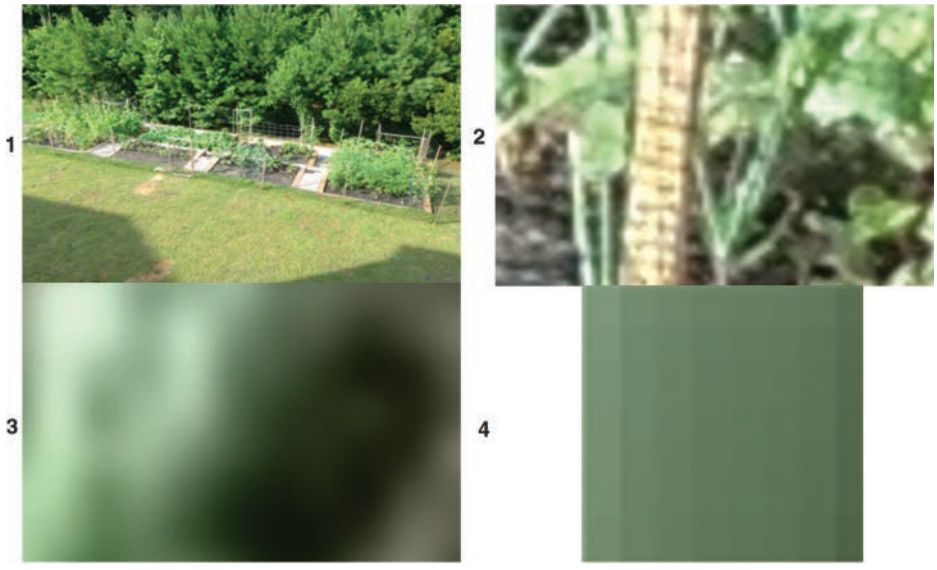

Figure 3-4 *Pixels Within an Image*

In Figure 3-4, image 1 is the original image of a backyard garden taken at some distance away. When this image is zoomed in, as can be seen in image 2, only the pixels within the image are enlarged. The images in this picture still resemble plants in the garden, but the image now appears blurry. The third graphic is zoomed even closer, and now the pixels are so large that the original picture is indistinguishable. This leads to the fourth image, which is a single pixel from within the original picture. When the pixel saturation is increased in the image, closer images can be made clearer.

Key Topic

The number of pixels within a digital frame is called *resolution*. Frames are made up from lines of pixels. When you see a resolution, such as 1280×720, the first number identifies how many pixels exist in each line, and the second number identifies how many lines exist per frame. The total pixel saturation is the two numbers multiplied together.

The conversation now comes around to the *aspect ratio* of an image, which describes the proportional relationship between an image's width and its height. An aspect ratio is commonly expressed as two numbers separated by a colon. For an x:y aspect ratio, the first number, x, represents the width, and the second number, y, represents the height. No matter how big or small the image is, if the width is divided into x units of equal length and the height is measured using this same length unit, the height will be measured to be y units. For example, in a group of images that all have an aspect ratio of 16:9, one image might be 16 inches wide and 9 inches high, another 16 centimeters wide and 9 centimeters high, and a third might be 8 yards wide and 4.5 yards high. Thus, aspect ratio concerns the *relationship* of the width to the height, not an image's actual size, but the aspect ratio can influence the pixel saturation.

One missing piece of the resolution puzzle is how these lines of pixels are populated on a display screen. There are two scan types used on all digital displays:

Key Topic

- **Progressive scanning:** This type of scanning begins in the top-left corner of the screen and populates the pixels across line 1, then moves down to line 2, again beginning on the left side of the screen, and so on until all lines within the frame have been populated. Then on the next frame, the scanning begins again, working from line 1 to 2 to 3 and so on. The main aspect of progressive scanning is that each frame is populated with all lines.
- **Interlaced scanning:** In contrast to progressive scanning, interlaced scanning populates only the odd lines on the first frame, then all the even lines on the second frame. The third frame is populated with the odd lines again, and the fourth frame with even lines, and so on.

When you look at the resolution for video, you will see a *p* or an *i*, which indicates whether progressive or interlaced scanning is being used. For example, if you were to see 1280 × 720p30, this can be read as 1280 pixels per line, 720 lines per frame, progressively scanned at 30 frames per second. Figure 3-5 illustrates the scanning process of progressive versus interlaced.

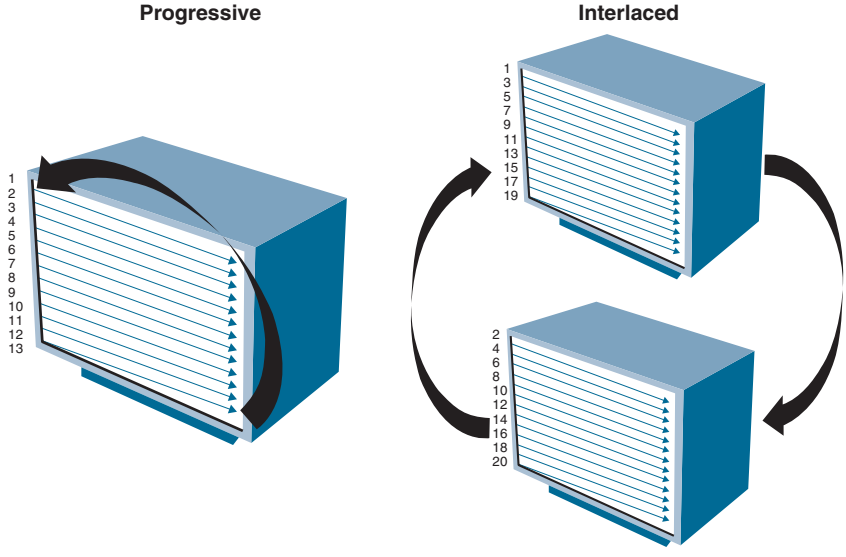

Figure 3-5 *Progressive vs. Interlaced Scanning*

There has been much debate as to which scanning method is better, but there do seem to be unique circumstances when one method is used over the other. Home TVs usually ship with interlaced set as the default scan source. Video endpoints usually ship with progressive set as the default. It seems that playback devices, such as home TVs, tend to gravitate toward interlaced, and live video systems, such as video endpoints, tend to gravitate toward progressive. Both system types provide the capability to change these settings. In my opinion, you can decide for yourself whether one scan source is better than the other. When you have time,

go to your television and change between interlaced and progressive to see which you prefer. Be warned; you may not see any difference between them.

Common Video Encoding Techniques

Throughout this chapter, I've referred to NTSC and PAL. Analog TV transmission systems, which dominated the airwaves for most of the 20th century and the first decade of the 21st century, use one of three different types of signal encoding, depending on country or region: NTSC is used throughout North America, PAL is used throughout Europe and Australia, and SECAM is used in France and the former Soviet Union. NTSC and PAL are the most pervasive of the three, and these two standards are defined as follows:

- NTSC
  - 525 lines of transmitted signal, of which 480 lines are actual visible picture content
  - 29.97 (or 30) fps
  - 75-ohm termination
- PAL
  - 625 lines, of which 576 lines are actual visible picture content
  - 25 fps
  - 75-ohm termination

Just as audio has standard techniques for digitization, such as pulse-code modulation (PCM), so too video has standardized uncompressed video formats. ITU-R BT.601 defines the color space, resolutions, and frame rates for encoding interlaced analog video signals into digital video form. A signal that conforms to the BT.601 standard can be regarded as if it is a digitally encoded analog component video signal, such as including data for the horizontal and vertical sync and blanking intervals. BT.601 has been reused in several other standards, such as MPEG. ITU BT.709 defines the color space, resolutions, and frame rates of widescreen high-definition television using the 16:9 aspect ratio. Common intermediate format (CIF) was designed by the ITU as a compromise between NTSC and PAL resolutions for digital video transmission, particularly regarding video communication. Source input format (SIF) was defined by the ISO as part of MPEG-1. Often referred to as a *constrained parameters bitstream*, SIF defines the minimum specifications any decoder should be able to handle to provide a decent balance between quality and transmission performance. Sometimes SIF is also referred to as standard intermediate format, albeit incorrectly. CIF and SIF exist because of the differences between PAL and NTSC cameras and displays; however, it is quite possible that you will not run into either of these today due to end-to-end all-digital systems that now exist. Table 3-2 identifies some of the common encoding techniques used in digital video communication. Take note of some of the PAL resolution similarities across CIF and SIF. All of the digitization techniques mentioned here use the 4:3 aspect ratio.

Key Topic

Table 3-2 Common Encoding Techniques Used in Digital Video Communication

CIF (Common Intermediate Format)	SIF (Source Input Format)
SQCIF = 128 × 96	N/A
QCIF = 176 × 144	QSIF = 176 × 140
SCIF – 256 × 192	SIF (NTSC/525) = 352 × 240
CIF = 352 × 288	SIF (PAL/625) = 704 × 480
DCIF = 528 × 384	N/A
2CIF = 704 × 288	N/A
N/A	4SIF (NTSC/525) = 704 × 480
4CIF = 704 × 576	4SIF (PAL/625) = 704 × 576
16CIF = 1408 × 1152	16SIF = 1408 × 960

When digital images are transmitted, the color and brightness information is actually coded into numerical values that contain all the information about each individual pixel within the image. Similar to converting analog audio to digital format, the more samples taken of a video image, the more accurate the digital representation will be. When digitizing an image, you are basically just dividing the image into tiny little regions, which are referred to as pixels. The more pixels, the better the resolution.

Several digital television (DTV) formats are in existence around the world, but most standard definition formats are based around NTSC or PAL resolutions so that they can be displayed easily on conventional TV screens. The most common DTV resolutions are NTSCs 480i and PALs 576i. The actual usable resolution of each, respectively, is 704 × 480 and 704 × 576, as only the center 704 horizontal pixels carry actual image. Regarding digital 480i content, where pixels actually determine the entire resolution, 480i resolution would be 640 × 480. In the case of CIF, as you can see from Table 3-2, the resolution is 352 × 288, or 101,376 pixels.

Digital television transmissions can take advantage of recent advancements in data compression and video display technology to deliver higher-quality images than standard analog TV formats. Digital TV provides various alternative options for TV formats, including progressive scanning and 16:9 widescreen aspect ratios. New formats for digital television broadcasts use the MPEG-2 video codec and include the following:

- **ATSC:** USA, Canada, Korea
- **Digital Video Broadcasting (DVB):** Europe
- **ISDB:** Japan
- **ISDB-Tb:** Uses the MPEG-4 video codec. Brazil, Argentina
- **Digital Multimedia Broadcasting (DMB):** Korea

ATSC replaced much of the analog NTSC television system in the United States as of June of 2009. In July 2008, ATSC was updated to support the ITU-T H.264 video codec. The new standard supports 1080p at 50, 59.94, and 60 frames per second; such frame rates require H.264/AVC High Profile Level 4.2, while standard HDTV frame rates require only Levels 3.2 and 4, and SDTV frame rates require only Levels 3 and 3.1.

Standard Video Codecs

Many video compression standards exist today, and each standard has its purpose because "video" can mean many different things. Much of this chapter has focused on broadcast video because the entertainment industry has driven a lot of the development in video transmission. Since high-speed Internet was introduced, many more applications for video have taken root and grown into their own subset of video. YouTube, Netflix, Amazon Prime, and Hulu have revolutionized streaming video. Major broadcasting networks have even joined the race to stream video over the Internet. Streaming video requires different standards from broadcast video. Many companies offer proprietary protocols for streaming video as well, such as Apple and Microsoft. Video communication requires a whole other set of standards regarding how to compress and send live video streams across the Internet. As cloud services become more prevalent in today's technological world, even newer standards are being developed. Before getting into the main standards that exist today, it is important to understand how video compression works.

Video Compression

In video communication, a frame is broken down into several components for the purpose of video compression and prediction. The first unit in an image division is called a *macroblock*. These units are a collection of pixels generally 16 × 16 in size but can be divided into 8 × 8 and 4 × 4 sizes as well. Each macroblock can be broken down into smaller units.

One such unit a macroblock can be broken down into is called a *transform block*. These transform blocks serve as input to a linear block transform. In the YCbCr color space, each single 16 × 16 macroblock consists of 16 × 16 luma (Y) samples and 8×8 chroma (Cb and Cr) samples. These samples are split into four Y blocks, one Cb block, and one Cr block. Figure 3-6 illustrates how frames can be broken down into macroblocks and transform blocks.

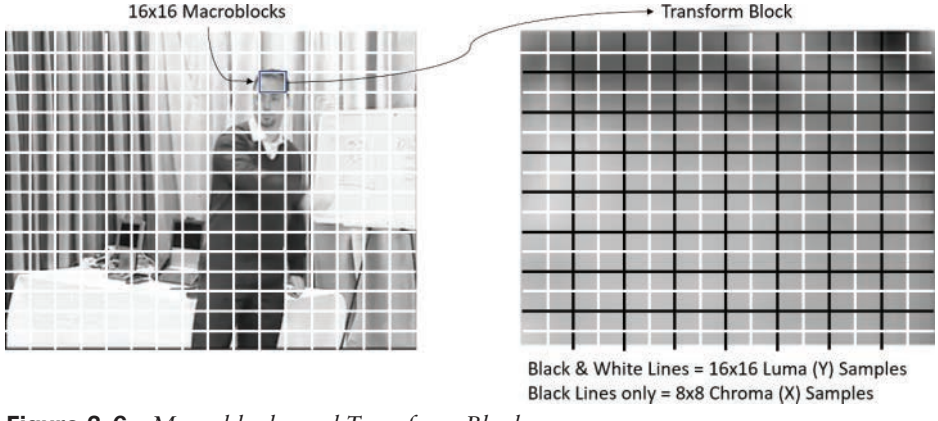

Figure 3-6 *Macroblocks and Transform Blocks*

Distinct from transform blocks, a macroblock can be split into prediction blocks. In early standards, motion compensation was performed with one motion vector per macroblock. In more modern standards, a macroblock can be split into multiple variable-sized prediction blocks, called *partitions*. In an interpredicted macroblock, a separate motion vector is specified for each partition. Correspondingly, in an intrapredicted macroblock, where samples are predicted by extrapolating from the edges of neighboring blocks, the prediction direction is specified on a per-partition basis. These prediction partition sizes range from 4×4 to 16×16 samples for both interprediction (motion compensation) and intraprediction.

Key Topic

Sending uncompressed video requires a lot of bandwidth, which is why video compression is essential, especially for live video communication. Once the macroblocks of a frame have been mapped out, video codecs search for changes in the pixels of each macroblock both before and after each frame. These techniques are known as *spatial and temporal redundancy*. Video compression involves sending only the macroblocks where change has been detected; therefore, less bandwidth is required to maintain the video connection. As motion increases and more of each frame needs to be refreshed, more bandwidth is required to compensate for the increased amount of data that needs to be sent. You can test this if you have access to a video endpoint. Place a video call and sit really still for the first minute or so. The video should come in very clear. Then start waving your hand in front of the camera, and you should see the video degrade as the codec tries to keep up with the amount of data that needs to be sent. You can also try this using different connection rates. Try it with a 768 kbps call, 512 kbps call, 384 kbps call, and a 256 kbps call. You should notice that the higher bandwidth rates can keep up with the increase in motion longer than the lower bandwidth rates. But all bandwidth rates will degrade, which is why video compression is so essential.

Key Topic

Video Quality

Three primary video codecs are used across different mediums. They are H.261, H.263, and H.264. Some of the mediums they are used with include ITU H.320, ITU H.323, and IETF SIP. Some streaming applications leverage H.264 as well, such as Netflix, YouTube, and other similar websites.

- **H.261** is the lowest video standard and was the first of the ITU video codecs. This codec will support QCIF and CIF formats, and uses 64 kbps to 2 mbps of bandwidth to transmit and receive video. Today, this standard is typically only used by legacy devices.
- **H.263** came out after H.261 and offers superior advantages. H.263 has better compression, especially in lower bit rate range, and uses basically the same bandwidth. H.263 also offers support for SQCIF 4CIF and 16CIF at a little less than 30 fps, hence a crisper image.
- **H.264**, sometimes called MPEG-4, was introduced at a time when HD communication was being more readily used. This standard was created by the ITU in cooperation with the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). It currently is the most often used for high-definition video. It is based off MPEG-4 and delivers video that is at the same quality as H.263. The reason it is so favored currently is its capability to deliver video at half the bandwidth usage as H.263.

Video Container Formats and Codecs

Video codecs can refer to either a physical device or coded software that enables video compression and decompression of digital video. An endpoint is a video codec because it does all the coding and decoding of the data. The previous codecs mentioned, such as H.264, are also video codecs because they define how data should be coded and decoded. Like audio, video was historically stored as analog signals on magnetic tape, but the digital revolution made it feasible to begin storing and using video in digital form.

There is always a complex balancing act happening between several factors regarding the quality of encoded video, such as the data bit rate, the complexity of the compression algorithms, robustness to data error correction, how easy it is to edit the compressed version, the capability of random access, transmission delay, and so on. A variety of video compression formats can be implemented on PCs and in consumer electronics equipment, and it is likely that multiple codecs are available within the same product, thus avoiding the need to choose any one specific format. Compatibility is key for ubiquity.

Most common video codecs use standard video compression formats, which make them compatible with each other. For example, video created with a standard MPEG-4 Part 2 codec such as Xvid can be played back using any other standard MPEG-4 Part 2 codec, such as Ffmpeg MPEG-4 or DivX Pro. As with audio, there are also container formats, some of which can be linked to the formats they support. The .MPEG/.MPG file container only supports MPEG-1 and MPEG-2 format standards. Others are not so clear, like .FLV file containers.

H.264 Compared to H.265 HEVC

H.264 Advanced Video Coding (AVC) is a block-oriented motion-compensation-based video compression standard that also goes by the name MPEG-4 Part 10, Advanced Video Coding (MPEG-4 AVC). At the time this book was written, it is one of the most commonly used formats for the recording, compression, and distribution of video content. The purpose of developing H.264 AVC was to create a standard capable of providing great video quality at substantially lower bit rates than previous standards without increasing the complexity of design so much that it would be impractical or expensive to implement. An additional goal was to provide enough flexibility to allow the standard to be applied to a wide variety of applications on a wide variety of networks and systems. Some of the more prevalent applications that use H.264 AVC include broadcast video, DVD storage, RTP/IP packet networks, and multimedia telephony systems. H.264 AVC supports low and high bit rates, as well as low- and high-resolution video. H.264 is typically used for lossy compression, although it is also possible to create truly lossless-coded regions within lossy-coded pictures or to support rare use cases for which the entire encoding is lossless.

The H.265 High Efficiency Video Codec (HEVC) is the newest draft compression standard ratified in 2013. It's a logical successor to H.264 AVC, which is aimed at reducing the bit rate significantly, and it leverages new compression and prediction techniques. In many ways, HEVC is an extension of the concepts in H.264 AVC. Both work by comparing different parts of a frame of video to find areas that are redundant, both within a single frame and between consecutive frames. These redundant areas are then replaced with a short description instead of the original pixels. The primary changes for HEVC include the expansion of the pattern comparison and difference-coding areas from 16 × 16 pixel to sizes up to 64 × 64, improved variable-block-size segmentation,

improved intraprediction within the same picture, improved motion vector prediction and motion region merging, improved motion compensation filtering, and an additional filtering step called sample-adaptive offset filtering. Now comes the bad news. To process all this data, a higher dependency on the hardware is required. As previous versions of video codecs were introduced, a mere software upgrade was all that was needed for those endpoints to support the newer codec. With H.265 HEVC, more signal processing capability for compressing the video is needed, so a simple upgrade patch will not render an older endpoint capable of supporting this newer codec. For this reason, the H.264 AVC codec is still the prominent codec used today, but the market is releasing a whole new line of products with a superior experience for users. Table 3-3 compares the H.264 AVC codec to the H.265 HEVC codec.

Key Topic

Table 3-3 H.264 AVC Compared to H.265 HEVC

	H.264 AVC	H.265 HEVC
Name	MPEG 4 Part 10, AVC	MPEG-H Part 2 HEVC
Approved date	2003	2013
Progression	Successor to MPEG-2 Part	Successor to H.264/AVC
Key improvement	■ 40%–50% bit rate reduction compared with MPEG-2 Part ■ Available to deliver HD sources for Broadcast and Online	■ 25%–50% bit rate reduction compared with H.264 at the same visual quality ■ It is likely to implement Ultra HD, 2K, 4K for Broadcast and Online (OTT)
Highest Resolution Supported	Supports up to 4K	Supports up to 8K
Highest Frame Rate Supported	Support up to 59.94 fps only	Supports up to 300 fps

Cisco was the first company to introduce video communication endpoints that support H.265 HEVC. The first endpoint Cisco released is called the Cisco SX80 Integrator Video Telepresence Endpoint. Based on the technology in the SX80, Cisco later came out with the MX700 and MX800 endpoints, which are run using an SX80 built into them. Cisco also released a new immersive telepresence endpoint called the IX5000 series, which also supports H.265 HEVC. Since then, Cisco has made some changes to its endpoint portfolio. These changes will be explained in more detail in Part II of this book, but essentially, all Cisco Webex endpoints now support H.265 HEVC.

Content Channels

One last aspect of video communication must be explained. One of the great advantages of communicating over video is the ability to share content though the video systems. For content from an external device, other than a camera, to be shared to the far end of the conference, one of two things must happen. Either the video stream from the camera must be replaced with the media device from which content will be shared, or an additional media stream must be added to the overall package to allow for the far end to see both the presenter and the image being shared. The latter option requires yet another protocol to

be employed. In SIP communications, the protocol to be employed is called Binary Floor Control Protocol (BFCP). In H.320 and H.323 communications, the additional protocol to be employed is called H.239.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 3-4 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 3-4 Key Topics for Chapter 3

Key Topic Element	Description	Page Number
Paragraph	Visible Spectrum of Light Explained	36
Paragraph	Chrominance and Luminance Explained	37
List	Component Video vs. Composite Video	38
List	Light Filters in Cameras	40
Paragraph	Perceived Motion at Certain Frame Rates	41
Paragraph	NTSC and PAL Frame Rates	41
Paragraph	Resolution	42
List	Progressive vs. Interlaced Scanning	43
Table 3-2	Common Encoding Techniques Used in Digital Video Communications	45
Paragraph	How Video Compression Works	47
Section	Video Quality	47
Table 3-3	H.264 AVC Compared to H.265 HEVC	49

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

3CCD, Bayer, BFCP, CCD, Chrominance, CIF, Component Video, Composite Video, CMOS, EMR, Foveon X3 Sensors, Frame, Frame Rate, H.239, H.261, H.263, H.264, H.265 HEVC, HDMI, Interlaced Scanning, ITU BT.709, ITU-R BT.601, Luminance, Macroblock, Pixel, Pixel Saturation, Prediction Blocks, Progressive Scanning, RBG, Resolution, RCA, SIF, S-Video, Transform Blocks, VGA, Visible Spectrum, YCbCr, YPrPb, YUV

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. Explain the difference between *composite* and *component*.
2. Identify the two main digital cameras available today and the three main light-filtering techniques.
3. Transform blocks are broken down into what two main components?
4. What are the two main content sharing protocols that are used in video communications?

CHAPTER 4

Collaboration Endpoint Components and Environment

This chapter covers the following topics:

Physical Components: This topic will examine the physical components that make up an endpoint.

Sound Behavior: This topic will discuss the behavior of sound in a room and the tools available to adjust how audio is perceived. The topic will then examine audio input and output devices, speaker and microphone placement, and issues to watch out for when designing a room for audio communication.

Light Behavior: This topic will examine video input and output devices, along with lighting conditions. Focus also will be placed on factors such as camera angle, video etiquette, and special conditions surrounding immersive telepresence.

Now that we've established a fundamental understanding of sound and light, the focus will turn toward the environmental conditions and equipment that affect audio and video quality. The standards and codecs can go only so far in providing good conditions for communications. Many factors within an environment itself can also negatively or positively impact the user experience. Topics discussed in this chapter include the following:

- Physical Components
- Sound Behavior
  - Microphone Types and Transducers
  - Pickup Patterns and Positioning
  - MIC Level and Line Level
  - Speakers: Active versus Passive
  - Audio Cables and Connectors
  - AEC (Acoustic Echo Canceller)
  - Microphone and Speaker Placement
  - Room Design for Noise Reduction
- Light Behavior
  - Camera Field of View, Depth of Field, and Zoom
  - White Balance

- Lighting Conditions
- Room and Environment Considerations
- Displays: Monitors and Projectors
- Video Cables and Connectors
- Immersive Telepresence
- Video Etiquette

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 4-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 4-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Physical Components	1
Sound Behavior	2–7
Light Behavior	8–12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is referred to as a codec?
   - **a.** Microphones and speakers
   - **b.** Cameras and displays
   - **c.** Endpoint
   - **d.** Network
2. Which of the following conditions is most likely to absorb the energy of a sound wave?
   - **a.** Lower frequencies
   - **b.** Higher frequencies
   - **c.** Concrete walls
   - **d.** Gypsum board walls

3. Sound waves reflecting off many complex surfaces and dispersing radiated energy so that it is less direct or coherent is referred to as what?
   a. Reflection
   b. Diffusion
   c. Absorption
   d. Echo
4. What is the difference between a condenser mic and a dynamic mic?
   a. Dynamic mics do not require an external power source.
   b. Condenser mics do not require an external power source.
   c. Dynamic mics have a tighter pickup area.
   d. Condenser mics have a tighter pickup area.
5. Which of the following mics has a 360-degree polar pickup pattern?
   a. Cardioid
   b. Supercardioid
   c. Unidirectional
   d. Omnidirectional
6. Professional-grade products use a measurement scale called decibel-unloaded, or dBu. What is the dBu range for line-level inputs?
   a. −10 dBu
   b. −4 dBu
   c. +4 dBu
   d. +10 dBu
7. Which of the following is an unbalanced audio connector?
   a. TS
   b. TRS
   c. XLR
   d. HDMI
8. The distance from the back of the lens to the frame sensor in a camera is referred to as what?
   a. Aperture
   b. Zoom
   c. Focal length
   d. Magnification
9. What tint will an image take on when the color temperature is cooler?
   a. Red
   b. Blue
   c. White
   d. Green

10. What is the maximum color temperature in kelvins recommended for video communication rooms?
    a. 3200 k
    b. 3900 k
    c. 4100 k
    d. 5000 k
11. Which of the following is the ideal location of a camera in a video communication room?
    a. Centered just above the display
    b. Centered just under the display
    c. Centered high on a wall above the display
    d. Multiple cameras placed around the room for different perspectives
12. Which of the following is an example of a composite video connector?
    a. HDMI
    b. DVI-D Dual Link
    c. DVI-D Single Link
    d. Y/C

Foundation Topics

Physical Components

Key Topic

Endpoints encompass many different form factors and span from software-based applications used on a laptop to personal deskphones, meeting room systems, and immersive Telepresence rooms (which are entire rooms equipped for and used solely for video communication). Whereas the term *endpoint* can refer to an audio-only phone or a video communication system, the term *video endpoint* refers specifically to any system that is used to make a video call. Regardless of the size, type, or purpose of an endpoint, all video endpoints share certain basic components, including a monitor, camera, microphone, speakers, codec, power, and some sort of network connection. Each video endpoint can differ in complexity, but all these components must exist for an endpoint to be fully functional. Each of these components can be divided into one of five categories: audio input, audio output, video input, video output, and cables and linkage. Figure 4-1 illustrates the physical components and form factors of endpoints.

The monitor is the video output device used to display video data received during a video call. It can also display data, such as a presentation received during a meeting. In addition, it displays the menus for the endpoint itself. The endpoint menu can be used to place or answer a call; change administrator settings; pan, tilt, or zoom the camera; and share content during a call. Some video endpoints have a monitor built in, but for integrator-type systems that do not come with a monitor, careful consideration should be taken to select an appropriate display.

Figure 4-1 *Physical Components of Endpoints in Various Form Factors*

Cameras are video input devices used to capture the video data that the endpoint will be transmitting. The quality of the image sent is limited, in part, to the camera's capability to capture an image. A built-in laptop camera might not produce the same quality image that a 1080p camera is capable of producing. The camera position is also an important consideration. Often the camera can be panned, tilted, or zoomed in or out depending on the preference of the user. Typically for video communication purposes, these will be Pan, Tilt, Zoom (PTZ) cameras or might even incorporate voice tracking and other automated features.

Microphones are audio input devices that capture the audio for the data. A software endpoint running on a computer, smartphone, or tablet will typically utilize the microphone built into the system. This is also true for endpoints that have a microphone built into them, but an integrator system will use external microphones that can be placed strategically around a room. The type and placement of the microphones play important roles in the quality of the sound that is heard by the far-end participants in the call.

Speakers are audio output devices used to broadcast the audio at the far end of the call. Just as with some of the other components mentioned, speakers are included in some endpoints, whereas they may be external to the endpoint system on others, such as with integrator systems. Much to the same effect as microphones, speakers should be strategically placed within a room to enhance the quality of audio heard.

Cables are used to connect the different systems together, supply power to the system, and link the network connections so that transmissions can be sent to a remote destination. All-in-one endpoints have only one or two cables that need to be connected. Some endpoints, such as the Cisco DX80, require a power cable and an Ethernet cable. Other endpoints, such as the Cisco SX10 or Cisco 8865 video phone, require only the Ethernet cable because it can supply both Power over Ethernet (PoE) and the data connection over the same cable. Integrator endpoints will require more cable connections to support all the peripheral devices, such as the microphones, speakers, cameras, and monitors along with power and Ethernet. Older legacy endpoints may even support BRI or PRI ISDN connections natively on the endpoint.

The term *codec* refers to the device used in the encoding and decoding of the audio, video, data, and control streams sent and received during a call. The codec is software within the endpoint that receives incoming audio and video from the microphones and cameras, encodes it, and sends it out across the data network. When already-encoded data comes into an endpoint from a far-end destination, the codec will decode this information before sending it out the speakers and monitor. As mentioned in the preceding chapter, the term *codec* is also used to describe the coding standard used for a signal, such as H.264. An endpoint uses codecs like H.264, so it too is referred to as a codec because it is doing the actual coding and decoding.

Sound Behavior

As discussed at some length in Chapter 1, "Introduction to Collaboration," sound waves will continue to move radially outward from the original source unless obstructed by an obstacle or until the energy runs out. In general, when a sound wave hits an obstacle, part of the wave will be reflected away and part will be absorbed, or transmitted, through the obstacle. To what degree sound waves are reflected or absorbed depends on physics, but basically, it's determined by the density and texture of the obstacle. The proportion in which a sound is reflected, absorbed, or transmitted depends on the shape and density of the material and the frequency of the sound.

When a boundary or obstacle, such as a wall, ceiling, or column, is encountered by a sound wave, some of the sound energy will be absorbed within the material. Absorption is similar to transmitting through an object, except that with absorption, the sound waves will completely dissipate within the object of obstruction. Different materials reflect some frequencies more efficiently than others, due to their roughness or absorbency characteristics. Also, lower frequency waves have an easier time being absorbed into solid surfaces than high frequencies. This is why your neighbors always complain about your bass rather than your treble.

Reflection is caused when an object of obstruction causes sound waves to bounce, or reflect, into another direction. Most hard surfaces will reflect sound waves. The path taken by reflected waves works much like you would expect a billiard ball that has banked off the bumpers of a pool table. The law of reflection is "the angle of incidence equals the angle of reflection." In other words, the angle after the impact will be equal to the angle before the impact. Think of how the balls on a pool table behave as they bounce, or reflect, off the walls of the table. Those bank shots are difficult for an inexperienced player to achieve due to the need to match the angle the ball must travel to the wall with the angle it will reflect, but the principle is the same as with sound waves.

As sound waves are reflected, diffusion occurs too. Diffusion is caused by sound waves reflecting off many complex surfaces and is the process of dispersing radiated energy so that it is less direct or coherent. Coherent reflections, or reflections you can distinguish, are what tend to cause problems in a listening environment. The plastic cover over a fluorescent light acts as a diffuser, making the light spread out in a more randomized way so it is less harsh. A textured wall, such as brick, would be better at diffusing sound than a completely flat surface, such as a concrete wall.

Reflections will happen all over an average room. So why don't all those reflections make it difficult to hear the original source? In some cases, that is exactly what happens. However,

sound reflections actually help humans to aurally perceive the size of spaces we are in and are important to our hearing in general. Experiments have been conducted in anechoic chambers to understand human reactions to environments containing no ambient sounds. An anechoic chamber is a room massively insulated with layers of concrete and steel to block out exterior sources of noise. It's also internally lined with crosshatched buffers that absorb all sound; even the floors are typically suspended mesh to stop any sound of footfalls. These chambers are like black holes for sound. Staying inside one longer than 15 minutes has been known to cause extreme symptoms in some people, from claustrophobia and nausea to fear, panic attacks, and aural hallucinations. There is a reason that sensory deprivation is considered an act of torture. The presence of ambient sounds and reflections communicates to the brain the normalcy of an environment. When ambient sound is absent, that signals the brain there is some kind of problem. Therefore, too much reflection can be bad, and not enough reflection can also be bad. The key is finding the right balance of reflection, which is referred to as *sound balancing* a room.

The reception of multiple reflections off walls and ceilings within a few milliseconds of each other causes reverberations, which is the prolonging of a sound. Reflections can actually be categorized into four main groups:

- Direct sound
- Early reflections
- Coherent late reflections, also known as echoes
- Incoherent late reflections, also known as reverberation

Direct sound is pretty obvious; it's the first and primary sound waves that hit your ears. In audio production environments, a direct sound can be referred to as being "dry," and all other sound is referred to as being "wet."

Early reflections are sound waves that bounce off obstructions but arrive at your ears at almost precisely the same time as the direct sound coming from the sound source. These reflections are measured in milliseconds, and human brains do not distinguish theses early reflections as "new" or separate sounds. People hear the reflections instantly as part of the richness of the original sound wave. Early reflections must usually arrive at your ear in less than 30–40 milliseconds for your brain to consider them to be indistinguishable from the original sound.

Late reflections fall into two types: the diffuse and incoherent type we think of as reverb, or reverberation, and the more coherent type we generally call echoes. Reflections from surfaces that stand out from normal reverberation levels are the ones we typically identify as echoes. In this case, the arrival of the late reflection sound waves has passed a certain millisecond tolerance, so those waves will be perceived as a second sound rather than the prolonging of the first sound. Reverb is essentially a bunch of echoes, where the reflections are so close and mixed together that the results are perceived as a single prolonged sound. Reverb generally reduces articulation of sounds such as speech, but is desirable when listening to music, as it is perceived to add a warmth or richness.

Reverberated sounds will eventually lose energy and drop below the level of perception. The amount of time a sound takes to die away is called the *reverb time*. A standard measurement of an environment's reverb time is the amount of time required for a sound to fade to –60 dB.

A sound that remains audible for a long time before the room absorbs it has a long reverberation time. Reverberation time is controlled by the size and shape of a room; however, the objects in the room also have an influence. Reflective objects such as hardwood floors increase reverberation time because the sound waves have more opportunity to reflect. By contrast, materials such as carpet and drapery are absorbent and decrease reverberation time because they absorb the sound waves and do not give them a good opportunity to bounce. Even air qualities such as humidity can affect reverb time. Most building materials are given a noise reduction coefficient (NRC) rating that informs the degree that a substance absorbs sound energy. Should the need arise in a room where reverberations are noticeable, acoustic paneling can be used to remediate the issue. Cisco recommends NRC ratings greater than 0.75 for acoustic paneling.

Another issue to keep in mind when designing a room is background ambient noise. This could include air flow through the HVAC system, the buzz of lights, noise coming through the walls from a neighboring room, or hallway traffic outside the room. Sounds coming from outside could also cause ambient noise issues, such as street traffic, planes, or subways and trains. Best practice is to control the environment such that any ambient noise is less than 45 dB.

The sound behavior described up to this point could occur whether technology is being used or not. An engineer designing a room for technological applications, such as a conference room, should be even more mindful of sound behavior because microphones can be more sensitive to sound in a room than the human ear. A person sitting in a room may not notice the air blowing out a vent, but if ceiling-mounted microphones were used in that room, the microphones could pick up that noise and magnify it before sending it to the destination. This could cause a real distraction to the participants at the other end of the call. There is only so much that can be altered within a room to balance the sound. Therefore, some technology applications assist in controlling sound picked up by a microphone.

Gain refers to the capability of a system to increase the power or amplitude of a signal between the input and the output of a given circuit. Gain can come in the form of amplification or attenuation for either a digital or analog process. No change in the signal as it passes through the microphone is called *unity gain*, or just unity. Headroom can be thought of as a safety zone for unintended peaks of a signal. It becomes particularly important when a signal may go through several opportunities for gain adjustment within a given system. Microphones, mixers, and amplifiers should be adjusted to always allow adequate headroom to avoid clipping.

Clipping is a form of distortion that occurs when a signal exceeds the maximum dynamic range of an audio channel. When you view a clipped sound wave, it will look as though someone chopped off all the peaks. In the analog world, musicians sometimes desire clipping because the distortion isn't a clean slice like in the digital realm. Digital clipping sounds very harsh and is usually avoided at all costs. The clipping referred to here is also known as *hard clipping*. Soft clipping is less harsh and is commonly called *overdrive*. You may have seen an overdrive knob on a guitar, for example. Soft clipping is not common in the audio and video communication world. Automatic gain control (AGC) makes dynamic adjustments to gain, typically on a microphone signal, to maintain an optimal level for different speakers. A great example is conference phones because you can hear the volume level change as each new person begins speaking, especially if the speakers are different distances from the microphone pickup area. The term *good levels* refers to an audio signal that is significantly

higher than the noise floor, but not so strong as to cause clipping, and may also indicate leaving appropriate headroom. Figure 4-2 illustrates how clipping can appear and how the gain controls should be set in an ideal environment.

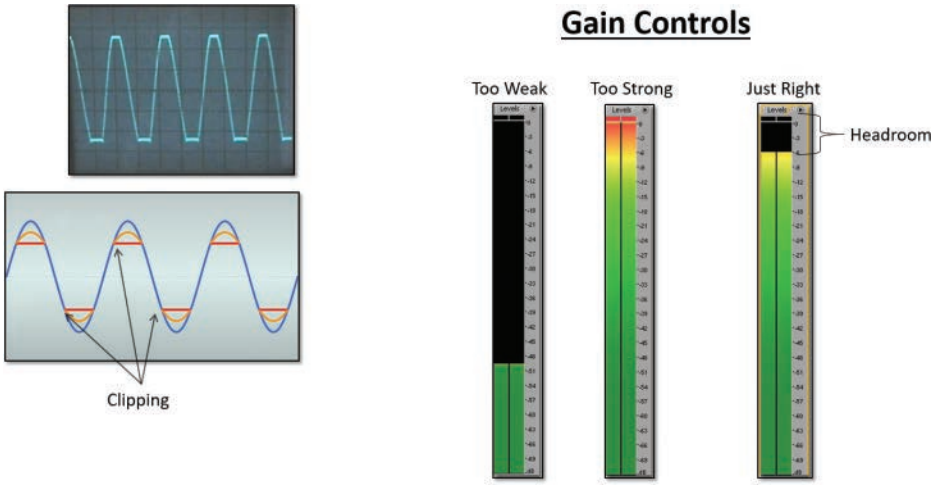

Figure 4-2 *Gain Controls and Clipping*

Microphone Types and Transducers

Because a microphone converts acoustic energy into an electrical signal, it is another form of a transducer. There are many different types of microphones with many different applications, a few of which will be mentioned in the content that follows. Some of the specifications most commonly found with microphones include the following:

- **Dynamic Range:** The range of amplitudes that can be accurately captured by the microphone. This range is usually represented in dB at a certain frequency, typically 1 kHz.
- **Frequency Response:** The range of frequencies accurately captured.
- **Polar Pattern:** The directionality pattern of highest sensitivity for the microphone element.

For best results, microphones should always be used for the applications for which they are designed. Many mics work well only in certain environments and deliver very poor results when used outside these areas. Most mics have transducer elements that fall into one of two categories: dynamic and condenser.

Key Topic

Dynamic microphones are also known as passive microphones because they do not require an external power source to operate. Instead, as sound waves strike the diaphragm, its vibrations move a magnet surrounded by a stationary coil, resulting in an electrical signal being transferred through the cable and ultimately to the pre-amp. Rather than employing a moving coil using an electromagnetic principle, condenser mics operate on electrostatic principles in which two conductive plates in close proximity to each other exchange a charge as the plates vibrate. This operation requires that an external power source be used to supply the charging voltage to the element, usually in the form of phantom power. A variation on the condenser mic is the *electret* mic, which utilizes a small battery to charge

the diaphragm. Because condenser mics require a power source, most modern mixers have the option to turn on phantom power, which supplies a voltage, usually +48v, to any mic that requires it. It is called *phantom power* because it doesn't show up on dynamic mics, thereby protecting the element from possibly damaging voltage. It is advisable to check the requirements for your mics of choice to ensure that your mixer or digital signal processor (DSP) can effectively supply the required voltage. If not, external phantom power supplies also are available. Figure 4-3 illustrates the physical differences between condenser microphones and dynamic microphones.

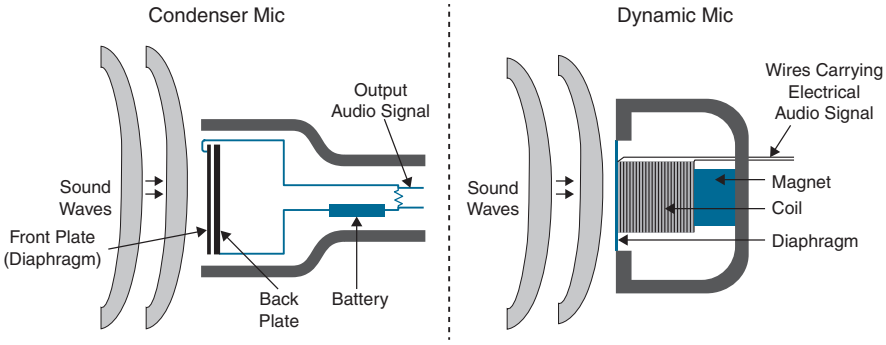

Figure 4-3 *Condenser Mic and Dynamic Mic Physical Differences*

Many different types of microphones are available on the market today. Some of the more common types of microphones are the types you would typically see in a conference room setting:

- **Handheld mics** are the most common style microphones. These are typically dynamic mics and as a result have lower sensitivity. They can be both wireless and wired.
- **Lapel mics**, also known as tie-clip or lavalier mics, are almost always condenser mics and generally are more sensitive and have a tighter polar pattern than handheld mics because they are intended to be used further from the speaker's mouth.
- **Desktop** or **podium mics** are usually condenser style and work well at greater distances from the speaker. They can be considered more forgiving for the inexperienced user, but as a result can require more processing to avoid feedback situations.
- **Ceiling-mounted mics**, also known as choir mics, are best utilized in multiuse rooms with moveable furniture or in cases where tabletop mics interfere with meeting operations. Significant design considerations need to be taken into account with these condenser-style mics due to their extreme sensitivity.
- **Boundary (or PZM) mics** are designed with a low-profile form factor for tabletop use. With these, as with other highly sensitive condenser mics, processing may be required if multiple units are in use.

Pickup Patterns and Positioning

The polar pickup patterns of all microphones can be broken down into two major types:

- **Directional, also called Unidirectional:** Directional mics are intended for pointing at a more selective group of audio sources. They typically have a greater sensitivity, and resistance to feedback is achieved with this class of mics.

- **Omnidirectional:** Omnidirectional mics are intended to be at the center of a group of audio sources, such as in the center of a conference room table, with a 360-degree pickup area. An omnidirectional mic can pick up sound equally in all directions. As a result, these mics tend to have lower sensitivity and are also more prone to feedback due to the nature of sound reverberation and other acoustical factors.

Pickup patterns of directional mics can be broken down further into several categories. The *cardioid* pickup pattern is shaped like a heart. The mic's effective pickup area primarily focuses on the area in front and to the sides of the microphone. Vocal mics are typically one of the cardioid styles. The *supercardioid* polar pickup pattern has a narrower range of focus, but as a result of this narrowing, a small lobe develops behind the mic and may require consideration. *Hypercardioid* has an even narrower focal range, and a larger lobe at the rear of the mic develops.

Bidirectional mics pick up equally in two directions, usually at 180 degrees opposed. Shotgun mics have a very narrow pickup range but can also pick up sound from the greatest distances. The distance factor is the ability of a microphone to pick up sound from a distance as related to an omnidirectional mic. Critical distance is the distance between the person who is speaking and the microphone as it relates to other active mics. A good rule of thumb is a 1:3 ratio. So, if the person speaking is
3 feet from the target mic, the next closest mic should be no closer than 9 feet.

Table 4-2 identifies the pickup patterns of different microphones with other relevant information.

Key Topic

Table 4-2 Microphone Pickup Patterns

Polar Pattern Name	Omnidirectional	Cardioid	Supercardioid	Hypercardioid	Bidirectional
Polar Pattern	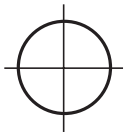	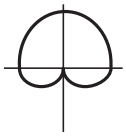	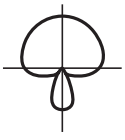	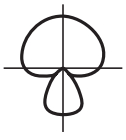	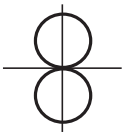
Angle of Coverage	360	130	112	103	90
Null Angle (Angle of Maximum Rejection)	N/A	180	120	108	90
Rear Rejection	0	23 dB	14 dB	7 dB	0
Ambient Sensitivity	100%	32%	26%	24%	32%
Distance Factor (in Meters)	1	1.8	1.9	2.1	1.6

Mic Level and Line Level

There are some considerations to heed when configuring and setting up systems with a variety of sources. Among the primary concerns during configuration are the various levels of audio signals in play and how those signals need to be handled. Dynamic mics have a lower voltage output that require amplification to boost their signal. This amplification can be performed with a preamp, mixer, or other amplifier device that possesses a gain control. Condenser mics have a higher output and might not need any amplification, or else much less. Most of the time these different mics offer a plug-and-play experience because line-level plugs are different from mic-level plugs. However, in some circumstances an engineer may need to manually configure a solution that requires mic or line levels. Certain Cisco endpoints, for example, allow the mic inputs to be configured as mic or line level. When the microphones plug directly into the endpoint itself, mic level can be used, which is the default setting. When more microphones are being used than the endpoint can support, such as a series of ceiling-mounted microphones in a conference room, an amplifier external to the endpoint may be required. The amplifier can be plugged into the same input on the endpoint where a microphone would have been plugged, but the setting for that input must be changed to line level to support the amplifier being connected.

Line level is usually the output of any device with an internal pre-amplifier, including MP3 players, mixers, TVs, and CD players. Mixer outputs are usually line level by default, but many are configurable to output mic level when needed. Line level can be expressed in a few ways, each with its own context. The decibel-volt, or dBv, is usually used with consumer audio equipment, and line level for that range of products is approximately –10 dBv. Professional-grade products use a different measurement scale called the decibel-unloaded, or dBu. Line levels for these products range around +4 dBu. In absolute terms, the actual peak-to peak voltage of a –10 dBV signal is about .447 volts, whereas a signal at +4 dBu equals a p-p voltage of around 1.7 volts. As you can see, there is quite a range here, depending on the equipment you choose, and allowances are needed to prevent overdriving the amplifier. Speakers also require amplifiers to project sound.

Speakers: Active versus Passive

Just as there are two types of microphones, there are also two types of speakers:

- **Active speakers** have amplifiers built into the speaker body. Active speakers need to have both a power supply and line-level audio wires connected to them. Computers, televisions, and endpoints use active speakers. Advantages of using active speakers include a self-contained form-factor, portability, and the number of speakers supported is not limited by external amplification limits. Field cables are only line level, thereby resolving cable isolation issues caused by unshielded speakers.
- **Passive speakers** do not have built-in amplification. They require an external amplifier to supply the appropriate signal; however, only unshielded audio cables need to be run to the speakers themselves because wall power is run only to the amplifier. Advantages of passive speakers include lighter weight for portability and, typically, better fidelity, which allows the "steering" of the signal with

customized crossovers and bi-amplification. Concert halls typically use passive speakers. In video conferencing, ceiling-mounted speakers are usually passive speakers.

Audio Cables and Connectors

An audio system has different signal types depending on both the source and destination of the signal and the way the signal is referenced to ground. For these signals to get from one device to another with minimal noise interference, the industry has created various connectors with which to make solid and trouble-free connections at each end of the cable. These connectors can be broken down into two primary classes: unbalanced and balanced. Figure 4-4 illustrates the differences between unbalanced and balanced audio cables.

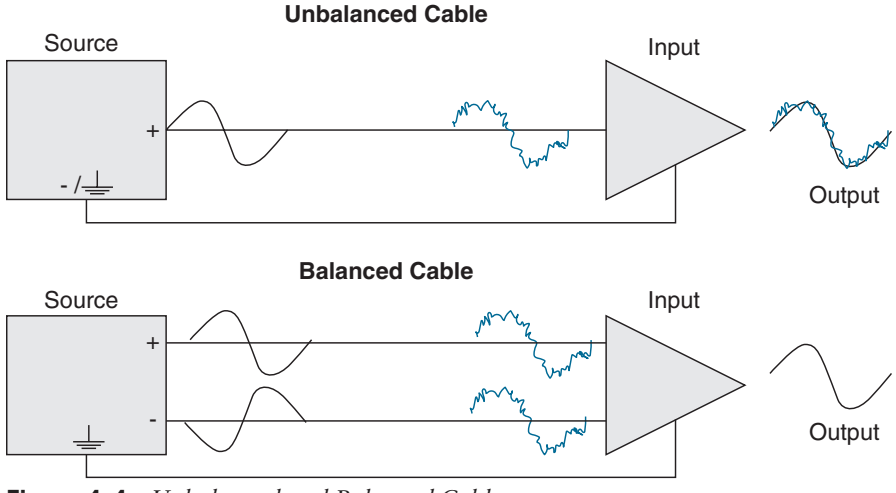

Figure 4-4 *Unbalanced and Balanced Cables*

Key Topic

Unbalanced cables have two wires: one carries the positive (+) side of the signal, and the other wire shares both the negative (–) side of the signal and the grounding shield. These cables are more typically found on consumer-level items. Inside the cable itself, the signal wire is typically in the center of the cable with the ground wire surrounding it. The ground wire serves two functions. It carries part of the audio signal and serves to shield the main signal wire to some degree from outside interference from noise. It does help reject some noise, but the wire itself also acts like an antenna and picks up noise. Unbalanced cables work fine in short runs, but they should have a maximum length of 15–20 feet, or 4–6 meters. RCA connectors are always used on unbalanced cables. Tip-sleeve (TS) connectors are unbalanced but resemble tip-ring-sleeve (TRS) connectors, which are balanced. An easy way to tell them apart is to look at how many rings are on the end of the connector. If it has one or two rings, the cable is unbalanced. If it has three rings, the cable is balanced.

Balanced cables are characterized by three wires in the cable, two of which carry the identical signal 180 degrees out of phase with each other. The third is the ground that encompasses the other two wires to protect them from outside noise. This allows for better isolation of the signal from EMI noise. The positive conductor carries the original audio signal, and the negative conductor carries the inverse of the original audio

signal. If you sum two signals that are identical but are reversed in polarity, the signals cancel out, leaving you with silence. When a positive is added to its negative counterpart, the outcome will always be zero, such as +20 added to –20 equals 0. Normally, you would not want audio gear that flips the polarity of your signal, but in this case you do. Both copies of the signal, positive and negative, pick up the same noise as they travel along the cable, and that noise is identical on the two wires in the cable. The component receiving the audio signal will flip the inverted signal back into its original orientation and invert the noise riding on the negative wire. Flipping the polarity of what arrives at the receiving gear will produce the original signal intact, and the noise, which now has reversed polarity, will be removed. What you end up with is a welcome result: a signal that's preserved and noise that's canceled. Because of this function within balanced cables, they can support much longer cable runs at 50–100 feet, or 15–30 meters. Standard connectors designed for use with balanced signals are XLR and TRS. Figure 4-5 identifies the most common balanced and unbalanced connectors available on the market today.

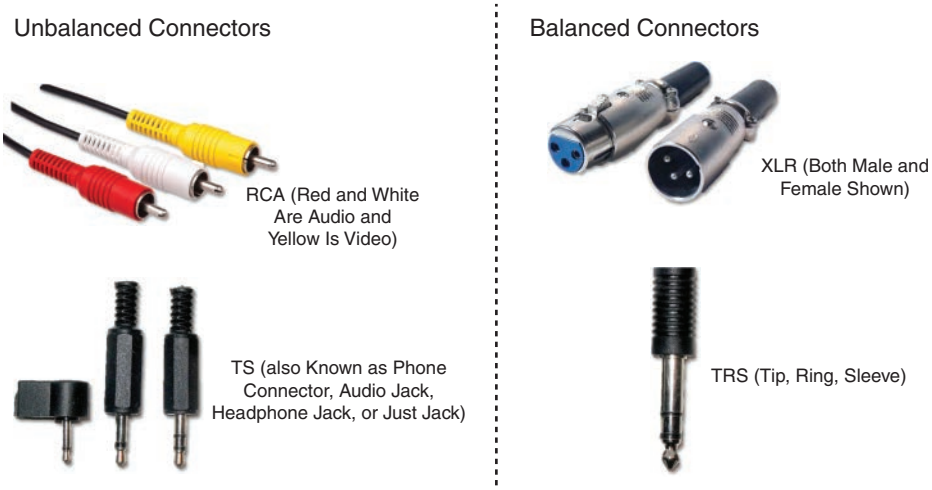

Figure 4-5 *Balanced and Unbalanced Connectors*

Acoustic Echo Canceller (AEC)

Imagine two conference rooms, each with ceiling-mounted microphones and ceiling-mounted speakers. As a person in one of those rooms speaks, the microphone picks up the audio. That sound is carried from the mic to the endpoint, across the network, to the endpoint on the far end where the audio is ultimately delivered to the amplifier, which sends it out the speakers. The speed that this audio is able to travel, whatever that distance might be, is almost instantaneous. However, a delay does in fact occur, and that delayed audio is detected by the far-end microphones. The audio picked up from the speakers on the far-end microphone would be sent back to the original location, causing an echo at the near end. You may have experienced something like this before, and it can be very distracting, sometimes to the point that it is intolerable. It doesn't matter whether the microphones and speakers are ceiling mounted or not; this unnatural echo will occur no matter what. All analog mics produce echo because echo is leakage of the analog signal in the RX path to the TX wire. The echo becomes noticeable to the listener as the leakage increases.

Key Topic

This is where Acoustic Echo Canceller (AEC) comes into play. AEC works by comparing the audio input from the near-end mic against the audio input from the far-end mic and subtracting the common delayed audio. This is the significance of that delay. Participants can be speaking at the same time from both locations, and AEC is able to filter out the delayed audio and reduced echo. This process happens on both ends of the audio communication by the endpoint reducing the local loudspeaker contribution to the local microphone. It is important to understand that AEC reduces echo, but it does not completely eliminate the echo. Figure 4-6 illustrates how AEC works to prevent echo from occurring.

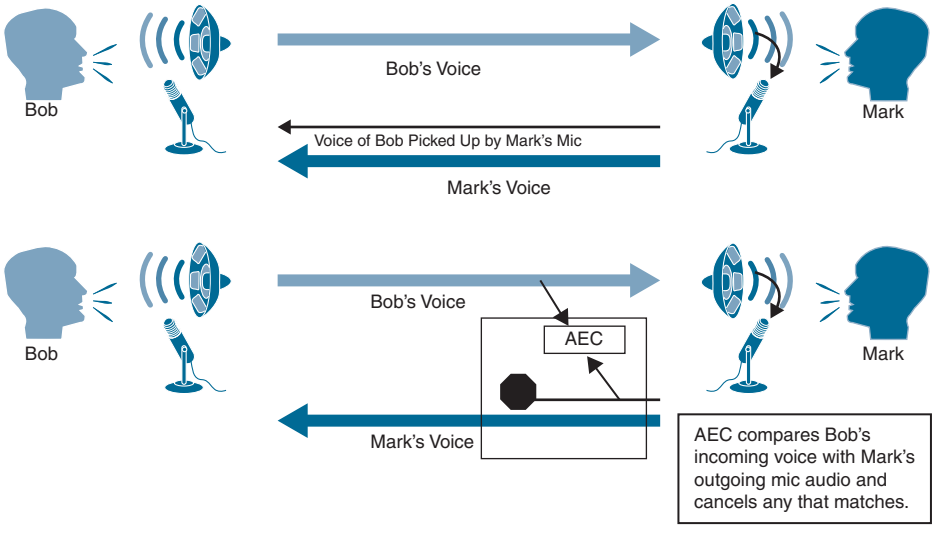

Figure 4-6 *Acoustic Echo Cancellation (AEC)*

AEC samples the signal being sent through the loudspeaker and establishes this as the reference signal. From this signal, the AEC compares the signals with the use of a predictive filter and learns to determine the difference between the near-end and far-end signals. This process is referred to as *training* or *convergence*. It usually takes a few sampling cycles for the algorithm to get sufficient sampling for effective filtering, and there are notable limitations that should be accounted for. AEC cannot correct direct acoustical anomalies in the room. It has a limited capability to correct the effect of room acoustics to the far end. Presenters need to be 1–2 two feet from the target mic for AEC to operate effectively. Finally, in situations of high network latency, sometimes as low as 200 ms, AEC can become ineffective. Many different devices use AEC, and if two devices in the same room are trying to cancel out echo, they will cancel each other out; thus, echo will be heard. In the first example provided, if ceiling-mounted microphones are being used, an amplifier or mixer external to the endpoint may have AEC enabled. The endpoint also has AEC enabled. Therefore, an engineer must disable AEC on one of the devices for it to work properly.

Microphone and Speaker Placement

The placement of microphones and speakers is critical in determining the quality of an audio call. The location of the microphone should be directly related to the type used and as close to equal distance from all participants as possible. As discussed previously, different microphones pick up audio in different patterns, so special emphasis should be placed on how

each microphone should be positioned based on the polar pattern. Equal to the importance of microphone placement, speaker placement is also essential.

Key Topic

A microphone's polar pattern is the pattern in which a microphone picks up sounds. An omnidirectional microphone's polar pattern is nearly spherical. Omnidirectional microphones can be found on desktops or hanging from ceilings. This type of microphone should be placed in the middle, or directly above the middle, of the participants so that no participants are outside the sphere of where the microphone can pick up sound. A directional microphone, of which the cardioid microphone is the most common, is typically used as a desktop microphone. It has a polar pattern that could be described as kidney shaped, with a dead zone behind the microphone. Therefore, a cardioid microphone should be placed at the end of a table, with the dead zone directed away from where participants are positioned. When both directional and omnidirectional microphones are positioned in environments where multiple microphones are being used, the polar patterns should overlap. Figure 4-7 illustrates the positioning of these two types of microphones in meeting rooms.

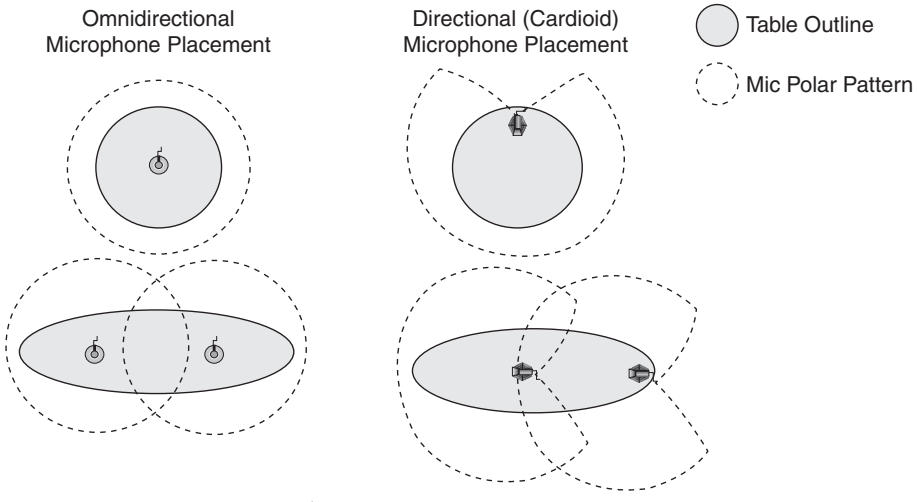

Figure 4-7 *Microphone Placement in Meeting Rooms*

Key Topic

Where speakers are placed within a meeting room could also impact the quality of audio. Many audio-only and video endpoints have the speakers built into the system. Video endpoints may use a third-party display with built-in speakers for the endpoint. In these cases, the placement of the speakers is less of a concern. Audio-only speakerphones are generally placed in the center of a table, and that is ideal. Video endpoints that use speakers built into the endpoint itself or the speakers built into a display should be positioned at the front of a room, which is the superlative location for video calling. The speakers are behind and out of range of the microphones, which will prevent feedback. Also, in a video call, the sound comes from the direction in which the far-end participants are seen. This will provide a more natural flow of sound. If someone were sitting in front of you speaking, you wouldn't expect to hear the person's words coming from behind or above you. If external speakers were set up around the room, the audience would be looking at the person speaking in front of them but hear the speakers from perhaps their right or left side, or behind them. This could be disruptive. Likewise, you would never want to set up surround sound for a meeting room

because of the same unnatural result. The one exception to this would be in a large theater-style setting. Due to the room's size, additional speakers should be put on the sides but still face out from the person speaking. This position helps to keep the listeners' orientation facing forward while achieving the volume level needed for the entire group to hear the person speaking clearly. In some custom meeting room integrations, the room will be designed with ceiling-mounted speakers above each participant chair in the room. Although this is not the ideal positioning of speakers within a room, for larger meeting room settings, this setup does distribute the audio more evenly throughout the room.

Cisco has a website at https://projectworkplace.cisco.com that provides many room design ideas. There, you can find pictures to help you visualize what a room will look like, and you can click through various customization options to change the room based on endpoint selection, participant capacity, and general purpose of the room. You can even open a schematic of the room to scale that will illustrate a two-dimensional drawing of the room's layout with measurements and total room layout, and you can then send these plans to an architect for official blueprint design. Figure 4-8 illustrates one of the schematic drawings available on Cisco's project workplace website.

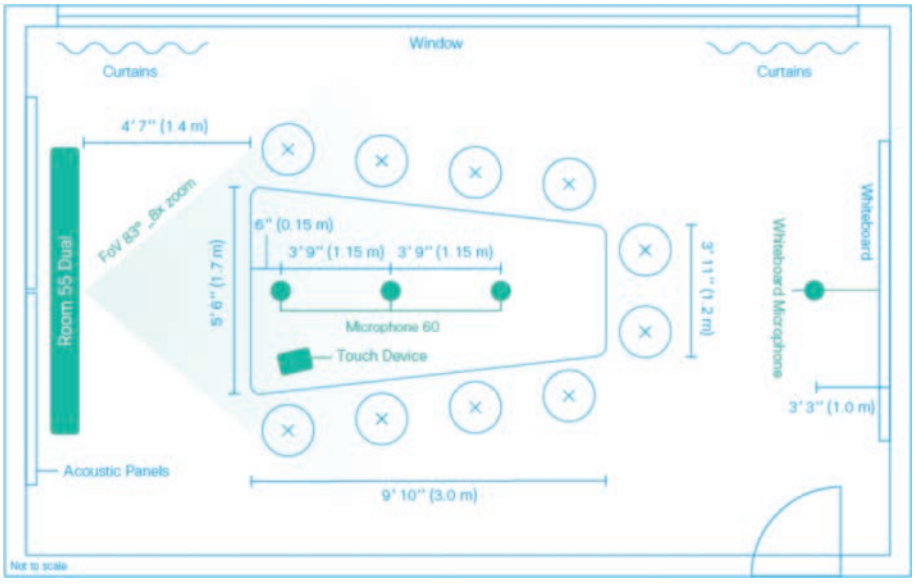

Figure 4-8 *Project Workplace Meeting Room Design*

Room Design for Noise Reduction

You learned about sound behavior at the beginning of the chapter. You should now understand that sound behavior can help identify external noises that can ruin a call, and it can help identify appropriate measures to take so that ideal conditions are met for premium audio quality within a meeting room. I placed a great deal of emphasis on the technical components that help improve sound quality. Such components include using appropriate microphones, speakers, and cables. You also learned about using balanced audio cables instead of unbalanced, and proper placement of microphones and speakers within a room. Beyond these technical measures, you can take many other steps to improve audio quality if you understand the behavior of sound.

It is important to take steps to eliminate as much external noise as possible. First, the location of the meeting room is very important. Location may not always be within your control, but when designing a new room or preparing for a buildout of an office space, you should look for some key aspects in the meeting room location. When you have the option, select a location that is as quiet as possible. Usually, this will be an interior room, with no windows, and away from a main walkway. The less foot traffic that exists outside the room, the less chance of unwanted noise from the office interfering with a call. Locating the room on the interior of the building will keep it away from windows where noise from outside the building can interfere, such as from a siren, car horn, or airplane. Special consideration should be given to the size of the room as well. The larger and more open the room, the more likely that sounds will tend to echo and become distracting. If this room is being custom built, other considerations can be included, such as insulating the walls and providing a solid door without sidelights, which are windows in the sides of the door.

Most often, customizing these aspects of a room is not possible; however, several other modifications can be made to a meeting room to improve the sound quality:

- If there are windows in the room, hanging curtains will help reduce outside noise.
- Replacing ceiling tiles with tiles rated higher for noise reduction will provide a noticeable change, along with acoustical panels that can be hung on the walls.
- Use padded placemats on conference tables to help remove some of the echo in a room.
- The floors should always be carpeted, but if they're not, placing area rugs will also help with echo and reverberation.
- Believe it or not, placing a plant or two in a meeting room not only will liven up the place but also will help with the sound quality. Plus, plants love when you talk to them, so they should be happy in a room built for talking.

Key Topic

Knowledge is power is time is money. Or, if you have watched the TV sitcom *Parks and Recreation*, then "time is money, money is power, power is pizza, and pizza is knowledge." Either way, you cannot dispute the importance of knowledge. When it comes to soundproofing a room, educating the office staff is an essential step. Congregating outside the meeting room should be discouraged. The door should be closed when the room is in use to help quiet external sounds. Notifying others that a meeting is in progress will help keep people mindful of what is happening around them. If the meeting room is in a location where external sounds cannot be eliminated, educating staff to mute the microphone unless speaking will help prevent distractions. Not all calls are from a quiet meeting room with a door. If staff members are using a software endpoint or a phone that supports headphones, the use of headphones and a small microphone can help reduce the external sounds.

Light Behavior

Up to this point, this chapter has covered sound behavior, what audio equipment should be used to improve audio quality, and room remediations that can improve audio quality. In like manner, I will turn the focus of this chapter over to light behavior. The next several sections will examine cameras, display cables, room remediation, and video communication etiquette.

Camera Field of View, Depth of Field, and Zoom

Chapter 3, "Video Basics," delved into the inner workings of how a camera operates. Now we will redirect our attention to the actual image detected and transmitted by the camera, more specifically the lens of the camera. Three primary parameters to be familiar with are field of view, depth of field, and zoom.

Key Topic

The *field of view* is exactly what it sounds like, which is the width and height of the image. Although it can also affect the amount of detail that is visible in the observed image, the field of view is what can be seen through the camera's lens. Some estimates suggest the human eye has a horizontal field of view of about 210 degrees and a vertical field of view of about 150 degrees. Other more conservative estimates put the horizontal and vertical fields of view at 120 degrees. These estimates do not take into account the movement of the eye, which does not change the field of view, but it does quickly change the perspective. The field of view on a camera can be adjusted with zoom or with a different lens if the camera has a fixed focal-length lens. *Focal length* refers to the capability of a lens to magnify the image of a distant subject. The focal length, or size of a camera lens, is the distance from the back of the lens to the frame sensor. Smaller-lens cameras are sometimes referred to as wide-angle because they produce a greater field of view than cameras with a larger lens, despite there being specific lenses for such wide-angle applications. Cameras with larger lenses are conversely referred to as *narrow-angle* as they have a smaller field of view. Calculating field of view is a topic that can go much deeper than is discussed here.

Depth of field refers to the objects, from nearest to farthest, that are in sharp focus. You may have noticed while watching a television show that the camera is focused on someone close up, but a person standing in the background is blurry. This effect is caused by a shallow depth of field. Then the camera will adjust the focus on the person in the background, and everything in the foreground will become blurry. This effect is caused by a deep depth of field. The following three main factors control the depth of field:

Key Topic

- **Aperture:** This factor refers to the amount of light allowed through the lens to the camera sensors. The size of your aperture, which is the diameter of the hole through which light enters the camera, controls the amount of light entering your lens. Using the aperture of your lens is the simplest way to control your depth of field. A larger aperture, or more light, will produce a shallower depth of field. A smaller aperture, or less light, will produce a deeper depth of field.
- **Distance from the object to the camera:** Distance from the object to the camera can also affect the depth of field. The closer an object is to the camera, the shallower the depth of field becomes. Therefore, moving farther away from the object will deepen the depth of field.
- **Focal length of the lens on your camera:** Focal length was discussed previously regarding field of view. This topic can be very complex, but the simple answer is that the longer you set your focal length, the shallower the depth of field.

The third parameter that affects a camera image is *zoom*. The adjustability of a lens that allows the illusion of bringing objects closer or increasing magnification is referred to as zoom. This is very similar to focal-length lenses, except that focal-length lenses are typically fixed, whereas a camera with zoom capability can be adjusted based on the environmental conditions. When you are looking at specifications for a digital camera, both the optical

and digital zoom measurements are listed as a number and an *X*, such as 3X or 10X. A larger number signifies a stronger magnification capability. Cameras are also identified by their focal length, which could be a single number, such as 35 mm, or a range, such as 28 mm to 280 mm. In most cases, a 50 mm lens measurement is considered normal with no magnification and no wide-angle capability. Putting these two measurements together, the multiplier is the difference between the smallest and largest focal-length measurements of the lens. For example, if a 10X optical zoom lens on a digital camera has a minimum focal length of 35 mm, the camera would have a 350 mm maximum focal length. However, if the digital camera offers some additional wide-angle capabilities and has a minimum 28 mm equivalency, then the 10X optical zoom would have a maximum focal length of only 280 mm. Obviously, this can impact both the depth of field and the field of view. Figure 4-9 illustrates the differences between a 50 mm zoom lens and a 200 mm zoom lens.

Figure 4-9 *50 mm Zoom Lens versus 200 mm Zoom Lens*

White Balance

Key Topic

Although the mind cannot always perceive it, objects in different kinds of light are affected with regard to the color temperature in which they are seen. If the light source is cooler, the subject being viewed will appear to have a bluish tint. This effect is common with fluorescent lighting. If the light source is warmer, the subject being viewed will appear to have a reddish tint. This effect is more common with natural lighting, such as the sun. Incandescent lighting can give off a yellowish tent. This effect is not often observed by the human eye because the brain corrects this anomaly; however, without some kind of filter, cameras will pick up the aberration and transmit it to the destination.

White balance is a setting used in cameras to set the reference value for white so that color anomalies caused by color temperature can be corrected. These settings can be either manual or automatic, depending on vendor and model. The manual setting is usually more accurate. One way to set white balance manually is to place a totally white card in front of the camera. When the white balance button is pressed, the camera identifies the color observed as white, so when cool fluorescent lights make an image appear blue, the camera can adjust the image accordingly.

Automatic white balance sets the camera to a known reference value, regardless of room lighting. The caveat is that if the room lighting is different from the camera's reference, the color correction could be inaccurate. In some cases, rooms have mixed-lighting environments with both daylight sources and fluorescent sources in different parts of the room. If a camera pans to different parts of the room, the color values can change dramatically. This is why, in part, it is recommended to locate meeting rooms that use video communication systems in the interior of a building so that there are no windows. If you cannot avoid windows in the meeting room, you can use shades or curtains to block out the natural light from coming into the room during the meeting. Another point of consideration is to use fluorescent lighting with bulbs that produce a temperature of 4000–4100 K.

Lighting Conditions

Lighting enhances the perception of the environment around you, and it can be used to set a mood to an environment. A soft candlelit room is soothing and relaxing as light dances on the walls from the flicker of the candle's flame. Warm reddish lights used in theater environments give audiences the sensation of being at home. Office environments use brighter but cooler lights to accentuate an atmosphere of business. Equal consideration should be pondered when designing a meeting room for video communication. Lighting in meeting room settings is rarely ideal, but the same three-point lighting technique still applies. Three-point lighting consists of key light, fill light, and back light.

Key Topic

Key light in a conference room can be the overhead fluorescent lights, but some track lighting systems allow positioning the lights to better suit the needs of the room. For most office environments, the ideal lighting should be indirect fluorescent lighting. Some vendors recommend a range of 3200 K to 4100 K for the color temperature of the light. Direct light can be harsh, especially while in a video call. On a sunny day, people wear sunglasses to diffuse the light before it enters their eyes. When fluorescent overhead lighting is used for key lighting, you can use diffused lighting systems in meeting rooms where video communication will occur. Some lighting systems already use reflection as a method of making light less direct. The ability to dim the lights with a dimmer switch is another great way to diffuse the light in a meeting room. Also, you should avoid direct down lights when possible. If track lighting is being used, angle the lights away toward the floor or a wall behind the participants.

Fill light, also referred to as soft light, is a very cool light that is directed toward the participants in the meeting. Some video communication systems come equipped with soft facial lighting. The light from the monitor can act as a fill light, as well as reflected light from the meeting room table. For this reason, it is recommended that the wood or surface of the table be a brighter color to allow the light in the room to be reflected back up toward the participants. Darker colors will absorb the light and not provide the fill lighting needed for video calls. However, you should be mindful of surface reflections. Remember that just like sound waves, light waves will be reflective off a hard surface at an angle equal to the incoming angle. Knowing this should make it relatively simple to predict where reflections could be an issue.

Although *back light* is not the most essential lighting needed, it is ideal in rooms that use video endpoints for live video calls. Back light is used to give the participants soft highlights and definition from behind where they're positioned. In this way, the participants are separated from the background, and the lighting gives them a more natural 3D presence during the call. The ultimate goal is to allow the technology surrounding participants to melt away so that they feel as if they are in the same room as the people on the other end of the call. Back lighting takes the experience of a video call one step closer to this goal.

The optimal lighting is a combination of key, soft, and back lighting. Any one type of light on its own is problematic, but when they're combined, the end result is excellent. Key lighting helps illuminate the face, while a soft light helps cut down the contrast, and back lighting can help make the person in the video feel more real and less two-dimensional. Shoulder lighting should not exceed two times the facial lighting, and the lighting should not fluctuate more than 100 lux within the camera's field of view. Figure 4-10 illustrates each of these lighting techniques working independently and together.

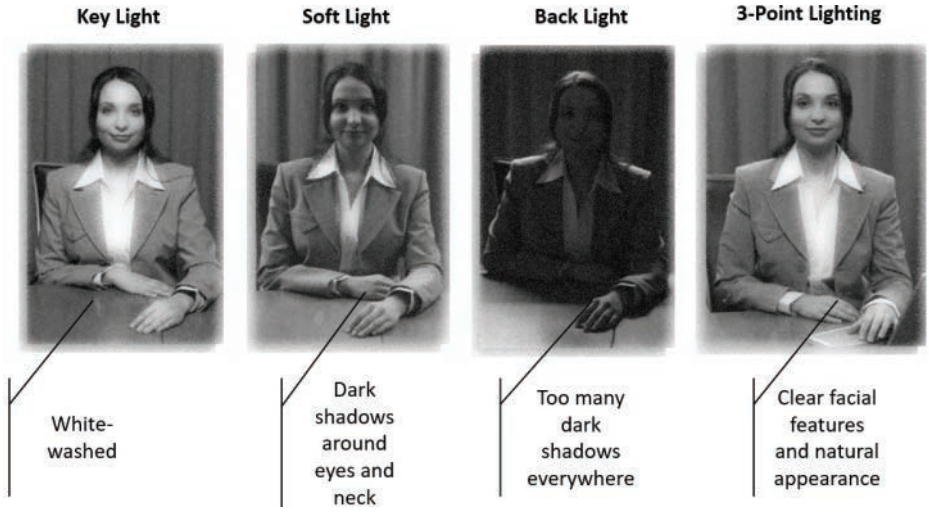

Figure 4-10 *Three-Point Lighting Technique*

Control of light in general is important. Be careful not to oversaturate the subjects you are lighting. Too much light can wash out your subjects on camera, as shown in Figure 4-10 with key light only. Natural light coming through a window is bad for video communication. Rooms with lots of windows are often equipped with shades or blackout curtains to control the amount of exterior light penetrating the room. You also should avoid pointing the camera toward windows with blinds or other hanging treatments that may allow light seepage. Ensure that participants do not sit with their backs toward an uncovered window. The brightness of the background will make the participants seem darkened, and facial features will be harder to make out. The same effect as the back-light-only view in Figure 4-10 will occur. Blinds are not ever recommended in rooms where video communication will take place. Shadows cast by the window frame or blinds can cause a problem as well, throwing off the white balance of the video. Over a video call, those shadows are enhanced, and you could appear to have stripes caused by the shadows. If the meeting room has windows, it is a best practice to cover them with a heavy drape that blocks out the light completely.

Room and Environment Considerations

Beyond lighting, several other considerations should be taken into account when designing a room for video communication. Careful reflection should center around tables and other furniture, as well as camera positioning, such as keeping access points out of the camera's field of view, and properly positioning the camera angle to the participants within the room. Other aspects include backgrounds, wall color, and even carpeting.

Table surfaces should be considered carefully. Light-colored tops are favorable, such as natural maple or walnut wood tops, without stain, and varnished to a nice sheen. You should avoid choosing bold colors and wood grains that present bold patterns. These patterns could have a negative effect with the cameras and cost a lot in bandwidth. An oval-shaped, or racetrack-style, table is usually suggested for rooms with up to six people. Larger rooms can benefit from trapezoidal tables. These table styles reduce the effect of forced perspective and allow everybody in the room to be visible from the camera. Chairs should be low backed to avoid blocking the view of people in the back when the room layout is designed with rows of participants. Other furniture located in the meeting room designed to hold audiovisual equipment should have adequate ventilation. Most electronic equipment produces a lot of heat. Routing cables for table power access, microphones, control pads, and display connectors should be carefully routed to avoid tripping hazards when possible.

There are many aspects to consider when setting up cameras in a meeting room. Cameras can be wall mounted or set on a credenza. There are also mounting brackets designed to hold cameras on top of the display. The following conditions will help determine where to best position cameras in the meeting room. Cameras should be positioned so that they are opposite the entrance doorways, but these entryways should be just outside the cameras' field of view. As mentioned previously, motion should be limited when possible, because excess motion will cost more bandwidth as the video system is forced to refresh more data within each frame. Doors opening or people moving behind a sidelight may cause these unwanted distractions, not to mention the fact that passersby poking their heads in a room while a meeting is in progress can be distracting to participants on both ends of the call. How participants are positioned around a table is important as well. You should avoid room layouts that would place anyone with their backs to the camera. If this position cannot be avoided, consider using dual cameras in a room, with one positioned at each end of the room.

During an endpoint installation, the position of the camera is a critical component. This is especially true on custom installation jobs that use integrator video endpoint kits. The idea is to position the camera and display in association with one another so that participants in the room can maintain eye contact with the participants at the far end of the call. The "eyes" of the participants on the far end of the call will always be the camera, not the display. However, it is very unnatural and uncomfortable to stare into a camera when a video of the person talking is shown on the display. Assuming that the participants in a room are looking at the camera, if the camera is set too high, the participants sitting in front of the camera will have to crane their necks to maintain eye contact with the participants on the far end. If the camera is set too high and participants are looking at the display, it will appear to the far-end participants that everyone in the call is looking down. Alternatively, if the camera is set too low and participants are looking at the display, the far-end participants get the "up-the-nose" view of whomever they are conversing with. Ideally, the monitor should be positioned so that the far-end participants are at eye-contact height with the participants displayed. The camera should be located just above the monitor and centered. This position creates a gaze angle that allows the participants in the room to look the far-end participants in the eye, or at the display, and still maintain eye contact with the far-end participants on the display. Figure 4-11 illustrates how this gaze angle should appear.

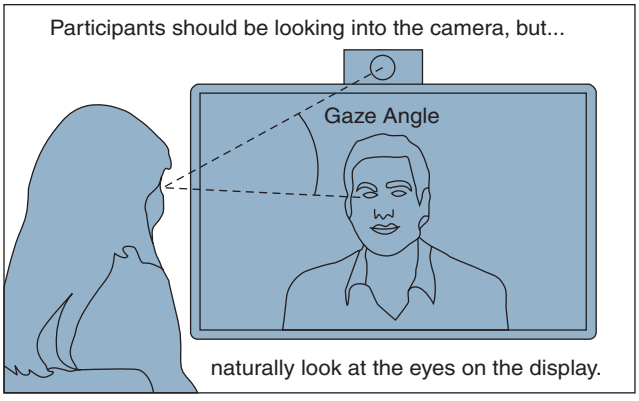

Figure 4-11 *Gaze Angle*

The finer aesthetics of a meeting room can have just as much of an impact on video communication quality as the other aspects of room and environmental conditions that have been discussed thus far. The background seen through the camera's eye should be carefully considered. We've already described how an endpoint will only update the part of a picture that has changed from frame to frame to conserve bandwidth. The more movement within a room, the more bandwidth needed to update the picture. It is best practice to eliminate any unnecessary motion in the background to minimize this issue; however, motion in the eye of a camera can be caused by more than just physical motion. Have you seen a picture such as a black-and-white spiral that appears to be moving even though you know in your mind it is not? This effect happens more frequently with a camera, especially when you're trying to focus on a busy background. Cameras are constantly trying to find a single point to focus on. When there are numerous patterns in the camera's field of view, the aperture of the camera is constantly adjusting to find that single focal point. The result seen on the far-end display is a false motion occurring in the background. Not only is this "motion" very distracting, but video quality is compromised because the video codec requires more bandwidth to update more of each frame. You should avoid a background that has busy wallpaper or many different objects. Pictures are always a nice touch in a room, but they should be hung outside the camera's field of view. However, a simple company logo could be hung on the back wall of a meeting room, as long as it does not have a busy pattern.

Key Topic

Another best practice is to try to minimize contrasting colors. It is preferable to have one solid color as a background, preferably painted with a flat matte or eggshell finish. Equal to the importance of using a solid color, the actual color choice can affect video quality and bandwidth utilization. Cisco recommends using earth tones, such as beige, tan, brown, or forest green. All of these examples are safe colors to use, and they look really good over a video call, but some industrywide studies have been conducted, and the two highest recommended colors are not on Cisco's colors of choice. The color that takes the number-two spot is gray. The number one choice is blue—not just any shade of blue, but a shade that is consistent with Cisco blue. I believe that these colors did not make Cisco's list because most professional people do not want a gray or blue environment to work in. Therefore, Cisco's selections are softer tones that accommodate both video communications and day-to-day operations within an office environment. All industry experts are united on the worst wall color to use for video communications, and that color is white. It's interesting that white is the most widely used color within office environments because it is neutral. The problem

with using white for video communication has to do with the way light and shadows react with white. Studies have shown that white walls use significantly more bandwidth than any other color in the spectrum. Therefore, the short answer to the question of what color a video meeting room should be painted is any color but white.

If you want to reference some of the color palettes Cisco recommends, check out the "Cisco Telepresence Room Design Palettes Quick Reference Guide": https://www.cisco.com/c/dam/en/us/solutions/collateral/collaboration-endpoints/color_ref_guide_c07-642558.pdf.

One final note about room and environment conditions has to do with the very clothes a person wears during a video meeting. Striped, flowery, or any other patterned article of clothing worn by participants in a video call can have the same effect with the camera as a busy background. Although this topic overlaps a little bit into the last section on video etiquette, office employees should be encouraged to wear solid-color clothing on days when they know they will be participating in a video call.

Displays: Monitors and Projectors

Much of this section on light behavior has been dedicated to cameras, how light behaves around cameras, and how to position cameras in a meeting room used for video communication. Now the focus will turn to video output using monitors and projectors. Three different types of monitors can be used in a video communications room: CRT, plasma, and LCD. A fourth option is to use a projection system for display. The Cathode Ray Tube (CRT) option is never used in corporations and rarely seen even in consumer homes anymore. These old televisions were large and heavy, and the standard definition resolution they supported is far inferior to the higher-quality HD resolutions available everywhere today. Therefore, other than the honorable mention here, these types of displays will not be discussed further. As far as the other options are concerned, so many variables are involved in making decisions about projection versus monitors that we could spend a great deal of time on just this one topic. Solutions you might find yourself designing will require you to do your homework before deciding on what would be the best possible solution to use. The paragraphs that follow offer some brief descriptions about the differences between each of these options.

Plasma display panels (PDPs) are a type of flat-panel display that uses small cells containing plasma, which is ionized gas that responds to electric fields. When plasma screens first came out, the clarity of these screens was far superior to anything else on the market, including LCD displays. They range in size from 30 inches (diagonally) up to 150 inches. However, plasma displays had run their course by 2014 due to two main factors. Probably the most prominent factor was the price. Plasma displays were much more expensive than LCD displays. As more people bought LCDs, the prices dropped even more, to the point plasma displays could no longer compete. The second factor that led to the demise of plasma displays was the short life they lived. After about three years of use, plasma displays were prone to burn-in, which is where the images displayed on the screen would leave a permanent ghosting effect. This became a major deterrent away from this type of monitor.

Liquid crystal display (LCD) was a fierce competitor to the plasma display. Though they could not compete with plasma in quality at first, they could compete in price, which gave them an advantage. Plus, their quality was far superior to that of the CRT monitor, which was what the dominant market owned before they upgraded their older systems to this much-improved option. As time progressed, LCD displays began using LED back lighting

to improve the image quality to that of plasma displays, putting the metaphorical nail in the coffin of plasma displays. LED monitors also reduce power consumption and space requirements. LCD displays are now reaching mammoth sizes, and prices continue to drop.

It would seem that LCD displays are the logical choice for video communication systems, but you should not rule out a fourth option just yet. Projection systems can now offer the same great quality of an LCD or plasma screen. Historically, there were three main issues with projection systems. First, more affordable projection systems had a lower quality, and the higher-resolution systems were priced too far out of reach. Second, the lumens on these projection systems were so low that room lights would have to be turned off before the display could be seen. This was not ideal for video communication because no one in the room would be seen if the lights were off. Third, the bulb in the projector had a finite number of burn hours before it went out. Then the cost of replacing the bulb was almost as high as the projector itself. Since those days of old, projectors have come a long way. You can now buy projectors that support full HD resolutions at about the same price as an LCD display. Both forward and rear projection systems are available, and they now use higher lumens to project on a screen in the daylight hours. Improvements have also been made to the bulbs in these projectors so that they last much longer, and they are not nearly as expensive to replace as they used to be. All of these factors make projectors a great alternative to an LCD display.

So, what is the best display to use in your video meeting room? For that answer, you will have to do your homework. Be aware of all the ways a display may be used, all the possible places a presenter might stand, and where viewers might sit or stand. Consider the angles of view, both horizontally and vertically. Look for possible sources of reflected light or ambient brightness on a nice day. Ask yourself a few probing questions to determine which display type is best for the solution you're designing:

- Will the display be on all the time?
- Will it show the same thing over and over?
- What type of content is most likely to be viewed: Excel spreadsheets or live video?

Here are a couple guidelines to help in your meeting room design. Don't use short-throw projectors with roll-up projection screens. The slightest air movement will cause the image on the screen to distort in a way that participants will find disturbing. LCDs are great for spaces with too much ambient light, but be aware of the potential for poor off-axis viewing angles.

Video Cables and Connectors

Just as with audio cables, two types of video cables can be used. Chapter 3 explained luminance, which establishes the brightness, or the light and dark portions of the picture, and chrominance, which establishes all the color information for the picture. Chrominance can be further broken down into the three main colors of the color spectrum, which are red, green, and blue components. These individual components are separated and combined in different manners depending on the connection method. Figure 4-12 illustrates the various composite and component video connectors.

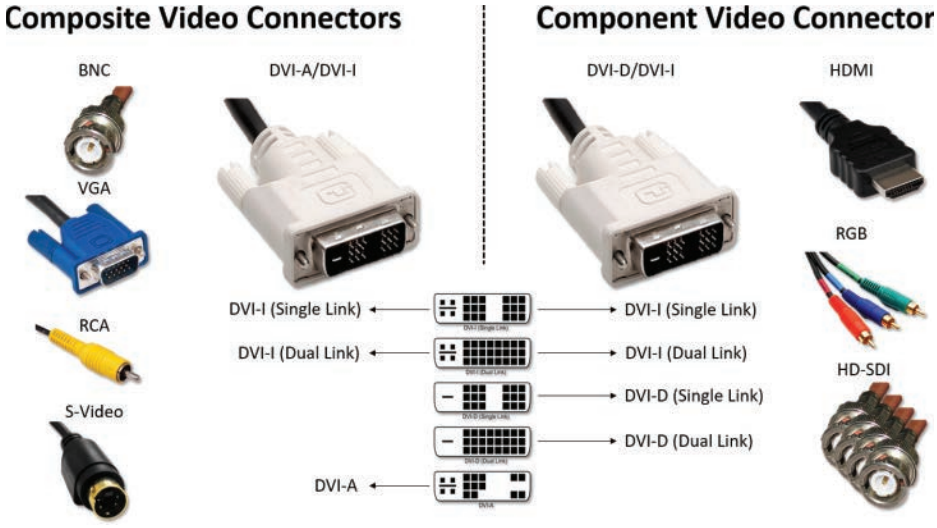

Figure 4-12 *Composite and Component Video Connectors*

Key Topic

Composite video is the simplest and lowest quality of all the varieties. All components that make up the video are combined and transmitted along a single cable. Video Graphics Array (VGA) is the most commonly used composite video connector. Each component is individually separated and given its own pin in the connector. Bayonet Neill-Concelman (BNC) is another composite video connector often found on the ends of coaxial cables. This type of connection is common with cable television, where a single coaxial cable runs into the back of your television set or cable service box. RCA connectors, which were discussed in the "Audio Cables and Connectors" section, often come with three connectors. The yellow connector offers composite video over a single wire, and the red and white connectors offer unbalanced audio for left and right speakers. These cables are commonly used when connecting a VHS or DVD player to a display. Y/C, also known as Separate Video (S-Video), is a connector that separates the Y, or luminance, and C, or chrominance components and transmits these signals individually. However, S-Video is still considered composite video because the chrominance components are not broken out into the three primary colors.

Key Topic

Component video improves on composite video quality by separating each of the color components of the chrominance portion of the signal, in addition to the luminance. Also known as YPbPr, the luminance (L) is integrated with the green portion of the signal. For a recap on how the green is calculated, refer to the "Chrominance and Luminance" section in Chapter 3, which discusses this procedure. RGB connector cable, pictured in Figure 4-12, is a type of RCA cable that will come with three or five connectors on each cable. The three video cables will be colored green for luminance, or Y; blue for primary blue, or Pb; and red for primary red, or Pr. Although BNC connectors can be used for a composite video connection, component video can also be achieved with three coax cables. This type of connection is frequently used in commercial applications with High Definition-Serial Digital Interface (HD-SDI) video formats. The most widely used component video connection globally is High-Definition Multimedia Interface (HDMI). HDMI is a fully digital standard that also has the added benefit of carrying balanced audio, power, and control on the same cable. HDMI uses YCbCr to calculate the color space for video, which allows a higher color depth. HDMI

has evolved from version 1.0 to the recent update 2.1 standard, and now has a maximum supported resolution of 10k at 120 Hz.

The last video connector worth mentioning is Digital Visual Interface (DVI). I saved this particular connector till last because five different DVI connector types can support either analog composite video or digital component video. The Digital Visual Interface-Analog (DVI-A) connector is similar to the VGA connector. These connectors are not often found today due to the mass adoption of HD. More commonly found is the Digital Visual Interface-Digital (DVI-D) connectors. These digital-only connectors support HD resolutions comparable to HDMI. Digital Visual Interface-Integrated (DVI-I) allows either a DVI-A or DVI-D connector to join, thus supporting either analog or digital signals. There are two different connector types for DVI-D and DVI-I:

- Single-link DVI employs a single 165 MHz transmitter that supports resolutions up to 1920 × 1200 at 60 Hz, or 2560 × 1600 at 30 Hz.
- Dual-link DVI adds six pins, at the center of the connector, for a second transmitter increasing the bandwidth and supporting resolutions up to 2560 × 1600 at 60 Hz, or 3,840 × 2,400 at 30 Hz.

As previously explained, the five connectors in the DVI category include DVI-A, DVI-D Single Link, DVI-D Dual Link, DVI-I Single Link, and DVI-I Dual Link.

Immersive Telepresence

Immersive Telepresence room systems represent the height of video communication today. They combine the best audio, video, and networking components available on the market to create an in-room experience that's as close to a face-to-face meeting as is possible with modern technology. Immersive Telepresence rooms of the past required strict implementation guidelines to create the environment to have a high-quality experience and allow the system to function at optimal levels. The legacy Tandberg T3 Immersive Telepresence room was quite literally a "room within a room" configuration that required the integrator to fully install the walls, ceiling, lighting, and flooring within an existing room to adequately control the conditions required to provide an exceptional video experience. In comparison, the legacy Cisco CTS 3000 was a standalone system that was installed into an already-completed room, but it still required extensive room remediation with proper lighting, temperatures, wall colors, and sound dampening in place before installation of the endpoint system and components was allowed.

Key Topic

The Cisco IX5000 was a game-changing, revolutionary Immersive Telepresence system that changed what was considered essential for room remediation with these types of systems. It is the first Immersive Video endpoint to require no room remediation prior to the installation, although some room remediation is still recommended. There are still recommendations and best practices that surround video communication, as mentioned earlier in the "Room and Environment Considerations" section of this chapter. For example, Cisco recommends lighting temperatures of 4000 K to 4100 K and a color rendering index of 82 or greater for best results with any telepresence endpoint. Also, the top-down lighting onto the subject's shoulders should not exceed two times the light value of the lighting that can be measured on the subject's face. The IX5000 was designed to operate with much wider room tolerances while still being able to control the customer experience. The IX5000 addresses the preceding two examples with built-in facial lighting elements mounted directly over the display

monitors. This extra facial lighting, which can be adjusted manually, allows the illumination on the participants to be better balanced based on each room's needs.

In October 2019, the IX5000 endpoint was end-of-sale, marking the dawn of the next phase of Immersive Telepresence solutions. Cisco has launched a new endpoint called the Cisco Webex Room Panorama. This system now supports three simultaneous video streams using a single codec, just like the IX5000, plus two additional content video streams at the same time. What's more, this integrator system allows the room to be customized to the specific needs of the customer, taking Immersive Telepresence one step further. This endpoint is offered at a significantly lower price. Plus, all the furniture that was required with the IX5000 is no longer required with the Webex Room Panorama, and installation is much easier as well.

Video Etiquette

People communicate with more than just voices. One of the great advantages of video communication over audio-only communication is the ability to read someone's body language, see emotional responses, and share content. Communicating this way also opens the gate for other distractions. As our world and workplaces become more enriched with video solutions, we need to keep in mind basic etiquette for using video to communicate. Some of the etiquette suggestions that follow should already be familiar from more than a century's use of telephones, although others will be unique to video.

Prior to a scheduled meeting, you should try a test call to make sure there will be no connection issues at the scheduled time of the call. Enter the meeting room early to make sure it is clean and the chairs are all pushed under the table. Connect to the meeting early if you can and turn on self-view to see how the room and participants will appear to the far end of the call. Adjust the camera, if needed, so that all participants in the room can be seen clearly.

When using video communication devices, be aware of what the far-end participants are hearing from your end. This is especially true when tabletop microphones are used. Clicking the end of your pen, tapping on the table, or rustling papers can be very distracting because these noises are magnified over the communication systems. Placing papers surreptitiously over the microphone can have a reverse effect and muffle the audio of people speaking within the room. Most people have the common decency today to set their cellphones to vibrate during a meeting. However, placing mobile phones on a meeting room table will create vibrations that the mic will pick up when they inevitably go off. The general rule of thumb is to try eliminating external sounds as much as possible. Side conversations should be discouraged as well. A best practice during an audio or video meeting is to mute the audio altogether, but always assume the audio feed is active. You do not want to be caught complaining about your boss on the other end of the call because you thought the microphone was muted. This also goes for the beginning and end of a call. Audio channels always open before video channels, and audio channels are always the last to close. That means that you may be sending audio between endpoints even though you can't see video.

What you are wearing may affect the way people see you in more than one way when it comes to video communication. Stripes, plaids, and other multicolored designs can appear more pixelated on video or cause a false sense of motion. This effect can be distracting to people at the other end of the call, and their focus may be on the clothing rather than the person speaking. It is preferable to stick with solid colors. This suggestion should also be considered for the purpose of the background of a video call. Just as you would in a face-to-face

conversation with someone, try to maintain eye contact. If the meeting room has been designed properly, doing so should be easy enough. Remember that the camera is the eyes of the participants at the far end of the call. Last to be mentioned, but certainly not the least to consider, try to avoid other distracting actions. Avoid eating food or chewing gum while using an endpoint, and don't swivel in your chair. Always try to maintain a professional presence for the duration of the call. There will be plenty of time to goof off when the call ends.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 4-3 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 4-3 Key Topics for Chapter 4

Key Topic Element	Description	Page Number
Paragraph	Physical Components of an Endpoint	55
Paragraph	How Reflections Affect Hearing	58
Paragraph	Dynamic Mics versus Condenser Mics	60
Table 4-2	Microphone Pickup Patterns	62
Paragraph	Balanced versus Unbalanced Audio Cables	64
Paragraph	AEC Operation	66
Paragraph	Microphone Placement	67
Paragraph	Speaker Placement	67
Paragraph	Audio Communication Etiquette	69
Paragraph	Field of View and Focal Length	70
List	Factors Controlling Depth of Field	70
Paragraph	White Balance for Color Temperature Correction	71
Section	Three-Point Light Technique	72
Paragraph	Gaze Angle	74
Paragraph	Wall Colors in Video Meeting Rooms	75
Paragraph	LCD and Projector Displays	76
Paragraph	Composite Video	78
Paragraph	Component Video	78
Paragraph	Immersive Telepresence Lighting Requirements	79

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Absorption, ACG, AEC, Aperture, Balanced Audio Cable, Bidirectional, Cardioid, Clipping, Codec, Coherent Late Reflections, Condenser Mics, Critical Distance, CRT, Depth of Field, Diffusion, Direct Sound, Directional Mics, Distance Factor, Dynamic Range, Early Reflections, Endpoint, Field of View, Frequency Response, Gain, Gaze Angle, Good Levels, Headroom, Hypercardioid, LCD, Line Level, Mic Level, NRC Rating, Omnidirectional Mic, PDL, PoE, Polar Pattern, PTZ, Reflection, Reverb Time, Reverberations, Shotgun, Supercardioid, Unbalanced Audio Cable, White Balance, Zoom

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. What are the six components that make up a video system?
2. What are the five functions of the menu settings on an endpoint?
3. What are the four categories of reflections?
4. List the three positions used in the three-point lighting technique.

This page intentionally left blank

CHAPTER 5

Communication Protocols

This chapter covers the following topics:

PSTN Communication: This topic will discuss the various PSTN connections options, including POTS, ISDN via BRI, and ISDN via PRI.

H.323 Communication: This topic will discuss the H.323 umbrella standard, H.323 gatekeeper registration, call flow without a gatekeeper, and call flow with a gatekeeper.

SIP Communication: This topic will cover the IETF SIP, discuss basic SIP registration, and address SIP call setup using early offer and delayed offer.

NAT and Firewall Traversal Solutions: This topic will discuss the different IETF NAT traversal protocols, including STUN, TURN, and ICE. This topic will also discuss the Cisco proprietary traversal protocol, ASSENT, and the ITU standard for traversal, H.460.

Up to this point, this book has focused on the intricate aspects of passing audio and video media between two nodes in an audio or video call. This chapter will turn the focus to the actual medium used to transport this data. Two mediums can be used for communication. Circuit-switching technology is the standard telephony mode of communication that has been used for over a century across the PSTN. Packet-switching is a relatively new technology that allows voice and video communication to travel across the digital network. Both modes of communication are incredibly complex, so this chapter will serve only as an introduction to each of these technologies. Topics discussed in this chapter include the following:

- PSTN Communication
- H.323 Communication
  - H.323 Gatekeeper Registration
  - H.323 Call Flow without a Gatekeeper
  - H.323 Call Flow with a Gatekeeper
- SIP Communication
  - SIP Registration
  - SIP Call Setup
  - Delayed Offer
  - Early Offer
- NAT and Firewall Traversal Solutions
  - STUN
  - TURN

- ICE
- ASSENT and H.460

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

1.3 Configure these network components to support Cisco Collaboration solutions:

- 1.3.a DHCP
- 1.3.c CDP
- 1.3.d LLDP
- 1.3.e LDAP
- 1.3.f TFTP

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 5-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 5-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
PSTN Communication	1–2
H.323 Communication	3–6
SIP Communication	7–10
NAT and Firewall Traversal Solutions	11–14

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. How many channels form an E1 link?
   a. 23
   b. 24
   c. 31
   d. 32

2. Which of the following control mechanisms uses IMUX to multiplex ISDN channels together?
   a. H.221
   b. H.460.19
   c. BONDING
   d. Assent
3. Which of the following are the minimum required standards for H.323 compliance? (Choose four.)
   a. G.711
   b. G.729
   c. G.722
   d. H.239
   e. H.261
   f. H.264
   g. H.221
   h. H.225
   i. H.245
   j. T.120
4. Which of the following RAS messages is used to broadcast communication for gatekeeper discovery?
   a. ARQ
   b. GRQ
   c. BRQ
   d. RRQ
5. Which of the following RAS messages is used to initiate an H.323 call when no gatekeeper is used?
   a. GRQ
   b. RRQ
   c. ARQ
   d. None of the above
6. Which of the following RAS messages is used to initiate an H.323 call through a gatekeeper?
   a. GRQ
   b. RRQ
   c. ARQ
   d. None of the above
7. Which of the following organizations is responsible for the management of SIP?
   a. IEEE
   b. IETF
   c. ITU-T
   d. ICANN

8. Once an endpoint has powered on and loaded the locally stored image, what is the next step for the endpoint to register to a CUCM?
   a. DHCP Discovery
   b. Send TFTP Get
   c. Send CDP
   d. Register to CUCM
9. Which protocol is responsible for exchanging capabilities during a SIP call setup?
   a. SIP
   b. SDP
   c. H.225
   d. H.245
10. In the call setup process, what SIP response does the destination endpoint use to send SDP information to the source endpoint?
    a. Invite
    b. Trying
    c. Ringing
    d. OK
11. In what private IP address class is the address 172.20.198.18 included?
    a. This is a public IP address, not a private IP address.
    b. Class A addresses
    c. Class B addresses
    d. Class C addresses
12. Which of the following network characteristics requires a STUN server to be used?
    a. Asymmetric network
    b. Symmetric network
    c. Both symmetric and asymmetric networks
    d. Neither symmetric nor asymmetric networks
13. What is the default listening port for TURN servers running on the Cisco Expressway?
    a. 5060
    b. 5061
    c. 3478
    d. 24000–29999
14. How many media ports need to be open for Assent to pass audio and video data through the firewall?
    a. 1
    b. 2
    c. 2400
    d. 30,000

Foundation Topics

PSTN Communication

The International Telecommunication Union, or ITU, is part of the United Nations. The purpose of the ITU is to help coordinate the use of information and communication technologies. Specifically, it coordinates global usage of the radio spectrum, facilitates cooperation between governments and corporations with assignment of satellite orbits, and most importantly for the purpose of this book, helps in the interconnection of networks and technology.

In particular, the sector most pertinent to this book is the ITU Telecommunication Standardization Sector, or ITU-T. The ITU-T is tasked with defining standards for any audio-visual communications on a circuit-switched or packet-switched network. Without a uniform standard, problems arise between systems that try to communicate but are using different formats to store and transmit data. Technically speaking, the main products of ITU-T are recommendations (ITU-T Recs)—standards defining how telecommunication networks operate and interwork. ITU-T Recs have nonmandatory status until they are adopted in national laws. The level of compliance is nonetheless high due to international applicability and the high quality guaranteed by ITU-T's secretariat and members from the world's foremost information and communication technology (ICT) companies and global administrations.

The ITU-T has created umbrella standards, which contain categorized lists of codecs that must be used to be able to communicate within a particular switching technology. The H.320 umbrella standard encompasses circuit-switched technologies, and the H.323 umbrella standard encompasses IP-based packet-switched technologies. Because the ITU-T has standardized how devices communicate, an H.320-supported device can communicate to any other H.320-supported device because they are all required to support the same minimum standards to be H.320 compliant. This is also true for H.323 devices. Defining what audio, video, data, and control codecs are being used in the conversation is crucial for devices being able to communicate.

Circuit switching is a method of implementing a telecommunications network in which two network nodes establish a dedicated communications channel, or circuit, through the network before the nodes may communicate. The circuit guarantees the full bandwidth of the channel and remains connected for the duration of the communication session. The circuit-switched network was first implemented in the old analog telephone network. When a call is placed from one telephone to another, switches within the telephone exchanges create a continuous wire circuit between the two telephones for the entire duration of the call across the public switched telephone network, or PSTN.

For call setup and control, it is possible to use independent dedicated signaling channels from the phone to the network. Integrated Services Digital Network, or ISDN, is a type of circuit-switched networking service that uses an independent signaling channel while plain old telephone service, or POTS, does not. The method of transmitting and receiving independent signals over a common signal path by means of synchronized switches at each end of the transmission line so that each signal appears on the line only a fraction of time in an alternating pattern is known as time-division multiplexing, or TDM. It is used when the bit rate of the transmission medium exceeds that of the signal to be transmitted. This form of signal multiplexing was developed in telecommunications for telegraphy systems in the late 19th century but found its most common application in digital telephony in the second half of the 20th century.

Key Topic

The two main types of ISDN to be familiar with are BRI and PRI. BRI, which stands for Basic Rate Interface, is primarily used for subscriber circuits similar to the same voice-grade telephony service that has been used for the last century. This allows BRI ISDN connections to use the same existing telephony exchange infrastructure at business locations. Each BRI ISDN circuit contains two Bearer channels, or B-channels, and one Delta channel, or D-channel. BRI Bearer channels support 64 kbps bandwidth and are used to carry the actual audio data. The BRI Delta channel supports 16 kbps and is used to send all the control signaling, such as call setup messaging, call teardown messaging, and timing for TMD. The Delta channel carries the address, informational, and supervisory signaling messages for all of the B-channels. By contrast, PRI, or Primary Rate Interface, circuits offer more B-channels.

Key Topic

PRI is primarily used for carrying multiple DS0 transmissions. PRI is the main standard used today for providing telecommunication services to businesses over a circuit-switched network. There are two different types of carriers for PRI ISDN services available globally. Throughout the United States, Canada, and Japan, the T-carrier is used, although Japan refers to its carrier as J-carrier. The T1 or J1 PRI ISDN circuit will carry 23 Bearer channels and one Delta channel. Although the purpose of these channels is the same as with BRI, they all support 64 kbps, including the D-channel. Multiple channels can be multiplexed together so that one T1 circuit can provide up to 1.544 Mbps of bandwidth for ISDN communication. Common throughout Europe, Australia, South America, and pretty much the rest of the world is the E-carrier. One E1 circuit consists of 30 Bearer channels and two Delta channels. Much like the T1 PRI circuit, each of these channels supports 64 kbps, including the two D-channels. Therefore, one E1 line can provide up to 2.048 Mbps of bandwidth for ISDN communication. Both T1 and E1 PRI ISDN lines can be leased from a telephone service provider, or TSP, with all channels available, or only a part of the channels available. These are called Fractional T1 or Fractional E1, and sometimes require special treatment when they are set up.

Audio communication has been available through ISDN since the H.320 standard was introduced in 1986. A single audio call required only 64 kbps bandwidth or less, so a single channel could easily be used to support an audio call. When video conferencing was first introduced, the medium was based on ISDN video endpoints that were connected via BRI or PRI circuits to the PSTN. Three BRI connections were common for low-resolution endpoints that supported up to 320 kbps for video and 64 kbps for audio. The common video format used was CIF. Higher-resolution endpoints could use fractional PRIs or a full PRI for up to 1.554 Mbps with T1 PRI and up to 2 Mbps with E1 PRI; however, this created a problem. The H.221 control standard, also referred to as Clear Channel Dialing, required the E.164 alias associated with each line to be dialed before the call could be placed. Therefore, to place a 384 Kbps call, up to six phone numbers had to be dialed at the same time.

An organization known as the Bandwidth On Demand Interoperability Group came up with a solution to this issue called BONDING. If two ISDN phones had a piece of software installed on them called an IMUX, then only one number needed to be dialed to initiate the call. The IMUX on the receiving phone would send a request for more channels to be opened. The IMUX on the source phone would proceed to place the remaining calls needed to support the bandwidth in the original call request. When all channels were opened, IMUX would allow the call setup to continue. This multiplexing of B-channels within ISDN based

on BONDING was so effective that the ITU-T adopted BONDING into the H.320 umbrella standard. Figure 5-1 illustrates the difference between Clear Channel Dialing and BONDING.

The still-lingering issue with ISDN-based calling using video communication was that many calls were long distance or international. Each channel that was opened had a toll charge associated with it. This resulted in high PSTN costs, even for CIF-based video communication calls between two locations. Therefore, video communication was not widely accepted into the business world until the less-expensive IP-based packet-switched communication was available with the premiere of high-speed Internet.

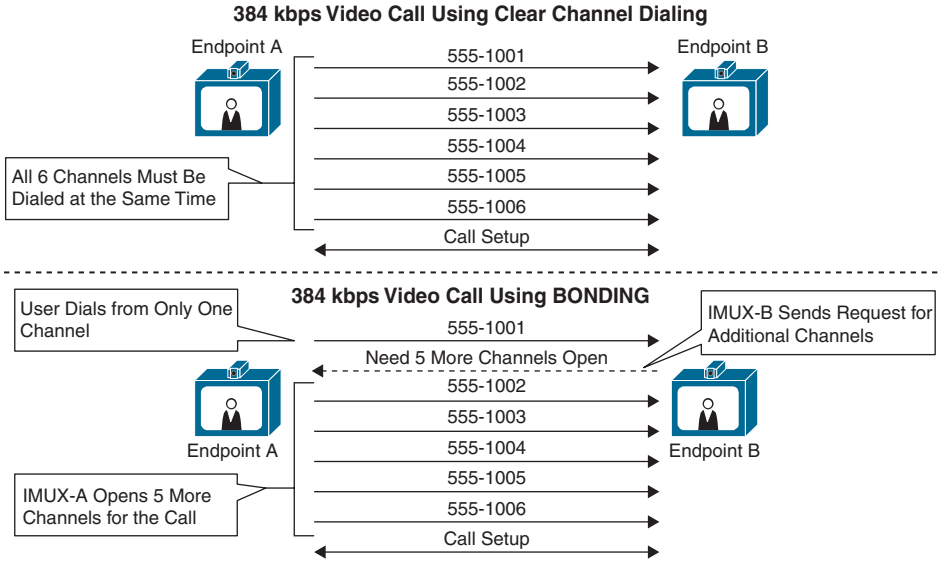

Figure 5-1 *Clear Channel Dialing vs. BONDING*

To be able to communicate with other devices that are using H.320, each endpoint must support certain minimum standards. Endpoints that support these minimum standards are referred to as being *standards compliant*. The minimum audio codec to support is G.711, but many other audio codecs can be supported as well, as outlined in Chapter 2, "Audio Basics." The minimum video codec that must be supported is H.261, although more video codecs are available as well. Refer to Chapter 3, "Video Basics," for more information on video codecs. The minimum control codec, which is the codec that controls call setup and call teardown and maintains the timing during the call, is H.221. A substandard under the H.221 standard is known as Q.931. This is the mechanism known for the initial handshake during call setup and for terminating the channels during call teardown. BONDING is also a control codec and is more commonly used today for video calls. These are the minimum codecs that must be supported to be H.320 compliant. However, other codecs used for other purposes are available under the H.320 umbrella standard, such as T.120 or H.239, which are both used for content sharing. Other codecs support features such as chair control during conferences, mute, call hold, call transfer, far-end camera control, and many more. Table 5-2 identifies the minimum standards for H.320 compliance.

Key Topic

Table 5-2 Minimum Standards for H.320 Compliance

Capability	Codec
Audio	G.711
Video	H.261
Data Sharing	T.120
Control	H.221

Private branch exchanges (PBXs), which were discussed in Chapter 1, "Introduction to Collaboration," were often used within enterprise environments. PBXs performed local call control for phones. They also allowed different enterprise locations to be connected via a leased service connection. This type of connection was usually less expensive than a connection through a telephony service provider via ISDN lines.

H.323 Communication

Key Topic

H.323 is an umbrella standard developed and maintained by the ITU-T for IP-based packet-switched communication. It is referred to as an *umbrella* standard because it encompasses many substandards. To be H.323 compliant, the device must support a minimum set of substandards, such as G.711 for audio, H.261 for video, and H.225 and H.245 for control. Many of the standards included in H.323 are taken from the H.320 umbrella standard for circuit-switched communication, such as Q.931, which is a substandard under H.225, all of the audio and video codecs, and the H.239 codec for content sharing.

H.323 devices can place calls without a central call control device by dialing the IP address of another system. However, a central call control device for H.323 endpoints, which is known as a *gatekeeper*, can be used to extend the management functions of an H.323 environment in three capacities: registration, security, and call control.

Key Topic

Registration allows the use of E.164 aliases, H.323 IDs, and prefixes in addition to IP address dialing. E.164 aliases are numeric-only values containing 1–15 digits that are assigned to an endpoint. They work in the same manner as any phone number would in a typical telephony environment. H.323 IDs can use any combination of numbers, letters, and/or special characters, but spaces are not allowed. Because of this ability, an H.323 ID can be in the form of a URI. However, an H.323 ID is not a URI because it is not dependent on the domain being a fully qualified domain name (FQDN). Prefixes are a feature of H.323 dial plan architecture that allows easy access to services such as Multipoint Conferencing Units (MCUs) and gateways.

Security regarding the gatekeeper does not refer to call encryption. That type of security is a function built into endpoints and can be used with or without a gatekeeper. Security within the context of a gatekeeper refers to the capability to determine which devices can and cannot register to the gatekeeper. This function helps secure access and control features within the communications network.

Call control is the ability to determine which devices can call each other and administer the bandwidth that can be utilized when calling between different devices and locations. More advanced call control policies can be configured, such as transforming the dialed aliases, redirecting call traffic, and restricting calls through ISDN gateways from an outside source, which is known as *hair-pinning*.

H.323 Gatekeeper Registration

Call Setup mode is an H.323 setting configured on an endpoint that can be set to either Direct or Gatekeeper. If Call Setup mode is set to Direct, the endpoint will never attempt to register to a gatekeeper and will only be able to dial by IP address. When Gatekeeper mode is used, the endpoint is completely subservient to a gatekeeper and will perform no function until it has registered. Once registration has occurred, the device must request permission from the gatekeeper before it will attempt any action including placing or answering calls. The ITU created the communication protocol known as RAS (Registration, Admission, and Status) that identifies all messaging schemes between any device and a gatekeeper. RAS is a substandard under the H.225 standard for call setup and is used only by endpoints that register to the gatekeeper.

Another setting on an endpoint that affects its registration is known as Discovery mode. Discovery mode determines how an endpoint will locate the gatekeeper to which it will attempt to register. Discovery mode can be configured to either Automatic or Manual. If Discovery mode is set to Automatic, the endpoint will broadcast a message that is known as a Gatekeeper ReQuest (GRQ) to the entire network broadcast domain to which the device belongs. The first gatekeeper to respond with a Gatekeeper ConFirm (GCF) is where that endpoint will send a Registration ReQuest (RRQ). If the Discovery mode is set to Manual, then the address of the gatekeeper must be entered into the device. This address becomes the address that the endpoint will direct the RRQ. The gatekeeper that receives the RRQ will respond with a Request In Progress (RIP). This message allows the gatekeeper to process though various security policies to assess whether the endpoint is allowed to register. If a security policy prohibits registration, the gatekeeper will respond to the device with a Registration ReJect (RRJ). If there are no configurations prohibiting the registration, the gatekeeper will respond with a Registration ConFirm (RCF). Figure 5-2 illustrates the H.323 registration process.

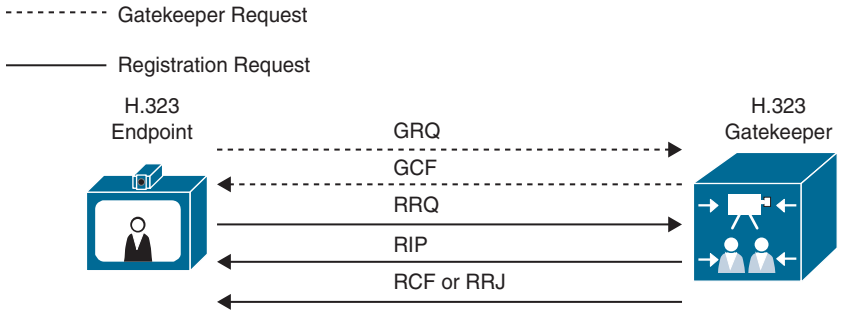

Figure 5-2 *H.323 Registration Process*

H.323 Call Flow without a Gatekeeper

Both H.323 and SIP use the three-way handshake method of the TCP/IP network for initial call setup. If you are not familiar with this method, this brief review should explain the purpose and process of the method appropriately. Assume that a client, such as a computer, is going to transfer files to an FTP server. It is vital that the information being shared is sent securely and completely, without any loss of data. Therefore, before the client will send any of the important information, a three-way handshake is used to establish a TCP connection with the server. The client initiates this process by sending a SYN message to the server. The server sends back a SYN/ACK message to the client acknowledging the receipt

of the SYN, and then the client sends an ACK message to the server establishing a TCP connection. Now the client can begin sending the file packets to the server and will receive an acknowledgment on each packet sent as they are received. If any packets sent by the client do not receive an acknowledgment in return, the client knows to resend those packets. This process is the same for IP telephony calling. A TCP connection must be established between two endpoints before important call capability information is exchanged. In the IP telephony world, this three-way handshake is known as *call setup*. Figure 5-3 illustrates the three-way handshake method.

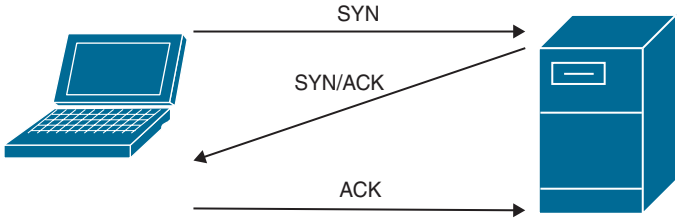

Figure 5-3 *Three-Way Handshake Method*

Key Topic

As mentioned before, H.323 devices can place calls without a central call control device by dialing an IP address of another system. The endpoint will not register to a gatekeeper, Call Setup mode should be set to Direct, and the endpoint will dial by IP address. The source endpoint will send a Q.931 call setup message to the destination endpoint. Q.931 contains the source and destination IP address, in hexadecimal format, and any crypto-hash token if the call is to be encrypted. Both Q.931 and RAS messaging are part of the H.225 control standard for call setup. Because a gatekeeper is not being used in this scenario, there is no RAS messaging either. H.225 performs the same function as the SYN message in a three-way handshake. The destination endpoint can now respond to the Q.931 call setup message it received. The first response this endpoint will send out is the Alerting message. This process allows the destination endpoint to ring and sends a ring-back tone to the source endpoint. The Alerting message is equivalent to the SYN/ACK message in a three-way handshake. Once the call is answered, either manually or automatically, a Connect message is sent to the source endpoint. The Connect message is equivalent to the ACK message in a three-way handshake. Both the Alerting and Connect messages are part of the Q.931 messaging system.

After the call setup messaging is complete, the two endpoints must now go through the H.245 negotiation process to ensure that the most appropriate codecs are selected and to identify the UDP ports that will be used for communications. This is the important data that the three-way handshake is used for to ensure all data is exchanged properly between two endpoints. First, each endpoint must send a Terminal Capabilities Set, sometimes called Capabilities Exchange or CapEx, containing all the codecs that each endpoint is capable of using. These codecs are the audio and video compression algorithms, such as G.711 and H.261, and other capabilities like dual video, far-end camera control, chair control, and so on. Because this communication is all TCP communication, there will be acknowledgments for each communication sent. Once this process is completed, the master/slave determination needs to be made. The master is simply the endpoint that will select the codecs to be used and allocate the UDP ports for RTP and RTCP communications. Based on several criteria, it can be either endpoint. Most often the endpoint that initiated the call will be the master, but that is not always true. In some instances, the endpoint with the most capabilities or the highest capabilities will be

the master. In every call that involves a bridge, such as the Cisco Meeting Server, the bridge will be the master. Once the master is selected, that endpoint will initiate the next step in the process, opening logical channels. Obviously, actual channels are not being used for an IP-based call. This terminology was taken from the H.320 umbrella standard. The logical channels referred to here are the UDP ports that will be used for sending the media packets and signaling between each endpoint. Audio ports will always be opened first, then video, then ports for each additional capability. The master will also open the Real-time Transport Control Protocol (RTCP) ports used for signaling first, which will always be odd-numbered. Then the Real-time Transport Protocol (RTP) ports will open for the actual media, and they will always be even-numbered. Once each set of ports has been opened, the endpoints will begin to exchange media and the call will be set up. Figure 5-4 illustrates the H.323 call setup process without a gatekeeper.

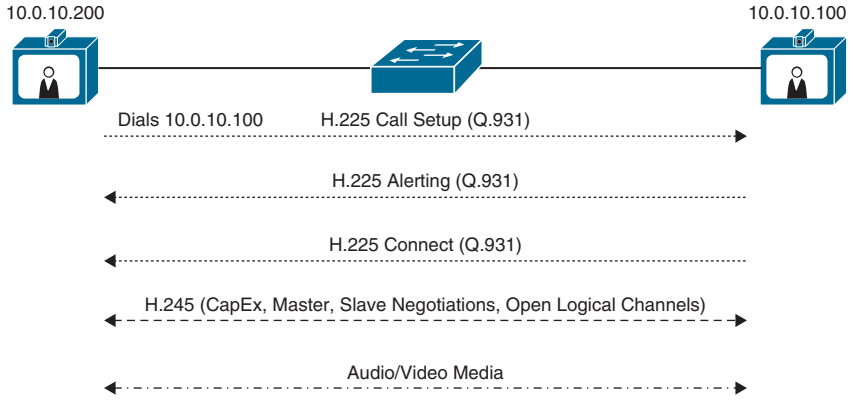

Figure 5-4 *H.323 Call Setup Process Without a Gatekeeper*

H.323 Call Flow with a Gatekeeper

Before I explain H.323 call flow with a gatekeeper, it is important to understand two methods that exist for H.323 call setup. These two methods are called *slow start* and *fast start*. In most cases on a Cisco Expressway running as an H.323 gatekeeper, slow start is used. Therefore, the explanations here are based on this method of call setup. Fast start simply combines some of the H.245 functions into the H.225 call setup process. Fast start goes beyond the scope of this book, but it is worth mentioning because it is still seen in Cisco Unified Communications Manager environments when H.323 gateways are being used.

Key Topic

Endpoints that are registered to a gatekeeper follow the same calling process as endpoints that do not register to a gatekeeper, but some extra messaging must take place between the endpoint and the gatekeeper first. This is where RAS messaging comes into play. RAS is commonly referred to as Registration, Admission, and Status. During an H.323 call setup involving a gatekeeper, a call admission process must occur. When the source endpoint dials the alias of the destination endpoint, it sends an Admission ReQuest (ARQ) to the gatekeeper. The gatekeeper responds with a RIP message because there may be call control policies that prohibit or restrict the call or certain aspects to the call. If the policies restrict access to the dialed alias, or the gatekeeper cannot locate the dialed alias, the return message will be an Admission ReJect (ARJ). Provided the destination endpoint is located and there are no restrictions prohibiting the call, the return message will be an Admission ConFirm (ACF),

which also includes the IP address of the destination endpoint and the bandwidth that is allowed for this call.

The source endpoint then uses the preceding information to send a Q.931 call setup message to the destination endpoint. When the destination endpoint receives the call setup message, it must first request permission to answer the call from the gatekeeper. Therefore, the destination endpoint sends an ARQ message to the gatekeeper. The gatekeeper responds with a RIP message and proceeds to check whether there are any bandwidth restrictions. Provided there are no restrictions, the gatekeeper responds with an ACF. The destination endpoint can now respond to the Q.931 call setup message it received with the same Alerting and Connect message discussed previously.

After the call setup messaging is complete, the two endpoints must now go through the H.245 negotiation process to ensure that the most appropriate codecs are selected and to identify the UDP port that will be used for communications. As mentioned before, the communication sent using H.245 includes a terminal capabilities set, master/slave determination, and opening logical channels, or ports. Once each set of ports has been opened, the endpoints will begin to exchange media and the call will be set up. Figure 5-5 illustrates the H.323 call setup process with a gatekeeper.

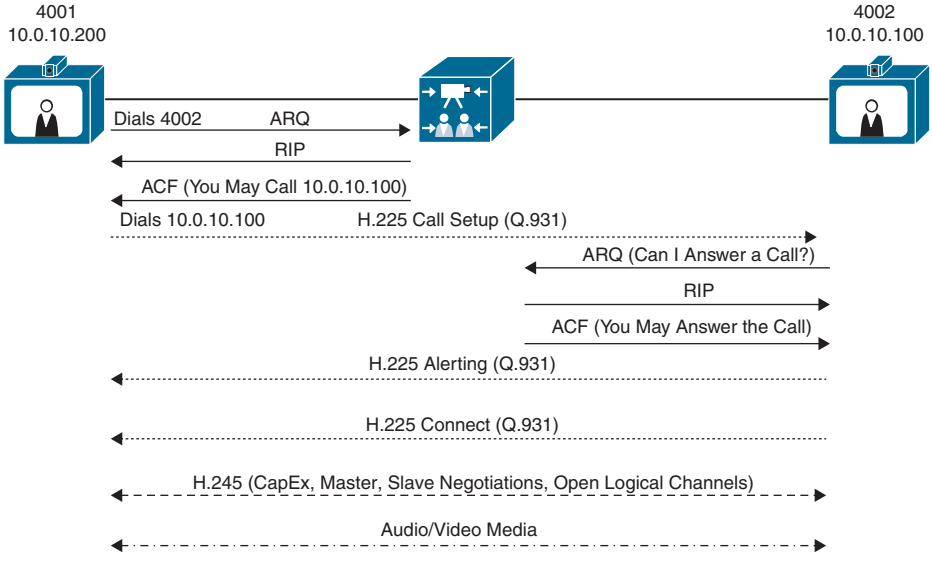

Figure 5-5 *H.323 Call Setup Process with a Gatekeeper*

SIP Communication

Key Topic

Session Initiation Protocol, or SIP, is a communication protocol that is created and managed by the Internet Engineering Task Force, or IETF. Unlike H.323, SIP endpoints cannot operate without a call control device. The SIP call control device is generically called a SIP server and functions at the center of all SIP-related call control activities. The SIP server performs two functions: Registrar and Proxy. The SIP Registrar maintains a table that is composed of IP addresses and Uniform Resource Identifier (URI) addresses taken from devices at the time of registration. URI addresses are the only type of alias SIP uses for communication other than IP addresses. Because a URI must be in the format Host@FQDN (fully qualified domain

name), the domain must exist within the SIP Registrar for successful registration to occur. SIP addresses can also be in the form of Host@IP_Address. The SIP Proxy helps in locating and connecting two devices at the time of call setup. Cisco has two devices that fulfill the role of SIP server: the Cisco Unified Communications Manager and the Expressway.

When SIP was originally designed, the intention was not focused specifically on voice and video communication over IP. In fact, it would be a surprise if this use was even conceived at the time SIP was first drafted. SIP came out at a time when the Internet was young, and bandwidth rates were slow. Part of what caused data to lag so badly was the size of the payload data on each packet being sent. Therefore, the IETF drafted a proposal for how packets could be sent that would minimize the payloads and in effect reduce the time it took to send data. Because SIP was not written for one specific use, many vendors use SIP, even today, for various applications. By the late 1990s and early 2000s, many companies began using SIP for voice and video communication. It was used so heavily that the IETF amended some of the original RFCs to accommodate this specific application of use. Because RFCs are just recommendations for how the protocol should be used, they can be changed by independent vendors to accommodate their own proprietary brand of SIP. Now there are multiple variants with SIP, such as Microsoft. Microsoft released OCS, which was later changed to Lync, which was changed again to Skype for Business, and is now known as Microsoft Teams. Cisco also released a brand of SIP known as Skinny Client Control Protocol, or SCCP. SCCP is a very lightweight protocol that has a simplified message structure, whereas SIP has a range of different messages with each having a lot of additional data. Other vendors followed suit, and this did present a problem with cross-vendor compatibility in the beginning. However, interoperability between vendors is improving. Cisco has made a deliberate move away from SCCP since 2010 and supports SIP as the IETF recommends its use to encourage better interoperability using Cisco Collaboration products.

Basic SIP Registration

The SIP Registration process is very simple without a lot of complex messaging being sent back and forth. The endpoint sends an alias with its IP address to the SIP Registrar. The Registrar will reply with a "200 OK" message if the registration is accepted or an error message, such as "404," if it is not. However, the process to get to this point of registration is quite different depending on the type of SIP Registrar used. The Cisco Unified Communications Manager and the Expressway-C have a different process that leads to the endpoint trying to register.

The registration process on the Cisco Unified Communications Manager seems lengthy, but it actually makes deploying large quantities of phones and endpoints much easier. Every part of this process is automated using different tools on the Cisco routers, switches, and the Cisco Unified Communications Manager to enable phones to communicate and register with the Cisco Unified Communications Manager without any human interaction outside of plugging the patch cable into the back of the phone. The steps to this process are as follows:

1. The endpoint obtains power from the switch.
2. The endpoint loads the locally stored image (Phone-Load).
3. After the locally stored image is loaded, the first communication that the endpoint will send out is a Cisco Discovery Protocol (CDP) frame with a Voice VLAN query to the switch. This CDP communication is used to obtain the Voice VLAN information if no local Voice VLAN ID (VVID) is configured already on the phone. If a non-Cisco

phone or non-Cisco switch is used, then the LLDP-MED protocol can be used for the same purpose.

4. If the Cisco Catalyst switch has a Voice VLAN configured, it sends back a CDP frame with the VVID. Note that the VVID is used for both voice and video traffic.
5. When VLAN discovery is complete, the endpoint will send a DHCP discovery message to the DHCP server. Typically, the DHCP server is a router, but the Cisco Unified Communications Manager, as well as other DHCP server types, can also fulfill this role. A limitation in using the Cisco Unified Communications Manager is that it allows support for only up to 1000 devices. However, in either case, an option is made available for the TFTP server address to be discovered at the same time. This option is called Option 150. When the DHCP server receives the DHCP discovery, it responds with a DHCP offer. The DHCP offer includes an IP address, subnet mask, and default gateway address at a minimum. Additionally, a TFTP server address (with use of Option 150) and possibly one or more DNS addresses can also be provided. The endpoint responds to the DHCP offer with a DHCP request for the specific information sent in the DHCP offer. The DHCP server will then send a DHCP acknowledgment authorizing the use of the DHCP information exchanged and end the DHCP session.
6. Now that the endpoint has appropriate IP address information and the TFTP server address, it can send a TFTP Get message to the TFTP server. This message is typically sent over HTTP when using current endpoints, but TFTP signaling could be used as well. The communication that the endpoints sent to the TFTP server contains their MAC addresses because that is what the Cisco Unified Communications Manager uses to identify the endpoint's configuration file.
7. The Cisco Unified Communications Manager will first exchange a certificate trust list (CTL) file. The CTL file contains a set of certificates and is used only when Cisco Unified Communications Manager cluster security has been enabled. Next, the Cisco Unified Communications Manager will send the configuration files. After the configuration file has been downloaded, the endpoints will verify they are running the requested load, or firmware version. If the version they are running is different from the current version on the TFTP server, or different from a version specified in the configuration file, the Cisco Unified Communications Manager will send the current system load files and upgrade the firmware. Once upgraded, the endpoints will reboot. All information obtained up to this point will be retained.
8. The final step in the process is for the endpoints to register to the Cisco Unified Communications Manager. This is the SIP registration part of the process. The endpoint will send its IP address and alias information to the Cisco Unified Communications Manager and request registration.
9. The Cisco Unified Communications Manager will respond with the SIP message "200 OK." Now the registration process is complete.

Figure 5-6 illustrates the preceding steps when registering SIP endpoints to the Cisco Unified Communications Manager.

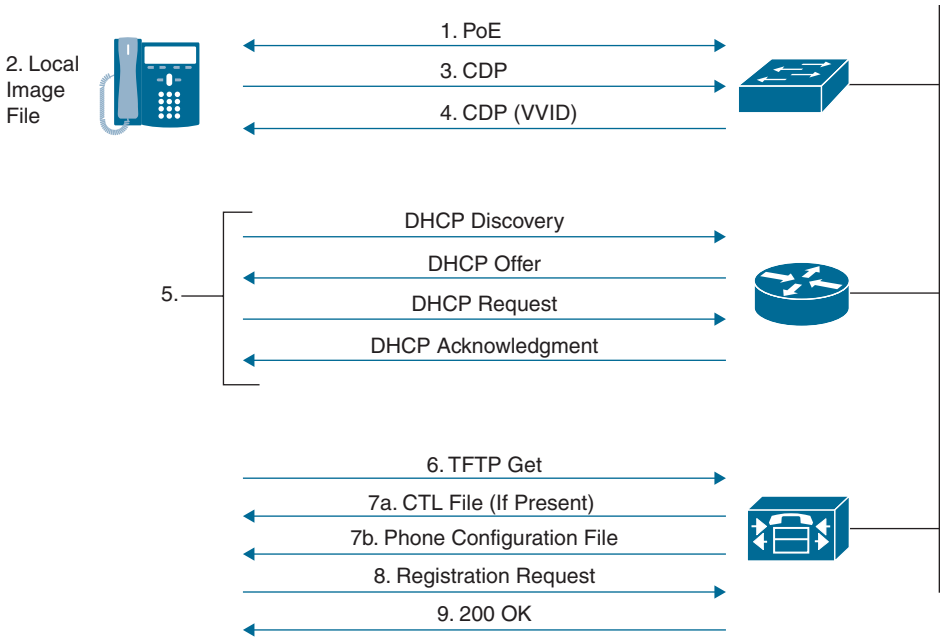

Figure 5-6 *SIP Registration Process to the Cisco Unified Communications Manager*

By contrast, the SIP registration process to a Cisco Expressway-C is not automated as it is when registering to the Cisco Unified Communications Manager. VLAN discovery through CDP is not essential for endpoints registering to the Expressway-C. Also, instead of configuration settings being sent to the endpoint using TFTP, these settings must be configured manually on the endpoint by an administrator. Therefore, there is no need for Option 150 to be configured on the DHCP server. The following steps outline the registration process on the Cisco Expressway-C:

Key Topic

1. Most Cisco Telepresence endpoints that register to the Cisco Expressway obtain local power from the power cube rather than through PoE. The exception would be the Cisco SX10 endpoint. Therefore, when an endpoint is going to register to the Cisco Expressway, it must first obtain power from the power cube.
2. As the endpoint is powering on, it will load the locally stored image, much like an endpoint would within a Cisco Unified Communications Manager environment.
3. Cisco CE software-based endpoints can use CDP for VLAN discovery, but VLAN discovery will not impact Expressway-C registration as long as the VLAN the device is on can route to the Expressway.
4. When VLAN discovery is complete, or if the endpoint does not use CDP, the endpoint will send a DHCP discovery message to the DHCP server. Once the DHCP server receives the DHCP discovery, it responds with a DHCP offer. The DHCP offer includes an IP address, subnet mask, default gateway address, and possibly one or more DNS addresses. TFTP addresses are not used in an Expressway registration deployment. The endpoints respond to the DHCP offer with a DHCP request for the specific information that is sent in the DHCP offer. The DHCP server will then send a DHCP acknowledgment authorizing the use of the DHCP information that is exchanged and end the DHCP session.

5. The SIP URI and Expressway IP address are typically configured manually on the endpoint. These configurations can also be provisioned through Cisco TMS.
6. The final step in the process is for the endpoints to register to the Cisco Expressway. The endpoints will send their IP address and alias to the Cisco Expressway in the request registration. The alias must be in the form of a URI (name@FQDN), and the domain of the alias must match one of the domains configured in the domain database of the Expressway.
7. If there are no configured restrictions on the Expressway, it will respond with the SIP message "200 OK." The registration process is now complete.

Figure 5-7 illustrates the SIP registration process to an Expressway-C.

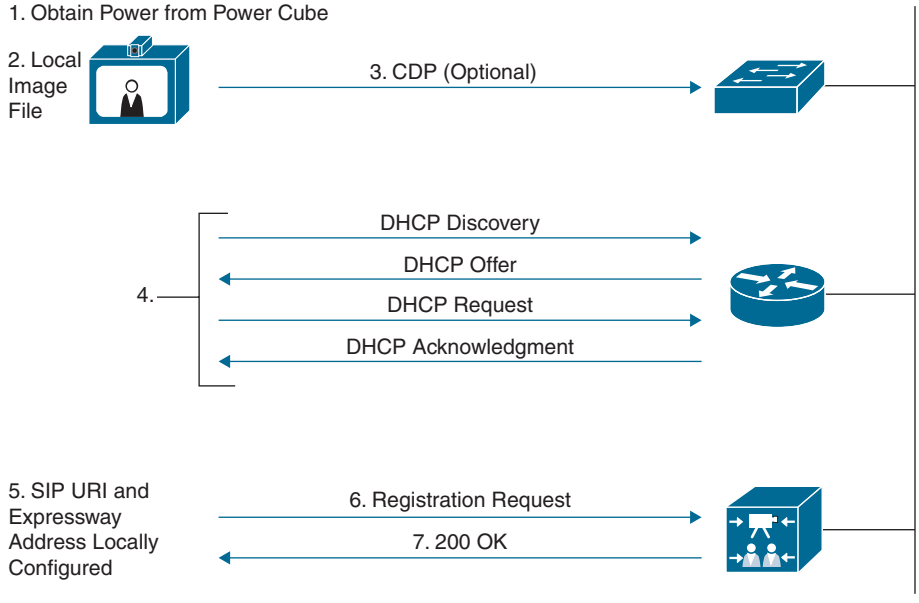

Figure 5-7 *SIP Registration Process to the Expressway Core*

SIP Call Setup

Key Topic

Both the Cisco Unified Communications Manager and the Cisco Expressway process SIP calls in a similar fashion. Both call control servers have Call Admission Control (CAC) elements that can be leveraged, although what these elements are and how they are leveraged mark the differences between each call control server. There are two methods to process a SIP call: early offer and delayed offer. Both early offer and delayed offer use Session Description Protocol (SDP), which is the mechanism SIP uses to exchange codec capabilities and identify the UDP ports that are needed for RTP media. The Cisco Unified Communication Manager can use either early offer or delayed offer, depending on certain settings and other configuration elements. Early offer is the only method that is used in the Cisco Expressway.

A third method of call setup is called early media. The early media feature is supported for SIP calls. Early media is the capability of two user agents to communicate before a call is actually established. Support for early media is important both for interoperability with the public switched telephone network (PSTN) and billing purposes. Early media is defined when media begins to flow before the call is officially connected. Media channels are set up prior to the call connection. These channels are used to provide the ring tone that the caller hears and are not generated by the caller's endpoint or other queuing services, such as Music-On-Hold. Early media is supported on Cisco IOS gateways.

Delayed Offer

SIP delayed offer begins when a source endpoint dials the destination alias using SIP, and an Invite message is sent to the SIP server. The SIP Proxy function of the SIP server examines the table that is created by the SIP Registrar to determine the destination endpoint's IP address using the alias that is dialed. The SIP server will then proxy the Invite message to the destination endpoint. At the same time the invite message is proxied, the SIP server will respond to the source endpoint with a Trying message. The Trying message contains the destination endpoint's IP information. After the Trying message is received, the source endpoint now possesses the source and destination IP addresses. The Invite message of the SIP call setup process is equivalent to the SYN message of the three-way handshake.

When the destination endpoint receives the Invite message, that endpoint now has the source and destination IP address information as well. The destination endpoint will then respond with two messages. The first message is the Ringing message, which tells the destination endpoint to ring and sends a ring-back tone to the source endpoint. The Ringing message is equivalent to the SYN/ACK message of the three-way handshake. When the user of the destination endpoint answers the call, an OK message is sent to the source endpoint. The OK message contains call connection status and is equivalent to the ACK message of the three-way handshake.

After the source endpoint receives the OK message from the destination endpoint, the SDP information can be sent to the destination endpoint. The destination endpoint is now aware of the ports specified by the source endpoint for the media communication, these ports are opened, and the destination endpoint can now receive audio and video media over these UDP ports. The destination endpoint will send back an acknowledgment that the SDP information was received, followed by that endpoint's SDP information. The source endpoint will open the ports specified by the received SDP communication and return an acknowledgment to the destination endpoint, and the call will be set up.

As you can see, a lot of back-and-forth communication occurs when delayed offer is used. For this reason, early offer is the preferred method to use. The whole purpose of SIP is to simplify the data messages sent between two nodes. Figure 5-8 illustrates the SIP delayed offer call setup process.

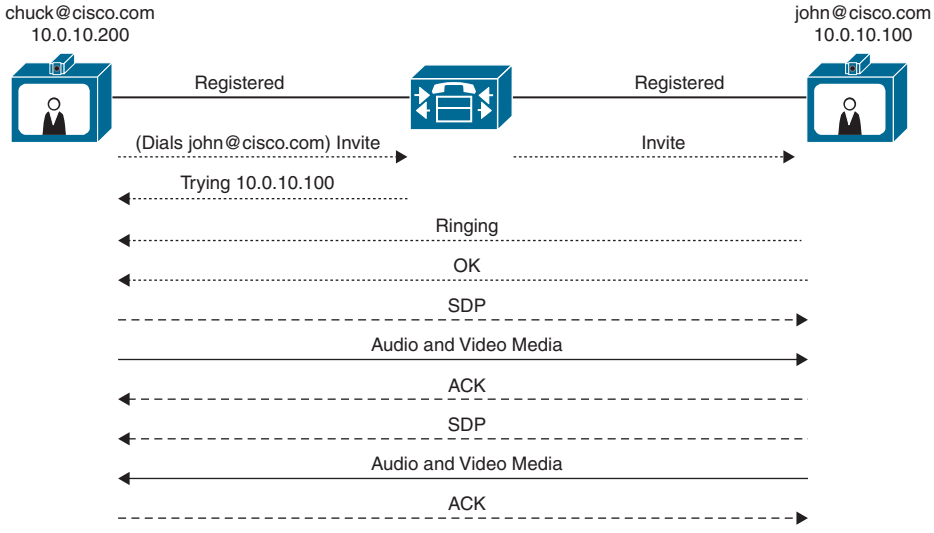

Figure 5-8 *SIP Delayed Offer Call Setup Process*

Early Offer

The big differentiator between delayed offer and early offer happens when the SDP information is sent. The SIP early offer begins when a source endpoint dials the destination alias using SIP, and an Invite message is sent to the SIP server along with the SDP information of the source endpoint. The SIP proxy function of the SIP server examines the table that is created by the SIP Registrar to determine the destination endpoint's IP address using the alias that is dialed. The SIP server will then proxy the Invite message with the SDP packets to the destination endpoint. At the same time the invite message is proxied, the SIP server will respond to the source endpoint with a Trying message. The Trying message contains the destination endpoint's IP information. After the Trying message is received, the source endpoint now possesses the source and destination IP addresses.

When the destination endpoint receives the Invite message, that endpoint will have the source and destination IP address information as well. The destination endpoint will then respond with two messages. The first message is the Ringing message, which tells the destination endpoint to ring and sends a ring-back tone to the source endpoint. Once the user of the destination endpoint answers the call, an OK message is sent to the source endpoint. The OK message contains call connection status, acknowledgment that SDP information has been received, and the destination endpoint's SDP information. Since the destination endpoint is now aware of the ports specified by the source endpoint for the media communication, these ports are opened, and the destination endpoint can now receive audio and video media over these UDP ports. Once the source endpoint receives the OK message from the destination endpoint, the UDP ports that are specified from that communication will be opened so that audio and video media can be received from the destination endpoint. An acknowledgment will be sent to the destination endpoint and the call will be set up. Figure 5-9 illustrates the SIP early offer call setup process.

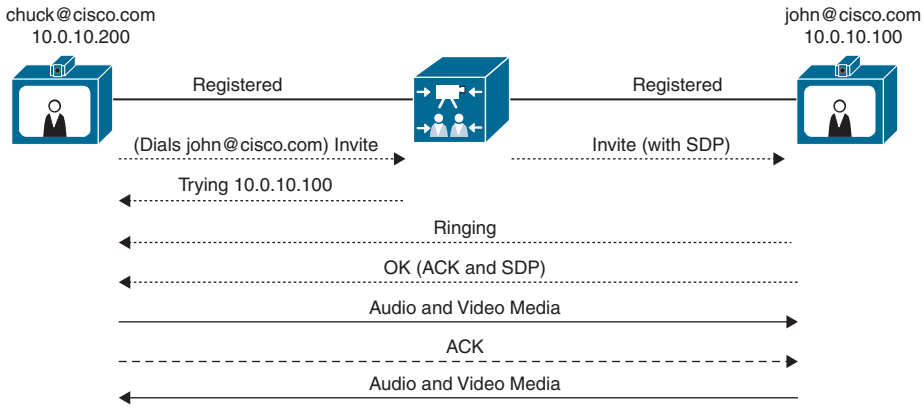

Figure 5-9 *SIP Early Offer Call Setup Process*

NAT and Firewall Traversal Solutions

Communication over IP has come a long way in a short time. One huge obstacle that had to be overcome before companies could really start migrating away from PSTN-based communication to IP-based communication was how to securely route network traffic through firewalls and across NAT servers.

The purpose of a firewall is to control IP traffic entering your network. Firewalls generally block unsolicited incoming requests, meaning that any communication originating from outside your network will be prevented. However, firewalls can be configured to allow outgoing requests to certain trusted destinations and to allow responses from those destinations. Allowing traffic in both directions prevents the firewall from doing its job. Therefore, a firewall exists to protect the inside of a corporate network from outside attack.

Most firewall ports are set to allow traffic to be sent from inside the network to an outside destination and accept a reply on the same port. This means that any communication must be started by an internal system. An example of this type of communication could be when a user within an enterprise network browses to Google.com from a web browser. A TCP communication is sent through the firewall, which marks the packets, and then forwards the request to the Google web server. When the Google server sends back the response, the firewall is already expecting this response, so the communication is allowed to come in and is redirected to the requesting application.

Key Topic

In a similar manner, a video call made from an endpoint inside a network through a firewall to an outside endpoint begins with a TCP communication. The firewall will allow the exchange of TCP call setup information and allow UDP media packets originating from an inside endpoint to be sent to an outside endpoint. However, the media traffic in the opposite direction does not come back on the same ports. The firewall observes these communication packets as originating from outside the network and blocks those packets from ever reaching the internal endpoint. This is a common issue, and it is referred to as *one-way audio* and *one-way video*. Notice also that to even get to this point, the call had to originate from inside the network. Inbound calls originating from outside the network are out of the question. Two endpoints located behind two different firewalls would never be able to call each other. The seeming resolution to this issue would be to open the ports

needed for two-way communication, but this will not work for two reasons. First, opening the ports on the firewall will leave the network vulnerable to attacks from outside the network. This is never a good idea. Second, there is a second issue with NAT that could prevent endpoints from communicating with each other.

The Institute of Electrical and Electronics Engineers (IEEE) first introduced communication using packet-switched technology in 1974. They experimented with several Internet Protocol versions (IPv1-3) until the predominant protocol called IPv4 was established circa 1981. At that time, engineers couldn't imagine the four billion addresses made available with IPv4 would ever run out. Initially, anyone could purchase IP addresses in pools, and they would own them for life. Telco companies and universities were some of the main consumers of these IP addresses. As the number of devices that required an IP address greatly increased, and the World Wide Web began to expand, they realized that the number of people and devices requiring an IP address would soon eclipse the finite number of IPv4 addresses available. One solution to this problem was the introduction of IPv6, which contains 340 undecillion addresses. Some say you could assign an IPv6 address to every grain of sand on Earth and still not run out of addresses. However, IPv6 introduces other issues, such as how to migrate hundreds of millions of devices that are already established under IPv4 over to an IPv6 network.

Another resolution that came about around the same time as IPv6 was Network Address Translation (NAT). The IETF came up with RFC 2663 outlining the basic use of NAT. For NAT to work, IP addresses first had to be divided into two pools: public and private IP addresses. Internet Corporation for Assigned Names and Numbers (ICANN) was created in 1998 to assume responsibility for managing IP addresses and domains. Private IP addresses are designated in the following categories, which anyone can use, but are not routable across the public Internet. The ranges for private IP addresses are

- Class A addresses 10.0.0.0–10.255.255.255 have 16,777,216 available addresses.
- Class B addresses 172.16.0.0–172.31.255.255 have 1,048,576 available addresses.
- Class C addresses 192.168.0.0–192.168.255.255 have 65,536 available addresses.

Public IP addresses are routable across the public Internet and can be leased from an Internet service provider. Today different versions of NAT can be used based on many different factors. However, the basis of how NAT works is that a private IP address is masqueraded with a public IP address when a device needs to route across the public Internet. Your router will mark the packets going out with a virtual port number to enable routing return traffic that comes back from the desired destination. For example, if a computer assigned a private IP address of 10.10.1.14 tries to navigate to Google.com, the edge router will masquerade that private IP address with its assigned public IP address of 209.165.201.1:12345. When the Google server returns communication to the computer, return traffic goes to the public-facing port of the router, but the port :12345 tagged at the end of the IP address indicates to the router where to send the return traffic based on a table the router keeps. The router will then change the destination address from 209.165.201.1:12345 to the private IP address of the endpoint, 10.10.1.14, and route the packets sent from Google to the computer that initiated the communication. Similar to NAT is another protocol called Port Address Translation (PAT). Sometimes NAT and PAT can be used together.

Key Topic

NAT becomes an issue with collaboration devices for two reasons. First, NAT doesn't allow communication to be initiated from outside the private network because the virtual ports can change with each new transmission that is created. So, if two video endpoints behind different NATs wanted to communicate, one would never be able to discover the other. For example, if a device were to try routing to the private IP address of another endpoint, the transmission would fail at the source router because private IP addresses are not publicly routable. Alternatively, if the source device tried to route to the public IP address of the far-end router, after the packets arrived, the far-end router wouldn't know to which device the packets should be routed.

The second issue that comes with NAT has to do with User Datagram Protocol (UDP) transmissions. Whereas TCP communications require a response, UDP communications are unidirectional. Once video calls are set up using TCP, the audio and video packets are sent using UDP. Since each UDP packet sent is essentially a new transmission, a different virtual port is used, and transmissions will never reach their targeted destination.

The IETF, which came up with the SIP communications protocol and NAT, also came up with the first solution that allowed communication between private networks through a NAT server. That protocol is known as STUN, which stands for Session Traversal Utilities for NAT. After creating the RFC for STUN, the IETF came up with two other RFC protocols known as TURN (Traversals Using Relays around NAT) and ICE (Interactive Connectivity Establishment).

STUN, TURN, and ICE are methods that assume certain behavior from the NAT/firewall and do not work in all scenarios. The control is removed from the firewall, which has to be sufficiently opened to allow users to create the pinholes needed to let the communication through. Therefore, STUN, TURN, and ICE offer only a NAT traversal solution and are not firewall traversal solutions at all. Also, these solutions can only operate within the SIP communications protocol, so they offer no solution for H.323 communications.

Key Topic

STUN

STUN requires a STUN client, which could be the phone or some other device, which sends packets to a STUN server on the Internet. The STUN server replies with information about the IP address and ports from which the packets were received and detects the type of NAT device through which the packets were sent. The STUN client can then use the public IP and assigned port in constructing its headers so that external contacts can reach the client without the need for any other device or technique. Once the STUN server assigns a port, it is no longer involved in the line of communication.

STUN requires that the NAT server allow all traffic that is directed to a particular port be forwarded to the client on the inside. This means that STUN works only with less-secure NATs, so-called full-cone NATs exposing the internal client to an attack from anyone who can capture the STUN traffic. STUN may be useful within asymmetric network environments but is generally not considered a viable solution for enterprise networks. In addition, STUN cannot be used with symmetric NATs. This may be a drawback in many situations because most enterprise-class firewalls are symmetric. For more information about STUN, see RFC 5389. Figure 5-10 illustrates how STUN works within a network.

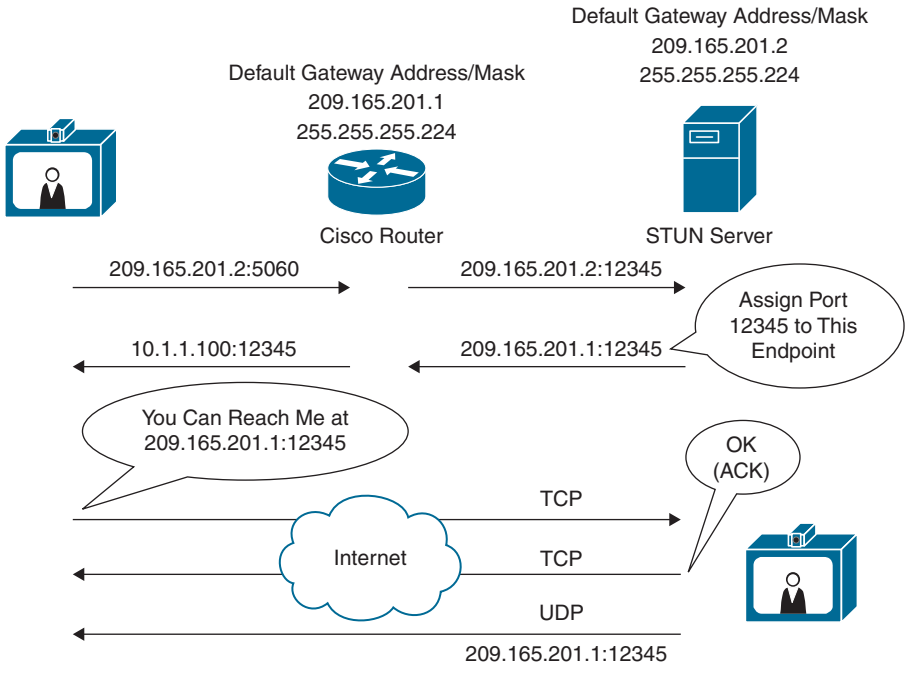

Figure 5-10 *STUN Operation Within an Asymmetric Network*

Key Topic

TURN

TURN operates similarly to STUN, but it allows an endpoint behind a firewall to receive SIP traffic on either TCP or UDP ports. This solves the problems of clients behind symmetric NATs, which cannot rely on STUN to solve the NAT traversal issue. TURN connects clients behind a NAT to a single peer. Its purpose is to provide the same protection as that created by symmetric NATs and firewalls. Symmetric NATs use dynamic ports that often change. Therefore, the TURN server acts as a relay so that any data received is forwarded on to the client, and port allocation can be updated on the fly. The client on the inside can also be on the receiving end, rather than the sending end, of a connection that is requested by a client on the outside.

This method is appropriate in some situations, but since it essentially allows inbound traffic through a firewall, with only the client in control, it has limited applicability for enterprise environments. It also scales poorly since the media must traverse through the TURN server. Also, since all media must traverse the TURN server, the server supporting TURN must be robust enough to handle high volumes of traffic. For more information about TURN, see RFC 5766. Figure 5-11 illustrates how TURN works within a network.

TURN relay services are the only IETF services available on the Expressway-E. To use TURN services, you need the TURN Relay option key. This controls the number of TURN relays that can be simultaneously allocated by the TURN server. The TURN page found under the **Configuration > Traversal > TURN** menu is used to configure the Expressway-E's TURN settings. Table 5-3 identifies the configurable options for TURN on the Cisco Expressway.

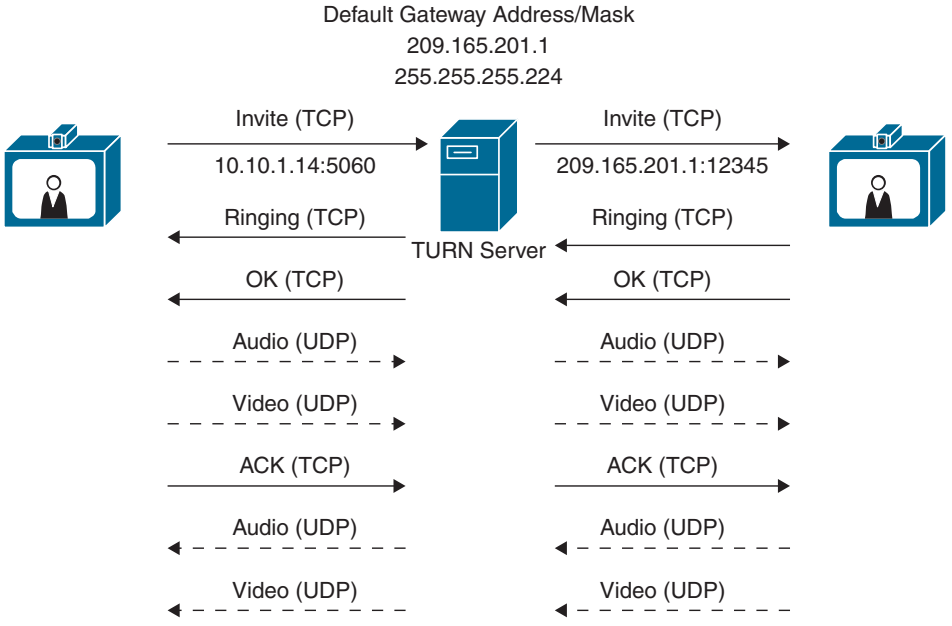

Figure 5-11 *TURN Operation Within a Symmetric Network*

Key Topic

Table 5-3 Configurable Options for TURN on the Cisco Expressway

Field	Description	Usage Tips
TURN Services	Determines whether the Expressway offers TURN Services to traversal clients.	
TURN Requests Port	The listening port for TURN requests. The default is 3478. On large VM deployments, you can configure a range of TURN request listening ports. The default range is 3478–3483.	To allow endpoints to discover TURN Services, you need to set up DNS SRV records for _turn._udp. and _turn._tcp. (either for the single port or a range of ports as appropriate). If you need to change the TURN requests port (or range, for large systems) while the TURN Services are already On, do the following: 1. Change TURN Services to Off and click Save. 2. Edit the port number/range. 3. Change TURN Services to On and click Save. The reason is that changes to the port numbers do not take effect until the TURN Services are restarted.

Field	Description	Usage Tips
Authentication Realm	This is the realm sent by the server in its authentication challenges.	Ensure that the client's credentials are stored in the local authentication database.
Media Port Range Start/End	The lower and upper port in the range used for the allocation of TURN relays. The default TURN relay media port range is 24000–29999.	

A summary of the TURN server status is displayed at the bottom of the TURN page. When the TURN server is active, the summary also displays the number of active TURN clients and the number of active relays. Click the active relay links to access the TURN relay usage page, which lists all the currently active TURN relays on the Expressway. Further details of each TURN relay can be reviewed, including permissions, channel bindings, and counters.

Key Topic

ICE

ICE provides a mechanism for SIP client NAT traversal. ICE is not a protocol, but a framework that pulls together a number of different techniques such as TURN and STUN. It allows clients residing behind NAT devices to discover paths through which they can pass media, verify peer-to-peer connectivity via each of these paths, and then select the optimum media connection path. The available paths typically depend on any inbound and outbound connection restrictions that have been configured on the NAT device. Such behavior is described in RFC 4787. ICE essentially incorporates all of the methods proposed for NAT traversal of SIP that do not rely on the firewall or NAT device. ICE is a complex solution to the problem of NAT traversal, but because it encompasses multiple solutions, it is regarded as one that will always enable the connection, regardless of the number of NATs involved. However, ICE still relies on client/server-based approaches and removes control from the enterprise. Due to its complexity, there is very limited client support for ICE today.

When a client reaches out to the ICE server, it can determine what type of NAT is being used; whether it's in an asymmetric or symmetric network environment. The ICE server will then establish a connection with the client using STUN or TURN, depending on what the situation calls for. If STUN is used, then the ICE server will assign a port to the client and step out of the line of communication. If TURN is used, then the ICE server will act as the relay between client communications. For more information about ICE, see RFC 5245. Figure 5-12 illustrates how ICE works within a network.

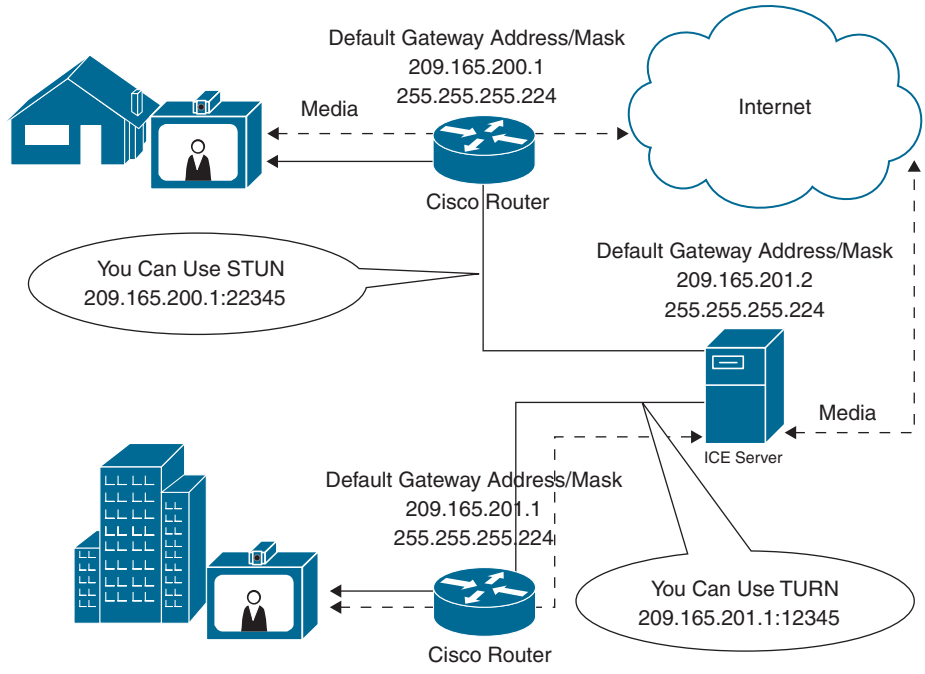

Figure 5-12 *ICE Operation Within Asymmetric and Symmetric Networks*

ASSENT and H.460

Although the IETF overcame many problems, there were still many more to overcome. Their solutions, although good, were incomplete. Tandberg was a company that had been a leader in video telepresence for many years prior to its acquisition by Cisco. Tandberg climbed this ladder to success much the same way as Cisco, by acquiring key companies that possessed the technology it needed for the time. In 2004 Tandberg acquired Ridgeway Systems and Software, which was a UK-based software company specializing in firewall and NAT traversal. Ridgeway had developed a unique proprietary solution that is known today as Assent. This protocol has revolutionized the way IP communication traverses firewalls and NATs and has become the standard the ITU followed for an open standard all companies can use.

Key Topic

Assent requires two components to work: a traversal server and a traversal client. The traversal server resides outside the firewall or in a demilitarized zone (DMZ). The traversal client resides inside the firewall and initiates communication with the traversal server. Ports do need to be opened on the firewall, but they cannot be used unless a communication is initiated from inside the firewall. This is where the magic happens. The traversal client sends a keepalive message to the traversal server, essentially asking, "Do you have any calls for me?" Should someone initiate a call from outside the firewall through the traversal server, that server can respond to the keepalive message sent from the traversal client. As far as the firewall is concerned, the communication was initiated from inside the firewall with the keepalive message. Now the ports allocated to this solution can be used after the call setup has completed. Even better, though, are the ports needed for media using Assent. Only two ports are required to be opened on the firewall because Assent will multiplex the media so that all RTP traffic uses one port and all RTCP traffic uses a second port. In addition to the firewall traversal capabilities of Assent, NAT traversal is built into the protocol as well. Also, Assent can be used with both the SIP and H.323 communication standards.

Key Topic

Assent is such a powerful tool that the ITU used it as the basis to develop H.323 Traversal standards. By the summer of 2005, the standards were completed and in full use. H.460.17 was the traversal standard used prior to the Assent-based standards. H.460.17 performs firewall traversal by carrying the media over TCP ports instead of UDP. H.460.18 works just like Assent, except it requires demultiplexed ports 50000 to 52400 to be opened on the firewall. H.460.19 works as a layer on H.460.18 to allow multiplexing the media ports so only two ports need to be opened for RTP and RTCP media streams. In this, H.460.18 and H.460.19 accomplish together what Assent is capable of independently. It is important to note that the ITU standards for firewall traversal only support the H.323 communication standard.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 5-4 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 5-4 Key Topics for Chapter 5

Key Topic Element	Description	Page Number
Paragraph	TDM	88
Paragraph	BRI	89
Paragraph	PRI	89
Table 5-2	Minimum Standards for H.320 Compliance	91
Paragraph	Minimum Standards for H.323 Compliance	91
Paragraph	H.323 Aliases	91
Paragraph	Q.931 and H.245	93
Paragraph	H.323 Call Setup with RAS Messaging	94
Paragraph	Functions of SIP Server	95
List	CUCM Registration Process for SIP	96
List	Expressway Registration Process for SIP	98
Paragraph	Differentiator Between Early Offer and Delayed Offer	99
Paragraph	Firewall Dilemma	102
List	Private IP Address Classes	103
Paragraph	NAT Dilemma	104
Section	STUN	104

Key Topic Element	Description	Page Number
Section	TURN	105
Table 5-3	Configurable Options for TURN on the Cisco Expressway	106
Section	ICE	107
Paragraph	Assent Operation for H.323 and SIP	108
Paragraph	H.460 Operation for H.323	109

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Assent, BRI, CAC, Call Setup Mode, CDP, CTL, DHCP, Discovery Mode, E.164 Alias, Firewall, FQDN, Gatekeeper, GRQ, H.239, H.245, H.320, H.323, H.323 ID, H.460.17, H.460.18, H.460.19, HTTP, ICANN, ICE, IEEE, IETF, ITU, LLDP-MED, NAT, Option 150, PAT, Prefix, PRI, Q.931, RAS, RFC, RIP, RRQ, SCCP, SDP, SIP, SIP Proxy, SIP Registrar, SIP Server, STUN, TCP, TDM, TFTP, TURN, UDP

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the two main types of ISDN with the channels supported, carrier types, and bandwidth supported on each channel.
2. List and define the three types of aliases that can be used with H.323 registered endpoints.
3. What are the nine steps in the SIP registration process to a CUCM?
4. List the three IETF NAT traversal solutions, the Cisco firewall, and NAT traversal solution and the three ITU firewall and NAT traversal solutions.

This page intentionally left blank

CHAPTER 6

Cisco Solution for Converged Collaboration

This chapter covers the following topics:

Introduction to Cisco Endpoints: This topic will introduce the various Cisco voice and video endpoints available on the market today, including UC phones, soft clients, and Telepresence endpoints.

Introduction to Cisco Call Control: This topic will introduce the Cisco infrastructure that can be used for call control in a Collaboration solution, including the Cisco Unified Communications Manager, Cisco Unified Communications Manager Express, Cisco Expressways, and Webex Control Hub.

Introduction to Cisco Applications: This topic will introduce common Cisco infrastructure applications that enhance the user experience in a collaboration deployment, including Cisco Unity Connection Server, Cisco IM and Presence Service, Cisco Meeting Server, and Management software.

Designing a Cisco Collaboration Solution: This topic will overview the various aspects to designing a Cisco Collaboration solution, including licensing, sizing, bandwidth management, high availability, disaster recovery, dial plan, security, and quality of service.

Establishing a foundation for audio and video communication up to this point has been important. We will revisit key information shared through the first five chapters of this book throughout the rest of the chapters ahead. This chapter, as well as the rest of this book, will focus on specific products in the Cisco Collaboration product portfolio. This chapter will not cover an exhaustive list of all Cisco Collaboration products, and not every product mentioned in this chapter will be covered in later chapters. The purpose of this chapter is merely to provide a high-level overview of the main components available in a Cisco Collaboration solution. Topics discussed in this chapter include the following:

- Introduction to Cisco Endpoints
  - UC Phones
  - Soft Clients
  - Telepresence Endpoints
- Introduction to Cisco Call Control
  - Cisco Unified Communications Manager
  - Cisco Unified Communications Manager Express

- Cisco Expressways
- Webex Control Hub

- Introduction to Cisco Applications
  - Cisco Unity Connection Server
  - Cisco IM and Presence Service
  - Cisco Meeting Server
  - Management Software
- Designing a Cisco Collaboration Solution
  - Licensing
  - Sizing
  - Bandwidth
  - High Availability
  - Disaster Recovery
  - Dial Plan
  - Security
  - QoS

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

1.1 Describe the key design elements of the following, pertaining to the Cisco Collaboration architecture as described in the SRND/PA:

- 1.1.a Licensing (Smart, Flex)
- 1.1.b Sizing
- 1.1.c Bandwidth
- 1.1.d High availability
- 1.1.e Disaster recovery
- 1.1.f Dial plan
- 1.1.g Security (certificates, SRTP, TLS)
- 1.1.h QoS

1.3 Configure these network components to support Cisco Collaboration solutions:

- 1.3.a DHCP
- 1.3.b NTP
- 1.3.c CDP

- 1.3.d LLDP
- 1.3.e LDAP
- 1.3.f TFTP
- 1.3.g Certificates

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 6-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 6-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Introduction to Cisco Endpoints	1–2
Introduction to Call Control	3–6
Introduction to Cisco Applications	7–9
Designing a Cisco Collaboration Solution	10–14

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What soft clients does Cisco offer in its current product portfolio? (Choose three.)
   - **a.** Jabber Client
   - **b.** Jabber Video for Telepresence
   - **c.** Webex (Meet, Team, Call)
   - **d.** CMS
   - **e.** WebRTC
   - **f.** CMA
2. Which of the following Cisco Telepresence endpoint category markings is used for personal endpoints?
   - **a.** DX
   - **b.** MX
   - **c.** SX
   - **d.** IX
   - **e.** Webex endpoint

3. What is the maximum number of DHCP addresses that can be delivered and managed by a CUCM?
   a. 10
   b. 100
   c. 1000
   d. 10,000
4. What is the maximum user capacity on the Cisco Unified Communications Manager Express?
   a. 150
   b. 250
   c. 350
   d. 450
5. What is the primary difference between a Cisco VCS and a Cisco Expressway Server?
   a. Cisco VCS has user-based licenses, but an Expressway has device-based licenses.
   b. Cisco VCS has device-based licenses, but an Expressway has user-based licenses.
   c. VCS is used autonomously, and the Expressway is used with CUCM exclusively.
   d. VCS is used exclusively with the CUCM, and the Expressway is autonomously.
6. Which Webex tool can be used to call into a Webex Meeting?
   a. Webex Meeting application
   b. Webex Teams application
   c. Webex Calling application
   d. All the above
7. Which of the following applications can offer voicemail services to UC phones? (Choose two.)
   a. CUCCE
   b. CUCCX
   c. CUC
   d. CUE
   e. IMP
   f. CUCM
   g. CMS
   h. Expressway
8. Which of the following is a service offered through CMS?
   a. Voicemail integration
   b. Manual creation of users
   c. Microsoft Skype for Business Integration
   d. Direct endpoint registration to the server

9. Which of the following management tools allows conference meetings to be scheduled?
   a. Telepresence Management Suite
   b. Prime Collaboration Provisioning
   c. Prime Collaboration Assurance
   d. Prime Collaboration Analytics
   e. Telepresence Management Suite and Prime Collaboration Provisioning
10. Which of the following options is included with a CUCL Enhanced license?
   a. Unity Connection
   b. Expressway Firewall Traversal
   c. PMP Basic
   d. Webex Conferencing
11. How many endpoints can a BE6000H server that is running CUCM support?
   a. 1000
   b. 1200
   c. 2500
   d. 5000
12. What is the maximum number of peers supported in an Expressway Cluster?
   a. 4
   b. 6
   c. 8
   d. 10
13. What protocol is used to secure SIP signaling across the CUCM?
   a. HTTPS
   b. SSL
   c. TLS
   d. AES128
14. What is Layer 2 QoS called?
   a. Cost of service
   b. Class of service
   c. DiffServ
   d. IntServ

Foundation Topics

Introduction to Cisco Endpoints

Cisco has spent the last several years restructuring its Collaboration endpoint portfolio to bring out a line of endpoints that offer cutting-edge technology in a sleek design at a reasonable price. All Cisco Collaboration endpoints can be divided into two main categories: Unified Communications (UC) endpoints and Telepresence endpoints. Some UC endpoints are voice-only, and others support high-definition (HD) video. Some of the soft clients available fall into the UC category as well. All Telepresence endpoints are HD video-capable and offer

more features for the end user. The Cisco Unified Communications Manager treats these two types of collaboration endpoints differently in regard to QoS.

UC Phones

Cisco voice over IP (VoIP) phones are user-friendly and full-featured to meet the needs of entire organizations. They range from a company lobby phone to the desks of the busiest managers and C-level employees. Cisco's VoIP phones provide all the features companies use from their office deskphones, such as speakerphone, transfer, hold, and voicemail access, as well as interactive video collaboration and the capability for a PC to use the same network connection as the phone. Some different models support Bluetooth, USB, Wi-Fi, and other advanced features. Some phones have a built-in camera with HD capabilities. Each phone connects back to the Cisco UCM using the Session Initiation Protocol (SIP) and comes equipped for Power over Ethernet (PoE), which can be supplied by many Cisco switches. Alternatively, a power supply can be used in conjunction with the phone. Cisco has narrowed its VoIP phone product portfolio to three categories of phones: the 3900 series, the 7800 series, and the 8800 series phones. DX series phones fit in this category as well if they are still running the Android software. Support for Android software on the DX endpoints was end of life as of October 1, 2018. Cisco strongly recommends migrating the software on these phones to the Cisco Telepresence CE software.

Soft Clients

Another product in the Cisco UC phone category is the Cisco Jabber client. This software phone can be installed on any Microsoft Windows or Apple Mac computer. An app version is also available on Android and Apple IOS devices, including phones and tablets. Jabber client phones can be configured in the Cisco Unified Communications Manager using the Cisco Unified Client Services Framework (CSF) for PC clients. Cisco Jabber client applications provide instant messaging (IM), presence, voice and video communication, voice messaging, desktop sharing, and other collaborative workspace capabilities that support 1080p30 high-definition video interoperability. The CPU must be Intel Core i5 or later, with a bandwidth of between 2 and 4 Mbps, and the client must be running version 12.6 or later. The Cisco Jabber client uses the Cisco Precision Video Engine and ClearPath technology to optimize video media. The Cisco Precision Video Engine uses fast video-rate adaptation to negotiate optimum video quality, based on network conditions. These clients can be used on premises or through a Hosted Collaboration Solution (HCS) in the cloud. No matter what platform you choose to operate this client from, Cisco Jabber clients provide a consistent experience across devices.

Cisco Unified Client Services Framework is a software application that combines several services into an integrated client. An underlying framework is provided for integration of Cisco UC services, including audio, video, web collaboration, visual voicemail, and more, into an IM and Presence application. As mentioned previously, Cisco Jabber is based on the Cisco Unified Client Services Framework and combines advanced collaborative media features with Cisco UC. Cisco Jabber uses SIP for call control, XMPP for IM and Presence, and Computer Telephony Integration (CTI) for desktop IP phone control. You can use CTI to take advantage of computer-processing functions while making, receiving, and managing telephone calls. CTI applications allow you to perform such tasks as retrieving customer information from a database using a caller ID, or working with the information gathered by an interactive voice response (IVR) system to route a customer's call, along with that caller's information, to the appropriate customer service representative.

Key Topic

Cisco Jabber operates in one of two modes: deskphone mode or softphone mode. In deskphone mode, the Cisco Jabber client controls the Cisco IP phone of the user. Should a call be placed from the Jabber soft client, the video phone connected will actually launch the call and use its own resources for audio and video. The same thing is true for answering incoming calls using Jabber in deskphone mode. For an IP phone without a camera, the video input and output are processed on the Cisco Jabber client platform, but the voice input and output are processed on the IP phone. A protocol known as Cisco Audio Session Tunnel (CAST) is used to split the audio and video media between the two destinations. In softphone mode, the Cisco Jabber client behaves like any other IP phone and originates and terminates all audio and video communication interactions using the computer resources upon which it is running.

The Cisco Meeting Application, or CMA, is another soft client application that can run on a computer, smartphone, or tablet. CMA is built on the WebRTC protocol and is dependent on the Cisco Meeting Server, or CMS. CMA is capable of audio-only calling, HD video calling, and instant messaging. Through proper integrations configured on the Cisco Meeting Server, CMA is completely interoperable with Microsoft Skype for Business. Cisco is no longer developing CMA, but development and support of WebRTC will continue. Cisco has been bolstering the Cisco Jabber application to replace CMA. For more information on these added capabilities to Jabber, see Chapter 26, "Users and Cisco Jabber Soft Clients."

Key Topic

Cisco Webex Teams is a cloud-based application that can be installed as an application on a computer, smartphone, or tablet. Originally, Webex Teams was called Cisco Spark, but Cisco decided to combine its two cloud-based solutions into a single and more powerful solution to bring more control and capabilities into the hands of users. Although the primary purpose of the Webex Teams application is persistent group and point-to-point messaging, it is capable of voice and video calling as well. The scope of this book does not cover Cisco cloud collaboration products and solutions, but watch for more information on this topic in other Cisco Press books or go to Cisco's website to find support documentation for more information.

Telepresence Endpoints

The categorization of Cisco Telepresence endpoints has changed many times over the years Cisco has been developing these products. Currently, Cisco Telepresence endpoints are divided into five categories. The SX endpoints are Solutions Experience endpoints intended for integrators to customize meeting rooms for businesses. The MX endpoints are Meeting Experience endpoints and are all-in-one meeting room solutions that don't require room remediation or integration. These endpoints offer a simple plug-and-play setup that anyone can accomplish. DX endpoints are Desktop Experience endpoints that offer a more personal user experience during meetings. In recent years, a new line of endpoints has been created within Cisco's Telepresence endpoint product portfolio. Webex endpoints, formerly known as Spark endpoints, are contained in a class all their own, although they have some similarities to the MX and SX endpoints. All of these endpoints share a common base code known as Cisco Telepresence Collaboration Endpoint (CE) software, which is based on the legacy Telepresence Codec (TC) software. Therefore, no matter what endpoint is used within an enterprise solution, configuration of each endpoint is the same, and the user experience is the same.

The fifth category of Telepresence endpoints is the Immersive Experience, or IX, room endpoints, which are complete room integration systems that offer 6–18 participants an "immersive experience" as close to an in-person meeting as current technology will allow. These

IX room systems are complex to install and cannot be moved after installation is complete. However, the user experience with these systems is as simple to use as any other Cisco Telepresence endpoint. This category of endpoints is the only Cisco Telepresence endpoint that does not use the Cisco Telepresence CE software. IX endpoints use a Cisco Telepresence System (CTS) software as the base code. There were once many endpoints in the Cisco product portfolio based on the CTS software, but they are all end of sale (EoS). The IX endpoints are the last and only endpoints that use this software, and they went end of sale in October 2019. Table 6-2 outlines all of the current endpoints in each of these categories.

Key Topic

Table 6-2 Cisco Telepresence Endpoint Product Portfolio

DX	MX	SX	IX	Webex Endpoints
DX70 (EoS August 16, 2018)	MX200G2	SX10	IX5000 (EoS October 2019)	Webex Room Kit Series
DX80	MX300G2	SX20	IX5200 (EoS October 2019)	Webex 55 (Single/Dual)
	MX700 (Single/Dual)	SX80		Webex 70 (Single/Dual)
	MX800 (Single/Dual)			Webex Board

Introduction to Cisco Call Control

What good are all these endpoints without a call control system to register to? The Cisco Collaboration solution is a complete end-to-end solution. Cisco is the only solution on the market today that can offer all the switching and routing needs, along with the call control and premium endpoints to connect users regardless of their location. Just as the endpoint portfolio is robust enough to meet the needs of any size company, the Collaboration infrastructure also is a robust and extensive solution with many facets to suit the needs of any customer. On-premises solutions offer call control through the Cisco Unified Communications Manager, the Cisco Expressway series, and the Cisco Unified Communications Manager Express (CME). The on-premises deployment can be extended with additional supporting infrastructure. Unity Connection provides unified messaging services to remote and on-premises endpoints through the Cisco Unified Communications Manager. Unity Express is a unified messaging service that runs on the Cisco Integrated Services Routers (ISRs). The Cisco Unified Communications Manager IM and Presence Service provide native standards-based, dual-protocol, enterprise instant messaging and network-based presence as part of Cisco Unified Communications. For small to medium-sized businesses, Cisco offers subscription-based cloud call control and services. Through partners of Cisco, customers can choose a Hosted Collaboration Solution (HCS), which offers the same on-premises products outlined previously, except hosted in the cloud. Alternatively, Cisco has a relatively new solution to offer to customers. The Cisco Webex Cloud solution (formerly Cisco Spark) allows the registration of endpoints, call control, and IP-to-PSTN connectivity, all managed from the cloud.

Cisco Unified Communications Manager

Cisco Unified Communications Manager extends enterprise telephony features and functions to packet telephony network devices. These packet telephony network devices include Cisco IP phones, media-processing devices, VoIP gateways, and multimedia applications. Additional data, voice, and video services, such as converged messaging, multimedia conferencing, collaborative contact centers, and interactive multimedia response systems, interact with the IP telephony solution through the Cisco Unified Communications Manager API. Cisco Unified Communications Manager provides call processing. Call processing refers to the complete process of routing, originating, and terminating calls, including any billing and statistical collection processes. Cisco Unified Communications Manager sets up all the signaling connections between endpoints and directs devices such as phones, gateways, and conference bridges to establish and tear down streaming connections. The dial plan is a set of configurable lists that Cisco Unified Communications Manager uses to determine call routing. Cisco Unified Communications Manager provides the ability to create scalable dial plans for users. Cisco Unified Communications Manager also extends services such as hold, transfer, forward, conference, speed dial, last-number redial, call park, and other features to IP phones and gateways. Cisco Unified Communications Manager uses its own database to store user information. You can authenticate users either locally or against an external directory. You also can provision users by directory synchronization. With directory synchronization, you can automatically add users from the directory to the local database. Cisco Unified Communications Manager allows synchronization from the following directories to the database:

- Microsoft Active Directory 2003 R1/R2(32-bit)
- Microsoft Active Directory 2008 R1(32-bit)/R2(64-bit)
- Microsoft Active Directory Application Mode 2003 R1/R2 (32-bit)
- Microsoft Active Directory 2012
- Microsoft Lightweight Directory Services 2008 R1(32-bit)/R2(64-bit)
- Microsoft Lightweight Directory Services 2012
- Sun ONE Directory Server 7.0
- OpenLDAP 2.3.39
- OpenLDAP 2.4
- Oracle Directory Server Enterprise Edition 11gR1

Cisco Unified Communications Manager provides a programming interface to external applications such as Cisco Jabber, Cisco Unified IP IVR, Cisco Personal Assistant, and Cisco Unified Communications Manager Attendant Console.

Cisco Unified Communications Manager uses different signaling protocols to communicate with Cisco IP phones for call setup and maintenance tasks, including SIP, SCCP, or even H.323 gateway services. When a calling endpoint dials the number of a called endpoint, dialed digits are sent to Cisco Unified Communications Manager, which performs its main function of call processing. Cisco Unified Communications Manager finds the IP address of the called endpoint and determines where to route the call. Using SCCP or SIP, Cisco Unified

Communications Manager checks the current status of the called party. If Cisco Unified Communications Manager is ready to accept the call, it sends the called party details and signals, via ring back, to the calling party to indicate that the destination is ringing. Cisco IP phones require no further communication with Cisco Unified Communications Manager until either the calling or called endpoint invokes a feature, such as call transfer, call conferencing, or call termination. After the call setup is finished, media exchange normally occurs directly between Cisco IP phones using RTP to carry the audio and potentially video stream.

Key Topic

Cisco Unified Communications Manager depends on some additional network elements. In particular, the Cisco Unified Communications Manager cluster uses external NTP and DNS servers plus DHCP and TFTP services that are used by the endpoints. NTP is a protocol for synchronizing computer system clocks over IP networks. NTP has a hierarchical organization that is based on clock strata. Stratum 0 is an extremely precise clock source, such as an atomic clock or radio clock. A stratum 1 server is directly connected to a stratum 0 clock and can provide time information to other (stratum 2) devices, which in turn serve stratum 3 devices. Cisco Unified Communications Manager typically uses stratum 1 NTP to obtain time information from a time server. Only the publisher sends NTP requests to the external NTP server or servers. Subscribers synchronize their time with the publisher. NTP must be enabled and configured during installation of Cisco Unified Communications Manager. At least one external NTP reference must be reachable and functioning when installing the Cisco Unified Communications Manager publisher to complete the installation. Cisco recommends using a minimum of three external NTP servers in a production environment.

It is extremely important that all network devices have accurate time information because the system time of Cisco Unified Communications Manager is relevant in the following situations:

- Cisco IP phones display date and time information. This information is obtained from the device pool on the Cisco Unified Communications Manager.
- CDR and CMR, which are used for call reporting, analysis, and billing, include date and time information.
- Alarms and events in log files, as well as trace information in trace files, include time information. Troubleshooting a problem requires correlation of information that is created by different system components (Cisco Unified Communications Manager, Cisco IOS gateway, and so on). This problem solving is possible only if all devices in the network have the same correct time information.
- Some Cisco Unified Communications Manager features are date-based or time-based and therefore rely on the correct date and time. These features include time-of-day routing and certificate-based security features.

To ensure that all network devices have the correct date and time, it is recommended that all network devices use NTP for time synchronization.

Key Topic

The Cisco Unified Communications Manager DHCP server is designed to serve IP phones in small deployments with a maximum of 1000 devices. It provides a subset of Windows, Linux, or Cisco IOS DHCP server functionality that is sufficient for IP phones, but it should not be used for other network devices, such as PCs. The DHCP server of Cisco Unified Communications Manager must not be used with deployments of more than 1000 registered

devices. Even if there are fewer devices, the CPU load of the services must be watched closely. If high CPU load is experienced, the DHCP service should be provided by other devices, such as a dedicated DHCP server, switch, router, or another server. Multiple DHCP services can be configured per Cisco Unified Communications Manager cluster. Each Cisco Unified Communications Manager DHCP server can be configured with multiple subnets. In nonattached subnets, DHCP relay must be enabled so that the DHCP requests that were sent out by the clients are forwarded to the DHCP server.

Cisco Unified Communications Manager can use IP addresses or names to refer to other IP devices in application settings. When names are used, they need to be resolved to IP addresses by DNS. Both methods have some advantages. The system does not depend on a DNS server, which prevents loss of service when the DNS server cannot be reached. When a device initiates a connection for the first time, the time that is required to establish the connection is shorter because a DNS lookup sent to the DNS server and a DNS reply sent back from the server are not required. By eliminating the need for DNS, there is no danger of errors caused by DNS misconfiguration. Troubleshooting is simplified because there is no need to verify proper name resolution.

When DNS is used, management is simplified because logical names are simpler to manage than 32-bit addresses. If IP addresses change, there is no need to modify the application settings because they can still use the same names; only the DNS server configuration has to be modified in this case. IP addresses of Cisco Unified Communications Manager servers can be translated toward IP phones because the IP phone configuration files include server names, not the original server IP address, which should appear differently to the IP phone. As long as these names are resolved to the correct address when IP phones send out DNS requests, the NAT is not a problem. If certificates are being used to secure the communications environment, DNS will be required because the DNS FQDN is an integral part of certificates. Although historically Cisco has recommended that DNS not be used, with the increasing need for secure connections, it is best practice to use DNS throughout a Cisco Collaboration environment.

The Cisco Unified Communications Manager is a very powerful tool with many call features to offer to companies of any size. Volumes of books have been dedicated to the many facets the Cisco Unified Communications Manager has to offer. Chapters 15 through 19 of this book will delve into these features, including initial setup considerations, LDAP integrations, registration methods, CAC, and globalized call routing through the Cisco Unified Communications Manager.

Cisco Unified Communications Manager Express

Cisco Unified Communications Manager Express (CME) provides call processing to Cisco Unified IP phones for distributed enterprise branch-office environments. Even branch offices within the same enterprise can have different needs and requirements when it comes to unified communications. Cisco Unified CME delivers on this need by providing localized call control, mobility, and conferencing alongside data applications on Cisco Integrated Services Routers (ISRs). Because the solution is Cisco IOS software-based, Cisco Unified CME is easy to configure and can be tailored to individual site needs. It is feature-rich and can be combined with Cisco Unity Express and other services on the Cisco ISR to provide an all-in-one branch-office solution that saves valuable real estate space. Cisco Unified CME is ideal if you are looking for an integrated, reliable, feature-rich unified communications system for up to 450 users.

Cisco Expressway

The Cisco Expressway Series call control components are based on the Cisco Telepresence Video Communication Server (VCS). In April 2010, Cisco closed on the acquisition of a company called Tandberg, which had a call control solution called the VCS. This call control system is different from the Cisco Unified Communications Manager in many ways; namely, the VCS provides call control only for video devices, whereas the Cisco Unified Communications Manager provides call control for both voice and video endpoints. However, the VCS solution possesses a capability that does not exist in the Cisco Unified Communications Manager. The VCS is capable of true firewall and NAT traversal between the internal network and the public Internet. In an effort to capitalize on this capability, Cisco released the Expressway series that is built on the same operating system (OS) as the VCS. The difference was that endpoints could not register directly to the Expressway series servers. The Expressway series existed to secure proxy registration requests from endpoints outside of the corporate network to the Cisco Unified Communications Manager inside the network without the use of a VPN. This function is known as Mobile and Remote Access (MRA).

Key Topic

In August 2016, Cisco announced that the Expressway series servers would support device registration directly to the Expressway with appropriate licenses. At this point, this announcement confused a lot of people as to what the distinction was between an Expressway and a VCS. The menus were already identical, and the only distinction before was the registration capabilities. This is still the distinction between these two servers. Although registration is allowed on both products, the Cisco VCS allows for device-based licensing, whereas the Expressway allows for user-based licensing. Endpoints can register directly to the Cisco VCS Control or the Cisco VCS Expressway via SIP or H.323. The Expressway Core and Expressway Edge can now also support endpoint registrations directly via SIP or H.323, but the Expressway Edge can also proxy registration requests to the Expressway Core or the Cisco Unified Communications Manager. However, it can only proxy SIP registration requests, not H.323. Table 6-3 compares the differences between the Cisco Expressway and the Cisco VCS.

Key Topic

Table 6-3 Comparison of the Cisco Expressway and the Cisco VCS

Feature	Cisco Expressway	Cisco VCS
Server Components	Expressway Core Expressway Edge	VCS Control VCS Expressway
Registration Licensing	Included with CUCL and CUWL user licenses (Registration supported on X8.9 or later)	Device Registration Licenses required (2500 max per server)
Call Licensing	Internal and mobile calling included Rich Media Session (RMS) Licenses required for B2B and B2C calling	Nontraversal Call Licenses required Traversal Call Licenses required
Microsoft Interop License	Requires RMS licenses	Requires Option Key
FindMe License	Available	Requires Option Key
Device Provisioning License	Requires Option Key (Free)	Requires Option Key (Free)
Clustering Capabilities	Up to 6 servers	Up to 6 servers

I do not know what Cisco's future plans are for these products. At the time this book was written, Cisco had not announced any plans to make either of these products end of sale, and they are both marketed as viable solutions for customers.

Webex Control Hub

Although cloud collaboration goes beyond the scope of this book, it is worth mentioning that Cisco does have a cloud-based call control solution for customers as well. Cisco Webex is a multifaceted cloud-based solution that warrants a deeper explanation on each of the Webex products Cisco has to offer. However, this section will only introduce Cisco Webex as a call control platform. The Cisco Webex solution can be divided into three categories: Meeting, Collaborating, and Calling.

Webex Meetings is the same powerful tool that has been used for years to allow multipoint conferencing in the cloud. Participants can join via a Webex Meeting client, through a browser, using Webex Teams, using Webex Calling, or using a phone or Telepresence endpoint. All the same tools that have traditionally been used with Webex Meetings are still available; plus, Cisco has added a few more enhancements. Webex Meetings allows for high-quality voice and HD video communication, content sharing, polling, annotation, and many more supported features. CMR Cloud through Webex Meetings can be extended to a CMR hybrid meeting deployment using the Video Mesh Node.

Webex Teams, formerly known as Spark, is a client that can be installed on a Windows or Apple computer, tablet, or smartphone. Webex Teams allows for point-to-point messaging or group messaging in Spaces. This highly secure messaging solution allows conversations to be escalated to a voice or video meeting where content can be shared, and a whiteboard application can be leveraged. All whiteboard notations can be saved into a Webex Teams Space so that collaboration can continue after the meeting ends. Webex Teams also supports the upload and download of documents, and many other integrations and bots can be leveraged through this application.

Webex Calling is a tool that has always been available with the Webex solution since the inception of Spark. However, since Cisco acquired BroadSoft, it has been working diligently to bring the BroadSoft calling features into Webex. Webex Calling allows certain phones to register to the Webex Control Hub, and from those phones, users can call out over IP or through the PSTN. Alternatively, a Webex Calling application can be used from any Windows or Mac computer, tablet, or smartphone. Additionally, Webex Calling powered by BroadSoft offers many more calling features than were previously available with Webex Calling.

The Webex Control Hub is the centralized, cloud-based management tool for all Cisco Webex–related products. Calling features, such as endpoint and phone registration and voicemail, are all managed through the Webex Control Hub. Users are also managed from here. Users can be imported through an LDAP integration, and single sign-on can be set up as well. User privileges can be assigned as needed, including additional administrators and security compliance officers. All Webex Meetings, Teams, and Calling features can be configured and managed on an individual basis, based on location or as an organization as a whole. Many different hybrid integrations are available between an on-premises deployment of Cisco Collaboration and the Webex cloud, which are all initiated through the Webex Control Hub. There are too many feature capabilities available through the Webex Control Hub to list here, but one last capability worth mentioning is the many analytics and

troubleshooting tools. They allow administrators to track, manage, and control the Cisco cloud collaboration solution. Cisco Webex is not just for small and medium-sized businesses; it is for any sized business that wants to extend and enhance its global collaboration efforts.

Introduction to Cisco Applications

UC applications are used to unify your voice, video, data, and mobile applications for collaboration within the Cisco Collaboration solutions. Applications include communication gateways, voicemail, and unified IM and Presence services. Other customer collaboration applications, which will not be discussed in this book beyond this chapter, are used to create the foundation for strong customer relationships. These products include contact center and voice self-service products. Media service applications are used to enable collaboration anywhere with more security and high-quality integrated voice, video, and content sharing. These applications include video conferencing products, web conferencing applications, and conferencing management tools.

Cisco Unity Connection Server

Key Topic

Cisco Unity Connection (CUC) is a robust unified messaging and voicemail solution that accelerates collaboration by providing users with flexible message access options and the IT department with management simplicity. You can access and manage messages from your email inbox, web browser, Cisco Jabber, Cisco IP phone, smartphone, or tablet with Cisco Unity Connection. You also can easily prioritize messages and respond quickly to colleagues, partners, and customers. Mobile users, or anyone who simply prefers to do so, can use the speech-activated tools for hands-free message retrieval.

For IT, CUC is an "integrated by design" extension of Cisco Unified Communications Manager. It is easy to manage using Cisco Prime Collaboration, Cisco's single application for unified management of the entire voice and video deployment. Cisco Prime Collaboration simplifies deployment, provisioning, monitoring, and system management. Chapter 23, "Understanding Cisco Unity Connection," and Chapter 24, "Cisco Unity Connection End-User and Voice Mailbox," will delve much deeper into CUC.

Key Topic

Cisco Unity Express (CUE) offers industry-leading integrated messaging, voicemail, fax, automated attendant, interactive voice response (IVR), time-card management, and a rich set of other messaging features on the Cisco Integrated Services Router (ISR) platform. It provides these integrated services specifically designed for small and medium-sized office environments or enterprise branch offices. With Cisco Unity Express, you can easily and conveniently manage your voice messages and greetings right through your web browser using Web Inbox, traditional intuitive telephone prompts, an easy-to-use visual voicemail interface (which is called the Cisco Unity Express VoiceView Express application), email access to messages, and a straightforward GUI that allows simple administration and management.

Cisco Unity Express is an essential component of either a Cisco Unified Communications Manager or Cisco Unified CME Solution. In a Cisco Unified Communications Manager environment, Cisco Unity Express provides local storage and processing of integrated messaging, voicemail, fax, automated attendant, and IVR for branch offices with limited WAN connectivity, thereby alleviating concerns about WAN bandwidth and quality of service (QoS). Additionally, Cisco Unified Communications Manager customers with Cisco Unity Connection unified messaging solutions at their larger locations can use Cisco Unity

Express at their branch-office locations and network the solutions so that employees can easily send messages between locations. In a Cisco Unified CME environment, customers deploy a single Cisco ISR platform with Cisco Unity Express installed to meet their office telephony and messaging needs, as well as their other business communications needs.

Cisco IM and Presence Service

Cisco Unified Communications Manager IM and Presence Service, or IMP, provides native standards-based, dual-protocol, enterprise instant messaging and network-based presence as part of Cisco Unified Communications. This secure, scalable, and easy-to-manage service within Cisco Unified Communications Manager offers feature-rich communications capabilities both within and external to the enterprise.

IM and Presence Service is tightly integrated with Cisco and third-party-compatible desktop and mobile presence and IM clients, including the Cisco Jabber platform, Cisco Webex Social, and Cisco Jabber SDK. It enables these clients to perform numerous functions such as instant messaging, presence, click-to-call, phone control, voice, video, visual voicemail, and web collaboration. IM and Presence Service offers customers and partners the flexibility of rich, open interfaces that enable IM and Cisco's rich network-based presence, as well as IM and presence federation for a wide variety of business applications. Chapter 25, "CUCM IM and Presence Service," will delve much deeper into the Cisco IM and Presence Service.

Cisco Meeting Server

Conferencing is an essential component of any collaboration solution, especially when serving remote users or a large user base. Cisco Rich Media Conferencing offers features such as instant, permanent, and scheduled audio and video conferencing, as well as content sharing. Conference bridges provide the conferencing function. A conference bridge is a resource that joins multiple participants into a single call. It can accept any number of connections for a given conference, up to the maximum capacity allowed for a single conference on that device. The output display for a given party shows all connected parties minus the viewer's own input. Cisco Rich Media Conferencing solutions utilize various infrastructures to provide audio and video conferencing capabilities and content sharing. The conferencing infrastructure can be Cisco Unified Communications Manager using software or DSP resources, Cisco Meeting Server, or Cisco Webex Collaboration Cloud. Cisco Rich Media Conferencing solutions are available as on-premises, cloud, or hybrid deployments. This allows organizations to integrate with the Collaboration solution in which they have already invested or, alternatively, to implement a service that is hosted in the cloud. This is one of the more important distinctions between the various solutions, and it is the first decision point when determining which solution is the best fit for an organization. Cisco Webex Software as a Service (SaaS) offers a completely off-premises solution, while Cisco Collaboration Meeting Rooms (CMR) Hybrid is a solution with a mix of on-premises and off-premises equipment. Organizations that have deployed Cisco Collaboration Systems Releases (CSRs) will benefit most from leveraging an on-premises solution. This section focuses specifically on introducing the CMS product.

Cisco Meeting Server (CMS) is a Rich Media Conferencing product for on-premises deployments only. It cannot be part of a CRM-Hybrid deployment. This robust CMR solution can be deployed as an appliance server or as a virtual machine. Virtual deployments use VMware ESXi hypervisors for deployment. The Cisco Meeting Server appliance server has several options available. Customers who purchased the Acano solution before the Cisco acquisition

can still use the Acano-X appliance server. Newer customers who wish to deploy the Cisco Meeting Server as an appliance can use the Cisco UCS platforms CMS1000 or CMS2000.

CMS offers a consistent one-meeting experience with many features available. However, exactly what features are available may depend on the device that is used to connect to the meeting. Calls can be made using the Cisco Meeting app, a physical endpoint, or through a third-party application. CMS supports both Telepresence Interoperability Protocol (TIP) and non-TIP supported devices. Once connected to a Space, which is what a virtual meeting room is called on CMS, users can set camera and microphone settings, mute or activate a personal microphone, share a screen or an application, chat, change devices, change screen layout, and see caller information to name a few. Bring your own device (BYOD) allows users to use their own devices to see presentations, chat with other participants, or even transfer a call from an endpoint to a smartphone or tablet using CMA.

Additional features of CMS include a seamless integration with Microsoft Skype for Business (S4B), support for WebRTC, and clustering the Database and Call Bridge services for a scalable and resilient deployment. CMS is a rich and robust CMR solution for on-premises deployments. Cisco Meeting Server is a topic that goes outside the scope of this book, so I will not go any deeper into this topic.

Management Software

Management software is a suite of tools that are used to unburden the IT administrator from some of the overwhelming day-to-day maintenance tasks of managing a collaboration network. Management tools are not essential for collaboration solutions to function, but as a network grows, these management tools do ease the stress among the people responsible for keeping all the components of the network functional. Think of what kind of car you need to take a long road trip. You don't need the comforts of leather seats, air conditioning, XM radio, or fine-tuned suspension. All you "need" is something that can get you from point A to point B. However, all those extra amenities sure do make that journey a lot more comfortable. You may even find yourself quite refreshed upon reaching your destination. Management software simply adds extra comforts to the management side of collaboration. There are two management products available in the Cisco Collaboration product portfolio: Telepresence Management Suite (TMS) and Prime Collaboration.

Cisco TMS offers everything from complete control and management of multiparty conferencing, infrastructure, and endpoints, to centralized management of the telepresence network. Flexible scheduling tools are designed to meet the needs of basic users for quick conference creation, including integration with Microsoft Exchange for scheduling through Outlook clients, and to provide advanced conference booking options for IT administrators. This includes One-Button-to-Push meeting access, which is supported in Cisco TMS Version 13.1 and later. Phonebooks can be created and pushed out to all Cisco Collaboration endpoints, and each phonebook can contain a multitiered layer of directories within it. Additional services provided by TMS include backup and restore features, scheduled system upgrades, configuration templates, dial-plan management, and many troubleshooting tools and reports.

Cisco TMS is provided as a software-based application for installation on a customer-provided Microsoft Windows Server with a SQL back end. Any physical server running an appropriate version of Microsoft Windows Server can run TMS, but Cisco recommends using the Cisco UCS server. The Cisco TMS user interface is a web browser-based

application that uses Microsoft Internet Information Services running on the .NET framework. The caveat to using TMS is that it only supports Telepresence products. UC phones and services are not supported through TMS. However, third-party Telepresence endpoints can be managed by TMS. Cisco Prime Collaboration is the management software to use for support of UC services and phones.

Cisco Prime Collaboration helps enterprises address the continuous transformation of their networks as they invest in next-generation collaboration technologies with integrated video and voice deployments. It empowers IT departments to effectively manage this transformation and the video network lifecycle as they meet demands from end users for high-quality solutions everywhere and at all times. At the same time, it addresses the need to reduce operating expenses and optimize limited resources. Prime Collaboration provides simplified, unified management for video and voice networks. This management solution helps ensure superior quality experiences for end users and lower operating expenses for supporting video and voice communication. Prime Collaboration removes management complexity and provides automated, accelerated provisioning, real-time monitoring, proactive troubleshooting, and long-term trending and analytics in one integrated product. This solution delivers a premier operations experience through an intuitive user interface and automated workflows that ease implementation and ongoing administration. Self-Provisioning and Self-Care features allow users to provision their own phones, like the 7800 and 8800 series phones, and change phone settings such as names, directories, and speed-dials. The three modules to contribute to an enhanced management experience from Prime Collaboration include

- Prime Collaboration Provisioning
- Prime Collaboration Assurance
- Prime Collaboration Analytics

As mentioned earlier, Cisco Prime Collaboration Provisioning unifies administration of your UC environment to one interface, including Cisco Unified Communications Manager, Cisco Unified Communications Manager Express, Cisco Unified Communications Manager IM and Presence service, Unity Connection, and Unity Express. It provides one interface for provisioning video, call control, messaging, and presence for a single cluster implementation. It has an intuitive interface that provides a single view of a user and the user's services, as well as a consolidated view of users across the organization. With these capabilities, Cisco Prime Collaboration significantly accelerates site rollouts and dramatically reduces the time required for ongoing moves, adds, and changes by facilitating the delegation of these tasks. This allows organizations to optimize IT resources, resulting in exceptional productivity gains and lowered operating expenses. Prime Collaboration Provisioning comes in Standard and Advanced versions. The Advanced version comes with everything in the Standard version and also includes multicluster and multiversion support for Cisco Unified Communications Manager and Unity Connection, advanced RBAC, ordering workflow, templates, and API support.

Prime Collaboration Assurance provides additional monitoring tools for the Cisco Unified Communications environment, including Cisco Unified Communications Manager, Unity Connection, and Video. It can continuously monitor in real time and do advanced troubleshooting of the environment, sending a notification when a problem arises so that issues can be proactively resolved. For video, it allows viewing of the end-to-end session paths

including jitter and packet loss, with a web-enabled interface for fault monitoring of video components, including a dashboard view. Prime Collaboration Assurance provides efficient, integrated service assurance management through a single, consolidated view of the Cisco video and voice collaboration environment and is offered in Standard and Advanced versions. The Advanced version includes all the components of the Standard version, plus diagnostics, reporting features, multiple cluster management, five access levels, and higher levels of discovery, inventory, fault management, and dashboards.

Cisco Prime Analytics enables network managers to maximize the value of the massive amounts of information available about their network traffic by

- Continuously monitoring live data
- Instantaneously processing queries
- Providing real-time analysis and action
- Efficiently using compute resources

The Analytics module provides historical reporting of key performance indicators and helps enable IT network managers to analyze trends for capacity planning, resource optimization, and quality of service. The solution helps track collaboration technology adoption rates in the network and provides metrics to help analyze how users are actually using the collaboration endpoints on a daily basis. It also provides insights into key collaboration network resource usage trends. With historical data and many reporting options with easy customization, IT managers have access to actionable information, simplifying the long-term planning process contributing to ongoing technology investment decisions, and helping to optimize the network configuration for an improved experience quality for end users. Cisco Prime Collaboration Analytics provides real-time monitoring and support capabilities for real-time communication.

Designing a Cisco Collaboration Solution

There are many factors to consider when designing a Cisco Collaboration solution for a customer. The Collaboration solution could be an on-premises solution, a cloud-based solution, or a hybrid of the two. The type of Collaboration solution being designed will directly impact the licensing used for that solution. The types of services offered to the customer will also impact the licensing model. Then there are sizing considerations that have to be taken into account. When sizing a solution, you should not forget to leave room for expected growth within the customer organization. Bandwidth allocations and CAC need to be planned so that call loads during peak hours do not overtax the network. Sometimes systems break. Therefore, high availability needs to be designed into the Collaboration solution so that appropriate redundancies are in place in the event of key system failures. When systems fail, data can be lost. Disaster recovery components will help ensure data retention during these outages. Then there's the dial plan, which may be one of the most important aspects to designing a Collaboration solution. The dial plan is used during every call, and the complexity of the dial plan will impact many aspects of the overall implementation and usability of the solution. Security is a growing concern in any networked solution. Designing appropriate security measures into the Collaboration solution is equally important. Finally, there is the quality of service design that will coincide with the Collaboration solution. QoS controls how data traffic is routed through the network during high congestion times. As you can

see, there are many aspects to designing a Collaboration solution that must be taken into consideration. Although each one of these topics could fill a chapter or more on its own, the following sections will dip into each one of them to provide a roadmap of the various aspects to consider when designing a Cisco Collaboration solution.

Licensing

Businesses today are as diverse as two snowflakes. These differences bring with them different needs in a workplace that is constantly changing. Layer this complication with a multitude of products ranging from call control devices such as the Cisco Unified Communications Manager, to conferencing applications, such as the Cisco Meeting Server, to cloud solutions, such as Cisco Webex, and contriving a singular license plan that covers the many needs of a business becomes a colossal task. Cisco rose to this charge and devised a licensing solution that can be tailored to any company's needs, regardless of the size of the organization, the solution it's using, or the components needed to meet its needs. As with any Cisco product, understanding the requirements for licensing your product at the time of installation is important.

A great example of the licensing obstacles Cisco has had to overcome, due in part to acquisitions, is the licensing differences between the Cisco VCS and the Cisco Unified Communications Manager. Historically, the Cisco VCS used "device-based" licenses, where licenses were purchased based on the number of devices that were allowed to register. Then additional call licenses were required based on the number of concurrent calls the VCS would allow, complicated even more by the type of call, and traversal or nontraversal call licenses. Cisco devised the Cisco Expressway series that could perform all the functions of the VCS because it is essentially a VCS and uses the same user-based licenses as the Cisco Unified Communications Manager. These licenses are the Cisco Unified Workspace Licensing, or CUWL.

CUWL licenses are broken down into two varieties, with a third possibility being the Cisco User Connect Licensing, or CUCL, which are designed for voice-only solutions in the Cisco Collaboration suite of devices. This program is based on users rather than devices and allows customers to simply purchase licenses based on the number of users they wish to service, with each user having access to multiple devices or services as part of the program. Table 6-4 identifies the CUWL and CUCL licenses and the capabilities included with each license as well as the purchasable options available for each license.

Key Topic

Table 6-4 CUWL and CUCL Licensing Model

	CUCL Essentials	CUCL Basic	CUCL Enhanced	CUWL Standard	CUWL Professional
Number of Devices Supported	One	One	One or Two	Multiple	Multiple
Cisco Prime Collaboration	Included	Included	Included	Included	Included
Jabber/IMP	Included	Included	Included	Included	Included
Jabber UC	N/A	N/A	Included	Included	Included

	CUCL Essentials	CUCL Basic	CUCL Enhanced	CUWL Standard	CUWL Professional
Expressway Firewall Traversal	N/A	N/A	Included	Included	Included
Unity Connection	Optional	Optional	Optional	Included	Included
Webex Conferencing	Optional	Optional	Optional	Optional	Included
PMP Basic	N/A	N/A	Optional	Optional	Included
PMP Advanced	N/A	N/A	Optional	Optional	Optional

Table 6-4 displays nine different components that can come with different licensing levels. Cisco Prime Collaboration is a Prime Collaboration Provisioning standard license that is included with the Cisco Unified Communications Manager. The difference between Jabber/IMP and Jabber UC is that Jabber/IMP refers to the desktop client on Microsoft Windows or Apple Mac computers, and Jabber UC refers to the Jabber application on smartphones and tablets. Expressway Firewall Traversal allows for firewall traversal licenses through the Expressway Core and Expressway Edge servers. This includes MRA capabilities and one RMS license per user. Unity Connection allows voicemail and other services that come with Unity Connection. Webex Conferencing includes one named user license for both Webex Meeting Center and Webex Meeting Server. Webex can be used for cloud-based meeting or Hybrid meetings using the Video Mesh server installed on premises.

To use the on-premises Cisco Meeting Server for multipoint conferences, multiparty licenses are required. Two types of multiparty licenses are available: Shared Multiparty Plus (SMP) and Personal Multiparty Plus (PMP). With the Cisco licensing model outlines in Table 6-4, PMP licenses can be purchased as Basic or Advanced. PMP Basic provides one PMP license per user that will support host meetings on the Cisco Meeting Server for up to four participants in each meeting. PMP Advanced provides one PMP license per user that will support host meetings on the Cisco Meeting Server for an undefined number of participants in each meeting. The number of participants for PMP Advanced licenses is limited only by the infrastructure that has been installed. A huge advantage to PMP licenses is the amount of security and control they offer to the meeting. Other users cannot join personal meeting spaces that have been assigned to each user on the Cisco Meeting Server unless the host to whom the space belongs has joined the meeting. Once the host drops out of a meeting hosted in that space, all other participants will be dropped as well. This feature prevents the meeting resources from being abused by other people. SMP licenses are not part of the licensing model outlined in the table because they are not assigned to any one user. When SMP licenses are added to the Cisco Meeting Server, any user can create and join a meeting using this license. Because SMP does not share the same level of control that PMP licenses do, they should be purchased and used sparingly.

There are some other differences between CUWL and CUCL licenses that should be noted. CUWL standard and professional licenses support multiple endpoints. For example, this license allows for two desktop endpoints for a single user license, allowing for a single user to have an endpoint both at the office and at home. CUCL packages are designed for voice-only solutions, whereas the Expressway series is only used to provide VPN-less traversal services for voice endpoints. Video can be used over CUCL using Jabber or video UC phones, such as 8845 or 8865, but this is not the designed purpose of these licenses. Telepresence endpoints cannot register under the CUCL model. Notice that CUCL supports only one or two devices per user. The idea here is that a user may need to register a VoIP phone and Jabber, or that user may just use a VoIP phone.

Cisco has been using CUWL and CUCL licenses for many years because it has led the market in on-premises infrastructure. In more recent years, a new market has opened up in cloud-based offerings, and Cisco has been working diligently to dominate this market as well. With the need to be able to deliver collaboration services cost-effectively, using on-premises infrastructure or cloud-based services, depending on the needs of employees, Cisco has devised a new layer to its licensing model known as Flex.

The Cisco Collaboration Flex Plan entitles people to use Cisco's industry-leading collaboration tools with one simple subscription-based offer. It helps with transitions to the cloud, and investment protection, by including cloud, premises, hosted, and hybrid deployments, with the flexibility to use them all. Companies can choose to equip employees with meetings, calling, or both, and add more licenses at the time they're needed. Companies can also easily add Contact Center capabilities, which are also included in the Collaboration Flex Plan. One agreement covers software, entitlements, and technical support for cloud-based and on-premises services. Companies simply choose the services they need today and grow at their own pace. There is no need to manage complex agreements. And you can mix and match meetings and calling subscriptions for flexibility and value. With the Flex Plan, you can choose the right subscription based on your business size and needs. Each option includes technical support. Choose from the following purchasing models:

- For enterprisewide deployment, Cisco Enterprise Agreement customers can purchase via the Cisco Collaboration Flex Plan. You can gain maximum value by enabling services for everyone in your organization for meetings or calling or both.
- To purchase meetings according to usage: Cisco Collaboration Flex Plan—Active User Meetings: Anyone can host a meeting, and you pay only for those who use the entitlement.
- To provide meetings and/or calling services to individuals, teams, or departments: Cisco Collaboration Flex Plan–Named User: Your purchase is based on the number of people who need services. You can grow at your own pace.
- To provide contact center services to your service agents: Cisco Collaboration Flex Plan—Concurrent Agent: Your purchase is based on the number of agents simultaneously using services at your peak busy hour. Again, you can grow at your own pace.

At the same time, you can seamlessly drive enhanced team collaboration with Cisco Webex Teams, which is included at no additional charge. Cisco Webex Teams is a great tool to collaborate with other coworkers for ongoing work. Teams can be used on every device, in

every place, to move work forward. You can enable services for selected individuals, teams, or departments, or for your entire organization. And you have the flexibility to add services as adoption grows. To learn more about the Cisco Collaboration Flex Plan, visit

https://cisco.com/go/collaborationflexplan

Cisco introduced a new way to add licenses to on-premises collaboration products called Smart Licensing. Smart Licensing was introduced as an option with Cisco Unified Communication product version 11.5, but it is required for licensing products from version 12.0 onward. Cisco is transforming the end-to-end software lifecycle to make the customers' experience better and easier. A major part of this change is a move away from Product Activation Key (PAK) licenses to Smart Licensing to make the license registration process faster and more flexible. At the heart of the transformation is Smart Licensing and Smart Accounts, which offer streamlined purchasing and software administration. Smart Licensing is a flexible software licensing model that simplifies the way you activate and manage licenses across your organization. The Smart Licensing model makes it easier for you to procure, deploy, and manage your Cisco software licenses. To use Smart Licensing, you must first set up a Smart Account.

A Smart Account is a central repository where you can view, store, and manage licenses across the entire organization. Comprehensively, you can get access to your software licenses, hardware, and subscriptions through your Smart Account. Smart Accounts are required to access and manage Smart License–enabled products. Creating a Smart Account is easy and takes less than five minutes. You can create a Smart Account on cisco.com. Smart Accounts offer a simple-to-use, centralized, and organized solution to license management. With a Smart Account, you get full visibility and insight into all of your Cisco software assets deposited into the Smart Account, including PAK licenses and Enterprise Agreements. When Smart Accounts are used with Smart Licenses, the benefits include

Key Topic

- **Real-Time Visibility:** You can view all of your software licenses, entitlements, and users across the organization.
- **Centralized Management:** A single location enables authorized users to see all license entitlements and move licenses freely through the network as needed.
- **Cost-Effectiveness:** You can drive down the cost of license management with reduced overhead, better utilization management, and more efficient planning.
- **Organization:** Virtual Accounts provide the flexibility to organize licenses by department, product, geography, or other designation, whatever makes the most sense for your company.

Sizing

Capacity planning involves sizing a solution to meet all of the current needs of an organization and have room to scale based on projected growth. This will determine the type of server that should be used for the installation. To illustrate how a server can be selected based on the capacity needed, this chapter will go into some of the Cisco UCS servers to a limited degree. The two primary questions that should be asked when sizing a solution are, "What is the maximum number of users who will utilize these services?" and "What is the maximum number of endpoints that will be used?"

For small and medium-sized businesses (SMB) that need only voice services, Cisco offers a great product to meet these needs. The Cisco Business Edition 4000 (BE4K) is a cloud-managed system that can support up to 200 VoIP phones. It can be preconfigured for the customer by the Cisco partner reseller through the cloud management portal and can be managed by the partner or customer after it has been installed and provisioned. This system will support the 7800 and 88X1 series phones and comes equipped with support for 120 hours of voicemail messages. You can add a 1-, 2-, or 4-port T1/E1 PRI card, or a 2- or 4-port BRI card for digital PSTN connections. You can add a 2- or 4-port FXO card, a 2- or 4-port FXS card, or a 2-port FXS with 4-port FXO combination card for analog PSTN connections. Alternatively, you can build a SIP trunk to your service provider for SIP-to-PSTN connections using the built-in CUBE gateway services.

Another UCS server Cisco offers is the Business Edition 6000 (BE6K), which comes in a Medium density (BE6000M) or High density (BE6000H) platform. The BE6K servers come preloaded with VMware ESXi and all the OVAs and ISOs needed to complete an installation of the VMs once the server has been installed and powered on. The BE6000M has a max capacity of 1000 users per cluster and 1200 endpoints per cluster. The BE6000H has a max capacity of 1000 users per cluster and 2500 endpoints per cluster. Extra nodes can be added for redundancy, but the capacity limits do not change. One of the great benefits to using a BE6K server is that the sizing tool is not needed because capacity limits are preset.

If higher capacity is needed than what the BE4K or BE6K can offer, Cisco has made available a third option called the Business Edition 7000 (BE7K). Like the BE6K, the BE7K comes preloaded with VMware ESXi and all the OVAs and ISOs needed to complete an installation of the VMs after the server has been installed and powered on. The BE7K can be purchased in a Medium density (BE7000M) or High density (BE7000H) platform. It is sized using the sizing tool, and capacity limits increase when extra nodes are added to this server. The BE7K can support up to 10,000 users per node and up to 40,000 users per cluster. It can also support up to 10,000 endpoints per node and up to 40,000 endpoints per cluster.

It should be plain to see that as the complexity of installing the server increases, the capabilities of the server increase as well. Installing the BE4K server is very easy, but it has a limited calling capability, whereas installing the BE7K is much more complex, but it has a significantly higher calling capability. What has not been mentioned before is that in the progression of each server, an increase in features is also supported. The link for the sizing tool is https://tools.cisco.com/cucst. You will have to log in with a CCO account and be associated with a Cisco partner company to access the sizing tool.

Bandwidth

Three main aspects to bandwidth must be considered when designing a Collaboration solution. First, you must consider the audio and video components being used. This includes everything discussed in Chapters 3–5 and much more. Codecs used, scan rate, compression algorithms, environmental conditions, and many other aspects can positively or negatively impact bandwidth. Second, QoS can affect bandwidth utilization. This topic will be discussed later in this chapter and in more depth in Chapter 12, “Cisco Core Network Components.” Third, provisioning and admission control allow you to set up parameters within the collaboration environment so that you can more closely observe and manage bandwidth.

Provisioning bandwidth and ensuring that the correct bit rate is negotiated between various groups of endpoints are important aspects of bandwidth management. In a Cisco Unified Communications Manager environment, bit rate is negotiated via Cisco Unified Communications Manager, which uses a concept of regions to set the maximum audio and maximum video bit rates for any given call flow. Cisco Unified Communications Manager locations work in conjunction with regions to define the characteristics of a call flow. Regions define the type of compression or bit rate that is used between any two devices. Location links define the amount of available bandwidth for the path between devices. Each device and trunk in the system is assigned to a region, by means of a device pool, and a location, by means of a device pool or by direct configuration on the device itself.

- Regions allow you to set the per-call bandwidth of voice and video calls. The audio limit on the region can result in filtering out codecs with higher bit rates. However, for video calls, the video limit constrains the quality (resolution and transmission rate) of the video.
- Locations define the amount of total bandwidth available for all calls to another location. When a call is made on a link, the regional value for that call must be subtracted from the total bandwidth allowed for that link.

Building a region matrix to manage maximum voice and video bit rate (video resolution) for groups of devices can assist in ensuring that certain groups of devices do not oversaturate the network bandwidth. When creating a region matrix, you should group devices into maximum video bit rate categories. The smaller the number of groups, the easier it is to calculate bandwidth requirements. Also, you should consider the default region settings to simplify the matrix and provide intra-region and inter-region defaults. There are other region considerations for bandwidth provisioning. The first consideration is whether to have different intra-region settings versus inter-region settings. This will determine whether per-site regions are required or not. The concept here is that if intra- and inter-regional audio or video bit rates are to be different, per-site regions will be required. This augments the configuration of regions to the number of sites (N) multiplied by the number of video groups (X): N*X = number of regions required on average. If intra- and inter-regional audio and video bit rates will be the same, only the regions for the video groups are required (X).

Another consideration is to reuse regions configured for audio-only IP phones when possible. Audio codec configuration is shared, so if video calls need to use different audio codecs, you need to configure new regions. For example, if voice-only devices use the G.729 audio codec over the WAN and G.711 or G.722 on the LAN while video devices always use G.711 or G.722, the voice-only and video endpoints cannot share a region. Thus, each site would require a region per group of devices. Sites = N, and video region groups = 4 + voice-only region group; then N*4 is the number of regions required. You can use the Prime Collaboration Provisioning tool or the Bulk Administration Tool as configuration aids. Per-site regions might not be needed if a single audio codec is used for both intra-region and inter-region calls as well as voice-only calls. If both audio and video endpoints use G.711 or G.722 over the WAN and LAN for voice-only or video calls, voice-only IP phones and video endpoints could use the same region. You should consider the default region settings to simplify the matrix.

The Call Admission Control (CAC) function can be an important component of a Collaboration system, especially when multiple sites are connected through an IP WAN and limited

bandwidth resources are available for audio and video calls. Consider for a moment that traditional TDM-based PBXs operate within circuit-switched networks, where a circuit is established each time a call is set up. As a consequence, when a legacy PBX is connected to the PSTN or to another PBX, a certain number of physical trunks must be provisioned. When calls have to be set up to the PSTN or to another PBX, the PBX selects a trunk from those that are available. If no trunks are available, the call is rejected by the PBX and the caller hears a network-busy signal.

Now consider an IP-connected Unified Communications system. Because it is based on a packet-switched IP network, no circuits are established to set up an IP telephony call. Instead, the IP packets containing the voice samples are simply routed across the IP network together with other types of data packets. Quality of service (QoS) is used to differentiate the voice packets from the data packets, but bandwidth resources, especially on IP WAN links, are not infinite. Therefore, network administrators dedicate a certain amount of "priority" bandwidth to voice traffic on each IP WAN link. However, after the provisioned bandwidth has been fully utilized, the IP telephony system must reject subsequent calls to avoid oversubscription of the priority queue on the IP WAN link, which would cause quality degradation for all voice or video calls. This function is known as Call Admission Control, and it is essential to guarantee good voice and video quality in a multisite deployment involving an IP WAN. To preserve a satisfactory end-user experience, the CAC function should always be performed during the call setup phase so that, if network resources are not available, a message can be presented to the end user, or the call can be rerouted across a different network (such as the PSTN).

High Availability

The next consideration when planning a call processing deployment should be high availability and redundancy within the solution. This involves planning clusters of the call agents being used and configuring proper redundancy. This also involves planning redundancy in power being supplied to the hosting servers along with uninterruptible power supply (UPS) sources. Finally, this involves planning high availability in the network connectivity.

As it pertains to call processing, clustering is a grouping of call agents that work together as a single call processing entity with higher capacity. Multiple Cisco Unified Communications Managers can be clustered together, multiple VCSs can be clustered together, and multiple Cisco Expressways can be clustered together. However, there is no cross-cluster between different call agents, such as the VCS and Cisco Expressway, the VCS and Cisco Unified Communications Manager, or the Expressway and the Cisco Unified Communications Manager. However, different call agent clusters can be trunked together in order to unify communications. The Cisco Unified Communications Manager cluster can be trunked to a VCS cluster, Cisco Expressway cluster, and another Cisco Unified Communications Manager cluster.

The Cisco Expressways and VCSs can support up to six peers in a cluster. One of the peers in the cluster is designated as the master, and all settings under the Configuration menu on the master are replicated to each of the other peers in the cluster. If any of these configuration settings are changed from any peer in the cluster that is not the master, those changes will immediately be overwritten by the master of the cluster. Settings under the Applications menu can be configured from any peer in the cluster, and these settings will be replicated to

all other peers in the cluster. Therefore, a round-trip delay time between any peer in the cluster should not exceed 30 ms, or 15 ms each one-way direction. In the event the master goes down, the next subsequent peer listed will assume the role of the master. When configuring a cluster of Cisco Expressways or VCSs, you need to configure each peer with a unique IP address, URL, and system name, but they should share the same cluster name, which should be in the form of a URL. The system URL should resolve to a DNS A-record for that particular call agent. The cluster URL should resolve to all call agents in the cluster using round-robin for distribution and load balancing.

Key Topic

Clusters in the Cisco Unified Communications Manager behave differently than that of the VCSs or Cisco Expressways. A Cisco Unified Communications Manager cluster is made up of two different service node types: the publisher and the subscriber. The publisher is essentially the master of the cluster, and each cluster can have only one publisher. Each subsequent node in a cluster is referred to as a subscriber. If you are installing a Cisco Unified Communications Manager for the first time, and you do not plan to establish a cluster, the one node is still designated as the publisher. You can then establish a cluster at a later time by setting up the required number of subscriber nodes. The publisher node is the only node in a cluster with full read-write access to the configuration database. Should the publisher go down, all services will continue to operate normally, and user-facing configuration changes to the database can be made during a publisher outage. Information is then synced to the publisher when connectivity is re-established. No other services can be written to the database when the publisher is down. Peers in a Cisco Unified Communications Manager cluster are referred to as *service nodes* because they can be grouped based on services they offer. Common service groupings include Call Processing, TFTP, Media Resources, Computer Telephony Integration (CTI) Manager, and Unified CM Applications. Figure 6-1 illustrates a Cisco Unified Communications Manager cluster with various service nodes grouped together.

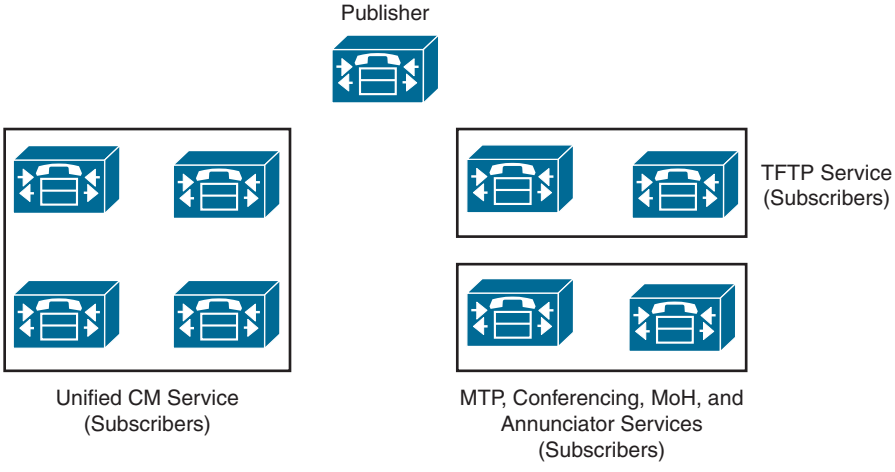

Figure 6-1 *CUCM Cluster with Service Nodes*

Key Topic

In addition to establishing service groups within a Cisco Unified Communications Manager cluster, you can also configure redundancy groups for call processing failover in the event a Cisco Unified Communications Manager within the cluster crashes. There are two options for configuring redundancy groups: a two-to-one (2:1) option and a one-to-one (1:1) option.

With the 2:1 option, there is one shared backup call processing subscriber for every two primary call processing subscribers. In the event one of the primary call processing subscribers crashes, the backup will immediately assume the role until such time the failed subscriber can be restored. However, should both primary call processing subscribers fail, the backup will be able to assume the role of only one of the primary call processing subscribers; therefore, a loss in call processing capabilities will be encountered. When the cluster is being upgraded, these redundancy groups can help maintain call processing while each primary subscriber is being rebooted. The upgrade order of sequence is to fully upgrade and reboot one of the primary subscribers first, then do the same for the second primary subscriber, and last should be the backup subscriber. Figure 6-2 illustrates how a 2:1 group of call processing subscribers can be used within the Cisco Unified Communications Manager dependent on the size of the deployment. Only two examples are provided, but many combinations exist based on the capacity of the cluster being deployed.

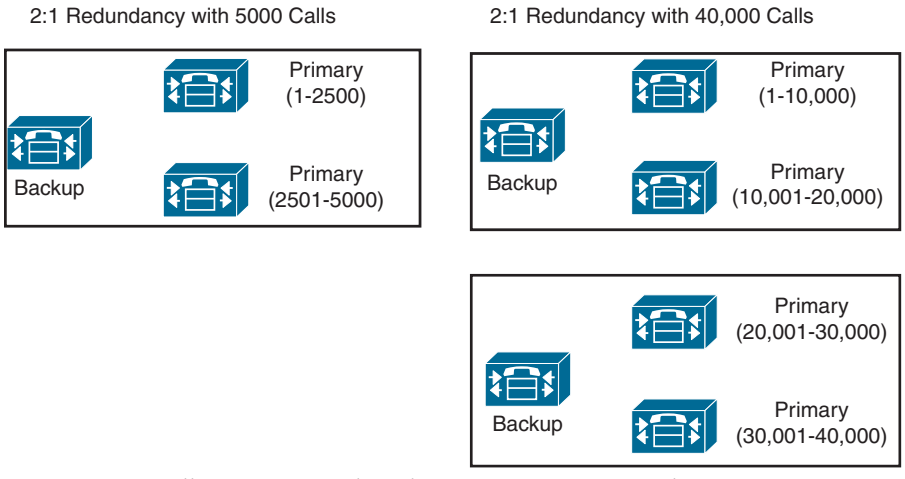

Figure 6-2 *Call Processing Subscriber Groups in a 2:1 Deployment*

Key Topic

With the 1:1 option, there is a single backup call processing subscriber for each primary call processing subscriber. This does require more server space and processing because more Cisco Unified Communications Manager deployments are needed. However, in the event any of the primary call processing subscribers crashes, the backup will immediately assume the role until such time the failed subscriber can be restored. This is a more robust solution that helps mitigate the prospect of downtime in your call center. In a 2:1 group, device registration and call processing services are available only on the primary subscribers unless a subscriber crashes; then the backup will kick in. However, in a 1:1 group, registration and call processing can be load-balanced between the primary and backup subscriber. Figure 6-3 illustrates how a 1:1 group of call processing subscribers can be used within the Cisco Unified Communications Manager dependent on the size of the deployment. Only two examples are provided, but many combinations exist based on the capacity of the cluster being deployed.

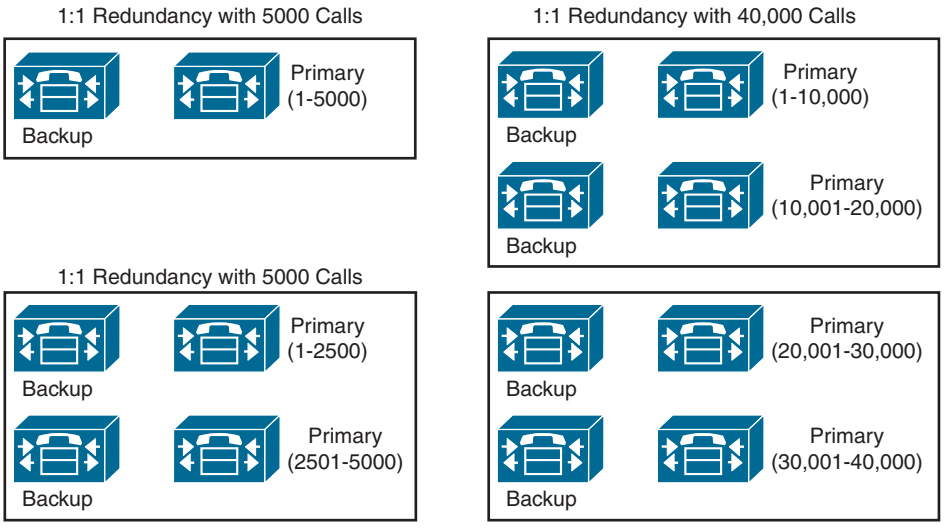

Figure 6-3 *Call Processing Subscriber Groups in a 1:1 Deployment*

When you are choosing the server(s) that will host the call agent VMs, it is best practice to choose a platform that supports dual power supplies. You should plug each power supply into different power sources so that in the event one of the power circuits fails, there is still constant power being supplied to the hosting server. Extra measures can be taken when combining dual power supplies with an UPS source. Should a power outage occur at the facility where the server is located, constant power will continue to be provided to the server, whereby phone services will also continue.

Several measures can be taken to provide high availability in the network connectivity. The speed and duplex used for network connectivity are essential to ensuring high availability. Many devices will be communicating with the call agent and will require a lot of bandwidth. Cisco recommends using a 1 Gbps or 10 Gbps throughput rate on the NIC the server is connected to. Voice and video communications should always use full duplex. If 1 Gbps or 10 Gbps throughput is used, full duplex is automatic. If a lower throughput is used, you should check the duplex configuration to ensure that it is set to full duplex, and not automatic or half duplex. Auto duplex on a Cisco switch will default to half duplex. Network redundancy can be achieved by using two Ethernet connections at the server. Each connection should be connected to a different switch so that in the event one of the switches fails, network connectivity will be maintained. This same redundancy can be implemented between switches within the network by physically distributing the network connections between different physical network switches within the same location. On the server, within the hypervisor there is a virtual switch with multiple uplinks. Therefore, a single virtual NIC defined in the call agent OVA settings is sufficient. In the VMware vSphere virtual switch, you can configure NIC teaming for the switch uplink.

Disaster Recovery

The Cisco Webex Global Site Backup architecture handles power outages, natural disaster outages, service capacity overload, network capacity overload, and other types of service interruptions. Global Site Backup supports both manual and automatic failover. The manual

failover mode is typically used during maintenance windows. The automatic failover mode is used in case of real-time failover due to a service interruption.

Global Site Backup is automatic and transparent to the end users, it is available for all users, and it imposes no limits on the number of users who can fail over. Global Site Backup consists of the following main components:

- **Global Site Service:** Is responsible for monitoring and switching traffic at the network level
- **Database Replication:** Ensures that the data transactions occurring on the primary site are transferred to the backup site
- **File Replication:** Ensures that any file changes are maintained in synchronization between the primary and the backup site

For disaster recovery, you can configure a cold-standby system in a second data center. If the primary system is configured for high availability, you can optionally choose to configure high availability for the disaster recovery system. Cisco Prime Collaboration Deployment does a direct migration, whereas previous migration methods involved more steps with a "server recovery" Disaster Recovery System relying on an initial upgrade followed by a restore from backup.

Dial Plan

The dial plan is one of the key elements of a Unified Communications and Collaboration system, and it is an integral part of all call processing agents. Generally, the dial plan is responsible for instructing the call processing agent on how to route calls. The dial plan performs many functions.

Endpoint addressing is one of the main functions of a dial plan. For destinations registered with the call processing agent, addresses are assigned to provide reachability. These internal destinations include all endpoints, such as IP phones, video endpoints, soft clients, and analog endpoints, as well as applications, such as voicemail systems, auto attendants, and conferencing systems. Path selection is another function of a dial plan. Depending on the calling device and the destination dialed, a path to the dialed destination is selected. If a secondary path is available, this path will also be considered if the primary path fails. Calling privileges is a third function of the dial plan. Different groups of devices can be assigned to different classes of service, by granting or denying access to certain destinations. For example, lobby phones might be allowed to reach only internal and local PSTN destinations, whereas executive phones could have unrestricted PSTN access. Dial plans can affect the manipulation of dialed destinations. On the path from the dialing device to the dialed destination, the dial plan can apply manipulations to the dialed destination. For example, users in the United States might dial 9011496901234 to reach a destination in the PSTN in Germany, while a user in France might be able to reach the same destination by dialing 000496901234. This dialed destination would need to be presented as 011496901234 to a PSTN trunk on a gateway in the U.S. and as 00496901234 to a PSTN trunk on a gateway in France. The dial plan can also affect calling numbers; for example, calls across the WAN may display a caller ID as 4111, but when the call is routed across the PSTN, it may appear as 3055554111. Presentation of information about identities involved in the call is also part of a dial plan. During session establishment and also while in the call, on both the calling and the called device,

information about the other device is displayed. Depending on call state and direction, this includes calling, diverting, alerting, and connected party information. The dial plan can define mappings that influence the format and content of information displayed.

Dial plan and number normalization considerations must be taken into account when deploying software-based endpoints. Jabber desktop clients typically use the directory for searching, resolving, and adding contacts. The number that is associated with those contacts must be in a form that the client can recognize, resolve, and dial.

Deployments may vary, depending on the configuration of the directory and Cisco Unified Communications Manager. In cases where the directory contains E.164 numbering, such as +18005551212, for business, mobile, and home telephone numbers, and Cisco Unified Communications Manager also contains an E.164 dial plan, the need for additional dial rules is minimized because every lookup, resolution, and dialed event results in an E.164-formatted dial string. If a Cisco Unified Communications Manager deployment has implemented a private dial plan, such as 5551212, then translation of the E.164 number to a private directory number needs to occur on Cisco Unified Communications Manager and possibly on the IOS gateways as well. Outbound calls can be translated by Cisco Unified Communications Manager translation patterns that allow the number being dialed, such as +18005551212, to be presented to the endpoint as the private number 5551212. Inbound calls can be translated by means of directory lookup rules. This allows an incoming number of 5551212 to be presented for reverse number lookup caller identification as 18005551212.

Private numbering plan deployments may arise, where the dial plan used for the company and the telephone number information stored in the LDAP directory may require the configuration of translation patterns and directory lookup rules in Cisco Unified Communications Manager to manage number format differences. Directory lookup rules define how to reformat the inbound call ID to be used as a directory lookup key. Translation patterns define how to transform a phone number retrieved from the LDAP directory for outbound dialing.

Cisco Unified Communications Manager uses translation patterns to manipulate both the calling and called numbers before a call is routed, and they are handled strictly by Cisco Unified Communications Manager. Application dialing rules can be used as an alternative to translation patterns to manipulate numbers that are dialed. Application dialing rules can automatically strip numbers from or add numbers to phone numbers that the user dials. Application dial rules are configured in Cisco Unified Communications Manager and are downloaded to the client from Cisco Unified Communications Manager. Translation patterns are the recommended method for manipulating dialed numbers.

Directory lookup rules transform caller identification numbers into numbers that can be looked up in the directory. A directory lookup rule specifies which numbers to transform based on the initial digits and the length of the number. Directory lookup rules are configured in Cisco Unified Communications Manager and are downloaded to the client from Cisco Unified Communications Manager. Before a call is placed through contact information, the client application removes everything from the phone number to be dialed, except for letters and digits. The application transforms the letters to digits and applies the dialing rules. The letter-to-digit mapping is locale-specific and corresponds to the letters found on a standard telephone keypad for that locale. Users cannot view or modify the client transformed numbers before the application places the call.

Security

Key Topic

Firewalls and ACLs are security capabilities that exist on the router to help secure your network by providing a first line of defense from attacks outside your network trying to access data inside. Unfortunately, that is not always enough to protect information inside the network from malicious attacks. If a user were to log in to his bank account across the public Internet, what is to stop a hacker from obtaining that user's login credentials and emptying the bank account? If that communication were sent over a nonsecure connection, the login information is in plain text. All the hacker would need is a packet sniffer to capture and view that information, and then would have access. To hide important information, the data needs to be encrypted. Think "Da Vinci Code" with text ciphers, only for a digital world. As long as your computer and your bank's server are the only two devices with the text ciphers, no other device will be able to read your information. Two security mechanisms can provide this level of encryption. Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols that provide communications security over a computer network for TCP and UDP traffic. Although SSL is rarely used anymore, the TLS protocol aims primarily to provide privacy and data integrity between two communicating hosts or applications.

Client/server applications such as web browsers, email, and VoIP commonly use the TLS protocol to prevent eavesdropping and tampering of information. The easiest way to segregate the information is to use different port numbers for unencrypted traffic and encrypted traffic, such as port 80 for HTTP or port 443 for HTTPS. The connection is secure because symmetric cryptography is used to encrypt the transmitted data. The keys for this symmetric encryption are generated uniquely for each connection and are based on a shared secret negotiation at the start of the session. The server and client negotiate the details of which encryption algorithm and cryptographic keys to use before the first byte of data is transmitted. Identification is usually in the form of digital "certificates" that contain the server name, the trusted certificate authority (CA), and the server's public encryption key. The identity of the communicating parties can be authenticated using this public-key cryptography (asymmetric cryptography) to ensure only the intended recipient can decrypt the traffic. The negotiation of a shared secret is both secure and reliable against eavesdroppers and attacks, including man-in-the-middle attacks. The connection ensures integrity because each message transmitted includes a message integrity check using a message authentication code to prevent undetected loss or alteration of the data during transmission.

After the client and server have agreed to use TLS, they negotiate a stateful connection by using a handshake procedure. Figure 6-4 shows a general overview of how a TLS handshake takes place.

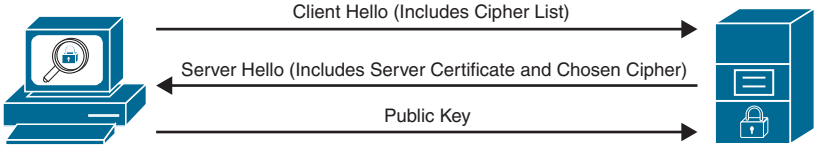

Figure 6-4 *TLS Handshake Overview*

The handshake begins when a client connects to a TLS-enabled server requesting a secure connection and presents a list of supported ciphers and hash functions. From this list, the server picks a cipher and hash function that it also supports and informs the client of the decision. The server then identifies itself with its digital certificate, which can contain the server name, the trusted certificate authority, and the server's public encryption key. The client then

validates the certificate before proceeding. Public-key encryption is used to share the pre-master secret via the use of RSA or Diffie-Hellman key exchange. This process generates a random and unique session key for encryption and decryption that has the additional property of forward secrecy, which protects past sessions against future compromises of secret keys or passwords.

Remember that the server is validated because the client initiates the secure connection. The client side confirms that the server is who it claims to be and whether it can be trusted with the use of certificates. Figure 6-5 illustrates the elements contained within a certificate that can be used to verify the certificate holder is authentic.

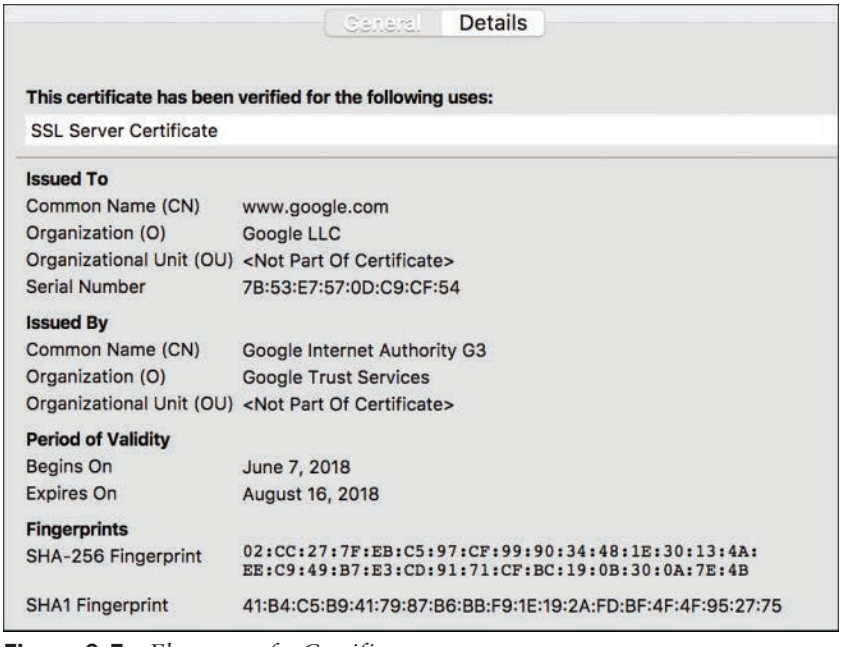

Figure 6-5 *Elements of a Certificate*

Key Topic

The client receives the digital certificate from the server side of the TLS negotiation, but the identity must be verified before proceeding. As seen in Figure 6-5, when Google's server sends its certificate, it contains the name of the certificate holder. This name is checked against the Common Name (CN) or the Subject Alternative Name (SAN), which is www.google.com in this instance. Also, it contains additional information like a serial number, expiration dates or Period of Validity, revocation status (not applicable in this figure), a copy of the certificate holder's public key (SHA-256 Fingerprint used for encrypting messages and digital signatures), and the digital signature of the certificate-issuing authority (SHA1 Fingerprint). If you trust this certificate authority, you can verify (using the CA's public key) that it really did sign the server's certificate. To sign a certificate yourself, you need the private key, which is known only to the CA of your choice. This way, an attacker cannot falsely sign a certificate and claim to be Google.com. When the certificate has been modified, the signature will be incorrect, and the client will reject it.

Although this form of TLS encryption is very secure, you can still take additional measures to ensure an even higher level of security. This is known as Mutual TLS, which is

synonymous with TLS Verify. In Mutual TLS authentication, both parties authenticate each other through verifying the provided digital certificate so that both parties are assured of the others' identity. Mutual TLS is similar to the normal process of the client handling the verification of the server's certification but includes the additional step of the client providing a certificate to the server for verification. This process allows the server side to authenticate the client, allowing both parties to trust each other. Figure 6-6 illustrates how a Mutual TLS negotiation would take place.

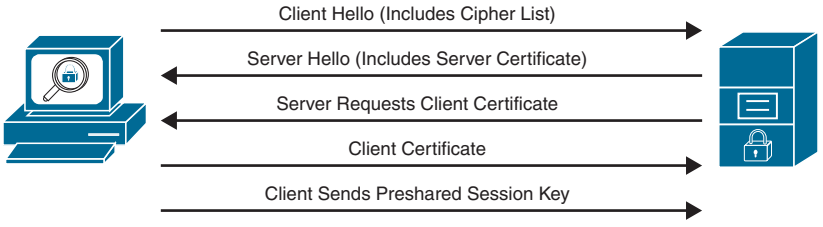

Figure 6-6 *Mutual TLS Negotiation*

Server-to-server connections rely on Mutual TLS for mutual authentication. In the Cisco Collaboration infrastructure, some examples would be a secure connection between endpoints and the Cisco Unified Communications Manager, referred to as TLS Verify, Cisco Unified Communications Manager intercluster trunks to other clusters, and even Cisco Unified Communications Manager SIP trunks to a Cisco Expressway or a Video Communication Server (VCS).

QoS

Key Topic

QoS is a marking system for network traffic that allows packets to be prioritized during high congestion times, so that drop-sensitive packets can be sent first and drop-insensitive packets are sent last. For example, an email is sent using TCP, which will resend the packets in the event they are not received at the destination. Voice and video packets are sent over UDP and will not be resent if they are dropped, which could cause media issues at the receiving end of the call. Therefore, voice and video traffic should be provided with a higher priority than email traffic. When it comes to QoS, it is best practice to mark packets as close to the source as possible. Most devices, such as computers and servers, cannot mark their own packets and should not be trusted even if they can. Cisco phones, however, can mark their own packets and can be trusted with the QoS markings they provide. Therefore, QoS trust boundaries should be set up so that the switch will trust the QoS markings that phones place on their own packets. Layer 2 QoS uses a mechanism called class of service (CoS), which operates on the 802.1q VLAN. Unlike Layer 3 QoS mechanisms, CoS does not ensure network performance or guarantee priority in packets being delivered. Therefore, after packets are marked with CoS, they will need to be converted to DSCP using the cos-to-dscp map, which is built into all Cisco switches. By default, QoS on a Cisco access switch is disabled. Once QoS is enabled, the switch does not trust QoS settings from a phone. Two simple commands can be entered under the global menu on a switch to enable QoS and change the trust boundary. Once QoS is enabled, you can use a **show** command to verify these settings. Example 6-1 illustrates the QoS Enable and Trust Boundary Commands and the **show** verification command.

Example 6-1 *QoS Enable and Trust Boundary Commands*

```
Switch(config)# mls qos
Switch(config)# interface fastethernet 0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# end
Switch# show mls qos interface fastethernet 0/1
FastEthernet0/1
trust state: trust cos
trust mode: trust cos
cos override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
```

Obviously, this is the simplest design, and there are many other concerns to consider, along with many other settings that can be configured. This example is intended to provide a basic understanding of QoS at the Layer 2 level. For more information on QoS, refer to the *Enterprise QoS Solution Reference Network Design Guide*.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 6-5 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 6-5 Key Topics for Chapter 6

Key Topic Element	Description	Page Number
Paragraph	Framework for Cisco Jabber	117
Paragraph	Deskphone Mode and Softphone Mode for Jabber	118
Paragraph	Webex Teams	118
Table 6-2	Cisco Telepresence Endpoint Product Portfolio	119
List	Databases Supported for Directory Synchronization	120
Paragraph	Signaling Protocols Supported Through the CUCM	120
Paragraph	NTP Used by CUCM	121
Paragraph	DHCP from CUCM	121

Key Topic Element	Description	Page Number
Paragraph	Licensing Differences Between the VCS and Expressway Products	123
Table 6-3	Comparison of the Cisco Expressway and the Cisco VCS	123
Paragraph	CUC	125
Paragraph	CUE	125
Paragraph	IM and Presence service	126
Paragraph	CMS	126
Paragraph	Additional Features Supported on CMS	127
Paragraph	TMS Features	127
List	Three Models of Prime Collaboration	128
Paragraph	Prime Collaboration Provisioning	128
Paragraph	Prime Collaboration Assurance	128
List	Prime Collaboration Analytics	129
Table 6-4	CUWL and CUCL Licensing Model	130
Paragraph	PMP and SMP Licenses	131
List	Purchasing Models for Flex Licensing	132
List	Benefits of Smart Accounts	133
Paragraph	Business Edition Series UCS Servers	134
List	Regions and Locations	135
Paragraph	Expressway Cluster Requirements	136
Paragraph	CUCM Cluster Behavior	137
Paragraph	2:1 Redundancy Group in CUCM	137
Paragraph	1:1 Redundancy Group in CUCM	138
Paragraph	Functions of a Dial Plan	140
Paragraph	TLS and SSL Comparison	142
Paragraph	Certificate Checking Process	143
Paragraph	QoS Best Practice for L2 Marking	144

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

CAST, CE, CMA, CME, CMS, CTI, CUBE, CUC, CUCL, CUCM, CUCSF, CUE, CUWL, DNS, DX, Flex, FXO, FXS, HCS, IM, IMP, IX, Locations, MRA, MTLS, MX, NTP, PMP, PoE, Regions, RMS, SMP, SSL, SX, TC, TIP, TLS, UC, URI, URL, VCS, VoIP, VPN, Webex Endpoints, WebRTC, XMPP

Command Reference to Check Your Memory

This section includes the most important configuration and EXEC commands covered in this chapter. It might not be necessary to memorize the complete syntax of every command, but you should be able to remember the basic keywords that are needed.

To test your memory of the commands, cover the right side of Tables 6-6 with a piece of paper, read the description on the left side, and then see how much of the command you can remember.

The CLCOR (350-801) exam focuses on practical, hands-on skills that are used by a networking professional. Therefore, you should be able to identify the commands needed to configure and test QoS settings on a switch.

Table 6-6 Cisco Meeting Server MMP Commands

Task	Command Syntax
Enables the multilayer switching quality of service	Switch(config)# **mls qos**
Enters the configuration field if fast Ethernet switch port 1	Switch(config)# **interface fastethernet 0/1**
Configures the switch to trust all ingress traffic	Switch(config-if)# **mls qos trust cos**
Takes the admin out of global configuration mode in the switch	Switch(config-if)# **end**
Displays the QoS settings configured on the switch from above	Switch# **show mls qos interface fastethernet 0/1**
Displays output from the previous **show** command	FastEthernet0/1 Trust state: **trust cos** Trust mode: **trust cos** CoS override: **dis** Default COS: **0** DSCP Mutation Map: Default DSCP Mutation Map Trust device: **none**

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the four main categories of Webex endpoints in the Telepresence product portfolio.
2. List the licensing differences between an Expressway and a VCS.
3. List the three models of Prime Collaboration.
4. List the switch commands to enable QoS at Layer 2.

Part II

Endpoints

This part covers the following topics:

- **Chapter 7, Cisco Unified Communications Phones:** This chapter will introduce the current Cisco UC phones available and discuss the features and capabilities of these phones. Specifically, the focus of this chapter will center around the 7800 series and 8800 series phones. As companies begin migrating to the cloud, the software running on the phones changes as well. This chapter will explain the differences between Enterprise software and Multiplatform Phone (MPP) software.
- **Chapter 8, Cisco Telepresence Endpoints:** This chapter will introduce the current Cisco Telepresence endpoint portfolio. All of the endpoints discussed in this chapter support the CE software operating system. Cisco has introduced many new endpoints and omitted a few from its product line. Therefore, the focus of this chapter will be on the current Cisco Telepresence endpoints at the time this book was written.
- **Chapter 9, Endpoint Registration:** This chapter will delve into the registration and call setup settings that exist on Cisco Telepresence endpoints. UC endpoints are controlled entirely from the Cisco Unified Communications manager, so they will not be discussed in this capacity. However, Cisco Telepresence endpoints have so much intelligence built into them that many features are available to the users from the interface of the endpoint itself. Unlike UC endpoints, Telepresence endpoints can register to the Cisco Unified Communications Manager, or the Cisco Expressway, or even a third-party call control system. This chapter will help you understand the differences in registering to these different call control systems.
- **Chapter 10, Call Settings on Cisco CE Software-Based Endpoints:** This chapter will describe how to access and configure various call settings on Cisco CE software-based endpoints, such as calling options, content sharing options, and several other options.
- **Chapter 11, Maintaining Cisco Endpoints:** This chapter will explain how to perform various maintenance tasks from the Cisco Telepresence endpoints. These maintenance tasks include upgrading the endpoint software, performing a backup and restore of the endpoint configurations, and accessing logs on Telepresence endpoints. Although the main focus of this chapter is on Telepresence endpoints, a short discussion of how to access logs on UC endpoints will also be included in this chapter.

CHAPTER 7

Cisco Unified Communications Phones

This chapter covers the following topics:

7800 Series Phones: This topic will introduce the Cisco 7800 series VoIP phones along with features and capabilities.

8800 Series Phones: This topic will introduce the Cisco 8800 series VoIP and video phones along with features and capabilities.

Software Versions for Phones: This topic will explain the differences between the Enterprise software for phones and the Multiplatform Phone (MPP) software available on Cisco UC phones.

Cisco Unified Communications (UC) phones are closer to what would normally be referred to as business phones. These are voice over IP (VoIP) phones that reflect the typical image of a phone with a handset and numeric dial pad. They operate entirely on an IP network and depend on an intelligent IP PBX in order to function. In the Cisco world of collaboration, that PBX is the Cisco Unified Communications Manager (or CUCM). Cisco offers two main series of phones to customers: the 7800 series and the 8800 series. These phone models can operate using two distinct operating systems. This chapter will introduce you to the current Cisco Unified Communications phones and the software versions they support. Topics discussed in this chapter include the following:

- 7800 Series Phones
- 8800 Series Phones
- Software Versions for Phones

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 7-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 7-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
7800 Series Phones	1–2
8800 Series Phones	3–6
Software Versions for Phones	7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What size is the screen on the Cisco IP Phone 7811?
   a. 3.28″
   b. 3.5″
   c. 4.78″
   d. 5.0″
2. How many lines does the Cisco IP Phone 7861 support?
   a. 8
   b. 10
   c. 12
   d. 16
3. Which of the following is a difference between the Cisco IP Phone 8811 and 8841?
   a. 8841 supports Bluetooth; 8811 does not.
   b. 8841 supports a color display; 8811 does not.
   c. 8841 supports Intelligent Proximity; 8811 does not.
   d. 8841 comes with one USB port; 8811 comes with none.
4. Which of the following comes with an 8851 phone but not with the 8851NR?
   a. USB Port
   b. Bluetooth
   c. Touchscreen
   d. Wi-Fi
5. Which of the following phones comes with video support and Wi-Fi?
   a. 8845
   b. 8865
   c. 8865NR
   d. Both B and C
6. What video resolution is supported on the Cisco IP Phone 8800 series, which supports video communication?
   a. 360p30
   b. 480p30
   c. 720p30
   d. 1080p30

7. Which Cisco IP video phone model supports the MPP firmware load?
   a. 7861
   b. 8845
   c. 8861
   d. All of these answers are correct.

Foundation Topics

7800 Series Phones

Cisco IP 7800 series VoIP phones are a lower-cost option for customers who are setting up a new UC solution and are not concerned about the features their phones offer to end users. The Cisco IP 7800 series phones are all voice-only and support a Class 1 Power over Ethernet (PoE), as well as other basic phone features, such as hold, call forward, and call transfer. The Cisco IP 7800 series phones can be used in an on-premises deployment and register to the Cisco Unified Communications Manager, or they can register to the Webex Control Hub in the cloud. These phones can also register to the Cisco Unified Communications Manager from a remote location without a VPN by using the Mobile and Remote Access (MRA) option through the Expressway Series servers. Table 7-2 identifies some of the features supported on the Cisco IP 7800 series phones.

Table 7-2 Cisco IP 7800 Series Phone Models and Features

Feature	7811	7821	7832	7841	7861
Screen	3.28"	3.5"	3.4"	3.5"	3.5"
Ethernet Switch	10/100	10/100	10/100	10/100/1000	10/100/1000
Line Keys	1	2	1	4	16
Backlit	No	Yes	Yes	Yes	Yes
Wideband Audio	Optional	Yes	Yes	Yes	Yes
Field-Replaceable bezel	No	Yes	No	Yes	Yes
PoE	Class 1	Class 1	Class 2	Class 1	Class 1
Cloud Ready	Yes	Yes	Yes	Yes	Yes
Power Save Plus	No	Yes	No	Yes	Yes

The line keys on each model are fully programmable. You can set up keys to support either lines, such as directory numbers, or call features, such as speed dialing. You can also boost productivity by handling multiple calls for each directory number, using the multi-call-per-line appearance feature. Tri-color LEDs on the line keys support this feature and make the phone simpler and easy to use. The Cisco IP Phone 7811 and the 7832 speakerphone support one line, and these phones are available only in a charcoal gray color. All other phone models in this series are available in either charcoal gray or white. The Cisco IP Phone 7821 supports two lines, the Cisco IP Phone 7841 supports four lines, and the Cisco IP Phone 7861 supports

16 lines. None of the Cisco IP 7800 series phones support the Key Expansion Module (KEM). Fixed function keys on all models give you one-touch access to services, messaging, directory, hold and resume, transfer, and conference features. The full-duplex Cisco 7832 speakerphone lets you set up clear multiparty conferences for flexible, productive collaboration.

The Cisco IP 7800 series phones set a new standard in usability and deliver a context-sensitive user experience. This series features a high-resolution 3.5" (396 × 162) grayscale display with white backlighting on IP Phones 7821, 7841, and 7861, and a 3.2" (384 × 106) display without backlighting on IP Phone 7811, for easy reading. The 7832 speakerphone features a 3.4" backlit, monochrome, pixel-based display with an antiglare bezel to make viewing and interaction easier. Localized language support, including right-to-left onscreen text, meets the needs of global users. In fact, 38 different languages are supported on the Cisco IP 7800 series phones. Refer to the "Cisco IP Phone 7800 Series Data Sheet" at Cisco.com for a specific language support listing.

The Cisco IP Phone 7800 series is also more energy efficient and eco-friendly, to support your green initiatives. Each phone supports PoE Class 1, Cisco's EnergyWise, and is Energy Star certified. A standard power-save option is available on Cisco IP Phones 7821, 7841, and 7861 to reduce power consumption during off hours, save money, and maximize energy efficiency.

The Cisco IP 7800 series phones portfolio is ideal for any mid-sized to large enterprise company that wants to update its phone system from a traditional analog or digital-based system to an IP communications system. It's also an excellent choice if you're seeking to expand your voice communications support with your current Cisco Unified Communications solution. Small businesses that have interest in the Cisco IP 7800 series phones but have investment in or are considering third-party hosted call control services are also candidates for the Cisco IP 7800 series phones. Figure 7-1 illustrates the different key options available on the Cisco IP Phone 7841 model.

Figure 7-1 *Key Options on the Cisco IP Phone 7841 Model*

Key Topic

The feature and line buttons may light up in three different colors—green, amber, or red—and may be steady or flashing. A steady green LED indicates an active call or two-way intercom call. A flashing green LED indicates a call on hold. A steady amber LED indicates privacy is in use, a one-way intercom call is in use, or that line is logged into a hunt group. A flashing amber LED indicates an incoming call or reverting call. A steady red LED indicates a remote line is in use, such as a shared line, or Do Not Disturb (DND) has been activated. A flashing red LED indicates that a remote line has been placed on hold. The phone screen shows information about your phone, such as directory number, active call and line status, softkeys, speed dials, placed calls, and phone menu listings. The screen is made up of three sections: the header row, the middle section, and the footer row. Figure 7-2 illustrates the three sections of the Cisco IP 7800 Series Phone screen.

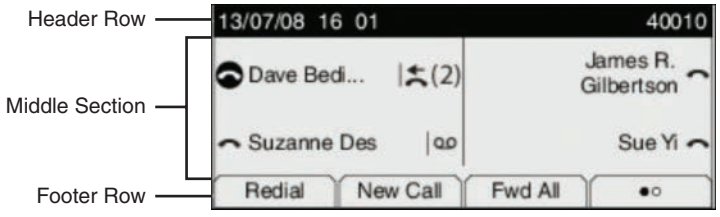

Figure 7-2 *Three Sections of the Cisco IP 7800 Series Phone Screen*

At the top of the screen is the header row. The header row displays the phone number, current date and time, and a number of icons. The icons displayed indicate when certain features are active. The middle of the phone screen displays the information associated with the line and feature buttons on the phone. The bottom row of the screen contains the softkey labels. Each label indicates the action for the softkey button directly below the indicated place on the screen. An administrator can reduce the amount of power a phone screen uses when the phone is not being used. Two energy-saving levels can be set up on a Cisco IP 7800 series phone:

Key Topic

- **Power Save:** The backlight or screen turns off when the phone is inactive for a set interval. The backlight can be managed.
- **Power Save Plus:** The phone screen turns on and off at times that are based on the employee's work schedule. If that employee's work hours or work days change, an administrator can reconfigure the phone.

For example, a Cisco IP 7800 phone can be set to alert the user 10 minutes before it turns off. The user will see the Select button light up and receive a message on the screen that the phone is turning off soon. Notifications can be set up at the following intervals:

- Four rings at 10 minutes before power off
- Four rings at 7 minutes before power off
- Four rings at 4 minutes before power off
- Fifteen rings at 30 seconds before power off

When a phone is active, it waits until it has been inactive for a set interval before it notifies the user of the pending power shutdown. Note that the Cisco IP Phone 7811 doesn't

support Power Save or Power Save Plus. When a phone does turns off to save energy, the phone screen goes blank, and the Navigation button lights up. A user simply needs to press this button to turn the phone back on.

8800 Series Phones

A step up from the 7800 series phones are the 8800 series phones. These phones are the latest and greatest in the Cisco UC Phone product portfolio and offer a more feature-rich experience for end users. The 8800 series phones can be used in an on-premises deployment and register to the Cisco Unified Communications Manager, or they can register to the Webex cloud. These phones can also use MRA from remote locations to register back to the Cisco Unified Communications Manager at a central office location so that a VPN does not need to be utilized. Table 7-3 identifies some of the features supported on the 8800 series phones.

Key Topic

Table 7-3 8800 Series Phone Models and Features

Feature	8811	8841	8851	8851NR	8861	8845	8865	8865NR
Screen	Grayscale	Color	Color	Color	Color	Color	Color	Color
HD Video 720p	No	No	No	No	No	Yes	Yes	Yes
Bluetooth	No	No	Yes	No	Yes	Yes	Yes	No
Cisco Intelligent Proximity (MV)	No	No	Yes	No	Yes	Yes	Yes	No
USB Ports	0	0	1	1	2	0	2	2
KEM	0	0	2	2	3	0	3	3
Wi-Fi	No	No	No	No	Yes	No	Yes	No

Key Topic

With the Cisco IP Phone 8811, you can increase personal productivity through an engaging user experience that is both powerful and easy to use. The Cisco IP Phone 8811 combines an attractive new ergonomic design with wideband audio for crystal clear voice communications, "always-on" reliability, encrypted voice communications to enhance security, and access to a comprehensive suite of unified communications features from Cisco on-premises and hosted infrastructure platforms and third-party hosted call control. The Cisco IP Phone 8811 supports five programmable line keys. You can configure keys to support either multiple directory numbers or calling features such as speed dial. You can also boost productivity by handling multiple calls for each directory number, using the multi-call-per-line feature. Fixed-function keys give you one-touch access to applications, messaging, directory, and often-used calling features such as hold/resume, transfer, and conference. Backlit acoustic keys provide flexibility for audio path selection and switching. The Cisco IP Phone 8811 offers a 5" high-resolution (800 × 480) widescreen backlit grayscale display. Localized language support, including right-to-left onscreen text, meets the needs of global users. This phone supports a built-in Gigabit Ethernet switch for your PC connection. Support for Cisco EnergyWise technology makes the Cisco IP Phone 8811 more energy efficient and eco-friendly; the phone is qualified by the Energy Star organization. The Cisco IP Phone 8841 supports all the same features and capabilities as the Cisco IP Phone 8811, except that this phone offers a 5" high-resolution (800 × 480) widescreen VGA backlit color display.

Key Topic

The Cisco IP Phone 8851 supports all the same features as the Cisco IP Phones 8811 and 8841, plus some additional features and capabilities. Cisco Intelligent Proximity for Mobile Voice (MV) brings the worlds of desk and mobile together for you when you are using your mobile device at the desk for your work. You can move the audio path over to the Cisco IP Phone 8851 during active mobile calls to take advantage of its superior audio acoustics. An example would be to share a conversation with a colleague you want to listen in on the call. This capability gives you greater flexibility and a superior user experience when at your desk. The Cisco IP Phone 8851 also comes standard with one USB port, so you can connect a headset or charge your smartphone while at your desk. The Cisco IP Phone 8851 offers a 5" high-resolution (800 × 480) widescreen VGA backlit color display. Up to two optional IP Phone 8800 Key Expansion Modules with up to 56 additional line and feature keys are supported. The Cisco IP Phone 8851NR is a No Radio variant of the 8851 model that can be used in secure environments such as government and military buildings. There are some feature differences between the 8851 and the 8851NR; for example, the Cisco IP Phone 8851NR does not support Bluetooth or Intelligent Proximity. All other features and settings are the same, however. The Cisco IP Phone 8861 adds four extra capabilities beyond those of the Cisco IP Phone 8851. First, the 8861 phone offers two USB ports, one on the side just like the 8851 offers, plus an additional port on the back of the phone. Second, this phone supports a wireless network connection with 802.11a/b/g/n/ac WLAN enabled. Third, this phone offers up to three optional IP Phone Key Expansion Modules supporting up to 108 additional line and feature keys. Finally, the Cisco IP Phone 8861 offers a 5" high-resolution (800 × 480) widescreen VGA backlit color *touchscreen* display.

Key Topic

The three IP phones Cisco offers with video capability are the Cisco IP Phones 8845, 8865, and the 8865NR. The Cisco IP Phone 8845 can help users increase personal productivity through powerful and easy user experiences. It combines an attractive ergonomic design with 720p30 HD video capabilities in addition to the wideband audio for crystal-clear video and voice communications. The 8845 encrypts video and voice communications for security and offers access to a comprehensive suite of unified communications features. Offering capabilities above the Cisco IP Phone 8841 beyond just video, the 8845 also supports Cisco Intelligent Proximity MV and Bluetooth. The Cisco IP Phone 8865 offers all the same great features as the Cisco IP Phone 8861, plus the HD 720p30 video capabilities. The Cisco IP Phone 8865NR offers the same features as the 8865, except the Cisco IP Phone 8865NR does not support Bluetooth, Wi-Fi, or Intelligent Proximity. All phones in the Cisco IP Phone 8800 series are available in two color options: charcoal gray and white.

The phone buttons on the Cisco IP Phone 8800 series are similar to the Cisco IP Phone 7800 series, with some minor differences. All Cisco IP Phones in the 8800 series have identical buttons to one another. Based on the preceding descriptions, there are obviously two hardware types available in the Cisco IP Phone 8800 series. The Cisco IP Phones 8811, 8841, 8851, 8851NR, and the 8861 do not have a camera. The Cisco IP Phones 8845, 8865, and 8865NR do have cameras. The cameras can be manually tilted but have no pan or zoom capabilities. A shutter on the end of the camera can be closed for those people who fear being watched even when a call is not in session. (You know who you are.) Figure 7-3 illustrates the different key options available on the Cisco IP Phone 8865.

Figure 7-3 *Key Options on the Cisco IP Phone 8865*

The Cisco IP Phone 8800 series and the Cisco IP Phone 7800 series have many similar characteristics. All phones in the 8800 series support Power Save and Power Save Plus. The LED colored lighting on the 8800 series phones is the same as it is on the 7800 series phones. The phone screen on the Cisco IP Phone 8800 series shows the same basic information about your phone as the Cisco IP Phone 7800 series, such as directory number, active call and line status, softkeys, speed dials, placed calls, and phone menu listings. The screen is made up of the same three sections: the header row, the middle section, and the footer row. The difference between these two phones is the size of the screen and the feel of the phone. Also, the screen on the 8845, 8865, and 8865NR phones will display the video communication from the far-end endpoint during a call. Figure 7-4 illustrates the three sections of the Cisco IP 8800 series phone screen.

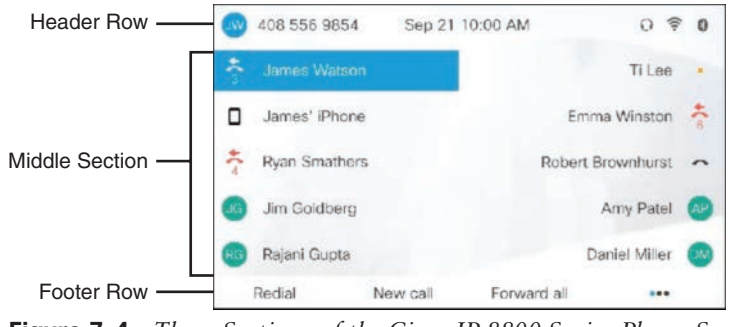

Figure 7-4 *Three Sections of the Cisco IP 8800 Series Phone Screen*

Three other phones in the 8800 series are worth mentioning. The Cisco Unified IP speakerphones 8831 and 8832 are audio-only phones with a single programable line key; they use the DECT microphones. These devices are great for smaller huddle rooms where small teams can place audio calls to clients. The main difference between these two speakerphones is that the 8831 has a dial pad in a separate device that is attached to the speakers

with a cable, whereas the 8832 is a redesigned speakerphone with a new look that includes an all-in-one unit between the speaker and the dial pad. Additionally, the 8832 has a physical USB port.

The last model worth mentioning is the Cisco Wireless IP Phone 8821. This phone is a ruggedized, resilient, and secure 802.11 wireless LAN handset that delivers cost-effective, on-premises, comprehensive voice over Wireless LAN (VoWLAN) communications for the highly mobile in-campus worker. While the 8821 is sleek and lightweight, the design is hardened for users. It is Ingress Protection standard (IP54) rated and is sealed for protection against dust, splashes, and water. The device is also MIL-STD-810G tested, with a dozen drops onto concrete from heights of up to 5 feet (1.5 m), to help ensure shock resistance and avoid breakage if dropped. The 8821 enhances security and simplifies configuration management. Stronger encryption is supported for certificate management and policy enablement with the support of Secure Hash Algorithm 2 (SHA-2). Simple Certificate Enrollment Protocol (SCEP) eases IT administration by enabling automatic certificate management on the device.

Software Versions for Phones

Key Topic

Currently, two different types of firmware can be used on the Cisco IP phones in the 7800 and 8800 series. Phones that register to the Webex Control Hub must be running the Multiplatform Phone, or MPP, firmware. Phones that register to the Cisco Unified Communications Manager must be running the Enterprise firmware. Cisco IP phones can be ordered with the desired firmware already installed, which is ideal for greenfield deployments. Alternatively, the software can be migrated to the needed platform on existing phones used in an enterprise, which is ideal for brownfield deployments.

For partners who provide voice and video services to end customers registering to third-party call control platforms, Cisco offers MPP firmware loads that support these platforms. The platforms that support the MPP firmware include Asterisk, Webex Calling (formerly Broadcloud), Broadworks, Centile, and Metaswitch. The feature set provided by this firmware is not identical to that of the Enterprise firmware designed and built for use with Cisco Unified Communications Manager, but there are many similarities. The features and information about the Cisco IP phones shared up to this point in this chapter have been based on the Enterprise firmware. For more information on specific feature support using Enterprise firmware, refer to the data sheet on the Cisco phone model you wish to inquire about. Table 7-4 identifies some of the features supported on phones running the MPP firmware.

Key Topic

Table 7-4 MPP Firmware Feature Support

Security	Applications	Call Control and Audio Features	Directory	Management
802.1x authentication	Cisco XML Services Interface (XSI)	Busy Lamp Field (BLF)	Local phonebook	Configuration: Browser Phone Auto Provision

Security	Applications	Call Control and Audio Features	Directory	Management
Media encryption via SRTP	UC-One Presences	Call forwarding	XML/LDAP remote directory	Auto Provision via TFTP/HTTP/HTTPs for mass deployment
Transport Layer Security (TLS)		Call hold	Intelligent search	Encrypted HTTP data in plain HTTP transmissions
Encrypted configuration files		Call pickup	Call history	Packet Capture, Problem Reporting Tool (PRT), and upload of PRT
Digest authentication		Call park	Reverse address lookup in all directories	Remote generation and upload of PRT data
Password login		Call transfer	Intelligent Proximity MV	Configuration report to provisioning server
HTTPs secure provisioning		Call waiting		
Mandatory/ Optional Secure Call		Do Not Disturb (DND) Extension Mobility/Hot Desking Intercom Music on Hold		

Many more call control and audio features are available with the MPP firmware, but this table should be sufficient to provide an understanding of the types of features that are available. In addition to the many features supported though this firmware, there is also support for 24 different languages. By enabling phones to support both the Enterprise firmware and the MPP firmware, Cisco is extending its reach to companies of every size, shape, and purpose.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 7-5 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 7-5 Key Topics for Chapter 7

Key Topic Element	Description	Page Number
Table 7-2	Cisco IP 7800 Series Phone Models and Features	152
Paragraph	Lines Supported on Different 7800 Series Phone Models	152
Paragraph	Color Indicator LED Reference	154
List	Power Save and Power Save Plus on 7800 Series Phones	154
Table 7-3	8800 Series Phone Model Features	155
Paragraph	Difference Between 8811 and 8841 Phones	155
Paragraph	Difference Between 8851, 8851NR, and 8861	156
Paragraph	Difference Between 8841 and 8845	156
Paragraph	Enterprise and MPP Phone Firmware	158
Table 7-4	MPP Firmware Feature Support	158

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Intelligent Proximity MV, KEM, LED, MPP, NR, PoE, Power Save, Power Save Plus, USB, WLAN

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the number of lines supported on each of the 7800 series phones.
2. List the number of KEMs supported on each of the Cisco IP Phone 8800 series models.
3. List five of the call control audio features supported through the MPP firmware.

This page intentionally left blank

CHAPTER 8

Cisco Telepresence Endpoints

This chapter covers the following topics:

CE Software: This topic will discuss the back-end software that runs on all the Cisco Telepresence endpoints mentioned in this chapter.

DX Series: This topic will introduce the DX series endpoint that is intended to be used as a personal desktop video communications device.

SX Series: This topic will introduce the SX series integrator endpoints that are intended to be used by professional integrators in multipurpose rooms.

MX Series: This topic will introduce the MX series all-in-one meeting room endpoints that are an easy installation option for meeting rooms where multiple participants will gather for local meetings as well as video call meetings with a remote location.

Webex Series: This topic will introduce the most recent endpoints in the Cisco product portfolio, the Webex series endpoints. These endpoints overlap with the SX and MX series but offer cutting-edge technology in the cameras, displays, speakers, microphones, and processing endpoints to deliver the best user experience in a video call that is available in the market today.

IX Series: This topic will introduce the only immersive product in the Cisco Telepresence endpoint portfolio, the IX5000. This room-within-a-room system offers best-in-class audio, video, and overall user experience.

Cisco has made its mission to be the top company in the IT industry that delivers the best products available in the market. It has not failed its mission with the development of the Cisco Telepresence endpoint product portfolio. These endpoints possess such advanced technology that they continually create a “wow” factor for anyone who has the chance to use them and see them perform. At the same time, Cisco has managed to keep the everyday end user in mind by complementing these endpoints with an easy-to-use touch controller that offers a seamless and consistent user experience no matter which of these endpoints is used. With Cisco Telepresence endpoints, the power is truly in the hands of the user. Topics discussed in this chapter include the following:

- CE Software
- DX Series
- SX Series
- MX Series

- Webex Series
  - Cisco Webex Room Kit Mini
  - Cisco Webex Room Kit
  - Cisco Webex Room Kit Plus
  - Cisco Webex Room Kit Pro
  - Cisco Webex 55 (Single and Dual)
  - Cisco Webex 70 (Single and Dual)
  - Cisco Webex Board
- IX Series

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 8-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 8-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
CE Software	1
DX Series	2
SX Series	3
MX Series	4
Webex Series	5–11
IX Series	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. In CE9.7.1 Cisco has added a graphical equalizer in the Audio Console to simplify usage and modification. How many different equalizer setups can be configured and attached to audio inputs and outputs?
   - **a.** 1
   - **b.** 2
   - **c.** 4
   - **d.** 8

2. Which of the following is a feature that comes with the Cisco DX80 endpoint?
   a. Touch 10 controller
   b. Document camera
   c. TRC 5 remote control
   d. IEEE 802.11AC wireless capability
3. Which of the following SX series endpoints supports the H.265HEVC codec?
   a. SX10
   b. SX20
   c. SX80
   d. All of these answers are correct.
4. What is the maximum number of multisite participants that can be supported on the MX700 endpoint at 1080p30 resolution?
   a. 2 + 1
   b. 3 + 1
   c. 4 + 1
   d. Multisite is not supported on these endpoints.
5. What is the horizontal field of view on the Cisco Webex Room Kit Mini?
   a. 120°
   b. 150°
   c. 170°
   d. 83°
6. How many participants does Cisco recommend in a space using the Webex Room Kit endpoint?
   a. 5
   b. 7
   c. 10
   d. 14
7. What is the zoom capability on the Cisco Webex Room Kit Plus cameras?
   a. 3x zoom
   b. 4x zoom
   c. 5x zoom
   d. 6x zoom
8. What type of audio input connectors exist in the back of the Cisco Webex Room Kit Pro for microphones?
   a. XLR
   b. Mini-jack
   c. USB
   d. Euroblock

9. Which of the following is a new technology that exists on the Cisco Webex Room 55 and not on the MX300G2?
   a. Touch 10 controller
   b. Multisite capabilities
   c. SIP and H.323 support
   d. Speaker tracking
10. How many participants are recommended in a space using the Cisco Webex 70 G2 endpoint?
   a. 24
   b. 12
   c. 28
   d. 14
11. Which of the following statements about the Cisco Webex Board is true?
   a. The Cisco Webex Board can now register to the CUCM and share whiteboarding during calls.
   b. The Webex Board cannot register to the CUCM but can be used for whiteboarding locally.
   c. The Webex Board can be used for whiteboarding locally or during a call only if it is registered to the Webex Control Hub.
   d. The Cisco Webex Board can now register to the CUCM but can share whiteboarding only during local meetings.
12. The Cisco IX5000 endpoint comes with two three-headed dongles. Which of the following are connections supported on these dongles? (Choose three.)
   a. Mini Display Port
   b. USB
   c. Thunderbolt
   d. HDMI
   e. Display Port
   f. VGA

Foundation Topics

CE Software

Tandberg was a leader in the video communication market for many years. The MXP endpoint product line the company offered was feature rich and simple to use; plus, it supported the best quality in video communications during that time. However, technology is always changing, and improvements were needed to continue leading the industry. Tandberg achieved a key asset when it acquired Codian. Although the main product behind this acquisition was the media resources Codian offered, Tandberg quickly used its technology to produce a new product line called the C-Series, which ran on the Tandberg/Codian (TC) software. After Cisco acquired Tandberg, it was not difficult for the company to see the value it had in these endpoints. Cisco continued to develop the software, eventually replacing its

own CTS software-based endpoints and has since come out with new endpoints based on this software code. Prior to developing the current product line of endpoints, Cisco made some improvements to the TC software so that when it launched the upgrade version to the eighth major adaptation, it also changed the name of the code to Collaboration Endpoint, or CE. With the vision of a new suite of products that would support this software, Cisco announced a lot of its video endpoints to go end of sale. At the time this chapter was written, the current CE software version is CE9.10, which was released on January 8, 2020. To see the current versions of Cisco Collaboration equipment, go to https://software.cisco.com.

If you have older equipment running TC software, and you want to upgrade, you must first make sure the endpoints are upgraded to TC7.3.6. From this TC version, you can proceed to upgrade to CE8.x or CE9.x. You can also downgrade from CE9 directly to CE8.x or TC7.3.6. The Cisco Room Kit Mini is initially shipped running version CE9.6, which is a special release for this system only; however, you can upgrade the software on this system, which Cisco does recommend. If you are using the Audio Console, which is a feature introduced with CE9.5, and you are upgrading to CE9.6.1 or later, you should make a note of the Audio Console setup because these settings will not transfer over to the newer versions, and a backup under TC9.5 cannot be restored after the upgrade is complete. You can manually configure Audio Console settings on one endpoint that has been upgraded and use a backup of that endpoint on other endpoints after they are upgraded.

With CE9.7.1 Cisco has introduced several new features, including added support for accessing the xAPI using a WebSocket. A WebSocket is a bidirectional persistent connection between a client and the server where information can flow back and forth without the overhead of initiating a new TCP connection or authentication for every request. From CE9.7.1 the room device will act as a WebSocket server with direct access to the xAPI using JSON RPC 2.0 as the data transport. After establishing a WebSocket connection to the room device, you can register feedback, execute configurations and commands, or get the status of the device by sending JSON documents over the WebSocket connection. The client will also receive unsolicited data from the server, such as feedback events if registered. This feature mainly targets integrators and is a modern alternative to access the xAPI compared to SSH or serial. For more information on how to get started with xAPI over WebSocket, refer to the official xAPI over WebSocket guide found at www.cisco.com/c/en/us/support/collaboration-endpoints/spark-room-kit-series/products-command-reference-list.html.

Ambient noise reporting is another new feature that uses the device microphones to estimate the ambient or background noise. The value is an A-weighted decibel value of the ambient noise level (dBa). Note that the value is not a calibrated sound pressure level (SPL), so it has to be evaluated as a relative value of the ambient sound level in the room. The feature is disabled by default and can be enabled by setting the ambient noise reporting to On in the xAPI using the command **xConfiguration RoomAnalytics AmbientNoiseEstimation Mode: On**. The estimated value is accessible via the xAPI with the command **xStatus RoomAnalytics AmbientNoise Level**.

Some other added features include Privacy mode, Room Kit Mini support for 1080p video, and editing favorites in the on-screen display (OSD). Privacy mode is a feature that adds a new button on the Touch 10 or in the on-screen UI, allowing you to disable video while in a call. The far-end participants will see only a placeholder image indicating that you have disabled the video on your device. On the local end, an icon will be displayed on the screen,

indicating that you are currently not sending video. With CE9.7.1 the Room Kit Mini will now support 1080p video while used as a USB camera. In previous versions, only 720p video resolution was supported. The DX70 and DX80 endpoints can now edit favorites from the OSD. This feature allows users to edit the contact information in their local favorites from the OSD interface of the DX70 and DX80 in the same way that was introduced for the Touch 10 in CE9.6.1. Note that this feature is not available when using the TRC6 remote control.

Key Topic

Audio Equalizer is available in previous software versions via the xAPI only. Since CE9.7.1, Cisco has added a graphical equalizer in the Audio Console to simplify usage and modifications. The graphical equalizer setup is available in the Audio Console app on the room device web interface. This new graphical equalizer allows you to create up to eight different equalizer setups and attach them to an output or a microphone input. For example, an equalizer can be used in scenarios where users want to tweak the audio experience on microphone inputs or if the output equipment is expressing too much bass or treble. The graphic equalizer makes it easier for users to customize the overall audio experience on the analog line inputs and outputs. Figure 8-1 illustrates the Audio Console for CE software-based endpoints.

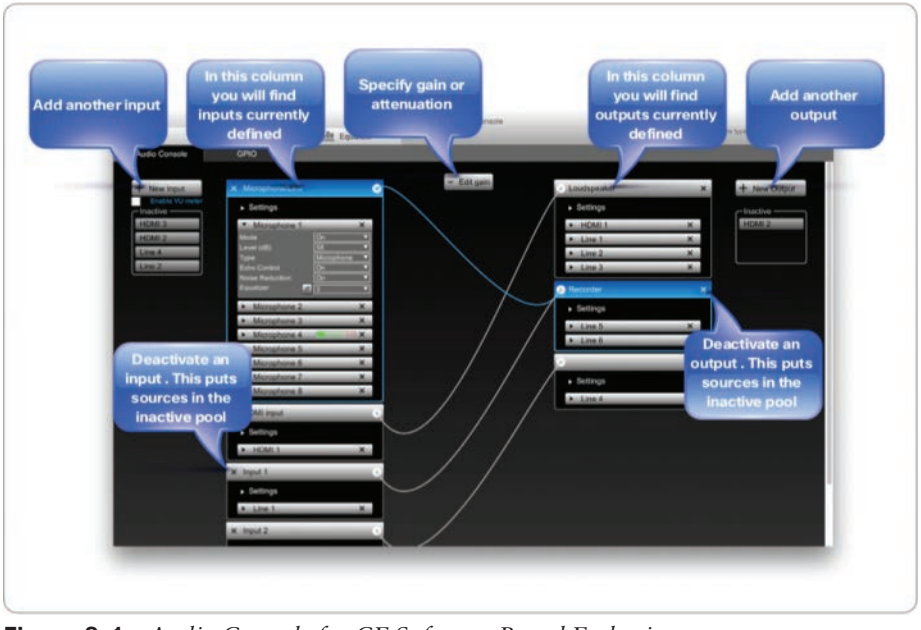

Figure 8-1 *Audio Console for CE Software-Based Endpoints*

Key Topic

One last feature worth mentioning that CE software supports is Intelligent Proximity for Content Sharing, which should not be confused with Intelligent Proximity for Mobile Voice. Intelligent Proximity for Content Sharing is a Cisco proprietary protocol that uses an ultrasonic audio tone, unheard by the human ear, that pairs the endpoint with the Intelligent Proximity application. This application can be installed on a smartphone, tablet, Mac computer, or Windows computer. Once paired, the application will use a Wi-Fi signal, which must be on the same network as the endpoint, to establish communication. Then the

Intelligent Proximity app can be used to view and select participants to call from the directories on the endpoint, launch calls, answer incoming calls, and view content being shared during a call. You can scroll back and view previously shared information even when the presenter is sharing something different, and you can take snapshots of the content to peruse after the call ends. When using Intelligent Proximity from a Mac or Windows computer, you can also share content through the application. This bring-your-own-device (BYOD) application has revolutionized collaboration as we know it, putting the power of information into the hands of the users within their own devices. Figure 8-2 illustrates how the Intelligent Proximity app can be used on a smartphone.

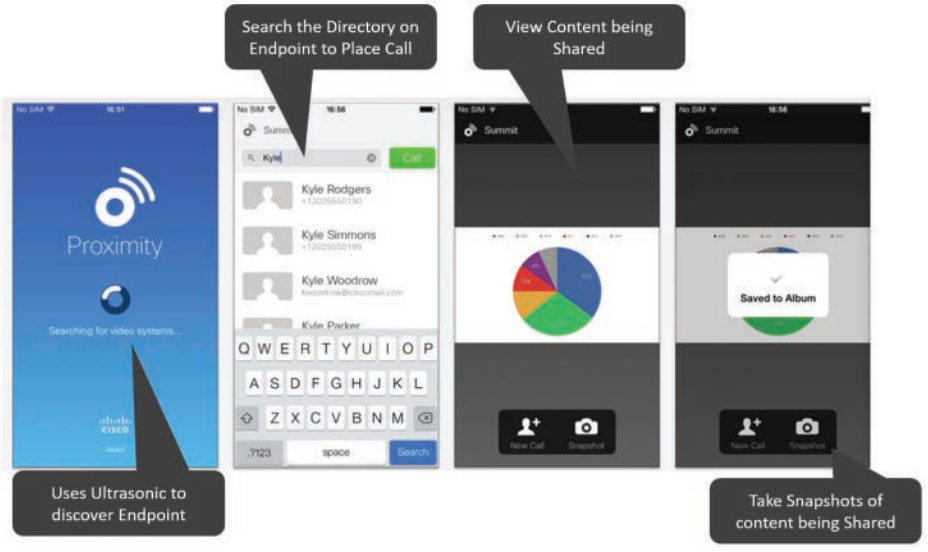

Figure 8-2 *Intelligent Proximity for Content Sharing on a Smartphone*

Cisco Telepresence is all about the user experience. It allows users to virtually be in two or more places at once while leveraging full-spatial audio, high-definition video, and interactive desktop communication. This life-like experience allows for face-to-face communication between users regardless of their location. The Cisco Telepresence product portfolio offers several platforms of products to meet the needs of any company, regardless of the number of participants or room size. Chapter 6, "Cisco Solution for Converged Collaboration," provided an introduction to the Cisco Telepresence product portfolio. This chapter will provide a more detailed explanation for each of the Cisco Telepresence endpoints. Some of the endpoints mentioned in this chapter have recently been announced as end of sale. Because these announcements were fairly recent, and some people are still using these endpoints in production environments, they will still be mentioned in their related sections. Since CE software is being used, it doesn't matter which product is used; the experience will be consistent with great audio and video quality.

DX Series

There are two products in the Desktop Experience (DX) portfolio from which companies can choose. These endpoints are targeted at management personnel and typically have one participant in the camera field of view. The endpoints in this portfolio include the Cisco DX70 and the Cisco DX80. At the time this chapter was written, the Cisco DX70 has become end

of sale but is still under Cisco support. When these endpoints first came out, they supported an Android operating system (OS), would only register to the Cisco Unified Communications Manager, and were treated as a Cisco UC phone. On April 20, 2017, Cisco announced the end of sale for the Android OS and encouraged companies to migrate the OS on these DX endpoints to CE software. With the CE software running on the DX endpoints, they are treated as a Telepresence endpoint on the Cisco Unified Communications Manager, will register to the Webex Control Hub via SIP, and will register to the Expressway via SIP or H.323.

Key Topic

The DX70 endpoint is a midsized personal desktop endpoint with a 14-inch multitouch capacitive display that supports a 1920 × 1080 resolution. The DX80 endpoint is a large personal desktop endpoint with a 23-inch multitouch capacitive display that supports a 1920 × 1080 resolution. Because both endpoints support the multitouch capacitive display, neither of them requires a Touch 10 or TRC remote; however, the menus on the display resemble the menus on the Touch 10, so user adoption should be easy and transferable to other CE software-based endpoints. These endpoints can double as a second monitor for a computer using the HDMI-in port, which helps reduce the number of devices on an office desktop. This same connection can be used to easily share content with far-end participants during a call through the simple touch of a button on the screen. They also support the Intelligent Proximity for Content Sharing application for easy synchronization of a person's smartphone, tablet, or computer. The manual HD camera attached to the system has a privacy shutter that can be opened or closed when needed. It also has manual tilt capability for positioning the camera, but no pan or zoom capability. If you point the camera straight down at your desktop, it will become a document camera that inverts the image so that it can be presented as content. Figure 8-3 illustrates the Cisco DX80 endpoint.

Figure 8-3 *Cisco DX80*

The Cisco Telepresence DX70 endpoint supports full HD resolutions up to 1080p30, and up to 48 kHz sampling rate for audio. It also includes a built-in acoustic echo canceller, automatic gain control, automatic noise reduction, and support for Bluetooth headsets using Bluetooth 3.0 (HFP, A2DP). Video codecs supported on the DX70 include H.263, H.263+, H.264, and H.264AVC. Audio codecs supported on the DX70 include AAC-LD,

OPUS, G.722, G.722.1, G.711mu, G.711a, G.729ab, and G.729. Lip synchronization is a feature built into H.323 but not SIP. However, the DX70 endpoints have active lip synchronization built into the endpoint. Content can be shared up to 1080p15 using H.239 over H.323, or BFCP over SIP. When the DX70 is registered to the Webex Control Hub, a whiteboard feature can be used to illustrate concepts or annotate on content being shared. This feature also works outside of a call no matter where the endpoint is registered. Additional features include +E.164 dialing support, adjustable ring and volume controls, adjustable display brightness, auto-answer, headset autodetection, call forward, caller ID, corporate directory with call history lists, ad hoc conferencing capabilities (conferencing services required through Cisco Unified Communications Manager), DND, Extension Mobility, hold, MWI, audio and video mute, self-view, OBTP, shared line, SNR, call transfer, and voicemail.

The Cisco Telepresence DX80 supports all the same features as the DX70, but the DX80 endpoint offers some added benefits. The DX80 supports a larger screen at 23 inches with high-contrast LED backlighting for a better user experience. The Automatic Wake-up feature is one of the added benefits to the DX80. This Collaboration endpoint can automatically detect when someone enters a room. It will wake up, say hello, and provide guided instructions, making it effortless for users to start using the device. The multisite feature on the DX80 allows multipoint meeting to be hosted directly on the endpoint itself without the use of an external bridge. The multisite feature supports connections to three participants—this system plus two others, or 2 + 1. Table 8-2 identifies some of the feature differences between the DX70 and DX80 endpoints.

Key Topic

Table 8-2 DX70 and DX80 Feature Differences

	DX70	DX80
Screen Resolution	14" 1920 × 1080	23" 1920 × 1080
Multisite	No	2 + 1
Contrast Ratio	700:1	1000:1
Wi-Fi Capable	802.11a, b, g, n	802.11a, b, g, n
EoS	August 16, 2018	No

Cisco announced at the Partner Summit in November 2019 the launch of two new endpoints: the Cisco Webex Desk Pro and the Cisco Webex Room Panorama. The Cisco Webex Room Panorama will be discussed more in the IX5000 section. The Cisco Webex Desk Pro is an AI-powered collaboration device for the desk. It is purpose-built for collaboration and features a stunning 4k display, advanced cognitive collaboration capabilities like Webex Assistant and facial recognition, and creative applications like digital whiteboarding. You can easily pair your device wirelessly or dock your laptop and quickly join or start your meeting with one button to push. With a USB-C connection, the Webex Desk Pro becomes an all-in-one primary monitor and collaboration device that supports your videoconferencing software of choice. The Webex Desk Pro is designed for personal desk-based collaboration and focus rooms that accommodate one or two people. Packed with all the workplace and workflow capabilities included within Cisco's larger meeting room devices, the Webex Desk Pro is the ultimate desk-based collaboration device. At the time the Cisco Webex Desk Pro

was announced, Cisco stated that it did not intend to replace the DX80 endpoint with this new addition.

SX Series

The Cisco Telepresence Solutions Experience (SX) products are integrator solutions that come with the codec, microphone, camera, and cables. The customer needs to supply the monitor and speakers. These products allow meeting environments to be customized to the specific needs of the customer. The extent to how extensive an environment can be customized depends on the product within the SX series that is being used, and there are three SX products from which to choose: SX10, SX20, or SX80.

The Cisco Telepresence SX10 Quick Set is a low-cost high-quality endpoint. The endpoint comes with the codec, internal microphone, and camera in compact packaging that is mountable on the top of most flat-screen displays. The Cisco SX10 is designed for a small meeting room with up to six participants. At the same price as roughly that of a computer, the SX10 is a low-cost entry point for organizations looking to purchase their first Telepresence devices or to extend their current topology. These endpoints can register to the Cisco Unified Communications Manager, the Cisco Expressway Core, or to the Webex cloud using SIP only. Included with the purchase of the Cisco SX10 is the TRC6 remote to navigate the physical interface and cables for a basic installation. Alternatively, the Intelligent Proximity for Content Sharing application can be used to control the SX10. A microphone is built into the endpoint, but an external Cisco Telepresence Precision MIC 20 can be used in addition to the built-in microphone of the SX10. The camera for the SX10 is integrated directly into the unit as well. The camera has an optical zoom of 2.65x and a total zoom of 5x.

Key Topic

There are two choices in regard to powering the SX10. Either an external power cube, which is included with the system, or Power over Ethernet (PoE) can be used to supply the 12 watts needed for operation. This is the only Cisco Telepresence endpoint that can be powered by PoE currently. Therefore, one of the nice benefits of the Cisco Telepresence SX10 Quick Set is that a single Ethernet cable can be used for both power and Ethernet connectivity. Add in a single HDMI cable to the monitor of your choice, and only two cables are needed to complete the installation of this endpoint. The Cisco Telepresence SX10 Quick Set supports high-definition video with up to 1080p30 resolution at up to 3 Mbps. Figure 8-4 illustrates all the components that come with the Cisco Telepresence SX10 Quick Set endpoint.

Key Topic

The Cisco Telepresence SX20 Quick Set is a low-price multipurpose set for simple and flexible meeting room installations that will accommodate up to 12 participants. This device comes standard with the Cisco Telepresence SX20 codec, a Cisco Telepresence Precision Camera, a Cisco Telepresence Precision MIC 20 (which can be connected through a mini-jack port), a Cisco Telepresence TRC5 remote control, and basic cables with power supply. Optional hardware that can be ordered with this endpoint include a Cisco Touch 10 control pad, wall-mount kit, one additional Cisco Telepresence Precision MIC 20, camera-mount bracket, or a spare TRC5 remote control. The codec is based on video-conferencing standards and can register to the Cisco Unified Communications Manager, Expressway, or Webex Control Hub via SIP, or to the Cisco Expressway or a third-party call control system via H.323. Video resolutions up to 1080p60 are supported on the SX20 endpoint using H.264AVC up to 6 Mbps per call. The Cisco Telepresence SX20 Quick Set has three camera options: The Cisco Telepresence PrecisionHD camera comes in 4x or 12x zoom-capable models, or the Precision 40 Camera with 8x zoom and 4x optical plus digital capable model. The Cisco Telepresence SX20 Quick

Figure 8-4 *Cisco Telepresence SX10 Quick Set*

Set also supports the option for dual-display and support for a four-way (3 + 1) multipoint call using the Multisite option key. The Cisco Telepresence SX20 codec supports an HDMI video input with VISCA far-end camera control to connect the Cisco Telepresence PrecisionHD Camera. A DVI-I input and an HDMI input are available for the connection to a PC for presentation sharing. The Cisco Telepresence SX20 codec supports audio on the HDMI output port for the monitor, but a one-line audio output is available for external speakers if the speakers in the monitor are not adequate. The network connection on the Cisco Telepresence SX20 supports Gigabit Ethernet. Figure 8-5 illustrates all the components that come with the Cisco Telepresence SX20 Quick Set endpoint.

Figure 8-5 *Cisco Telepresence SX20 Quick Set*

Key Topic

The Cisco Telepresence SX80 is a flexible integrator video endpoint for medium to large meeting rooms and boardrooms. It could also be used in larger auditorium-style rooms. The SX80 is sold with three different integration packages. For small deployments, the SX80 is sold with

the Precision HD 1080p 4x camera. For larger rooms, the SX80 can be sold with the Precision 60 camera or the Speaker Track 60 dual-camera system, which is the same Speaker Track system sold with the MX700 and MX800 endpoints. The SX80 supports either H.323 or SIP single-stack call registration, meaning it supports either the H.323 or SIP protocols for registration, but not both concurrently. The SX80 supports the multisite conferencing option key with capacity to support five participants (4 + 1) in a single multipoint conference. All current Cisco Telepresence devices support the H.264 video standard. The SX80 is the first Telepresence endpoint to be able to support the H.265 High Efficiency Video Codec (HEVC) video standard. Because the MX700 and MX800 are built using the SX80 codec, they too can support the H.265 HEVC video codec. The SX80 codec is a powerful audio and video engine. It incorporates high-definition video collaboration applications into large meeting rooms, boardrooms, and purpose-built or vertical application rooms such as training, briefing, demo rooms, and auditoriums. The SX80 delivers up to 1080p60 end-to-end high-definition video. It offers industry-first support for H.265, which lays the foundation for future bandwidth efficiencies that the new standard makes possible. The codec offers a rich input and output set and flexible media engine. It supports three screens to help enable various use cases that are adaptable to your specific needs. The SX80 can support up to four HD cameras and eight microphones directly connected to the codec. It has an eight-port audio mixer and eight separate acoustic echo cancellers for each microphone connection. Figure 8-6 illustrates the front and back views of the Cisco Telepresence SX80 endpoint.

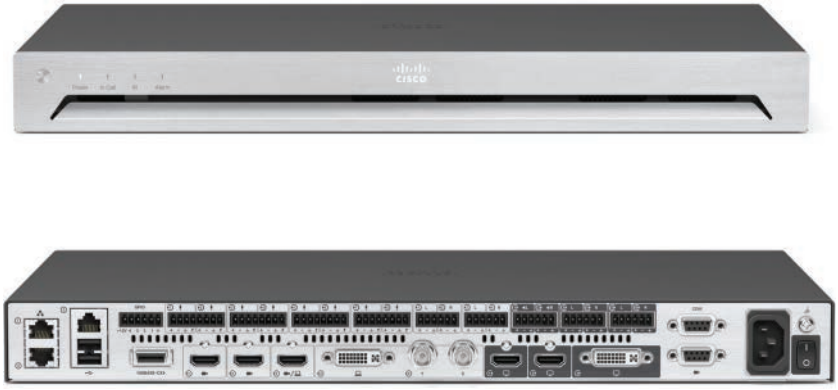

Figure 8-6 *Cisco Telepresence SX80*

Based on the preceding descriptions, there are some pretty obvious and significant differences between these three integrator solutions. The SX10 is a small but easy-to-install solution that brings powerful HD video capabilities to small meeting rooms for businesses. The SX20 is a little bit more complex to set up but brings with it more powerful tools to enhance the meeting experience and extend the number of participants to a medium-sized meeting space. The SX80 extends the level of room customization to astronomical levels. With extensive tools built into the SX80 endpoint, the reach of local and far-end participants goes beyond what either of the previous products can provide. Due to the introduction of several integrator Webex endpoints, which will be discussed later in this chapter, the SX20 and SX80 endpoints were scheduled to go end of sale in October 2019. The SX10 reached end of sale in January 2020; however, Cisco will continue to support these powerful products for five years beyond the EoS date. Table 8-3 illustrates some of the differences between these three SX integrator products.

Key Topic

Table 8-3 SX Endpoint Feature Differences

	SX10	SX20	SX80
Video Resolution	1080p30	1080p60	1080p60
Protocol Support	SIP	SIP/H.323	SIP/H.323
Multisite	N/A	2 + 1 at 720p30 3 + 1 at 576p30	4 + 1 at 720p30 3 + 1 at 1080p30
Video Codec	H.264AVC	H.264AVC	H.265HEVC
Display Support	1	2	3
Bandwidth Support	3 Mbps point-to-point	6 Mbps point-to-point or multipoint	6 Mbps point-to-point 10 Mbps multipoint
EoS	January 28, 2020	October 29, 2019	October 29, 2019

MX Series

The Multipurpose Experience (MX) portfolio is made up of four main products that can be deployed quickly and easily in any multipurpose conference room. These units can be leveraged to support local meetings or conference calls out to remote locations. Because these units are an all-in-one system, installation is as simple as removing the unit from the box; setting up the mounting option you chose; and connecting the power, Ethernet, microphone, and touch controller. No other cables need to be connected. Which unit should be installed depends on the size and dimensions of the room within which it will be utilized.

The Cisco Telepresence MX 200G2 endpoints come with a 42-inch display that supports resolutions up to 1920 × 1080. The Cisco Telepresence MX 300G2 endpoints come with a 55-inch display that supports resolutions up to 1920 × 1080. Both systems support a DVI-I connected device and an HDMI-connected device for content sharing and support of PC-input resolutions from SVGA (800 × 600) to 1080p (1920 × 1080). These endpoints also support an embedded four-way multisite conferencing option (3 + 1). The Cisco Telepresence MX200G2 and MX300G2 endpoints have one integrated full-range microphone and an integrated full-range speaker system. The system also supports up to two instances of the Cisco Telepresence Precision MIC 20 so that the audio input range can be extended. One Cisco Telepresence Precision MIC 20 microphone is included with the MX200G2, and a second mic can be purchased as an add-on option. Two Cisco Telepresence Precision MIC 20 microphones are included with the Cisco Telepresence MX300G2.

The following mounting options are available for these endpoint solutions:

Key Topic

- Floor stand
- Table stand (Cisco Telepresence MX200G2 only)
- Wall mount (which doubles as a bracket for VESA mount systems; VESA is a standard for video-display mount systems)
- Wheel base

The Cisco Telepresence MX200G2 and MX300G2 endpoints are easy to install. Each purchase includes an all-in-one monitor, codec and camera unit, and cables, plus your choice of mounting option. These endpoints can register to the Cisco Unified Communications Manager using SIP. They can register to the Cisco Expressway Core using either SIP or H.323. They can also register to the Webex cloud using SIP. The MX200G2 and the MX300G2 endpoints were announced to be end of sale on May 2, 2018. The replacement product is the Cisco Webex Room 55. Figure 8-7 illustrates the MX300G2 with mounting options.

Figure 8-7 *Cisco Telepresence MX300G2 with Mounting Options*

Key Topic

The Cisco Telepresence MX700 and MX800 represent the performance line in the Cisco integrated video collaboration room systems. The MX700 and MX800 systems come standard with a built-in amplifier and speaker system for high-fidelity sound. You can choose from a powerful single camera or an intelligent dual-camera speaker-tracking solution. Both cameras provide 1080p60 resolution and support 20x total zoom (10x optical, 2x digital zoom.) The MX700 and MX800 are driven by the SX80 codec integrated into the system. Should an engineer need to access the connectors on this codec, such as to connect an external display, the cover on the left side of the system can be removed. The cover is fastened with magnets. Premium resolution and dual display are also standard features on the MX700 and MX800. The intuitive Cisco Telepresence Touch 10 provides an easy-to-use interface for both MX700 and MX800 systems. The Cisco Telepresence MX700 has the options of one or two 55-inch (1.4m) TFT-LCD monitors. The Cisco Telepresence MX800 has the options of one or two 70-inch TFT-LCD monitors. The displays on both systems have a resolution of 1920 × 1080 (16:9) with a contrast ratio of 4000:1. Both systems also have 3 HDMI and 1 DVI-I video inputs and 15 audio inputs. Mounting options include either a wall-mounted solution or a floor-stand mounting solution. The options for the monitors, camera, and mounting must all be selected at the time of purchase because these units are customized on a per-order basis. These endpoints can register to the Cisco Unified Communications Manager using SIP. They can register to the Cisco Expressway Core using either SIP or H.323. They can also register to the Webex cloud using SIP. Figure 8-8 illustrates the single-screen and

dual-screen options, as well as the single camera and dual camera with speaker-tracking options on the Cisco Telepresence MX800 endpoint.

Figure 8-8 *Cisco Telepresence MX800 with Display and Camera Options*

As with other Cisco endpoint product portfolios, there are some significant differences between the MX series endpoints mentioned in this section. The first endpoint that Cisco released after the Tandberg merger was the MX200 (G1), which was quickly followed by the MX300 (G1). These two endpoint systems were all-in-one contained units that allowed companies to deploy a room system in 10 minutes or less. Just take them out of the box, attach the mounting system, plug in the cables, and you were ready to use the systems. Improving on the display and codec capabilities of these two endpoints, a few years later Cisco released the MX200G2 and MX300G2. They were followed shortly by the MX700 and MX800, which bring 4k display and H.265HEVC video compression along with them. The MX700 and MX800 also have higher multisite capability, the Speaker Track 60 dual-camera option, and the option for a built-in second display for content sharing. They offer a more uniform look to room integration and offer much higher capabilities to enhance the user experience. Table 8-4 illustrates some of the feature differences between the MX series endpoints.

Key Topic

Table 8-4 Cisco Telepresence MX Series Feature Differences

	MX200G2	MX300G2	MX700	MX800 (Single or Dual)
Display Size	42" 1920 × 1080 with 1300 contrast ratio	55" 1920 × 1080 with 4000:1 contrast ratio	55" 1920 × 1080 with 4000:1 contrast ratio	70" 1920 × 1080 with 4000:1 contrast ratio
Mounting Options	Floor stand, wheel base, table stand, wall mount	Floor stand, wheel base, table stand, wall mount	Floor stand, wall mount	Floor stand, wall mount
Multisite	2 + 1 at 720p30 3 + 1 at 576p30	2 + 1 at 720p30 3 + 1 at 576p30	4 + 1 at 720p30 3 + 1 at 1080p30	4 + 1 at 720p30 3 + 1 at 1080p30
Video Codec	H.264AVC	H.264AVC	H.265HEVC	H.265HEVC

	MX200G2	MX300G2	MX700	MX800 (Single or Dual)
Camera	2.5x optical zoom (5x with digital)	4x optical zoom (8x with digital)	20x total zoom (10x optical, 2x digital)	20x total zoom (10x optical, 2x digital)
Bandwidth Support	6 Mbps point-to-point or multipoint	6 Mbps point-to-point or multipoint	6 Mbps point-to-point 10 Mbps multipoint	6 Mbps point-to-point 10 Mbps multipoint
EoS	May 2, 2018	May 2, 2018	April 1, 2019	April 1, 2019

Webex Series

When Cisco released the cloud solution called Cisco Spark, it began developing a line of endpoints with cloud support in mind. The first product it released was the Spark Board, which brought whiteboarding capabilities into the video meeting, along with annotation capabilities. The original Spark Board had a unique OS built specifically for its intended purpose; therefore, the Spark Board would only register to Cisco Spark in the cloud. Shortly after the Spark Board came out, Cisco began changing the current endpoints in its portfolio so that CE software-based endpoints could register to on-premises infrastructure or to the cloud. This brought about the development of the Spark Room Kit, which is based on the same CE software. When Cisco changed its cloud collaboration solution from Spark to Webex, the name changed on the endpoints as well, but the functions all stayed the same. Since then, Cisco has continued to develop a whole line of products called the Collaboration Room Endpoints, some of which have already replaced the Cisco Telepresence products discussed earlier in this chapter. For the record, all Cisco CE software-based endpoints can register to the Webex Control Hub in the cloud or to on-premises infrastructure, such as the Cisco Unified Communications Manager, Expressway, or other standards-based third-party call control systems. The following subsections of this Webex Series topic will delve into the more recent Collaboration Room Endpoints Cisco has developed.

Cisco Webex Room Kit Mini

The Webex Room Kit devices are all integrator-style endpoints, similar to the SX series. However, Cisco has added a lot of enhanced features and capabilities to each of these devices. Launched in March 2019, the latest endpoint in the Webex Series is the Cisco Webex Room Kit Mini, which is an artificial intelligence (AI)-powered video-conferencing system custom-designed for the huddle workstyle, and is easy to use, deploy, and manage. Don't worry; the Webex Room Kit Mini isn't going to take over the world. It combines codec, camera, microphones, and speakers into a single device that integrates with a 4k display supplied by the customer to bring more intelligence and usability to all of your huddle rooms and spaces. Room Kit Mini is rich in functionality and experience, while priced and designed to be easily scalable. The Room Kit Mini is ideal for huddle spaces with three to five people because of its wide 120-degree field of view, which allows everyone in a huddle space to be seen. It also offers the flexibility to connect to laptop-based video-conferencing software via USB. The Mini is tightly integrated with the industry-leading Cisco Webex platform for continuous workflow and can register on premises to the Cisco Unified Communications Manager via SIP, the Expressway via SIP or H.323, or to Cisco Webex in the cloud. Figure 8-9 illustrates the Cisco Webex Room Kit Mini.

Figure 8-9 *Cisco Webex Room Kit Mini*

Key Topic

The Cisco Webex Room Kit Mini brings intelligent views to smaller rooms with a discreet, integrated camera. The system will "wake up" automatically when someone walks into the room and will recognize who entered the room through their mobile device syncing with Intelligent Proximity or Webex Teams enabled. The system can be easily controlled with the Cisco Touch 10 controller as well. When the meeting begins, the camera will automatically detect meeting participants and provide an ideal framing base on their location within the room. The integrated microphones and speakers provide a great audio experience during the meeting as well. Add to that the automatic noise suppression, which reduces disruptive sounds coming from the meeting room, such as typing, paper rustling, pencil tapping, or other such noises. Not only are meetings smarter, but presentations are smarter too. Content can be shared using a wired or wireless connection to your PC or other device, and you can share clearer content with 4k content-sharing capabilities. The AI integrations bring a smarter room all around. There are built-in metrics that count people in the room, enabling analytics for better resource planning. The system can connect to the network using Wi-Fi, and it supports Bluetooth as well. In-room controls are also built into the system so that peripherals, such as lights and blinds, can be controlled through the Touch 10 controller. Table 8-5 outlines some of the features supported on the Webex Room Kit Mini.

Key Topic

Table 8-5 Webex Room Kit Mini Features

Feature	Description
Bandwidth	Up to 6 Mbps point-to-point
Resolution	Up to 4k video input and output at 30 fps or 1080p60
Audio Features	High-quality 20 kHz audio Automatic gain control Automatic noise reduction Active lip synchronization
Content Sharing	H.239 and BFCP up to 3840 × 2160p5
Wireless Sharing	Webex Teams App Webex Meetings App Intelligent Proximity
Multipoint Support	2 + 1 up to 1080p30 3 + 1 up to 720p30
Protocols	SIP, H.323, and Webex
Camera	4k UltraHD 2x zoom, autoframing, autobrightness and white balance, focus distance 1 m to infinity

Cisco Webex Room Kit

Key Topic

The Cisco Webex Room Kit was one of the first products created in the Cisco Webex endpoint portfolio, and many of the advanced features included in the Room Kit Mini were first incorporated in the Cisco Webex Room Kit, such as the AI functionality. Like the Cisco Telepresence SX80 endpoint, the Webex Room Kit endpoint supports the H.265 HEVC standard. The Cisco Webex Room Kit delivers the unmatched video and audio experience customers have come to expect from Cisco. In addition, new capabilities enable even smarter meetings, smarter presentation capabilities, and smarter room and device integrations—further removing the barriers to usage and deployment of video in small to medium-sized rooms. The Room Kit includes camera, codec, speakers, and microphones integrated in a single device. It also comes with a Touch 10 controller, but the monitor and speakers are not provided. It is ideal for rooms that seat up to seven people. It offers sophisticated camera technologies that bring speaker-tracking capabilities to smaller rooms. The product is rich in functionality and experience but is priced and designed to be easily scalable to all of your small conference rooms and spaces. Although the Room Kit was built to enhance the user experience in the cloud, it can be registered on the premises to the Cisco Unified Communications Manager using SIP, or to the Cisco Expressway Core using SIP or H.323, or to Cisco Webex in the cloud using SIP. The OS used by this system is the same CE software found on all the aforementioned Cisco Telepresence endpoints. Figure 8-10 illustrates the Cisco Webex Room Kit bottom and front views.

Figure 8-10 *Cisco Webex Room Kit*

The camera on the Cisco Webex Room Kit is a 4k Ultra HD camera with 3x digital zoom and an 83-degree horizontal field of view and a 51.5-degree vertical field of view. It supports resolutions up to 1080p60 and features autoframing by speech detection and facial recognition, autofocus, autobrightness, and autowhite-balancing. The connectors on the back of the Room Kit allow for one or two displays to be connected; the second display is used for content. Content can be shared over an HDMI cable through the video input, wirelessly through Intelligent Proximity, or wirelessly through the Webex Teams app when registered to the Webex Control Hub. The Cisco Webex Room Kit does have a built-in microphone with a six-element microphone array to provide accurate speaker tracking; plus, there are two 4-pin mini-jack audio inputs for table mics to be added. Five integrated high-quality speakers in balance have a frequency response from 70 Hz to 20 kHz, 24 W of amplifier power, and a max output level of 86 dB. This all translates to best-in-the-industry audio and video quality. Table 8-6 outlines some of the features supported on the Webex Room Kit.

Key Topic

Table 8-6 Webex Room Kit Features

Feature	Description
Bandwidth	Up to 6 Mbps point-to-point
Resolution	Up to 4k video input and output at 30 fps or 1080p60
Audio Features	High-quality 20 kHz audio Subwoofer line out Automatic gain control Automatic noise reduction Active lip synchronization
Content Sharing	H.239 and BFCP up to 3840 × 2160p5 or 1080p30
Wireless Sharing	Webex Teams App Webex Meetings App Intelligent Proximity
Multipoint Support	2 + 1 up to 1080p30 3 + 1 up to 720p30
Protocols	SIP, H.323, and Webex
Camera	5k UltraHD 3x zoom, 15.1 MP image sensor, 1/1.7 CMOS, autoframing, autobrightness and white balance, focus distance 1 m to infinity

Cisco Webex Room Kit Plus

Key Topic

The Cisco Telepresence Room Kit Plus was released at the same time as the Cisco Webex Room Kit and is similar in nature to the Room Kit but supports a quad-camera bar with integrated speakers and microphones, making it ideal for rooms that seat up to 14 people. It offers sophisticated camera technologies that bring speaker-tracking and autoframing capabilities to medium or large-sized rooms. Because the video capabilities of the Room Kit Plus are much more advanced than the Room Kit, the codec is external to the rest of the unit itself. The product is rich in functionality and experience but is priced and designed to be easily scalable to all of your conference rooms and spaces—whether registered on the premises or to Cisco Webex in the cloud. Figure 8-11 illustrates the Cisco Webex Room Kit Plus.

Figure 8-11 *Cisco Webex Room Kit Plus*

One of the prominent differentiators between the Room Kit Plus from the Room Kit or the Room Kit Mini is the quad-camera capabilities and the separate codec system that brings more processing power to this particular solution. However, the Room Kit plus still brings the same intelligence to medium- and large-sized meeting rooms. While other companies are still struggling to insert advanced features such as speaker tracking, wireless sharing, and 4k content into their high-end products, Cisco is already delivering these innovations to meeting rooms of all sizes in a cost-effective and simple way. With the Room Kit Plus, Cisco is helping customers experience smarter meetings, enable smarter presentations, and create smarter room and device integrations. These features were previously the domain of higher-end video-conferencing rooms but can now be brought to every room and every team. And when registered to Cisco Webex, additional cloud-based functionalities are enabled that enhance the user experience and team workflow and further simplify deployment. The Cisco Webex Room Kit Plus comes standard with the Cisco Webex Codec Plus, the Cisco Webex Quad Camera and Sound Bar, a Cisco Touch 10 controller, and a wall-mount kit for the quad camera. It supports the H.265 High Efficiency Video Codec, and two HDMI video outputs for 4k video up to 30 frames per second or 1080p60. A third HDMI output will support 1080p60 for content sharing. Because the Cisco Webex Room Kit Plus was designed for larger meeting rooms, there are three 4-pin mini-jacks for external microphones in addition to the six-element microphone array built into the system. Table 8-7 outlines some of the features supported on the Webex Room Kit.

Key Topic

Table 8-7 Webex Room Kit Plus Features

Feature	Description
Bandwidth	Up to 6 Mbps point-to-point
Resolution	Up to 4k video input and output at 30 fps or 1080p60
Audio Features	High-quality 20 kHz audio Subwoofer line out Prepared for inductive loop (line out) Automatic gain control Automatic noise reduction Active lip synchronization
Content Sharing	H.239 and BFCP up to 3840 × 2160p5 or 1080p30
Wireless Sharing	Webex Teams App Webex Meetings App Intelligent Proximity
Multipoint Support	2 + 1 up to 1080p30 3 + 1 up to 720p30
Protocols	SIP, H.323, and Webex
Camera	5k UltraHD 5x zoom, 15.1 MP image sensor, 1/1.7 CMOS, autoframing, autobrightness and white balance, focus distance 1 m to infinity

Cisco Webex Room Kit Pro

Key Topic

The Cisco Webex Room Kit Pro provides a powerful and flexible platform for creating the ultimate video collaboration experience. Similar to the SX80, the Room Kit Pro can be integrated in large custom video rooms, including boardrooms, auditoriums, and purpose-built rooms for vertical applications. The Room Kit Pro is built with integrators in mind and enables flexibility and creativity for customized video collaboration rooms that delight customers. The Room Kit Pro acts as the audio and video engine for UltraHD video collaboration applications and AV integrations in which up to three screens, multiple cameras, and multiple content sources can be leveraged. Camera options include the Cisco Webex Quad Camera used with the Cisco Webex Room Kit Plus, Precision 60 camera, or SpeakerTrack 60 dual camera. The codec that drives this system is the component that sets it apart from the rest. The Room Kit Pro continues with the same artificial intelligence capabilities already offered on the rest of the Cisco Webex Room Series, including intelligent views, noise suppression, voice commands, and people count. Figure 8-12 illustrates the Cisco Webex Room Kit Pro codec connection ports.

Figure 8-12 *Cisco Webex Room Kit Pro*

The Room Kit Pro delivers up to 2160p60 end-to-end UHD video. The codec's rich set of video and audio inputs, flexible media engine, and support for up to three screens enable a variety of use cases, adaptable to your specific needs. The Room Kit Pro can register on premises or to Cisco Webex in the cloud. The Room Kit Pro supports intelligent cameras and functionality to enable dynamic viewing capabilities in video meetings. The Cisco Webex Quad Camera system with four embedded digital cameras enables the best overview and speaker-tracking capabilities. With the quad camera, the Room Kit Pro can also deliver analytics such as people count. The Cisco SpeakerTrack 60 dual-camera system features a unique direct fast-switching approach for speaker tracking with two Precision 60 cameras. Cisco PresenterTrack makes it easier for presenters to move around the front of the room, using the Precision 60 camera to follow a presenter within a defined zone. With its powerful media engine, the Room Kit Pro lets you build the video collaboration room of your dreams. Cisco offers three options to purchase the Cisco Webex Room Kit Pro:

Key Topic

- **Room Kit Pro:** Codec Pro, Quad Camera, and Touch 10
- **Room Kit Pro Precision 60:** Codec Pro, Precision 60 camera, and Touch 10
- **Codec Pro:** Codec only

The Cisco Webex Room Kit Pro is video innovation in a box, bringing more intelligence and usability to your large specialized collaboration rooms and spaces. The Room Kit Pro supports up to six simultaneous video inputs: three 4k and three 1080p. It will also support up to eight microphones directly connected to the endpoint. More can be added by changing the audio input from mic level to line level and adding an external equalizer or amplifier. All microphone connections use the Euroblock port so that cable lengths can be more easily customized. Table 8-8 identifies some of the features supported on the Webex Room Kit Pro.

Key Topic

Table 8-8 Webex Room Kit Pro Features

Feature	Description
Bandwidth	Up to 6 Mbps point-to-point up to 15 Mbps multisite
Resolution	Up to 4k video input and output at 30 fps or 1080p60
Video Inputs	2 HDMI up to 1080p60 3 HDMI up to 3840 × 2160p30 1 3G-SDI/HD-SDI up to 1080p60
Video Outputs	2 HDMI up to 3840 × 2160p60 1 HDMI up to 3840 × 2160p30
Audio Features	High-quality 20 kHz audio 8 separate acoustic echo cancellers 8-port audio mixer 8 assignable equalizers Automatic gain control Automatic noise reduction Active lip synchronization
Audio Inputs	8 microphones, 48V phantom powered, Euroblock connector, mic level or balanced line level 3 HDMI outputs
Audio Outputs	6 balanced line-level outputs, Euroblock connector 3 HDMI outputs HDMI Input #1 supports Audio Return Channel (ARC) audio output to Cisco Webex Quad Camera 1 Line out for Subwoofer (Cisco Webex Quad Camera)
Multipoint Support	2 + 1 up to 1080p30 3 + 1 up to 720p30 4 + 1 up to 720p30
Network Interfaces	1 Ethernet 10/100/1000 for LAN 2 Ethernet 10/100/1000 for direct pairing with camera 2 Ethernet 10/100/1000 with PoE, 1 dedicated for direct pairing with Touch 10 Wi-Fi 802.11a/b/g/n/ac 2.4 GHz and 5 GHz for LAN 2 × 2 Multiple Input and Multiple Output (MIMO) Bluetooth 4.0 LE

Cisco Webex 55 (Single and Dual)

Key Topic

The Cisco Webex Room 55 is the replacement product for the Cisco Telepresence MX200G2 and MX300G2 endpoints. The Room 55 includes camera, codec, display, speaker system, and microphones integrated in a single device and is optimized for rooms that seat up to seven people. It is an all-in-one, integrated system that's easy to install, use, deploy, and manage. It's crafted with high-quality components: professional 4k display for longevity and minimal

latency, powerful digital zoom camera for discreet tracking, and sophisticated speaker system and amplifier that deliver rich sound. The light industrial design combines aluminum and fabric for a sustainable and humanizing effect. The Cisco Webex Room 55 has the Cisco Webex Room Kit as its base technology, bringing new capabilities such as speaker tracking, best overview, automatic wake-up, and people count to enable even smarter meetings, smarter presentation capabilities, and smarter room and device integrations. The Room 55 is rich in functionality and experience but is priced and designed to be easily scalable to all of your meeting rooms and spaces, whether registered on-premises to the Cisco Unified Communications Manager via SIP, or the Cisco Expressway Core via SIP or H.323, or to Cisco Webex cloud via SIP. Figure 8-13 illustrates a Cisco Webex Room 55.

Figure 8-13 *Cisco Webex Room 55 Dual*

The Cisco Webex Room 55 was granted the Red Dot award for innovation in design in 2017. This unit utilized the same 5k UltraHD camera as the Cisco Webex Room Kit and a 4k display to provide powerful and clear video up to 4kp30. This H.265-compliant system can be ordered with a single display or with a dual display. The integrated Room Kit is located above the display for better camera angles and for better audio distribution throughout the meeting room. The integrated speakers offer stereo quality with a dedicated center speaker for optimal voice pick-up, and there is a built-in amplifier that delivers rich sound. The speakers are covered in a fabric for a more natural and inviting feel. The frame of the system is built using aluminum for a lighter design and sustainability. The height is increased over the MX300G2 to accommodate taller tables. The Webex Room 55 continues to support the same mounting options as the MX300G2 endpoints with your choice between a floor stand, wheel base, or wall mount. All three of these order options come with a Touch 10 controller, two additional table mics, and all the other cables needed to initially set up the system.

Cisco Webex 70 G2 (Single and Dual)

Key Topic

The Cisco Webex Room 70 G2 is similar to the Cisco Telepresence MX800. They both present a 70-inch display, they both support H.265HEVC, and they can both be ordered with a single or dual display on a floor stand or wall mount. However, there are some distinct differences between the two as well. Where the MX800 is equipped with the Precision 60 or SpeakerTrack

60 cameras, the Cisco Webex Room 70 G2 is built off the Cisco Webex Room Kit Plus. This system offers a powerful codec, a quad camera, and 70-inch single or dual 4k display(s) with integrated speakers and microphones, making it ideal for rooms that seat up to 14 people. It offers sophisticated camera technologies that bring speaker-tracking and auto-framing capabilities to medium and large-sized rooms, plus all of the AI integrations offered with the Cisco Webex Room Kit series of endpoints. The product is rich in functionality and experience but is priced and designed to be easily scalable to all of your conference rooms and spaces—whether registered on the premises to the Cisco Unified Communications Manager via SIP, or the Cisco Expressway Core via SIP or H.323, or to Cisco Webex in the cloud. Figure 8-14 illustrates the Cisco Webex Room 70 G2 product.

Figure 8-14 *Cisco Webex Room 70 G2*

Cisco Webex Board

With the Cisco Webex Board, you can wirelessly present, whiteboard, video- or audio-conference, and even annotate shared content. It has everything you need for team collaboration at the touch of a finger. You can use the Cisco Webex Teams app to connect with virtual team members through the devices of their choice. This product comes in three models: the Cisco Webex Board 55s, the Cisco Webex Board 70s, or the new Cisco Webex Board 85. The Cisco Webex Board 55s is a fully self-contained system on a high-resolution 4k 55-inch LED screen with an integrated 4k camera, embedded microphones, and a capacitive touch interface. The Cisco Webex Board 70s is a fully self-contained system on a high-resolution 4k 70-inch LED screen with an integrated 4k camera, embedded microphones, and a capacitive touch interface. The Cisco Webex Board 85 is a fully self-contained system on a high-resolution 4k 85-inch LED screen with an integrated 4k camera, embedded microphones, and a capacitive touch interface. Figure 8-15 illustrates the three Webex Boards available for use today.

Figure 8-15 *Cisco Webex Board 85, 70s, and 55s*

Key Topic

With the original release of the Cisco Spark Board, which was later changed to the Cisco Webex Board, this series would only register to Cisco Webex in the cloud and required activation and registration to use the features available on these systems. This was due to the operating system that drove this endpoint. Since then, Cisco has changed the OS to the CE software so that the Webex Board will now register to the Cisco Unified Communications Manager via SIP or to the Cisco Expressway via SIP or H.323. You can also still register the Webex Board to the Webex Control Hub in the cloud. There are some limitations to registering the Webex Board on premises. At the time this chapter was written, the Webex Board can only be used for whiteboarding and annotation during local meetings when registered on premises. These whiteboard sessions cannot be saved, and whiteboarding and annotation are not supported during a call from the Webex Board. This capability is roadmapped for an undefined time in the future. However, you no longer have to unlock the whiteboarding functionality to use it, and these functions still work the same when registered to the Webex Control Hub. Table 8-9 illustrates some of the feature differences and commonalities between the three different models of the Cisco Webex Board.

Key Topic

Table 8-9 Cisco Webex Board Features

	Webex Board 55s	Webex Board 70s	Webex Board 85
Display	55" LED LCD 4k	70" LED LCD 4k	85" LED LCD 4k
Camera	Fixed lens Infinite focus 4kp60 resolution 83° horizontal field of view 55° vertical field of view	Fixed lens Infinite focus 4kp60 resolution 83° horizontal field of view 55° vertical field of view	Fixed lens Infinite focus 4kp60 resolution 83° horizontal field of view 55° vertical field of view
Participants	Up to 5 people	Up to 7 people	Up to 14 people
Dimensions (H × W × D)	36.2 × 55.7 × 7.5 in (919 × 1416 × 191 mm)	47.5 × 73.8 × 9.6 in (1207 × 1875 × 245 mm)	48.1 × 77.4 × 3 in (1221 × 1966 × 76 mm)
Audio Features	High-quality 20 kHz audio Acoustic echo cancellation Automatic gain control Autonoise reduction Active lip synchronization Mic array with voice tracking	High-quality 20 kHz audio Acoustic echo cancellation Automatic gain control Autonoise reduction Active lip synchronization Mic array with voice tracking	High-quality 20 kHz audio Acoustic echo cancellation Automatic gain control Autonoise reduction Active lip synchronization Mic array with voice tracking

IX Series

The Cisco Immersive Experience IX Series is designed for "Immersive Collaboration Beyond the Boardroom." The Cisco Telepresence IX Series is the industry's first H.265 triple-screen product. It requires zero room remediation and takes only half the installation time, power, and bandwidth of existing products for better total cost of ownership (TCO) and faster deployment. This system runs on a single, powerful codec, which is reduced from four codecs on the TX9000, and an intelligent software platform delivers unmatched high-fidelity audio, video, and an overall superior rich media collaboration experience. This codec is the first to handle five simultaneous streams of data using the H.264 video codec. Three media streams are primary video, and two streams are content streams. Rich multimedia presentation capabilities give users the most vivid collaboration experience with flexibility to share content on any of the three 70-inch screens. Content being shared with this system is shown on the Touch 10 controllers, of which nine controllers can be connected to the system. Video dewarping allows a whiteboard to be placed almost anywhere in the room. There are up to 18 mobile devices supported for Intelligent Proximity for Content Sharing with no degradation in bandwidth. Shared audio content automatically uses system speakers for high-fidelity sound. The system components comprise the following:

- Triple 4k camera cluster with automated alignment
- Three thin bezel 70-inch LCD displays
- One powerful single codec, allowing five simultaneous streams
- H.265 HEVC capable of supporting 1080p 60 fps
- Eighteen positioning microphones
- Nineteen special audio speakers
- Two three-headed dongles that support HDMI, Display Port, and Mini Display port connections
- An integrated LED lighting bar
- Requires only a single 10/15 amp circuit requiring 950 watts of power to run the whole system

The IX series runs on Cisco Telepresence System (CTS) software and can only register to the Cisco Unified Communications Manager via SIP. This product can be ordered as the IX5000, which comes with seating for 6 participants around the table, or it can be ordered as the IX5200, which comes with a second row of seating for an additional 12 participants bringing the total number to 18 participants. This is truly an extraordinary video-conferencing room solution. Figure 8-16 illustrates the Cisco Telepresence IX5000 endpoint product.

As you can see, Cisco has a broad range of Telepresence products from which customers can choose. Which products to choose depends entirely on the purpose the endpoints will serve and the call control product they will operate through. However, no matter which product is used, the experience will be consistent with great audio and video quality.

Figure 8-16 *Cisco Telepresence IX5000*

As mentioned before, Cisco has released a new product called the Cisco Webex Room Panorama, which is a first-class collaboration system built for the modern C-level employee. The Room Panorama provides an immersive video experience, rich content sharing, and a cocreation experience within the Cisco Webex Room Series portfolio. These fully integrated systems combine beautiful design and powerful functionality into an all-in-one solution for medium to large rooms. The Room Panorama features a powerful engine to build more complex and specialized video scenarios. It will transform your meeting spaces into a video collaboration hub, whether for connecting teams across the world or for local meetings. The Room Panorama will be shipped as a complete technology package from Cisco. This includes everything you need to make it operational, including the following components:

- Two Samsung 82-inch 8K QLED displays for video and one Samsung 65-inch 4k display for presentations
- Codec and four 5k cameras at 60 fps
- Touch 10 controller
- Twelve-loudspeaker system with directional audio capabilities
- Four bass modules
- Cooling system
- Audio amplifier
- Complete power system
- Three table microphones
- Presentation cables
- Wall structure (Customers can design their own wall finish to match an existing room, or they can buy the Cisco wall finish in light oak as an option.)

The Cisco Webex Room Panorama runs on the same CE software that the other Webex endpoints use, allowing this executive room system to register via SIP to Webex in the cloud, or on premises to the Cisco Unified Communications Manager or Expressway. It also supports H.323 for registration to the Expressway. The powerful, integrated cameras deliver intelligent view capabilities, such as panorama video, automatic framing, and speaker tracking. Automatic noise suppression reduces meeting disruptions. Room Panorama supports up to three screens, dual-content sources, wireless sharing, and 4k content for great presentations. The people count feature offers usage metrics and resource allocation. Tight integrations with screens enhance user interactions, and APIs and macros allow for meeting personalization. With the enhanced capabilities and extensive registration options, there is no wonder that this solution is the next generation to replace the IX5000.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 8-10 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 8-10 Key Topics for Chapter 8

Key Topic Element	Description	Page Number
Paragraph	xAPI and WebSockets	166
Paragraph	Audio Console with Graphical Equalizer	167
Paragraph	Intelligent Proximity for Content Sharing	167
Paragraph	DX80 Endpoint Overview	169
Table 8-2	DX70 and DX80 Feature Differences	170
Paragraph	SX10 Cabling	171
Paragraph	SX20 Camera Options	171
Paragraph	SX80 Capabilities	172
Table 8-3	SX Endpoint Feature Differences	174
List	Mounting Options for MX200 and MX300	174
Paragraph	MX700 and MX800 Displays	175
Table 8-4	Cisco Telepresence MX Series Feature Differences	176
Paragraph	Room Kit Mini AI Capabilities	178
Table 8-5	Webex Room Kit Mini Features	178
Paragraph	Webex Room Kit Components	179

Key Topic Element	Description	Page Number
Table 8-6	Webex Room Kit Features	180
Paragraph	Room Kit Plus Components	180
Table 8-7	Webex Room Kit Plus Features	181
Paragraph	Webex Room Kit Pro Components	182
List	Purchase Options for Webex Room Kit Pro	182
Table 8-8	Webex Room Kit Pro Features	183
Paragraph	Webex Room 55 Components	183
Paragraph	Webex Room 70 Components	184
Paragraph	Webex Board Registration	186
Table 8-9	Cisco Webex Board Features	186
List	IX5000 Components	187

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

3G-SDI, AI, Autoframing, BFCP, CE, DND, DX, Euroblock, H.239, HD-SDI, Inductive Loop, Intelligent Proximity for Content Sharing, JSON, MIMO, MWI, MX, OBTP, OS, OSD, PoE, Precision Camera, PrecisionHD Camera, Precision MIC 20, SNR, Speaker Track 60, SX, Touch 10, TRC, UltraHD, VISCA, Webex Meeting App, Webex Teams App, WebSocket, xAPI

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List all of the endpoints in the DX, SX, and MX product lines.
2. List the mounting options for the MX300 G2 endpoint.
3. List all of the endpoints in the Webex endpoint product line.
4. What are the three purchase options for the Cisco Webex Room Kit Pro?
5. List all the system components that make up the IX5000 endpoint room system.

This page intentionally left blank

CHAPTER 9

Endpoint Registration

This chapter covers the following topics:

SIP Registration to the Cisco Unified Communications Manager: This topic will explain how to configure settings on UC phones and Telepresence endpoints for SIP registration to the Cisco Unified Communications Manager.

SIP Registration to the Cisco Expressway Core: This topic will explain how to configure settings on Telepresence endpoints for SIP registration to the Cisco Expressway Core.

H.323 Registration to the Expressway Core: This topic will explain how to configure settings on Telepresence endpoints for H.323 registration to the Cisco Expressway Core.

For any old-school voice engineers who might be reading this book, and to the general audience of people who may or may not understand this simple truth, the Cisco Unified Communications approach to provisioning and controlling phones has been to keep the intelligence in the call control systems, such as the Cisco Unified Communications Manager, and keep the phones dumb. This is a practice taken from tradition telephony over analog or digital circuit-switched systems and continues to be practiced cross-vendor in the IP telephony world. Cisco has added some intelligence to its UC phones, such as with Dial Rules, but on the norm, it continues to keep the intelligence at the hub of the call control systems. This idea has changed with Telepresence endpoints. There is a shared intelligence between the endpoint and the call control system, both in function and capabilities. What you will discover throughout this chapter is the much more advanced configuration options that exist on Telepresence endpoints versus their counterparts, UC phones. Topics discussed in this chapter include the following:

- SIP Registration to the Cisco Unified Communications Manager
  - PoE
  - CDP and LLDP-MED
  - DHCP
  - TFTP
  - SIP Registration
  - ITL, CTL, and CAPF
- SIP Registration to the Expressway Core
  - DHCP versus Static IP
  - Manual Configuration of SIP Settings

- H.323 Registration to the Expressway Core
  - H.323 Aliases
  - Manual Configuration of H.323 Settings

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.1.g Security (certificates, SRTP, TLS)
- 1.3.g Certificates
- 2.4 Deploy SIP endpoints

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 9-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 9-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
SIP Registration to the Cisco Unified Communications Manager	1–6
SIP Registration to the Expressway Core	7–8
H.323 Registration to the Expressway Core	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. How much power can be supplied to a phone using Cisco prestandard PoE?
   - **a.** 5.5W
   - **b.** 7W
   - **c.** 14.5W
   - **d.** 15W

2. When you are configuring a VLAN on a switchport, which of the following commands will configure the phone to use IEEE priority tagging for voice traffic and to use the default native VLAN to carry all traffic?
   a. 802.11n
   b. 802.1q
   c. 802.1p
   d. 802.3af
3. When the CUCM is used as the DHCP server, how many devices can it support?
   a. 100
   b. 500
   c. 1000
   d. 10,000
4. The Maximum Serving Count service parameter specifies the maximum number of requests that can be concurrently handled by the TFTP server on the CUCM. What is the default value of this parameter?
   a. 2500
   b. 5000
   c. 7500
   d. 10,000
5. Which SIP registration header field contains the address of record of the person responsible for the registration?
   a. Request URI
   b. To
   c. From
   d. Call-ID
6. Which of the following are considered Security By Default mechanisms on the CUCM? (Choose two.)
   a. ITL
   b. CTL
   c. CAPF
   d. TLS
   e. TLS Verify
   f. TVS
7. Which of the following is a required setting when configuring static network information on an endpoint?
   a. Gatekeeper address
   b. DNS
   c. TFTP
   d. Default gateway address

8. Which of the following commands is required before the SIP endpoint can register to the Cisco Expressway when configuring the CE endpoint using CLI?
   a. **xConfiguration NetworkServices SIP Mode: On**
   b. **xConfiguration Provisioning Mode: VCS**
   c. **xConfiguration Provisioning Mode: TMS**
   d. **xConfiguration SIP DisplayName:** *name*
9. Which of the following aliases is acceptable as an H.323ID?
   a. 555 - 1212
   b. John Hancock
   c. John 1212
   d. John@Hancock
10. An administrator is using the CLI to configure a DX80 endpoint to register to a Cisco Expressway Core. While configuring the alias, the administrator configures an H.323 ID when it should not have been. How can the administrator remove the H.323 ID using the CLI?
    a. **xconfiguration h323 h323alias remove all**
    b. **xconfiguration h323 h323alias id:** *leave blank*
    c. **xconfiguration h323 h323alias id: delete**
    d. **xconfiguration h323 h323alias id: “”**

Foundation Topics

SIP Registration to the Cisco Unified Communications Manager

Chapter 5, “Communication Protocols,” provided a high-level overview of the SIP registration process to the Cisco Unified Communications Manager. The basic elements of the whole SIP registration process can be summarized in the following order:

1. The endpoint obtains power.
2. The endpoint loads the locally stored image file.
3. The endpoint communicates with the switch using CDP or LLDP-MED.
4. The endpoint negotiates DHCP with the DHCP server.
5. The endpoint is issued CTL (optional) and ITL certificates from the Cisco Unified Communications Manager.
6. The endpoint communicates with the TFTP server.
7. The endpoint registers to the CUCM.

Although many moving parts must be configured for this registration process to work, after each component is configured, very little must be done on the end user’s part for devices to register to the Cisco Unified Communications Manager. This process is designed to ease the registration process and also allow engineers to mass-deploy tens of thousands of endpoints at the same time. When the components involved with this process are deployed properly, little to no configuration needs to be set up on the endpoint itself. You can just power the

system on and let these services provision all the various settings on the phones. The following sections describe how to configure these various components so that this provisioning process will operate correctly.

PoE

An endpoint can receive power in two ways: Power over Ethernet (PoE) or through a power cube, which is the traditional cable you plug into the system and the wall outlet. All Telepresence endpoints require a power cube to be used, with the one exception of the Cisco Telepresence SX10. All of the Cisco UC phones can support either a power cube or PoE. Because PoE is the expected form of power, none of the UC phones come standard with a power cube; this component must be ordered separately. A big difference between some of the phones as related to PoE is the type or class of PoE supported.

PoE, also referred to as *inline power*, is the capability for the LAN switching infrastructure to provide power over a copper Ethernet cable to an endpoint or powered device. Cisco developed and first delivered this capability in 2000 to support the emerging IP telephony deployments. IP telephones, such as desktop PBX phones, need power for their operation, and PoE enables scalable and manageable power delivery and simplifies deployments of IP telephony. Because these early PoE-capable devices were basic phones without a lot of features added to them, the power requirements were pretty low. Cisco's prestandard PoE, which was called inline power, supported only 7 watts (7W) of power. The IEEE quickly recognized the contribution Cisco made to the IT industry and began working on a standardization for PoE. The first standard for PoE that could be used industrywide and applicationwide was 802.3af. With the advent of 802.3af came many more devices that could support PoE, including wireless access points, video cameras, point-of-sale devices, security access control devices such as card scanners, building automation, and industry automation, to name a few.

The IEEE 802.3 standard outlines two types of devices: power sourcing equipment (PSE) and powered devices (PDs). Power sourcing equipment provides power to the powered devices. A PSE can support power delivery Type A, Type B, or both. Type A involves sending power over two unused pairs or wires on a CAT3, CAT5, or CAT5E cable. This works well for links up to 100 Mbps, but Gigabit Ethernet uses all the copper pairs in a CAT5E cable. Therefore, Type B uses a "phantom power" technique to send power and data over the same pairs. When an Ethernet device is connected to a PSE, the PSE initially applies a low voltage (2–10 volts) to sense whether the device is a PoE PD. If it is, the PD will send a return current back to the PSE, and 48 volts will be supplied so the device can power on and load its locally stored image file. If no return current is sent back to the PSE, the device connected is not a PD, such as a computer connected over Ethernet, and no power will be supplied. The maximum power that a PSE will supply down an Ethernet cable using 802.3af is 15 watts; however, due to possible losses on a long cable, the maximum power that a PD can receive down an Ethernet cable using 802.3af is 12.95 watts.

The IEEE 802.3at standard, also known as PoE+, supports up to 25.5W of power on the ports, allowing devices that require more than 15.4W to power on when connected to the PoE+ ports. Several Cisco switches support 802.3at PoE, including the Meraki MS series switches. Based on the classification currently used by the device, the Meraki MS switch will classify the device as a Class 0, 1, 2, 3, or 4 type device and apply the proper standards-defined behaviors to the port. Table 9-2 describes these five classifications on Meraki switches.

Key Topic

Table 9-2 Meraki Switch PoE Classifications

Class	Usage	Classification Current [mA]	Power Range [watt]	Class Description
0	Default	0–4	0.44–12.94	Classification unimplemented
1	Optional	9–12	0.44–3.84	Very low power
2	Optional	17–20	3.84–6.49	Low power
3	Optional	26–30	6.49–12.95	Mid power
4	Valid for 802.3at (Type 2) devices, not allowed for 802.3af devices	36–44	12.95–25.50	High power

When a PD is connected to an 802.3at switch port, a lower power voltage can be supplied because 802.at is backward compatible. There is a limit to how much power may be drawn across wires before electrical damage occurs due to overheating within connectors and cable bundles. There is also a concern over signaling interference with this protocol. Some of these issues are resolved by using multiple pairs to deliver the necessary power. At the moment, 802.3at limits the number of pairs that can carry power to two. A current limit of 720mA is being considered, allowing 29.5W per pair; however, the IEEE is working on Draft 3.0 to reduce this to 600mA giving 25W per pair, or 50W per device. The IEEE is also looking at mandating Category 5 cables and later to be used with 802.3at so that you do not have to worry about supporting Category 3 cabling. With 802.3at the maximum power that can be delivered to the application is 50W. The first detection pulse, or *classification pulse*, from the PSE will be the same as 802.3af, to which the 802.3af PDs will respond normally. A second PoE Plus pulse then will cause an 802.3at PD to respond as a Class 4 device and draw the Class 4 current. After this has happened, there will be a data exchange where information such as duty-cycle, peak, and average power needs will be shared. Other features to be catered for in 802.3at include dynamic power assignment, leading to more efficient power supply designs and consequent power saving.

For a Cisco UC phone to receive PoE from a switch, it is essential to connect the patch cable coming from the switch into the appropriate port on the phone. You should be aware of three switch ports in Cisco phones. There are two physical switch ports on the back of each phone. One of the switch ports is called a *network port*, and this is the port to which the patch cable that connects back to the switch should be connected. The second physical port on the back of a Cisco phone is called the *PC port*, and it is the port that a computer can be connected to in order to receive network access. PoE will not be supplied to the phone if the network cable from the switch is connected to the PC port, nor will the phone be able to communicate with the switch. These two physical interfaces are easy to recognize because there is a graphic below each port to describe that port's purpose. The computer port has the graphic of a computer monitor, and the network port graphic displays three squares interconnected with lines, signifying the network. The third port on a Cisco phone that you should be aware of is a virtual port located in the phone's software, which bridges between the computer port and the network port. The phone uses this port to control how packets

are marked for the purpose of quality of service (QoS) before they are sent to the switch. PoE promises to create a new world of networked appliances as it provides power and data connectivity over existing Ethernet cables. Table 9-3 identifies the three PoE types discussed in this section. Bear in mind that Cisco has not created or sold prestandard PoE devices since the late 2000s. The information on prestandard PoE is for comparison purposes only.

Key Topic

Table 9-3 Three PoE Types Supported on Cisco Switches

Prestandard Inline PoE	802.3af PoE	802.3at PoE+
Cisco Proprietary	IEEE standard	Backward compatible with 802.3af; PoE+ just adds an additional class of power to the 802.3af standard
10/100 only	15.4W per port	30W per port
7W per port	Compatible with Gigabit Ethernet	Relatively new; currently only Cisco is shipping PoE+ phones
Incompatible with all non-Cisco devices that accept Power over Ethernet	PoE devices are not compatible with Cisco prestandard PoE; the power negotiation process is completely different	
	Cisco PoE switches are backward compatible with prestandard PoE	
	Enough power for most IP phones and wireless access points from all manufacturers	

No settings need to be configured on a Cisco phone to enable PoE. These settings are hard-coded into the phone itself. Most Cisco PoE-capable switches come preconfigured to support PoE as well; however, some settings on a Cisco switch can be configured to disable or enable PoE on certain ports and even boost the power capabilities on specific ports. The following examples illustrate how to configure PoE on a Cisco Catalyst 4500 series switch. Depending on the model number you are configuring, the commands you use might vary slightly. You can enter the following commands into a Cisco switch to enter configuration mode and configure PoE on a specific switchport:

Key Topic

```
Switch# configure terminal

Switch(config)# interface {fastethernet|gigabitethernet}
(slot/port)

Switch(config-if)# power inline {auto[max milli-watts] | never |
static [max milli-watts]}

Switch(config-if)# end

Switch# show power inline {fastethernet|gigabitethernet} slot/port
```

An example of the preceding configuration could be as follows:

```
Switch# configure terminal

Switch(config)# interface fastethernet 0/1

Switch(config-if)# power inline auto 15.4

Switch(config-if)# end

Switch# show power inline fastethernet 0/1
```

As you can see from the preceding example, you can configure three settings on a switchport for PoE: auto, static, and never. The auto setting is the default value, and this setting allows the supervisor engine on the switch to direct the switching module to power up the interface only if the switching module discovers the phone and the switch has enough power. This mode has no effect if the interface is not capable of providing PoE. The static setting is recommended on high-priority PoE interfaces. The supervisor engine will pre-allocate power to the interface, even when nothing is connected, guaranteeing that there will be power for the interface. If the switch does not have enough power for the allocation, the command will fail. The supervisor engine directs the switching module to power up the interface only if the switching module discovers the powered device. Both of these modes allow you to specify the maximum wattage that is allowed on the interface. If you do not specify a wattage, the switch will deliver no more than the PSE hardware-supported maximum value. The never setting will disable PoE on that particular switchport. This setting is typically used when the interface is intended to be used only as a data interface. The supervisor engine never powers up the interface, even if an unpowered phone is connected. This mode is needed only when you want to make sure power is never applied to a PoE-capable interface. The switch can measure the actual PoE consumption for an 802.3af-compliant PoE module and displays this in the **show power inline** *module* command from the privileged EXEC mode.

One other command worth knowing can boost the power on a specific PoE port if not enough power is available: **power inline delay shutdown**. You might use this command when a prestandard PoE switch is being used to try powering up an 802.3af phone. You also can use this command on an 802.3af switch trying to power up an 802.3at device. However, it is important to understand that when this command is used, you are essentially "borrowing from Peter to pay Paul," as the saying goes. Power is being taken from another switchport to supply extra power to that particular port being configured. If you use this command, you will not be able to support as many phones on that switch as it was originally designed to support. The command you can use from the configuration mode prompt is as follows:

```
Switch(config-if)# power inline delay shutdown 20 initial 300
```

The **power inline delay shutdown** command configures the port to delay shutting down. This command is useful when a phone requesting more power than the port is originally designed to support would normally go into a cyclical reboot. The initial time period—in this example, 20 seconds—begins when the IEEE-compliant powered device is detected by the switch. If link-down occurs on the connected device during the initial time period, the shutdown time, **initial 300**, determines how long the switch continues to provide power to the device.

CDP and LLDP-MED

The next step in the registration process to the Cisco Unified Communications Manager, after the phone has powered on and loaded the locally stored image file, is to communicate with the switch. Two protocols can be used for this type of communication: Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED).

CDP is a proprietary protocol that will work only when a Cisco device is communicating to a Cisco device, such as a Cisco phone communicating with a Cisco switch.

If either one of the devices, the phone or the switch, is not a Cisco product, the LLDP-MED protocol will need to be used instead. Several LLDP protocols are available, but LLDP-MED is specifically designed for voice and video communication. Now that it has been established that communication must occur between the phone and the switch, exactly what information is being exchanged is equally important to understand.

A PSE has a PoE "power budget," which is simply the total power the switch can supply down Ethernet cables. Each time a PD is plugged into the PSE, the PSE subtracts the PD's maximum power usage from its power budget. If the power budget is all used up, and another PD is connected, that PD will not receive any power. To avoid this problem, 802.3af specifies a method whereby a PD can indicate to the PSE what its maximum power usage will be, by indicating that it complies with one of four "power classes" specified in the standard. For example, if a PD indicates that it complies with power class 2, the PSE knows that it only has to subtract 3.84 watts from its power budget for that PD, thus leaving more power budget for other PDs. In this manner, after the phone powers on, the phone is capable of sending to the switch the required amount of power needed to sustain its core systems. CDP, or LLDP-MED, is the mechanism used to communicate these power adjustments to the switch.

Before the phone obtains its IP address, the phone must also determine which VLAN it should be in by means of the CDP communication that takes place between the phone and the switch. This communication allows the phone to send packets with 802.1Q tags to the switch in a "voice VLAN" so that the voice data and all other data coming from the PC behind the phone are separated from each other at Layer 2. Voice VLANs are not required for the phones to operate, but they provide additional separation from other data on the network. Voice VLANs can be assigned automatically from the switch to the phone, thus allowing for Layer 2 and Layer 3 separations between voice data and all other data on a network. A voice VLAN also allows for a different IP addressing scheme because the separate VLAN can have a separate IP scope at the Dynamic Host Configuration Protocol (DHCP) server. Applications use CDP messaging from the phones to assist in locating phones during an emergency call. The location of the phone will be much more difficult to determine if CDP is not enabled on the access port to which that phone is attached. There is a possibility that information could be gathered from the CDP messaging that would normally go to the phone, and that information could be used to discover some of the network. As mentioned before, not all devices that can be used for voice or video with Cisco Unified Communications Manager are able to use CDP to assist in discovering the voice VLAN. Third-party endpoints do not support CDP. To allow device discovery when third-party devices are involved, you can use LLDP-MED. This protocol defines how a switch port transitions from LLDP to LLDP-MED if it detects an LLDP-MED-capable endpoint. Support for both LLDP and LLDP-MED on IP phones and LAN switches depends on the firmware and device models.

To determine if LLDP-MED is supported on a particular phone or switch model, you can check the specific product documentation, release notes, and whitepapers.

VLANS can be configured on specific switchports as a Voice VLAN ID (VVID) or as a data VLAN, or both as a voice and data VLAN. Video will typically use the VVID along with audio. It is not necessary to create a separate voice and video VLAN. Because VLANs will be used to communicate the Layer 2 class of service (CoS) to the phone, it is important to also set up CoS tagging on the switch for specific VLANs. CoS and QoS will be discussed in more depth in Chapter 13, "Layer 2 and Layer 3 QoS Parameters." The following commands outline how to configure VLANs on the switchports and enable CoS. There are several ways this can be done, and the commands may vary depending on the switch being used. These examples are based on current Cisco Catalyst Switch series.


```
Switch# configure terminal

Switch(config)# vlan number

Switch(config-vlan)# name name

Switch(config-vlan)# exit

Switch(config)# interface {fastethernet|gigabitethernet}
(slot/port)

Switch(config-if)# mls qos trust cos

Switch(config-if)# switchport voice {detect cisco-phone
[full-duplex] | vlan {vlan-id | dot1p | none | untagged}}

Switch(config-if)# end

Switch# show vlan
```

Issuing the **configure terminal** command allows you to enter the configuration mode.

Next, you can create a VLAN using the **vlan** *number* command and assign a **name** to the VLAN for a description.

After exiting VLAN configuration mode, you issue the **interface {fastethernet| gigabitethernet} (***slot/port***)** command to enter the interface configuration mode. Next, you issue the **mls qos trust cos** command so that this switchport will trust the CoS-to-QoS mapping embedded in the switch. Then you enter the **switchport voice {detect cisco-phone [full-duplex] | vlan {***vlan-id* **| dot1p | none | untagged}}** command. The **detect** part of the command will configure the interface to detect and recognize a Cisco IP phone. The **cisco-phone** is the only option allowed when you initially implement the **switchport voice detect** command. The default is **no switchport voice detect cisco-phone [full-duplex]**. The **full-duplex** command is optional. You can configure the switch to accept only a full-duplex Cisco IP phone. This setting is highly recommended because you do not want voice or video traffic to use half-duplex. Calls will consume twice the bandwidth, and connections will be spotty at best. The **vlan** *vlan-id* command will configure the phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP phone forwards the voice traffic with an IEEE 802.1Q priority of 5. Valid VLAN IDs are 1 to 4094. The **dot1p** command will configure the phone to use IEEE 802.1p priority tagging for voice traffic and to use the default native VLAN (VLAN 0) to carry all traffic. Alternatively, the

none command will allow the phone to use its own configuration to send untagged voice traffic. Another alternative is the **untagged** command, which will configure the phone to send untagged voice traffic. Using **none** or **untagged** is not recommended for voice traffic.

You type **end** to exit configuration mode and then issue the **show vlan** command to verify the VLAN configuration. The following example shows how these commands should appear:

```
SW1(config)# vlan 200

SW1(config-vlan)# name VVID

SW1(config-vlan)# exit

SW1(config)# Interface fastethernet 0/1

SW1(config-if)# mls qos trust cos

SW1(config-if)# switchport voice vlan 200

SW1(config-if)# end
```

As an alternative to the **switchport voice vlan 200** command, you could use **switchport voice detect cisco-phone full-duplex** or **switchport voice vlan 802.1p**. The 802.1p will tag voice traffic with a CoS priority of 5 but use the default VLAN to send traffic. The **detect** command will send untagged traffic to the switch. Neither one of these options is ideal in a voice or video deployment, which is why the example uses the **voice vlan** command. In the preceding example, only the voice VLAN was applied to the switchport. Both the voice VLAN and the data VLAN can be applied to the switchport as another option for VLAN configuration. The phone will then use these VLANs to tag all communication originating from the phone with the voice VLAN, and all information originating from a connected computer with the data VLAN. The following example demonstrates how this might look on a switchport:

```
SW1(config)# vlan 200

SW1(config-vlan)# name VVID

SW1(config-vlan)# exit

SW1(config)# vlan 100

SW1(config-vlan)# name Data

SW1(config-vlan)# exit

SW1(config)# Interface fastethernet 0/1

SW1(config-if)# switchport mode access

SW1(config-if)# switchport access vlan 100

SW1(config-if)# switchport voice vlan 200

SW1(config-if)# end
```

Key Topic

A computer connected to the phone will normally send packets as untagged. The command **switchport mode access** will force the port into access mode. The phone will tag computer traffic as data and will tag phone traffic as voice (or video when applicable). Then the command **switchport access vlan 100** is used to identify that vlan100 should be used for access tagging of all data traffic originating from the computer. The command **switchport voice vlan 200** is used to identify that vlan200 should be used for tagging all voice and video traffic originating from the phone. If CoS Priority 5 should be used for Layer 2 QoS, you will still need to add the command **mls qos trust cos**, as mentioned previously. When third-party phones are being used, two VLANs on a single switchport may not work for appropriate tagging. LLDP-MED and most other vendor phones do not support this feature.

VLAN discovery is automatic on all Cisco IP phones and Telepresence endpoints, so no settings have to be configured on these devices for VLAN tagging to occur. However, some circumstances may prevent the phone or endpoint from discovering the VLAN. Therefore, you can manually configure a VLAN on a phone or endpoint. The following information on how to configure VLAN settings manually on a Cisco phone is applicable to the 7800 and 8800 series phones. These settings also are available on other Cisco phones, but the menu options to get to these settings may be different.

Key Topic

Step 1. Press the **Settings** button on your phone. This is the button with the gear icon.

Step 2. Use the circular navigation button on the phone to choose the **Admin Settings** menu option, or select the number associated with the menu option using the numeric keypad on the phone.

Step 3. Choose the **Network Setup** menu option in the same manner as above.

Step 4. In the next set of menus that appear, choose **Ethernet Setup**. There are several menu options under Ethernet Setup.

Step 5. Scroll down to the Operational VLAN ID field. This field is not configurable from the user interface but will display the Voice VLAN ID if one was provided from the switch.

Step 6. If the Operational VLAN ID field is blank, you will need to configure the Admin VLAN ID field with the appropriate VLAN ID.

Step 7. The PC VLAN field is another field that is not configurable from the user interface, unless the Admin VLAN ID was configured first. It will display the data VLAN ID if one was provided from the user interface. Otherwise, you will need to configure one.

Step 8. Press the **Apply** softkey to save the configured settings. Figure 9-1 illustrates these menu options under the Ethernet Configuration menu.

VLAN settings can also be configured on a Cisco CE software-based endpoint. It is important to note here that CE software-based endpoints have auto VLAN discovery enabled by default, but TC software-based endpoints do not. When CE endpoints are registering to a Cisco Unified Communications Manager, no settings have to be preconfigured on the endpoint for that endpoint to register. However, with legacy TC endpoints, you may have to preconfigure some settings on the endpoint before the endpoint can register to the Cisco

Unified Communications Manager. The instructions that follow outline how to configure VLAN settings on a Telepresence endpoint and are based on CE software. Some menus may be different if you are configuring a Telepresence endpoint running the TC software. Also note that if the **NetworkServices > CDP Mode** is set to **Off**, VLAN discovery will not work, even though Auto Discovery may be enabled.

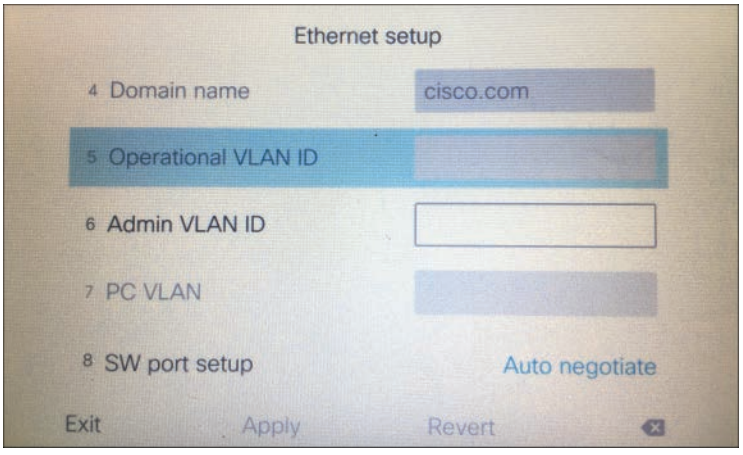

Figure 9-1 *Ethernet Configuration Menu Options on Cisco 7800 and 8800 Series Phones*

Key Topic

Step 1. Log in to the web interface for the CE endpoint you want to configure.

Step 2. Navigate to **Setup > Configuration > NetworkServices** and verify the CDP Mode is set to **On**. This setting should be on by default. Notice that LLDP-MED is not supported on CE software-based endpoints.

Step 3. In the left column, click the **Network** menu and scroll to the bottom of the page. Under the VLAN Voice section, you can set the VLAN Mode to **Auto**, **Manual**, or **Off**.

Step 4. Below the Mode setting is the VlanId setting. This will default to the value 1, and can be set between 1 and 4094. This field does not apply if the Mode is set to Auto. If you change the Mode to Manual, you will also need to enter the VlanId value that is appropriate to the voice VLAN configured on the switch.

Figure 9-2 illustrates how to configure CDP and VLAN settings on a Cisco CE software-based endpoint.

DHCP

After a phone has adjusted the power consumption from the switch and requested the VLAN information, the next step in the process is to establish appropriate network information through the DHCP process. The minimum amount of network information any device must possess in order to communicate across a network is an IP address, subnet mask, and default router address, also called a *default gateway address*. Additional network information that can be obtained using DHCP include Domain Name System (DNS) addresses and Trivial File Transfer Protocol (TFTP) server addresses. The TFTP server address can be obtained by configuring Option 66 or Option 150. Option 66 is an open

standard that will operate on any vendor's DHCP server. Option 150 is a Cisco proprietary protocol that will operate only on Cisco's and a few other vendors' DHCP servers. Additionally, Option 66 will support only one TFTP address, whereas Option 150 will support up to two addresses, making it a more redundant option to Option 66. There are many pros and cons to using DNS, but one great advantage in this context is that the TFTP server address can be configured as the URL of the Cisco Unified Communications Manager cluster so that more than one address can be delivered to the device trying to register.

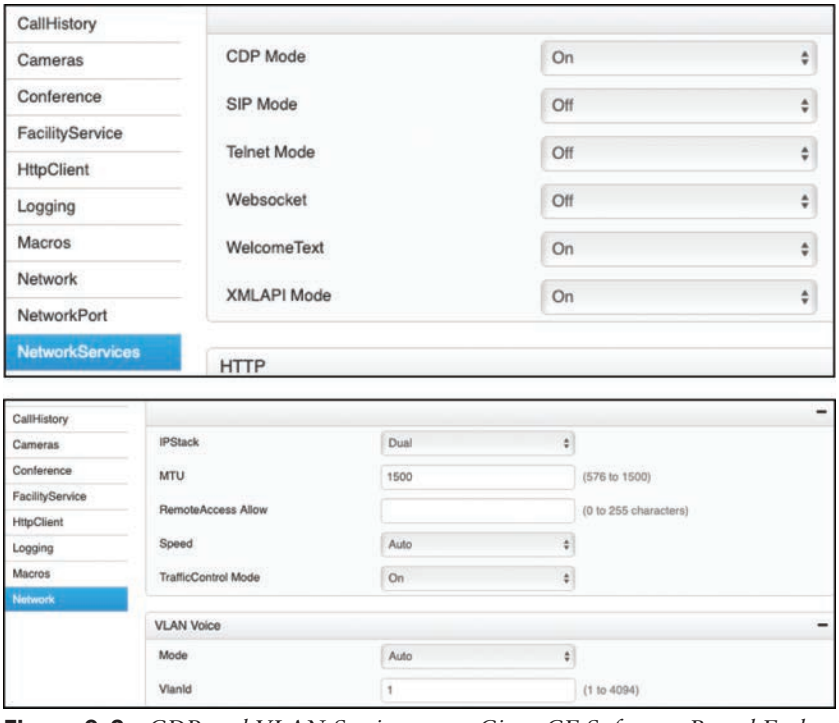

Figure 9-2 *CDP and VLAN Settings on a Cisco CE Software-Based Endpoint*

Many types of DHCP services can be used. The Cisco Unified Communications Manager has a built-in DHCP service, but there are limitations to using this option. It is recommended to use the DHCP service only for phones and Telepresence endpoints, but not for computers or other such devices. There is a capacity limit of 1000 devices that can receive DHCP information, and it is more complex to set up the DHCP service on the CUCM as opposed to another device because a helper address needs to be configured and certain devices will need to be restricted from using this service. A commonly used DHCP service is on the router itself. There, a greater pool of addresses can be set up, it can be used for all devices communicating on the network, and it is relatively easy to set up. The following example demonstrates how to set up DHCP on a router to support IP addressing, along with DNS address and TFTP address distribution. You should understand that this is not the only way to configure DHCP on a router, but one possible way it can be configured.

Key Topic

```
Router# configure terminal
Router(config)# ip dhcp pool caret&stic
```

```
Router(dhcp-config)# network 10.0.0.0 255.255.255.0
Router(dhcp-config)# default-router 10.0.0.1
Router(dhcp-config)# option 150 ip cucm-cluster.caret&STIC.com
Router(dhcp-config)# dns-server 8.8.4.4
Router(dhcp-config)# domain-name caret&STIC.com
Router(dhcp-config)# lease 7
Router(dhcp-config)# exit
```

The **ip dhcp pool** *name* command creates a pool from which IP addresses can be issued to devices that send a DHCP request. The *name* field can be any name you want to give to the pool. In the preceding example, **caret&STIC** is a fictitious company name that has also been assigned to the DHCP pool. The next command, **network** *IP address subnet mask*, establishes all the available addresses within a pool that can be used for DHCP assignment. This is where you will need to be careful. Servers will often be assigned static IP addresses, which should not be included in the pool. You can either configure a smaller subnet of addresses here, or you can go back and issue an exclusion range of addresses that will not be used in DHCP assignments. An example of how to issue a range of excluded addresses is as follows:

```
Router(config)# ip dhcp excluded-address 10.0.0.1 10.0.0.99
```

In this example, all the IP addresses from 10.0.0.1 to 10.0.0.99 are excluded from the DHCP pool. The first IP address that will be provided to a device that sends a DHCP request will be 10.0.0.100. The **default-router** *IP address* command establishes the default gateway address that will be assigned to devices. Devices trying to route data to another part of the network will send that data to this router address, and the default router will forward that traffic toward the intended destination across the network. The **option 150 ip** *TFTP server address* command will assign the TFTP server address to the endpoint. If DNS is used, the TFTP server address could be the URL address of the destination TFTP server. Otherwise, the IP address of the TFTP server can also be used. You could also use the **option 66 ip** *address* command here, but Cisco recommends using Option 150 when Cisco routers and phones are being leveraged. The **dns-server** *IP address* command allows up to three DNS server addresses to be listed. All addresses will be sent to devices in the order listed here, and they must be in the form of an IP address. To enter more than one DNS address, you must use a space between the different addresses. The **domain-name** *name* command will assign the domain to devices through DHCP. This allows for easier searching within a domain, such as registering Jabber to the IM and Presence server. When a user signs into Jabber with *user_name@domain*, the Jabber client will search for a *cup_login* SRV record associated with the *domain*. When DNS returns the address of the IM and Presence server, the Jabber client will try to register with the *user_name* part of the URI. Jabber will be covered in more depth in Chapter 26, "Users and Cisco Jabber Soft Clients." Finally, the **lease** *n* command determines how long a leased address can be used by a device before a new lease has to be requested. By default, the duration of a lease is one day. When you enter the **lease** *n* command, this duration will be extended to that number *n* of days. Actually, three values can be entered here to extend the duration to days, hours, and minutes. For example, if you wanted leased addresses to be available

for 8 hours and 30 minutes, you could enter the command **lease 0 8 30**. This establishes DHCP leases to 0 days, 8 hours, and 30 minutes.

DHCP can be enabled or disabled on Cisco UC phones, but this setting is enabled by default. If the setting is disabled, static IP settings must be configured manually on the Cisco phone. The following steps outline how to disable DHCP and configure static IP settings on Cisco 7800 and 8800 series phones. If you are configuring a different phone model, the menus to configure these settings may be different.

Key Topic

Step 1. Press the **Settings** button on your phone. This is the button with the gear icon.

Step 2. Use the circular navigation button on the phone to choose the **Admin Settings** menu option, or select the number associated with the menu option using the numeric keypad on the phone.

Step 3. Choose the **Network Setup** menu option in the same manner as above.

Step 4. In the next set of menus that appear, choose **Ethernet Setup**. There are several menu options under Ethernet Setup.

Step 5. Choose **IPv4 Setup** to enter into the IPv4 menu options.

Step 6. The DHCP field is set to On by default. Change this field to **Off** to manually configure the IP information.

Step 7. Enter the IP Address, Subnet Mask, and Default Router information in the appropriate fields.

Step 8. The DNS Server 1, DNS Server 2, and DNS Server 3 fields allow you to enter up to three DNS addresses.

Step 9. The Alternate TFTP field indicates whether the phone is using an alternate TFTP server address. If this setting is set to Off, which is the default value, the phone is expecting the TFTP server address to come from Option 150 (or Option 66). If you change the value to **On**, you must manually enter a TFTP server address.

Step 10. The TFTP Server 1 and TFTP Server 2 fields allow you to enter a primary (1) and backup (2) TFTP server address.

Step 11. Click the **Apply** soft key to save these changed settings. The phone must be reset for the new network information to bind to the phone.

Figure 9-3 illustrates how to configure these network options on a Cisco 8865 phone.

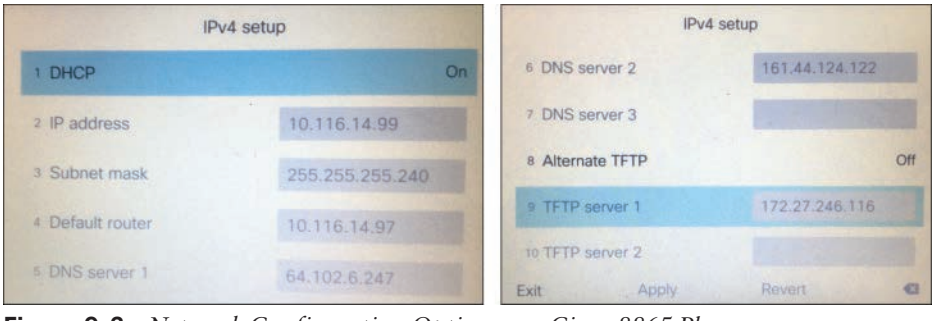

Figure 9-3 *Network Configuration Option on a Cisco 8865 Phone*

DHCP can be enabled or disabled on Cisco CE software-based endpoints, but this setting is enabled by default. If the setting is disabled, then static IP settings must be configured manually on the CE endpoint. The steps that follow describe how to disable DHCP and configure static IP settings on Cisco Telepresence CE software-based endpoints. If your endpoint is running the legacy TC software, the menus to configure these settings might be different.

Key Topic

Step 1. Log in to the web interface for the CE endpoint you want to configure.

Step 2. Navigate to **Setup > Configuration > Network** and locate the DNS section (second section down from the top).

Step 3. Enter the Domain in the Domain Name field. Enter up to three DNS addresses in the Server 1 Address, Server 2 Address, and Server 3 Address fields.

Step 4. Choose **Save** in the bottom-right corner of this section, and then scroll down to the IPv4 section.

Step 5. Change the Assignment field from DHCP to **Static**.

Step 6. Enter an IP address for the endpoint in the Address field. Enter the Default Router address in the Gateway field and enter the Subnet Mask address in the SubnetMask field.

Step 7. Choose **Save** in the bottom-right corner of this section. You will have to restart the endpoint before these changes will take effect.

Figure 9-4 illustrates the network setting menus on a CE software-based endpoint.

Figure 9-4 *Network Setting Menus on a CE Software-Based Endpoint*

Step 8. To configure the TFTP server address, choose the **Provisioning** menu in the left column.

Step 9. In the top section, ensure the Connectivity field is set to Auto, and then change the Mode to **CUCM**.

Step 10. Choose **Save** in the bottom-right corner of this section.

Step 11. In the ExternalManager section, enter the address of the TFTP server in the Address field. You can enter a second AlternateAddress as well, or just use a URL address for the Cisco TFTP Server cluster in the first Address field.

Step 12. Ensure the Protocol field is set to **HTTP**, and then choose the **Save** button in the bottom-right corner of that section.

Figure 9-5 illustrates how to configure the Provisioning settings on a CE software-based endpoint.

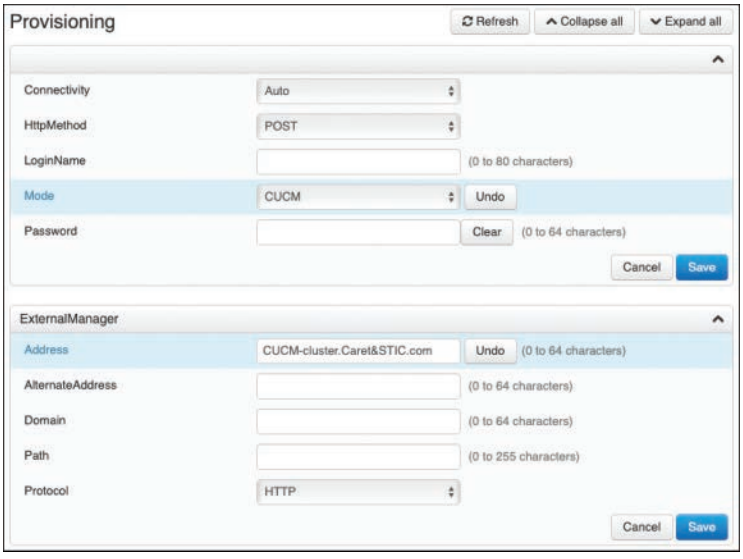

Figure 9-5 *Configure Provisioning Settings on a CE Software-Based Endpoint*

TFTP

IP phones and Telepresence endpoints rely on a TFTP-based process to acquire configuration files, software images, and other endpoint-specific information. The Cisco Unified Communications Manager runs a TFTP service as a file serving system, and this service can run on one or more Cisco Unified Communications Manager servers. It builds configuration files and serves firmware files, ringer files, device configuration files, and so forth, to endpoints. The TFTP file systems can hold several file types, such as the following:

Key Topic

- Phone configuration files
- Phone firmware files
- Certificate trust list (CTL) files
- Identity trust list (ITL) files
- Tone localization files
- User interface (UI) localization and dictionary files
- Ringer files
- Softkey files
- Dial plan files for SIP phones

The TFTP server manages and serves two types of files: those that are not modifiable, such as firmware files for phones, and those that can be modified, such as configuration files. A typical configuration file contains a prioritized list of Cisco Unified Communications Managers for a device, the TCP ports on which the device connects to those

Cisco Unified Communications Managers, and an executable load identifier. Configuration files for selected devices contain locale information and URLs for the messages, directories, services, and information buttons on the phone. When a device's configuration changes, the TFTP server rebuilds the configuration files by pulling the relevant information from the Cisco Unified Communications Manager database. The new file is then downloaded to the phone after the phone has been reset. An example of a file being downloaded to a phone again could be if a single phone's configuration file is modified, such as during Extension Mobility login or logout. Only that file is rebuilt and downloaded to the phone. However, if the configuration details of a device pool are changed, such as if the primary Cisco Unified Communications Manager server is changed, all devices in that device pool need to have their configuration files rebuilt and downloaded. For device pools that contain large numbers of devices, this file rebuilding process can impact server performance.

The TFTP server can also perform a local database read from the database on its co-resident subscriber server. A local database read not only provides benefits such as the preservation of user-facing features when the publisher is unavailable but also allows multiple TFTP servers to be distributed by means of clustering over the WAN. The same latency rules for clustering over the WAN apply to TFTP servers as apply to servers with registered phones. This configuration brings the TFTP service closer to the endpoints, thus reducing latency and ensuring failure isolation between the sites.

When a device requests a configuration file from the TFTP server, the TFTP server searches for the configuration file in its internal caches, the disk, and then remote Cisco TFTP servers, if specified. If the TFTP server finds the configuration file, it sends the file to the device. If the configuration file provides Cisco Unified Communications Manager names, the device resolves the names by using DNS and opens a connection to the Cisco Unified Communications Manager. If the device does not receive an IP address or name, it uses the TFTP server name or IP address to attempt a registration connection. If the TFTP server cannot find the configuration file, it sends a "file not found" message to the device.

A device that requests a configuration file while the TFTP server is processing the maximum number of requests will receive a message from the TFTP server that causes the device to request the configuration file later. The Maximum Serving Count service parameter, which is configurable, specifies the maximum number of requests that can be concurrently handled by the TFTP server. You can use the default value of 2500 requests if the TFTP service is run along with other Cisco CallManager services on the same server. For a dedicated TFTP server, you should use the following suggested values for the Maximum Serving Count: 2500 for a single-processor system or 3000 for a dual-processor system. The Cisco Unified IP Phones 7800 series and 8800 series request their TFTP configuration files over the HTTP protocol (port 6970), which is much faster than using the TFTP protocol.

Every time an endpoint reboots, an angel gets his wings. Not really; just wanted to make sure you're still paying attention. Every time an endpoint reboots, the endpoint will send a TFTP GET message for a file whose name is based on the requesting endpoint's MAC address. For a Cisco Unified IP Phone 8861 with a MAC address of abcdef123456, the filename would be SEPabcdef123456.cnf.xml. The received configuration file includes the version of software that the phone must run and a list of Cisco Unified Communications

Manager servers with which the phone should register. The endpoint might also download ringer files, softkey templates, and other miscellaneous files to acquire the necessary configuration information before becoming operational. If the configuration file includes software file version numbers that are different from those the phone is currently using, the phone will also download the new software files from the TFTP server in order to upgrade the firmware on the phone or endpoint. The number of files an endpoint must download to upgrade its software varies based on the type of endpoint and the differences between the phone's current software and the new software.

Option 150 allows up to two IP addresses to be returned to phones as part of the DHCP scope. The phone tries the first address in the list, and it tries the subsequent address only if it cannot establish communications with the first TFTP server. This address list provides a redundancy mechanism that enables phones to obtain TFTP services from another server even if their primary TFTP server has failed. Cisco recommends that you grant different ordered lists of TFTP servers to different subnets to allow for load balancing. An example of how this might look could be as follows:

- In subnet 10.0.0.0/24: Option 150: TFTP1_Primary, TFTP1_Secondary
- In subnet 10.1.1.0/24: Option 150: TFTP1_Secondary, TFTP1_Primary

Under normal operations, a phone in subnet 10.0.0.0/24 will request TFTP services from TFTP1_Primary, while a phone in subnet 10.1.2.0/24 will request TFTP services from TFTP1_Secondary. If TFTP1_Primary fails, phones from both subnets will request TFTP services from TFTP1_Secondary. Load balancing avoids having a single point of failure, where all phones from multiple clusters rely on the same server for TFTP service. TFTP load balancing is especially important when phone software loads are transferred, such as during a Cisco Unified Communications Manager upgrade, because more files of larger size are being transferred, thus imposing a bigger load on the TFTP server and on the network.

SIP Registration

As mentioned in Chapter 5, everything discussed up to this point in this chapter is not actually part of the SIP registration process; however, these processes had to transpire in order for SIP registration to the Cisco Unified Communications Manager to occur successfully. After the TFTP process has completed and the endpoint or phone has obtained all the system configuration information, SIP registration can occur. SIP endpoints must send a REGISTER request to a SIP server with a URI and a Call-ID, which is the IP of the endpoint registering. According to the IETF RFC 3621, a URI is made up of a host portion and a fully qualified domain name (FQDN) portion separated by an @ symbol, such as andy.dwyer@caret&stic.com. This is important to understand because the SIP registrar function of the SIP server breaks down these components individually to successfully process the REGISTER request. Each component is broken down as follows:

Key Topic

- **Request-URI:** The Request-URI names the domain of the location service for which the registration is meant, such as *sip:caret&stic.com*. The *userinfo* and @ components of the SIP URI must not be present. The domain in the URI must be qualified by the SIP server for registration to occur.

- **To:** The To header field contains the address of record whose registration is to be created, queried, or modified. The To header field and the Request-URI field typically differ because the former contains a username. This address of record must be a SIP URI or SIPS URI. Based on the preceding example, the request URI would be sip:andy.dwyer@caret&stic.com.
- **From:** The From header field contains the address of record of the person responsible for the registration. The value is the same as the To header field unless the request is a third-party registration. Therefore, the From header will also be sip:andy.dwyer@caret&stic.com.
- **Call-ID:** All registrations from a user agent client (UAC) should use the same Call-ID header field value for registrations sent to a particular registrar. If the same client were to use different Call-ID values, a registrar could not detect whether a delayed REGISTER request might have arrived out of order.
- **CSeq:** The CSeq value guarantees proper ordering of REGISTER requests. A user agent (UA) must increment the CSeq value by one for each REGISTER request with the same Call-ID.
- **Contact:** REGISTER requests may contain a Contact header field with zero or more values containing address bindings.

If you are already familiar with the Cisco Unified Communications Manager, you might see an issue here based on the IETF components needed for registering to the SIP registrar. The Cisco Unified Communications Manager uses directory numbers (DNs) for endpoint registration, not URIs. First, it's important to understand that the IETF RFCs are just guidelines for how SIP should operate. They are not hard-set boundaries you have to follow. If you want to "color outside the lines with SIP, you can certainly do so. Second, a DN on the Cisco Unified Communications Manager is essentially a URI in its base form.

Suppose you have a phone you were trying to register to a Cisco Unified Communications Manager at 10.0.0.30 with the DN 2001. The "domain" would be the IP address of the Cisco Unified Communications Manager, so the DN would actually be 2001@10.0.0.30. If DNS were used in a Cisco Unified Communications Manager environment, and the URL for the Cisco Unified Communications Manager was cucm1.caret&stic.com, the DN would actually be 2001@caret&stic.com.

When calls are placed between the Cisco Unified Communications Manager and the Expressway Core, incoming requests will use this format and so must be changed accordingly. The Expressway will be discussed further in Part IV. What is different within a Cisco Unified Communications Manager environment is the dialing behavior on the back end. The Cisco Unified Communications Manager will treat DNs dialed differently than URIs dialed within the same cluster. Bringing the subject back to registration, even though DNs are technically used on the Cisco Unified Communications Manager, for the purpose of registration they are treated the same as any other URI. Consider how the Cisco Unified Communications Manager will separate the different components of a REGISTER request from an endpoint with the DN 2001:

- **Request URI:** 10.0.0.30
- **To:** sip:2001@10.0.0.30

- **From:** sip:2001@10.0.0.30
- **CallID:** aef12b80-d1e1367c-2b2f-db8512c6@10.0.0.30
- **CSeq:** (101 REFER, 101 NOTIFY OR 101 SUBSCRIBE)
- **Contact:** N/A

Verifying phone registration from the Cisco IP phone is quite simple. The easiest way to verify that a phone is registered is to look on the display screen. If you see a message across the top of the screen and a spooling circle, the phone is not registered. If you see the assigned phone DN, such as 2001, beside Line 1 and at the top of the screen, the phone is registered. You can verify other registration information by pressing the Settings button and then selecting the Phone Information menu. On this screen, you can see the following information:

- **Model Number**, such as *CP8845*
- **IPv4 Address**, such as *10.0.0.100*
- **Host Name**, which is the *SEP<MAC Address>*
- **Active Load**, which is the Enterprise phone load version, such as *sip8845_65.11-7-1-17*
- **Last Upgrade**, such as *07/04/2019 2:34pm*
- **Active Server**, such as *ucm-sub1.caret&stic.com*
- **Stand-by Server**, such as *ucm-pub.caret&stic.com*

The last two options shown here signify that the phone is registered. These fields will be blank if the phone is not registered. Figure 9-6 illustrates an 8845 phone that has been registered to the Cisco Unified Communications Manager.

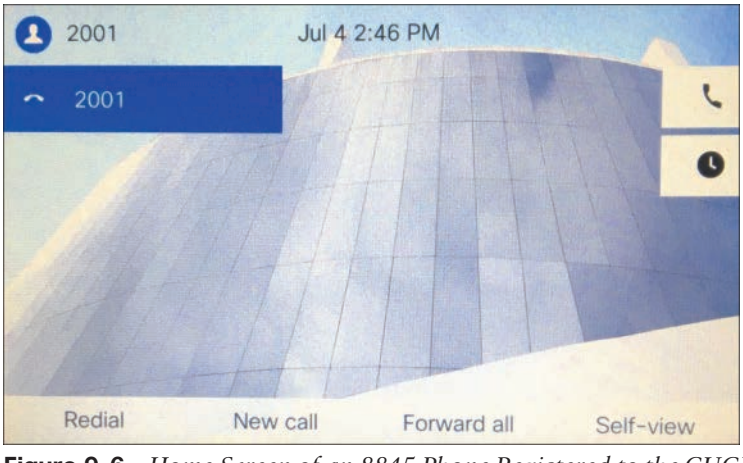

Figure 9-6 *Home Screen of an 8845 Phone Registered to the CUCM*

Verifying SIP registration on Cisco CE software-based endpoints is just as simple as verifying registration on Cisco IP phones. If you see a "Not Registered" message in the top-left corner of the screen, that is the indicator the phone is not registered. If you see the assigned

phone DN, such as 2002, in the center of the screen or the Touch 10 controller, the phone is registered. You can verify other registration information by pressing the Settings button in the top-left corner of the screen and then selecting the System Information menu. On this screen, you can see the following information:

- **Video Address**, such as *2002@caret&stic.com*
- **IP Address**, such as *10.0.0.101*
- **MAC Address**, such as *AB:CD:EF:12:34:56*
- **SIP Proxy**, such as *10.0.0.30 (Registered)*
- **Software**, such as *ce 9.1.5 d1c67fb 2017-11-16*
- **Device**, such as *Cisco Telepresence DX80 TANDBERG*

The SIP Proxy setting is another identifier that the endpoint is registered. For one thing, it says Registered. For another, this field will be blank or show Not Registered if there is an issue or if the phone has never been provisioned. Figure 9-7 illustrates a Cisco Telepresence DX80 that has been registered to the Cisco Unified Communications Manager.

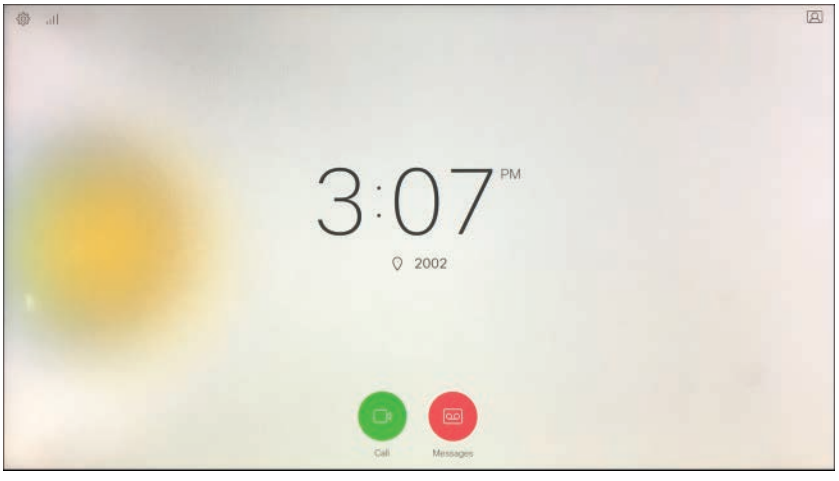

Figure 9-7 *Home Screen of DX80 Endpoint Registered to the CUCM*

ITL, CTL, and CAPF

Even though both an endpoint and the Cisco Unified Communications Manager might be located on the same LAN or WAN, it is still important to secure communications between these two products. Cisco offers several security measures, depending on what level of security you desire for the network you're managing. The lowest level of security available on the Cisco Unified Communications Manager for endpoint registration is the identity trust list (ITL). A step above this option is the certificate trust list (CTL). The highest level of security available is the Certificate Authority Proxy Function, or CAPF. In order to better understand each of these types of security on the Cisco Unified Communications Manager, it is important to formulate a basic understanding of how secure communication over IP works in its base form. If security certificates is an unfamiliar topic to you, review the "Security" subsection under the "Designing a Cisco Collaboration Solution" section in Chapter 6.

Key Topic

Returning to the topic of securing communications between an endpoint or phone and the call control system, a Security By Default (SBD) mechanism called ITL and Trust Verification Service (TVS) is enabled on all Cisco Unified Communications Manager installments. IP phones contain a limited amount of memory, and there can also be a large number of phones to manage in a network. The Cisco Unified Communications Manager acts as a remote trust store via TVS so that a full certificate trust store does not have to be placed on each IP phone. Any time the phone cannot verify a signature or certificate via the CTL or ITL files, it asks the TVS server for verification. This central trust store is easier to manage than if the trust store were present on all IP phones. A number of files must be present on the Cisco Unified Communications Manager itself. The most important piece is the TFTP certificate and the TFTP private key. The TFTP certificate is located under **OS Administration > Security > Certificate Management > CallManager.pem**. The Cisco Unified Communications Manager uses the CallManager.pem certificate's private and public keys for the TFTP service, as well as for the Cisco CallManager (CCM) service. All phones can use the TFTP public key in the CallManager.pem certificate to decrypt any file encrypted with the TFTP private key and to verify any file signed with the TFTP private key.

The presence of an ITL file will direct the phone to request a signed TFTP configuration file from the Cisco Unified Communications Manager TFTP server. The ITL file allows the phone to verify that the configuration file came from a trusted source. After the phone boots and obtains an IP address and the address of a TFTP server, it asks for the ITL files first. After the ITL file is downloaded, it must be verified. Several permutations could transpire here, including the following possibilities:

- The phone has no ITL file present. In this state, the phone blindly trusts the next ITL file downloaded and uses this signature henceforth.
- The phone already has an ITL file. In this state, the phone verifies that the recently downloaded files match the signature in either the ITL or TVS server.

After the signed configuration file is downloaded, the phone must authenticate it against the function for CCM+TFTP inside the ITL. After the phone receives the ITL files from TFTP successfully and they have been authenticated, the phone asks for a signed configuration file. This is the TFTP Get message discussed previously. With ITL files present on phones, configuration files must be signed by a trusted TFTP server. The file is plain text on the network while it is transmitted but comes with a special verification signature. The phone requests SEP<*MAC Address*>.cnf.xml.sgn to receive the configuration file with the special signature. This configuration file is signed by the TFTP private key that corresponds to CallManager.pem on the Operating System Administration Certificate Management page.

ITL is a leaner version of a CTL file. The CTL file contains a server certificate, public key, serial number, signature, issuer name, subject name, server function, DNS name, and IP address for each server. The CTL file can also contain entries for the following servers or security tokens:

- System Administrator Security Token (SAST)
- Cisco CallManager and Cisco TFTP services that are running on the same server

- Certificate Authority Proxy Function (CAPF)
- TFTP server(s)
- ASA firewall

After you create the CTL file, you must restart the Cisco CallManager and Cisco TFTP services in Cisco Unified Serviceability on all nodes that run these services. The next time that the phone initializes, it will download the CTL file from the TFTP server. If the CTL file contains a TFTP server entry that has a self-signed certificate, the phone requests a signed configuration file in .sgn format. If no TFTP server contains a certificate, the phone requests an unsigned file. After the Cisco CTL client adds a server certificate to the CTL file, you can display the certificate in the CTL client GUI. When you configure a firewall in the CTL file, you can secure a Cisco ASA firewall as part of a secure Cisco Unified Communications Manager system. The Cisco CTL client displays the firewall certificate as a CCM certificate. Cisco Unified Communications Manager Administration uses an e-token to authenticate the TLS connection between the Cisco CTL client and Cisco CTL provider. The process of obtaining a CTL file is the same as with an ITL file, except that the CTL request must occur first. The CTL file is then used to verify the ITL when it is requested. If the phone already has a CTL but no ITL, the phone trusts an ITL only if it can be verified by the CCM+TFTP function in the CTL file.

The Certificate Authority Proxy Function, which automatically installs with the Cisco Unified Communications Manager, performs several tasks depending on your configuration. It can be used to authenticate via an existing manufacturing installed certificate (MIC), locally significant certificate (LSC), randomly generated authentication string, or optional less secure "null" authentication. It issues locally significant certificates to supported Cisco Unified IP phones. It upgrades existing locally significant certificates on the phones. CAPF also retrieves phone certificates for viewing and troubleshooting. During installation, a certificate that is specific for CAPF gets generated. This CAPF certificate, which the Cisco CTL client copies to all Cisco Unified Communications Manager servers in the cluster, uses the .0 extension.

When the phone interacts with CAPF, the phone authenticates itself to CAPF by using an authentication string, existing MIC or LSC, or "null"; generates its public key and private key pair; and then forwards its public key to the CAPF server in a signed message. The private key remains in the phone and never gets exposed externally. CAPF signs the phone certificate and then sends the certificate back to the phone in a signed message. Before you use CAPF, ensure that you have performed all necessary tasks to install and configure the Cisco CTL client. To use CAPF, you must activate the Cisco Certificate Authority Proxy Function service on the first node.

SIP Registration to Expressway Core

Up to this point, we've discussed the registration process to the Cisco Unified Communications Manager using both the Cisco Unified IP phones and the Cisco Telepresence endpoints. Throughout the rest of this chapter, the focus will be on the Cisco Telepresence endpoints because the Cisco Unified IP phones cannot register to the Cisco Expressway. A proxy registration function using Mobile and Remote Access (MRA) allows Cisco Unified IP phones to register to the Cisco Unified Communications Manager through the Expressway Core and Edge, but in this type of deployment, the registration is still to the Cisco Unified

Communications Manager. This section of the chapter will focus on direct registration to the Cisco Expressway, which is intended for Cisco Telepresence endpoints and third-party endpoints.

Unlike the Cisco Unified Communications Manager, the Cisco Expressway does not have an easy provisioning mode available natively through which endpoints can register. There is a way to provision registration to a Cisco Expressway through another application called the Cisco Telepresence Management Suite (TMS), but TMS will not be discussed in this book. Therefore, endpoints that will register to the Cisco Expressway must be manually configured before they can register. This is one of the reasons the Cisco Unified Communications Manager is a preferred call control system. However, some circumstances require endpoints to register to the Expressway. A company may be using legacy Tandberg endpoints or third-party endpoints that cannot register to the Cisco Unified Communications Manager. Another reason could be that a company is already using a legacy VCS call control system and wants to upgrade to Expressway and continue to use the same system. Whatever the reason, this section will describe how to register SIP endpoints to the Cisco Expressway.

DHCP versus Static IP

Chapter 5 discussed the registration process to an Expressway using SIP. The following is a summation of that process for a quick review:

1. An endpoint obtains local power from the power cube and loads the locally stored image file. The exception is the SX10, which could obtain power from a power cube or PoE.
2. Cisco CE software-based endpoints can use CDP for VLAN discovery, but VLAN discovery will not impact Expressway-C registration.
3. When VLAN discovery is complete, or if the endpoint does not discover the VLAN, the endpoint will send a DHCP discovery message to the DHCP server. Alternatively, the endpoint may be configured with static network settings.
4. The SIP URI and Expressway IP address are configured manually on the endpoints.
5. The final step in the process is for the endpoints to register to the Cisco Expressway. The endpoints will send an IP address and alias to the Cisco Expressway in the REGISTER request. The alias must be in the form of a URI (name@FQDN), and the domain of the alias must match one of the domains configured in the Domain database of the Expressway.
6. If there are no configured restrictions on the Expressway, it will respond with the SIP message "200 OK." The registration process is now complete.

The SIP registration process to the Cisco Expressway is much easier to explain because there are not as many moving parts. As is evident in the preceding steps, no PoE settings need to be configured because most of the endpoints do not support PoE. No VLAN settings have to be configured because the default VLAN can be used to register to the Cisco Expressway. The Expressway will tag voice and video packets, so the voice VLAN can be used but is not required for QoS purposes. This brings the process to the network addressing step. DHCP can be used on Telepresence endpoints, but it is recommended to use static IP addressing for several reasons. First, there is always the possibility that the IP address can change when DHCP is being used. This can negatively impact several aspects of Telepresence

endpoint registration. The Expressway identifies endpoints registered by the IP address, not the MAC address, as does the Cisco Unified Communications Manager. Therefore, if the IP address changes, there could be duplicate registration entries for a single endpoint on the Cisco Expressway, causing call-routing issues. The IP address changing will also prevent, or at least slow down, an administrator from accessing the Telepresence endpoint remotely, which brings up the second reason static addresses should be used. The Cisco Telepresence endpoints support a high-functioning web interface and a command-line interface that can be accessed via the IP address. Cisco Unified IP phones do not have as robust of an interface; therefore, administrators do not need constant access to those addresses. All of the features a unified IP phone can support are managed through the Cisco Unified Communications Manager. A Cisco Telepresence endpoint, on the other hand, has many features supported directly on the endpoint itself, and accessing many of those features requires an IP address to access the Web interface of the endpoint, or access it through the CLI. Another reason static IP addressing is recommended on Cisco Telepresence endpoints is the newer API integrations that can be configured on newer versions of the CE software-based endpoints. There may be many more reasons why static addressing should be used over DHCP, but at the end of the day, you should use whatever method works best for your environment. Both methods will work in the end.

If you choose to use static IP addressing, the following instructions will provide some direction to change the default DHCP value to Static and configure all the necessary settings to render the Cisco Telepresence endpoint functional. Note that changing the IP address on a Cisco Telepresence endpoint will not impact access until after the system has been rebooted. A binding process must take place between the new IP address and the endpoint while the old address is purged from the system, which can occur only during a reboot. Therefore, you can make all the necessary changes over the web interface or CLI without losing the connection to the endpoint. However, after the reboot process has started, you will need to use the "new" IP address to regain access to the system. This first set of instructions will guide you through changing the IP settings using the web interface. This process is similar to the previously described process under the "DHCP" section.

Step 1. Log in to the web interface for the CE endpoint you want to configure.

Step 2. Navigate to **Setup > Configuration > Network** and locate the DNS section (second section down from the top).

DNS settings are not required for SIP registration to the Cisco Expressway. There may be reasons you may or may not want to configure these settings, so they are included for you to decide whether to use them.

Step 1. Enter the Domain in the Domain Name field. Enter up to three DNS addresses in the Server 1 Address, Server 2 Address, and Server 3 Address fields.

Step 2. Choose **Save** in the bottom-right corner of this section, and then scroll down to the IPv4 section.

Step 3. Change the Assignment field from DHCP to **Static**.

Step 4. Enter an IP address for the endpoint in the Address field. Enter the Default Router address in the Gateway field, and enter the Subnet Mask address in the SubnetMask field. All three of these fields are required for proper routing across the network.

Step 5. Choose **Save** in the bottom-right corner of this section.

Step 6. You will have to restart the endpoint before these changes will take effect. To restart the endpoint, choose **Maintenance > Restart**. Click the blue **Restart Device** button. When a popup window appears, click the red **Restart** button, and the system will go through a restart process. This will take about 60 seconds to complete.

One of the nice aspects to the design of the Cisco Telepresence CE software-based endpoints is the continuity between the web interface menu structure and the CLI command structure. To access the configuration menus from the web interface, you were instructed previously to choose **Setup > Configuration**. To access the configuration menus from the CLI, you need to enter the command **xConfiguration**. By the way, you can use abbreviations, such as **xconfig**, and you can use the Tab key to finish commands. Also, if you are not sure of what the next command should be, use the question mark and press Enter. A list of possible commands will be displayed. The next menu on the web interface to select is the Network menu from the left column. In the CLI, the next word to enter after **xConfiguration** is **Network**. On the web interface, if you wanted to configure the domain name, you would scroll down to the DNS section and find the Domain Name field. From the CLI, if you wanted to configure the domain name, you would enter the command **xConfiguration Network 1 DNS Domain Name:** followed by the domain name you want to configure. Figure 9-8 illustrates the continuity of these two interfaces with a side-by-side comparison between the CLI and the web interface configuration options for the network settings.

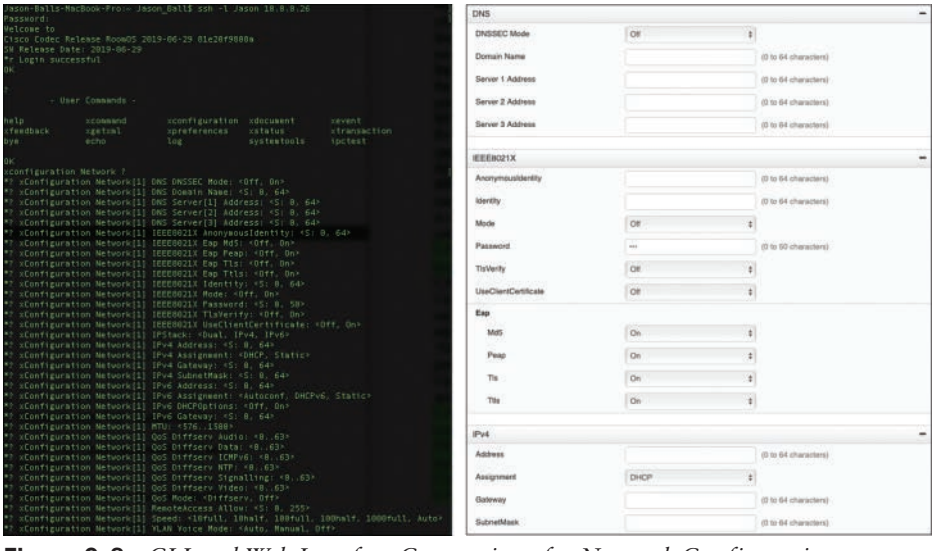

Figure 9-8 *CLI and Web Interface Comparison for Network Configuration*

To access the CLI of a CE software-based endpoint, you can open the terminal emulator of your choosing. PuTTY is a popular tool to use from Windows-based computers, and it is free to download. The preceding example is taken from Terminal, which comes with Apple Mac computers. You also can use other emulator tools, so do your homework if you do not already have a preference. Some will be free, and others you will have to pay for. You will need to use SSH to access the endpoint's CLI. Telnet is supported, but it is disabled by default. For security reasons, it is recommended to leave Telnet disabled. From PuTTY, open the emulator and enter the IP address in the Host Name (or IP address) field. Choose the

SSH radio button below that field, and then click the Open button. From Terminal, enter the command **ssh -l admin** *IP address* and then press Enter. Both tools will prompt you to validate the certificate and then prompt you to log in. With Terminal, the username is already added to the access command, so you will only need to enter the password. With PuTTY, you will need to enter the username first followed by the password. After you are logged in, you can enter the commands to configure the endpoint. The following is a list of network commands that can be entered through the CLI. Use the up-arrow key to recall the last command entered. Commands in the CLI are not case sensitive.

- **xConfiguration Network 1 DNS Domain Name***: domain*
- **xConfiguration Network 1 DNS Server [1-3] address***: IP of DNS*
- **xConfiguration Network 1 IPv4 Assignment***:* [**DHCP** | **Static**]
- **xConfiguration Network 1 IPv4 Address***: IP address for endpoint*
- **xConfiguration Network 1 IPv4 Gateway***: default gateway IP address*
- **xConfiguration Network 1 IPv4 SubnetMask***: subnet mask IP address*
- **xCommand Boot**

Because the last command does not configure any settings on the endpoint, it does not use the **xConfiguration** level of command. **xCommand** is the structure level that tells the system to execute a function. In this case, the endpoint must go through a reboot process to bind the IP address with the endpoint. After the reboot process begins, the connection to the emulator will terminate. You will need to use the new IP address to access the system remotely again.

Manual Configuration of SIP Settings

After the network settings have been configured on the CE software-based endpoint, you can configure the SIP settings so that the endpoint can try registering to the Cisco Expressway. When a CE software-based endpoint is first booted, a setup wizard runs on the endpoint. You can choose whether the endpoint will register to Cisco Webex or to Other Services. After you select Other Services, the options presented are Cisco UCM, Cisco UCM via Expressway, VCS (VCS can also be an Expressway Core), or Advanced Setup. When you select the Cisco UCM option, a field called Enter Server Address will populate so that you can configure the TFTP server address. When you select the Cisco UCM via Expressway option, three fields will be populated: Username, Passphrase, and Domain. When you select the VCS option, four fields will be populated: Host Server Address, Username, Passphrase, and Domain. When you select the Advanced Setup option, the next screen will provide a warning: "Advanced setup is an option for administrators to directly set up the system with their own software. If you continue you will not be able to call until you register with a service." The steps outlined here assume Advanced Setup is the option you chose on the endpoint. Two methods can be used to manually configure a CE software-based endpoint to an Expressway Core. The administrator can use the web interface or the CLI, as mentioned previously with the network settings options. This section will first examine the web interface options for configuring the SIP registration settings on a CE endpoint.

First, you need to obtain the IP address from the endpoint and enter this address in the address bar on a web browser. You can use either HTTP or HTTPS to access the endpoint by default, although you may want to disable HTTP for security reasons. When the login screen appears, enter the Username admin and leave the Passphrase field bank. CE software-based endpoints do not have a password set on them by default. The administrator will need to set this password. If you want to go ahead and set a password on the endpoint, click the Admin link in the top-right corner of the screen and select the Change Passphrase menu option. You can also navigate to **Security > Users** and select the Admin user account to gain access to this screen. You will be prompted to enter the Current Passphrase, Passphrase, and Repeat Passphrase. Here, leave the first field blank, enter a password in the second and third fields, and then click the Change Passphrase button. Now you are ready to configure the SIP registration settings to the Expressway Control. Navigate to **Setup > Configuration** and follow these steps:

Key Topic

Step 1. Choose **NetworkServices** from the menu column on the left of the screen, and ensure the SIP Mode setting is set to On. This is the default setting.

Step 2. Choose **Provisioning** from the menu on the left and change the Mode to **VCS**. This is an optional step because registration to the Expressway Core will occur whether this setting is changed or not. The mode could also be set to TMS or Auto and still register to the Expressway Core. Click **Save** after changing the setting.

Step 3. Choose **SIP** from the menu column on the left. You can configure several settings under this menu. Some of those settings are as follows:

- **a.** **DefaultTransport:** This setting defines the default mechanism used when SIP communications are sent. The values can be set to Auto, TCP, TLS, or UDP. The default value is Auto. You can leave this setting as Auto or change it to another value. For security reasons, you may want to change this setting to TLS.
- **b.** **DisplayName:** This is an optional setting. If the display name is set, when a call is placed from or to this endpoint, this name will be displayed to the participants on the other end of the call. If this value is not configured, the URI address will be displayed as the name.
- **c.** **Proxy 1 Address:** This is a required field to be configured. This is where you will need to enter the address of the Expressway Core. If DNS is configured on this endpoint, you can enter the URL of the Expressway Core. If DNS is not being used, you will need to enter the IP address of the Expressway Core. In either case, some setting must be configured in this field.
- **d.** **TLS Verify:** This setting should be enabled only if TLS Verified is being used. You will first need to upload a signed certificate to this endpoint, DNS will have to be enabled, and the Proxy 1 Address will have to be a URI. The default value for this setting is Off.

e. **Type:** Two types of SIPs can be used on CE software-based endpoints. Standard is the default value, and the value that should be selected when registering to the Expressway Core. Cisco is the other value option and should be used only when registering to the Cisco Unified Communications Manager. When the **Provisioning > Mode** is set to CUCM, this value will change automatically to Cisco.

f. **URI:** This is the URI address assigned to this endpoint. This is a required field because the Expressway Core will check the domain portion of the URI address against a domain database within the Expressway before allowing registration. The URI should be in the form of Host@FQDN.

Step 4. After you configure all the settings under this section, click the **Save** button.

Registration should occur instantaneously. To verify the endpoint has registered successfully, choose Home from the menus at the top of the screen. On the right side of the screen under the SIP Proxy 1 heading, you should see the registration status of the endpoint. Figure 9-9 illustrates the SIP Configuration menus on a DX80 endpoint.

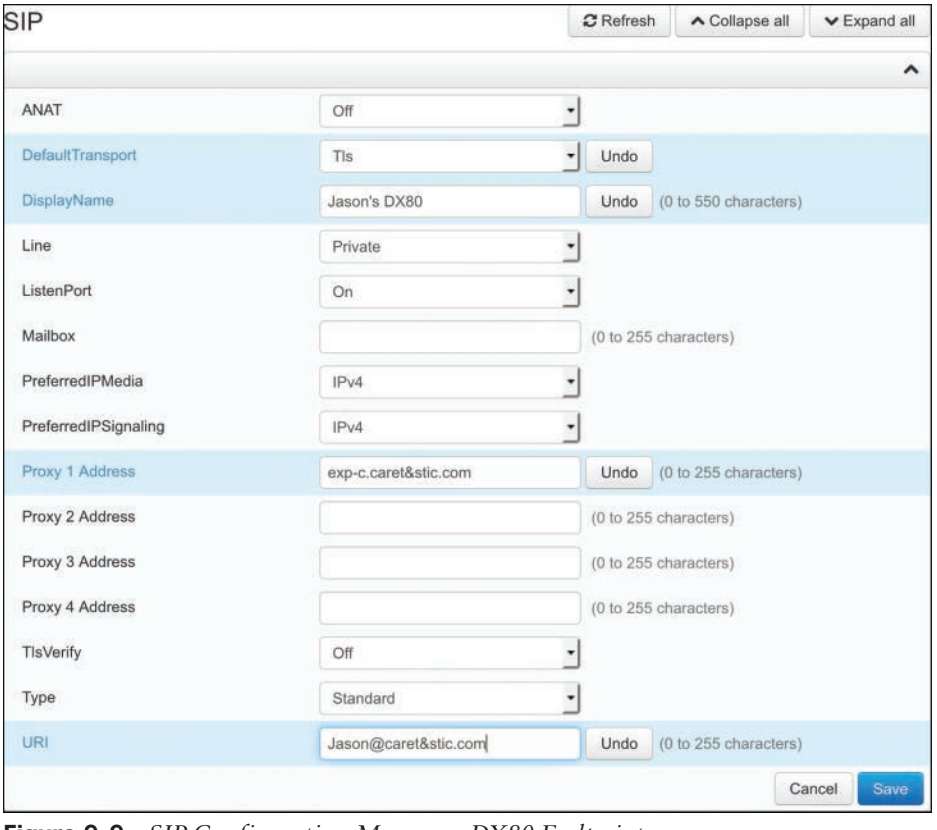

Figure 9-9 *SIP Configuration Menus on DX80 Endpoint*

Another method that you can use to manually configure a CE software-based endpoint to register with an Expressway Core is to use the CLI. The preceding section discussed different terminal emulators that can be used for CLI access and explained the process to log in to the endpoint using the PuTTY and Terminal emulators. The following CLI commands can be used to configure all of the settings previously mentioned using the web interface:

- **xConfiguration NetworkServices SIP Mode***: On*
- **xConfiguration Provisioning Mode***: VCS*
- **xConfiguration SIP DefaultTransport***: tls*
- **xConfiguration SIP DisplayName***: Andy_Dwyer_DX80*
- **xConfiguration SIP Proxy 1 Address***: exp-c.caret&stic.com*
- **xConfiguration SIP URI***: andy.dwyer@caret&stic.com*
- **xStatus SIP**

All of the settings in italic in the preceding list are examples of the values that you can add. Obviously, you would want to configure your own unique values in each of these fields. The last command, **xStatus SIP**, is how you can check the registration status for SIP on the endpoint. This will display all the settings configured for SIP with their values, the Registration 1 Status, and the Registration 1 URI address.

H.323 Registration to the Expressway Core

In addition to supporting SIP registration, the Cisco Expressway can also support H.323 registration. As discussed in Chapter 5, H.323 is an ITU-T standard for packet-switched communication over IP. H.323 and SIP cannot communicate with one another unless there is a gateway to bridge the differences between the two communication protocols. Therefore, the Expressway has a built-in SIP to the H.323 gateway that is enabled by default. This gateway can also bridge communications between IPv4 and IPv6. When communicating between SIP and H.323, it is important to pay attention to the aliases being used because the different alias mechanisms between the two protocols can cause further complications in bridging the communication chasm. Other tools within the Expressway can aid in bridging this gap, but they will not be discussed in this book.

H.323 Aliases

An alias can be generally defined as any identification method that is not an IP address. When video endpoints in a network are known to the rest of the environment only by their IP addresses, many limitations can occur. These limitations can be summarized by saying that having some flexibility and customization available to the administrator serves everyone well in terms of overall system administration, security, dial plan effectiveness, and so on. The two primary conferencing protocols, SIP and H.323, use very different alias methodologies and when not implemented properly can create unnecessary workload or configuration challenges.

H.323 aliases can take three forms. First, H.323ID is an alphanumeric string of characters that includes special characters but no spaces, such as *room231* or *helpdesk*. The H.323ID can also be configured in the form of a URI, such as *andy.dwyer@caret&stic.com*. Make

no mistake: this is an H.323 ID and not a URI. Remember that the FQDN in the URI for SIP must be qualified against a domain configured in the Expressway. However, with an H.323 ID, the supposed "domain" part of the alias is not qualified against anything. The second alias supported with H.323 is an E.164 alias. This type of alias can only be configured with numeric values consisting of 1–15 digits. Incidentally, the E.164 protocol is a holdover from PSTN protocols and serves to unify legacy ISDN endpoints with newer IP-based endpoints. The third type of alias is a routing prefix. These prefixes are registered aliases that have been configured on gateways and bridges. They serve to route all calls that were dialed with that prefix number to the server that registered the routing prefix. For example, if an ISDN gateway were to register with a routing prefix of 9, and a user dialed 919195551001, the Expressway would identify the first digit as a routing prefix, ignore all the numbers that followed, and route the call to the ISDN gateway. The call connection would then depend on the routing rules configured in the ISDN gateway to match the call attempt and connect it across the PSTN to the number dialed. Another example could be an MCU of some type that registered a routing prefix of 814. When a user dialed 8144001, the Expressway would match the 814 prefix and route the call to the MCU regardless of remaining digits dialed. This second example serves to illustrate that a prefix does not have to be a single digit. It could contain multiple digits, but it should be kept simple so that dialing among employees does not become a cumbersome task. Endpoints do not use prefixes, so this alias type would not apply to them.

The previous information that a single endpoint can have as many as two different aliases is significant because the Expressway identifies endpoints by their alias. The Expressway will, therefore, make all its security, access, and bandwidth management decisions based on the alias of an endpoint, and will only be concerned with IP addresses when regarding routing the signaling and media. When an administrator is planning and configuring the dial plan for Cisco Expressway call routing, how the Expressway will treat each alias must be given careful consideration.

Manual Configuration of H.323 Settings

Just as with SIP registration configurations on CE software-based endpoints, H.323 registration configurations can be made using the web interface or the CLI. The following steps outline how to configure H.323 settings on CE software-based endpoints for registration to the Gatekeeper function on the Cisco Expressway Core. H.323 and SIP cannot be used simultaneously on CE software-based endpoints. Therefore, SIP will need to be disabled when H.323 is enabled on the endpoint.

Step 1. Log in to the web interface for the CE endpoint you want to configure.

Step 2. Choose **Setup > Configuration**.

Step 3. Choose **NetworkServices** from the menu column on the left side of the screen.

Step 4. Change the SIP Mode setting to **Off**, and then change the H.323 Mode to **On**. This setting is Off by default. Click **Save** when finished.

Step 5. Choose **Provisioning** from the menu on the left and change the Mode to **VCS**. This is an optional step because registration to the Expressway Core will occur whether this setting is changed or not. The mode could also be set to TMS or Auto and still register to the Expressway Core. Click **Save** after changing the setting.

Step 6. Choose **H323** from the menu on the left. Several settings under this menu can be configured. Some of those settings are as follows:

a. **CallSetup Mode:** This setting can be set to Gatekeeper, which is the default, or Direct. If Gatekeeper is chosen, the endpoint must register to an H.323 Gatekeeper before it can place and receive calls. If the endpoint is set to Direct, it will never try to register but can place calls based on IP address dialing.

b. **Gatekeeper Address:** This is where the IP address or URL of the Cisco Expressway Core needs to be entered for the endpoint to register.

c. **Authentication:** Three settings pertain to authentication for H.323 registration:

   i. **LoginName:** This is the login name used when Authentication services are enabled on the Expressway Core. If Authentication services are not enabled on the Expressway Core, this field can be left blank.

   ii. **Mode:** This setting can be set to On or Off (default is Off). Authentication mode does not need to be turned on unless Authentication services are enabled on the Expressway Core. If this mode is turned on and Authentication services are not enabled on the Expressway Core, it will not impact registration. This setting will just be ignored by the Expressway Core.

   iii. **Password:** This setting coincides with the Authentication username when the Authentication services are enabled on the Expressway Core.

d. **H323Aliases:** Two aliases can be configured on CE software-based endpoints.

   i. **E164:** This is the numeric alias assigned to H.323 endpoints. Digits can range between 0 and 9, and aliases can be composed of up to 15 digits in length.

   ii. **ID:** This alias can contain alphanumeric and special characters up to 32 characters in length.

There are several more settings than what has been listed here, but the remainder of the settings are seldom used and therefore do not warrant any discussion. Figure 9-10 illustrates the previously described settings for H.323 configuration under the H.323 menu on a Cisco DX80 endpoint.

Another method that you can use to manually configure a CE software-based endpoint to register with an Expressway Core is to use the CLI. The CLI commands that can be used to configure all of the H.323 settings previously mentioned using the web interface are as follows:

- **xConfiguration NetworkServices SIP Mode***: Off*
- **xConfiguration NetworkServices H323 Mode***: On*
- **xConfiguration Provisioning Mode***: VCS*
- **xConfiguration H323 CallSetup Mode***: Gatekeepr*

- **xConfiguration H323 Gatekeeper Address**: *exp-c.caret&stic.com*
- **xConfiguration H323 Authentication LoginName**: “”
- **xConfiguration H323 Authentication Mode**: *Off*
- **xConfiguration H323 Authentication Password**: “”
- **xConfiguration H323 H323Alias E164**: *4002*
- **xConfiguration H323 H323Alias ID**: *4002@caret&stic.com*
- **xStatus H323**

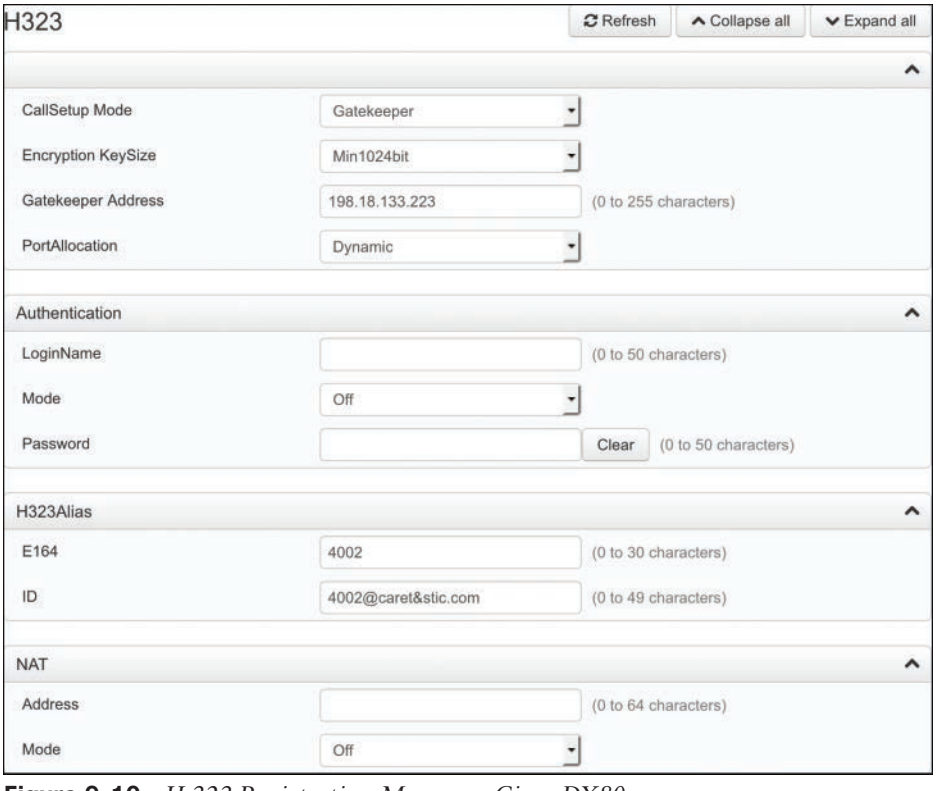

Figure 9-10 *H.323 Registration Menus on Cisco DX80*

All of the settings in italic in the preceding list are examples of the values that you can add. Obviously, you would want to configure your own unique values in each of these fields. Because the Authentication Username and Password are not needed in this example, quotation marks are used to show there is no value in that field. Quotation marks can be used to remove a setting as well. The last command, **xStatus H323**, is how you can check the registration status for H.323 on the endpoint. This will display the H.323 Gatekeeper Address, Port, Mode as Enabled, and Status as Registered. Aliases will not be displayed with the **xStatus H323** command. To see the alias configured on this system, type **xConfiguration H323** and press Enter. All H.323 settings will be displayed as they have been configured.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 9-4 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 9-4 Key Topics for Chapter 9

Key Topic Element	Description	Page Number
Table 9-2	Meraki Switch PoE Classifications	197
Table 9-3	Three PoE Types Supported on Cisco Switches	198
Commands	Enabling PoE on a Cisco Switch Port	198
Commands	Enabling CoS on Cisco Switches	201
Paragraph	Dual VLAN Tagging on a Switchport	203
List	Configure VLANs on Cisco 8800 Series Phones	203
List	Configure VLANs on Cisco CE Software-Based Endpoints	204
Commands	Configure DHCP with Option 150 on Cisco Router	205
List	DHCP Settings on Cisco 8800 Series Phones	207
List	DHCP Settings on Cisco CE Software-Based Endpoints	208
List	TFTP System File Types	209
Paragraph	TFTP GET Process Between Phone and CUCM	210
List	Components of a SIP REGISTER Request	211
Paragraph	TVS Explained	215
Paragraph	CAPF Explained	216
Paragraph	Static IP versus DHCP on Telepresence Endpoints	217
List	CLI Commands on CE Endpoints to Configure Network Settings	220
List	Steps to Configure SIP Registration to an Expressway via Web Interface	221
Paragraph	Three Forms of H.323 Aliases	223
List	Steps to Configure H.323 Registration to an Expressway via Web Interface	224

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

802.1p, 802.1Q, 802.3af, 802.3at, Asymmetric Cryptography, CA, CAPF, CDP, Classification Pulse, CLI, CoS, CTL, Data VLAN, DHCP, Diffie-Hellman Key Exchange, DN, DNS, E.164 alias, FQDN, H.323 ID, HTTP, HTTPS, IEEE, ITL, LLDP-MED, Mutual TLS, Option 150, Option 66, PD, PoE, PoE Power Budget, Prestandard PoE, PSE, QoS, Routing Prefix, RSA, SIPS, SSL, Symmetric Cryptography, TFTP, TLS, TLS Verify, TVS, URI, URL, VLAN, Voice VLAN, VVID, WAN

Command Reference to Check Your Memory

This section includes the most important configuration and EXEC commands covered in this chapter. It might not be necessary to memorize the complete syntax of every command, but you should be able to remember the basic keywords that are needed.

To test your memory of the commands, cover the right side of Table 9-5 with a piece of paper, read the description on the left side, and then see how much of the command you can remember.

The CLCOR (350-801) exam focuses on practical, hands-on skills that are used by a networking professional. Therefore, you should be able to identify the commands needed to configure and test switch commands for PoE and QoS, and router commands for DHCP and Option150.

Table 9-5 Cisco Meeting Server MMP Commands

Task	Command Syntax
This command enters configuration mode.	Switch# **Configure Terminal**
This command selects the interface to configure.	Switch(config)# **interface {fastethernet|gigabitethernet} (slot/port)**
The **auto** keyword sets the interface to automatically detect and supply power to the powered device. This is the default configuration. The **static** keyword sets the interface to higher priority than auto. If necessary, you can use the **max** keyword to specify the maximum wattage allowed on the interface (4000 to 15,400 milliwatts). You can use the **never** keyword to disable detection and power for the PoE-capable interface.	Switch(config-if)# **power inline {auto[max milli-watts] | never | static [max milli-watts]}**
This command exits configuration mode.	Switch(config-if)# **end**
This command displays the PoE state for the switch.	Switch# **show power inline {fastethernet|gigabitethernet} slot/port**
This command configures the port to delay shutting down.	Switch(config-if)# **power inline delay shutdown 20 initial 300**

Task	Command Syntax
This command creates a VLAN and associated number value.	Switch(config)# **vlan** *number*
This command provides a description of the VLAN.	Switch(config-vlan)# **name** *name*
This command enables the switchport to trust the CoS-to-QoS mapping embedded in the switch.	Switch(config-if)# **mls qos trust cos**
The **detect** command configures the interface to detect and recognize a Cisco IP phone. The **cisco-phone** option is the only one allowed when you initially implement the **switchport voice detect** command. The default is **no switchport voice detect cisco-phone [full-duplex]**. The optional **full-duplex** command configures the switch to accept only a full-duplex Cisco IP phone. The **vlan-id** command configures the phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP phone forwards the voice traffic with an IEEE 802.1Q priority of 5. Valid VLAN IDs are 1 to 4094. The **dot1p** command configures the phone to use IEEE 802.1p priority tagging for voice traffic and to use the default native VLAN (VLAN 0) to carry all traffic. The **none** command allows the phone to use its own configuration to send untagged voice traffic. The **untagged** command configures the phone to send untagged voice traffic.	Switch(config-if)# **switchport voice {detect cisco-phone [full-duplex] | vlan {vlan-id | dot1p | none | untagged}}**
This command is used to hard-code the port into access mode.	SW1(config-if)# **switchport mode access**
This command is used for access VLAN tagging all data originating from the computer.	SW1(config-if)# **switchport access vlan number**
This command is used for voice VLAN tagging all voice and video traffic originating from the phone.	SW1(config-if)# **switchport voice vlan** *number*
This command creates a pool from which IP addresses can be issued to devices that send a DHCP request. The *name* field can be any name you want to give to the pool.	Router(config)# **ip dhcp pool** *name*
This command establishes all the available addresses within a pool that can be used for DHCP assignment.	Router(dhcp-config)# **network** *starting IP address subnet mask*
This command establishes the default gateway address that will be assigned to devices.	Router(dhcp-config)# **default-router** *default gateway address*

Task	Command Syntax
This command assigns the TFTP server address to the endpoint. You could also use the **option 66 ip** *address* command here, but Cisco recommends using **Option 150** to add additional redundancy.	Router(dhcp-config)# **option 150** *TFTP server address*
This command allows you to list up to four DNS server addresses. You must separate DNS addresses with a space.	Router(dhcp-config)# **dns-server** *DNS address*
This command enables you to assign the domain to devices through DHCP.	Router(dhcp-config)# **domain-name** *name*
This command determines how long a leased address can be used by a device before a new lease has to be requested. By default, the duration of a lease is one day. When you enter the **lease** *n* command, this duration will be extended to that number *n* of days. You can enter three values here to extend the duration to days, hours, and minutes.	Router(dhcp-config)# **lease** *n*
This command issues an exclusion range of addresses that will not be used in DHCP assignments.	Router(config)# **ip dhcp excluded-address** *starting IP address ending IP address*

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the three PoE protocols supported on Cisco switches and the maximum watts each can support.
2. Assuming the data VLAN ID is 100 and the VVID is 200, list the commands required to apply both the voice and data VLANs to a switchport.
3. What are the two options that can be configured in a DHCP server so that the TFTP server address can be discovered by an endpoint?
4. List the nine different file types that the TFTP file system can hold.
5. List the six different fields found in a SIP REGISTER header.

This page intentionally left blank

CHAPTER 10

Call Settings on Cisco CE Software-Based Endpoints

This chapter covers the following topics:

Calling Options: This topic will discuss the various ways calls can be placed from CE software-based endpoints.

Content Sharing Options: This topic will discuss the various options available for sharing content through a CE software-based endpoint both locally and during a call.

Other Options: Various other options can be leveraged from CE software-based endpoints. This topic will introduce the function of these other options and how they can be configured.

The previous chapter mentioned settings that can be configured on CE software-based endpoints that cannot be configured on Cisco Unified IP phones. H.323 and SIP registration settings are only some of the features that set these intelligent systems apart. This chapter will delve into some of the other features that uniquely identify the superiority of CE software-based endpoints. Topics discussed in this chapter include the following:

- **Calling Options:**
  - Call by Alias
  - Call by Directory
  - Multipoint Calling
  - One Button to Push (OBTP) and Scheduled Conferences
- **Content Sharing Options:**
  - Direct Sharing Content
  - Using Intelligent Proximity for Content Sharing
- **Other Options:**
  - Audio Settings
  - Encryption Mode
  - AutoAnswer
  - Far-End Camera Control (FECC)
  - Phonebook
  - Video Settings

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 10-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 10-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Calling Options	1–4
Content Sharing Options	5–6
Other Options	7–11

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following user interface control devices can be used on the Cisco DX80 endpoint?
   - **a.** TRC5 Remote Control
   - **b.** Touchscreen
   - **c.** Touch 10 control pad
   - **d.** Web interface
2. How many directory entries can a Cisco CE software-based endpoint hold locally within the endpoint database?
   - **a.** 35
   - **b.** 350
   - **c.** 3500
   - **d.** 35,000
3. Which of the following terms is defined by any call involving three or more participants?
   - **a.** Multipoint
   - **b.** Multisite
   - **c.** Multiway
   - **d.** Ad hoc

4. Which of the following devices can be used to schedule OBTP meetings?
   a. CUCM
   b. VCS
   c. TMS
   d. CMS
5. Which of the following is the ITU standard for content sharing?
   a. DuoVideo
   b. People+Content
   c. H.224
   d. H.239
   e. BFCP
6. Which of the following devices can initiate content being shared via the Proximity application?
   a. Windows computer
   b. Smartphone
   c. Tablet
   d. Cisco Unified IP phone
7. Which of the following statements is true?
   a. The Microphone Mute Enabled setting will mute the microphone when a call is set up.
   b. The Microphone Mute Enabled setting will not mute the microphone when a call is set up.
   c. The Input Microphone Mode setting will mute the microphone when a call is set up.
   d. The Input Microphone Mode setting determines which microphone is set as the primary.
8. Which of the following is the default Encryption Mode setting on a CE software-based endpoint?
   a. AES 128
   b. AES 256
   c. On
   d. Best Effort
9. Which of the following is the standard for FECC?
   a. H.224
   b. H.239
   c. T.150
   d. BFCP

10. Which of the following directories requires the endpoint to send a subscribe message before phonebook entries can be retrieved?
    a. Local directory
    b. Global directory
    c. Corporate directory
    d. Both global and corporate directories
    e. No directory requires a subscribe message before phonebooks can be received.
11. The RGB Quantization Range setting overrides devices that do not follow the CEA-861 standard in order to provide the perfect image with any display. Which of the following settings is used to set the RGB quantization range based on the RGB Quantization Range bits (Q0, Q1) in the AVI infoframe?
    a. Auto
    b. Full
    c. Limited
    d. Manual
    e. None of these answers are correct.

Foundation Topics

Calling Options

Because the primary purpose of having an endpoint is to be able to place and answer calls, it is important to understand how to use these systems once they have registered to the call control system of your choosing. A call can be placed or answered from CE software-based systems in essentially three ways:

Key Topic

1. Call out or answer an incoming call from the user interface. This could be the Touch 10 controller, a remote control, or a touchscreen on the endpoint, depending on which device is being used. This could also be a personal device, such as a smartphone or tablet, through the Proximity application.
2. Use the web interface.
3. Use the command-line interface (CLI).

An administrator typically uses the web interface and CLI. Because end users will be sitting in front of the system, they would interact with the user interface.

Calls from a CE software-based endpoint can be made by dialing the destination alias or by selecting a participant from the directory, or phone book. Multipoint conference meetings can be arranged by dialing into a bridge that will host the meeting or by utilizing the Multisite option on the CE endpoint itself. The Multisite option is a licensed feature that must be added to the CE software-based endpoint before this feature can be used. Scheduled meetings can be accessed in a variety of ways as well, such as using a feature called *One Button to Push* (OBTP). The following sections will delve into each of these dialing behaviors to provide a more thorough understanding of call behavior and how to configure settings related to each of these calling components.

Call by Alias

Since the invention of automatic telephony switches, the most common means of initiating communication with another party, whether through a PSTN telephone, IP phone, or video phone, has been by dialing the alias of the destination. Traditional PSTN phones and most IP phones use E.164 aliases to dial, which is more commonly known as the *telephone number*. In the case of the Cisco Unified Communications Manager, these aliases are known as *directory numbers (DNs)*. However, as explained already in previous chapters, other alias types can be used, and they are growing in popularity. The most common type of alias used outside of E.164 aliases is the SIP URI. Regardless of the alias type, the endpoint used to place calls must possess the ability to dial the destination alias. All Cisco CE software-based endpoints share a common interface so that no matter what Cisco Telepresence endpoint is being used, the experience will be the same for every user. The following figures and descriptions are based on the Cisco DX80 endpoint, but they can be applied to any DX, MX, SX, IX, or Webex Telepresence endpoint.

Key Topic

Users can use three different control mechanisms available on Cisco CE software-based endpoints to interact with the different systems. The Cisco SX10 endpoint comes with a TRC5 remote control. The DX80 endpoint has a touch control screen, and all other Cisco endpoints come with a Touch 10 control pad. Regardless of which control device is being used on the Cisco endpoints, the screen layouts and menu options are the same. Figure 10-1 illustrates the menu layout options on a Cisco DX80 touchscreen.

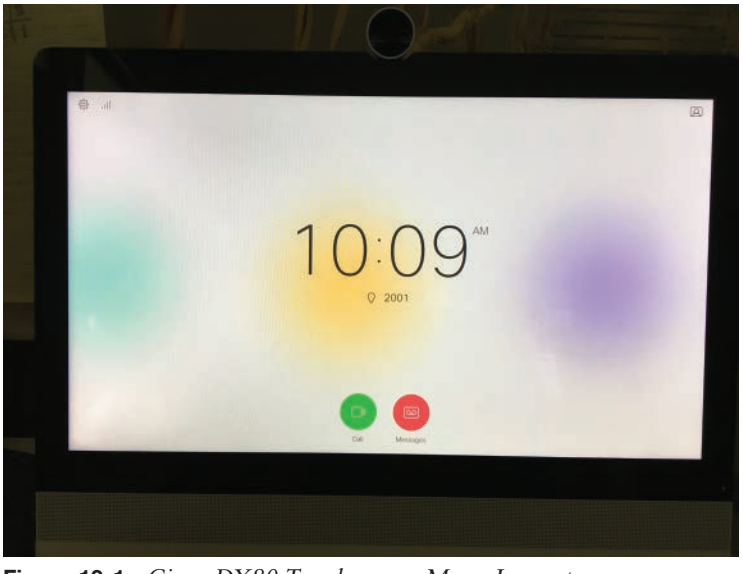

Figure 10-1 *Cisco DX80 Touchscreen Menu Layout*

On the user interface, the time of day will always be displayed on the center of the screen. There is an option to set this to a 12-hour or 24-hour clock when the system is set up for the first time. In the top-right corner of the screen is an icon with a person displayed on a screen. If you select this option, a self-view window will appear. Once displayed, the self-view window can be moved to six different points around the screen: the top-left corner, top-right corner, middle-left edge, middle-right edge, bottom-left corner, or bottom-right corner. The gear icon in the top-left corner will bring up some selective menu options. The

menus available from the top down are Do Not Disturb, Light Adjustment Bar, Forward All Calls To…, Forward All Calls to Voicemail, System Information, and Standby. The System Information menu will provide the endpoints' video address, IP address, MAC address, SIP Proxy, software version, and device type. The Settings button along the bottom of the screen will allow you to view and change some extended settings on the endpoint. However, advanced settings need to be configured from the web interface or the CLI.

Key Topic

At the bottom of the main screen, located under the clock, are two circles. The red circle is a direct access to a voicemail box, and the green circle is used for calling. If you select the green Call button, a dial box will display. When you select the dial box, a QWERTY keyboard with a numeric dialpad will display at the bottom of the screen. This allows for easier dialing of both SIP URIs and E.164 aliases. You can enter the alias of the destination and click the green Call button that appears to the right. This Call button will not appear until you start typing a destination alias. Once the Call button is pressed, the destination alias will ring, and the call will connect when the far-end participant answers the call. Figure 10-2 illustrates the call settings for dialing by alias from the user interface.

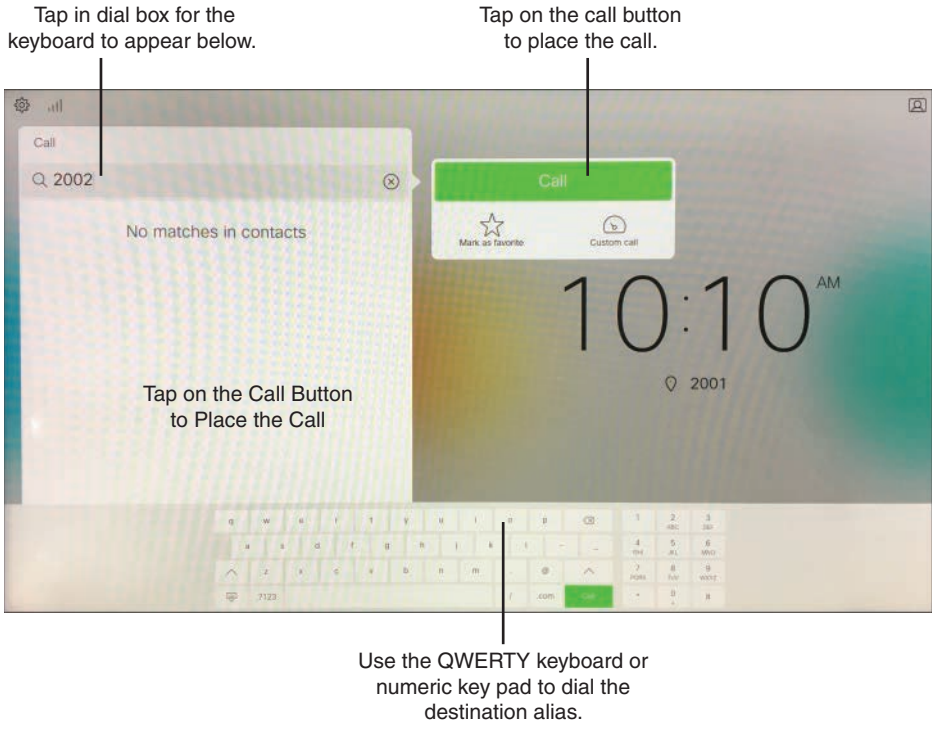

Figure 10-2 *Cisco DX80 Dial by Alias Settings from User Interface*

Another way to dial from an endpoint is to use the Intelligent Proximity for Content Sharing application from a Windows or Mac computer, smartphone, or tablet. Once the Proximity application pairs with the endpoint, a dial box will appear in the middle of the screen. When you tap inside this box, a keyboard will display at the bottom of the screen. You can dial the alias of the destination endpoint and tap the green Call button. The endpoint, not the Proximity app, will dial out to the destination. Once the call has connected, you can use the Proximity app to adjust the volume on the endpoint, receive content being shared, scroll

through previously shared content, and save content to the Photos library on your device. There is also an End button on the Proximity app to hang up a call, and when the endpoint is being called, an Answer button will appear. Figure 10-3 illustrates some of the functionality that exists with Intelligent Proximity.

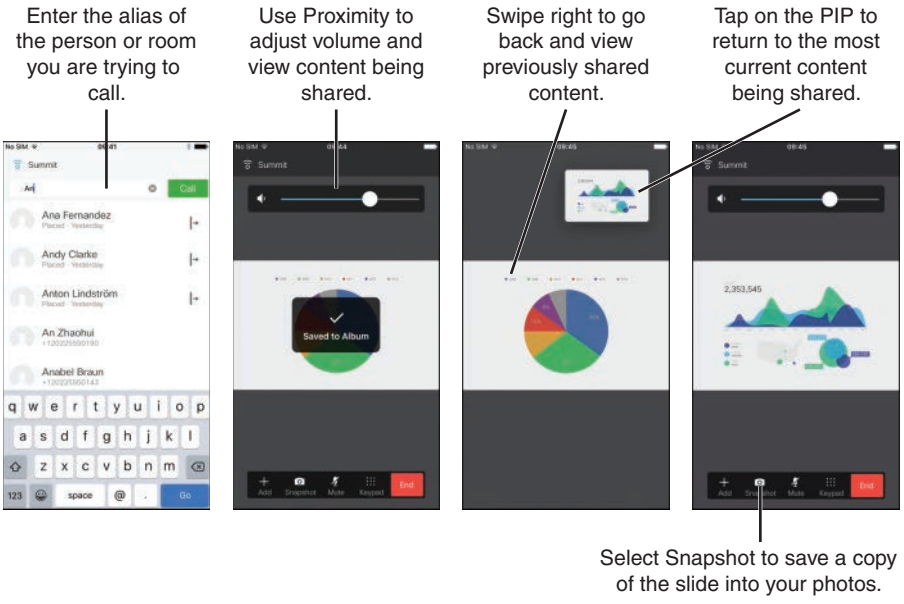

Figure 10-3 *Dialing from an Endpoint Using Proximity*

Key Topic

The next two options for dialing out from a Cisco CE software-based endpoint using the alias of a destination are only used typically by an administrator. The web interface or CLI allows an administrator to dial out from an endpoint without being physically present with the endpoint. This serves a great purpose when troubleshooting call setup issues from remote locations. This capability could also prove useful when an administrator needs to dial out on behalf of a user at the endpoint's location. Although this second point of reasoning may not make sense to everyone, many companies and organizations do not want users dialing out from conference endpoints. Therefore, a conference administrator will dial out on behalf of participants at the time a meeting is set to begin. From the web interface of a CE software-based endpoint, you click the Call Control menu from across the top of the screen. Under the Contacts section, click in the dial box and enter the alias of the destination. Once an alias has been entered, a green Call button will display. Below the Call button there is an option called Show Call Settings. This allows the person dialing to change the Call Rate, which is the requested bandwidth for this call attempt, and the Protocol, which could be SIP, H.323, or H.320. Only the protocol enabled and used on the endpoint will display under the Protocol section. Once the call connects, you can display call details on the web interface by selecting the i button. Beside this button there is also a Hold and Disconnect button. Figure 10-4 illustrates how to dial out from the web interface of a Cisco DX80 endpoint.

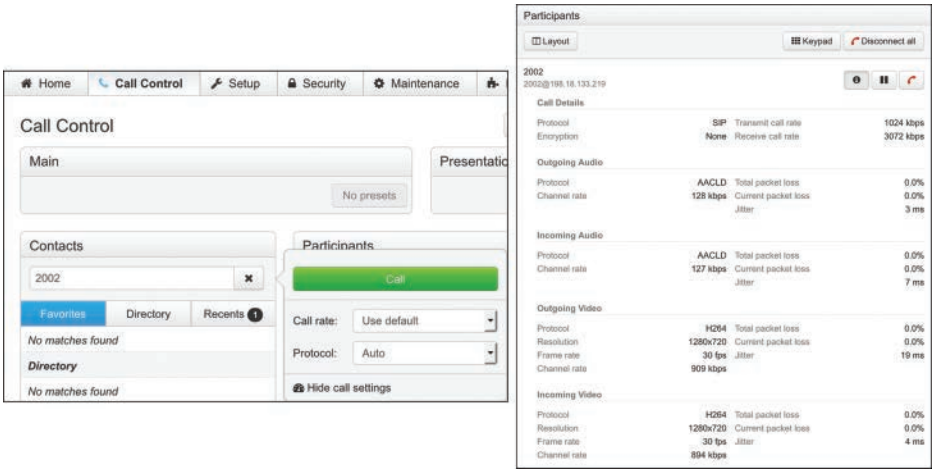

Figure 10-4 *Dialing by Alias from the Web Interface of a DX80*

Dialing out from the CLI requires a simple command. You open an SSH connection to the endpoint and then enter the following command:

```
xCommand dial number: alias
```

The *alias* should be the alias of the destination endpoint, such as 2002. Once the call connects, you can enter the following commands to view call connection status and to disconnect the call:

```
xStatus call
```

```
xCommand call disconnect
```

For more information about API commands on Cisco CE software-based Telepresence endpoints, use the following API reference guide: https://www.cisco.com/c/dam/en/us/td/docs/telepresence/endpoint/ce97/collaboration-endpoint-software-api-reference-guide-ce97.pdf.

Call by Directory

Key Topic

Another means of dialing out from an endpoint is to use the *directory*. A directory is a phonebook that the endpoint has access to but does not necessarily reside on the endpoint itself. Three different types of directories are available to endpoints: a local directory, a corporate directory, and a global directory.

The local directory is a collection of aliases that have been saved directly on the endpoint itself. The endpoint is the sole source of these directory entries; therefore, the endpoint will always contain these entries unless an administrator intentionally deleted them from the system. To add an entry into the local directory, a call attempt must be placed first. After a call has been attempted, whether the call successfully connected or not, locate the dialed alias in the Recents section of the call settings. When you select the entry, a list of options will appear. Select the Add to Local Contact option, and the entry will be added to the Favorites section, which is the local directory. Figure 10-5 illustrates how to add an alias to the local directory.

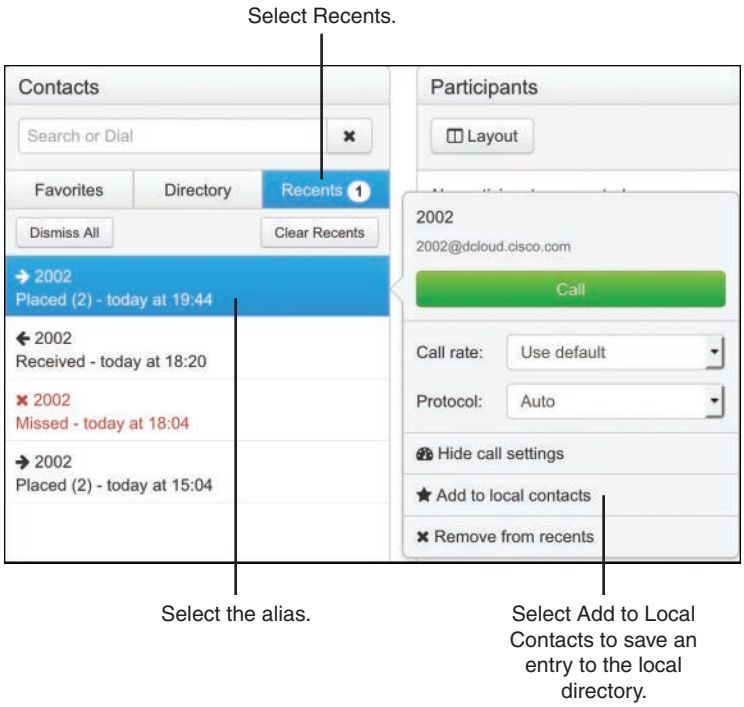

Figure 10-5 *Adding an Alias to the Local Directory*

Key Topic

The corporate directory is the most common directory used in Cisco collaboration. The corporate directory is a phonebook that an endpoint subscribes to when an entry needs to be looked up. These directories do not live on the endpoint itself; rather, they exist solely on the system designed to deliver the aliases when requested. Two systems in a Cisco collaboration solution are capable of delivering a corporate directory to Cisco endpoints. They are the Cisco Unified Communications Manager and the Cisco Telepresence Management Suite, or TMS. Because the focus of this book is on the Cisco Unified Communications Manager, we will not discuss TMS at this time. Endpoints have a finite amount of storage. In fact, a Cisco Collaboration Telepresence endpoint could hold only about 350 directory entries locally. Therefore, the idea of a corporate directory is to create a much larger depository of directory entries on a server built to sustain the greater load and make these entries available to the endpoint as users need the information. The endpoint must have the corporate directory location configured so that it knows where to send the subscription request when enquiring about an entry. When an endpoint registers to the Cisco Unified Communications Manager, the corporate directory address is included in the TFTP Get information sent to the endpoint. You can verify this setting is configured on the endpoint by navigating in the web interface to **Setup > Configuration > Phonebook**. The Type should be set to CUCM, and there may or may not be a URL configured.

The global directory is similar to a corporate directory, in that it originates on a server outside the endpoint itself. However, a global directory is pushed out to the endpoint, so the directory entries live on the endpoint just as the local directory entries live on the endpoint. Obviously, the limitation to this type of directory is the same limitation to the number of directory entries that can exist on a Cisco Telepresence endpoint. However, global

directories can be used with a corporate directory. A limitation to the corporate directory is that if the endpoint loses the connection to the corporate directory, then the endpoint also loses those phonebook entries. However, important aliases can be pushed to an endpoint using a global directory, so that those phonebook entries will never be lost to the endpoint. Unfortunately, the Cisco Unified Communications Manager does not support a global directory. This is a function that only the Cisco TMS can provide.

Regardless of the directory choice, users can be dialed by the directory on the endpoint. You initiate the dialing behaviors the same way as described in the previous section. As you type letters in a name or numbers in a DN, entries will populate the screen. This is the nature of a corporate directory on the endpoint. When you see the name of the person or room you wish to call, select that entry and press the green Call button. Alternatively, you could select the Directory tab and see an alphabetized list of contacts from the combined directories that exist on your endpoint. Figure 10-6 illustrates the use of the directory feature for dialing from a Cisco DX80.

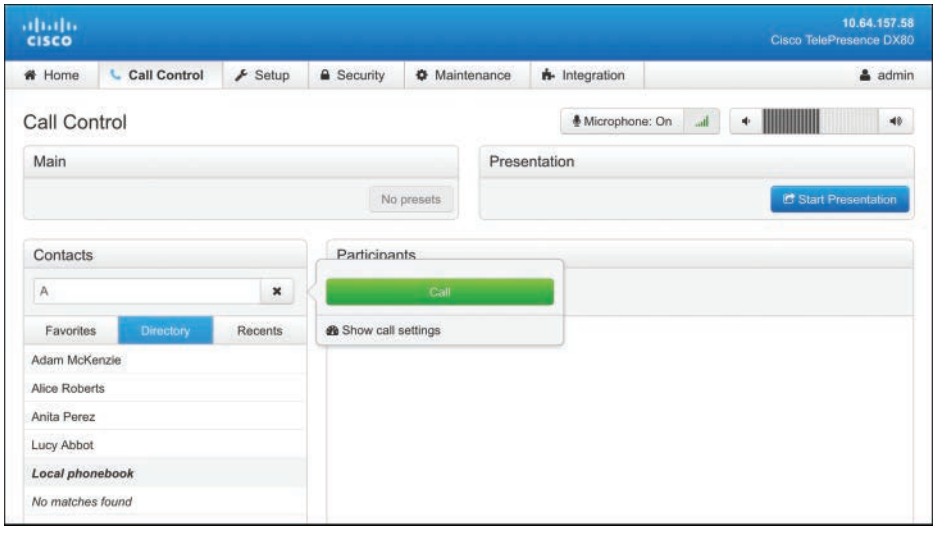

Figure 10-6 *Dialing by Directory on a Cisco DX80*

Multipoint Calling

Up to this point all the calling options discussed involve a point-to-point call. When you are connecting with a multipoint call, the options can be slightly different because different multipoint calling options exist. In a Cisco collaboration solution, three terms related to multipoint calling must be defined. They are *multipoint*, *multisite*, and *multiway*.

Key Topic

Multipoint is an industrywide term used to describe any call that involves three or more participants. Many conferencing products are available to host a multipoint call. Some of these products will be discussed momentarily. *Multisite* is a Cisco-specific term that came from the Tandberg acquisition. Multisite is the option key available on CE software-based endpoints that enables the endpoint to host a multipoint call. When multisite is used, no external conferencing resource is required. However, there are a lot of limitations to using the multisite option over an external conferencing resource. Depending on the endpoint being used for the multisite call, the number of participants allowed to join the call is

limited. All participants in a multisite call must connect at a common bit rate, which is usually the lowest common bit rate among all the participants. Also, a lot more call options are available on a conferencing resource that are not available on the endpoint using multisite to host a call. The third term related to multipoint calling that needs to be defined is *multiway*. Like multisite, *multiway* is a Cisco-specific term that was adopted with the Tandberg merger. Multiway is simply call escalation from a point-to-point call to a multipoint call hosted on a Multipoint Conferencing Unit (MCU). Multiway is a function used by the Cisco VCS through a setting called the Conference Factory. It can be used only with a Cisco Telepresence MCU, which are end-of-life products.

A call escalation function is available through a Cisco Unified Communications Manager as well, although it is called *ad hoc conferencing*, not multiway. Ad hoc conferencing operates in the same capacity as multiway, allowing a point-to-point call to escalate to a multipoint call using an external conferencing resource to host the call. Any current Cisco conferencing resource can be used to support ad hoc calling on the Cisco Unified Communications Manager. Another conferencing option through the Cisco Unified Communications Manager is called rendezvous conferencing. Think of this conferencing type as an always-on conference space that can be joined at any time. Different settings must be configured on the Cisco Unified Communications Manager for ad hoc or rendezvous conferences, but both can exist at the same time. A third conferencing option is scheduled conferences, but these types of meetings cannot be configured from the Cisco Unified Communications Manager. Scheduled conferences must be configured through Cisco TMS.

Basically, three different Multipoint Conferencing Resources are available in a Cisco solution external to the multisite option on Cisco endpoints. Each of these conferencing options can be divided into an on-premises solution, a cloud-based solution, or a hybrid solution between the two. The on-premises solution Cisco offers is called the Cisco Meeting Server, or CMS. Licenses can be added to CMS for Personal Multiparty (PMP) or Shared Multiparty (SMP). PMP licenses are assigned to individual users, and no other party can join their personal space on CMS until the owner of the space has joined. SMP licenses are used to create a shared space into which any user can initiate a call. Therefore, it is recommended to protect SMP licensed meeting spaces on CMS with PINs in order to restrict who can utilize those resources. SMP licenses are also needed for scheduled meetings through TMS. When scheduled meetings are created, no participant can join the meeting until TMS initiates the conference. TMS can also create a private PIN, which participants must enter before joining the call.

Cloud-based meetings are hosted through the Cisco Webex Meeting Center. This is the same powerful tool that has been used for years to allow multipoint conferencing in the cloud. Participants can join via a Webex Meeting client, through a browser, using Webex Teams, or using a unified IP phone or Telepresence endpoint. Physical devices located on-premises, such as the Unified IP phones or Telepresence endpoints, require Expressway Core and Expressway Edge to be configured for firewall traversal to the Webex cloud before meetings can be joined from these devices. All the same tools that have traditionally been used with Webex Meetings are still available, such as high-quality voice and HD video communication, content sharing, polling, and annotation. Additionally, Cisco has added a few more enhancements to Webex Meeting Center, such as cognitive collaboration features.

Bandwidth limitations and network constrictions may negatively impact cloud-hosted meetings through the Webex Meetings solution. Therefore, Cisco has developed a new method to

allow a hybrid service using both Webex Meetings and an on-premises conferencing service called the Video Mesh Node (VMN). The VMN is a virtual server that must be installed on-premises but operates in conjunction with Webex Meeting Center in the cloud. When Webex Meetings are scheduled, on-premises endpoints call into the VMN instead of calling into Webex directly. Then the VMN will send a single stream out to Webex with a composite of all the audio and video of the participants connected. Webex will send a composite of any participants connected directly to the cloud back to the VMN. In this manner, all participants are able to see and hear one another as if they were all connected to the same conferencing unit. The single stream sent between the VMN and Webex limits the bandwidth consumed across the edge network, and reduces the network constraints, creating a better user experience all around. Much more can be said for the hybrid conferencing solution, but that topic will have to be saved for another book.

One Button to Push (OBTP) and Scheduling Conferences

You should now have a basic understanding of the terms *multipoint*, *multisite*, and *multiway*, as well as *ad hoc*, *rendezvous*, and *scheduled conferences*. Furthermore, you should comprehend the differences between an on-premises conferencing solution compared to a cloud-based conferencing solution and a hybrid conferencing solution and be able to identify the different products used for each of these solutions. Bringing the subject back to the topic of dialing behaviors, each of these circumstances around multipoint communication can impact how participants are connected to meetings.

Similar to how participants call one another in a point-to-point call, participants can dial into a multipoint meeting if they know the associated alias. If multisite is used, this will be the alias of the endpoint hosting the call. If CMS or Webex are hosting the call, then the alias of the meeting space must be provided to the attendees before they will be able to dial in.

Multiway and ad hoc calls do not require the meeting ID to be known at all. One of the participants in a point-to-point call will place the second endpoint on hold while calling a third endpoint. With ad hoc, the Conference button must be selected to escalate the call. With multiway, the Join or Connect button must be selected to escalate the call. At that point the call control system, whether it is the Cisco Unified Communications Manager or the Expressway Core, is responsible for transferring the endpoints to the conferencing solution.

Scheduled multipoint calls can utilize any of the conferencing resources: multisite, CMS, or Webex with or without VMN. Many different circumstances influence how dialing behaviors will be impacted, but the following is simply a list of the possible ways through which calls can be connected. TMS can initiate calls from endpoints to the conferencing solution at a scheduled time. The user will not have to dial anything; the endpoint will just connect. TMS can also initiate calls from the conferencing solution out to the endpoint. In this case the user will need to answer the incoming call. A combination of these two options can also be configured, where TMS will automatically dial from some of the endpoints into the meetings and dial from the meeting out to other endpoints. TMS can start the meeting but not connect any participants to the meeting. The participants will need to know the alias to call into the meeting and dial in manually. A final, and more commonly used, option for connecting participants to meetings is the use of a tool call One Button to Push (OBTP).

Key Topic

One Button to Push is a communication option Cisco created prior to the Tandberg acquisition. After the merger was completed, Cisco quickly incorporated this feature into all the Tandberg products acquired. As new products have been developed, Cisco has ensured this highly sought-after feature is continually updated and supported across all of Cisco's communications product portfolio. OBTP is a call connection option that functions only when meetings are scheduled. Although technology systems are very punctual, people are not always as prompt. If a scheduled meeting begins before the participants are ready to communicate, late participants entering the meeting room could be disruptive to others already connected to the meeting. The idea behind OBTP is that a Join button will appear on the endpoint scheduled for a meeting at the time the meeting is scheduled to begin. When the participants in the room of the endpoint are settled and ready to join the meeting, they select that Join button and are connected to the meeting at that time. There is no need know the meeting alias to join, and the endpoint will not be joined to the meeting before the participants are ready. It is no wonder this seamless solution is so widely adopted.

Content Sharing Options

Content sharing is a feature that allows media from a device external to the video endpoint to be displayed on or through the video endpoint. The content sharing feature has changed drastically over the last 30 years. Since video communication is a relatively new technology, the concept of sharing content over great distances had to be conceived and developed over time. The first standard that resembled content sharing was an old ITU standard that was part of the circuit-switched umbrella standard H.320. This standard was called T.150, and it was the *Terminal Equipment and Protocols for Telematic Services*, otherwise known as the *Telewriting Terminal Equipment*. This standard was originally drafted in 1983 and was later revised in 1993.

Key Topic

No other protocol was introduced for content sharing until 1999, when the Olivetti and Oracle Research Lab in Cambridge, UK, released an open-source protocol called virtual network computing, or VNC. VNC is a graphical desktop-sharing system that uses the Remote Frame Buffer (RFB) protocol to remotely control another computer. It transmits the keyboard and mouse events from one computer to another, relaying the graphical-screen updates back in the other direction, over a network. Although VNC can be used for remote desktop control, it was commonly used for content sharing over an IP call, since there was not a viable standard for content sharing yet.

This situation led to two proprietary protocols that were developed and released in 2000, specifically designed for content sharing. Tandberg released the DuoVideo protocol first, and later that year PictureTel released the People+Content protocol. Polycom acquired PictureTel in 2001 and assumed the continued development of People+Content. By 2002 Polycom was offering People+Content for any vendor to use royalty free.

In 2003, the ITU used both the Tandberg and Polycom content sharing protocols to develop a standard that could be used across all H.323 IP calls known as H.239. As SIP was growing in popularity for voice and video calls, a standard had to be drafted for content sharing over SIP as well. In 2006 the IETF drafted RFC 4582 for the Binary Floor Control Protocol (BFCP), which is used for content sharing over SIP. Today, H.239 and BFCP are the two predominant content sharing protocols used industrywide. However, Cisco has been working on other means of sharing content for audio-only participants while in a call. Who knows what other crazy idea might shape the future of content sharing standardizations!

Sharing Content

Content sources could be any devices capable of connecting to the video endpoint. Common content sources could be a dedicated desktop computer; personal laptop computer; smartphone; tablet; VHS, DVD, or Blu-ray player; or even a digital media player (DMP), such as an Apple TV or TV tuner. Content can be shared locally, meaning only in the same room as the video endpoint when not in a call, or content can be shared remotely through the video endpoint to the remote destination at the other end of the call. Most Cisco CE software-based endpoints will display the local content on the screen as soon as the content source is connected. Typically, this is through an HDMI, VGA, or DVI cable, but other connection systems are available although they will not be discussed in this book.

Sharing content while in a call is a manual process that must be executed with intention. This is not a difficult task to impose, but there is no possible way content can be shared by accident. If you were sharing content locally and someone were to call the endpoint you were sharing content through, the content sharing would cease the moment the call was answered. To share content while in a call, you must first select the Show PC button. This will only share the content locally. This feature gives the presenter a chance to see the content that will be presented prior to actually sharing the content. When you are ready to share the content, select the Share button, and the content will be sent in a video stream to the participants at the far end of the call. Tap the Stop Sharing button to end the content sharing session. If you disconnect the call before stopping the content sharing, you will still cease to share content as the call is terminated. If the other party in a call tries to share content while you are currently sharing content, your content sharing session will end, and that person's content will be displayed as a replacement. The same behavior would occur if you were to try sharing content again while the far end was still sharing content.

Using Intelligent Proximity for Content Sharing

Earlier in the chapter, we discussed receiving content using the Proximity application. Any device running the Proximity application can receive content being shared; however, there is also a way to share content using the Proximity application. Content can be shared only if the Proximity application has been installed on a computer running Microsoft Windows or Apple Mac OS. Linux-based operating systems do not support the Proximity application, and smartphones and tablets do not support sharing content, only receiving content.

To install Proximity on your computer, navigate to www.proximity.cisco.com. Choose the appropriate operating system and download the executable file. When the download is complete, launch the installer, agree to the EULA, and begin using Proximity. Assuming the computer being used to share content through Intelligent Proximity meets the previous criteria, there are two ways to initiate the content sharing. Each way is slightly different depending on whether you are sharing content from a Windows PC or a Mac.

When you open the Proximity application on your Mac computer, a Proximity window will appear on the screen. In the app window is a Share Screen button. Tap this button to start sharing your screen. If you have two screens connected to your computer, select the Video System menu from the top left of the screen and choose **Select Screen > <choose screen>**. To stop sharing, bring the app back to the front of the screen and click the Stop Sharing button. When Proximity is installed, a Proximity icon also appears in the top bar on the screen. If you click this bar, you will see all the different menu options available. Go to the Select Screen menu to choose the appropriate display to present from. Then choose the Share

Screen menu option to start sharing content. Choose Stop Sharing when you are finished sharing.

There are two ways to share content from a Microsoft Windows computer as well. Sharing content using the Proximity app on a Windows computer is basically the same as sharing content through the app on a Mac computer. Alternatively, you could press Alt+F12 to start sharing content. Proximity must be running and connected to the endpoint before content can be shared through the application. If another user is sharing content and you start sharing from the Proximity app on your computer, your shared content will override that user's content, and they will stop sharing. The same would be true if someone else started sharing content while you were sharing. Figure 10-7 illustrates the two ways to share content through the Proximity application from a Mac.

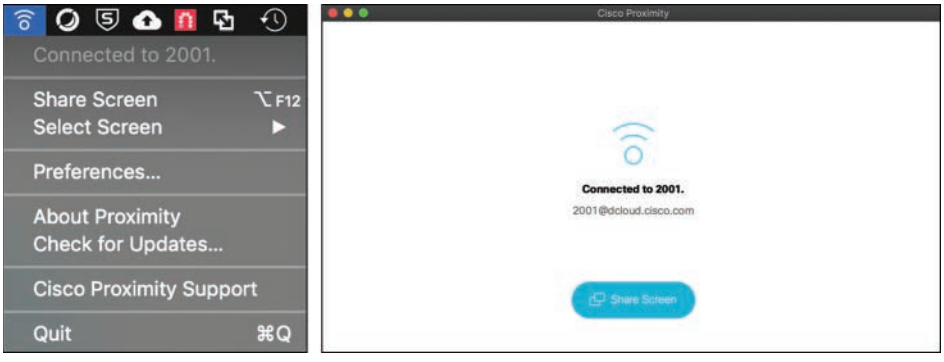

Figure 10-7 *Content Sharing from a Mac Using Proximity*

Other Features

Many calling features available through CE software-based endpoints provide administrators with granular control over the video Telepresence environment. There are more features than what is covered in the sections that follow; however, the following examples are some of the more commonly used settings on these endpoints. None of these settings can be configured from the user interface. They must all be configured from the web interface or the CLI. When an administrator accesses the **Setup > Configuration** menus of a CE software-based endpoint through the web interface, all of the following settings can be configured by selecting the appropriate menu in the left column. The sections that follow cover those menu options in order from the top down. Encryption Mode, AutoAnswer, and Far-End Camera Control are all features listed under **Setup > Configuration > Conference**.

Audio Settings

Some of the audio settings on certain CE software-based endpoints are more extensive than the audio settings on other endpoints. For example, the Cisco Telepresence DX80 endpoint does not have as many audio inputs and outputs as the Cisco Webex Room Kit Pro endpoint, which was built for custom integrator solutions. Therefore, the audio menu options on a Cisco Webex Room Kit Pro are much more extensive than the audio menu options on the Cisco Telepresence DX80. Depending on what endpoints you are planning to deploy and support, you should spend some time in the deployment guides to familiarize yourself with

the different audio options available. However, all CE software-based endpoints share a few common audio settings.

Basic audio settings include audio output settings, such as DefaultVolume, and input settings, such as Input MicrophoneMode, which sets the microphone pickup area, or Microphone Mute Enabled. The DefaultVolume setting comes preset to 50 and can be changed to any value between 0 and 100. The Input MicrophoneMode is preset to Wide so that more participants can be accommodated around a single microphone, but it can be changed to Focused if the room is designed for fewer participants. The Microphone Mute Enabled setting can be set to True or InCallOnly. This setting does not mute the microphone; rather it enables the user to mute the microphone. Therefore, this setting can be configured to always allow the microphone to be muted, or it can prevent the microphone from being muted unless there is an active call in session. Some people like to mute the microphone prior to placing a call; therefore, the default value of True should be left unchanged.

Another section of audio settings that are consistent among all CE software-based endpoints is the SoundAndAlerts settings. The two settings that can be configured in this section are RingTone and RingVolume. RingTone is the audio tone a user will hear when an incoming call is being attempted. The following 12 different ringtones are available on Cisco Telepresence endpoints:

- Sunrise (default value)
- Mischief
- Ripples
- Reflections
- Vibes
- Delight
- Evolve
- Playful
- Ascent
- Calculation
- Mellow
- Ringer

RingVolume is different from DefaultVolume, in that it only impacts how loud the ringing signal will alert. DefaultVolume impacts the actual audio of the far-end participants as their dulcet tones are projected from the system speakers. Much like DefaultVolume, RingVolume defaults at 50 and can be modified between 0 and 100. Figure 10-8 illustrates these audio settings from a Cisco Telepresence DX80 endpoint.

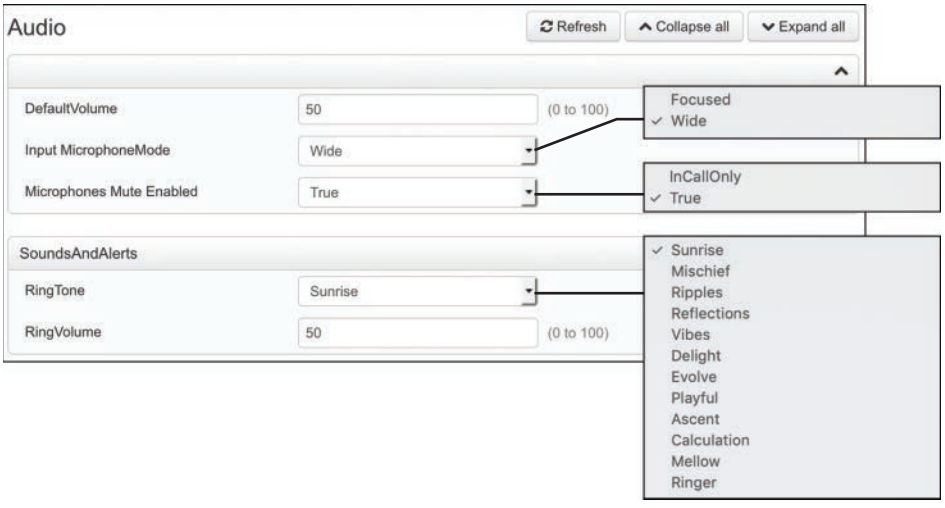

Figure 10-8 *Audio Settings on a Cisco Telepresence DX80*

Key Topic

Encryption Mode

Encryption Mode, which is a setting located under the Conference menu, has to do with call encryption during call setup. Call encryption can occur for both H.323 and SIP calls, but the encryption they use is slightly different. Both use TLS for TCP and UDP packet encryption. However, H.323 supports 56-bit DES encryption or 128-bit AES encryption. SIP supports AES 128- or 256-bit encryption. Also, the means by which H.323 and SIP perform the handshake for secure transmission is different. Because so many different encryption options exist for the two call control protocols, only three settings on CE software-based endpoints revolve around Encryption Mode: BestEffort, On, and Off.

BestEffort is the default setting. The idea behind BestEffort is that the endpoint will try to encrypt when the call is first set up. If the far-end endpoint does not support the same encryption algorithms, or if the far-end endpoint has encryption disabled, then the call will continue as an unencrypted call. When the EncryptionMode is set to on, the call can proceed only if the destination endpoint is also configured to support call encryption and the two algorithms match. Otherwise, the call will fail. The same is true when the Encryption Mode is set to Off. The endpoint will never try to encrypt, so if the far-end endpoint requires encryption, then the call attempt will fail.

AutoAnswer

The second section located under the Conference menu is called AutoAnswer. This feature performs exactly as the name implies. When an endpoint is called and the AutoAnswer feature is enabled, the endpoint will answer the call on its own volition without any human interaction. This capability might seem scary to some people because with AutoAnswer enabled, a user may be caught unaware during a call. You should understand that this feature is disabled by default, so you would have to intentionally enable it before calls could be answered automatically. However, there was a time when AutoAnswer was the norm, and a lot of people had home video systems. People have been caught sleeping, coming out of the shower, and entrenched in many other precarious situations.

Key Topic

The idea behind why this feature exists came out of a time when the conference meeting bridge would dial out to all the scheduled endpoints at the time the meeting was supposed to start. Whether the participants were in attendance at their assigned location or not, the endpoint needed to answer the call to ensure the call connected at the scheduled time. Then Cisco released the OBTP feature discussed previously in this chapter, and just like that, the AutoAnswer feature became much less commonly used. If the AutoAnswer feature is going to be used, some best practice tips need to be taken into consideration. First, you should enable this feature only on meeting room endpoints. Do not enable the feature on personal video endpoints, and definitely do not enable this feature on home video devices. Second, if AutoAnswer is enabled, it is a good idea to enable the Mute on Answer feature. This way, the call will connect as it should, but while people are still settling in the meeting room, the noise will not be disturbing to other participants in the call at remote locations.

The AutoAnswer section consists of three settings: Delay, Mode, and Mute. Delay is measured in seconds and determines the time duration the incoming call should ring before the endpoint will answer the call automatically. The default value is 1, and this setting can be set to any number between 0 and 50. Mode is the setting used to enable the AutoAnswer feature. The default value is Off, and this setting can be changed to On. Mute can be set to Off, which is the default value, or On. When this feature is enabled along with Mode, the endpoint will mute the microphone(s) at the time the call is answered. It is strongly recommended to enable the Mute feature if AutoAnswer is enabled. Figure 10-9 illustrates the AutoAnswer settings in a DX80 endpoint.

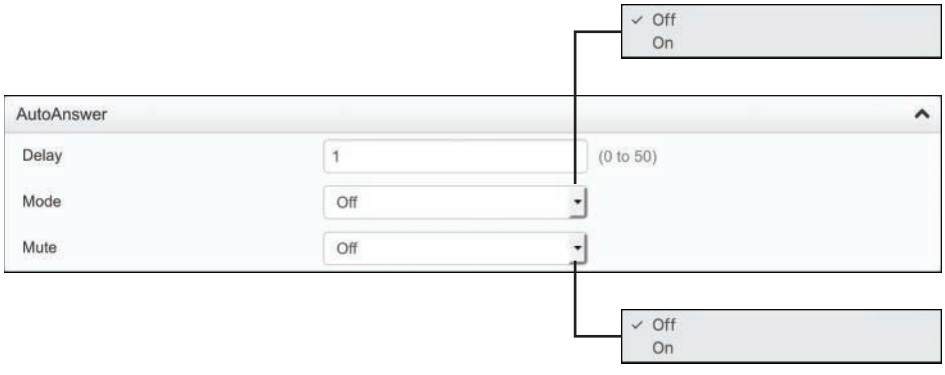

Figure 10-9 *AutoAnswer Feature Settings on a DX80*

Far-End Camera Control (FECC)

Key Topic

A third setting located under the Conference menu that is commonly used is the Far-End Camera Control (FECC) settings. FECC is a setting that allows the camera on your video endpoint to be controlled remotely by another endpoint while in a call. Obviously, camera control can occur only if the camera on the endpoint is an auto PTZ camera, such as the Cisco Precision 60 camera. The camera built into the DX80 endpoint is a manual tilt and zoom camera, with no panning capability. Therefore, enabling FECC on this endpoint would be pointless. The same would be true for the Webex Room Kit series endpoints because the cameras integrated into those endpoints operate differently than the traditional auto PTZ cameras. The capability to control a camera on the far end of a call does not require any setting to be enabled. This capability is built into all Cisco Telepresence endpoints automatically and is always available.

To configure FECC on a Cisco CE software-based endpoint from **Setup > Configuration > Conference**, scroll down to the bottom of the page and look for the FarEndControl section. The Mode setting should be set to On by default. This means that FECC is automatically enabled on the endpoint. To disable FECC, change the Mode to Off. Another setting under the FarEndControl section, called SignalCapability, can be set to On (default) or Off. The standard for FECC is H.224, so this setting enables or disables the use of H.224 for FECC. Because the SignalCapability setting performs the same essential function as Mode, it should be configured using the same setting as Mode for the desired service of FECC. In other words, if Mode is set to On, then SignalCapability should be set to On. If Mode is set to Off, then SignalCapability should be set to Off.

Phonebooks

The earlier "Call by Directory" section discussed the differences between the local directory, corporate directory, and global directory. Local directories are created on the endpoint itself, and global directories are initiated from the Phonebook service; therefore, no settings need to be configured on the endpoint to receive global directories. However, corporate directories require the endpoint to initiate a Subscription message to the phonebook service for each individual phonebook lookup in order to receive entries in reply. This allows the phonebook service to supply only listings based on the characters or numbers that have been entered. You can configure subscription information on the Phonebook menu on Cisco Telepresence endpoints so that the endpoints know where to send the Subscription message. When an endpoint registers to the Cisco Unified Communications Manager, the Phonebook information is automatically provisioned on the endpoint, so no settings need to be configured. However, when the Cisco TMS services are used for Phonebook management, these settings may need to be changed.

There are three settings under the Phonebook > Server 1 section. The ID setting allows a name to be assigned to the phonebook. By default, no name is associated with the phonebook, and the phonebook will continue to function as normal if no name is assigned.

The Type setting allows an administrator to determine from where the phonebook source will come. The default value for this setting is Off, but it can be configured as any of the following:

- **CUCM:** Use this setting if the Cisco Unified Communications Manager is the source of the corporate directory. When the endpoint registers to the Cisco Unified Communications Manager, this setting will change automatically, and no other settings have to be configured on the endpoint.
- **Spark:** Use this setting if the Webex Control Hub is the source of the corporate directory. When the endpoint registers to the Webex Control Hub, this setting will change automatically, and no other settings have to be configured on the endpoint.
- **TMS:** Use this setting if Cisco TMS is the source of the corporate directory.
- **VCS:** Neither the Cisco VCS nor the Expressway can be the source of the corporate directory. However, if the endpoint is registered to one of these products, TMS can be the source for the phonebook. Therefore, Type can be set to VCS or TMS.

The third setting under Phonebook > Server 1 is the URL setting. In some instances, a URL address to the server providing phonebook services is required. This setting will never be

required when **Type > Spark** is selected, will only occasionally be required when Type > CUCM is selected, and will always be required when Type > TMS or Type > VCS is selected. An example of the URL that may be used when the Type is set to CUCM could be as follows:

https://*<cucm-host-name>*:8443/cucm-uds/users

An example of the URL that may be used when the Type is set to TMS or VCS could be as follows:

https://*<tms-host-name>*/tms/public/external/phonebook/phonebookservice.asmx

In both examples, the portion of the URL that is in bold could be the IP address or DNS A-record of the server the name references, TMS or CUCM. Notice that within both URLs specific directories are referenced. You could use these URLs in any production environment by simply replacing the bold portion with your specific server address information. Figure 10-10 illustrates the Phonebook settings on a Cisco Telepresence endpoint.

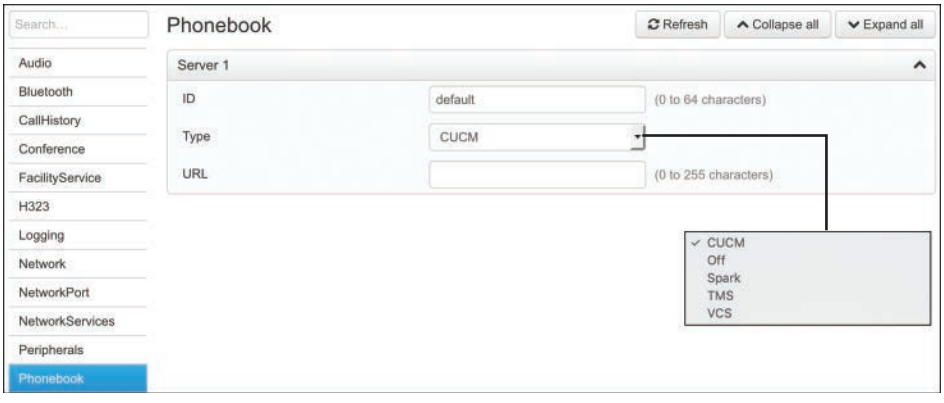

Figure 10-10 *Phonebook Settings on a Cisco Telepresence Endpoint*

Video Settings

The last group of settings that will be discussed in this chapter is the video settings. Under the **Setup > Configuration** section, Video is at the bottom of the menus in the left column. Much the same as audio settings, video settings may vary based on the endpoint's capabilities. Different endpoints have more or fewer video inputs and video outputs. Therefore, there may be more or fewer menu options based on the number of inputs and outputs. However, many common settings are consistent among all Cisco Telepresence endpoints regardless of the endpoint being used. The Video menu is divided into the following six sections:

- Top Main Video section
- DefaultLayoutFamily
- Input
- Output
- Presentation
- Selfview

The top main video section does not have a section title. The settings listed in this section are the main functional settings related to the endpoint. Some of the settings listed here include Active Speaker DefaultPIPPosition. PIP stands for Picture-in-Picture and refers to a video screen layout where a smaller video pane can exist within a larger video pane. The configuration options for this setting pertain to the positioning of the PIP, and they affect the layout only when a call is in session that uses a layout with a PIP overlay. The options for this setting can be any of the following:

- CenterLeft
- CenterRight
- Current
- LowerLeft
- LowerRight
- UpperCenter
- UpperLeft
- UpperRight

The next setting in the first section is the DefaultMainSource. This setting determines which display will be the main video display when a call is in session. Some of the Cisco CE software-based endpoints support two or more displays. In these environments the default main source displays the incoming video of the far-end participants, while the second display is used for content presentation only. When the endpoint is used for local meetings, both displays can support content sharing. The next setting in this section is called Monitors. Similar to the DefaultMainSource setting, the Monitors setting determines how many monitors are currently being supported from this endpoint.

The next section under the Video menu is the DefaultLayoutFamily. These settings pertain only to multipoint calls that use the multisite feature on the endpoint. Five different layouts are supported with the multisite feature, and they can be configured differently for Local and Remote. Local is the layout that participants at this endpoint will view during a call, and Remote is the layout that will be presented to other participants connected to the call. The five layouts supported using the multisite feature are as follows:

- **Auto:** The default layout family, as given by the local layout database, will be used as the remote layout.
- **Equal:** All video participants will have equal-sized panes as long as there is enough space on the screen.
- **Overlay:** The active speaker, or the presentation if present, will be shown in full screen, while the other participants will be small PIPs. Transitions between active speakers are voice switched.
- **Prominent:** The active speaker, or the presentation if present, will be a larger picture, while the other participants will be smaller pictures. Transitions between active speakers are voice switched.

- **Single:** The active speaker, or the presentation if present, will be shown in full screen. The other participants will not be shown at all. Transitions between active speakers are voice switched.

The third Video section is the Input section. These settings are the controls for all of the video input ports on the endpoint itself, and therefore, there could be more or fewer settings listed based on the type of endpoint. Video input settings control devices such as cameras or content devices. Some endpoints support daisychaining multiple cameras together in a single room environment, so an administrator may want to assign a camera ID to each camera in the room and provide a name for the cameras, such as Rear Camera or Whiteboard Camera. Content sources could be an in-room computer, table-connected laptop, DVD or Blu-ray player, document camera, Apple TV, or any other device through which content can be shared. Because more than one content source can be connected to the endpoint at a time, the administrator may be inclined to name the content sources as well. Figure 10-11 illustrates the first three video sections that have been mentioned up to this point.

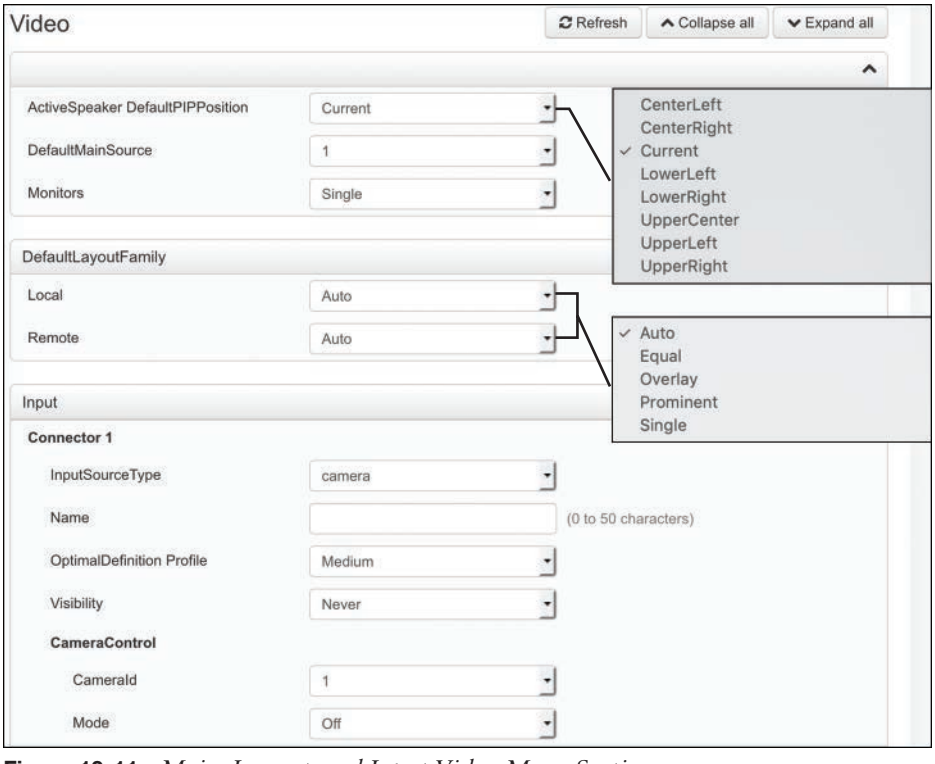

Figure 10-11 *Main, Layout, and Input Video Menu Sections*

Just as with the **Video > Input** section, the **Video > Output** section will have an equal number of Connector settings for configuration as the number of physical video output ports on the endpoint itself. So, the number of settings in this section could be more or fewer depending on the endpoint being used. Additionally, the setting options under each connector may differ based on the type of connection supported. Video outputs are the connector

ports that the displays connect to on the endpoint. Some of the different settings you may encounter in the Output section on the endpoint include the following:

- **Brightness:** This setting defines the brightness level.
- **Resolution:** This setting defines the resolution and refresh rate for the connected screen. When Auto resolution is selected, the endpoint will automatically try to set the optimal resolution based on negotiation with the connected monitor.
- **Whitebalance Level:** This setting defines the camera's white balance level.
- **CEC Mode:** This video output (HDMI) supports Consumer Electronics Control (CEC). When this setting is On, the system will use CEC to set the screen in standby mode when the endpoint itself enters standby. Likewise, the system will wake up the screen when the system itself wakes up from standby mode. Note that different manufacturers use different marketing names for CEC, such as Anynet+ (Samsung); Aquos Link (Sharp); BRAVIA Sync (Sony); HDMI-CEC (Hitachi); Kuro Link (Pioneer); CE-Link and Regza Link (Toshiba); RIHD (Onkyo); HDAVI Control, EZ-Sync, VIERA Link (Panasonic); EasyLink (Philips); and NetCommand for HDMI (Mitsubishi).
- **OverscanLevel:** Some monitors may not present the entire image that they receive. This means that the outer parts of the image that is sent from the video system may be cut off when displayed on the monitor. You can use this setting to instruct the video system not to use the outer part of the available frame. This part might be cut off by the monitor. Both the video and messages onscreen will be scaled in this case.
- **RGBQuantizationRange:** Devices connected to an HDMI output should follow the rules for RGB video quantization range defined in CEA-861. Unfortunately, some devices do not follow the standard, and this configuration may be used to override the settings to get a perfect image with any display. Most HDMI displays expect the full quantization range.
  - **Auto:** The RGB quantization range is automatically selected based on the RGB Quantization Range bits (Q0, Q1) in the AVI infoframe. If no AVI infoframe is available, the RGB quantization range is selected based on the video format according to CEA-861-E.
  - **Full:** Full quantization range. The RGB quantization range includes all code values (0–255). This is defined in CEA-861-E.
  - **Limited:** Limited quantization range. This RGB quantization range excludes some code values at the extremes (16–235). This is defined in CEA-861-E.
- **Location > HorizontalOffset/VerticalOffset:** These two settings are associated with each video output, and they are used to signal the relative position of the displays that are connected to these outputs. HorizontalOffset = 0 and VerticalOffset = 0 indicates that the display is positioned in the center, both horizontally and vertically. A negative horizontal offset indicates that the monitor is left of center, and a positive horizontal offset indicates that the monitor is right of center. A negative vertical offset indicates that the monitor is below center, and a positive vertical offset indicates that the monitor is above center. The magnitude of the offset indicates how far the display is from center (relative to other displays).

The next section is the Presentation section, and two settings here can be configured. The DefaultPIPPosition is the same setting as the ActiveSpeaker DefaultPIPPosition setting discussed earlier. All of the configuration options are the same, except that this setting pertains specifically and exclusively to content being shared. Again, if the layout choice does not use PIP, this setting will not apply. The DefaultSource setting allows an administrator to specify the video input that will act as the primary source for content sharing. This setting will always be a number value and will be based on the video inputs the endpoint supports. For example, the DX80 has only one video input dedicated to content sharing, and that port is hard-coded as video input 2. Therefore, the DefaultSource setting for presentation will always be 2, and this setting cannot be changed. However, a Cisco Webex Room Kit Pro has several video input ports, so this setting would be configurable from the web interface.

The Selfview section pertains to how the self-view window will appear when this feature is enabled. Self-view can be enabled, disabled, and positioned from the user interface, but the more advanced settings must be configured from the web interface or the CLI. The two subsections in the Selfview section are called Default and OnCall. The Default subsection contains the following parameters:

- **FullScreenMode:** This setting defines whether self-view should be shown in full-screen or as a small PIP after a call. The setting takes effect only when self-view is switched on.
  - **Off:** Self-view will be shown as a PiP.
  - **Current:** The size of the self-view picture will be kept unchanged when leaving a call; that is, if it was a PiP during the call, it will remain a PiP after the call; if it was full-screen during the call, it will remain full-screen after the call.
  - **On:** The self-view picture will be shown in full-screen.
- **Mode:** This setting defines whether self-view should be displayed onscreen after a call. The position and size of the self-view window are determined by the **Video > Selfview > Default > PIPPosition** and the **Video > Selfview > Default > FullscreenMode** settings, respectively.
  - **Off:** Self-view is switched off when leaving a call.
  - **Current:** Self-view is left as is; that is, if it was on during the call, it will remain on after the call; if it was off during the call, it will remain off after the call.
  - **On:** Self-view is switched on when leaving a call.
- **OnMonitorRole:** This setting defines which screen output to display the main video source for self-view after a call. The value reflects the monitor roles set for the different outputs in the **Video > Output > Connector > [n] > MonitorRole** setting. The setting applies both when self-view is displayed in full-screen, and when it is displayed as a PIP.
  - **Current:** When leaving a call, the self-view picture will be retained on the same output as it was during the call.

- **First:** The self-view picture will be shown on outputs with the **Video > Output > Connector > [n] > MonitorRole** set to First.
- **Second:** The self-view picture will be shown on outputs with the **Video > Output > Connector > [n] > MonitorRole** set to Second.
- **Third:** The self-view picture will be shown on outputs with the **Video > Output > Connector > [n] > MonitorRole** set to Third.

- **PIPPosition:** This setting defines the position onscreen of the small self-view PIP after a call. The setting takes effect only when self-view is switched on and full-screen view is switched off. All of the configuration options are the same as the ActiveSpeaker DefaultPIPPosition options, except that this setting pertains specifically and exclusively to self-view.

The OnCall subsection contains two parameters: Duration and Mode. **OnCall > Mode** is used to switch on self-view for a short while when setting up a call. The **Video > Selfview > OnCall > Duration** setting determines how long self-view will remain on at the beginning of a call. This setting applies when self-view in general is switched off. Duration defaults to 10 and can be set to any value between 1 and 60. Each numeric value represents one second. If Mode is set to On, then self-view will show momentarily at the beginning of a call. If Mode is set to Off, then self-view will not show at all during any point of the call. Figure 10-12 illustrates the video output settings, presentation settings, and selfview settings on a Cisco Telepresence endpoint.

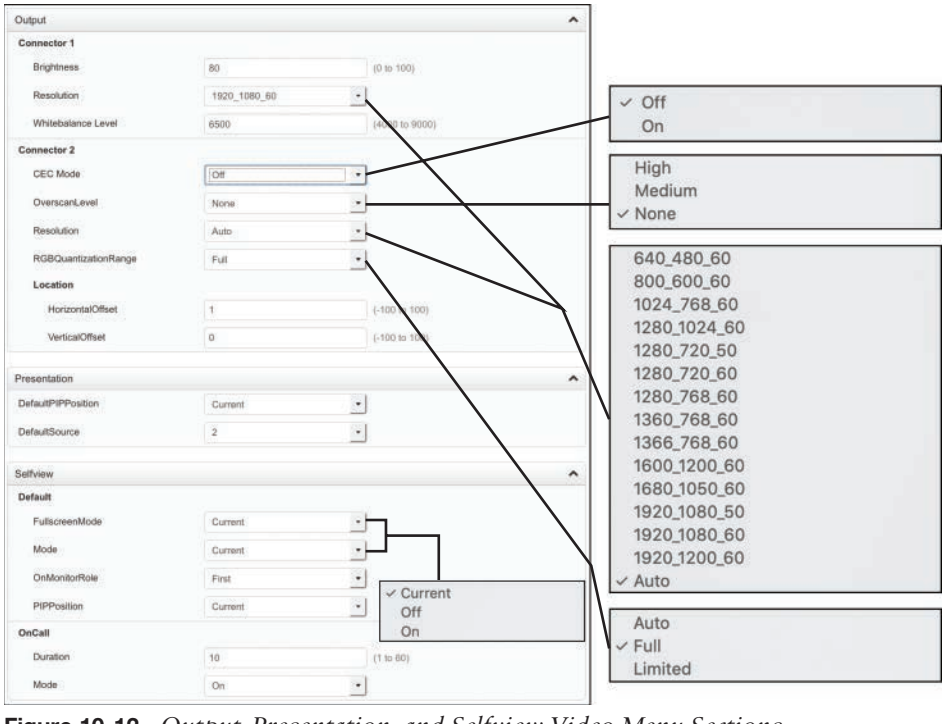

Figure 10-12 *Output, Presentation, and Selfview Video Menu Sections*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 10-2 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 10-2 Key Topics for Chapter 10

Key Topic Element	Description	Page Number
List	Three Ways to Place or Answer a Call on CE Endpoint	235
Paragraph	Three Control Mechanisms for CE Endpoints	236
Paragraph	Calling from CE Endpoint Using the On-Screen Display (OSD)	237
Paragraph	Reasons for Dialing from Web Interface or CLI	238
Paragraph	Three Types of Directories	239
Paragraph	Two Systems That Deliver Corporate Directories	240
Paragraph	Define Multipoint, Multisite, and Multiway	241
Paragraph	Explain CMS as the On-Premises Conferencing Solution	242
Paragraph	Explain Webex Meeting Center as the Cloud-Based Conferencing Solution	242
Paragraph	Explain VMN as the Hybrid Conferencing Solution	242
Paragraph	Call Connections Through Scheduled Multipoint Meetings	243
Paragraph	OBTP	244
Paragraph	Define VNC	244
Paragraph	How to Share Content while in a Call	245
Paragraph	How to Share Content Using Proximity App	245
Section	Encryption Mode	248
Paragraph	AutoAnswer Best Practices	249
Paragraph	Explain FECC	249
Paragraph	Corporate Directory Settings Explained	250
List	Active Speaker DefaultPIPPosition Settings	252
List	Five Layouts Supported by Multisite	252

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Ad Hoc, BFCP, CEC, CLI, CMS, Corporate Directory, DN, DuoVideo, E.164 Alias, FECC, Global Directory, H.224, H.239, Local Directory, Multipoint, Multisite, Multiway, People+Content, PIP, PMP, Precision Camera, PTZ, Rendezvous Conferencing, RGBQuantizationRange, SMP, T.150, TMS, URI, VNC

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the five control mechanisms that allow users and administrators to interact with Cisco CE software-based endpoints.
2. List five devices that can be used as a content resource on Cisco CE software-based endpoints.
3. What are the four Type options for a phonebook source on the Cisco Telepresence endpoint?
4. List and describe the five multisite layouts supported on the Cisco Telepresence endpoint.

This page intentionally left blank

CHAPTER 12

Cisco Core Network Components

This chapter covers the following topics:

LAN, WAN, and Wireless LAN: This topic will discuss the various layers of an enterprise network to cover the LAN, WAN, and wireless LAN as they relate to collaboration.

Gateways: This topic will introduce various types of gateways, both old and new, with special emphasis on IOS gateway services through the Cisco ISR routers and the features they support that impact the Cisco collaboration solution.

The network is the most important aspect to any business today because that is what connects people together across the world. In an episode of a funny British sitcom called *The IT Crowd*, the IT department's nontechnical boss was convinced that the "Internet" was a single black box. She was planning to use it as a prop in a presentation she was to present in hopes of wowing the audience. As I hope you are aware, the Internet, or any network for that matter, is not a single device. It is a combination of devices, software, and protocols that are the culmination of years of development by multiple vendors to become what it is today. It is a living entity in that it continues to grow and change as time passes, and it will continue to grow as long as people have a need for it. The vast embodiment of the Cisco core network components is too colossal to cover in one chapter. However, this chapter will introduce many concepts concerning the Cisco core network components as they pertain to collaboration. Topics discussed in this chapter include the following:

- LAN, WAN, and Wireless LAN
  - LAN (Access Layer, Distribution Layer, Core Layer)
  - WAN Aggregation Design
  - Wireless LAN (Basic Configuration and Design, High Availability, Capacity Planning, Design Considerations)
- Gateways
  - ISR, ASR, and IOS Software Comparisons
  - ISR Products Explained

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

1.3 Configure these network components to support Cisco Collaboration solutions

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 12-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 12-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
LAN, WAN, and Wireless LAN	1–6
IOS Gateways	7–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is an element configured on a Layer 2 switch?
   a. VLAN
   b. QoS Policing
   c. EIGRP
   d. OSPF
2. Which of the following does the IEEE 802.1d standard define?
   a. Layer 2 QoS
   b. Spanning Tree Protocol
   c. Rapid Spanning Tree Protocol
   d. Multiple Instance Spanning Tree Protocol
3. Which of the following protocols protects data traffic from a failed router or circuit, while also allowing packet load sharing between a group of redundant routers?
   a. HSRP
   b. ARP
   c. VRRP
   d. GLBP
4. Which of the following is a "best effort" bandwidth option for network connections?
   a. DSL
   b. ATM
   c. Leased Lines
   d. Frame Relay

5. Which of the following is the wireless router to which a device would connect for network access?
   a. WLC
   b. LAN
   c. LWAP
   d. LWAPP
6. What is the recommended cell boundary overlap to provide high availability in a wireless network?
   a. There should not be any cell boundary overlap.
   b. 20 percent
   c. 50 percent
   d. 100 percent
7. Which of the following is an analog station gateway?
   a. E&M
   b. FXO
   c. PRI
   d. FXS
8. Which of the following routers supports the IOS XE software?
   a. ISRv
   b. ISR 800 Series
   c. ISR 2900 Series
   d. ASR 9000 Series
9. Which of the following Cisco routers should an engineer choose for a customer who needs to support 1000 users at a specific location?
   a. ISR 800 Series
   b. ISR 1000 Series
   c. ISR 2900 Series
   d. ISR 4000 Series
10. Which Cisco router should be used for machine-to-machine and device-to-device deployments such as ATMs, point-of-sale kiosks, and vending machines (fixed platform)?
    a. ISR 800M
    b. ISR 810
    c. ISR 860
    d. ISR 890

Foundation Topics

LAN, WAN, and Wireless LAN

The most foundational components of any corporate communication solution are the network infrastructure components. One of the reasons Cisco is the leader in the collaboration market is that only Cisco can offer an end-to-end solution to its customers. Of course, providing superior collaboration products with extensive capabilities and beautiful designs helps contribute to the company's ability to hold that leading position. The purpose of this chapter is not to provide an extensive education on these network components and how to configure them. However, there is such a close dependency on Cisco collaboration products and the network that it is essential to have an understanding of the network to a certain level. To provide a deeper understanding of basic networking components, Cisco offers the CCNP Enterprise certification courses, which can also be studied using the Cisco Press material. These courses and the material will provide a more thorough understanding of what each network component is and how to configure it. For the purposes of this book, we will examine the foundational network infrastructure components because they relate directly to the Cisco preferred architecture for enterprise collaboration.

A network can be defined as a group or system of interconnected things. A local-area network (LAN) is a network of devices within a limited area. This could be a business office, school, or campus. A home network is a LAN that might interconnect computers, smartphones, tablets, smart TVs, printers, and other media devices. A wide-area network (WAN) is a network of devices within a wider area than the LAN. Imagine two LAN offices, one located in New York City and the other in Washington DC, but devices within each of these locations can communicate with one another as if they were within the same LAN. This is a WAN. Then there is the wireless local-area network (WLAN) or wireless LAN. Because different technologies exist within wireless technology as compared to a physical LAN, this type of network must be categorized independently. Most home networks use some sort of consumer wireless router, but the technology behind a commercial wireless LAN goes far beyond what is available to the everyday consumer.

Now that we've defined the different types of networks, let's examine some of the physical network components and how they might be used. The only network component needed to set up a LAN is a switch. A basic switch is a device with multiple physical ports to which multiple devices can be connected using an Ethernet cable so that communication between these devices can be established. Switches operate on Layer 2 of the OSI model. Cisco switches have a higher level of intelligence than a basic switch, so a network administrator can configure parameters that control how traffic flows through these switches. In fact, some of the switch models that Cisco offers can be configured with Layer 2 and Layer 3 capabilities. As switches pertain to collaboration, several configuration elements can be configured, including virtual LANs (VLANs), Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED), and quality of service (QoS), to name a few. As essential as a switch is within a network, you cannot access the public Internet or establish a WAN without the next network component—the router.

Routers are Layer 3 components of the OSI model and provide communication into and out of the LAN. Many services can be provided through a router. Routers are often configured to offer Dynamic Host Configuration Protocol (DHCP) services to devices on the network. DHCP provides devices with an IP address, subnet mask, and default gateway (also known as the default router) address at a minimum. It can also provide Domain Name System (DNS) address(es) and Trivial File Transfer Protocol (TFTP) server address(es). Devices connected to a switch know how to route traffic to a router using the Default Gateway Address, which is the internal IP address of the router. TFTP addresses can be provided using Option 66 or the Cisco proprietary Option 150.

Because a LAN operates using private IP addresses, which are not publicly routable, the router can masquerade these private IP addresses with a public IP address so that traffic can be routed out to the public Internet. The service used to masquerade these addresses is known as Network Address Translation (NAT) or Port Address Translation (PAT). Many devices on the network require the timing to be synchronized for services to operate properly, such as endpoints joining a scheduled meeting. Therefore, these networked devices rely on Network Time Protocol (NTP) to provide timestamp information. When the edge router is configured as the NTP authority, it can provide a Stratum 2+ NTP reference to these devices.

Firewall software is typically also available on routers. Some companies opt for a firewall server in lieu of, or in addition to, the firewall software available on the router. Firewalls protect nodes within your network from malicious attacks coming from outside your network. Think of firewalls as a first line of defense. Other defensive control mechanisms available on the router are access control lists (ACLs). ACLs are lists of protocols and port numbers that are allowed or not allowed to flow through a router. ACLs can be applied on an inbound or outbound (physical) port on the router. For example, an ACL could be configured on a router that allows TLS traffic on port 5061 but rejects TCP traffic on port 5060. The idea here is to allow encrypted SIP signaling and reject nonencrypted SIP signaling. ACLs can also be used as a stateless inspection of the traffic, which differentiates ACLs from firewalls.

Routers offer many more features, but one last feature worth mentioning is QoS. As mentioned previously with Layer 2 switches, QoS can be applied at Layer 3 on the router. In fact, Layer 3 QoS is even more critical than Layer 2 QoS because this is typically where you will find congestion in a network. Ideally, you want to mark packets as close to the source as possible; therefore, Layer 2 QoS is designed to mark packets early in the routing process. Layer 3 QoS prioritizes how traffic will flow during these high-congestion times. On the router, you need to convert Layer 2 QoS marking to Layer 3 QoS marking. Other Layer 3 tools for QoS include shaping, policing, queuing, and QoS type. Cisco has a lot of information available on QoS, and it is essential to research and understand QoS to work effectively in collaboration as a technician or engineer. QoS will be covered in a little more depth in the next chapter, although QoS is a very deep topic that could fill volumes of books all on its own.

Among Cisco routers, one stands out above the rest: the Cisco Integrated Services Router (ISR). The ISR has all the same services that other routers have, as mentioned previously. However, additional services and modules can be added to the ISR. Some of the collaboration services available on an ISR include Cisco Unified Communications Manager Express (CUCME), Survivability Remote Site Telephony (SRST), Cisco Unified Border Element (CUBE), and Cisco Unity Express (CUE). Modules that are supported in select models of ISRs include PRI cards (E1 and T1), FXS and FXO cards, and PVDM cards.

Key Topic

Collectively, Cisco is known as "The Network People" for a reason. It offers the best proven network products available on the market. Over 80 percent of the public Internet space consists of Cisco networking products. And the company is continually releasing software advancements on its network products that push the edge of what is possible. One such software advancement that provides added intelligence to your network is known as Medianet. Cisco Medianet can be defined as an end-to-end architecture for a network comprising advanced, intelligent technologies and devices in a platform optimized for the delivery of rich-media experiences. Medianet allows network devices to be media-aware so that they can detect and optimize different media and application types to deliver the best experience to the user, such as Telepresence, video surveillance, desktop collaboration, and streaming media, to name a few. Medianet also makes networking devices endpoint-aware to automatically detect and configure media endpoints. Finally, Medianet makes networking equipment network-aware so that it can detect and respond to changes in device, connection, and service availability. With the increasing adoption of new video and rich-media applications, Medianet technologies become critically important to address challenges associated with the transmission of video, voice, and data over the network, including ensuring predictability, performance, quality, and security. By accelerating deployment of applications, minimizing complexity and ongoing operational costs, increasing visibility into the network, and helping to scale the infrastructure for the best quality of experience, Medianet technologies help address these challenges. Check out the Cisco Medianet Data Sheet at Cisco.com for more information on Medianet.

Depending on the environment being configured, there might be a need for Layer 2 switches, Layer 3 switches, and Layer 3 routers. A large enterprise network can be divided into four layers at the central office and two layers at a branch office. The central office can be divided into the Access layer, Distribution layer, Core layer, and the WAN Aggregation layer. The Access layer is typically made up of Layer 2 switches. The Distribution layer is typically made up of Layer 3 switches. The Core layer can be made up of Layer 3 switches or Layer 3 routers. The WAN Aggregation layer is always a Layer 3 router. The branch office typically utilizes a branch router and a branch switch to form the two layers needed for communication. Figure 12-1 illustrates how a typical enterprise network infrastructure is designed.

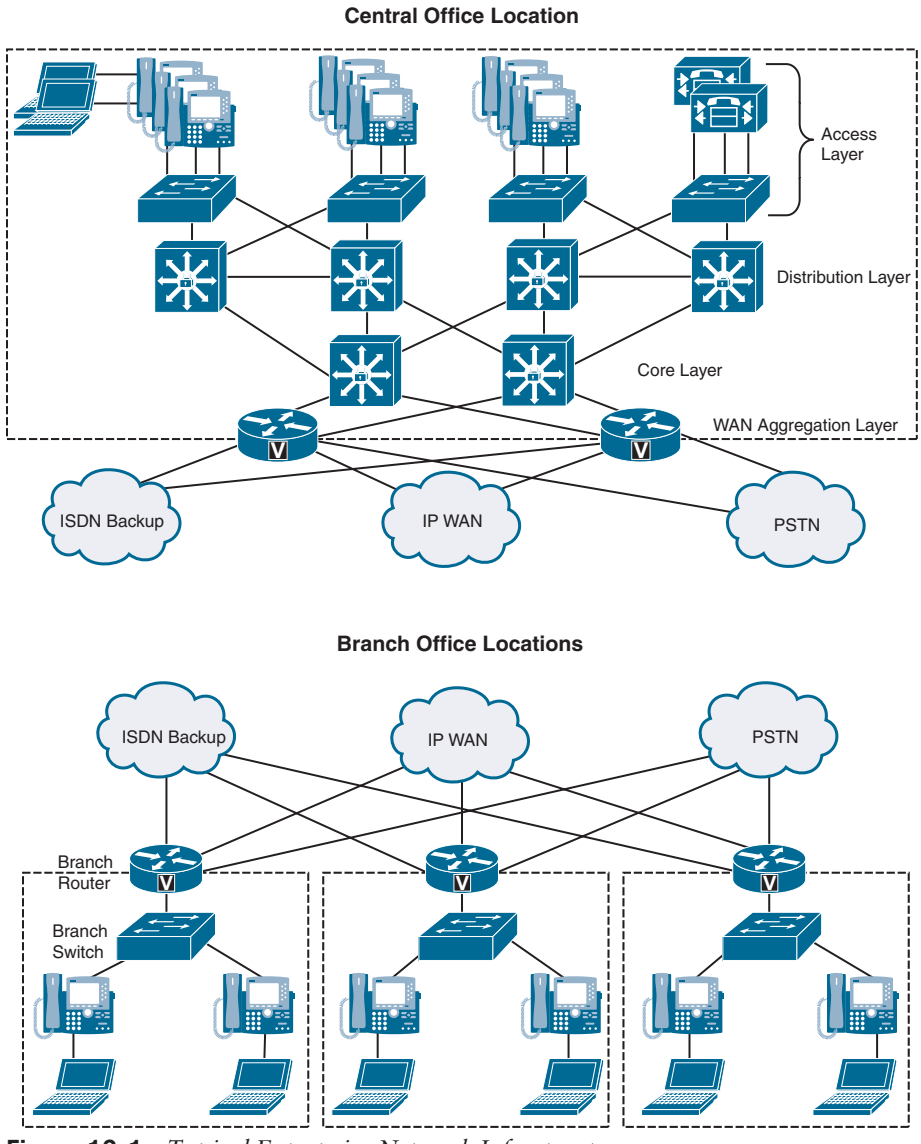

Figure 12-1 *Typical Enterprise Network Infrastructure*

LAN

A properly designed LAN will take into consideration the needs for high availability and quality of service. This will account for the Access layer, Distribution layer, and Core layer of the typical enterprise network infrastructure. The Access layer offers in-line power to the phones, multiple queue support, 802.1p and 802.1q, and fast link convergence. The Distribution and Core switches offer multiple queue support as well as 802.1p and 802.1q, the same as the Access layer, along with traffic classification and reclassification. An IEEE protocol, 802.1p refers to the support of QoS on Layer 2 switches. Also an IEEE protocol, 802.1Q refers to the support of virtual LANs on Layer 2 switches.

Access Layer

Key Topic

High availability can be configured on the Access layer by using the Spanning Tree Protocol (STP). STP is a Layer 2 protocol that runs on switches and is specified by the IEEE standard 802.1d. The purpose of STP is to prevent loops when configuring redundant paths within the network. In Figure 12-1, observe that each switch is connected to two or more other switches, so that if one path fails, there is a redundant path to the destination. On the switch ports, STP can be configured to block traffic on one port and forward traffic on the other port. In the event that the forwarding port can no longer send and receive communications, the state of the blocking port will change to allow the data to flow along the alternate path. This ensures that there is an alternate path for routing traffic but eliminates the chance of a loop occurring with two open ports. Different flavors of STP can be used, and each one requires different timing for convergence. Therefore, it is recommended that the same version of STP be used within a single environment. Some of the other Spanning Tree Protocols that exist include IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) and IEEE 802.1s Multiple Instance Spanning Tree Protocol (MISTP). These two can converge at much higher rates than the traditional STP.

Key Topic

A virtual local-area network, or VLAN, is another essential part of the Access layer switch, and it should be configured prior to setting up STP. A VLAN is a logical grouping of devices connected to the switch that allows data traffic in the network to be decoupled for access control and prioritization. VLANs can be used to group devices in several ways, such as device type or department. For example, there may be a server where accounting software resides. The accounting team needs access to this software, but the sales team does not need access. The accounting team also needs to be able to communicate with the sales team for handling expense reports. In this scenario, three VLANs can be used to control who can access the accounting software. The server where the software resides can be placed in VLAN110, the accounting team computers can be placed in VLAN120, and the sales team computers can be placed in VLAN130. Then the network administrator can create connections from VLAN120 to VLAN110, and from VLAN120 to VLAN130. This will allow the accounting team to access the accounting software and communicate with the sales team. However, the sales team will be restricted from accessing the accounting software.

In a Cisco collaboration environment, VLANs are essential to implementing proper QoS. At least two VLANs should be created in this environment: a data VLAN and a voice and video VLAN, which is typically signified as VVID (Voice VLAN ID). Voice and video data can traverse the same VLAN, even though they typically experience different QoS markings. The reason these two VLANs need to be created is that Cisco phones have a NIC connecting the phone to a switch and a computer NIC connecting a computer to the phone. The phone and the computer require different QoS treatment and therefore should belong in different VLANs. This is where the configuration gets really interesting. If the phone is connected to the switch and the computer is connected to the phone, how can they possibly be decoupled into different VLANs? On the port at the switch where the phone is connected, both the Data VLAN and the VVID can be assigned. When the phone boots up, CDP or LLDP-MED can be used to discover both of these VLANs. There is a third virtual NIC in the phone that exists to monitor egress traffic and determine which VLAN should be used. Data traffic sourced from the computer will use the Data VLAN. Voice and video data from the phone will use the VVID. If content is

being shared from the computer during a video call, then the VVID will be used for that particular data sourced from the computer.

One final offering that can be utilized at the Access layer needs to be mentioned here: inline power. Inline power, or Power over Ethernet (PoE), has already been discussed at great length. For a review of the information coving PoE, refer to Chapter 9, "Endpoint Registration." Table 12-2 outlines the different types of PoE and the maximum power available. Some examples of Cisco switches that support the different types of PoE can also be found in Table 12-2.

Key Topic

Table 12-2 PoE Types and Supported Power

PoE Type	PoE Power Capabilities	Example Switches
Pre-Standard Inline Power	6.3 Watts power	3550-24 or 48 ports
802.3af PoE	15.4 Watts power (Type 1)	3560-24 ports or 3670-48 ports
802.3at PoE	30 Watts power (Type 2)	2960-24 is Type 1 or Type 2
	60 Watts power (Type 3)	4500 supports all types of PoE
	100 Watts power (Type 4)	9000 supports all types of PoE

Distribution Layer

The Distribution layer switches can offer the same multiple queue support, 802.1p, 802.1q, and fast link convergence as the Access layer. However, the focus of the Distribution layer should be to offer Layer 3 routing, load balancing, and fault tolerance. These Distribution layer switches are the bridge between Layer 2 and Layer 3 of the enterprise network.

Key Topic

The Distribution layer switch can often serve as the Layer 3 default gateway for the Layer 2 devices. Should the Distribution layer switch fail, then many devices could lose communication across the network. Cisco initially released the Hot Standby Router Protocol (HSRP) to provide a fault-tolerant default gateway. The IETF developed a similar protocol called the Virtual Router Redundancy Protocol (VRRP) with RFC 5798. Although these two protocols are similar in nature and resolve the gateway redundancy issue, they are not compatible protocols and they each have some limitations. Cisco overcame these limitations when it released another protocol called the Gateway Load Balancing Protocol (GLBP). This protocol protects data traffic from a failed router or circuit, while also allowing packet load sharing between a group of redundant routers.

Endpoints use the Address Resolution Protocol (ARP) to learn the physical MAC address of their default gateway. With HSRP, a single virtual MAC address is provided to these endpoints. With GLBP, two virtual MAC addresses can be provided to the endpoints—one from the primary gateway and one from a peer gateway—which are distributed using round-robin technique.

Another way to ensure fast convergence, load balancing, and fault tolerance on the Distribution layer is to use Layer 3 routing protocols such as OSPF or EIGRP. You can use parameters such as routing protocol timers, path or link costs, and address summaries to optimize and control convergence times as well as to distribute traffic across multiple paths and devices. Cisco also recommends using the **passive-interface** command to prevent routing

neighbor adjacencies via the access layer. These adjacencies are typically unnecessary, and they create extra CPU overhead and increased memory utilization because the routing protocol keeps track of them. By using the **passive-interface** command on all interfaces facing the access layer, you prevent routing updates from being sent out on these interfaces, and therefore, neighbor adjacencies are not formed.

Core Layer

The Core layer operates entirely in Layer 3 of the enterprise network. This layer can consist of Layer 3 switches or routers. The purpose of the Core layer is to provide redundancy between different Distribution switches. In the event of network outages, the Core layer can redirect traffic along a more stable path. The types of redundancy that need to be provided at the Core layer include Layer 1 link paths, redundant devices, and redundant device subsystems, such as power supplies and module cards. The Cisco Catalyst switches with Virtual Switching System (VSS) provide a method to ensure redundancy in all of these areas by pooling together two Catalyst supervisor engines to act as one. This is why Cisco recommends using a Layer 3 switch at the Core layer. Routing protocols at the Core layer should again be configured and optimized for path redundancy and fast convergence. There should be no STP in the core because network connectivity should be routed at Layer 3. Finally, each link between the core and distribution devices should belong to its own VLAN or subnet and be configured using a 30-bit subnet mask.

In the campus LAN, bandwidth provisioning recommendations can be summarized by the motto "overprovision and undersubscribe." This motto implies careful planning of the LAN infrastructure so that the available bandwidth is always considerably higher than the load and there is no steady-state congestion over the LAN links. The addition of voice traffic onto a converged network does not represent a significant increase in overall network traffic load; the bandwidth provisioning is still driven by the demands of the data traffic requirements. The design goal is to avoid extensive data traffic congestion on any link that will be traversed by telephony signaling or media flows. Contrasting the bandwidth requirements of a single G.711 voice call (approximately 86 kbps) to the raw bandwidth of a Fast Ethernet link (100 Mbps) indicates that voice is not a source of traffic that causes network congestion in the LAN, but rather it is a traffic flow to be protected from LAN congestion.

WAN

The next layer in the enterprise network solution is the WAN Aggregation layer. These are the edge routers that allow different locations to communicate over the public Internet. There are general design considerations for deploying a WAN, as well as specific bandwidth considerations. There are also QoS tools at the WAN that will need careful design because the WAN presents the greatest potential for congestion. QoS topics will be discussed in the next chapter.

The Cisco recommendation when designing the WAN Aggregation layer is to establish multiple links for redundancy in case one of the links should fail. WAN designs include hub and spoke, full mesh, and partial mesh. The hub-and-spoke design contains one central "hub" router connected to multiple "spoke" routers. Each spoke is one hop away from the hub and two hops away from other spokes. Alternatively, a full mesh or partial mesh design could be implemented. In this type of design, multiple WAN links are established between locations so that each location has at least two WAN links, each link to a different router.

Redundancy should be built into the WAN Aggregation layer using multiple links. This will ensure connectivity in the event one link fails, and additional bandwidth can be provisioned for load-balancing network traffic. Another design consideration for WAN Aggregation is noncentralized resources so that these services are available to all locations in the event of a WAN failure. These resources include media resources, DHCP servers, voice gateways, and call-processing applications. Earlier in this chapter, we discussed some of the call-processing applications such as SRST or Cisco Unified Communications Manager Express. If you are unfamiliar with these call-processing applications, you may want to research them on your own because they are outside the scope of this book.

The two types of bandwidth options to choose from when designing the WAN Aggregation layer are best-effort bandwidth and guaranteed bandwidth. Examples of best-effort bandwidth include the public Internet, DSL, cable, satellite, and wireless. With best effort there is no QoS, and bandwidth availability is on a first-come basis. Although these types of links are suitable for home offices or commuters, voice and video traffic will suffer. Therefore, Cisco recommends that you do not use best effort for voice-enabled networks that require enterprise-class voice services and quality.

Alternatively, you can choose from available guaranteed bandwidth link options. Leased Lines, Frame Relay, Asynchronous Transfer Mode (ATM), and ATM/Frame-Relay Service Interworking are older technologies that use dedicated circuits through a telephony service provider. These were the best options available to corporations prior to the introduction of broadband and high-speed Internet. Today, these technologies are very expensive and offer lower bandwidth rates compared to other packet-switched solutions available.

Some other guaranteed bandwidth link options available include Multiprotocol Label Switching (MPLS), Cisco Voice and Video Enabled IP Security Virtual Private Network (IPSec V3PN), and Dynamic Multipoint Virtual Private Network (DMVPN). MPLS is a transport protocol that uses "labels" to route traffic rather than network addresses. Packets are forwarded based on the content of the label, so deciphering between voice, video, and data is simple. It is protocol agnostic, so it will function in circuit-switched or packet-switched networks. IPSec V3PN integrates three core Cisco technologies: IP Telephony, QoS, and IPSec VPN. This results in an end-to-end VPN service that can guarantee the delivery of latency-sensitive voice and video communications. DMVPN is a solution that provides an alternative to the complicated administrative setup and maintenance that comes with establishing a mesh network. Initially, the DMVPN is set up as a hub-and-spoke network. Once communication is established between the spokes and the hub, each spoke will dynamically discover each of the other spokes and establish a tunnel between one another. This will reduce the amount of traffic at the hub, preserving bandwidth and processing capacity limits. For information on the deployment of multisite DMVPN WANs with centralized call processing, refer to "Cisco Unified Communications Voice over Spoke-to-Spoke DMVPN Test Results and Recommendations," available at https://www.cisco.com/go/designzone. Figure 12-2 illustrates how a DMVPN network operates.

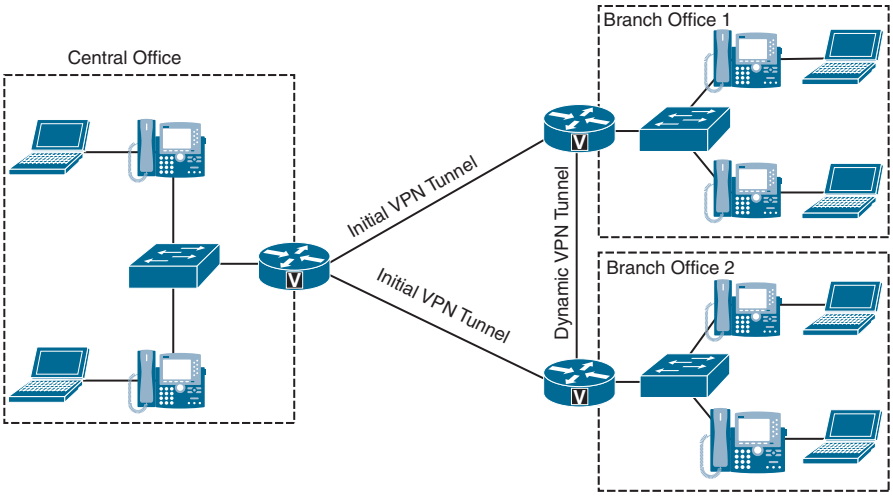

Figure 12-2 *DMVPN Network Operations*

WLAN

The wireless local-area network (WLAN) is a whole infrastructure deployment that hangs off the LAN but requires a unique set of configurations. WLAN infrastructure design becomes important when collaboration endpoints are added to the WLAN portions of a converged network. With the introduction of Cisco Unified Wireless endpoints, voice and video traffic has moved onto the WLAN and is now converged with the existing data traffic there. Just as with wired LAN and wired WAN infrastructure, the addition of voice and video in the WLAN requires following basic configuration and design best practices for deploying a highly available network. In addition, proper WLAN infrastructure design requires understanding and deploying QoS on the wireless network to ensure end-to-end voice and video quality on the entire network.

Basic Configuration and Design

Wireless IP network architectures enable IP telephony to deliver enterprise mobility by providing on-premises roaming communications to the users with wireless IP telephony devices. Wireless IP telephony and wireless IP video telephony are extensions of their wired counterparts, which leverage the same call elements. Additionally, wireless IP telephony and IP video telephony take advantage of wireless 802.11-enabled media, thus providing a cordless IP voice and video experience. The cordless experience is achieved by leveraging the wireless network infrastructure elements for the transmission and reception of the control and media packets. The architecture for voice and video over wireless LAN includes the following basic elements:

- Wireless access points
- Wireless LAN controllers
- Authentication database
- Supporting wired network
- Wireless collaboration endpoints
- Wired call elements

The wireless access points enable wireless devices to communicate with wired network elements. In the case of a Cisco Collaboration environment, these wireless devices include all the UC voice and video endpoints that support wireless communications. Access points function as adapters between the wired and wireless world, creating an entryway between these two media components. Cisco access points can be managed by a wireless LAN controller (WLC), or they can function in autonomous mode. When the access points are managed by a WLC, they are referred to as lightweight access points (LWAPs), and in this mode, they use the Lightweight Access Point Protocol (LWAPP) or Control and Provisioning of Wireless Access Points (CAPWAP) protocol, depending on the controller version, when communicating with the WLC.

Many corporate environments require deployment of wireless networks on a large scale. The wireless LAN controller (WLC) is a device that assumes a central role in the wireless network and helps make it easier to manage such large-scale deployments. Traditional roles of access points, such as association or authentication of wireless clients, are handled by the WLC. Access points, called lightweight access points in the Unified Communications environment, register themselves with a WLC and tunnel all the management and data packets to the WLCs, which then switch the packets between wireless clients and the wired portion of the network. All the configurations are done on the WLC. LWAPs download the entire configuration from WLCs and act as a wireless interface to the clients. Figure 12-3 illustrates the relationship between the WLC and the LWAP.

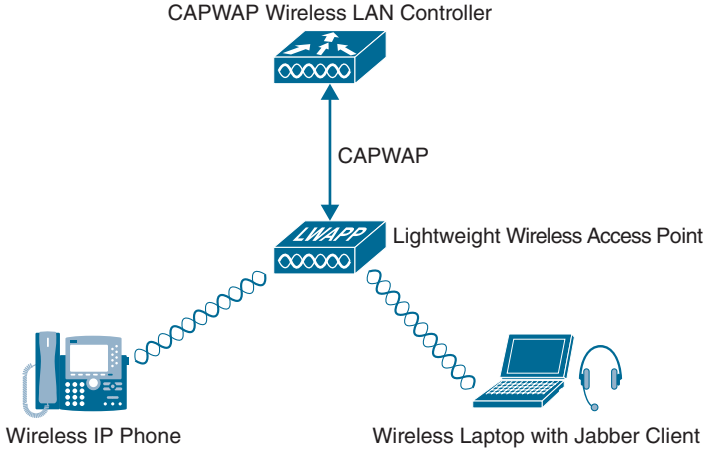

Figure 12-3 *WLC Deployment with LWAP*

The authentication database is a core component of the wireless networks, and it holds the credentials of the users to be authenticated while the wireless association is in progress. The authentication database provides network administrators with a centralized repository to validate the credentials. Network administrators simply add the wireless network users to the authentication database instead of having to add the users to all the wireless access points with which the wireless devices might associate. In a typical wireless authentication scenario, the WLC couples with the authentication database to allow the wireless association to proceed or fail. Authentication databases commonly used are LDAP and RADIUS, although under some scenarios the WLC can also store a small user database locally that can be used for authentication purposes.

The supporting wired network is the portion of the system that serves as a path between WLCs, WAPs, and wired call elements. Because the WAPs need to communicate to the wired world, part of the wired network has to enable those communications. The supporting wired network consists of the LAN switches, routers, and WAN links that work together to communicate with the various components that form the architecture for voice and video over WLAN.

Key Topic

The wireless collaboration endpoints are the user-facing voice and video nodes that operate over the WLAN, which are used for communication. These endpoints can be voice only or enabled for both voice and video. When end users employ the wireless communication endpoints to call a desired destination, the endpoints in turn forward the request to their associated call-processing server. If the call is allowed, the endpoints process the voice or video, encode it, and send it to the receiving device or the next hop of processing. Typical Cisco wireless endpoints are wireless IP phones, voice and video software clients running on desktop computers, mobile smartphones connected through wireless media, and mobile collaboration enterprise tablets. Specifically, the Cisco Unified IP Phones 7861, 8861, and 8865 support wireless communication. Any Windows or Mac computer running the Cisco Jabber application, or any smartphone or tablet running Jabber, could be connected wirelessly to the network. Also, the Cisco DX series endpoints support wireless connections to the network.

Whether the wireless collaboration endpoints initiate a session between each other or with wired endpoints, wired call elements are involved in some way. Wired call elements, such as gateways and call-processing entities, are the supporting infrastructure for voice and video endpoints coupled to that infrastructure. Figure 12-4 illustrates all of the components needed in a WLAN deployment.

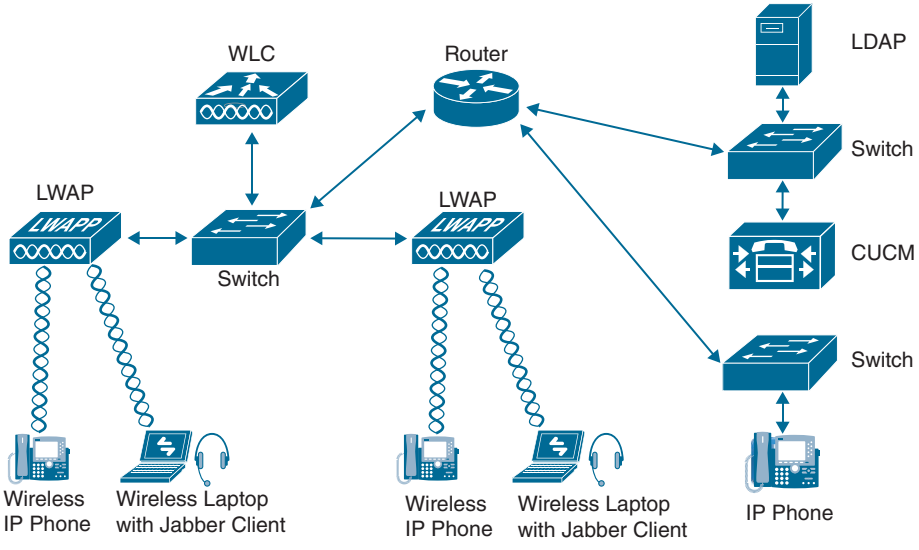

Figure 12-4 *Components of a Full WLAN Deployment*

High Availability

Providing high availability in collaboration solutions is a critical requirement for meeting the modern demands of continuous connectivity. Collaboration deployments designed for high availability increase reliability and uptime. Using real-time applications such as voice or

video over WLAN without high availability could have very adverse effects on the end-user experience, including an inability to make voice or video calls.

A unique aspect of high availability for voice and video over WLAN is high availability of radio frequency (RF) coverage to provide Wi-Fi channel coverage that does not depend on a single WLAN radio. The Wi-Fi channel coverage is provided by the AP radios in the 2.4 GHz and 5 GHz frequency bands.

Key Topic

The primary mechanism for providing RF high availability is cell boundary overlap. In general, a cell boundary overlap of 20 to 30 percent on nonadjacent channels is recommended to provide high availability in the wireless network. For mission-critical environments, at least two APs should be visible at the required signal level (–67 dBm or better). An overlap of 20 percent means that the RF cells of APs using nonadjacent channels overlap each other on 20 percent of their coverage area, while the remaining 80 percent of the coverage area is handled by a single AP. Furthermore, when determining the locations for installing the APs, you should avoid mounting them on reflective surfaces, such as metal, glass, and so forth, which could cause multipath effects that result in signal distortion. Figure 12-5 illustrates a high availability deployment of WLANs using overlapping channel cells.

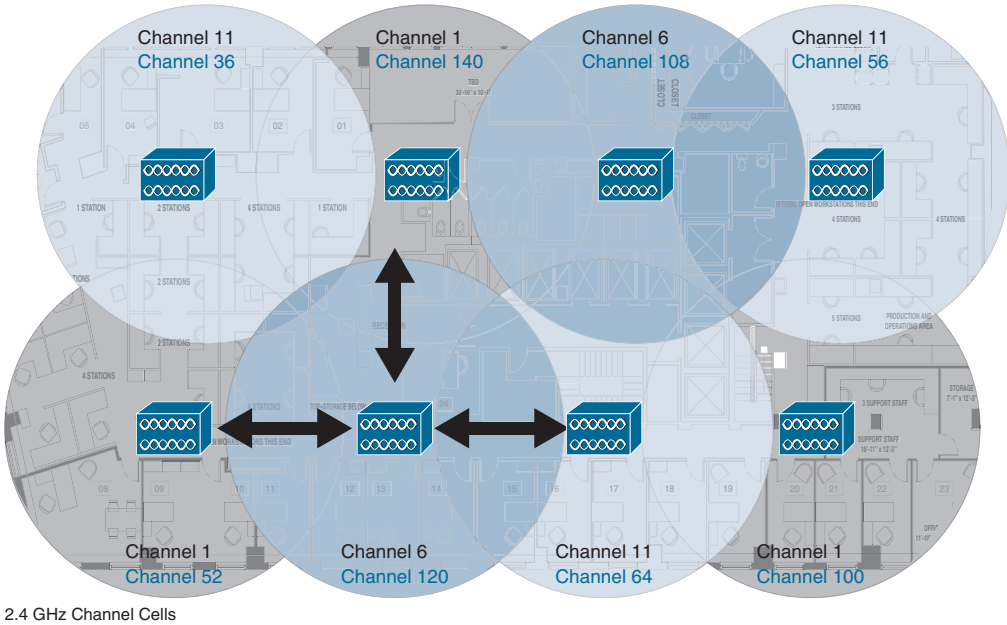

Figure 12-5 *High Availability Deployment of WLANs Using Overlapping Channel Cells*

Careful deployment of APs and channel configuration within the wireless infrastructure are imperative for proper wireless network operation. For this reason, Cisco requires customers to conduct a complete and thorough site survey before deploying wireless networks in a production environment. The survey should include verifying nonoverlapping channel configurations, Wi-Fi channel coverage, and required data and traffic rates; eliminating rogue APs; and identifying and mitigating the impact of potential interference sources. Additionally, you should evaluate utilizing a 5 GHz frequency band, which is generally less crowded and thus usually less prone to interference. If Bluetooth is used, it is highly recommended to use

a 5 GHz WLAN band (802.11a/n/ac) whenever possible for endpoint connectivity. Similarly, the use of Cisco CleanAir technology will increase the WLAN reliability by detecting radio frequency interference in real time and providing a self-healing and self-optimizing wireless network.

Capacity Planning

A crucial piece in planning for voice and video over WLAN is adequately sizing the solution for the desired call capacity. Capacity is defined as the number of simultaneous voice and video sessions over WLAN that can be supported in a given area. Capacity can vary depending on the RF environment, the collaboration endpoint features, and the WLAN system features. For instance, a solution using Cisco Unified Wireless IP Phones 8861 on a WLAN that provides optimized WLAN services would have a maximum call capacity of 27 simultaneous sessions per channel at a data rate of 24 Mbps or higher for both 802.11a and 802.11g. On the other hand, a similar solution with a wireless device such as a tablet running the Jabber client making video calls at 720p and a video rate of 2500 kbps on a WLAN, where access points are configured as 802.11a/n with a data rate index of Modulation and Coding Scheme 7 in 40 MHz channels, would have a maximum capacity of seven video calls and two bidirectional voice and video streams per channel. To achieve these capacities, there must be minimal wireless LAN background traffic and RF utilization, and Bluetooth must be disabled in the devices. It is also important to understand that call capacities are established per nonoverlapping channel because the limiting factor is the channel capacity and not the number of access points (APs). The call capacity specified by the actual wireless endpoint should be used for deployment purposes because it is the supported capacity of that endpoint.

Design Considerations

It is easy to understand that the design and implementation of a WLAN environment in the workplace is very complex and requires many considerations to be factored into the overall network design. Additionally, WLAN configuration specifics can vary depending on the voice or video WLAN devices being used and the WLAN design. Other design considerations for a proper WLAN deployment include the following:

- VLANs
- Roaming
- Wireless channels
- Wireless interference and multipath distortion
- Multicast on the WLAN
- Wireless AP configuration and design
- Wireless LAN controller design considerations
- WLAN quality of service

An entire series of books could be written just to cover the WLAN components and considerations that must be accounted for in a network design that supports voice and video communications. Cisco has a CCNP and CCIE certification track that deals with the many facets of a WLAN environment, and many resources are available to extend an engineer's

understanding of these solutions. To keep the contents of this section relevant to the scope of this book, we will not cover the preceding topics. However, the next chapter will delve into some of the WLAN QoS settings that need to be configured to support voice and video over a WLAN environment.

Gateways

Gateways provide a number of methods for connecting a network of collaboration endpoints to the public switched telephone network (PSTN), a legacy PBX, or external systems. Voice and video gateways range from entry-level and standalone platforms to high-end, feature-rich integrated routers, chassis-based systems, and virtualized applications.

During the 1990s and early 2000s, the only way for an enterprise to connect its internal voice and video network to services outside the enterprise was by means of time-division multiplexing (TDM) or serial gateways through the traditional PSTN. Cisco still offers a full range of TDM and serial gateways with analog and digital connections to the PSTN as well as to PBXs and external systems. TDM connectivity covers a wide variety of low-density analog (FXS and FXO), low-density digital (BRI), and high-density digital (T1, E1, and T3) interface choices. Starting around 2006, new voice and video service options to an enterprise became available from service providers, often as SIP trunk services. Using a SIP trunk for connecting to the PSTN and other destinations outside the enterprise involves an IP-to-IP connection at the edge of the enterprise's network. The same functions traditionally fulfilled by a TDM or serial gateway are still needed at this interconnect point, including demarcation, call admission control, quality of service, troubleshooting boundary, security checks, and so forth. For voice and video SIP trunk connections, the Cisco Unified Border Element and the Cisco Expressway Series fulfill these functions as an interconnection point between the enterprise and the service provider network.

There are two types of Cisco TDM gateways: analog and digital. Both types support voice calls, but only digital gateways support video. The two categories of Cisco analog gateways are station gateways and trunk gateways. Analog station gateways connect the Cisco Unified Communications Manager to plain old telephone service (POTS) analog telephones, interactive voice response (IVR) systems, fax machines, and voicemail systems. Station gateways provide foreign exchange station (FXS) ports. Analog trunk gateways connect the Cisco Unified Communications Manager to PSTN central office (CO) or PBX trunks. Analog trunk gateways provide foreign exchange office (FXO) ports for access to the PSTN, PBXs, or key systems, and E&M (recEive and transMit, or ear and mouth) ports for analog trunk connection to a legacy PBX. Analog Direct Inward Dialing (DID) and Centralized Automatic Message Accounting (CAMA) are also available for PSTN connectivity. Cisco analog gateways are available on the following products and series:

- Cisco Analog Voice Gateways VG204XM and VG300 Series (VG310, VG320, VG350) all support SCCP.
- Cisco Integrated Services Routers Generation 2 (ISR G2) 2900, 3900, 3900E, and 4000 Series (4300 and 4400) with appropriate PVDMs and service modules or cards. PVDM4s utilized by ISR 4000 Series do not support video today.
- Cisco Analog Telephone Adapter (ATA) 190 (SIP only) provides a replacement for the ATA188.

A Cisco digital trunk gateway connects the Cisco Unified Communications Manager to the PSTN or to a PBX via digital trunks such as Primary Rate Interface (PRI), Basic Rate Interface (BRI), serial interfaces (V.35, RS-449, and EIA-530), or through T1 Channel Associated Signaling (CAS). Digital T1 PRI and BRI trunks can be used for both video and audio-only calls. Cisco digital trunk gateways are available on the following products and series:

- Cisco Integrated Services Routers Generation 2 (ISR G2) 1900, 2900, 3900, 3900E, 4300, and 4400 Series with appropriate PVDMs and service modules or cards (The PVDM2 cards were end of life as of June 30, 2019. Cisco recommends using the PVDM3 or PVDM4 cards on appropriate routers.)
- Cisco Telepresence ISDN GW 3241 and MSE 8321 (These products were end of sale as of May 2, 2017.)
- Cisco Telepresence Serial GW 3340 and MSE 8330 (These products were end of sale as of May 2, 2017.)

The Cisco Telepresence ISDN link is a compact appliance that provides Cisco Telepresence endpoints direct ISDN and external IP network connectivity. This unit is supported on all Cisco Telepresence endpoints running TC or CE software. While traditional voice and video gateways are shared resources that provide connectivity between the IP network and the PSTN for many endpoints, each Cisco ISDN link is paired with a single Cisco endpoint.

ISR, ASR, and IOS Software Comparisons

Cisco has a line of routers that offer advanced services beyond what traditional routers can offer. These routers are ideal for use in voice and video environments as well as support of smaller branch office locations or hybrid communication to the cloud. They are the Integrated Services Routers (ISRs), Aggregation Services Routers (ASRs), and Cloud Services Routers (CSRs). The biggest difference between Cisco ASR and ISR routers is that ASR routers are for enterprises and service providers, whereas ISRs are for customers with small- or medium-sized networks. CSRs go beyond the scope of this book, so the focus will be on the ISR and ASR options.

Two factors that should be considered will influence which of these router product lines should be used within a particular customer environment. Sizing is the most important factor, followed by the software running on the router. Different software will provide different features and capabilities. The classic IOS software was used on all Cisco routers prior to 2007 and may still be found running on some routers in production today. However, most current products in the Cisco ISR and ASR routers use either the IOS XE or IOS XR software version.

Cisco IOS XE software is an open and flexible operating system optimized for a new era of enterprise networks. Its standards-based programmable interfaces automate network operations and give you deep visibility into user, application, and device behaviors. As the single OS for enterprise wired and wireless access, aggregation, core, and WAN, Cisco IOS XE reduces business and network complexity. Cisco IOS XE software is open because it includes the following open standards-based capabilities: NETCONF (RFC 6241) programmable interfaces, IETF YANG push telemetry, OpenConfig and IETF YANG data models, and Guest Shell Linux Containers (LXC). Yet the user interface is the same familiar CLI that engineers have been using throughout the lifecycle of the older classic IOS software. This

highly scalable software has been developed with resiliency in mind; Cisco IOS XE reduces planned and unplanned downtime. Service and software upgrades are more efficient, and Graceful Insertion and Removal lets you update or debug a switch without disrupting network traffic. Cisco IOS XE software also has built-in security and trust, which helps protect against modern cyberattacks. It assures that Cisco hardware and software are genuine and unmodified. And its enhanced platform integrity, security, and resilience mean you can be confident that data is trustworthy. Additionally, all IOS XE software-based routers support hybrid cloud services.

IOS XR is Cisco IOS software used on the high-end Network Converging System (NCS), carrier-grade routers such as the CRS Series, 12000 Series, and ASR9000 Series. In fact, the ASR9000 Series are the only ASR or ISR routers that support the IOS XR software. Cisco's IOS XR software shares very little infrastructure or feature support with the other IOS software options and is instead built on a preemptive, memory-protected, multitasking, microkernel-based operating system. The microkernel was formerly provided by QNX; versions 6.0 and up use the Wind River Linux distribution. IOS XR is an on-premises-only software solution with no hybrid cloud support, and it aims to provide the following advantages over the earlier IOS versions:

- Improved high availability (largely through support for hardware redundancy and fault containment methods such as protected memory spaces for individual processes and process restart-ability)
- Better scalability for large hardware configurations (through a distributed software infrastructure and a two-stage forwarding architecture)
- A package-based software distribution model (allowing optional features such as multicast routing and MPLS to be installed and removed while the router is in service)
- The ability to install package upgrades and patches (potentially while the router remains in service)
- A web-based GUI for system management (making use of a generic, XML management interface)

As mentioned previously, the number of users the router needs to support for sizing and the software version will determine the specific product needed for a given customer site. Table 12-3 identifies all the different series of ISR and ASR routers available with their sizing limitations and software availability.

Key Topic

Table 12-3 ISR and ASR Routers and Software Options

Router Model	Software Version	Sizing Limitations (SRTP/RTP Sessions)
ASR 900	IOS XE	Unknown
ASR 1000 Series	IOS XE	Unknown
ASR 9000 Series	IOS XR	Unknown
ISRv	IOS XE	Up to 1000
ISR 1000	IOS XE	75 to 100

Router Model	Software Version	Sizing Limitations (SRTP/RTP Sessions)
ISR 4000 Series	IOS XE	40 to 1500
CSR 1000v	IOS XE	225 to 800 (1 or 4 vCPU)
ISR 800 Series	Classic IOS	20
ISR 900 Series	Classic IOS	50
ISR 2900 Series	Classic IOS	Up to 200

ISR Products Explained

As Table 12-3 shows, you have many product lines to choose from when deploying the Cisco ISR. This section will delve into four specific ISRs: the ISRv (which is a virtual router), 800 Series ISR, 1000 Series ISR, and the 4000 Series ISR. The 800 Series ISR is the only router discussed in this section that runs the classic IOS software instead of the IOS XE software. However, this is a great router to use for remote office locations, and it supports DMVPN.

There are many benefits to using the ISRv. First, it supports rapid deployment and service automation. The virtual form factor accelerates deployment and eliminates hardware costs such as complete equipment upgrades and return materials authorization (RMA). It also supports single-tenant use. This feature allows a cloud service provider to provision a routing instance per tenant, simplifying service delivery and tenant management. It also helps the provider overcome VLAN scale limits, increasing tenant scale. The enterprise network extension to the cloud feature provides enterprises highly secure direct connections from their distributed sites to their cloud-hosted applications, improving application response time and user experience. The network consistency feature uses familiar enterprise-class Cisco IOS software features for consistent network operation across premises and the cloud, allowing the enterprise to view the cloud as just another node in its network. Network scalability allows scale beyond the limitations of 802.1q VLAN tagging by building a VXLAN network or extending Layer 3 routing deeper into the cloud environment. Finally, consolidation of network functions eliminates the facility requirements and complexity of physical network devices by consolidating multiple network functions onto a single piece of server hardware.

The following features are supported on the Cisco ISRv:

- **Routing:** Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Policy-Based Routing, IPv6, Virtual Route Forwarding Lite (VRF-Lite), Multicast, LISP, and Generic Routing Encapsulation (GRE)
- **Addressing:** Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), Network Address Translation (NAT), 802.1Q VLAN, Ethernet Virtual Connection (EVC), and VXLAN
- **VPN:** IPsec VPN, DMVPN, Easy VPN, SSL VPN, and FlexVPN
- **MPLS:** MPLS VPN, VRF, and Bidirectional Forwarding Detection (BFD)
- **Security:** Cisco IOS Zone-Based Firewall (ZBFW); access control list (ACL); authentication, authorization, and accounting (AAA); RADIUS; and TACACS+

- **High availability:** HSRP, Virtual Router Redundancy Protocol (VRRP), Gateway Load Balancing Protocol (GLBP), and box-to-box high availability for ZBFW and NAT
- **Traffic redirection:** AppNav (to Cisco Wide Area Application Services [WAAS]) and Web Cache Communication Protocol (WCCP) application visibility, performance monitoring, and control: quality of service (QoS), AVC, and IP service-level agreement (SLA)
- **Hybrid cloud connectivity:** OTV, Virtual Private LAN Service (VPLS), and Ethernet over MPLS (EoMPLS)
- **Management:** Command-line interface (CLI), Secure Shell (SSH) Protocol, NetFlow, Simple Network Management Protocol (SNMP), Embedded Event Manager (EEM), and RESTful application programming interfaces (APIs)
- **NFV:** Virtual route reflector (vRR), virtual broadband network gateway (vBNG), and virtual intelligent services gateway (vISG)

An affordable router series that supports important network services such as security is the Cisco 800 Series Integrated Services Routers. These ISRs deliver secure, reliable WAN connectivity that small offices and remote workers need. Additionally, they support built-in voice features, wireless, WAN optimization, and machine-to-machine communications. Different models in the series support different connection types to serve the specific needs of a small office. That could be xDSL, Wi-Fi, 4G LTE, Ethernet, fiber, or something else. Routing, switching, wireless, and intelligent IP network services are all bundled into one compact form factor that's quick to install using the Cisco Configuration Professional Express tool. They can all be managed centrally from a data center with Cisco Prime Infrastructure and LiveAction applications. The 800 ISRs provide comprehensive security—encryption, VPN, firewall, and cloud-based URL filtering—to help safeguard customers and data. Table 12-4 outlines the models in the 800 Series ISRs and includes deployment recommendations and top WAN speeds.

Key Topic

Table 12-4 800 Series ISRs

Model	Deployment Recommendation	Top WAN Speed with Services On
860	Home or small offices with up to 10 users	10 Mbps
880	Remote workers, small offices, and branch locations with up to 20 users	15 Mbps
810	Machine-to-machine and device-to-device deployments such as ATMs, point-of-sale, kiosks, vending machines (fixed platform)	15 Mbps
890	Enterprise remote offices with up to 50 users	>20 Mbps
800M	Microbranches, industrial, Internet of Things/IoT (modular platform)	Various Cellular Data Rates

The Cisco 1000 Series ISR platform with its small form factor is best suited for small and midsize businesses, enterprise branches, and as customer premises equipment in managed services environments. The routers come with four or eight LAN ports in various model options. They

have high performance with Gigabit Ethernet packet-forwarding capabilities. The multicore architecture has separate cores for the data plane and control plane. The 1000 Series ISRs support Power over Ethernet (PoE) and PoE+ to power branch devices such as IP phones and cameras. They are easy to deploy with zero-touch provisioning using Plug-and-Play capability. There are multiple combinations to choose from, including LAN, WLAN, WAN, DSL, LTE, and pluggable, depending on your branch needs. The 1000 Series can be used in ATMs, retail stores, and kiosks, as well as for various other purposes. The 1000 Series ISRs address the demands of increased mobility with LTE Advanced and 802.11ac (Wave 2) Wi-Fi. It has a comprehensive set of WAN connectivity options such as Ethernet, Fiber, LTE, and the latest DSL technologies, like G.fast. The routers provide a great return on investment, allowing you to save on operating expenses by reducing WAN link costs with software-defined WAN capability and transport independence using Cisco SD-WAN. You can also reduce capital expenses using pay-as-you-grow licensing for IPsec performance. The 1000 Series ISRs answer the latest security threats to networking devices with advanced features such as zone-based firewall, Trustworthy Systems, Cisco Umbrella security, and Encrypted Traffic Analytics.

The Cisco 1000 Series ISRs include the following models:

- Cisco 1100-8P ISR with LTE Advanced
- Cisco 1100-4P ISR with DSL
- Cisco 1101-4P
- Cisco 1101-4PLTEP

As you build out the digital capabilities in your enterprise branch offices, you should consider the full-service sophistication of the Cisco 4000 Series Integrated Services Routers. The 4000 Series ISRs consolidate many must-have IT functions, including network, security, compute, storage, and unified communications. So, you get everything you need in a single platform. That means significant savings in capital, operational, and management expenses for lower total cost of ownership. The platform is modular and upgradable, so you can add new services without changing equipment. It supports multiple application-aware services concurrently while maintaining WAN performance of up to 2 Gbps, even during heavy traffic loads. The backplane architecture supports high-bandwidth, module-to-module communication at speeds up to 10 Gbps. The 4000 Series includes Cisco Trust Anchor Technologies that help mitigate modern cyberattacks by verifying platform integrity and providing protection from counterfeit and unauthorized modification of hardware and software.

The 4000 Series runs Cisco Intelligent WAN (IWAN), a comprehensive set of traffic control and security features. IWAN includes all the business-grade capabilities of a Multiprotocol Label Switching (MPLS) VPN using other types of less-expensive links, such as per-application traffic management, WAN optimization, and VPN tunneling, which can be put to work across Internet, cellular, and other lower-cost services as connections are added. Additionally, new router services can be activated on demand through a simple licensing change. Local IT staff are not needed in the branch to deliver a fully comprehensive computing and networking experience with remote application installation and management capabilities.

Cisco IWAN features can now be configured in next to no time, thanks to Cisco's enterprise software-defined networking (SDN) controller, the Application Policy Infrastructure Controller Enterprise Module (APIC EM). APIC EM allows automation of lots of tasks across the

network. You can implement an SDN on the Cisco WAN infrastructure without having to upgrade any equipment; just install the no-charge APIC EM software-based controller between applications and network infrastructure. The controller translates business policy directly into network device-level policy for automatic compliance with any corporate and industry-mandated polices. For additional WAN management simplicity, you can also use the IWAN app for APIC EM. The app automates the configuration of Cisco Intelligent WAN features, such as quality of service, WAN optimization, and security, in Cisco branch and edge WAN routers. The app slashes what used to require 1000 CLI steps to just 10 mouse clicks per site. With the IWAN app's template functionality, the ability to configure, deploy, and manage large numbers of branch offices has never been easier. The 4000 Series ISR contains these platforms: the 4451, 4431, 4351, 4331, 4321, and 4221 ISRs.

APIC EM is now wrapped into DNA Center, which falls under SD-Access. It could also be described as "Cisco's Software-Defined Access (SD-Access) solution" because it's the blanket term of the software-defined LAN side. Also, to add to the confusion, APIC EM and IWAN are still being used, but more attention is focused on the Cisco SD-WAN solution that's based on the Viptela acquisition. Lastly, both SD-Access and SD-WAN fall under Cisco Digital Network Architecture (Cisco DNA), similar to how everything in collaboration now falls under Webex.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 12-5 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 12-5 Key Topics for Chapter 12

Key Topic Element	Description	Page Number
Paragraph	LAN Communication and Layer 2 Switches Overview	281
Paragraph	Layer 3 Routers and Basic Services They Provide	282
Paragraph	Medianet Explained	283
Paragraph	Spanning Tree Protocol Explained	285
Paragraph	VLANs Explained	285
Table 12-2	PoE Types and Supported Power	286
Paragraph	Comparison of HSRP, VRRP, and GLBP	286
Paragraph	WAN Aggregation Design Models	287
Paragraph	DMVPN Explained	288

Key Topic Element	Description	Page Number
List	Basic Elements of a WLAN	289
Paragraph	Examples of Wireless Voice and Video Endpoints	291
Paragraph	Cell Boundary Overlap	292
List	WLAN Design Considerations	293
Paragraph	Gateway Categorization	294
Table 12-3	ISR and ASR Routers and Software Options	296
List	ISRv Features Supported	297
Table 12-4	800 Series ISRs	298
List	Cisco 1000 Series ISR Models	299

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

802.1d, 802.1s, 802.1w, 802.1p, 802.1Q, ACL, ADDTS, Advanced Networking, Annunciator, ARP, BRI, CAMA, CAPWAP, CAS, CDP, Cell Boundary Overlap, CleanAir, DHCP, DID, DMVPN, DNS, EIGRP, FXO, FXS, GLBP, HSRP, IPSec V3PN, ISR, IVR, LAN, LLDP-MED, LWAP, LWAPP, MPLS, NAT, NTP, OSPF, Option 66, Option 150, PAT, PBX, POTS, PRI, PSTN, QoS, RF, STP, TDM, TFTP, VLAN, VRRP, VSS, VVID, WAN, WAP, WLAN, WLC

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. What are the three types of Spanning Tree that can be configured on Cisco switches?
2. What are three fault-tolerant default gateway protocols?
3. List three guaranteed bandwidth options for Internet connection that are currently being used.
4. List the four Cisco ISR 1100 Series routers.

CHAPTER 13

Layer 2 and Layer 3 QoS Parameters

This chapter covers the following topics:

QoS-Related Issues: This topic will discuss latency, jitter delay, and bandwidth issues that can be overcome with a good QoS solution.

Class Models for Provisioning QoS: This topic will discuss three different models for classifying network traffic using QoS: the 4/5 class model, 8 class model, and the 11 class model. All three are based on the QoS Baseline model.

QoS Requirements: This topic will begin to explain the layered components that make up a complete QoS solution, such as Layer 2 trust boundaries; congestion management tools; congestion avoidance tools; and policing, shaping, and link efficiency methods.

Traffic Classifications: This topic will break down different traffic classifications within the LAN, across the WAN, and over the wireless LAN.

Configure and Verify LLQ: This topic will explain how to configure and verify an LLQ QoS deployment through configuring a class map, policy map, and service policy.

Quality of service (QoS) refers to the capability of a network to provide improved service to selected network traffic over various underlying technologies. This chapter will explain what QoS is, review different components that make up a QoS solution, and provide a basic understanding of how to configure a QoS solution on an IOS router. Topics discussed in this chapter include the following:

- QoS-Related Issues:
  - Latency, Jitter, and Packet Loss
  - Bandwidth
- Class Models for Provisioning QoS:
  - 4/5 Class Model
  - 8 Class Model
  - QoS Baseline Model (11 Class)
- QoS Requirements:
  - QoS Trust Boundaries
  - Congestion Management
  - Congestion Avoidance

- Policing
- Shaping
- Link Efficiency Methods

- Traffic Classifications:
  - LAN and WAN Traffic Classifications
  - WLAN Traffic Classifications
- Configure and Verify LLQ:
  - Class Map
  - Policy Map
  - Service Policy

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.1.h QoS
- 5.1 Describe problems that can lead to poor voice and video quality
  - 5.1.a Latency
  - 5.1.b Jitter
  - 5.1.c Packet loss
  - 5.1.d Bandwidth
- 5.2 Describe the QoS requirements for these application types (voice and video)
- 5.3 Describe the class models for providing QoS on a network
  - 5.3.a 4/5 Class model
  - 5.3.b 8 Class model
  - 5.3.c QoS Baseline model (11 Class)
- 5.4 Describe the purpose and function of these DiffServ values as it pertains to collaboration
  - 5.4.a EF
  - 5.4.b AF41
  - 5.4.c AF42
  - 5.4.d CS3
  - 5.4.e CS4
- 5.5 Describe QoS trust boundaries and their significance in LAN-based classification and marking

- 5.6 Describe and determine location-based CAC bandwidth requirements
- 5.7 Configure and verify LLQ (class map, policy map, service policy)

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 13-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 13-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
QoS-Related Issues	1–3
Class Models for Provisioning QoS	4–6
QoS Requirements	7–9
Traffic Classifications	10–11
Configure and Verify LLQ	12–14

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is an example of drop-insensitive data?
   a. FTP data packets
   b. Call setup messages
   c. Email
   d. Voice media during a call
2. Which of the following is described as a variable of the delay over a period of time?
   a. Latency
   b. Jitter
   c. Packet loss
   d. Slow bandwidth
3. An engineer is trying to calculate how much bandwidth is being consumed across the network for video calls. All video calls are required to consume no more than 512 kbps bandwidth per call. How much bandwidth is actually being consumed per call?
   a. (512k + 40 + 32) × 50 pps × 8 bits/bytes = 524,800 bps or 524 kbps
   b. 512 kbps

c. 512 × 1.2 = 614.4 or 614 kbps

d. 512 × 2 (send and receive) = 1024 kbps

4. Which of the following classifications is part of the 4/5 class model?

a. Real Time

b. Audio

c. Video

d. Bulk Data

5. Cisco recommends that call signaling be marked with CS3 for proper QoS handling. However, older model phones do not use CS3 for call signaling. What is the other QoS marking for call signaling that may need to be accounted for?

a. EF

b. DF

c. AF31

d. AF41

6. When establishing the QoS Baseline 11 class model, how much bandwidth should be allocated for interactive video?

a. 13%

b. 23%

c. 25%

d. 33%

7. What command can be entered into a switch to enable QoS at the Layer 2 level?

a. **mls qos**

b. **mls qos interface fastethernet 0/1**

c. **mls qos trust cos**

d. No command is needed because QoS is enabled by default.

8. Which of the following congestion management mechanisms allows delay-sensitive data, such as voice and video, to be given preferential treatment over other traffic by letting this data be dequeued and sent first?

a. FIFO

b. PQ

c. CQ

d. WFQ

e. CBWFQ

f. LLQ

9. Which of the following is an early detection congestion avoidance mechanism that ensures high-precedence traffic has lower loss rates than other traffic during times of congestion?

a. CBWFQ

b. WRED

c. CAR

d. GTS

e. FRTS

10. The PHB QoS marking for voice-only packets is EF. What is the DSCP equivalent to this marking?
    a. 32
    b. 34
    c. 46
    d. 48
11. How many QoS queues do Cisco APs provide for downstream traffic being sent to wireless clients?
    a. 2
    b. 4
    c. 8
    d. 11
12. When it comes to Class-Based Weighted Fair Queuing, what advantage does LLQ offer over CBWFQ without LLQ?
    a. LLQ provides WFQ based on defined classes for CBWFQ.
    b. LLQ provides strict priority queueing for CBWFQ.
    c. The weight for a packet belonging to a specific class is derived from the bandwidth assigned to the class when LLQ is used.
    d. All packets are serviced fairly based on weight with LLQ.
13. Which of the following statements is true regarding traffic classification and traffic marking?
    a. Traffic marking can be viewed as an additional action, specified in a policy map, to be taken on a traffic class.
    b. Traffic classification can be viewed as an additional action, specified in a policy map, to be taken on a traffic class.
    c. Traffic marking allows you to organize packets into traffic classes on the basis of whether the traffic matches specific criteria.
    d. After the traffic is organized into traffic classes, traffic classification allows you to mark an attribute for the traffic belonging to that specific class.
14. You can configure class policies for as many classes as are defined on the router, up to what maximum value?
    a. 11
    b. 16
    c. 64
    d. 128

Foundation Topics

QoS-Related Issues

Quality of service (QoS) was developed out of necessity. Early networks suffered from bursty data flows due to the rate at which data came into the network. As data packets arrived on the network, they would try to consume as much bandwidth as possible. Access

was on a first-come, first-served basis. The data rates available to any one user depended on the number of users accessing the network at that given time. The networking protocols that were developed prior to QoS were intentionally designed to adapt to the bursty nature of the network so that packets being sent could survive bursty traffic and brief outages within the network. The nature of how TCP packets are sent is a great example of the ingenuity behind these earlier protocols. TCP packets require an acknowledgment for each packet sent. If the acknowledgment doesn't come within a given period of time, the TCP transmission will be resent. Email uses TCP. For this reason, an email may come into an inbox seconds after it is sent or several minutes later. The contents of the email are whole and intact, so the delivery time is irrelevant, though the arrival time of the email might be annoying. This type of traffic is referred to as *drop-insensitive data* because lost packets will not prevent the data from eventually transmitting.

Drop-sensitive data is negatively impacted when data packets are lost because the nature of the transmission does not allow for the packets to be resent. Drop-sensitive data is typically in real time, and the nature of UDP packets uses a one-way, one-time send delivery mechanism. It's like a shipping company that delivers a package to your front door when you are not home to receive it. The delivery person's job is to get the package to the door. If it is stolen before you get home, that is not the delivery company's issue. As companies began using the packet-switched network for drop-sensitive data, such as voice and video communications, a more permanent solution had to be developed to contend with the bursty and sporadic nature of the network. Early network engineers faced with these issues would develop and support nonintegrated networks, each designed to carry a specific type of traffic. However, these network setups were very complex, not scalable, and very difficult to support.

Key Topic

The concept of QoS allows for various data traffic types to be carried over a single converged network, and yet each type of traffic can be treated differently. Four factors in a network design can lead to poor audio and video quality. They are bandwidth capacity, latency (delay), jitter, and packet loss. The end goal of QoS is to provide better and more predictable network services by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics as required by the business applications. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network.

Key Topic

Latency, Jitter, and Packet Loss

Latency is the delay packets' experience when traversing across many different network devices. Jitter is similar to latency in that it is also a delay, but jitter is a variable of the latency that occurs over a period of time. If every packet sent between phone A and phone B took the exact same amount of time to traverse the network, there would be no jitter, but there could still be a delay. When the time between when packets are delivered is different, that time variance is referred to as *jitter*. Jitter can be overcome with buffers, but that will add to the overall latency that occurs. *Packet loss* refers to packets being dropped by the router due to congestion on a link. As mentioned before, TCP packets are drop insensitive because they will be retransmitted if an acknowledgment is not received by the initiating application. By contrast, UDP packets, such as voice and video, will not be retransmitted and are therefore drop sensitive.

Bandwidth

Bandwidth capacity limitations come into play when multiple flows through the router compete for limited bandwidth. QoS is not a substitute for bandwidth, and increasing the available bandwidth will not solve all these problems either. Before bandwidth issues can be resolved, an engineer needs to calculate the bandwidth being consumed across the network. Based on the information covered earlier in this book, calculating bandwidth consumption for voice and video is relatively easy.

To calculate bandwidth requirements for voice-only traffic, an engineer must first identify the codecs being used. Other codecs may be chosen to control or limit bandwidth consumption. Although the actual codec being used is the most important parameter in this assessment, it is still very important to know about packetization and the technologies that will be used. The payload size is comparable to the header size. The more packets being sent and the shorter the packetization time, the more overhead is being consumed. If the G.711 codec, which consumes 64 kbps bandwidth, is being used with a packetization period of 20 ms, and the voice payload is 160 bytes at 50 packets per second, the bandwidth for the call is 87.2 kbps. However, this calculation does not factor in the Layer 2 or Layer 3 overhead. If we take the overhead into account, the total bandwidth being consumed will be even higher. Some QoS tools also affect overhead bandwidth, such as Compressed Real-time Transport Protocol (cRTP) and link fragmentation and interleaving (LFI), but they will be discussed later in this chapter. All packets being equal, there are definitive bit counts to the Layer 2 and Layer 3 headers. Figure 13-1 illustrates the total bandwidth calculations for a G.711 call with Layer 2 and Layer 3 overhead included.

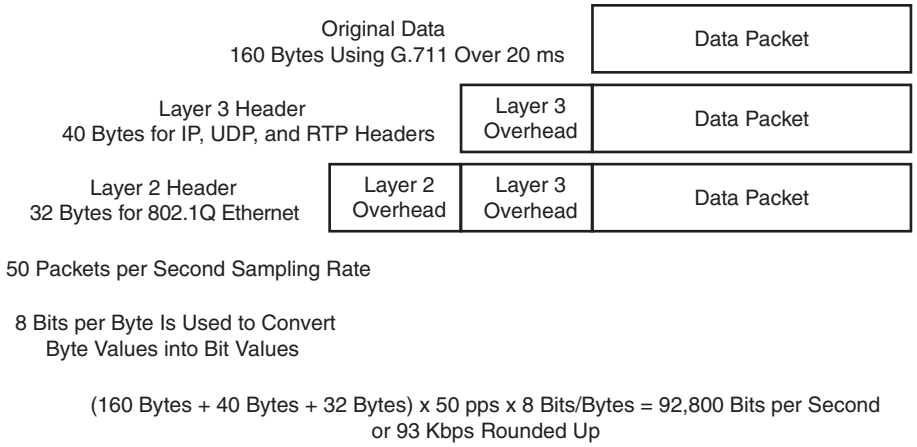

Figure 13-1 *Total Bandwidth Calculations for a Call with Overhead*

Calculating the bandwidth for a video call is similar to an audio call, except the math required for calculating the overhead is much simpler. First, an engineer must specify the codec being used along with the resolution of the call. These two factors together determine the bandwidth and compression of the video call. The codec alone cannot determine the bandwidth used for a video call, nor can the resolution alone. A user can place a video call at 720p30 using H.263 and will consume 1152 kbps bandwidth. The same call can be placed using the 720p30 resolution and the H.264 codec and consume only 768 kbps bandwidth. However, if the H.264 codec is used for a 480p30 call, the bandwidth required could range anywhere between 256 kbps and 512 kbps.

Another consideration when calculating bandwidth for a video call is that the total bandwidth allocated for the call is for both audio and video. If you were to place a 384 kbps call using the G.711 codec for audio and H.264 codec for video, 64 kbps would be allocated to audio, leaving 320 kbps for video. In that same call, if G.722 were used at 48 kbps, 336 kbps bandwidth would be available for video.

Key Topic

Once you know the total bandwidth available for the call and the codecs being used for audio and video, you can break down the Layer 2 and Layer 3 headers for audio and video separately and then add the total bandwidth back together. There is an easier way to calculate total bandwidth consumption for video calls. Simply add in 20 percent to the payload bandwidth for overhead. So, a 768 kbps call will consume approximately 922 kbps bandwidth with overhead. A 384 kbps call will consume 460 kbps bandwidth with overhead. You can do the math the long way, but you will find the actual numbers are very close to 20 percent overhead.

Once the bandwidth consumption has been calculated per call, some simple math based on actual call volume will determine the overall bandwidth consumed via voice and video communications. Determining the other limitations within a network will require other assessments of the network itself. A proper QoS design begins with a network audit and a business audit to determine the type of traffic that is running on the network and then determine the QoS requirements for the different types of traffic. Once that has been accomplished, the next step is to group the traffic into classes with similar QoS requirements. The third step occurs at the WAN Aggregation layer: to define QoS policies that will meet the QoS requirements for each traffic class.

Class Models for Provisioning QoS

Businesses should define the strategy and goals for different applications running in their network before deciding on a QoS plan or applying any QoS tools. The number of different traffic classes identified within the company's network should directly correlate to the end-to-end QoS objectives of the business. Three different QoS strategy models can be deployed, depending on the granularity of applications running within a company's network:

Key Topic

- 4/5 Class Model
- 8 Class Model
- QoS Baseline Model (11 Class)

Although the more classes you define, the more specific and granular traffic treatment will be per application, the selection of a certain strategy model must be based on application requirements coupled with the WAN provider QoS model. The following sections provide a detailed view of each of these QoS strategy models.

4/5 Class Model

The 4/5 class model is the simplest of the three models in terms of QoS polices and typically accounts for real-time communications, call signaling, critical data, best-effort data, and scavenger data. The call signaling and critical data are often grouped together as a "mission-critical" category, thus this model is called the 4/5 class model. There could be four or five classes depending on how services are grouped together. The mission-critical class can also be used for multimedia conferencing, multimedia streaming, and bulk data applications. The

4/5 class model is commonly used within small and medium-sized businesses (SMB) that have deployed VoIP telephony. The five traffic classes of QoS markings and guarantees are as follows:

- **Real Time:** Typically voice-only communications. Marked with EF and provisioned to leverage up to 33 percent of link bandwidth.
- **Call Signaling:** Marked with CS3 and provisioned to leverage a minimum of 7 percent of link bandwidth.
- **Critical Data:** Marked with AF31 and provisioned to leverage 35 percent of link bandwidth. When Signaling and Critical Data are combined, CS3 is used across the board.
- **Best Effort Data:** Marked with DF and provisioned to take advantage of 25 percent of link bandwidth.
- **Scavenger:** Marked with CS1 and provisioned to utilize any unused available link bandwidth. The Scavenger class does not have bandwidth directly provisioned, so packets marked CS1 will be sent only if bandwidth is available and will be the first packets to drop during high congestion times.

Voice and signaling guarantees must be selected based on the volume of voice calls and the VoIP codec that is used through the given link. Mission-critical data is selected based on the decision of the director of each company department who has given info about critical business application needs to the networking team. Platform-specific constraints or service-provider constraints may affect the number of classes of service. Businesses should consider a migration strategy to allow the number of classes to be smoothly expanded as future needs arise.

8 Class Model

As needs arise, businesses might need to expand their service groups to the 8 class model, which builds on the 4/5 class model by dividing three of the classes into two more granular classes each. The additions to this model include splitting the Real Time class into two distinct classes: Voice and Video. The Critical Data class is divided into Network Control and Critical Data. The explicitly defined Network Control traffic class is used for applications such as network routing protocol updates or network infrastructure control traffic such as operations, administration, and maintenance (OAM). Finally, the Best Effort class is divided into Bulk Data and Best Effort classes. The recommendations for each traffic class in this model are as follows:

- **Voice:** Marked with EF and limited to 10 percent of link bandwidth in a strict-priority queue.
- **Video:** Marked with AF41 or sometimes as EF and limited to 23 percent of link bandwidth in a strict-priority queue.
- **Call Signaling:** Marked CS3 and provisioned with a minimum of 2 percent of link bandwidth.
- **Network Control:** Marked with CS2 and provisioned with a minimum of 5 percent of link bandwidth.

- **Critical Data:** Marked with AF31 and provisioned with 25 percent of link bandwidth.
- **Bulk Data:** Marked with AF11 and provisioned with 10 percent of link bandwidth with WRED enabled.
- **Best Effort:** Marked with DF and provisioned with 25 percent of link bandwidth.
- **Scavenger:** Marked with CS1 and provisioned to utilize any unused available link bandwidth. The Scavenger class does not have bandwidth directly provisioned, so packets marked CS1 will be sent only if bandwidth is available and will be the first packets to drop during high congestion times.

Although Cisco does recommend configuring call signaling with CS3, some legacy Cisco Unified IP phone products still mark call signaling to AF31. Cisco has been working on a marking migration from AF31 to CS3 with its newer Cisco Unified IP phone models, but some businesses that use older phone models may still want to reserve both AF31 and CS3 for call signaling. In these cases, the critical data applications should be marked to a temporary placeholder nonstandard DSCP, such as 25. After companies migrate their phones to the newer IP phone models, the QoS Baseline marking recommendations of CS3 for call signaling and AF31 for critical data applications should be used.

QoS Baseline Model (11 Class)

Cisco has adopted a new initiative called the *QoS Baseline*. The QoS Baseline is a strategic document designed to unify QoS within Cisco, from enterprise to service provider and from engineering to marketing. The QoS Baseline was written by Cisco's most qualified QoS experts, who have developed or contributed to the related IETF RFC standards and as such are supremely qualified to interpret these standards. The QoS Baseline also provides uniform, standards-based recommendations to help ensure that QoS designs and deployments are unified and consistent. You can see the "QoS Baseline at a Glance" document at the following link: https://www.cisco.com/en/US/technologies/tk543/tk759/technologies_white_paper0900aecd80295a9b.pdf. Several books that go into more detailed information about this model also are available.

The QoS Baseline defines up to 11 classes of traffic that might be viewed as critical to a given enterprise. The 11 class QoS Baseline model builds on the 8 class model, and represents Cisco's interpretation of the RFC 4594 recommendation, which outlines 12 different classes of traffic. The recommendations for each traffic class in this model are as follows:

Key Topic

- **Voice:** Refers to voice only, and it is marked with EF and limited to 10 percent of link bandwidth in a strict-priority queue.
- **Interactive Video:** Refers to voice and video, and it is marked with AF41 or sometimes as EF and limited to 13 percent of link bandwidth.
- **Streaming Video:** Marked with CS4 or sometimes as EF and limited to 10 percent of link bandwidth.
- **Call Signaling:** Marked with CS3 and provisioned with a minimum of 2 percent of link bandwidth.
- **IP Routing:** Marked with CS6 and limited to 3 percent of link bandwidth.

- **Network Management:** Marked with CS2 and provisioned as guaranteed 2 percent of link bandwidth.
- **Mission Critical Data:** Marked with AF31 and provisioned with 15 percent of link bandwidth.
- **Transactional Data:** Marked with AF21 and provisioned with 10 percent of link bandwidth with Weighted Random Early Detection (WRED) enabled.
- **Bulk Data:** Marked with AF11 and provisioned with 10 percent of link bandwidth with WRED enabled.
- **Best Effort Data:** Marked with 0 and provisioned with 25 percent of link bandwidth.
- **Scavenger:** Marked with CS1 and provisioned to utilize any unused available link bandwidth. The Scavenger class does not have bandwidth directly provisioned, so packets marked CS1 will be sent only if bandwidth is available and they will be the first packets to drop during high congestion times.

Enterprises do not need to deploy all 11 classes of the QoS Baseline model. This model is intended to be a forward-looking guide that considers as many classes of traffic with unique QoS requirements as possible. Familiarity with this model can assist in the smooth expansion of QoS policies to support additional applications as future requirements arise. However, at the time of QoS deployment, the enterprise needs to clearly define its organizational objectives, which will correspondingly determine how many traffic classes will be required.

This consideration should be tempered with the determination of how many application classes the networking administration team feels comfortable with deploying and supporting. Platform-specific constraints or service-provider constraints may also affect the number of classes of service. At this point, you should also consider a migration strategy to allow the number of classes to be smoothly expanded as future needs arise. Figure 13-2 illustrates the strategy Cisco recommends for expanding the number of classes of service over time.

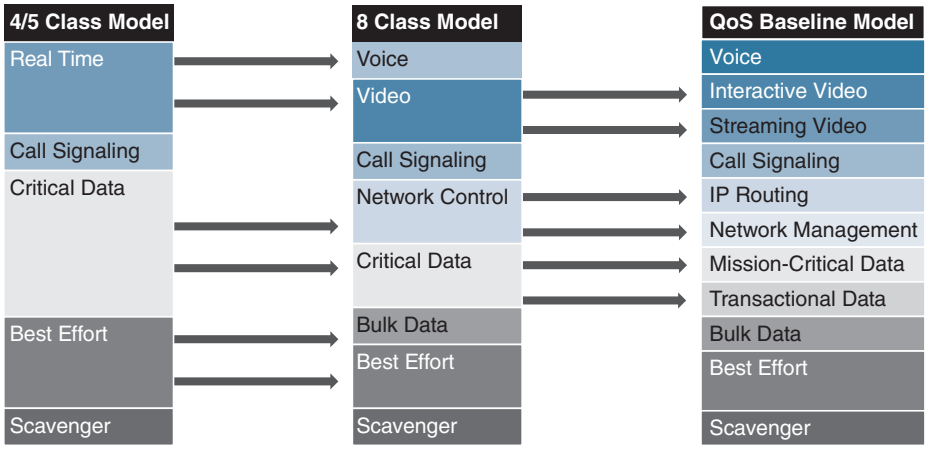

Figure 13-2 *Strategy for Expanding the Number of Classes of Service over Time*

QoS Requirements

Key Topic

After a network has been assessed and a class model for provisioning QoS has been chosen, many layers to a proper QoS solution must still be deployed. As mentioned earlier in this chapter, the end goal of QoS is to provide better and more predictable network services by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics as required by the business applications. Therefore, understanding what the voice and video QoS requirements are will help in deploying the other aspects of the QoS solution. The end goal of a properly designed voice and video network solution should be to maintain a delay of less than 150 ms, jitter less than 30 ms, and a packet loss of less than 1 percent.

The first step in designing a QoS solution is to identify and establish the QoS trust boundaries. Then you will need to set up the congestion management and the congestion avoidance. Additional tools that will need to be configured include policing and shaping. Optionally, you might want to also deploy some link efficiency methods. All of these topics will be discussed in depth in this section of the chapter. It is also important to bear in mind that QoS parameters will not be enforced until there is congestion over the network.

QoS Trust Boundaries

Key Topic

When it comes to QoS, it is best practice to mark packets as close to the source as possible. Most devices, such as computers and servers, cannot mark their own packets and should not be trusted even if they can. Cisco phones, however, can mark their own packets and can be trusted with the QoS markings they provide. Therefore, QoS trust boundaries should be set up so that the switch will trust the QoS markings that phones place on their own packets. Layer 2 QoS uses a mechanism called class of service (CoS), which operates on the 802.1Q VLAN. Unlike Layer 3 QoS mechanisms, CoS does not ensure network performance or guarantee priority in packets being delivered. Therefore, after packets are marked with CoS, they will need to be converted to DSCP using the cos-to-dscp map, which is built into all Cisco switches. By default, QoS on a Cisco access switch is disabled. Once enabled, the switch does not trust QoS settings from a phone. Two simple commands can be entered under the global menu on a switch to enable QoS and change the trust boundary. Once it is enabled, you can use a **show** command to verify these settings. Example 13-1 illustrates the QoS enable and trust boundary commands and the **show** verification command.

Example 13-1 *QoS Enable and Trust Boundary Commands*

```
Switch(config)# mls qos
Switch(config)# interface fastethernet 0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# end
Switch# show mls qos interface fastethernet 0/1
FastEthernet0/1
trust state: trust cos
trust mode: trust cos
cos override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
```

Obviously, this is the simplest design, and there are many other concerns to consider, along with many other settings that can be configured. This example is intended to provide a basic understanding of QoS at the Layer 2 level. For more information on QoS, refer to the "Enterprise QoS Solution Reference Network Design Guide" available at https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html.

Congestion Management

The WAN presents the greatest potential to be a bottleneck point within an enterprise network. WAN Internet speeds are usually much slower than LAN speeds; however, the network switches use hardware-based buffers, which compared to the interface speed are much smaller than those found on WAN interfaces in routers. This merely increases the potential for even short-lived traffic bursts on the LAN to cause buffer overflow and dropped packets. For these reasons, different QoS tools are available at the WAN than what exist within the LAN; however, all of these tools work together to provide an overall QoS implementation that has been carefully designed. The content that follows examines some of the key concepts to a proper QoS implementation.

Three models of QoS can be implemented in a Cisco network design:

- **Best Effort model:** This model uses no QoS and does not guarantee that packets will be delivered. Obviously, this is not the model most companies would choose to implement.
- **IntServ model:** This model uses RSVP to guarantee predictable behavior on the network for applications that have specific bandwidth and delay requirements.
- **DiffServ model:** This model operates on classes that require special QoS treatment. It is the model that has been discussed up to this point and will continue to be the focus of this dialogue. QoS components used in the DiffServ model include classification, marking, congestion management, congestion avoidance, policing and shaping, and link efficiency.

Classification was discussed previously in the examination of class models. Markings were discussed briefly with class models and Layer 2 CoS marking and Layer 2 to Layer 3 conversion mapping. A detailed explanation of these markings is available later in the section titled "Traffic Classifications."

Congestion management mechanisms use the marking on each packet to determine in which queue to place packets. Different queues are given different treatment by the queueing algorithm that is based on the class of packets in the queue. Generally, queues with high-priority packets receive preferential treatment. The Cisco IOS software for congestion management, or queuing, includes the following queuing methods:

- **First-In First-Out (FIFO):** Performs no prioritization of data packets on user data traffic. It entails no concept of priority or classes of traffic. When FIFO is used, ill-behaved sources can consume available bandwidth, bursty sources can cause delays in time-sensitive or important traffic, and important traffic may be dropped because less important traffic fills the queue.

- **Priority Queue (PQ):** Guarantees strict priority in that it ensures that one type of traffic will be sent, possibly at the expense of all others. For PQ, a low-priority queue can be detrimentally affected, and, in the worst case, never allowed to send its packets if a limited amount of bandwidth is available or if the transmission rate of critical traffic is high.
- **Custom Queuing (CQ):** Guarantees some level of service to all traffic because bandwidth can be allocated to all classes of traffic. You can define the size of the queue by determining its configured packet-count capacity, thereby controlling bandwidth access.
- **Weighted Fair Queuing (WFQ):** Does not require configuration of access lists to determine the preferred traffic on a serial interface. Rather, the fair queue algorithm dynamically sorts traffic into messages that are part of a conversation.
- **Class-Based Weighted Fair Queuing (CBWFQ):** Provides class bandwidth guarantee for user-defined traffic classes. It provides flow-based WFQ support for nonuser-defined traffic classes.
- **Low-Latency Queuing (LLQ):** A congestion management mechanism developed by Cisco to bring strict priority queuing (PQ) to Class-Based Weighted Fair Queuing (CBWFQ). LLQ allows delay-sensitive data, such as voice and video, to be given preferential treatment over other traffic by letting this data be dequeued and sent first.

Congestion Avoidance

Congestion avoidance mechanisms monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks. Congestion avoidance is achieved through packet dropping. Congestion avoidance mechanisms are typically implemented on output interfaces where a high-speed link feeds into a lower-speed link, such as a LAN feeding into a WAN. Weighted Random Early Detection (WRED) is an early detection congestion avoidance mechanism that ensures high-precedence traffic has lower loss rates than other traffic during times of congestion. WRED is not recommended for voice and video queues, and the network should not be designed to drop voice and video packets.

Policing

Policing is used to condition traffic before transmitting or receiving through the network. Policing controls traffic bursts by marking or dropping packets when predefined limits are reached. Policing mechanisms can drop traffic classes that have lower QoS priority markings. Policing tools include class-based policing and committed access rate (CAR). CAR services limit the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.

Shaping

Shaping mechanisms are used on output interfaces to help smooth out mismatches in the network and limit transmission rates. Although these mechanisms are typically used to limit the flow from a high-speed link to a low-speed link, shaping could also be used to manage the flow of traffic at a point in the network where multiple flows are aggregated. Cisco IOS software uses two traffic-shaping tools called Generic Traffic Shaping (GTS) and Frame Relay Traffic Shaping (FRTS) to manage traffic and congestion on the network.

Link Efficiency Methods

Link efficiency methods reduce the overhead that is associated with voice and video transportation. These bandwidth-saving mechanisms, such as compression and link fragmentation and interleaving (LFI), help support large amounts of traffic over a slower link. Compression is one of the link efficiency mechanisms that work in conjunction with queuing and traffic shaping to manage existing bandwidth more efficiently and predictably. Two types of compression are available. Compression of the payload of Layer 2 frames can be implemented using the Stacker or Predictor algorithm. The other compression available is cRTP, which can compress the IP, UDP, and RTP headers down from 40 bytes to 2–4 bytes. Compression should be used only on slow WAN links because the drawback is the consumption of computational resources on a hop-by-hop basis. Another link efficiency method is link fragmentation and interleaving. Interactive traffic, such as voice and video, is susceptible to increased latency and jitter when the network processes large packets. This susceptibility increases as the traffic is queued on slower links. LFI can reduce delay and jitter on slower-speed links by breaking up large data packets and interleaving low-delay traffic packets with the resulting smaller packets. LFI is typically used on slow WAN links to ensure minimal delay for voice and video traffic. Figure 13-3 summarizes all of the QoS components needed for an efficient network deployment.

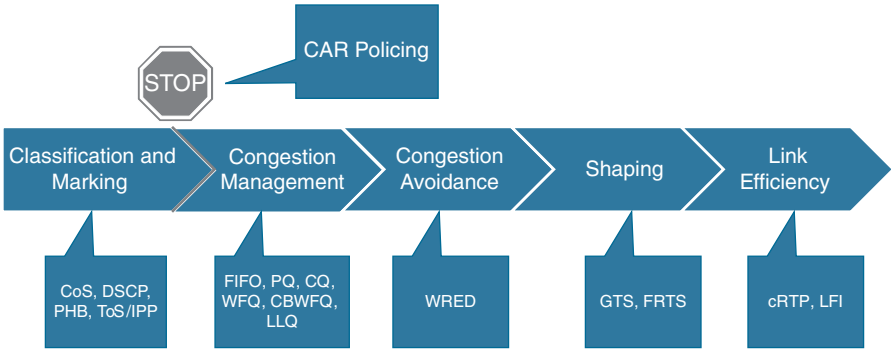

Figure 13-3 *QoS Components*

Traffic Classifications

QoS classification initially takes place in the Access layer; however, other QoS tools are needed to ensure voice and video quality is maintained throughout the network. In addition to traffic classification, queuing and bandwidth provisioning also ensure voice and video quality.

LAN and WAN Traffic Classifications

As mentioned previously, Layer 2 classification uses CoS markings for packets at the Access layer. The Distribution and Core layers use the existing CoS markings to map QoS to Layer 3 classifications. These Layer 3 classifications include differentiated services code point (DSCP), per hop behavior (PHB), and type of service (ToS) or IP Precedence (IPP). Table 13-2 summarizes each of these classifications and how they map to one another for different applications.

Key Topic

Table 13-2 Traffic Classification Map

Application	Layer 3 Classification			Layer 2 Classification
	ToS/IPP	PHB	DSCP	CoS
Routing	6	CS6	48	6
Voice Only	5	EF	46	5
Voice/Video	4	AF41	34	4
Telepresence Video	4	CS4	32	4
Streaming Video	3	CS4	32	3
Call Signaling	3	CS3	24	3
Transactional Data	2	AF21	18	2
Network Management	2	CS2	16	2
Bulk Data	1	AF11	10	1
Scavenger	1	CS1	8	1
Best Effort	0	0	0	0

In Table 13-2, voice only and voice with video traffic are separated into two different categories. The reason for this is so that when a video call is placed, both voice and video packets reach the destination at roughly the same time. There is no support for lip synchronization in SIP, so this will help preserve lip-syncing during the call. Also, when a video call is placed, the expectation is that both the voice and video media will share the same quality during the call. These reasons are why Cisco recommends deploying QoS in this manner. Notice, however, that this table is slightly different from the QoS Baseline model, although there are 11 classes in both design models. Table 13-2 provides a classification for Voice/Video and another for Telepresence Video. Cisco offers different endpoints that will utilize different QoS classifications based on type. Based on Cisco's current endpoint portfolio, the Voice/Video applications include the 8845, 8865, and Jabber client endpoints. If you have a DX endpoint running the legacy Android software, it would fall under this category too. The Telepresence Video applications include all endpoints running CE or CTS software, including the DX, MX, SX, IX, and Webex endpoints.

After packets have been classified for Layer 2 and Layer 3, the next step is to queue traffic based on the classification. Transmit interface buffers within a campus tend to congest in small, finite intervals as a result of the bursty nature of network traffic. When this congestion occurs, any packets destined for that transmit interface are dropped. The only way to prevent dropped voice traffic is to configure multiple queues on campus switches. By enabling multiple queues on campus switches, you can configure all voice traffic to use separate queues, thus virtually eliminating the possibility of dropped voice packets when an interface buffer fills instantaneously. For this reason, Cisco recommends always using a switch that has at least two output queues on each port and the ability to send packets to these queues based on QoS Layer 2 or Layer 3 classification. The majority of Cisco Catalyst switches support two or more output queues per port.

WLAN Traffic Classifications

Just as QoS is necessary for the LAN and WAN wired network infrastructure in order to ensure high voice and video quality, QoS is also required for the wireless LAN infrastructure. Because of the bursty nature of data traffic and the fact that real-time traffic such as voice and video are sensitive to packet loss and delay, QoS tools are required to manage wireless LAN buffers; limit radio contention; and minimize packet loss, delay, and delay variation. Unlike most wired networks, however, wireless networks are a shared medium, and wireless endpoints do not have dedicated bandwidth for sending and receiving traffic. While wireless endpoints can mark traffic with 802.1p CoS, ToS, DSCP, and PHB, the shared nature of the wireless network means limited admission control and access to the network for these endpoints.

As with the wired network infrastructure, it is important to classify or mark pertinent wireless traffic as close to the edge of the network as possible. Because traffic marking is an entrance criterion for queuing schemes throughout the wired and wireless network, marking should be done at the wireless endpoint device whenever possible. Marking or classification by wireless network devices should be identical to that for wired network devices. In accordance with traffic classification guidelines for wired networks, the Cisco wireless endpoints mark voice media traffic or voice RTP traffic with DSCP 46 (or PHB EF), video media traffic or video RTP traffic with DSCP 34 (or PHB AF41), and call control signaling traffic (SCCP or SIP) with DSCP 24 (or PHB CS3). Once this traffic is marked, it can be given priority of better than best-effort treatment and queuing throughout the network. All wireless voice and video devices that are capable of marking traffic should do so in this manner. All other traffic on the wireless network should be marked as best-effort or with some intermediary classification as outlined in wired network marking guidelines. If the wireless voice or video devices are unable to do packet marking, alternate methods such as port-based marking should be implemented to provide priority to video and voice traffic.

Key Topic

While 802.1p and differentiated services code point (DSCP) are the standards to set priorities on wired networks, 802.11e is the standard used for wireless networks. This is commonly referred as user priority (UP), and it is important to map the UP to its appropriate DSCP value. Table 13-3 compares the 802.11e QoS values compared to wired QoS values.

Key Topic

Table 13-3 QoS Value Comparison with 802.11e

Traffic Type	DSCP (PHB)	802.1p UP	802.11e UP
Voice	46 (EF)	5	6
Video	34 (AF41)	4	5
Voice and Video Signaling	24 (CS3)	3	4

Key Topic

After traffic marking has occurred, it is necessary to enable the wired network access points (APs) and devices to provide QoS queuing so that voice and video traffic types are given separate queues to reduce the chances of this traffic being dropped or delayed as it traverses the wireless LAN. Queuing on the wireless network occurs in two directions: upstream and downstream. Upstream queuing concerns traffic traveling from the wireless endpoint up to the AP, and from the AP up to the wired network. Downstream queuing concerns traffic traveling from the wired network to the AP and down to the wireless endpoint.

For upstream queuing, devices that support Wi-Fi Multimedia (WMM) are able to take advantage of queueing mechanisms, including priority queueing. As for downstream QoS, Cisco APs currently provide up to eight queues for downstream traffic being sent to wireless clients. The entrance criterion for these queues can be based on a number of factors, including DSCP, access control lists (ACLs), and VLAN. Although eight queues are available, Cisco recommends using only two queues when deploying wireless voice. All voice media and signaling traffic should be placed in the highest-priority queue, and all other traffic should be placed in the best-effort queue. This ensures the best possible queuing treatment for voice traffic.

To set up this two-queue configuration for autonomous APs, you can create two QoS policies on the AP. Name one policy Voice and configure it with the class of service *Voice < 10 ms Latency (6)* as the Default Classification for all packets on the VLAN. Name the other policy Data and configure it with the class of service *Best Effort (0)* as the Default Classification for all packets on the VLAN. Then assign the Data policy to the incoming and outgoing radio interface for the data VLAN(s) and assign the Voice policy to the incoming and outgoing radio interfaces for the voice VLAN(s). With the QoS policies applied at the VLAN level, the AP is not forced to examine every packet coming in or going out to determine the type of queuing the packet should receive.

For lightweight APs, the WLAN controller has built-in QoS profiles that can provide the same queuing policy. Voice VLAN or voice traffic is configured to use the Platinum policy, which sets priority queueing for the voice queue. Data VLAN or data traffic is configured to use the Silver policy, which sets best-effort queuing for the Data queue. These policies are then assigned to the incoming and outgoing radio interfaces based on the VLAN. The preceding configurations ensure that all voice and video media and signaling are given priority queuing treatment in a downstream direction.

To avoid exceeding the capacity limit of a given AP channel, some form of call admission control is required. Cisco APs and wireless Unified Communications clients now use Traffic Specification (TSPEC) instead of QoS Basic Service Set (QBSS) for call admission control. Wi-Fi Multimedia Traffic Specification (WMM TSPEC) is the QoS mechanism that enables WLAN clients to provide an indication of their bandwidth and QoS requirements so that APs can react to those requirements. When a client is preparing to make a call, it sends an Add Traffic Stream (ADDTS) message to the AP with which it is associated, indicating the TSPEC. The AP can then accept or reject the ADDTS request based on whether bandwidth and priority treatment are available. If the call is rejected, the client receives a Network Busy message. If the client is roaming, the TSPEC request is embedded in the reassociation request message to the new AP as part of the association process, and the TSPEC response is embedded in the reassociation response. Alternatively, endpoints without WMM TSPEC support, but using SIP as call signaling, can be managed by the AP. Media snooping must be enabled for the Service Set Identifier (SSID). The client's implementation of SIP must match that of the wireless LAN controller, including encryption and port numbers.

Configure and Verify LLQ

Without LLQ, CBWFQ provides WFQ based on defined classes with no strict priority queue available for real-time traffic. CBWFQ allows you to define traffic classes and then assign characteristics to that class, such as designating the minimum bandwidth delivered to the class during congestion. For CBWFQ, the weight for a packet belonging to a specific

class is derived from the bandwidth you assigned to the class when you configured it. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced fairly based on weight; no class of packets may be granted strict priority. This scheme poses problems for voice traffic that is largely intolerant of delay, especially variation in delay. For voice traffic, variations in delay introduce irregularities of transmission manifesting as jitter in the heard conversation.

LLQ provides strict priority queueing for CBWFQ, reducing jitter in voice conversations. Strict PQ allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued. Configured by the **priority** command, LLQ enables use of a single, strict priority queue within CBWFQ at the class level, allowing you to direct traffic belonging to a class to the CBWFQ strict priority queue. To enqueue class traffic to the strict priority queue, you specify the named class within a policy map and then configure the **priority** command for the class. (Classes to which the **priority** command is applied are considered priority classes.) Within a policy map, you can give one or more classes priority status. When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is enqueued to the same, single, strict priority queue.

One of the ways in which the strict PQ used within CBWFQ differs from its use outside CBWFQ is in the parameters it takes. Outside CBWFQ, you can use the **ip rtp priority** command to specify the range of UDP ports whose voice traffic flows are to be given priority service. Using the **priority** command, you are no longer limited to a UDP port number to stipulate priority flows because you can configure the priority status for a class within CBWFQ. Instead, all of the valid match criteria used to specify traffic for a class now apply to priority traffic. These methods of specifying traffic for a class include matching on access lists, protocols, and input interfaces. Moreover, within an access list, you can specify that traffic matches are allowed based on the IP differentiated services code point (DSCP) value that is set using the first six bits of the ToS byte in the IP header.

Although it is possible to enqueue various types of real-time traffic to the strict priority queue, Cisco strongly recommends that you direct only voice traffic to it because voice traffic is well behaved, whereas other types of real-time traffic are not. Moreover, voice traffic requires that delay be nonvariable in order to avoid jitter. Real-time traffic such as video could introduce variation in delay, thereby thwarting the steadiness of delay required for successful voice traffic transmission.

To configure network traffic marking, you use the modular quality of service (QoS) command-line interface (CLI), also referred to as MQC. The MQC is a CLI structure that allows you to complete the following tasks:

1. Specify the matching criteria used to define a traffic class.
2. Create a traffic policy (policy map). The traffic policy defines the QoS policy actions to be taken for each traffic class.
3. Apply the policy actions specified in the policy map to an interface, subinterface, or ATM PVC by using the **service-policy** command.

Figure 13-4 illustrates the process used to configure QoS within a Cisco environment.

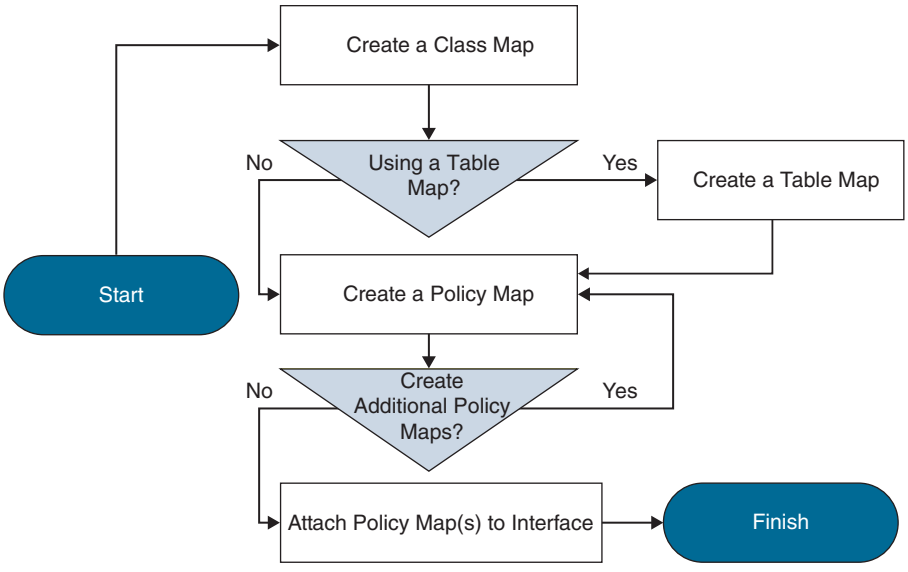

Figure 13-4 *QoS Configuration Process*

Class Map

Traffic classification and traffic marking are closely related and can be used together. Traffic marking can be viewed as an additional action, specified in a policy map, to be taken on a traffic class. Traffic classification allows you to organize packets into traffic classes on the basis of whether the traffic matches specific criteria. For example, all traffic with a CoS value of 2 is grouped into one class, and traffic with DSCP value of 3 is grouped into another class. The match criterion is user defined.

After the traffic is organized into traffic classes, traffic marking allows you to mark an attribute for the traffic belonging to that specific class. For instance, you may want to change the CoS value from 2 to 1, or you may want to change the DSCP value from 3 to 2. The match criteria used by traffic classification are specified by configuring a match command in a class map. The marking action taken by traffic marking is specified by configuring a set command in a policy map. These class maps and policy maps are configured using the MQC.

Key Topic

The following commands explain how to create a class map to define traffic classes. Within the class map, the appropriate match command is used to specify the matching criteria for the traffic classes. To create the class map and specify the matching criteria, complete the following steps:

```
Router> enable

Router# configure terminal

Router(Config)# class-map class-map-name
```

From this point, you can use a few options available within the class map to establish the search criterion against which packets are checked to determine if they belong to the class. Table 13-4 identifies four of these criteria.

Key Topic

Table 13-4 Four Criterion Matches for Packet Classification

MQC Command	Description
Router(config-cmap)# **match access-group** *access-group-name*	Specifies the name of the access control list (ACL) against whose contents packets are checked to determine if they belong to the class.
Router(config-cmap)# **match input-interface** *interface-name*	Specifies the name of the input interface used as a match criterion against which packets are checked to determine if they belong to the class.
Router(config-cmap)# **match protocol** *protocol*	Specifies the name of the protocol used as a match criterion against which packets are checked to determine if they belong to the class.
Router(config-cmap)# **match fr-dlci** *dlci-number*	Specifies the Frame Relay DLCI number as a match criterion against which packets are checked to determine if they belong to the class.

Policy Map

Key Topic

As previously mentioned, creating a table map is not required unless the desired outcome is to change some of the CoS or DSCP values. The table map contains the mapping scheme used for establishing the to-from relationship and equivalency between one traffic-marking value and another. The table map can be configured for use with multiple policy maps. The policy maps can then be configured to convert and propagate the traffic-marking values defined in the table map. Then the policy maps can be attached to the input or output interface of either the ingress or egress router, as appropriate, to serve the QoS requirements of your network. To create and configure the table map, enter the following MQC commands:

```
Router> enable

Router# configure terminal

Router# table-map name map from from-value to to-value
        [default default-action-or-value]
```

To configure a policy map and create class policies that make up the service policy, begin with the **policy-map** command to specify the policy map name. Then use one or more of the following commands to configure the policy for a standard class or the default class:

- **priority**
- **bandwidth**
- **queue-limit** or **random-detect**
- **fair-queue** (for class-default class only)

For each class that you define, you can use one or more of the commands listed to configure the class policy. For example, you might specify bandwidth for one class and both bandwidth and queue limit for another class. The default class of the policy map, commonly known as the class-default class, is the class to which traffic is directed if that traffic does not satisfy the match criteria of the other classes defined in the policy map.

You can configure class policies for as many classes as are defined on the router, up to the maximum of 64. However, the total amount of bandwidth allocated for all classes in a policy map must not exceed the minimum committed information rate (CIR) configured for the virtual circuit (VC) minus any bandwidth reserved by the **frame-relay voice bandwidth** and **frame-relay ip rtp priority** commands. If the minimum CIR is not configured, the bandwidth defaults to one-half of the CIR. If all of the bandwidth is not allocated, the remaining bandwidth is allocated proportionally among the classes on the basis of their configured bandwidth.

Key Topic

The following commands describe how to create and configure a policy map to use the class map and the table map. The policy map applies the appropriate QoS feature to the network traffic based on the traffic classification. To configure class policies in a policy map, use the MQC commands described in the following sections.

```
Router> enable

Router# configure terminal

Router(Config)# policy-map name

Router(Config-pmap)# class {class-name | class-default}

Router(Config-pmap-c)# set cos cos-value
```

or

```
Router(Config-pmap-c)# set cos dscp table name
```

The **policy-map** *name* command creates a policy map by the name provided and enters the policy-map configuration mode. The **class** command specifies the name of the class whose policy you want to create and enters policy-map class configuration mode. This class is associated with the class map created earlier. You can either enter the name of the class created earlier or enter the **class-default** keyword. The **class-default** class is used to classify traffic that does not fall into one of the defined classes. Even though the class-default class is predefined when you create the policy map, you still have to configure it. If a default class is not configured, then traffic that does not match any of the configured classes is given best-effort treatment, which means that the network will deliver the traffic if it can, without any assurance of reliability, delay prevention, or throughput. The **set cos** command and **set cos dscp table** *name* commands are examples of the **set** commands that can be used when marking traffic. Other **set** commands can be used as well. For a list of other **set** commands, refer to the "Cisco IOS Quality of Service Solutions Configuration Guide" at Cisco.com.

Other class policies can be configured here as well. To configure a class policy for a priority queue, enter the following command:

```
Router (config-pmap-c)# priority bandwidth in kbps
```

This command creates a strict priority class and specifies the amount of bandwidth in kbps to be assigned to the class.

To configure a class policy using a specified bandwidth, enter the following command:

```
Router (config-pmap-c)# bandwidth bandwidth in kbps
```

This command specifies the amount of bandwidth to be assigned to the class in kbps, or as a percentage of the available bandwidth. Bandwidth must be specified in kbps or as a percentage consistently across classes. Bandwidth of the priority queue must be specified in kbps.

To configure a class policy that specifies a number of queues, enter the following command:

```
Router (config-pmap-c)# fair-queue number-of-dynamic-queues
```

This command specifies the number of dynamic queues to be reserved for use by flow-based WFQ running on the default class. The number of dynamic queues is derived from the bandwidth of the interface.

You can create and configure as many policy maps as you need for your network. To create and configure additional policy maps, repeat the preceding commands per your network deployment plan. After all the policy maps have been created, they will need to be attached to the appropriate interface. The next section on service policy will explain how to attach the policy maps to the interfaces.

Service Policy

After you create the policy map, you must attach it to an interface. Policy maps can be attached to either the input or output direction of the interface. Depending on the needs of your network, you can attach policy maps to an interface, a subinterface, or an ATM permanent virtual circuit (PVC). To attach the policy map, enter the following commands in the MQC:

```
Router> enable

Router# configure terminal

Router(Config)# interface type number [name-tag]

Router(Config-if)# service-policy {input | output}
policy-map-name
```

This last command is the key. It attaches the specified service policy map to the output interface and enables LLQ, assuming LLQ has been configured throughout this process. Next, all these QoS settings will need to be verified.

Verify and Monitor LLQ Settings

After the QoS designs have been finalized and the proof of concept tested, it is vital to ensure that the networking team thoroughly understand the QoS features and syntax before enabling features on production networks. Such knowledge is critical for both rollout and subsequent troubleshooting of QoS-related issues. Furthermore, it is recommended to schedule network downtime in order to roll out QoS features. While QoS is required end-to-end, it does not have to be deployed end-to-end at a single instance. A pilot network-segment can be selected for an initial deployment, and pending observation, the rollout can be expanded in stages to encompass the entire enterprise. A rollback strategy is always recommended, to address unexpected issues arising from the QoS deployment.

From the Privileged Exec Mode on the router, you can use several **show** commands to verify that all the QoS settings have been configured correctly. Table 13-5 illustrates these **show** commands with a description of what information each command displays.

Key Topic

Table 13-5 IOS Router Show Commands for QoS

Command	Description
Router# **show policy-map**	Displays all configured policy maps.
Router# **show policy-map** *policy-map* **class** *class-name*	Displays the configuration for the specified class of the specified policy map. Enter the policy map name and the class name.
Router# **show frame-relay pvc dlci**	Displays statistics about the PVC and the configuration of classes for the policy map on the specified data-link connection identifier (DLCI).
Router# **show policy-map interface** *interface-name*	When LLQ is configured, displays the configuration of classes for all policy maps.
Router# **show policy-map interface** *interface-name* **dlci**	When LLQ is configured, displays the configuration of classes for the policy map on the specified DLCI.
Router# **show policy-map interface** *interface-name*	Displays traffic statistics of all classes configured for all service policies on the specified interface, subinterface, or PVC on the interface. When a policy map has multiple instances of the same class, and this policy map is attached to an interface, the following command returns only the first instance: **show policy-map interface** *interface_name* **output class** *class-name*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 13-6 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 13-6 Key Topics for Chapter 13

Key Topic Element	Description	Page Number
Paragraph	Four Factors Requiring QoS	307
Section	Latency, Jitter, and Packet Loss	307
Paragraph	Calculating Bandwidth for Voice Packets	308
Paragraph	Calculating Bandwidth for Video Calls	309
List	Three QoS Strategy Models	309
List	Five Traffic Classes of QoS Markings	310
List	Eight Traffic Classes of QoS Markings	310
List	Eleven Traffic Classes of QoS Markings	311
Paragraph	End goal of a properly designed QoS deployment for voice and video	313
Paragraph	COS and Trust Boundaries	313
List	Three models of QoS	314
Table 13-2	Traffic Classification Map	317
Paragraph	802.11e Explained	318
Table 13-3	QoS Value Comparison with 802.11e	318
Paragraph	Upstream and Downstream Queueing for Wireless Networks	318
Commands	Commands to Create a Class Map	321
Table 13-4	Four Criterion for Packet Classification	322
Commands	Commands to Create a Table Map	322
Commands	Commands to Configure Class Policies in a Policy Map	323
Commands	Commands to Attach a Policy Map to an Interface	324
Table 13-5	IOS Router Show Commands for QoS	325

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

802.11e, ADDTS, CAR, CBWFQ, Compression, Congestion Avoidance, Congestion Management, CoS, CQ, cRTP, DiffServ, DSCP, FIFO, FRTS, GTS, IntServ, IPP, Jitter, Latency, LFI, Link Efficiency, LLQ, OAM, Packet Loss, PHB, Policing, PQ, PVC, Shaping, SSID, Stacker, ToS, TSPEC, WFQ, WMM TSPEC, WRED

Command Reference to Check Your Memory

This section includes the most important configuration and EXEC commands covered in this chapter. It might not be necessary to memorize the complete syntax of every command, but you should be able to remember the basic keywords that are needed.

To test your memory of the commands, cover the left side of Tables 13-7 through 13-8 with a piece of paper, read the description on the right side, and then see how much of the command you can remember.

The 350-801 exam focuses on practical, hands-on skills that are used by a networking professional. Therefore, you should be able to identify the commands needed to configure and test QoS.

Table 13-7 Layer 2 QoS Commands

Command Syntax	Task
Switch(config)# **mls qos**	Enables QoS on the Layer 2 switch and CoS to DSCP mapping.
Switch(config-if)# **mls qos trust cos**	Establishes a trust boundary between a phone and the switch for QoS. Must be enabled on the actual switchport.
Switch# **show mls qos interface fastethernet 0/1**	Reveals that QoS has been enabled and a QoS trust boundary has been set up.

Table 13-8 Layer 3 QoS Commands

Command Syntax	Task
Router(Config)# **class-map** *class-map-name*	Creates a class map to be used for matching traffic to a specified class and enters class-map configuration mode. The class map name must be specified after the **class-map** command.
Router(config-cmap)# **match access-group** *access-group-name*	Specifies the name of the access control list (ACL) against whose contents packets are checked to determine if they belong to the class.
Router(config-cmap)# **match input-interface** *interface-name*	Specifies the name of the input interface used as a match criterion against which packets are checked to determine if they belong to the class.
Router(config-cmap)# **match protocol** *protocol*	Specifies the name of the protocol used as a match criterion against which packets are checked to determine if they belong to the class.

Command Syntax	Task
Router(config-cmap)# **match fr-dlci** *dlci number*	Specifies the Frame Relay DLCI number as a match criterion against which packets are checked to determine if they belong to the class.
Router(Config)# **table-map** *name* **map from** *from-value* **to** *to-value* [**default** *default-action-or-value*]	Creates a table map using the specified name and enters table-map configuration mode. Enter the name of the table map you want to create. Enter each value mapping on a separate line. Enter as many separate lines as needed for the values you want to map. The **default** keyword and *default-action-or-value* argument set the default value (or action) to be used if a value is not explicitly designated.
Router(config)# **policy-map** *name*	Specifies the name of the policy map created earlier and enters policy-map configuration mode. The policy map name must be entered with this command.
Router(config-pmap)# **class** *name*	Specifies the name of the class whose policy you want to create and enters policy-map class configuration mode. This class is associated with the class map created earlier. Enter the name of the class or enter the **class-default** keyword.
Router(config-pmap)# **class class-default** *name*	Specifies the default class so that you can configure or modify its policy.
Router(config-pmap-c)# **set cos** *cos-value*	(Optional) Sets the CoS value in the type of service (ToS) byte.
Router(config-pmap-c)# **set cos dscp table** *name*	(Optional) If a table map was created earlier, sets the CoS value based on the DSCP value (or action) defined in the table map.
Router(config-pmap-c)# **priority** *bandwidth in kbps*	Creates a strict priority class and specifies the amount of bandwidth, in kbps, to be assigned to the class.
Router (config-pmap-c)# **bandwidth** *bandwidth in kbps*	Specifies the amount of bandwidth to be assigned to the class, in kbps, or as a percentage of the available bandwidth. Bandwidth must be specified in kbps or as a percentage consistently across classes. Bandwidth of the priority queue must be specified in kbps.

Command Syntax	Task
Router (config-pmap-c)# **fair-queue** *number-of-dynamic-queues*	This command specifies the number of dynamic queues to be reserved for use by flow-based WFQ running on the default class. The number of dynamic queues is derived from the bandwidth of the interface.
Router(Config-if)# **service-policy** {**input** | **output**} *policy-map-name*	Attaches the specified service policy map to the output interface and enables LLQ.

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. Name the four factors in a network design that can lead to poor audio and video quality.
2. List the 11 classes in the QoS Baseline model.
3. The end goal of a properly designed voice and video network solution should be to maintain what parameters for delay, jitter, and packet loss?
4. Outline the QoS mapping for 802.11e User Priority to Layer 2 802.1p and Layer 3 DSCP for Voice, Video, and Signaling.
5. List all of the commands, from the switch to the router, required for a QoS deployment. Do not include optional commands.

CHAPTER 14

DNS, NTP, and SNMP

This chapter covers the following topics:

DNS Settings: This topic will explain various DNS settings that need to be configured to support a Cisco collaboration solution, including A-records, SRV records, and PTRs.

NTP Settings: This topic will explain the dependency between Cisco Unified Communications Manager and NTP.

SNMP Settings: This topic will discuss other unified communications components that use SNMP for collecting and organizing information from various devices within the UC environment.

This chapter will look at some of the network-related components that can impact the Cisco Unified Communications environment. Although the Cisco Unified Communications Manager can operate without a Domain Name System (DNS), many services cannot function without DNS in place. The Cisco Unified Communications Manager cannot function at all without Network Time Protocol (NTP), but the level of strata can influence the effectiveness of the Cisco Unified Communications Manager. Other services that can enhance the user experience in a Cisco Unified Communications Manager environment have a dependency on the Simple Network Management Protocol. These topics are all addressed in this chapter in the following sections:

- DNS Settings
  - A/AAAA Records
  - SRV Records
  - Reverse Pointer Record (PTR)
- NTP Settings
- SNMP Settings

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.5 Explain these components to support Cisco Collaboration solutions
  - 1.5.a SNMP
  - 1.5.b DNS

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 14-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 14-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
DNS Settings	1–4
NTP Settings	5
SNMP Settings	6

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a disadvantage to not using DNS in a Cisco Unified Communications Manager environment?
   - **a.** If you lose the connection with DNS, you lose the connection to the server.
   - **b.** Domain verification certificates cannot be used.
   - **c.** Management of the network is simplified.
   - **d.** NAT services do not work without DNS.
2. Which of the following statements best describes an AAAA-record?
   - **a.** AAAA-records allow four different IP addresses to resolve to the same domain.
   - **b.** AAAA-records allow a single IP address to resolve to four domains.
   - **c.** AAAA-records allow 32-bit IP addresses to resolve to a domain.
   - **d.** AAAA-records allow 128-bit IP addresses to resolve to a domain.
3. Which of the following SRV records is used for the Cisco Unified Communications Manager?
   - **a.** _sips._tcp.fqdn. 7200 10 10 5061 cucm.fqdn
   - **b.** _sip._tcp.fqdn. 7200 10 10 5060 cucm.fqdn
   - **c.** _cisco-uds._tcp.fqdn. 7200 10 10 8443 cucm.fqdn
   - **d.** _cucmlogin._tcp.fqdn. 7200 10 10 8443 cucm.fqdn

4. When manually configuring an RPT, what should you do first as an administrator after logging in to the DNS server and pulling up the menu?
   a. Click the New Pointer (PTR) menu option.
   b. Right-click the domain you want to create the PTR under and select New Pointer (PTR).
   c. Right-click the reverse lookup zone you want to create the PTR under and select New Pointer (PTR).
   d. Right-click the domain you want to create the PTR under and select Other New Record.
5. What stratum level does Cisco recommend using when installing a Cisco Unified Communications Manager Publisher?
   a. Stratum 1
   b. Stratum 3
   c. Stratum 4
   d. Stratum 9
6. Which of the following devices does not use SNMP to collect information from other devices?
   a. Cisco Unified Communications Manager
   b. Cisco Emergency Responder
   c. Cisco Paging Server
   d. Cisco Unified CVP

Foundation Topics

DNS Settings

There is much to be said about the Domain Name System (DNS), the different record types that can be used, and the different settings for each record type. Entire books have been dedicated to the extent of capabilities within DNS, which can be referenced for more information. I recommend *DNS Bind* as a reference resource that should exist in every IT engineer's library. This chapter will not go into the same detail as *DNS Bind* concerning DNS settings, but a basic explanation of A-records, SRV records, and reverse proxy settings is warranted.

Cisco collaboration products, such as the Cisco Unified Communications Manager, can use IP addresses or names to refer to other IP devices in application settings. When names are used, DNS needs to resolve them to IP addresses. Both methods have some advantages and disadvantages.

When using IP addresses, the systems do not depend on a DNS server, which can prevent loss of service when the DNS server cannot be reached. When a device initiates a connection to a server for the first time, the time that is required to establish the connection is shorter because a DNS query—a DNS lookup sent to the DNS server followed by a DNS reply sent back from the server—is not required. When the need for DNS is eliminated, there is no danger of errors that are caused by DNS misconfiguration. Troubleshooting is simplified because there is no need to verify proper name resolution. A big disadvantage is related

to certificate security. Many certificates used in the Cisco collaboration solution require domain verification, which in turn requires DNS. This codependency between the certificate server and DNS might require using DNS in your environment.

When DNS is used in the Cisco collaboration solution, management is simplified because logical names are simpler to manage than 32-bit addresses. If IP addresses change, there is no need to modify the application settings because they can still use the same names; only the DNS server configuration has to be modified in this case. IP addresses of Cisco Unified Communications Manager servers can be translated toward IP phones because the IP phone configuration files include server names, not the original server IP address, which should appear differently to the IP phone. As long as these names are resolved to the correct IP address when Cisco Unified IP phones send out DNS requests, the NAT is no problem. Also, most IP clients cache the IP address information that is received from the DNS servers to avoid subsequent name resolution requests for the same name. Although DNS provides an additional point of failure caused by configuration errors or unavailability of the service, Cisco recommends using DNS within the Cisco collaboration solution. This is due to the increased use of certificates for secure communication across networked devices.

By default, the Cisco Unified Communications Manager propagates the machine name and not the IP addresses of its active *Cisco CallManager Services*. These host names are part of TFTP configuration files for devices such as IP phones. DNS reliance refers to the requirement for IP phones to use DNS servers to resolve host names of Cisco CallManager Services. Some situations might require administrators to remove DNS reliance from the Cisco Unified Communications Manager. To remove DNS reliance, navigate to **System > Server** in Cisco Unified CM Administration, choose each available server from the list, and change the server name to the IP address. By default, host names are also used in phone URLs. When DNS reliance is removed, host names that are used in these phone URLs must also be replaced by IP addresses. Phone URLs are configured by using enterprise parameters. Enterprise parameters and their configuration will be explained in Chapter 15, "Cisco Unified Communications Manager Setup."

A/AAAA-Records

Key Topic

An A-record in DNS is a type of lookup record that resolves 32-bit IPv4 addresses to URLs. A Uniform Resource Locator, or URL, is a string of numbers, characters, or letters that identifies a resource. URLs are more commonly known as website addresses. If a user were to put the URL www.cisco.com into a web browser, DNS would resolve this URL to an IPv4 address. The web browser could now forward the query on to the resource at that associated IP address. A-records in DNS contain a host, domain, URL, and an IP address. The URL is the host and domain together. DNS will create the URL field automatically. A sample A-record may look something like this:

Host	Domain	URL	IP Address
cucm_pub	*cisco.com*	*cucm_pub.cisco.com*	*10.1.1.40*

Similar to a A-records, AAAA-records resolve 128-bit IPv6 addresses to URLs. This should be easy to remember because this record type has four *A*s. If each *A* represents 32 bits, then 32 times 4 equals 128, and IPv6 addresses are 128 bits in length. When DNS is used for B2B

communications over IP, whether A-records, AAAA-records, or both are used, a public DNS should be configured with a registered domain. The URL should resolve to a public IP address assigned to the Expressway-E or CUBE, so that incoming calls can be routed to that server.

Different companies offer different DNS software, and the menus to set up these services might be quite different; however, the DNS settings within each software offering are the same, as previously mentioned. The most common DNS used among businesses is the Microsoft Windows Server DNS services. To configure a DNS A-record using the Microsoft DNS, follow these steps:

Step 1. Log in to the Windows Server hosting the DNS services, and open the DNS application.

Step 2. Right-click the domain you wish to create the A-record under, and select **New Host (A or AAAA)…** from the menu that appears.

Step 3. When the New Host window pops up, enter the host name of the URL in the Name (Uses Parent Domain Name if Blank): field. The Fully Qualified Domain Name (FQDN): field will autopopulate as the host name is being typed.

Step 4. Enter the IP address of the server for which the record is being created. This will be the IPv4 address for A-records or the IPv6 address for AAAA-records.

Step 5. (optional) You can check the box beside the **Create Associated Pointer (PTR) Record** if you want DNS to automatically create a reverse DNS record for this server. This topic will be discussed after SRV records. It is recommended that all A-records have a reverse DNS record created as well.

Step 6. Once all fields above have been populated, click the **Add Host** button to create the A-record. The New Host popup will remain open so that other A-records can be created. When you are finished creating records, close the New Host window. Figure 14-1 illustrates the menus used to create DNS A-records as indicated in the preceding steps.

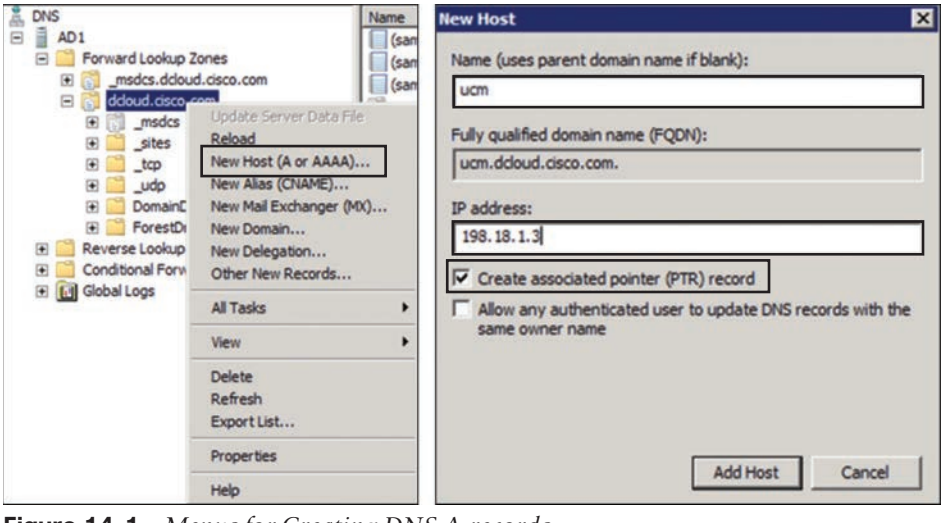

Figure 14-1 *Menus for Creating DNS A-records*

SRV Records

Key Topic

The DNS Service record, or SRV record, is a location service within DNS that can be used to identify protocols, port numbers, and host names of servers for particular services. SRV records should be configured after A-records because the A-record is configured as the Target address in an SRV record. An SRV record will always use the following format:

_ service. _ protocol.<fqdn>. TTL Priority Weight Port Target

SRV records must be created for B2B communication over the public Internet. An SRV record is primarily used for port association, so DNS knows how to handle incoming requests on these ports. However, SRV records can also be used for redundancy and load balancing. If both SIP and H.323 are being used, an SRV record must be created for every port that is associated with these two protocols. H.323 uses UDP port 1719 and TCP port 1720. SIP uses UDP port 5060, TCP port 5060, and TLS port 5061. Therefore, five SRV records will need to be created. The Cisco Unified Communications Manager and the IM and Presence servers use special protocols and ports for communication, which will require unique SRV records as well. Using the preceding A-record example, Table 14-2 outlines the different SRV fields that need to be configured for each of these seven SRV records. The TTL, Priority, and Weight are recommended basic setting values. Also, the service _sips. and protocol _tcp. can be replaced with the service _sip. and the protocol _tls. Pay close attention to the underscores and dots. These are essential characters when configuring SRV records.

Key Topic

Table 14-2 SRV Records Needed for SIP, H.323, CUCM, and IMP

_service.	_protocol.	FQDN.	TTL	Priority	Weight	Port	Target
_sips.	_tcp.	cisco.com.	7200	10	10	5061	exp1.cisco.com
_sip.	_tcp.	cisco.com.	7200	10	10	5060	exp1.cisco.com
_sip.	_udp.	cisco.com.	7200	10	10	5060	exp1.cisco.com
_h323ls.	_udp.	cisco.com.	7200	10	10	1719	exp1.cisco.com
_h323cs.	_tcp.	cisco.com.	7200	10	10	1720	exp1.cisco.com
_cisco-uds.	_tcp.	cisco.com.	7200	10	10	8443	ucm.cisco.com
_cuplogin.	_tcp.	cisco.com.	7200	10	10	8443	imp.cisco.com

To configure an SRV record using the Microsoft DNS, follow these steps:

Step 1. Log in to the Windows Server hosting the DNS services, and open the DNS application.

Step 2. Right-click the domain you wish to create the SRV record under, and select **Other New Records…** from the menu that appears.

Step 3. On the next popup window, scroll down to and select the **Service Location (SRV)** menu option.

Step 4. When the New Resource Record window pops up, enter the appropriate information for each field. The following information is an example based on the ucm.cisco.com SRV information from Table 14-2.

- **a.** **Domain:** (auto-populated)
- **b.** **Service:** _cisco-uds.

c. **Protocol:** _tcp.

d. **Priority:** 10

e. **Weight:** 10

f. **Port Number:** 8443

g. **Host Offering This Service:** ucm.cisco.com

Step 5. After all of the fields in Step 4a–g have been populated, click the **OK** button to create the SRV record. The New Resource Record popup will remain open so that other SRV records can be created. Once finished creating records, close the New Resource Record window. Figure 14-2 illustrates the menus used to create DNS SRV records as indicated in the preceding steps.

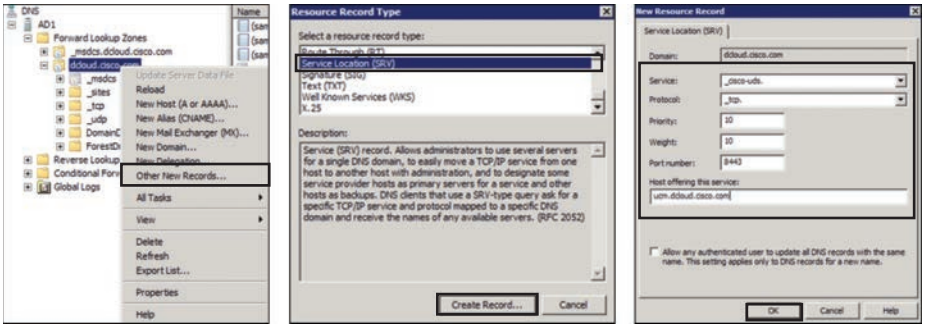

Figure 14-2 *Menus for Creating DNS SRV Records*

Reverse Pointer Record (PTR)

In a DNS lookup, sometimes referred to as a forward lookup, the DNS server resolves a URL to the associated IP address. A-records and AAAA-records are used for forward lookups within DNS. However, sometimes the inverse is required, where the IP address is provided to DNS so that the associated URL can be provided in response. This is referred to as the reverse lookup, and the setting that needs to be configured within DNS to provide reverse lookup is the reverse pointer record, or PTR.

Although many iterations of reverse lookup have been ratified over the years, RFC 2317 outlines Classless IN-ADDR.ARPA, which describes a way to do IN-ADDR.ARPA delegation on nonoctet boundaries for address spaces covering fewer than 256 addresses. In other words, when different companies have leased public IP addresses within the same 256-bit IP range that resolve to different URLs for each company, respectively, classless PTRs can be used to identify domain boundaries so that each company can manage its own domain space without transecting into another company's domain space. This is one of many examples illustrating how PTRs can be used within DNS.

Key Topic

There are no standards that require PTRs to be created for A-records, and there are many A-records that function without PTRs ever being created. However, it is recommended that a PTR be created for each A-record that is created. Reverse pointer records can be configured automatically at the time A-records are configured by simply checking the box titled Create Associated Pointer (PTR) Record. Refer to Figure 14-1 to see this setting. However, the reverse lookup zones that divide the PTRs into their respective groups must be established

before the auto setting will function. In other words, if an administrator selects the Create Associated Pointer (PTR) Record check box while creating an A-record, and no reverse lookup zones for PTRs have been established prior to the A-record creation, then the associated PTR will not be created. In some instances, PTRs can be created manually. To manually create a PTR, follow these steps:

Step 1. Log in to the Windows Server hosting the DNS services, and open the DNS application.

Step 2. Right-click the reverse lookup zone you wish to create the PTR under, and select **New Pointer (PTR)…** from the menu that appears.

Step 3. When the New Resource Record window pops up, enter the Host IP Address in the first field. Note the address that appears in the Fully Qualified Domain Name (FQDN) field. The IP address is listed in reverse order from the way it was originally entered, followed by in-addr.arpa.

Step 4. In the Host Name field, you can enter the A-record URL this PTR is associated with, but it may be more prudent to click the **Browse** button and select the A-record from the list.

Step 5. Once all fields above have been populated, click the **OK** button to create the PTR. The New Resource Record popup will remain open so that other PTRs can be created. When you are finished creating records, close the New Resource Record window. Figure 14-3 illustrates the menus used to create pointer records as indicated in the preceding steps.

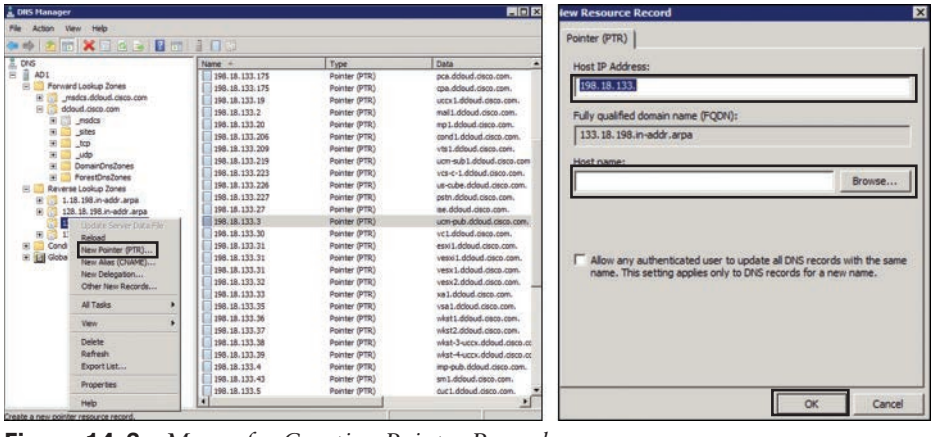

Figure 14-3 *Menus for Creating Pointer Records*

NTP Settings

The Network Time Protocol (NTP) is used to provide common and consistent timestamp information to networked devices. The Cisco Unified Communications Manager uses NTP to obtain time information from a time server; however, only the publisher sends NTP requests to the external NTP server or servers. Subscribers synchronize their time with the publisher.

NTP is a protocol for synchronizing computer system clocks over IP networks. NTP has a hierarchical organization that is based on clock strata. Stratum 0 is an extremely precise clock source, such as an atomic clock or radio clock. A stratum 1 server is directly connected to a stratum 0 clock and can provide time information to other (stratum 2) devices, which in turn serve stratum 3 devices and so on. Cisco Unified Communications Manager typically uses stratum 1 but can be set to a stratum 3 without installation failing, which is the recommended maximum stratum for the Cisco Unified Communications Manager. Companies that try to install the Cisco Unified Communications Manager with a stratum 4 or higher might experience issues during installation. If installation doesn't fail entirely, then production performance will most certainly ensue.

NTP must be enabled and configured during the installation of Cisco Unified Communications Manager. At least one external NTP server must be reachable and functioning when installing the Cisco Unified Communications Manager publisher to complete the installation. Cisco recommends using a minimum of three external NTP servers in a production environment. It is extremely important that all network devices have accurate time information because the system time of Cisco Unified Communications Manager is relevant in the following situations:

- Cisco IP phones display date and time information. This information is obtained from Cisco Unified Communications Manager.
- CDR and CMR, which are used for call reporting, analysis, and billing, include date and time information.
- Alarms and events in log files, as well as trace information in trace files, include time information. Troubleshooting a problem requires correlation of information that is created by different system components, such as Cisco Unified Communications Manager, Cisco IOS gateway, and so on. This problem solving is possible only if all devices in the network have the same correct time information.
- Some Cisco Unified Communications Manager features are date-based or time-based and therefore rely on correct date and time. These features include time-of-day routing and certificate-based security features.
- Certificates include a validity period. If a system that receives a certificate has an invalid future date, it may consider the received certificate to be invalid or expired.
- Endpoints joining a scheduled meeting may also have issues pertaining to when and if they are able to join the scheduled meeting when NTP is out of sync.

To ensure that all network devices have the correct date and time, it is recommended that all network devices use NTP for time synchronization. The master reference clock should be a stratum 0 or stratum 1 NTP server.

SNMP Settings

Simple Network Management Protocol (SNMP) is an Internet standard protocol for collecting and organizing information about managed devices on IP networks and for modifying that information to change device behavior. Devices that typically support SNMP include cable modems, routers, switches, servers, workstations, printers, and more. Within a Cisco collaboration network, there are a lot of different devices that use SNMP.

The primary method for location identification in Cisco Emergency Responder is the detection of an endpoint via Layer 2 discovery at the switch port level. Discovering an endpoint though Layer 2 Cisco Discovery Protocol (CDP) enables Emergency Responder to determine the exact physical location of the calling device based on the physical termination of the network cable to a network jack in a cubicle or office. Although the discovery mechanism of the connected device is reliable, the accuracy of the physical location relies on two main assumptions:

Key Topic

- The wired infrastructure of the enterprise is well established and does not change sporadically, and any wiring closet changes trigger notification to the Emergency Responder administrator indicating what changed.
- The infrastructure is available for Cisco Emergency Responder to browse; that is, Cisco Emergency Responder can establish SNMP sessions to the underlying network infrastructure and can scan the network ports for the discovery of connected phones.

The Cisco Paging Server allows users to send audio-only messages to groups of up to 50 IP phones in an organization. The Cisco Paging Server communicates with the Cisco Unified Communications Manager using SIP, SNMP, AXL, and CTI. When the Cisco Paging Server starts, and at configurable intervals after that, it connects with the Cisco Unified Communications Manager using SNMP. The Cisco Paging Server uses SNMP to find the other Cisco Unified Communications Manager cluster member IP addresses as well as a list of phones registered to each cluster member. Once the SNMP communications are complete, the Cisco Paging Server uses AXL to determine additional information regarding each registered phone, such as device name, description, device pool, calling search space, directory number, and location. This information can be used to build logical groups of phones, called recipient groups. In the Cisco Paging Server, recipient groups can contain a maximum of 50 phones.

Cisco Unified Contact Center Enterprise (UCCE) is managed with SNMP. UCCE devices have a built-in SNMP agent infrastructure that supports SNMP v1, v2c, and v3, and it exposes instrumentation defined by the CISCO-CONTACT-CENTER-APPS-MIB. This MIB provides configuration, discovery, and health instrumentation that can be monitored by standard SNMP management stations. Moreover, UCCE provides a rich set of SNMP notifications that alert administrators of any faults in the system. UCCE also provides a standard syslog event feed for those administrators who want to take advantage of a more verbose set of events. For more information about configuring the UCCE SNMP agent infrastructure and the syslog feed, refer to the SNMP Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted, available at https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-installation-and-configuration-guides-list.html.

Unified CVP health can be monitored by using any SNMP standard monitoring tool to get a detailed visual and tabular representation of the health of the solution network. All Unified CVP product components and most Cisco Unified Customer Voice Portal (CVP) solution components also issue SNMP traps and statistics that can be delivered to any standard SNMP management station or monitoring tool. Cisco Unified Contact Center Express, or UCCX, can also be managed with SNMP and a syslog interface.

Prime Collaboration's automated device discovery is based on a Cisco Discovery Protocol (CDP) table. Ping Sweep may be used instead of CDP, but IP phones discovered using Ping

Sweep are reported in "unmanaged" state. Another protocol that Prime Collaboration uses to monitor the Unified Communications elements is SNMP. SNMP is an application layer protocol using UDP as the transport layer protocol. There are three key elements in an SNMP managed network:

- **Managed devices:** Network devices that have an SNMP agent, such as Cisco Unified Communications Manager, routers, switches, and so on.
- **Agent:** A network management software module that resides in a managed device. This agent translates the local management information on the device into SNMP messages.
- **Manager:** Software running on a management station that contacts different agents in the network to get the management information, such as Prime Collaboration.

The SNMP implementation supports three versions: SNMP v1, SNMP v2c, and SNMP v3. SNMP v3 supports authentication, encryption, and message integrity. SNMP v3 may be used if security is desired for management traffic. Prime Collaboration supports all three versions of SNMP. SNMP v1 and v2c read/write community strings, or SNMP v3 credentials must be configured on each device for agent and manager to communicate properly. Prime Collaboration needs only SNMP read access to collect network device information. SNMP must also be enabled on network devices to allow Prime Collaboration to get information on network devices at configured polling intervals and to receive alerts and faults via trap notification sent by the managed devices. For more information on SNMP, refer to the Cisco Prime Collaboration documentation available at https://www.cisco.com/c/en/us/products/cloud-systems-management/prime-collaboration/index.html.

Cisco TMS can manage endpoints registered to both Cisco Expressway (or VCS) and Cisco Unified Communications Manager. There are two types of device management: direct managed and provisioned. Direct-managed devices are manually added into the Cisco TMS system navigator. Cisco TMS supports 5,000 direct-managed devices. Cisco TMS communicates with the endpoints directly via HTTP or SNMP protocols. When a direct-managed endpoint is registered to the Cisco Unified Communications Manager, the Cisco Unified Communications Manager handles most management capabilities such as software upgrades. When a direct-managed endpoint is registered to Cisco Expressway, Cisco TMS handles management and provisioning of the endpoint, including capabilities such as software upgrades.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 14-3 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 14-3 Key Topics for Chapter 14

Key Topic Element	Description	Page Number
Paragraph	DNS A-record Example and Definition	333
Paragraph	DNS SRV Record Example and Definition	335
Table 14-2	SIP, H.323, CUCM, and IMP Services and Ports for SRV Records	335
Paragraph	Automatic Creation of PTRs	336
List	NTP Relevance in a CUCM Environment	338
List	Dependencies for Emergency Responder Accuracy	339

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

A-records, AAAA-records, AXL, CDP, CER, CTI, DNS, NTP, PTR, SNMP, SRV Record, Stratum, TMS, UCCE, UCCX, Unified CVP, URL, VCS

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. What are the SRV records for the CUCM and IM and Presence servers?
2. List the five important reasons that NTP services must be synchronized correctly.
3. List seven devices that use SNMP in a Cisco UC environment.

CHAPTER 21

Mobile and Remote Access (MRA)

This chapter covers the following topics:

Requirements for MRA: This topic will examine the different prerequisites that must be configured before an MRA solution can be deployed. These requirements include DNS settings, firewall ports and considerations, certificate requirements, HTTPS reverse proxy settings, and the service discovery.

Cisco Unified Communications Manager Settings for MRA: This topic will examine settings that should be configured on the Cisco Unified Communications Manager in support of endpoints registering over MRA.

TLS Verify Requirements: This topic will identify the certificate requirements for an MRA deployment using the Cisco Unified Communications Manager, the Cisco Expressway Core, and the Cisco Expressway Edge servers.

Initializing MRA on Expressway Servers: This topic will walk through the steps to enable MRA on the Expressway Core and Edge servers.

Collaboration Traversal Zones and Search Rules: This topic will walk through the steps to configure the appropriate traversal zones required to support the MRA solution.

Virtual private networks have a long-time tradition of connecting remote locations with an enterprise network. However, VPNs are very complex and require a lot of resources to operate. Additionally, the modern workplace is not always located in an office environment. Companies are now encouraging their employees to work from home, which complicates how these employees communicate with each other and the outside world. Cisco's out-of-the-box thinking has brought about a solution to many of the problems created by the modern workplace mindset. The Mobile and Remote Access (MRA) solution is a unique deployment that incorporates many facets of the more traditional traversal solution discussed in the preceding chapter. Communication devices can operate from any network at any location without the use of VPNs. Employees can leverage these communication devices from a home office without the need to set up and store another router in their home. Additionally, all the features and capabilities that employees would have from a communications device located in an office are still at their disposal from their remote location using this MRA solution. Topics discussed in this chapter include the following:

- Requirements for MRA
  - DNS A-Records and SRV Records
  - Firewall Ports and Considerations
  - Certificate Requirements and Recommendations

- HTTPS Reverse Proxy Settings
- Service Discovery

- Cisco Unified Communications Manager Settings for MRA
- TLS Verify Requirements
  - Cisco Expressway Certificates
  - Cisco Unified Communications Manager Certificates
  - Creating Certificates for MRA
- Initializing MRA on Expressway Servers
- Collaboration Traversal Zones and Search Rules

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.2 Describe the purpose of Edge devices in the Cisco Collaboration architecture such as Expressway and Cisco Unified Border Element
- 4.4 Describe Mobile and Remote Access (MRA)

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 21-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 21-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Requirements for MRA	1–5
Cisco Unified Communications Manager Settings for MRA	6
TLS Verify Requirements	7–9
Initializing MRA on Expressway Servers	10–11
Collaboration Traversal Zones and Search Rules	12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is an SRV record that's needed on the public DNS for an enterprise MRA deployment?
   a. *_collab-edge._tls.<domain>*
   b. *_cisco-uds._tcp.<domain>*
   c. *_cuplogin._tcp.<domain>*
   d. *_sip._tcp.<domain>*
2. Which of the following ports needs to be opened between the DMZ and the public Internet for an MRA deployment?
   a. TCP 7001
   b. UDP 36000–36001
   c. TCP 2222
   d. TCP 8443
3. Which of the following certificate pairs are required for an MRA deployment? (Choose two.)
   a. Public or enterprise CA certificate chain used to sign Expressway Core certificate
   b. Public or enterprise CA certificate chain used to sign Expressway Edge certificate
   c. CUCM Tomcat certificates or CA chain
   d. CUCM CallManager certificates or CA chain
   e. IMP Tomcat certificate or CA chain
   f. CUCM CAPF certificates
4. What two ports are used with reverse proxy to allow inbound authenticated HTTPS requests for TFTP file download and SOAP API requests on the CUCM?
   a. TCP 2222 and TCP 8443
   b. TCP 6970 and TCP 8443
   c. TCP 7400 and TCP 8443
   d. TCP 5222 and TCP 8443
5. When an endpoint located outside the corporate network is configured to register to the CUCM using MRA, what is the first communication sent by that endpoint?
   a. TLS handshake with the Expressway Edge to establish a trusted certificate verification
   b. Registration request sent to the CUCM through the Expressway Core and Edge servers
   c. DNS SRV Lookup for *_collab-edge._tcp.<domain>*
   d. DNS SRV Lookup for *_cisco-uds._tcp.<domain>*
6. What service on the Cisco Unified Communications Manager should be enabled for MRA?
   a. Cisco CallManager Service
   b. Cisco TFTP Service
   c. Cisco AXL Web Service
   d. Cisco CTI Service

7. Which of the following certificate options should be used on the Cisco Expressways for an MRA deployment?
   a. Self-signed certificates
   b. Single host/domain certificate
   c. Multiple subdomain wildcard certificates
   d. All of these answers are correct.
8. What CUCM certificates are significant for Mobile and Remote Access? (Choose two.)
   a. Public or enterprise CA certificate chain used to sign Expressway Core certificate
   b. Public or enterprise CA certificate chain used to sign Expressway Edge certificate
   c. CUCM Tomcat certificates or CA chain
   d. CUCM CallManager certificates or CA chain
   e. IMP Tomcat certificate or CA chain
   f. CUCM CAPF certificates
9. What is the recommended format for certificates on the Expressway servers for an MRA deployment?
   a. .cer or .crt format
   b. DER-encoded or Base64-encoded format
   c. DER-encoded format
   d. Base64-encoded format
   e. Any format can be used; there is no recommended format.
10. Which of the following statements is true when configuring an MRA solution?
   a. Enabling MRA on the Expressway-C involves turning it on and configuring MRA Access Control settings, but enabling MRA on the Expressway-E only involves turning it on.
   b. Enabling MRA on the Expressway-E involves turning it on and configuring MRA Access Control settings, but enabling MRA on the Expressway-C only involves turning it on.
   c. MRA needs to be enabled only on the Expressway-C, not the Expressway-E.
   d. MRA needs to be enabled only on the Expressway-E, not the Expressway-C.
   e. Enabling MRA is exactly the same on both the Expressway-C and the Expressway-E.
11. When nodes are being discovered on the Expressway-C for an MRA deployment, which of the following statements is true?
   a. The CUCM and CUCM IM and Presence nodes will not show Active until the traversal zones are configured and active.
   b. The CUCM IM and Presence nodes will not show Active until the traversal zones are configured and active.
   c. The CUCM node will not show Active until the traversal zones are configured and active.
   d. The CUCM and CUCM IM and Presence nodes will show Active immediately after they are discovered.

12. What zone type should be selected on the Cisco Expressway Edge server when setting up the traversal component of the MRA solution?
    a. Traversal client zone
    b. Neighbor zone
    c. Traversal server zone
    d. Unified Communications Traversal
    e. DNS zone
    f. ENUM zone

Foundation Topics

Requirements for MRA

Cisco Collaboration Mobile and Remote Access (MRA) is a core part of the Cisco Collaboration Edge architecture. MRA allows endpoints, such as Cisco Jabber, UC phones, and CE software-based Telepresence endpoints, to securely utilize registration, call control, provisioning, messaging, and presence services that are provided by Cisco Unified Communications Manager when the endpoint is not within the enterprise network. Cisco Expressway series components are used to provide secure access and firewall traversal to the endpoints that register with the Cisco Unified Communications Manager.

Although the MRA solution operates in a similar fashion to a standard firewall traversal solution for B2B communications, there are some significant differences between them. First, MRA is only supported for SIP; there is no H.323 support in the MRA solution. Second, certificates are required for MRA; there is no way to build the traversal zones with a basic TLS or TCP SIP connection. TLS Verify is required for MRA. Third, some specific settings must be configured to enable MRA on the Expressway servers. Fourth, the zones created between the Cisco Expressway Core and Cisco Expressway Edge servers are not the same traversal client zone and traversal server zone used in a standard firewall traversal solution. Finally, the DNS SRV records that need to be created are different from what is required for a traditional firewall traversal solution.

MRA consists of four main components. One of the components is Firewall Traversal Services. MRA supports internal firewalls between Cisco Expressway Core and Cisco Expressway Edge, and an external firewall between Cisco Expressway Edge and the Internet. The firewall traversal capabilities of MRA use the same Assent traversal protocol of standard firewall traversal, but traversal chaining is not supported with MRA. Another component is DNS records. Internal and external DNS records are essential to enable endpoints to detect whether they should register directly with Cisco Unified Communications Manager or proxy registration through the MRA deployment. Certificates are another component of MRA. This solution provides secure communication over Transport Layer Security (TLS). Trust between TLS entities is established based on certificates. Implementing the necessary certificates for a public key infrastructure (PKI) is an important part of Cisco Collaboration MRA implementation. The last component in an MRA solution is reverse HTTPS proxy. To support secure data services, such as visual voicemail, contact photo retrieval, Cisco Jabber custom tabs, and so on, a reverse HTTPS proxy runs on the Cisco Expressway Edge server. If these services are not needed in an enterprise deployment of MRA, this component does

not need to be set up. Once these components are set up, the MRA deployment can support two main features:

- **Off-Premises Access:** Cisco Collaboration MRA offers a consistent experience to clients, such as Cisco Jabber; UC phones; and Cisco DX, MX, SX, and Webex series endpoints, regardless of whether they are in the internal network or on an external network.
- **Business-to-Business Communications:** Cisco Collaboration Mobile and Remote Access offers secure communication to other businesses.

DNS A-Records and SRV Records

The DNS A-records and SRV records required for an MRA deployment are different from those required for a traditional firewall traversal solution. Additionally, different records need to be configured on an internal DNS and an external DNS. Before deploying an MRA solution, you will need to set up DNS first. Certificates cannot be created until DNS is configured because the PKI certificates depend on the DNS records of the different servers.

Cisco endpoints, especially Jabber, are programmed to use DNS so that they always search for the CUCM SRV record first. If the CUCM cannot be reached, they search for the Expressway-E. The Expressway-E can be located with an SRV lookup whether the endpoint is internal or external to the network, but the endpoint should not search for the Expressway unless it's external to the corporate network. The endpoint should be able to locate the CUCM using an SRV lookup only if the endpoint is located within the corporate network. This is the reason that the endpoint will always search for the CUCM first. If that path fails, there is an alternative path for the endpoint to register through the Expressway servers using MRA.

The external DNS server must be configured with a *_collab-edge._tls.<domain>* service record so that external endpoints can discover that they should use Cisco Expressway-E for Mobile and Remote Access. Service records for secure SIP are also required, not specifically for Mobile and Remote Access, but for deploying a secure SIP service on the Internet. The service records must point to each cluster member of the Cisco Expressway-E server. Table 21-2 provides examples of the service records needed on a public DNS for two Cisco Expressway Edge servers clustered together.

Key Topic

Table 21-2 Public DNS SRV Records for Expressway-E Cluster

Service	Protocol	Domains	Priority	Weight	Port	Target
_Collab-edge.	_tls.	Cisco.com	10	10	8443	Exp-e1.cisco.com
_Collab-edge.	_tls.	Cisco.com	10	10	8443	Exp-e2.cisco.com
_sips.	_tcp.	Cisco.com	10	10	5061	Exp-e1.cisco.com
_sips.	_tcp.	Cisco.com	10	10	5061	Exp-e2.cisco.com

The internal DNS server must be configured with a *_cisco-uds._tcp.<domain>* SRV record so that internal endpoints can discover that they should use Cisco Unified Communications Manager for direct registration. When using Cisco Unified Communications Manager IM and Presence Services, a *_cuplogin._tcp.<domain>* SRV record is also required on the internal DNS server. Just as the public DNS SRV records must refer to the Cisco Expressway-E servers, the internal DNS SRV records must refer to all call processing nodes of a Cisco Unified

Communications Manager cluster, as well as with all Cisco Unified Communications Manager IM and Presence server SRV records. The internal DNS records must be available to all internal endpoints and to the Cisco Expressway Core. The Cisco Unified Communications Manager and Cisco Unified Communications Manager IM and Presence server SRV records must not be resolvable from outside the internal network. Otherwise, Cisco endpoints and soft clients will not use the necessary Mobile and Remote Access registration via Cisco Expressway-E. Table 21-3 provides examples of the SRV records needed on a private DNS for two Cisco Unified Communications Managers and two Cisco Unified Communications Manager IM and Presence servers clustered together.

Key Topic

Table 21-3 Private DNS SRV Records for CUCM and CUCM IMP Clusters

Service	Protocol	Domains	Priority	Weight	Port	Target
_cisco-uds.	_tcp.	Cisco.com	10	10	8443	cucm1.cisco.com
_cisco-uds.	_tcp.	Cisco.com	10	10	8443	cucm2.cisco.com
_cuplogin.	_tcp.	Cisco.com	10	10	5061	imp1.cisco.com
_cuplogin.	_tcp.	Cisco.com	10	10	5061	imp2.cisco.com

Firewall Ports and Considerations

Cisco MRA uses a firewall traversal connection to allow inbound and outbound-initiated packet exchange, such as registration and call setup messages. MRA uses the Cisco Expressway Edge as the traversal server that is installed in a demilitarized zone (DMZ), and the Expressway Core is the traversal client that is installed on the internal network. Firewall traversal offers secure communication across firewalls as follows:

Key Topic

1. The Cisco Expressway-C initiates an outbound traversal connection through the internal firewall to specific ports on the Cisco Expressway-E with secure authentication credentials to establish a connection between the two servers.
2. Once the connection has been established, the Cisco Expressway-C sends keepalive packets periodically to Cisco Expressway-E to maintain the connection.
3. When Cisco Expressway-E receives an incoming message, whether it's a registration or call setup message, from an outside endpoint, it sends the request to the Cisco Expressway-C through the existing traversal connection.
4. The Cisco Expressway-C then sends the message, such as a call setup request, to the Cisco Unified Communications Manager.
5. The Cisco Unified Communications Manager processes the call, and media streams are set up over the existing traversal connection.

For communication to flow through the firewall, appropriate ports must be opened to allow the flow of packets. The following ports must be opened on the internal firewall between the Expressway Core and the Expressway Edge:

Key Topic

- **SIP:** TCP 7001
- **Traversal Media:** UDP 36000 to 36001 (for small to medium VM deployments)
- **Extensible Messaging and Presence Protocol (XMPP):** TCP 7400

- **HTTPS (Tunneled over Secure Shell [SSH] between Expressway-C and Expressway-E):** TCP 2222

The following ports must be opened on the external firewall between the public Internet and the Cisco Expressway Edge in the DMZ:

- **SIP:** TCP 5061
- **HTTPS:** TCP 8443
- **XMPP:** TCP 5222
- **TURN Server Control and Media:** UDP 36012 to 59999 (if TURN relays are being used only)

The firewall administrator should open all of the ports from the preceding list before traversal zones are set up between the Expressway-C and the Expressway-E. Certificates must also be set up before traversal zones are created. Whether firewall ports are opened or certificates are established first doesn't matter as long as both tasks are completed before configuring the traversal zones.

Certificate Requirements and Recommendations

Six different certificate pairs can be configured in an MRA deployment. However, only two pairs are required to set up the solution. The other four exist in an ideal environment for absolute security pertaining to registration and calling. The first certificate required is a public or enterprise CA certificate chain used to sign the Expressway-C. This is required to establish the traversal client zone connection. The second certificate required is a public or enterprise CA certificate chain used to sign the Expressway-E. This is also required to establish the traversal server zone connection. Both are absolutely required for TLS Verify to operate successfully. The traversal zones used for an MRA deployment will not work without these two certificate pairs. The root CA certificate for the Expressway-C certificate should be added to both the Expressway-C and the Expressway-E. The root CA certificate for the Expressway-E certificate should be added to both the Expressway-E and the Expressway-C. If both servers were signed by the same CA, then they will use the same root CA certificate; therefore, it needs to be added to each server only once.

The next optional certificate is the Cisco Unified Communications Manager Tomcat certificate or CA chain. The Tomcat certificate is for Tomcat trust. This certificate is used for MRA only when the Expressway-C is configured to use TLS Verify mode on Cisco Unified Communications Manager discovery. The Tomcat CA should be added to the Expressway-C, and the root CA certificate for the Expressway-C should be added to the Cisco Unified Communications Manager. If TLS Verify is not used on the Expressway-C for Cisco Unified Communications Manager discovery, this certificate is not needed.

Another optional certificate is the Cisco Unified Communications Manager certificate or CA chain used when the Cisco Unified Communications Manager is in mixed mode for end-to-end TLS. If this certificate is used, the Cisco Unified Communications Manager CA should be added to the Expressway-C, and the certificate CA for the Expressway-C should be added to the Cisco Unified Communications Manager.

The Cisco Unified Communications Manager IM and Presence Tomcat certificate or CA chain is similar to the Cisco Unified Communications Manager Tomcat certificate or CA chain. This certificate is used only when the Expressway-C is configured to use TLS Verify mode on Cisco Unified Communications Manager IM and Presence discovery. The Tomcat CA should be added to the Expressway-C, and the certificate CA for the Expressway-C should be added to the Cisco Unified Communications Manager IM and Presence server. If TLS Verify is not used on the Expressway-C for Cisco Unified Communications Manager IM and Presence discovery, this certificate is not needed.

The last optional certificate is the Cisco Unified Communications Manager Certificate Authority Proxy Function (CAPF) certificate. This certificate is used when remote endpoints authenticate using a Locally Significant Certificate (LSC). By default, LSC is signed by the CAPF, so the CAPF is the CA for phones in this scenario. However, when the CAPF is signed by an external CA, then the CAPF in this scenario acts as a subordinate CA or intermediate CA. The difference between a self-signed CAPF and CA-signed CAPF is that the CAPF is the root CA to LSC when doing a self-signed CAPF, but the CAPF is the subordinate or intermediate CA to LSC when doing a CA-signed CAPF. Table 21-4 identifies each of the six certificate pairs used in an MRA deployment.

Table 21-4 Certificate Pairs Used in an MRA Deployment

Certificate Type	Core	Edge	Required
Public or enterprise CA certificate chain used to sign Expressway Core certificate	Yes	Yes	Yes
Public or enterprise CA certificate chain used to sign Expressway Edge certificate	Yes	Yes	Yes
CUCM Tomcat certificates or CA chain	Yes	No	No
CUCM CallManager certificates or CA chain	Yes	No	No
IMP Tomcat certificate or CA chain	Yes	No	No
CUCM CAPF certificates	No	Yes	No

HTTPS Reverse Proxy Settings

The Cisco MRA reverse proxy settings provide a mechanism to support visual voicemail access, contact photo retrieval, Cisco Jabber custom tabs, and other data applications. HTTPS reverse proxy is a function that is provided by the Cisco Expressway-E using port TCP 8443 for HTTPS traffic. Initial MRA configuration allows inbound authenticated HTTPS requests to the following destinations:

- TCP 6970 (TFTP file download) and TCP 8443 (SOAP API) to all discovered Cisco Unified Communications Manager nodes
- TCP 7400 (XCP router) and TCP 8443 (SOAP API) to all Cisco Unified Communications Manager IM and Presence nodes

Additional hosts can be added to the allow list on the Cisco Expressway-C.

Service Discovery

Before we show how to configure an MRA solution, we should more closely examine the MRA service discovery operation. This discussion will help administrators deploying this solution fully understand the dependencies between the components involved with an MRA solution. Figure 21-1 illustrates the way Cisco MRA service discovery operates on the public network. This example is for a Cisco Jabber client in phone-only mode. Additional steps involving the Cisco Unified Communications Manager IM and Presence Services would need to be included if additional Cisco Jabber services were being utilized.

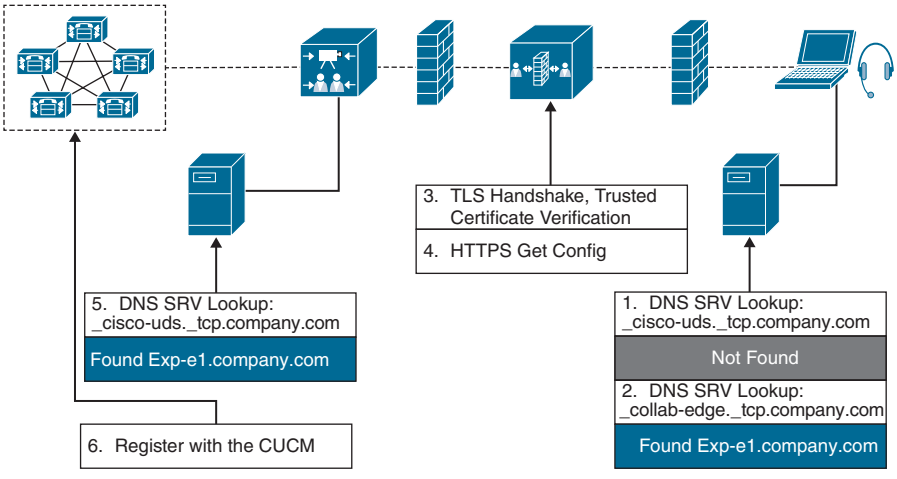

Figure 21-1 *Cisco MRA Service Discovery Operation*

Figure 21-1 assumes that the initiating endpoint, or Jabber client in this case, does not connect to the corporate network over a VPN. This is the reason that the initial DNS SRV lookup for the *_cisco-uds._tcp.domain* record fails. The service discovery occurs as follows. First, a Cisco Jabber client located outside the corporate network, and without a VPN connection, sends a DNS SRV record lookup for *_cisco-uds._tcp.company.com* to a public DNS server. The public enterprise DNS that manages company.com should not have such an SRV record, and therefore, the lookup fails. Next, the Cisco Jabber client sends another DNS SRV record lookup for *_collab-edge._tls.company.com*. This time the lookup is successful, and the address of the Cisco Expressway Edge is provided to the Jabber client in the DNS response.

Now the Cisco Jabber client can start the Mobile and Remote Access negotiation with the Cisco Expressway Edge server. A certificate is presented to Cisco Jabber and may need to be manually trusted by the user if it is not signed by a certificate authority server that the client PC already trusts. A TLS handshake is exchanged to establish a secure connection. The Cisco Expressway Edge will then act as a proxy for the Cisco Jabber client by passing messages that it receives from Cisco Jabber to Cisco Expressway Core through the firewall traversal connection and return messages from the Expressway Core to the Jabber client.

When a trusted connection between Cisco Jabber and Cisco Expressway Edge is established, Cisco Jabber tries to register to the services that are enabled on Cisco Expressway Core, which in this case is Cisco Unified Communications Manager. The Cisco Expressway Core will send a DNS SRV record lookup for *_cisco-uds._tcp.company.com* to the internal DNS. The internal DNS will respond with the address of the Cisco Unified Communications Manager. The Cisco

Expressway Core will then forward the registration request from the Cisco Jabber client to the Cisco Unified Communications Manager. The Cisco Expressway Core will act as the proxy for messages between the Cisco Unified Communications Manager and the Expressway Edge.

Cisco Unified Communications Manager Settings for MRA

After an administrator has ensured that all the prerequisite components have been configured, which were covered in the first section of this chapter, the process for configuring MRA in a corporate network begins on the Cisco Unified Communications Manager. The following seven steps must be completed to deploy Mobile and Remote Access endpoints.

Step 1. Make sure that the Cisco AXL Web Service is activated on the publisher node.

- **a.** From Cisco Unified Serviceability, navigate to **Tools > Service Activation**.
- **b.** From the **Server** drop-down menu, select the publisher node and click **Go**.
- **c.** Under the Database and Admin Services section, confirm that the Cisco AXL Web Service is Activated.
- **d.** If the service is not activated, check the corresponding check box and click **Save** to activate the service.

Step 2. Optionally, configure region-specific settings for MRA endpoints. The default settings may be sufficient in many cases, but if you expect MRA endpoints to use video, you may want to increase the Maximum Session Bit Rate for Video Calls within your region configuration. The default setting of 384 kbps may be too low for some video endpoints, such as the DX series.

- **a.** From Cisco Unified CM Administration, navigate to **System > Region Information > Region**.
- **b.** Perform any one of the following:
  - Click **Find** and select a region to edit the bit rates.
  - Click **Add New** to create a new region.
- **c.** In the Modify Relationship to Other Regions section, configure a new setting for the Maximum Session Bit Rate for Video Calls, such as 6000 kbps.
- **d.** Configure other fields in the Region Configuration window as necessary. For more information on the fields and their configuration options, see the system's online help.
- **e.** Click **Save**.

Step 3. After you have created a new region, assign your region to the device pool that your MRA endpoints use.

- **a.** From Cisco Unified CM Administration, navigate to **System > Device Pool**.

b. Do either of the following:

- Click **Find** and select the existing device pool to edit.
- Click **Add New** to create a new device pool.

c. Enter a device pool **Name**.

d. Select a redundant Cisco Unified Communications Manager Group.

e. Assign a Date/Time Group. This group includes the Phone NTP references that may have been set up for MRA endpoints.

f. Assign a region from the Region drop-down menu to the device pool that MRA endpoints will use.

g. Complete the remaining fields in the Device Pool Configuration window as necessary. For more information on the fields and their configuration options, see the system's online help.

h. Click **Save**.

Step 4. Use this procedure to set up a Phone Security Profile to be used by MRA endpoints. You must apply this profile to the phone configuration for each of your MRA endpoints.

a. From Cisco Unified CM Administration, navigate to **System > Security > Phone Security Profile**.

b. Click **Add New**.

c. From the Phone Security Profile Type drop-down list, select your device type, such as the Cisco Unified Client Service Framework for a Jabber application.

d. Click **Next**.

e. Enter a Name for the profile. For MRA, the name must be in FQDN format and must include the enterprise domain.

f. From the Device Security Mode drop-down list, select **Encrypted**. This field must be set to Encrypted; otherwise, Expressway will reject the communication.

g. Set the Transport Type to **TLS**.

h. Leave the TFTP Encrypted Config check box unchecked for the following phones because MRA will not work for these phones with this option enabled: DX Series, IP Phone 7800, or IP Phone 8811, 8841, 8845, 8861 and 8865.

i. Complete the remaining fields in the Phone Security Profile Configuration window. For more information on the fields and their configuration options, see the system's online help.

j. Click **Save**.

Step 5. This step is for Cisco Jabber only. Set up an MRA Access Policy for Cisco Jabber users. Cisco Jabber users must be enabled with MRA access within their user profiles in order to use the MRA feature. The Jabber desktop client includes Cisco Jabber for Windows users and Cisco Jabber for Mac users. The

Jabber mobile client includes Cisco Jabber for iPad and iPhone users and Cisco Jabber for Android users.

a. In Cisco Unified CM Administration, navigate to **User Management > User Settings > User Profile**.

b. Click **Add New**.

c. Enter a Name and Description for the user profile.

d. Assign a Universal Device Template to apply to users' Desk Phones, Mobile and Desktop Devices, and Remote Destination/Device Profiles.

e. Assign a Universal Line Template to apply to the phone lines for users in this user profile.

f. If you want the users in this user profile to be able to use the self-provisioning feature to provision their own phones, do the following:

- Check the **Allow End User to Provision Their Own Phones** check box.
- In the Limit Provisioning Once End User Has This Many Phones field, enter a maximum number of phones the user is allowed to provision. The maximum is 20.

g. If you want Cisco Jabber users associated with this user profile to be able to use the Mobile and Remote Access feature, check the **Enable Mobile and Remote Access** check box. By default, this check box is selected. When you uncheck this check box, the Jabber Policies section is disabled, and the No Service Client Policy option is selected by default. This setting is mandatory only for Cisco Jabber users. Non-Jabber users do not need this setting to be able to use MRA.

h. Assign the Jabber policies for this user profile. From the Jabber Desktop Client Policy and Jabber Mobile Client Policy drop-down list, choose one of the following options:

- **No Service:** This policy disables access to all Cisco Jabber services.
- **IM & Presence Only:** This policy enables only instant messaging and presence capabilities.
- **IM & Presence, Voice and Video Calls:** This policy enables instant messaging, presence, voicemail, and conferencing capabilities for all users with audio or video devices. This is the default option.

i. Click **Save**.

Step 6. This step is also for Cisco Jabber users only. The user policy that was set up previously must be applied to the appropriate end user. Remember from Chapter 16, "LDAP Integration with Cisco Unified Communications Manager," and Chapter 17, "Registering SIP Endpoints to the Cisco Unified Communications Manager," that end users can be set up manually, from an import using LDAP or using BAT. Because each method was covered in these chapters, we will not cover the steps to apply the user policy to the end user at this point in the book. Refer to these other chapters if a review is necessary.

Step 7. Configure and provision endpoints that will use the MRA feature. This step is achieved by ensuring the corresponding settings above are applied to the phone through TFTP when registration is attempted. Refer to Chapter 7, "Cisco Unified Communications Phones," Chapter 8, "Cisco Telepresence Endpoints," and Chapter 9, "Endpoint Registration," for a review in configuring endpoints registering to the Cisco Unified Communications Manager.

High volumes of Mobile and Remote Access calls may trigger denial of service thresholds on the Cisco Unified Communications Manager. The reason is that all the calls arriving at the Cisco Unified Communications Manager are from the same Expressway-C cluster. If necessary, Cisco recommends that you increase the level of the SIP Station TCP Port Throttle Threshold to 750 kbps. To make this change from Cisco Unified CM Administration, navigate to **System > Service Parameters**, and select the Cisco CallManager service.

Key Topic

Another requirement on the Cisco Unified Communications Manager must be configured before MRA is set up on the Expressway Core. An application user must be configured on the Cisco Unified Communications Manager that has been assigned the *AXL API Access* role. The Cisco Expressway Core will need this level of access into the Cisco Unified Communications Manager; otherwise, communication will fail. The default administrator account can be used, but Cisco recommends a new application user account be created for each application service that's set up with the Cisco Unified Communications Manager. The same is also true for MRA deployments that include the Cisco Unified Communications Manager IM and Presence services. The corresponding IM and Presence user must have the *Standard AXL API Access* role assigned.

TLS Verify Requirements

Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), both frequently referred to as "SSL," are cryptographic protocols that provide communications security over a computer network. This TCP protocol aims primarily to provide privacy and data integrity between two communicating hosts or applications.

Client/server applications such as web browsers, email, and VoIP commonly use the TLS protocol to prevent eavesdropping and tampering of information. The protocols these applications use must choose to use or not to use TLS. The easiest way to segregate the information is to use different port numbers for unencrypted traffic, such as port 80 for HTTP, and TLS-encrypted traffic, such as port 443 for HTTPS. The connection is secure because symmetric cryptography is used to encrypt the transmitted data. The keys for this symmetric encryption are generated uniquely for each connection and are based on a shared secret negotiation at the start of the session. The server and client negotiate the details of which encryption algorithm and cryptographic keys to use before the first byte of data is transmitted. Identification is usually in the form of digital "certificates" that contain the server name, the trusted certificate authority (CA), and the server's public encryption key. The identity of the communicating parties can be authenticated using this public-key cryptography (asymmetric cryptography) to ensure only the intended recipient can decrypt the traffic. The negotiation of a shared secret is both secure and reliable against eavesdroppers and attacks, including man-in-the-middle attacks. The connection ensures integrity because each message transmitted includes a message integrity check using a message authentication code to prevent undetected loss or alteration of the data during transmission.

Once the client and server have agreed to use TLS, they negotiate a stateful connection by using a handshake procedure. The handshake begins when a client connects to a TLS-enabled server requesting a secure connection and presents a list of supported ciphers and hash functions. From this list, the server picks a cipher and hash function that it also supports, and it informs the client of the decision. The server then identifies itself with its digital certificate, which can contain the server name, the trusted certificate authority, and the server's public encryption key. The client then validates the certificate before proceeding. Public-key encryption is used to share the pre-master secret via the use of RSA or Diffie-Hellman key exchange. This process generates a random and unique session key for encryption and decryption that has the additional property of forward secrecy, which protects past sessions against future compromises of secret keys or passwords.

Remember that the server is validated because the client initiates the secure connection. The client side confirms that the server is who it claims to be and whether it can be trusted with the use of certificates. The client receives the digital certificate from the server side of the TLS negotiation, but the identity must be verified before proceeding. The server certificate may contain the name of the certificate holder. This name is checked against the Common Name (CN) or the Subject Alternative Name (SAN). Also, additional information like a serial number, expiration dates, revocation status, a copy of the certificate holder's public key (which is used for encrypting messages and digital signatures), and the digital signature of the certificate-issuing authority. This information identifies to the client that the certificate is signed by a certificate authority. If you trust this certificate authority, you can verify using the CA's public key that it really did sign the server's certificate. To sign a certificate yourself, you need the private key, which is only known to the CA of your choice. This way an attacker cannot sign a certificate himself and falsely claim to be the server. When the certificate has been modified, the sign will be incorrect, and the client will reject it. Figure 21-2 illustrates the steps involved with a security handshake between a client and a server.

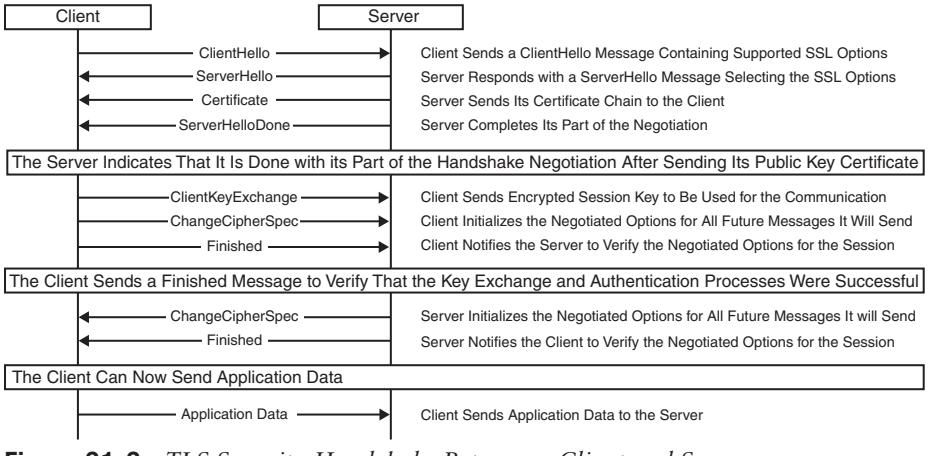

Figure 21-2 *TLS Security Handshake Between a Client and Server*

Mutual TLS authentication is also an option that can be chosen. In this type of authentication, both parties authenticate each other through verifying the provided digital certificate so that both parties are assured of the others' identity. Mutual TLS is very similar to the normal process of the client handling the verification of the server's certification but including the additional step of the client providing a certificate. This process allows the server side to authenticate the client, allowing both parties to trust each other.

Server-to-server connections rely on mutual TLS for mutual authentication. In the Cisco Collaboration infrastructure, some examples would be a secure connection between endpoints and the Cisco Unified Communications Manager, Cisco Unified Communications Manager SIP trunks to other clusters, and even Cisco Unified Communications Manager SIP trunks with a Cisco Expressway Core.

To secure voice and video traffic, you must understand multiple technologies. Remember, the most common VoIP communication used today is SIP. For secure transmissions of SIP messages, the protocol may be encrypted with TLS. Media identification and negotiation are achieved with the Session Description Protocol (SDP). SDP can also be used for the master key exchange. For the transmission of media streams, SIP employs the Real-time Transport Protocol (RTP) or Secure Real-time Transport Protocol (SRTP). Unencrypted SIP generally uses port 5060, whereas TLS-encrypted SIP utilizes port 5061.

Trusted certificates are very important to create secure connections. This part of the process is where the certificate authority comes into the solution. These certificate authorities are widely used both on the public Internet and private networks to issue digital certificates containing identity credentials binding them to SSL or TLS cryptography keys. However, because these CAs are trust anchors, they must conduct several checks into the identity of the applicant. The checks correlate to the class and type of certificate being applied for. Table 21-5 identifies each of these classes of certificates available and which of these certificates are supported on Cisco Collaboration servers.

Key Topic

Table 21-5 Classes of Certificates on Cisco Collaboration Servers

Options	Types			Support Info
	DV	OV	EV	
Single Host/Domain	Yes	Yes	Yes	Supported on all Cisco Collaboration servers
UCC/Multiple SAN/Cert	Yes	Yes	Yes	Supported on all Cisco Collaboration servers
Multiple Subdomain Wildcard Cert	Yes	Yes	No	Not supported at all on Cisco Expressways

For domain validation (DV) certificates, the CA checks only the right of the applicant to use a specific domain name. No company identity information is vetted, and no information is displayed other than encryption information within the Secure Site Seal. For organization validation (OV) certificates, the CA checks the right of the applicant to use a specific domain name and conducts some vetting of the organization. Additional vetted company information is displayed to customers when clicking the Secure Site Seal, giving enhanced visibility into who is behind the site and associated enhanced trust. For extended validation (EV) certificates, the certificate authority checks the right of the applicant to use a specific domain name, and in addition, it conducts a thorough vetting of the organization. The issuance process of EV certificates is strictly defined in the EV guidelines, as formally ratified by the CA/Browser Forum in 2007, that specify all the steps that are required for a CA before issuing a certificate. They include

- Verifying the legal, physical, and operational existence of the entity
- Verifying that the identity of the entity matches official records

- Verifying that the entity has exclusive rights to use the domain that is specified in the EV SSL certificate
- Verifying that the entity has properly authorized the issuance of the EV SSL certificate

EV certificates are available for all types of businesses, including government entities and both incorporated and unincorporated businesses. A second set of guidelines, the EV Audit Guidelines, specify the criteria under which a CA needs to be successfully audited before issuing EV certificates. The audits are repeated yearly to ensure the integrity of the issuance process.

Because people are constantly searching the Internet, the browsers constantly are checking the websites that are visited against a CA to authenticate web pages. As an example, web browsers like Google Chrome, Firefox, and Internet Explorer maintain lists of certificate authorities they consider trustworthy. When you access what should be a secure website, the site presents its security certificate to your browser. If the certificate is up-to-date and from a trusted certificate authority, you will see the trusted secure connection. If the certificate lacks any of the requirements, you will see that your web browser will not establish a connection until you accept the risks and proceed.

With the Cisco Collaboration products, only certain certificate options are supported. The Expressway-E will require a public certificate because it is the most public-facing part of the Cisco Collaboration solution and needs to be trusted by outside sources like clients and other businesses. A single host/domain certificate option will suffice with any type of validation desired. If multiple hosts, domains, or subdomains need to be covered, multiple Subject Alternative Name (SAN) certificates are required. Note that the wildcard certification is not supported on the Cisco Expressway series.

In the Cisco Collaboration architecture, it is easy to secure all enterprise network information simply by creating a private network behind the security of a firewall. However, companies may need to work with other businesses or clients for their day-to-day operations. To place voice and video calls outside the network and connect to other networks securely, Cisco offers the Expressway Series. This robust and secure solution comes equipped with Cisco's proprietary Assent protocol that allows secure firewall-traversal technology for any-to-any collaboration.

Cisco Expressway Certificates

Best practice for setting up the Expressway Series begins with the Expressway Edge, which is usually placed in between two firewalls in a separate network from the private enterprise network and the public outside network. This subnetwork is known as a demilitarized zone (DMZ). First, you can add authentication credentials and build a traversal server zone on the Expressway Edge. All zones also require search rules that will determine when and how they are searched. Next, you can configure the Expressway Core to use a traversal client zone. This zone initiates a secure traversal connection through the firewall on a specific keepalive port and authenticates against the authentication database configured on the Expressway Edge. Once a connection is established, the Expressway Core preserves the connection by continuously sending UDP keepalive packets over that same port to the Expressway Edge. This allows endpoints to both place calls out of the network and receive incoming calls as well. When a call comes into the network, the call setup is forwarded from the Expressway Edge to the Expressway Core and ultimately to the Cisco Unified Communications Manager to search for a user or endpoint. Once call setup has completed, both the call signaling and media will securely traverse through the firewall using the traversal zones previously created.

The process described here can use a standard TLS verification, which uses a single self-signed certificate on the Expressway Edge, or it can use the more secure TLS Verify mode where both Expressways have to validate each other's certificates. The TLS Verify mode can be set to enabled or disabled on the Zone settings page, which will decide which mode is used. When the zone is configured with the TLS Verify mode set to OFF, the Expressway Edge declines to verify the host name and signature of a certificate from the Expressway Core. The Expressway Core still verifies the certificate of the Expressway Edge's self-signed certificate, but this certificate does not use domain verification; therefore, this configuration also allows for the use of IP addresses for the peers. With TLS Verify mode set to ON, Mutual TLS (MTLS) is activated, and both client and server will match the CN or SAN against the peer address.

When using the TLS Verify mode in the ON configuration on an Expressway-E, the CA and SAN must match the TLS Verify Subject Name field in the zone configuration. As a result, this configuration is commonly used in closed or federated systems. For open B2B searches to other nonfederated enterprises, TLS Verify mode on the Expressway-E needs to be in the OFF mode. When you are setting up traversal zones for MRA, TLS Verify mode is required to be turned ON. MRA will not work if TLS Verify mode is set to OFF.

Cisco Unified Communications Manager Certificates

Two Cisco Unified Communications Manager certificates are significant for Mobile and Remote Access. They are the CallManager certificate and the Tomcat certificate. These certificates are automatically installed on the Cisco Unified Communications Manager, and by default, they are self-signed and have the same Common Name (CN). Cisco recommends using CA-signed certificates. However, if self-signed certificates are used, the two certificates must have different Common Names. The Expressway does not allow two self-signed certificates with the same CN. If the CallManager and Tomcat self-signed certificates have the same CN in the Expressway's trusted CA list, the Expressway can trust only one of them. This means that either secure HTTP or secure SIP between Expressway-C and Cisco Unified Communications Manager will fail.

Two IM and Presence Service certificates are significant if you use XMPP. They are the cup-xmpp certificate and tomcat certificate. Cisco recommends using CA-signed certificates. However, if self-signed certificates are used, these two certificates must also have different Common Names. The Expressway does not allow two self-signed certificates with the same CN. If the cup-xmpp and tomcat (self-signed) certificates have the same CN, Expressway trusts only one of them, and some TLS attempts between Cisco Expressway-E and IM and Presence Service servers will fail.

Although Expressway certificates were discussed in the previous section, some important settings on the Cisco Unified Communications Manager can affect how the Expressway Core certificates are set up. The Expressway Core server certificate needs to include the following elements in its list of subject alternate names:

- **Unified CM Phone Security Profile Names:** The names of the phone security profiles in the Cisco Unified Communications Manager that are configured for encrypted TLS and are used for devices requiring remote access. Use the FQDN format and separate multiple entries with commas. Having the secure phone profiles as alternative names means that the Cisco Unified Communications Manager can communicate via TLS with the Expressway-C when it is forwarding messages from devices that use those profiles.

- **IM and Presence Chat Node Aliases (Federated Group Chat):** The chat node aliases that are configured on the IM and Presence Service servers. These are required only for Unified Communications XMPP federation deployments that intend to support group chat over TLS with federated contacts. The Expressway-C automatically includes the chat node aliases in the CSR, providing it has discovered a set of IM and Presence Service servers. Cisco recommends using DNS format for the chat node aliases when generating the CSR. The same chat node aliases must be used in the Expressway Edge server certificate's alternative names (SANs).

The Expressway Edge server certificate needs to include the following elements in its list of Subject Alternative Names (SANs):

- **Cisco Unified Communications Manager Registrations Domains:** All of the domains that are configured on the Expressway Core for Cisco Unified Communications Manager registrations. Required for secure communications between endpoint devices and the Expressway Edge.
- **XMPP Federation Domains:** The domains used for point-to-point XMPP federation. These are configured on the IM and Presence Service servers and should also be configured on the Expressway-C as domains for XMPP federation. Select the DNS format and manually specify the required FQDNs. Separate the FQDNs with commas if you need multiple domains. Do not use the XMPPAddress format because your CA may not support it, and it may be discontinued in future versions of the Expressway software.
- **IM and Presence Chat Node Aliases (Federated Group Chat):** The same set of chat node aliases as entered on the Expressway-C's certificate. They are required only for voice and presence deployments that will support group chat over TLS with federated contacts. Note that you can copy the list of chat node aliases from the equivalent Generate CSR page on the Expressway-C.

The Cisco Unified Communications Manager registration domains used in the Expressway configuration and Expressway-E certificate are used by Mobile and Remote Access clients to look up the *_collab-edge* DNS SRV record during service discovery. They enable MRA registrations on the Cisco Unified Communications Manager and are primarily for service discovery. These service discovery domains may or may not match the SIP registration domains. It depends on the deployment, and they don't have to match. One example is a deployment that uses a *.local* or similar private domain with Cisco Unified Communications Manager on the internal network, and public domain names for the Expressway-E FQDN and service discovery. In this case, you need to include the public domain names in the Expressway-E certificate as SANs. There is no need to include the private domain names used on the Cisco Unified Communications Manager. Only the edge domain needs to be listed as a SAN. Select the DNS format and manually specify the required FQDNs. Separate the FQDNs with commas if you need multiple domains. You may select *CollabEdgeDNS* format instead, which simply adds the prefix *collab-edge* to the domain that you enter. This format is recommended if you do not want to include your top-level domain as a SAN.

Creating Certificates for MRA

The Cisco Unified Communications Manager does not require an MTLS connection to the Expressway Core for the MRA deployment; therefore, this section will cover only the detailed steps on how to sign and load certificates on the Expressway Core and Edge servers. Once all

MRA settings have been configured on the Expressway-C, traversal zones can be configured between the Expressway-E and the Expressway-C. TLS Verify is required for these zones, so certificates must be used. If the certificates come from the same CA, the root CA will be the same on both Expressway servers. However, if a different CA is used for the Expressway Core than what is used on the Expressway Edge, the root CA certificate must be uploaded to both Expressways to identify where the certificates were signed. Information required when signing certificates is the same regardless of the CA being used. The instructions provided in this book on how to sign certificates are through a Microsoft certificate server.

The Expressway-C and Expressway-E have a tool built into them that can generate a certificate signing request (CSR). The CSR contains all the information the CA will need to sign the certificate. When the CA signs the certificate, it's important that it is done with a template that contains the server and client authentication extensions. The certificate generation process seems to confuse a lot of customers and engineers. The basic requirement to get this set up is fairly straightforward. The first consideration is what to use for a CA. There are two commonly used approaches. One method is to use OpenSSL, and the other is to use Active Directory Certificate Services (ADCS). Setting up ADCS in a Microsoft environment is very complex and would warrant discussion that goes beyond the scope of this book. The OpenSSL method is well defined in the *VCS Certificate Creation and Use Deployment Guide*; therefore, this section will focus only on how to use an ADCS after it has already been configured on a Windows server.

Step 1. On the Expressway, generate a certificate signing request (CSR). For the most part, the following steps pertain to both the Expressway-C and Expressway-E servers.

a. On the Expressway, navigate to **Maintenance > Security > Server Certificate**.

b. Click **Generate CSR** to access the Generate CSR page.

c. Fill out the appropriate details, and then click the **Generate CSR** button:

- **Key Length:** (Recommended to use 2048 or higher)
- **Country:** (Optional) Abbreviations OK
- **State or Province:** (Optional) Abbreviations OK
- **Locality (Town Name):** (Optional)
- **Organization (Company Name):** (Optional)
- **Organization Unit:** (Optional)

Key Topic

Make a note of the Common Name that is autopopulated on the Generate CSR page. This is automatically created using the DNS settings, so the DNS settings on the Expressway must be accurate. The Common Name on the Expressway-C will be used as the Subject Name when the traversal zone is configured on the Expressway-E. The Common Name on the Expressway-E will be used as the Peer address when the traversal zone is configured on the Expressway-C. There are a few points of interest for the Expressway-C certificate. If the IM and Presence Service has already been added to the Expressway-C, a prepopulated Chat Node Alias will appear. This is required for XMPP federation deployments

that intend to use both TLS and group chat, such as *conference-2-StandAloneCluster5ad9a.%yourdomain%*. Also, if the solution is being deployed using TLS between the Expressway-C and Cisco Unified Communications Manager, ensure that the Subject Alternative Name on the certificate contains the names in FQDN format of all the phone security profiles in the Cisco Unified Communications Manager that are configured for encrypted TLS, such as *CSFJabber.tftp.com*.

d. Under the Certificate Signing Request section, click **Show (PEM file)**. Copy the entire contents of the PEM file to a notepad. Be sure to include the -----Begin Certificate Request----- and -----End Certificate Request----- lines. The contents of this PEM file will be used to sign the Expressway-C certificate using Microsoft ADCS. Figure 21-3 illustrates the settings that need to be configured when generating a CSR.

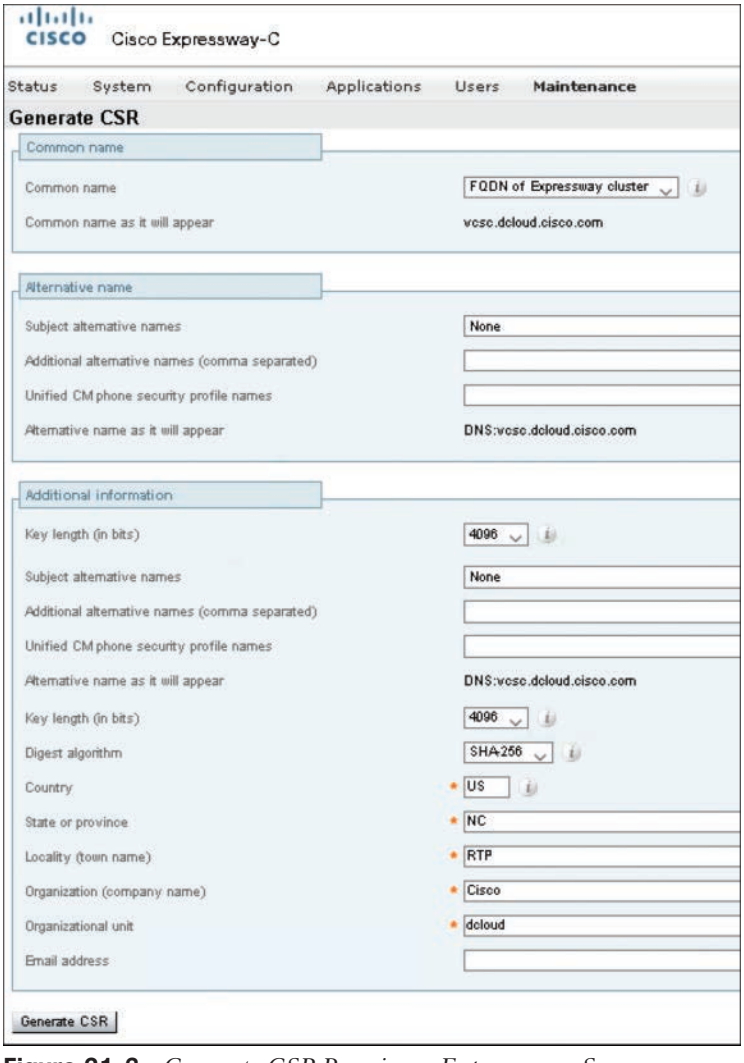

Figure 21-3 *Generate CSR Page in an Expressway Server*

Step 2. Sign the CSR with the Microsoft ADCS.

a. Browse to the ADCS web interface at **https://<*IP_address*>/certsrv**.

b. Log in and select **Request a Certificate**.

c. Click the **Submit a Certificate Request by Using a Base-64-Encoded CMC or PKCS #10 File**, or **Submit a Renewal Request by Using a Base-64-Encoded PKCS #7 File** option.

d. Paste the copied contents from the Expressway-C Server PEM file into the Base-64-Encoded Certificate Request box.

e. Some ADCS servers have a Certificate Template field. If this option is available, choose the most appropriate one for your certificate purpose, and then click **Submit**.

Key Topic

f. The next page that appears will allow you to download the signed certificate in either a DER-encoded or Base64-encoded format. DER stands for Distinguished Encoding Rules, which is a binary format. Base64 is an encoding method that converts binary to plain ASCII text. Some scenarios prevent copying and transferring data in binary, so plain text is needed. Therefore, it is recommended to choose the **Base 64 Encoded** option before downloading the certificate.

Key Topic

g. After the signed certificate has been downloaded, a copy of the root CA certificate will need to be downloaded as well. The root CA certificate establishes a trusted chain that begins at the root CA, or in this case the ADCS, through the root CA certificate, and ends at the certificate that was signed. The use of the root CA certificate provides an added level of security. From the ADCS page, click **Home** in the top-right corner of the screen.

h. Click the **Download a CA Certificate, Certificate Chain or CRL** link.

i. Select the **Base 64** radio button and click **Download CA Certificate**.

j. When prompted, save the certificate into the same location as the signed server certificate.

Step 3. After both certificates have been obtained, return to the Expressway and apply the certificates.

a. Navigate to **Maintenance > Security > Server Certificate**. This is the same page where the CSR request was generated.

b. Scroll to the bottom of the page, select **Browse**, and choose the server certificate that was just signed by ADCS.

c. Click **Upload Server Certificate Data**. Depending on what version of Expressway is being used, the web browser may prompt the administrator to reauthenticate again. A restart may also be required to complete the certificate installation. If so, do not restart the Expressway until the root CA has been installed. Figure 21-4 illustrates the Server Certificate page on an Expressway-C.

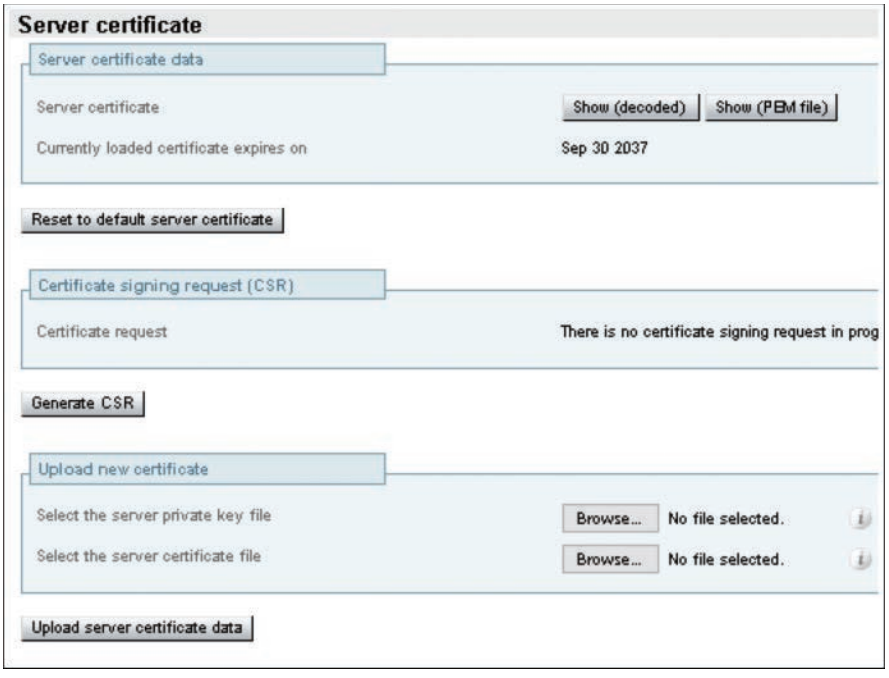

Figure 21-4 *Expressway-C Server Certificate Page*

d. Navigate to **Maintenance > Security > Trusted CA Certificate**.

e. Select **Browse** and choose the root CA certificate that was downloaded from ADCS.

f. Click the **Append CA Certificate** button. A restart of the server may be required.

g. Navigate to **Maintenance > Restart Options** and click **Restart**. When prompted to confirm, click **OK**. The restart will take two to three minutes on the Expressways. Figure 21-5 illustrates the Trusted CA Certificate page.

Key Topic

Now you need to repeat all these steps on the other Expressway. There are some points of interest for the Expressway-E certificate. If multiple domains are being used, be sure that each domain configured for the Cisco Unified Communications Manager is a part of the Subject Alternative Name on the certificate. As with the Expressway-C certificate, if you are deploying the solution with XMPP federation, the same chat node aliases will be required. For successful validation of received certificates, the Cisco Expressway servers must trust the CA that issued certificates and will be exchanged during the TLS handshake. Therefore, if a different CA was used for the Expressway Core and Expressway Edge, the root CA certificate must be added to the counterpart server. For example, if an ADCS was used to sign the Expressway-C CSR, and a public CA was used to sign the Expressway-E CSR, the public root CA certificate will need to be loaded on both servers, as well as the ADCS root CA certificate.

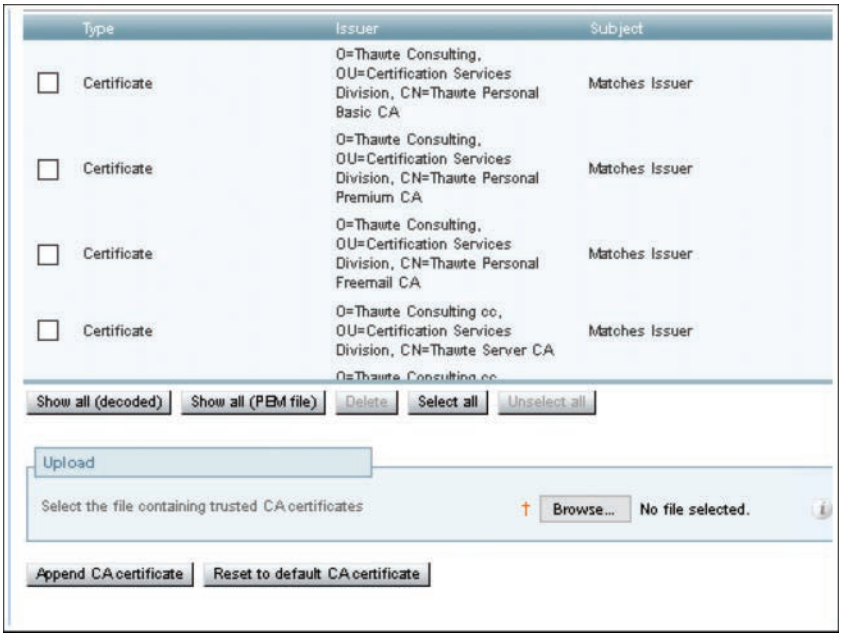

Figure 21-5 *Expressway-C Trusted CA Certificate Page*

Initializing MRA on Expressway Servers

Before configuring a Cisco MRA solution, make sure certain settings are already configured within the Collaboration environment. All endpoints being used with the MRA solution need to be running a version of software that supports this feature. Cisco Jabber 9.7 or later must be used. Starting with this version, Cisco Collaboration Mobile and Remote Access functionality is enabled by default, and the client can identify that it must connect through Cisco Expressway Edge when there is no response to the *_cisco-uds._tcp.<domain>* DNS SRV record lookup request. Administrators should also ensure that the local DNS and public DNS servers are configured with the required SRV records for MRA functionality. The *_cisco-uds._tcp.<domain>* and *_cuplogin._tcp.<domain>* SRV records must not be resolvable from outside the internal network. The Cisco Expressway-C and Cisco Expressway-E should be configured with initial configurations, such as system name, DNS, and NTP at a minimum. Cisco Unified Communications Manager should be configured to allow registrations from Cisco Jabber clients.

Most of the configuration steps to deploy an MRA solution are performed on the Cisco Expressway Core and Cisco Expressway Edge servers. The Cisco Unified Communications Manager Tomcat certificate that must be installed on the Cisco Expressway Core must be obtained from the Cisco Unified Communications Manager server or servers. These certificates are required only if MTLS is being used between the Expressway Core and the Cisco Unified Communications Manager. The following is an overview of the steps to configure MRA on the Expressway servers:

Key Topic

Step 1. Enable MRA on both Cisco Expressways (Core and Edge).

Step 2. Configure MRA on the Cisco Expressway Core.

- **a.** Configure a SIP domain to route registrations to the Cisco Unified Communications Manager.

b. Install the Cisco Unified Communications Manager Tomcat certificate (if TLS Verify is being used).

c. Discover the Cisco Unified Communications Manager from the Expressway Core.

Step 3. Configure a secure traversal zone connection between the Cisco Expressway Edge and the Cisco Expressway Core.

a. Generate a CSR on both Expressways.

b. Sign both CSRs.

c. Install the signed CA and root CA on each respective Cisco Expressway (Core and Edge).

d. Configure a Unified Communications Traversal (Server) Zone on the Expressway Edge.

e. Configure a Unified Communications Traversal (Client) Zone on the Expressway Core.

To enable the Cisco Collaboration Mobile and Remote Access on the Expressway-C, navigate to **Configuration > Unified Communications > Configuration**. Change the Unified Communications Mode to Mobile and Remote Access. All other settings can be left as their defaults. Click Save when finished. On the Expressway-E, the menu path is the same to enable MRA, but the menu options are slightly different. Change the Unified Communications Mode to Mobile and Remote Access and leave all other settings as their defaults. Figure 21-6 illustrates some of the settings available when MRA is enabled on the Expressway-C. Figure 21-7 illustrates the settings available when MRA is enabled on the Expressway-E. Use these figures to compare and contrast the differences between the settings.

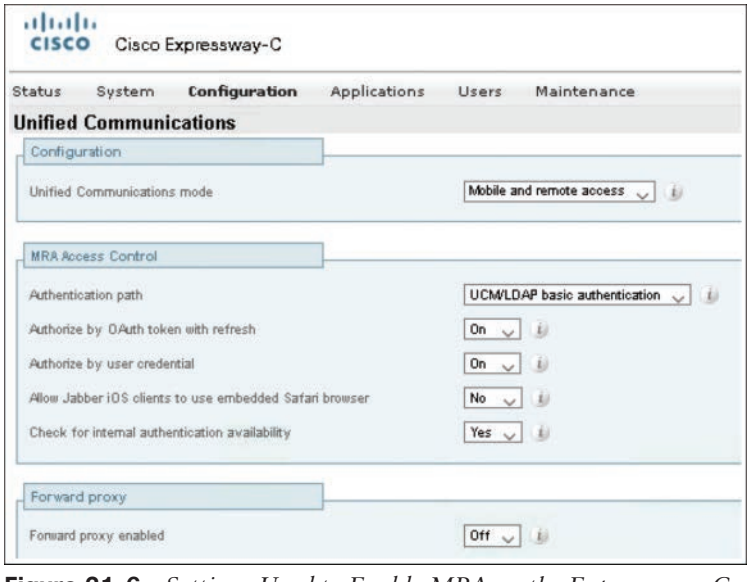

Figure 21-6 *Settings Used to Enable MRA on the Expressway-C*

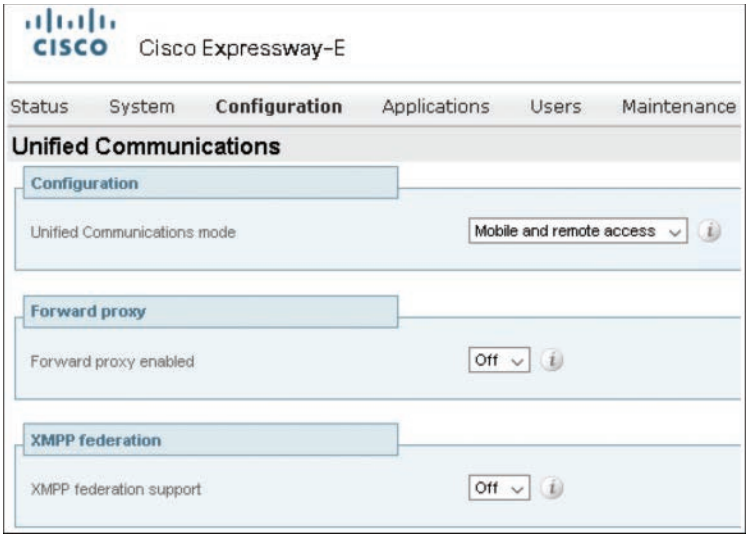

Figure 21-7 *Settings Used to Enable MRA on the Expressway-E*

No more MRA-specific settings have to be configured on the Expressway-E. However, several settings need to be configured on the Expressway-C. First, navigate to **Configuration > Domains**. If a domain was configured previously, an administrator can click that domain to edit the settings. If not, you will need to create a new domain by clicking the New button. When MRA is not in use, there is only one field to configure in the Domains menu; that is the domain itself. When MRA is enabled, several settings will need to be configured. First, configure the domain in the Domain Name field. In the next section, an administrator will need to enable all the services for this domain that will need to be supported using MRA. The options include SIP Registrations and Provisioning on Expressway, SIP Registrations and Provisioning on Unified CM, IM and Presence Service, and XMPP Federation. Figure 21-8 illustrates the DNS settings related to MRA on the Expressway-C.

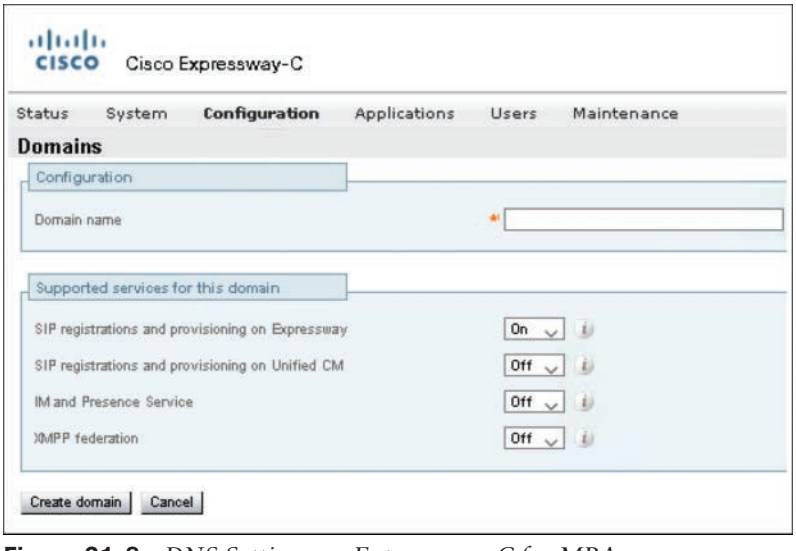

Figure 21-8 *DNS Settings on Expressway-C for MRA*

Before adding any servers to the Expressway-C, such as the Cisco Unified Communications Manager, you will need to add the Tomcat certificate first. This certificate needs to be added only if TLS Verify is being used for communication between the Cisco Unified Communications Manager and the Expressway-C. Navigate to **Maintenance > Security > Trusted CA Certificate**. Under the Upload section, click Browse and select the Tomcat certificate that's intended for this server. Click Open to return to the Trusted CA Certificate page and click the Append CA certificate button to load the certificates. Check the list of certificates to ensure the Tomcat certificate was uploaded successfully.

Now the Cisco Unified Communications Manager can be discovered by the Expressway-C. Navigate to **Configuration > Unified Communications > Unified CM Servers**. Click the Add button and enter the following parameters. If TLS Verify is being used, the Unified CM Publisher Address must be the URL of the Cisco Unified Communications Manager publisher. If TLS Verify is not being used, this address can be the URL or the IP address of the Cisco Unified Communications Manager Publisher. The Username and Password settings should correspond to the AXL application user credentials created on the Cisco Unified Communications Manager. Verify TLS Verify Mode is set to On if TLS Verify is being used. If not, change this setting to Off. When these settings are saved, the Expressway-C will automatically create a neighbor zone to the Cisco Unified Communications Manager. The last setting on this page is the AES GCM Support setting. If it is enabled, the neighbor zone generated for the Cisco Unified Communications Manager will support AES GCM algorithms to encrypt and decrypt media passing through the zone. The default is Off but can be switched to On depending on how the Cisco Unified Communications Manager is configured to handle media encryption. When finished, click the Add Address button. This will return you to the Unified CM Servers page. In the Currently Found Unified CM Nodes section, verify that the discovery status is displayed as Active. Figure 21-9 illustrates the settings that need to be configured when adding a Cisco Unified Communications Manager to the Expressway-C for discovery.

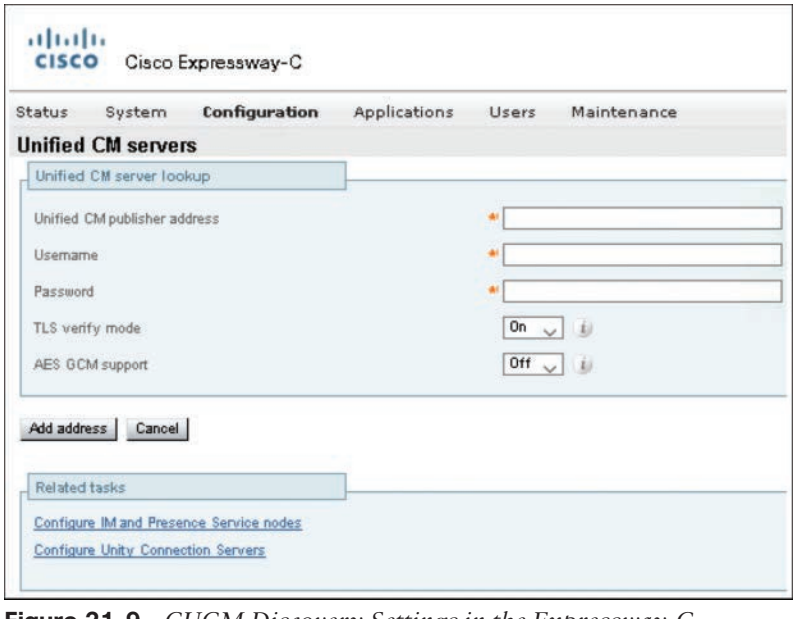

Figure 21-9 *CUCM Discovery Settings in the Expressway-C*

Key Topic

Administrators can navigate to **Configuration > Zones > Zones** and verify that the neighbor zone to the Cisco Unified Communications Manager has been created. Clicking into this zone will display all the settings, but they will be grayed out and cannot be changed. There is also a search rule associated with this zone that was automatically created. Navigate to **Configuration > Dial Plan > Search Rules** to verify this rule exists. These settings will also be grayed out. The Cisco Unified Communications Manager IM and Presence Service and Cisco Unity Connections servers can also be discovered by the Expressway-C if these servers are being used. However, be aware that the status will not show Active on these until the traversal zones are configured and active. Navigate to **Configuration > Unified Communications > IM and Presence Service Nodes** or **Configuration > Unified Communications > Unity Connection Servers** to configure the discovery settings for these servers in the same manner as the Cisco Unified Communications Manager.

Collaboration Traversal Zones and Search Rules

After the MRA settings have been configured and the certificates have been exchanged, the next step is to create traversal zones between the Expressway servers. Traversal zones should always be configured on the traversal server first, in this case the Cisco Expressway-E.

Step 1. Configure the traversal server zone on the Expressway E.

- **a.** On the Expressway-E, navigate to **Configuration > Authentication > Local Database** and add new credentials.
- **b.** Navigate to **Configuration > Zones > Zones** and add a new traversal server zone. Because this zone is specifically for MRA, the zone Type should be set to **Unified Communications Traversal**.
- **c.** Supply the username from the authentication database. Notice that no H.323 settings are available under this zone type. The reason is that H.323 is not supported using MRA. If H.323 calls are to be supported, another traversal zone must be established using the standard traversal zone setup.

  To enable Unified Communications Services on this traversal zone, you must configure the SIP settings to use TLS with TLS Verify Mode enabled, and Media Encryption Mode must be set to Force Encrypted. All of these settings default to the aforementioned parameters, and they cannot be changed. Therefore, these settings will not appear in the Unified Communications Traversal Zone settings page.
- **d.** Supply the SIP TLS Verify Subject Name. The Subject Name must match the subject name or the alternative subject name specified in the Cisco Expressway Core server security certificate. This is the Common Name you should have noted from the CSR of the Expressway Core. If you did not write it down, it is the full URL of the Expressway Core.
- **e.** Set the Accept Proxied Registrations setting to **Allow**.
- **f.** Set the Authentication Policy to **Treat as Authenticated**.
- **g.** Click **Create Zone** when finished.

Step 2. To route calls from the Expressway Edge to the Expressway Core through this traversal server zone, you need to configure a search rule as well.

- **a.** Navigate to **Configuration > Dial Plan > Search Rules**.

b. Click **New**, configure the settings, and then click **Create Search Rule.** Figure 21-10 illustrates some of the Unified Communications traversal zone settings on the Expressway-E.

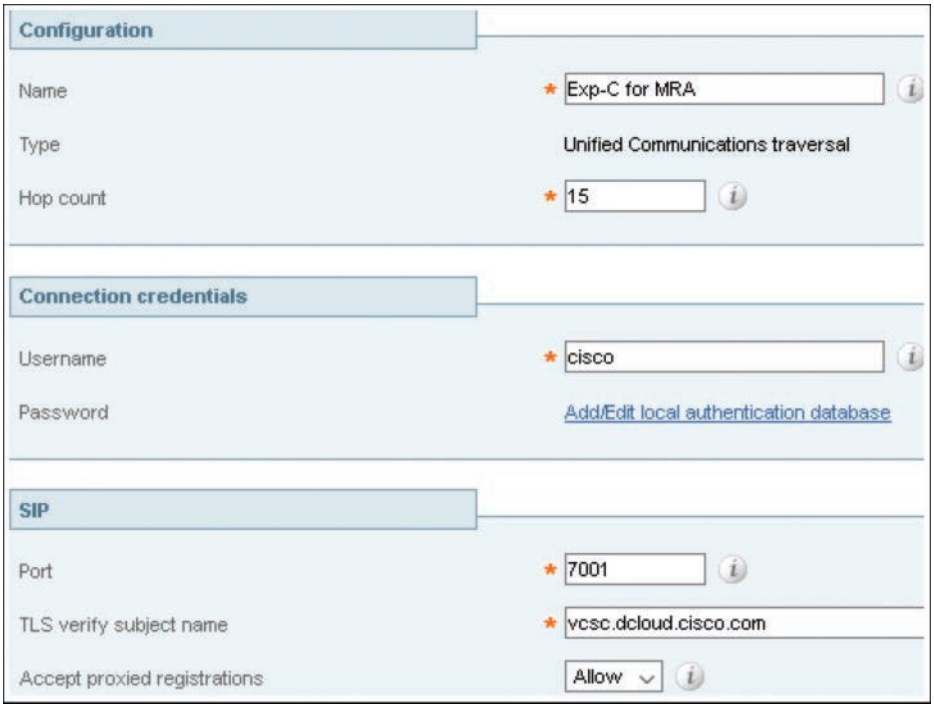

Figure 21-10 *Unified Communications Traversal Zone Settings on the Expressway-E*

Once you've configured the traversal server zone, you can configure the traversal client zone on the Cisco Expressway-C to initiate communication between the two servers.

Step 3. Configure the traversal server zone on the Expressway Core.

a. On the Expressway-C, navigate to **Configuration > Zones > Zones** and add a new traversal client zone. Again, since this zone is specifically for MRA, you should set Zone Type to **Unified Communications Traversal.**

b. Supply the Username and Password that were configured in the Authentication database on the Expressway-E.

c. Enter the port that will be used to establish a connection and keep the connection alive. This port needs to match the SIP keepalive port on the Expressway Edge.

d. Set the Accept Proxied Registrations setting to **Allow.**

e. Set the Authentication Policy to **Treat as Authenticated.**

f. In the Location Peer 1 Address field, enter the URL of the Expressway Edge.

Because TLS Verify is being used, this setting must be in the URL format. It must also match the Common Name you should have noted from the CSR of the Expressway Edge. If the IP address of the Expressway-E is used, communication will fail between the traversal client and the traversal server.

g. Click **Create Zone** when finished.

h. Verify that the state of the zone is Active after saving the client zone.

Step 4. To route calls through the Expressway Edge using this traversal client zone, you need to configure a search rule here as well. However, because calls may be routed to any possible destination, you can use an Any Alias rule. Figure 21-11 illustrates some of the Unified Communications Traversal Zone settings on the Expressway-C.

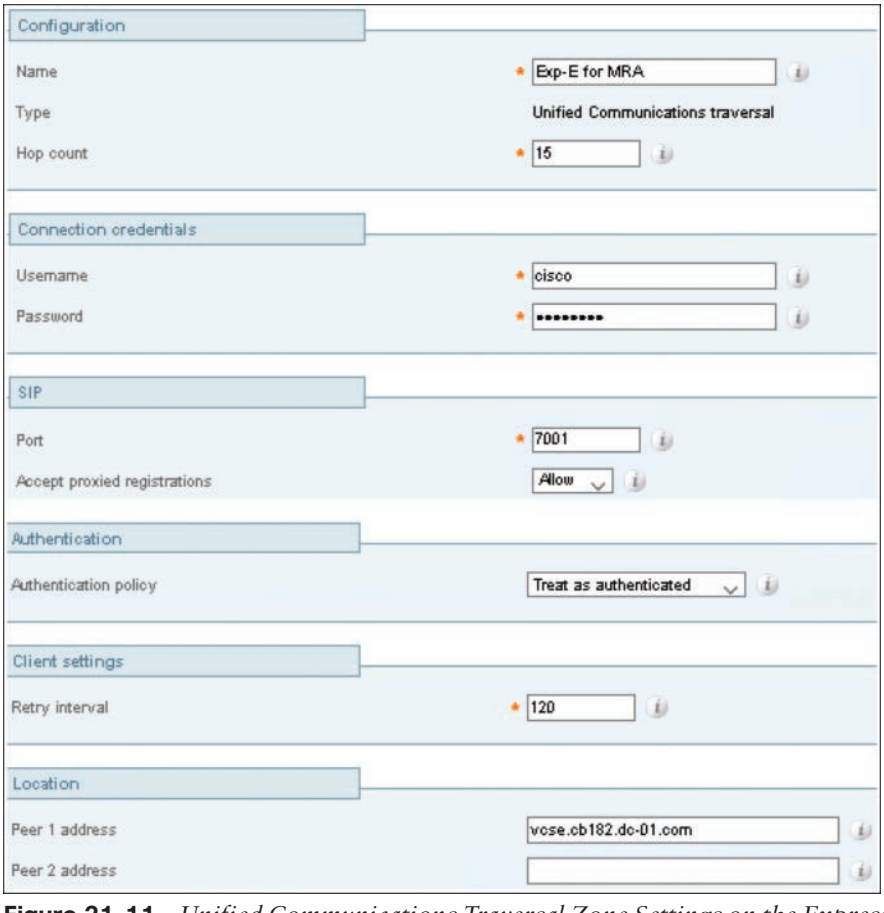

Figure 21-11 *Unified Communications Traversal Zone Settings on the Expressway-C*

The MRA deployment is now complete. You can test these settings by trying to register endpoints located both inside the enterprise network and outside the network. If the zones created will support calls as well, test the deployment by trying to place a few calls. Try calling between an internally registered endpoint and an externally registered endpoint. Then try calling between an internally registered endpoint and an endpoint located in another business network. Also, try calling between an externally registered endpoint and an endpoint located in another business network. You should try all call attempts from both directions: initiated from inside out, and initiated from outside in.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 21-6 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 21-6 Key Topics for Chapter 21

Key Topic Element	Description	Page Number
Paragraph	Standard Traversal Solution versus MRA Traversal Solution	478
Table 21-2	Public DNS SRV Records for Expressway-E Cluster	479
Table 21-3	Private DNS SRV Records for CUCM and CUCM IMP Clusters	480
List	Firewall Traversal Communications Process	480
List	Firewall Ports for MRA	480
Table 21-4	Certificate Pairs Used for MRA Deployments	482
List	Reverse Proxy Ports	482
Paragraph	AXL User Needed on CUCM for MRA	487
Table 21-5	Classes of Certificates on Cisco Collaboration Servers	489
Paragraph	Points of Interest for Certificates on Expressway-C	493
Paragraph	DER and Base64 Comparison	495
Paragraph	Root CN Chaining Explained	495
Paragraph	Points of Interest for Certificates on Expressway-E	496
List	General Steps to Configure MRA	497
Paragraph	Autoconfigured Neighbor Zones for MRA	501

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

A-Record, ADCS, Asymmetric Cryptography, Base64 Encoded Format, CA, CSR, DER Encoded Format, Diffie-Hellman Key Exchange, DNS, DV Certificates, EV Certificates, Firewall, HTTPS Reverse Proxy, MRA, OV Certificates, PKI, Root CA, Root CA Certificate, RSA, SRV, SSL, Symmetric Cryptography, TLS, TLS Verify, Traversal Chaining, Traversal Client Zone, Traversal Server Zone, Traversal Zone, Unified Communications Traversal Zone

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. What are the five differences between a standard firewall traversal solution and MRA?
2. What are the four prerequisites that MRA depends on to be configured?
3. What are the six steps in a basic service discovery from an endpoint using the MRA solution?
4. What are the three main categories for deploying an MRA solution?

CHAPTER 26

Users and Cisco Jabber Soft Clients

This chapter covers the following topics:

Registration Options for Jabber Client: This topic will introduce the various registration options for the Cisco Jabber client, including softphone mode and deskphone mode for on-premises deployments, Hybrid Message Service for Hybrid message and presence interworking with Webex Teams clients, and Team Messaging mode for Jabber registration to Webex Control Hub for messaging and presence.

Configure Cisco Unified Communications Manager for Jabber Client: This topic will examine all the configuration components required to support Jabber registration to the Cisco Unified Communications Manager. These components include configuring end-user accounts, configuring the Cisco Unified Client Services Framework (CSF) phone for Jabber, associating the Cisco Unified CSF phone with the end-user account, and configuring computer telephony integration (CTI) and other components with the Cisco Jabber account.

This chapter takes the information about the Cisco IM and Presence server from the preceding chapter and explains how to apply this information in a practical solution using the Cisco Jabber client. The playing field is changing in the Collaboration world to a more cloud-centric environment. Therefore, this chapter will introduce two new concepts that Cisco engineers have conceived, but due to the nature of this book, the primary focus throughout the rest of this chapter will remain on the deployment of Jabber on-premises. Topics discussed in this chapter include the following:

- Registration Options for Jabber Client
  - Softphone Mode
  - Deskphone Mode
  - Hybrid Message Service
  - Team Messaging Mode
- Configure Cisco Unified Communications Manager for Jabber Client
  - Configure End-User Accounts
  - Configure Cisco Unified Client Services Framework (CSF) Phone
  - Add CSF Phone to End-User Account
  - Computer Telephony Integration (CTI) and Other Considerations

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 6.4 Describe Cisco Unified IMP protocols and deployment
  - 6.4.a XMPP
  - 6.4.b High availability
- 6.5 Deploy Cisco Jabber on-premises

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 26-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 26-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Registration Options for Jabber Client	1–5
Configure Cisco Unified Communications Manager for Jabber Client	6–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is the highest video resolution that can be obtained using the Jabber client for video calls?
   - **a.** 480p30
   - **b.** 720p30
   - **c.** 720p60
   - **d.** 1080p30
2. What DNS SRV record should be used so that Jabber can discover the Cisco IM and Presence Service?
   - **a.** _sip._tcp.
   - **b.** _cisco-uds._tcp
   - **c.** _cuplogin._tcp.
   - **d.** _imp._tcp

3. What protocol does the Cisco Jabber client use to take control over a phone?
   a. CCMCIP
   b. XMPP
   c. SOAP
   d. CTIQBE
4. Which of the following message scenarios is possible with the new Cisco Jabber Hybrid Message Service?
   a. Webex Teams send messages to Jabber clients, but neither client is set up with the Hybrid Message Service.
   b. Jabber clients send messages to Webex Teams, but neither client is set up with the Hybrid Message Service.
   c. Webex Teams clients not enabled for the Hybrid Message Service send messages to Jabber clients set up with the Hybrid Message Service.
   d. Jabber clients not enabled for the Hybrid Message Service send messages to Webex Teams clients set up with the Hybrid Message Service.
5. Which of the following statements about Cisco Jabber Team Messaging mode is true?
   a. The Webex Control Hub replaces the IM and Presence Service.
   b. Calls are not possible with Team Messaging mode.
   c. Presence will suffer long delays updating status using Team Messaging mode.
   d. Calls are possible with Team Messaging mode, but voicemail is not available.
6. Which of the following settings are required when creating an end-user account in the CUCM?
   a. Password
   b. PIN
   c. Last Name
   d. Display Name
7. Which of the following settings cannot be configured in the CUCM until after an end-user account is created?
   a. Manager User ID
   b. Controlled Devices
   c. UC Service Profile
   d. Presence Gateway to be configured on CUCM IM and Presence server
8. Which of the following clients must be combined with Cisco Unified Video Advantage so that video capabilities can be extended to that soft client?
   a. Cisco Unified IP Communicator
   b. Cisco Unified Personal Communicator
   c. Cisco Webex Connect
   d. Movi

9. Which of the following statements is true?
   a. If the CSF phone is associated with an end user, that end user does not need to be associated with the phone.
   b. If an end user is associated with a CSF phone, that phone does not need to be associated with the end user.
   c. Both the end user and the CSF phone must be associated with each other.
   d. Only the Cisco Unified CSF phone has to be associated with an end user, which is why the end-user account must be configured first.
10. Which of the following is a required field when configuring CTI as a UC Service on the CUCM for Jabber deskphone control?
   a. Specify a name for the UC service.
   b. Configure the CTI service address.
   c. Specify the port number for the CTI service.
   d. All of these answers are correct.

Foundation Topics

Registration Options for Jabber Client

Cisco Jabber is a Cisco Unified Communications client application for Microsoft or Apple computers, tablets, and smartphones. Cisco Jabber client applications provide IM and presence, voice and video, voice messaging, desktop sharing, and other collaborative workspace capabilities that support 720p30 high-definition video interoperability. Cisco Jabber uses the Cisco Precision Video Engine and ClearPath technology to optimize video media. The Cisco Precision Video Engine uses fast video-rate adaptation to negotiate optimum video quality, based on network conditions. Cisco Jabber clients are available for many devices and can be used with on-premises or cloud-based Cisco Unified Communications services, or with the Cisco Webex solution in the cloud. Cisco Jabber clients provide a consistent experience across devices. These desktop clients include Cisco Jabber for Windows and Cisco Jabber for Mac. Jabber can also be used on smart devices such as Apple or Android tablets or smartphones.

Cisco Jabber takes advantage of Cisco Medianet-based networks and Cisco Unified Communications Manager call control to deliver secure, clear, and reliable communications. It streamlines communications and enhances productivity by securely unifying IM and presence, video, voice, voice messaging, desktop sharing, and conferencing capabilities into one client on your desktop. It also delivers highly secure, clear, and reliable communications. Cisco Jabber offers flexible deployment models, is built on open standards, and integrates with commonly used desktop applications. The Cisco Jabber client allows users to communicate and collaborate effectively from anywhere that has an Internet connection.

Cisco Jabber is based on the Cisco Unified Client Services Framework (CSF) and combines advanced collaborative media features with Cisco Unified Communications. It uses SIP for call control, XMPP for IM, and CTI for deskphone control. Cisco Jabber is a software-based service framework that supports multiple lines and enables the following client service applications:

Key Topic

- Audio and video calls with integrated multipoint conferencing control
- Deskphone control

- Instant messaging (IM)
- Presence
- Advanced voicemail control
- Communication history management

Cisco Jabber operates in one of two modes. In *deskphone mode*, the Cisco Jabber client controls the Cisco IP phone of the user. In *softphone mode*, the Cisco Jabber client behaves like an IP phone and originates and terminates all audio and video communication interactions on the software endpoint itself. Cisco Jabber supports video in both softphone and deskphone modes. Video resolution is supported from QCIF and CIF through VGA up to 720p high-definition video with up to 30 fps. Cisco Jabber also supports integrated ad hoc video conferencing control, movable self-view, audio and video mute, and desktop sharing by using standard video conferencing that is based on SIP BFCP content-sharing protocols. It can operate in one of the following video options:

- Cisco Jabber operating in softphone mode terminates both audio and video sessions on the host computer as mentioned before.
- Cisco Jabber operating in deskphone mode can use one of two options:
  - The first option uses a Cisco Unified IP phone that is not video capable. In this situation, the audio terminates on the Cisco Unified IP phone, and video terminates on the host computer. Audio and video streams are mutually synchronized with the Cisco Audio Session Tunnel (CAST) protocol between Cisco Jabber and the Cisco Unified IP phone. For CAST to work, the host computer must connect to the PC port on the Cisco Unified IP phone, and IP routing must work between the voice (deskphone) and data (PC with Cisco Jabber) VLANs.
  - The second option of deskphone mode uses a video-capable Cisco Unified IP phone, which terminates both audio and video sessions. On a Cisco Unified IP Phone 8845 or 8865, or on the Cisco DX70 or DX80 video endpoint, the Cisco Jabber desktop client detects a connected camera on the CTI-controlled IP phone, and both voice and video are serviced through that IP phone or Telepresence endpoint.

Figure 26-1 illustrates the three video options available to Cisco Jabber.

Softphone Mode

The first three steps of the registration process of a Cisco Jabber client are the same whether the client is operating in softphone mode or deskphone mode. Cisco Jabber takes all IP network parameters from the host computer. The IP address, network mask, default router, and DNS must be properly set up on the host computer before the client is launched. If DNS was configured on the back end with the SRV record for Cisco Jabber to discover the IM and Presence server (_cuplogin._tcp.), then when the client is launched, it requires only two parameters:

- **User ID:** *<username@domain>* (The domain is used to look up the SRV record of the Cisco Unified Communications Manager IM and Presence Service.)
- **Password**

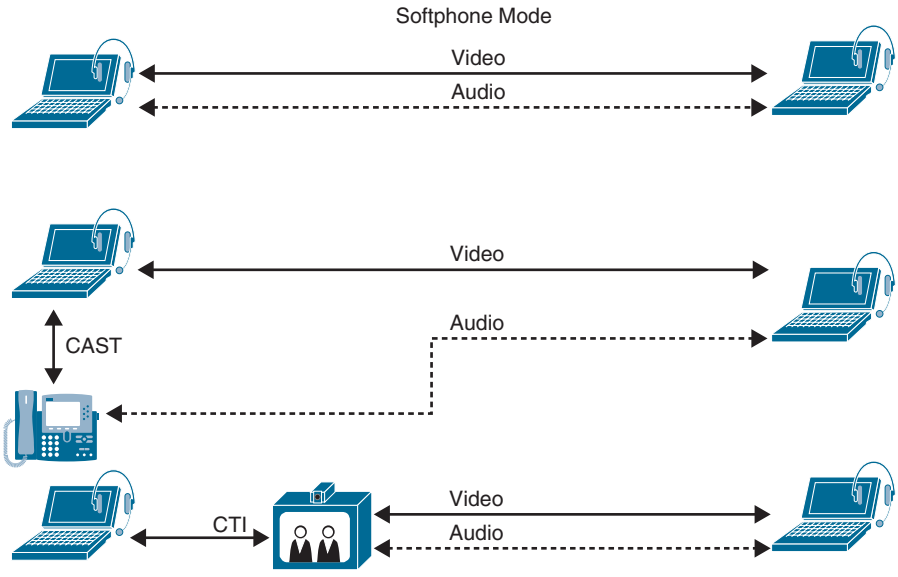

Figure 26-1 *Three Video Options Available to Cisco Jabber*

If the SRV record cannot be discovered, an Advanced Settings option will allow the user to enter the Cisco Unified Communications Manager IM and Presence Service IP address or name manually. Cisco Jabber sends the user ID and password to the indicated Cisco Unified Communications Manager IM and Presence Service after the address is obtained. The Cisco Unified Communications Manager IM and Presence Service verifies whether the user is licensed to use Cisco Jabber and then sends the user ID and password to Cisco Unified Communications Manager for authentication. Cisco Unified Communications Manager verifies the credentials locally, against its own database, or it uses an external LDAP directory server if integrated. After verification, authentication is acknowledged, and the registration process can continue as follows:

1. When Cisco Jabber launches and a user logs in, the client makes a secure connection using SOAP over HTTPS to the Cisco Unified Communications Manager IM and Presence Service and downloads the following information:
   - Contact list
   - Client profiles, such as CTI, Cisco Unified Communications Manager IP Phone Service, and LDAP
   - User settings
   - Portable user preferences
2. The client registers to receive presence information from the Cisco Unified Communications Manager IM and Presence Service. Via the Extensible Messaging and Presence Protocol (XMPP), the client receives presence information for each contact in the contact list.
3. Based on the Cisco Unified Communications Manager IP phone profile that Cisco Jabber downloaded initially, the client uses the Cisco Unified Communications

Manager IP Phone Service over HTTPS to contact Cisco Unified Communications Manager. The client then receives the list of endpoint devices that are associated with the end user. When Cisco Jabber discovers endpoint devices that are associated with the user, a Cisco Unified Client Services Framework device can define softphone functionality. In general, the Cisco Unified Client Services Framework is defined in Cisco Unified Communications Manager using a name, such as *jwhite*. The following steps are unique for the Cisco Jabber softphone mode registration process.

4. Cisco Jabber, now operating in softphone mode, contacts the Cisco Trivial File Transfer Protocol (TFTP) server to download the configuration for the softphone device. The TFTP address was obtained as part of the profiles that were downloaded using the Simple Object Access Protocol (SOAP). The name of the Cisco Unified Client Services Framework device in the Cisco Unified Communications Manager database is used to identify which configuration file to download. If the name was jwhite, this file would be called jwhiteclient.cnf.xml. If security was used, the file ctlsepjwhiteclient.tlv would be retrieved from the TFTP server before the configuration file. Cisco TFTP sends the jwhite.cnf.xml file to the Jabber client.
5. Cisco Jabber then uses that information to register using SIP to the primary Cisco Unified Communications Manager for call processing. Cisco Jabber registers as the video-capable softphone to Cisco Unified Communications Manager, using a SIP REGISTER message. Cisco Unified Communications Manager sends an acknowledgment, using a "SIP 200 OK" response. Figure 26-2 illustrates how the softphone registration process works for Cisco Jabber.

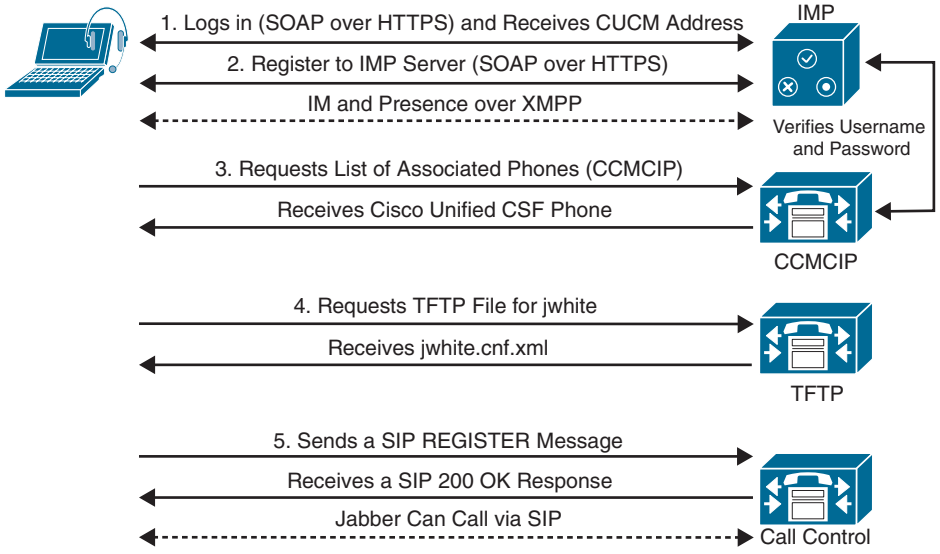

Figure 26-2 *Softphone Registration Process for Cisco Jabber Client*

After Jabber has successfully registered, you can click the computer icon at the bottom-right corner of Cisco Jabber to expand the menu. When Cisco Jabber, operating in softphone mode, has registered successfully, the menu shows a check mark next to the Use My Computer for Calls option, and the presence status at the top of the window is Not Unknown (gray color).

Deskphone Mode

The registration process for Cisco Jabber operating in deskphone mode is simplified and does not include the TFTP and SIP registration stages. However, it does initialize CTI communication to Cisco Unified Communications Manager CTI Manager for deskphone control. The deskphone mode process associates Cisco Jabber with a Cisco Unified IP phone that can support only audio or both audio and video. The Cisco Jabber client obtains all the IP network parameters, such as the IP address, network mask, default gateway, and DNS address if used, from the hosting computer or device, which typically obtains this information through a DHCP server. Therefore, obtaining this network information is not part of the Jabber registration process. The Cisco Jabber client will use the following steps to register using deskphone mode:

1. When Cisco Jabber launches and a user logs in, the client makes a secure connection using SOAP over HTTPS to the Cisco Unified Communications Manager IM and Presence Service and downloads the following information:
   - Contact list
   - Client profiles, such as CTI, Cisco Unified Communications Manager IP Phone Service, and LDAP
   - User settings
   - Portable user preferences
2. The client registers to receive presence information from the Cisco Unified Communications Manager IM and Presence Service. Via the XMPP protocol, the client receives presence information for each contact in the contact list.
3. Based on the Cisco Unified Communications Manager IP phone profile that Cisco Jabber downloaded initially, the client uses the Cisco Unified Communications Manager IP Phone Service over HTTPS to contact Cisco Unified Communications Manager. The client then receives the list of endpoint devices that are associated with the end user.
4. To control an associated deskphone, Cisco Jabber initializes CTIQBE to the Cisco Unified Communications Manager CTI Manager, which is usually collocated with Cisco Unified Communications Manager but can be a separate server for scalability reasons. If the user did not specify a preferred endpoint device in Cisco Jabber, the endpoint with the line that is configured in Cisco Unified Communications Manager Administration is used as the primary line if the user becomes the preferred CTI endpoint device.
5. If desktop video is shown on Cisco Jabber with a Cisco Unified IP phone that is audio-capable only, the Cisco Audio Session Tunnel protocol is used to synchronize video and audio media. In this case, Cisco Discovery Protocol is used at Cisco Jabber

start-up to discover whether the hosting computer connects to the PC port at the controlled IP phone. Figure 26-3 illustrates how the deskphone registration process works for Cisco Jabber.

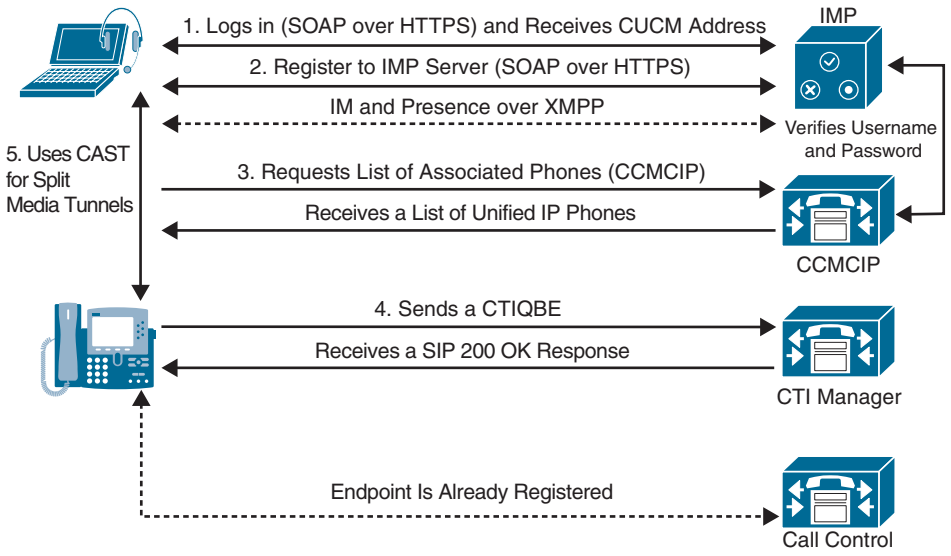

Figure 26-3 *Deskphone Registration Process for Cisco Jabber Client*

The phone being controlled using CTI does not have to be physically connected to the computer running Jabber with a patch cable as long as CAST isn't being used. After Jabber has successfully registered, you can click the computer icon at the bottom-right corner of Cisco Jabber to expand the menu. When Cisco Jabber, operating in deskphone mode, has registered successfully, the menu shows a check mark next to the Use My Phone for Calls option, and the presence status at the top of the window is Not Unknown (gray color).

Hybrid Message Service

Several new deployment options are available for businesses using Cisco Jabber. Historically, Cisco Jabber has been available through an on-premises deployment only. There is now a cloud-based deployment for Cisco Jabber, which uses Cisco Webex to host services. This option is referred to as *Team Messaging mode*. Another option allows Cisco Jabber clients deployed on-premises to integrate with the messaging functionality of Cisco Webex Teams. This service is referred to as the *Hybrid Message Service*.

The Cisco message to customers today is "Cloud First, not Cloud Only." Although Jabber continues to be Cisco's primary on-premises application offering, with Jabber version 12.5, Cisco has introduced a unified user experience aligning the Jabber user interface with that of Webex Teams. This will create easier transitions should companies wish to move their users from Cisco Jabber to a full Webex Teams deployment. Jabber 12.5, which was released in January 2019, offers that same great interface as Webex Teams, plus Team Messaging mode capabilities, meetings tools when used with Cisco Meeting Server, and media optimization for a better user experience during calls. Jabber 12.6 was released in April 2019 and offers all the same great tools as the version 12.5 release, plus JVDI and Intelligent Proximity support. Jabber 12.7 was released in October 2019 and has brought continued enhancements and support, along with additional features.

Cisco Jabber Softphone for VDI extends the Cisco Collaboration experience to virtual deployments. With a supported version of Cisco Jabber for Windows, users can send and receive phone calls on their hosted virtual desktops. The Cisco Jabber Softphone for VDI software detects the virtual environment and routes all audio and video streams directly from one endpoint to another, without going through the hosted virtual desktops. Figure 26-4 illustrates the similarities between the new Cisco Jabber client and the Cisco Webex Teams client.

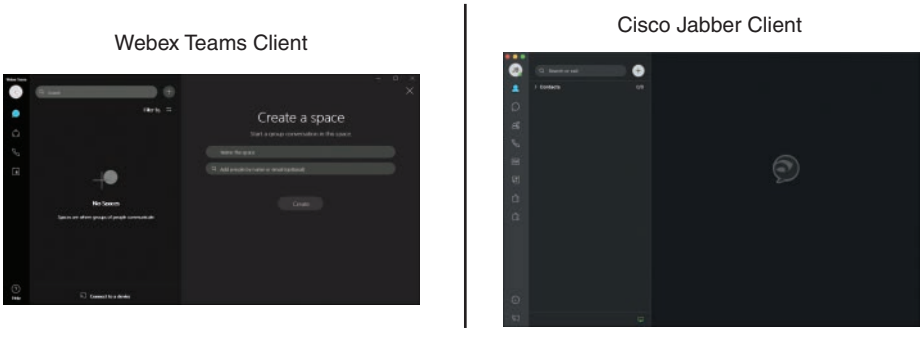

Figure 26-4 *Comparison of the Cisco Jabber and Cisco Webex Teams Clients*

Hybrid Message Service involves an on-premises deployment of Cisco Jabber integrating with a cloud deployment of Cisco Webex Teams. As with any other hybrid integration, a connector card along with other settings must be configured for the integration to perform correctly. After a Hybrid Message Service integration has been configured, the Jabber client can engage in chat messaging with Webex Teams applications, whether they are 1:1 conversations or group spaces. The following services are available in a hybrid cloud-based deployment that uses Webex Messenger service:

- The Cisco Webex Messenger service provides contact resolution.
- The Cisco Webex Messenger service allows users to publish their presence availability and subscribe to other users' presence availability.
- The Cisco Webex Messenger service allows users to send and receive instant messages between Jabber and Webex Teams.

Audio calling, video calling, conferencing, and voicemail services are all available and continue to operate through the Cisco Unified Communications Manager just as they always have.

As the Hybrid Message Service is implemented, a company's user population may progress as follows. At the start, all users are on Cisco Jabber. Then the company places an order for the Webex solution and implements it through the Webex Control Hub. After the Control Hub is set up, users can be integrated into Webex through an Active Directory integration. Assuming the users in the Cisco Unified Communications Manager are also integrated through Active Directory, this will ensure that the same users will have accounts with Jabber and with Webex Teams. The next task that must be performed is to enable Hybrid Message Service for some of the users. After this service has been enabled, there are four ways to *send* chat messages:

Key Topic

- A user without Message Service uses Webex Teams.
- A user without Message Service uses Cisco Jabber.

- A user with Message Service uses Webex Teams.
- A user with Message Service uses Cisco Jabber.

The recipients can also use those four ways to *receive* chat messages, which means there are 16 possible interactions. You can expect eight of them to work because they are interactions between the same clients, such as Jabber to Jabber or Webex Teams to Webex Teams. Of the remaining eight interop scenarios, you should expect three cases to fail. Figure 26-5 illustrates these 16 possible interactions between Cisco Jabber and Webex Teams using Hybrid Message Service.

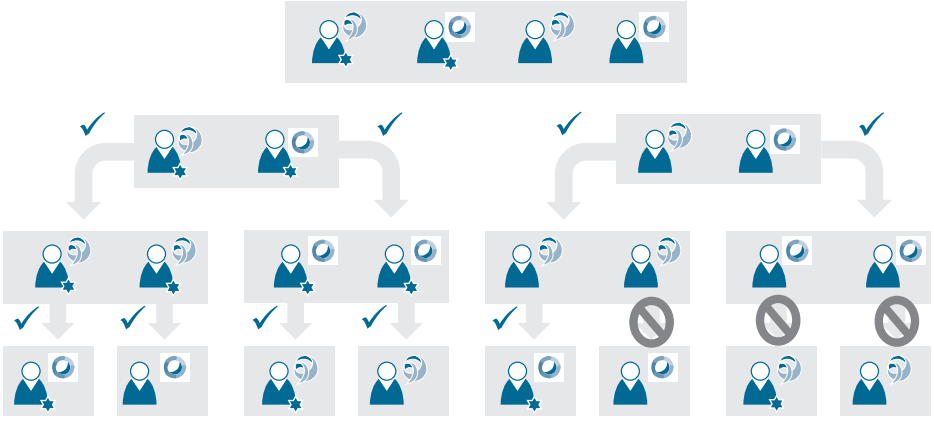

Figure 26-5 *Interactions Between Cisco Jabber and Webex Teams Using Hybrid Message Service*

In Figure 26-5, the box at the top of the image shows all the different meeting clients available in a Cisco Collaboration solution. From left to right, there are the Jabber client with Hybrid Message Service enabled, Webex Teams with Hybrid Message Service enabled, Jabber client without Hybrid Message Service enabled, and Webex Teams without Hybrid Message Service enabled. Moving down the figure are all of the one-way connections that are possible messaging attempts that could occur between each of these clients. Assuming the initiator of the chat message does have Hybrid Message Service enabled, the recipient will see the conversation regardless of whether Hybrid Message Service is enabled. Without Hybrid Message Service enabled, Webex Teams can send messages to other Webex Teams users, and Jabber clients can send messages to other Jabber clients.

However, if the sender uses Webex Teams to send a message to a Jabber client, and neither client is set up with Hybrid Message Service, the message will not be sent. The same is true if a user tries to send a message using Jabber to a Webex Teams client, and neither client is set up with Hybrid Message Service. If the Webex Teams client not enabled for Hybrid Message Service sends a message to a Jabber client that is set up with Hybrid Message Service, the message would still not be sent. This behavior is expected. Cisco designed the service this way to reduce load because it anticipated that organizations would enable Hybrid Message Service for the early adopters and that they would use Webex Teams as their primary chat client. Conversely, if the Jabber client not enabled for Hybrid Message Service sends a message to a Webex Teams client that is set up with Hybrid Message Service, the message will be sent.

Team Messaging Mode

By contrast, Team Messaging mode enables you to register Cisco Jabber directly to the Webex Control Hub. Five services are available in a cloud-based deployment using Team Messaging mode:

Key Topic

- **Cisco Team Messaging Mode:** Provides contact resolution through the Webex Control Hub.
- **Cisco Team Messaging Mode:** Lets users show their presence availability and see other users' availability.
- **Cisco Team Messaging Mode:** Lets users send and receive instant messages.
- **Cisco Webex Meetings Center:** Provides hosted meeting capabilities that can be joined through Jabber. This is not a change in service capability; it is available to Jabber users regardless of where the client is registered.
- **Voice and Video Calling:** Can be added by registering Jabber via SIP to the Cisco Unified Communications Manager after Team Messaging mode has been configured. This is the only piece that connects to an on-premises device, and calling cannot occur in Jabber without it.

Cisco Jabber Team Messaging mode offers an alternative path to the traditional Jabber on-premises deployment, which uses the IM and Presence server for XMPP messaging. Instead, Jabber Team Messaging allows the Webex Control Hub to be used as a replacement for the IM and Presence server, and all messaging and presence communication traverses through the same media path as Webex Teams messaging. This allows a seamless integration between Webex Teams clients and the Cisco Jabber client. This also allows for easier migration from Cisco Jabber to Webex Teams. Figure 26-6 illustrates the infrastructure used to support the Jabber Team Messaging mode solution.

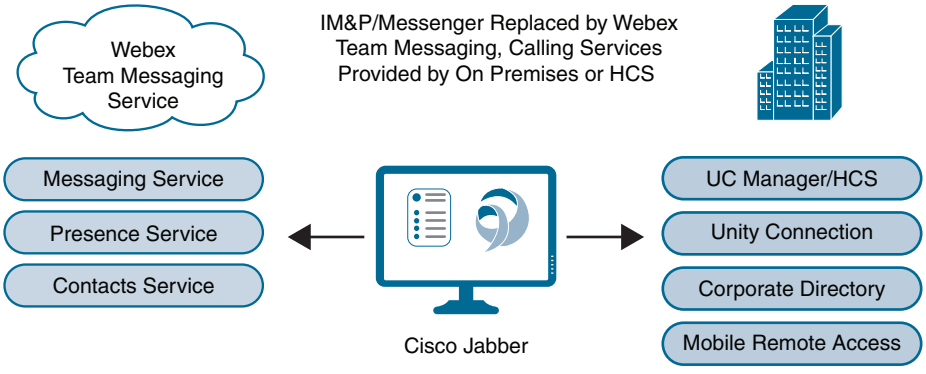

Figure 26-6 *Jabber Team Messaging Mode Infrastructure*

As mentioned earlier, the services provided to Jabber through Team Messaging mode include instant messaging, presence services, and contact services. If voice and video calling is still required through Jabber, the client can register with the Cisco Unified Communications Manager. With a CUCM registration, the Jabber client can still utilize all the benefits that coincide with a CUCM registration, such as voicemail through Unity Connection, MRA, and the corporate directory.

Because Jabber 12.5 and later releases provide a seamless user interface with the Webex Teams App, usability is easily transferable. There are no limitations in what you can do with Jabber as compared to Webex Teams. Team messaging through Jabber supports 1:1 conversations and group spaces, member lists, moderator controls for meetings, and participant lists. Presence status is a seamless integration as well. When Hybrid Message Service is used, there is a delay in sharing presence information. But when Jabber Messaging mode is used, presence status is updated and shared instantaneously.

Configure Cisco Unified Communications Manager for Jabber Client

Several settings must be configured before the Cisco Jabber client will be able to register and access the services on the Cisco IM and Presence server and the Cisco Unified Communications Manager. Many of these settings were discussed in the preceding chapter. The components discussed in this section, which must be configured on the Cisco Unified Communications Manager to support the Cisco Jabber client, are as follows:

- Configure End-User Accounts
- Configure Cisco Unified Client Services Framework (CSF) Phone
- Add CSF Phone to End-User Account
- Computer Telephony Integration (CTI) and Other Considerations

Configure End-User Accounts

Chapter 16, "LDAP Integration with Cisco Unified Communications Manager," briefly discussed the steps necessary to manually add an end user to the Cisco Unified Communications Manager and then discussed in depth how to import users from an LDAP directory service. This section will not discuss adding users using LDAP; instead, it will examine means of adding end users manually as an alternative to the method described in Chapter 16.

The settings to add an end user described in Chapter 16 required an administrator to navigate in the Cisco Unified Communications Manager web interface to **User Management > User/Phone Add > Quick User/Phone Add**. Another way to create an end user is to navigate to **User Management > End User** and click Add New. Some settings must be configured before the new end-user account can be saved, and other settings cannot be configured until the new account is saved. All settings with an asterisk are required settings, which include User ID and Last Name. Before saving the settings, notice that the Controlled Devices buttons do not exist yet, and the Permissions Information section at the bottom of the page is missing. These settings will appear after the end-user account is saved. For best practice, ensure the following settings are configured initially before saving the end-user account:

- User Information:
  - User ID
  - Password
  - Confirm Password

- Self-Service User ID (only if the user portal will be enabled and used for this end user)
- PIN (again, only if the user portal will be enabled and used for this end user)
- Confirm PIN
- Last Name
- First Name
- Display Name
- Directory Number (Refer to the section titled "Endpoint Addressing" in Chapter 18, "Cisco Unified Communications Manager Call Admission Control (CAC)," for a reminder of the function this field performs.)
- Manager User ID (for reporting purposes that will be discussed in later chapters)

- Service Settings:
  - Home Cluster (checked)
  - Enable User for Unified CM IM and Presence (Configure IM and Presence in the Associated UC Service Profile) (checked)
  - Include Meeting Information in Presence (Requires Exchange Presence Gateway to Be Configured on CUCM IM and Presence Server) (checked)
  - UC Service Profile (Select the appropriate service profile for Jabber operations, which is usually the system default.)

No other settings need to be configured on the End User Configuration page at this time. The administrator needs to return to the end-user account to finish configuration for Jabber client support after the phone is configured in a different menu. Figure 26-7 illustrates the end-user settings that should be configured when setting up a user account manually for the first time.

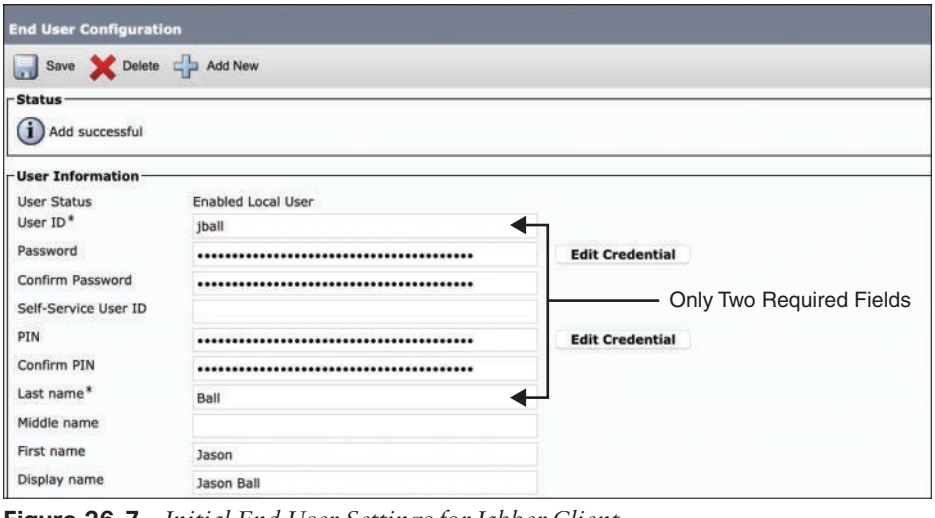

Figure 26-7 *Initial End User Settings for Jabber Client*

Configure Cisco Unified Client Services Framework (CSF) Phone

The Cisco Unified Client Services Framework (CSF) phone in Cisco Unified Communications Manager is a phone configuration template designed to deliver Cisco UC settings and services to soft clients, but it did not originate with Cisco Jabber. The Cisco Unified IP Communicator was Cisco's first attempt at a soft client. This Microsoft Windows PC–based softphone application uses Skinny Call Control Protocol (SCCP) as the call control protocol. This application was designed to bring the voice telephony functions of Cisco Unified IP phones to the desktop over a software application. The Cisco Unified Video Advantage application is a layered application over the Cisco Unified IP Communicator Softphone that brought video to the IP Communicator and Cisco Unified IP phones (via CAST) by using a web camera on the hosting computer.

Later, Cisco launched a new client application, which evolved from the Cisco Unified IP Communicator. This new application not only supported voice and video in a single application, but it also supported instant messaging and presence. This new client application was known as the Cisco Unified Personal Communicator, or CUPC, and it too was based on the Cisco Unified Client Service Framework.

CUPC lasted for a few years, but Cisco continuously strives to be number one with the products it offers business communities around the world. Therefore, Cisco made three key acquisitions that aided in providing the best soft client on the market to date. The first company Cisco acquired was called Jabber. The second company, called Tandberg, had a soft client called Movi. Jabber was a much more feature-rich client in terms of the number of services available, but Movi, which was later named Cisco Jabber Video for Telepresence, supports high-definition video and desktop- and application-sharing options. The third company was called Webex, and it had a soft client called Webex Connect. This client connected to the Webex Cloud and supported voice, video, IM, and presence.

Cisco took the best features of all these new soft clients, combined with the rich features of the Cisco Unified IP Communicator and CUPC, and built a soft client application that has endured the test of time and proven to be the best soft client available today. Based on its popularity, Cisco decided to dub this soft client Cisco Jabber, but it has evolved far from the original Jabber client that Cisco bought so many years ago.

The biggest competitor product to Cisco Jabber is the Microsoft soft client, which has gone by many names over the years. It began as Microsoft OCS, then later became Microsoft Lync, then Skype for Business, and now it goes by the name Microsoft Teams. The greatest downside to the Microsoft solution is that it still uses proprietary protocols, rendering them incompatible cross-vendor. Cisco Unified Communications Integration for Microsoft, which is available with the Cisco Unified Client Service Framework, integrates Cisco Unified Communications and Cisco Collaboration with Microsoft. This is just one more way Cisco is setting the standard for collaboration today.

As mentioned previously, Cisco Jabber is built on the Cisco Unified Client Services Framework, which is a software application that combines several services into a single integrated client. An underlying framework is provided for integration of Cisco Unified Communications services, including audio, video, web collaboration, visual voicemail, and so on, into a single IM and presence application.

The following steps outline how to configure a Cisco Unified CSF phone on the Cisco Unified Communications Manager for the Cisco Jabber client registration. To differentiate the phone that will be created for Jabber and the user associated with the phone, the administrator may want to use different names for the end-user account and the CSF phone.

Key Topic

Step 1. From the Cisco Unified Communications Manager Administration page, navigate to **Devices > Phones** and click **Add New**.

Step 2. From the Phone Type drop-down menu, choose **Cisco Unified Client Services Framework** and click **Next**.

Step 3. In the Device Name field, enter the Cisco Jabber client username, and configure the following fields:

- **Description:** (Optional field)
- **Device Pool:** (Required field; choose from the drop-down list)
- **Phone Button Template: Standard Client Services Framework**
- **Phone Common Profile: Standard Common Phone Profile**
- **Calling Search Space:** (Optional field; choose from a drop-down list)
- **Owner: USER**
- **Owner User ID:** (Required field for Jabber; choose the associated user from the drop-down list)
- **Device Security Profile: Cisco Unified Client Services Framework-Standard SIP Non-Secure**
- **SIP Profile: Standard SIP Profile**

Step 4. When you are finished, click **Save**.

Step 5. After the page refreshes, click the **Line [1] – Add a new DN** link in the left column.

Step 6. On the new page that appears, enter the Directory Number that should be assigned to Cisco Jabber for calling, and choose a Route Partition if required. Click **Save**.

When you are entering the Directory Number, the page needs to refresh before it can be saved. The refresh will happen automatically when you perform one of two actions. You can either press the Tab key to move to the next field or use the mouse to click anywhere else on the screen. This is also true when configuring the Route Partition setting. Figure 26-8 illustrates how these settings might look for a Cisco Unified CSF phone.

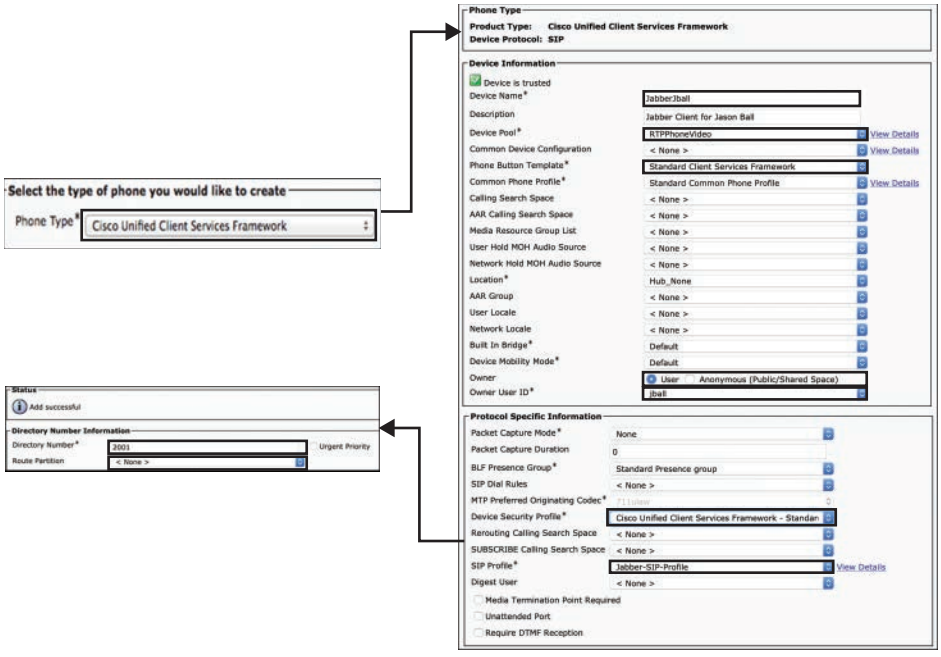

Figure 26-8 *Cisco Unified CSF Phone Settings on CUCM*

Add CSF Phone to End-User Account

After the end-user account has been created and the Cisco Unified CSF phone has been created and associated with the end user, the end-user account needs to be associated with the Cisco Unified CSF phone. You can use the following steps to associate the end-user account with the Cisco Unified CSF phone:

Step 1. From the Cisco Unified Communications Manager Administration page, navigate to **User Management > End Users**.

Step 2. Use the Find button to locate your username. When you find your username, click it.

Step 3. In the middle of the end-user page, locate the Device Information area and click **Device Association**.

Step 4. On the new page that appears, click **Find** and locate the Cisco Unified Client Services Framework device that you configured previously.

Step 5. Check the box beside it, and then click **Save Selected/Changes**.

Step 6. In the top-right section of the page, click **Go** next to the Related Links: Back to User drop-down menu to get back to the end-user page.

For your end user, set the primary extension to the earlier configured directory number. Locate the Directory Number Associations section. From the Directory Number* drop-down list, choose the entry with your directory number.

Step 7. Click **Save** to save the end-user configuration.

Your Jabber client should now be able to sign in and register to the Cisco Unified IM and Presence server and the Cisco Unified Communications Manager. Figure 26-9 illustrates the preceding outlined settings used to associate the Cisco Unified CSF phone with the end-user account.

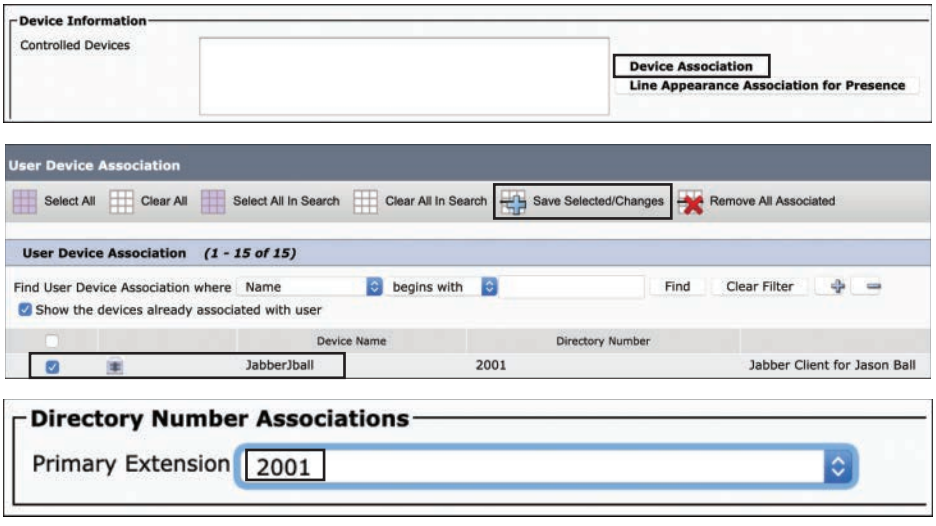

Figure 26-9 *Associating the Cisco Unified CSF Phone with an End-User Account*

Computer Telephony Integration (CTI) and Other Considerations

Key Topic

Computer telephony integration (CTI) is a service in the Cisco Unified Communications Manager that enables computer and telephone systems to interact together. CTI is most commonly used by call centers handling a large number of incoming calls. Call centers implementing CTI can use computers to manage all telephone calls, which in turn leads to increased efficiency and better results. A CTI port is a virtual port. It is analogous to a trunk line in a traditional ACD or PBX setting. One such use of a CTI port is to allow access to the post routing capabilities of the IP-IVR. A CTI route point is a virtual device that can receive multiple simultaneous calls for the purpose of application-controlled redirection. A CTI route point is another useful virtual device; this one can receive multiple calls at once and efficiently distribute them to the various CTI ports. The manner in which the calls are distributed is based on application-controlled redirection. A program called CTIManager includes the CTI components that interface with the applications that are separated out of Cisco Unified Communications Manager. Cisco IM and Presence Service CTI configuration (CTI Server and Profile) for use with Cisco Jabber is automatically created during the database synchronization with Cisco Unified Communications Manager. All Cisco Jabber CTI communication occurs directly with the Cisco Unified Communications Manager and not through the Cisco IM and Presence Service.

Adding the CTI service capabilities on the Cisco Unified Communications Manager for Cisco Jabber is not always a straightforward process. Different Cisco Unified IP phones may have slightly different requirements when it comes to configuring Jabber desk-phone control. However, some common settings should be configured for every instance

of Jabber deskphone control using CTI. You can use the following steps to configure these settings:

Step 1. Add a CTI UC service, and then add the CTI UC service to your service profile. The CTI UC service provides Jabber with the address of the User Data Services (UDS) device service. The UDS device service provides a list of devices associated with the user. After you add a CTI service on Cisco Unified Communications Manager, you must apply it to a service profile so that the client can retrieve the settings.

- **a.** From the Cisco Unified Communications Manager Administration page, navigate to **User Management > User Settings > UC Service**.
- **b.** Select **Add New**, and then in the Add a UC Service section, select **CTI** from the UC Service Type drop-down list.
- **c.** Click **Next**.
- **d.** Provide details for the instant messaging and presence service as follows:
  - Specify a name for the service in the Name field. The name you specify will be displayed when you add services to profiles. Ensure the name you specify is unique, meaningful, and easy to identify.
  - Specify the CTI service address in the Host Name/IP Address field.
  - Specify the port number for the CTI service in the Port field.
- **e.** Click **Save**.
- **f.** Navigate to **User Management > User Settings > Service Profile**.
- **g.** Find and select your service profile to access the Service Profile Configuration window.
- **h.** Scroll down to the CTI Profile section, and select up to three services from the following drop-down lists:
  - Primary
  - Secondary
  - Tertiary
- **i.** Click **Save**.

Step 2. Enable the device for CTI. If you want Cisco Jabber desktop clients to be able to control the user's deskphone, you must select the **Allow Control of Device from CTI** option when you create the device for the user.

- **a.** From the Cisco Unified Communications Manager Administration page, navigate to **Device > Phone** and click **Find** to search for the phone that is to be controlled by Cisco Jabber.
- **b.** In the Device Information section, check the box beside the Allow Control of Device from CTI setting.
- **c.** Click **Save**.

Step 3. Deskphone video capabilities let users receive video on their computers through the Jabber client for transmissions targeted toward their deskphone devices. To set up deskphone video, you must complete the following steps:

a. Physically connect the computer to the computer port on the deskphone device. You must physically connect the computer to the deskphone device through the computer port so that the client can establish a connection to the device. You cannot use deskphone video capabilities with wireless connections to deskphone devices.

b. Enable the deskphone device for video in Cisco Unified Communications Manager.

c. Install Cisco Media Services Interface on the computer. Cisco Media Services Interface provides the Cisco Discover Protocol (CDP) driver that enables the client to do the following:

- Discover the deskphone device
- Establish and maintain a connection to the deskphone device using the CAST protocol

You can download the Cisco Media Services Interface installation program from the download site at cisco.com. There are some considerations and limitations that should be taken into account before provisioning deskphone video capabilities for a user. You cannot use deskphone video capabilities on devices if video cameras are attached to the devices, such as a Cisco Unified IP Phone 8865. However, you can use deskphone video capabilities if you remove video cameras from devices that have a removable camera. Video desktop sharing, using the BFCP protocol, is not supported with deskphone video.

Legacy phones have some unique restrictions that should be considered as well. You cannot use deskphone video capabilities with devices that do not support CTI, so do your homework when using legacy phones. It is not possible for endpoints that use SCCP to receive video only. SCCP endpoints must send and receive video. Instances where SCCP endpoints do not send video result in audio-only calls. The 7900 series phones must use SCCP for deskphone video capabilities, but these 7900 series phones cannot use SIP for deskphone video capabilities. If a user initiates a call from the keypad on a deskphone device, the call starts as an audio call on the deskphone device. The client then escalates the call to video. For this reason, you cannot make video calls to devices that do not support escalation, such as H.323 endpoints. To use deskphone video capabilities with devices that do not support escalation, users should initiate calls from the Jabber client. A compatibility issue exists with Cisco Unified IP phones that use firmware version SCCP45.9-2-1S. You must upgrade your firmware to version SCCP45.9-3-1 to use deskphone video capabilities. Some antivirus or firewall applications, such as Symantec Endpoint Protection, block inbound CDP packets, which disables deskphone video capabilities. You should configure your antivirus or firewall application to allow inbound CDP packets. You must not select the Media Termination Point Required check box on the SIP trunk configuration for Cisco Unified Communications Manager. Deskphone video capabilities are not available if you select this check box.

Step 1. The client uses video rate adaptation to negotiate optimum video quality. Video rate adaptation dynamically increases or decreases video quality based

on network conditions. To use video rate adaptation, you must enable Real-time Transport Control Protocol (RTCP) on Cisco Unified Communications Manager. RTCP is enabled on software phone devices by default. However, you must enable RTCP on deskphone devices. You can enable RTCP on a common phone profile to enable video rate adaptation on all devices that use the profile.

a. From the Cisco Unified Communications Manager Administration page, navigate to **Device > Device Settings > Common Phone Profile**.

b. Specify the appropriate filters in the Find Common Phone Profile Where field and then select **Find** to retrieve a list of profiles.

c. Select the appropriate profile from the list.

d. Locate the Product Specific Configuration Layout section.

e. Select **Enabled** from the RTCP drop-down list.

f. Click **Save**.

Alternatively, you can enable RTCP on specific device configurations instead of a common phone profile. The specific device configuration overrides any settings you specify on the common phone profile.

g. Navigate to **Device > Phone**, click **Find**, and select the appropriate phone from the list.

h. Locate the Product Specific Configuration Layout section, and select **Enabled** from the RTCP drop-down list.

i. Click **Save**.

Step 2. When you associate a user with a device, you provision that device to that user. You then need to create and configure the Cisco Jabber devices as follows:

a. From the Cisco Unified Communications Manager Administration page, navigate to **User Management > End User**.

b. Click **Find** and select the appropriate end user from the list.

c. Scroll down to the Service Settings section and select **Home Cluster**.

d. Select the appropriate service profile for the user from the UC Service Profile drop-down list. This should be the service profile previously configured that contains the CTI UC Service.

e. Scroll down to the Device Information section and click on the **Device Association** button. The Device Association window will appear.

f. Select the devices to which you want to associate the user. Jabber supports only a single softphone association per device type. For example, only one CSF device can be associated with a user. However, a CSF soft phone and a Cisco Unified IP phone can be associated with the same user.

g. Click the **Save Selected/Changes** button.

h. In the top-right section of the page, click **Go** next to the Related Links: Back to User drop-down menu to get back to the end-user page.

i. Scroll down to the Permissions Information section and select the **Add to Access Control Group** section. The Find and List Access Control Groups dialog box opens.

j. Select the access control groups to which you want to assign the user. At a minimum, you should assign the user to the following access control groups:

- Standard CCM End Users
- Standard CTI Enabled

Remember, if you are provisioning users with secure phone capabilities, do not assign the users to the Standard CTI Secure Connection group. Certain phone models require additional control groups, as follows:

- For Cisco Unified IP Phone 9900, 8900, or 8800 series or DX series, select **Standard CTI Allow Control of Phones Supporting Connected Xfer and conf**.
- For Cisco Unified IP Phone 6900 series, select **Standard CTI Allow Control of Phones Supporting Rollover Mode**.

k. Click **Add Selected** and then click **Save** on the End User Configuration page.

Step 3. After you create and associate users with devices, you should reset those devices:

a. From the Cisco Unified Communications Manager Administration page, navigate to **Device > Phone**.

b. Click **Find** and select the appropriate device from the list.

c. Scroll down to the Association Information section and select the appropriate directory number configuration.

d. Click **Reset**, and the Device Reset dialog box opens. Click **Reset** again and then click **Close**.

When the Cisco Jabber client is registered, there are options available to tune the client to the environment. There are two options to tune video: one is at the operating system level, and the other is at the application level. An internal camera on a laptop can be used with Jabber. A USB camera must be connected, configured, and operational at the operating system level before it can be used with Cisco Jabber. Cisco Jabber can use multiple video peripherals, if available at the computer. From the Jabber client, navigate to **File > Options > Video > Advanced** and choose a camera to use from the Camera drop-down menu. Click the Advanced button to order the video peripherals in the priority sequence that is required for Cisco Jabber to use them. For example, a USB camera can be set up as a primary camera, and the built-in camera can be set up as a secondary camera. When Jabber is booted, if the USB camera is unavailable, the internal camera will be used instead. As soon as the USB camera is connected and detected, the camera being used by Jabber will switch to the preferred camera.

As in the case for video, Cisco Jabber can use multiple audio peripherals, if available at the computer. Navigate to **File > Options > Audio > Advanced** and choose the peripherals for audio output, audio input, and audio alerts to use from the drop-down menus. Click the

Advanced button to order the audio peripherals in the required priority sequence for audio output, input, and alerts. Figure 26-10 illustrates the camera and audio configuration options from the Cisco Jabber client.

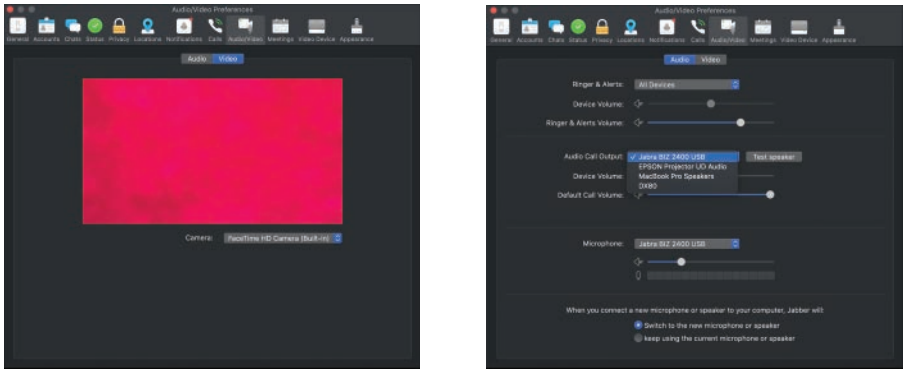

Figure 26-10 *Camera and Audio Configuration Options from Cisco Jabber Client*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 26-2 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 26-2 Key Topics for Chapter 26

Key Topic Element	Description	Page Number
List	Client Service Applications Supported by Cisco Jabber	611
List	Three Video Options for Cisco Jabber	612
List	Four Ways to Send Messages Using Hybrid Message Service	617
List	Five Services Available in Team Messaging Mode	619
Paragraph	End-User Settings Not Configurable Until After Saved	620
Steps	Configure CSF Phone in CUCM	623
Paragraph	Define CTI, CTI Port, and CTI Route Point	625

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

CAST, CCMCIP, Cisco Precision Video Engine, ClearPath Technology, CTI, CTI Port, CTI Route Point, Deskphone Mode, HTTP, HTTPS, IM, IMAP, LDAP, RTP, SIP, SOAP, Softphone Mode, SRTP, TFTP, XMPP

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the video options for Cisco Jabber.
2. What are the two main protocols used by Cisco Jabber in deskphone mode?
3. List the six soft clients that Cisco has had, which influenced the Jabber client today.

Part VII

Troubleshooting Collaboration Components

This part covers the following topics:

- **Chapter 27, Troubleshooting Endpoints:** This chapter will identify the logs available on Cisco Unified IP phones and Cisco CE software-based endpoints. This chapter will also cover common registration issues, call setup issues, and media issues related to these endpoint devices.
- **Chapter 28, Cisco Unified Communications Manager Reports:** This chapter will examine how the Dialed Number Analyzer and CAR tool can be used in the CUCM to troubleshoot issues related to registration, call setup, and media-related issues. Along with the CAR tool, this chapter will also examine how to use CDR and CMR reports to view user reports, system reports, and device reports on the Cisco Unified Communications Manager.
- **Chapter 29, Real-Time Monitoring Tool (RTMT):** The Real-Time Monitoring Tool is a complex but granular tool used by top-level engineers to troubleshoot Cisco UC and Telepresence systems. This chapter will provide a brief overview of the Real-Time Monitoring Tool and how to use this tool to monitor activity over the Cisco Unified Communications Manager.
- **Chapter 30, Understanding the Disaster Recovery System:** This chapter will explain how to create a backup of the Cisco Unified Communications Manager configuration and how to do a restore of those configuration settings in the event they need to be recovered after a disaster.
- **Chapter 31, Monitoring Voicemail in Cisco Unity Connection:** This chapter will identify how to generate reports on the Cisco Unity Connections server and through the Cisco Unified Serviceability page. This chapter will also identify how to use those reports to perform troubleshooting and maintenance on the Cisco Unity Connection server.

CHAPTER 27

Troubleshooting Endpoints

This chapter covers the following topics:

Accessing Logs on Cisco Unified IP Phones: This topic will provide a brief description of how to access logs on Cisco Unified IP phones.

Accessing Logs on CE Software-Based Endpoints: This topic will provide a detailed explanation about the different types of logs available on CE software-based endpoints and how to access these logs.

Call Signaling and Quality: This topic will explain how to use some of the tools on Cisco CE software-based endpoints to capture signaling and media information from the endpoint. This topic will also identify common issues related to registration, call setup, and media for Cisco Unified IP phones and Cisco Telepresence endpoints.

Troubleshooting Cisco Jabber: This topic will introduce some of the troubleshooting tools available on the Cisco Jabber client, as well as common registration, call setup, and media issues that Cisco Jabber can encounter.

This chapter draws off previous chapters to describe how Cisco Unified IP phones, Cisco Telepresence endpoints, and Cisco Jabber can be supported when registration, call setup, and media issues occur. Each device has some tools that can be used to troubleshoot issues, which will also be discussed in this chapter. Topics discussed in this chapter include the following:

- Accessing Logs on Cisco Unified IP Phones
- Accessing Logs on CE Software-Based Endpoints
  - Log Collection
  - Log Bundle
- Call Signaling and Quality
  - Signaling and Media Detailed Capture
  - Common Registration Issues
  - Common Call Setup Issues
  - Common Media Issues
- Troubleshooting Cisco Jabber

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.4 Troubleshoot these network components in a Cisco Collaboration solution
  - 1.4.a DNS (A/AAA, SRV, Reverse Pointer Record [PTR])
  - 1.4.b NTP
- 2.1 Troubleshoot these elements of a SIP conversation
  - 2.1.a Call set up and tear down
  - 2.1.b SDP
- 2.5 Troubleshoot collaboration endpoints

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 27-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 27-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Accessing Logs on Cisco Unified IP Phones	1
Accessing Logs on CE Software-Based Endpoints	2–5
Call Signaling and Quality	6–10
Troubleshooting Cisco Jabber	11

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What must an administrator do before accessing logs on the Cisco Unified IP phone through the web interface?
   - **a.** Nothing, the web interface is always available to the administrator.
   - **b.** Set the username and password for the phone on the phone.
   - **c.** Set the username and password for the phone on the CUCM.
   - **d.** Enable access to the web interface of the phone on the CUCM.

2. What menu location can an administrator use to access log files on a CE software-based endpoint web interface?
   a. **Maintenance > System Logs**
   b. **Status > Logs > Event Logs**
   c. **Maintenance > Current Logs**
   d. **Status > Current Logs**
3. On a CE software-based endpoint, what is the difference between Start Extended Logging and Include a Full Packet Capture?
   a. Start Extended Logging includes a packet capture with signaling only, but Include a Full Packet Capture includes a packet capture with signaling and media.
   b. Start Extended Logging includes a full debug level with no packet capture, but Include a Full Packet Capture includes a packet capture with signaling and media.
   c. Start Extended Logging includes a full debug level with signaling packet capture, but Include a Full Packet Capture includes a packet capture with signaling and media.
   d. Start Extended Logging includes a full debug level with no packet capture, but Include a Full Packet Capture includes a packet capture with signaling only.
4. Which of the following log files contain general information on all previous calls made by the endpoint?
   a. Status.txt
   b. Journal.log
   c. Call_history.txt
   d. Latest-provisioning.log
5. How many historical log files are stored in the log bundle when it is downloaded?
   a. 1
   b. 2
   c. 5
   d. 11
6. Which of the following statements is true about the following logs available on CE software-based endpoints?
   a. Call information from the Call Control page is available after the call ends.
   b. Call information from the Status page is available during a call or after the call ends.
   c. Call information from the Call Logs page is available after the call ends or if a call attempt fails to connect.
   d. All of these answers are correct.

7. When you are reading a detailed debug capture of a SIP call, what indictor marks the beginning of the SDP media capabilities exchange?
   a. M=
   b. V=
   c. CSeq: SDP
   d. A=
8. When an administrator suspects a Webex Room Kit endpoint is not registering to the CUCM due to a network reachability error, which of the following commands can be used to ping the CUCM from the CLI of the endpoint?
   a. **Systemtools network ping**
   b. **Xcommand network ping**
   c. **Xconfiguration network ping**
   d. **xstatus network ping**
9. When a user tries to place a call and only dead air is heard, what is the probable cause of this issue?
   a. Misconfigured dial plan
   b. CAC
   c. IP reachability issue to the CUCM
   d. Firewall or ACL blocking media
10. A user has called in a complaint that the audio and video were out of sync while in a video call from a SIP endpoint. What can an engineer do to prevent this issue from happening again?
    a. Reprovision the available bandwidth for that endpoint.
    b. Nothing; lip synchronization is not supported over SIP natively.
    c. Change the codec used in the CUCM Regions settings.
    d. Change the CAC settings in the CUCM.
11. Which of the following is a TCP port that Cisco Jabber uses for XMPP communication?
    a. 8443
    b. 5222
    c. 5060
    d. 2748

Foundation Topics

Accessing Logs on Cisco Unified IP Phones

Key Topic

Cisco Unified IP phones can display status messages that show the most recent events from the phone on the screen display itself. These messages can indicate settings such as if a DHCP server can be reached and the address that is assigned to the phone. The status messages are available on all types of Cisco Unified IP phones. To check the status messages on

a Cisco Unified IP Phone 8800 series, you can navigate to **Applications > Administrator Settings > Status**. Three types of status messages are available: Status Messages, Network Statistics, and Call Statistics. The Call Statistics menu allows you to check the Average (Avg) and Max Jitter, Receiver (Rcvr) Lost Packets, and Latency to verify that the call meets the general QoS requirements for the video traffic type. Beyond these status menus, no more logs are available from these phones on the endpoint itself. However, many more logs are available through the web interface of these phones.

Before you can access the logs available from the web interface of the Cisco Unified IP Phones 7800 series and 8800 series, you must configure some settings on the Cisco Unified Communications Manager first. From the Cisco Unified Communications Manager Administration page, navigate to **Device > Phone**. Select a Cisco Unified IP phone from the list by clicking on the MAC address, and then scroll down to the Web Access setting. Using the drop-down menu, change this setting from Disabled, which is the default, to Enabled. Then click Save and Apply Config from the buttons across the top of the screen. After the phone has been reset, you should be able to log in to the phone through the web interface. Figure 27-1 illustrates how to configure all these settings.

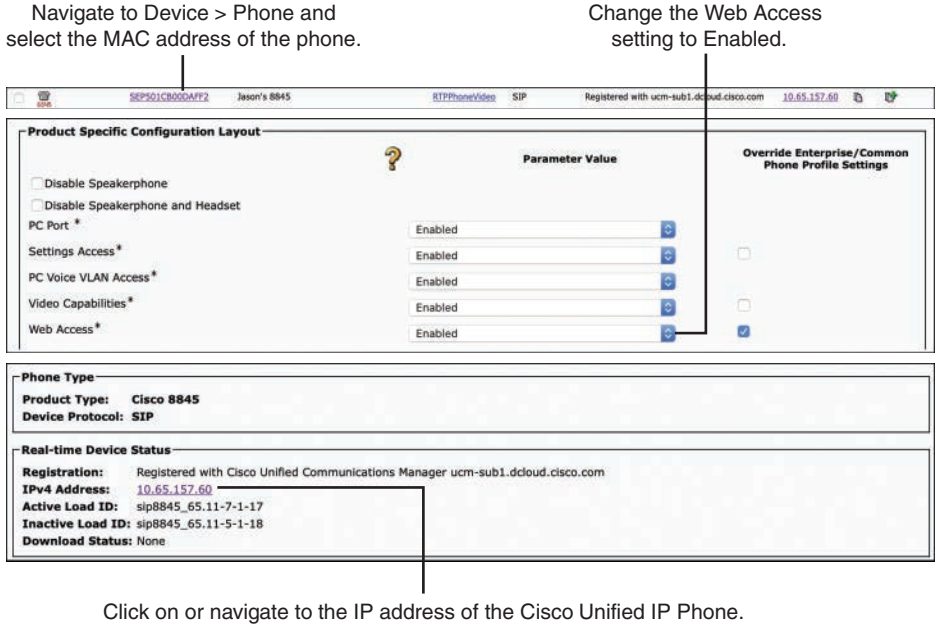

Figure 27-1 *Configure CUCM for IP Phone Web Interface Access*

In a new web browser tab, navigate to the IP address of the Cisco Unified IP phone. This IP address can be identified on the Cisco Unified IP phone or on the endpoint itself by navigating to **Application > Phone Information**. You do not need to use HTTPS, and you do not need to enter a username or password. No settings on the phone can be changed from the web interface. Only logs can be accessed from the web interface, but a lot more logs are available to the administrator. Figure 27-2 illustrates the web interface and logs available on a Cisco Unified IP Phone 8845.

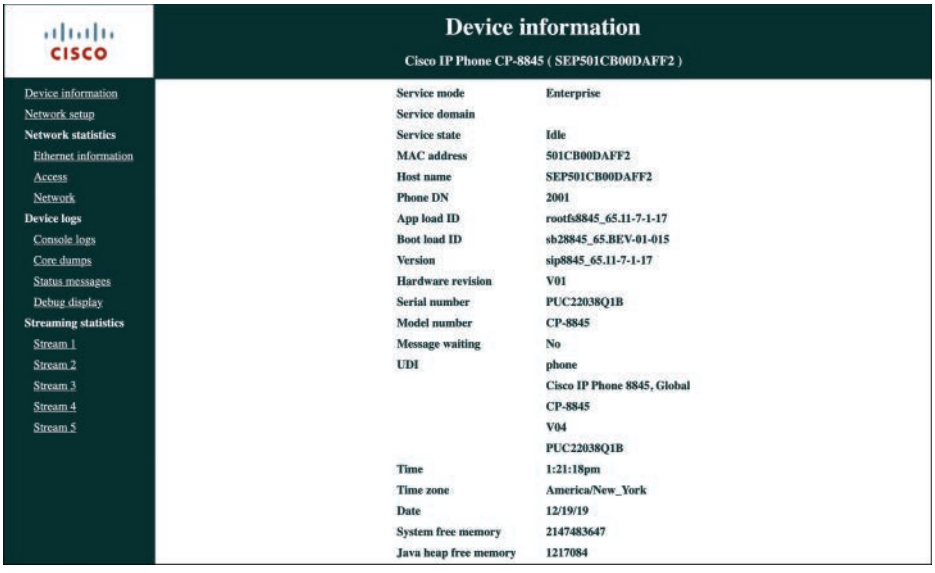

Figure 27-2 *Cisco Unified IP Phone 8845 Web Interface Logs*

Accessing Logs on CE Software-Based Endpoints

Many different types of logs are available on Cisco Telepresence endpoints running CE software. Historically, administrators had to access these logs from the CLI, and they often had to use other tools to extract these logs, such as the Microsoft WinSCP tool. Cisco has made many enhancements to the web interface of these endpoints, making access to these logs much easier. You can use the following steps to access the logs from the web interface of an endpoint.

Step 1. Open a web browser and enter the IP address of the endpoint in the address bar.

Step 2. From the login screen, enter the username and password assigned to the endpoint. By default, CE software-based endpoints use **admin** as the username and the password field is left blank.

Step 3. Once logged in, navigate to **Maintenance > System Logs**.

The System Logs page is divided into three main sections.

- Download Log Archive
- Start/Stop Extended Logging
- Manually Browse Current Logs or Historic Logs

Log Collection

Extended logging enables additional debugging including SIP tracing. Historically, this debug level had to be enabled from the CLI and then retrieved from the eventlog.application.log file. Now Cisco has added the ability to turn on this debug trace directly from the web

interface. When Start Extended Logging is selected, the extended logging is turned on and will last for 10 minutes. Two other options exist for enabling extended logging. The Include a Limited Packet Capture option will not only turn on the extended logging but will also capture log information in a separate pcap file. The "limited" component signifies that media information will not be captured, only signaling data. This option will also last for 10 minutes when enabled. If media information is needed because of media issues, you should select the Include a Full Packet Capture option. This option will collect all the data previously noted; plus, it will also capture all the RTP traffic. Because this type of trace will include much more information, this option will last only 3 minutes. Figure 27-3 illustrates the menu options on the CE endpoint web interface used to enable these extended logging options.

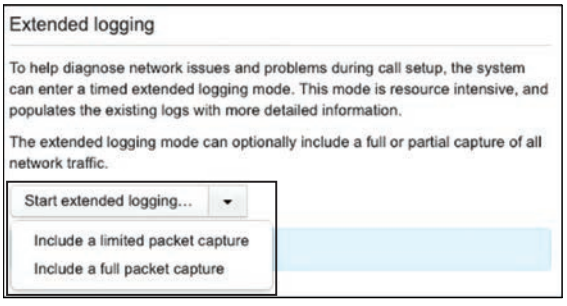

- Extended Logging
  - Enables additional debugs, including SIP tracing
  - Lasts 10 minutes
- Include a Limited Packet Capture
  - Starts pcap, which will filter out RTP media
  - Lasts 10 minutes
- Include a Full Packet Capture
  - Captures all traffic, including RTP
  - Lasts 3 minutes

Figure 27-3 *Extended Logging Options from CE Endpoint Web Interface*

When one of the options from Figure 27-3 is selected, an indicator flag will notify the administrator that extended logging is active and will indicate the total time allotment. Also, the packet capture file will be displayed in the GUI. This file can be downloaded directly by clicking the hyperlink, or it can be included in the log bundle when downloaded.

Much like with enabling extended logging, three options are available when downloading the logs archive:

Key Topic

- **Download Logs Archive:** Selecting this option will enable a standard download with most details still included for analysis. Call info is excluded from this standard download to ensure the caller's privacy.
- **No Call History:** Selecting this option will download the same log information as Download Logs Archive, except no "history" log information will be included to save on space.
- **Full Call History:** Selecting this option will download the exact same information as the standard download, including caller information. Only three lines are different between this option and the standard option.

Figure 27-4 illustrates the three log archives that can be downloaded and the three lines that distinguish the standard logs archive from the full call history.

Download Options for Log Archive

- Download logs archive: Standard download with most details still included for analysis (excludes caller info for privacy)
- No call history: No history is included to save on space
- Full call history: Full call history is included with caller info

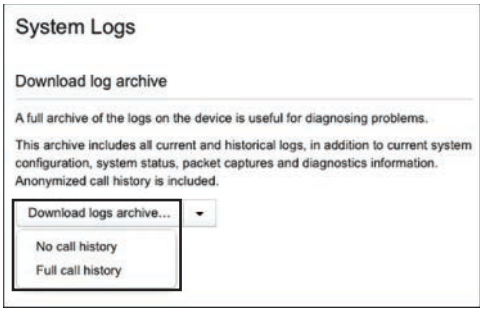

Standard vs. Full Call History

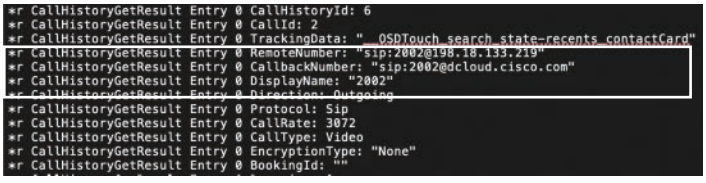

Figure 27-4 *System Logs Download Options from CE Endpoint Web Interface*

27

Log Bundle

When the log archive files have been downloaded, they will be zipped in a tar.gz file. After that file is uncompressed and opened, two or three other folders will be included. There will always be a *current* and an *old* folder. If extended logging was enabled with a packet capture before the log archive was downloaded, there will be a third folder that includes the pcap file that was created. The current folder contains all log files and information from the current boot of the system. The old folder contains historical log bundles created at shutdown. When the current folder is opened, the following files are available. This is not an exhaustive list; rather, it is a list of the most likely logs an administrator would use, with a description of each file.

Key Topic

- **Call_history.txt:** Contains general info on all previous calls made by the system
  - Protocol
  - Negotiated call rate
  - Start/end times
  - Disconnect information
  - Media statistics
  - Direction
- **Configuration.txt:** Includes all system configuration settings set through the web GUI and CLI
  - Same as running xConfiguration in the CLI
  - Good for quick reference of settings without needing to access system directly or if troubleshooting issue from some time ago
  - Always pulled at the time of log collection

- **Status.txt:** Includes all system status outputs
  - Same as running xStatus in the CLI
  - Useful as a quick health check for processes such as provisioning or peripheral connections
  - Always pulled at the time of log collection
- **Peripherals.txt:** Contains information on all connected peripherals
  - Same information included in status.txt
  - Good for quick reference without having to sort through other noise
- **Journal.log (also known as messages.log):** Contains low-level system information and boot processes
  - Also includes kernel messages from kernel.log
  - Useful for troubleshooting system crashes
- **Latest-provisioning.log:** Present if the system is provisioned by CUCM
  - Contains the most recently pushed cnf.xml file
  - Useful to validate configuration from CUCM without needing access to CUCM
- **Latest-valid-provisioning.log** contains config after parsing is complete

The eventlog folder, also included in the current folder, contains a directory of log files that offer more verbose information. The logging capacity on these endpoints is finite, but the system will take measures to ensure the most relevant information is stored. Eventlog files will use up to four file extensions before older data is lost. Figure 27-5 illustrates how these four file extensions are created within the CE software-based endpoint event log. The boxes in the figure illustrate how these file extensions work. Note that **x** in these examples indicates any eventlog file. The description that follows uses application.log as a point of reference and works from the top left toward the bottom right.

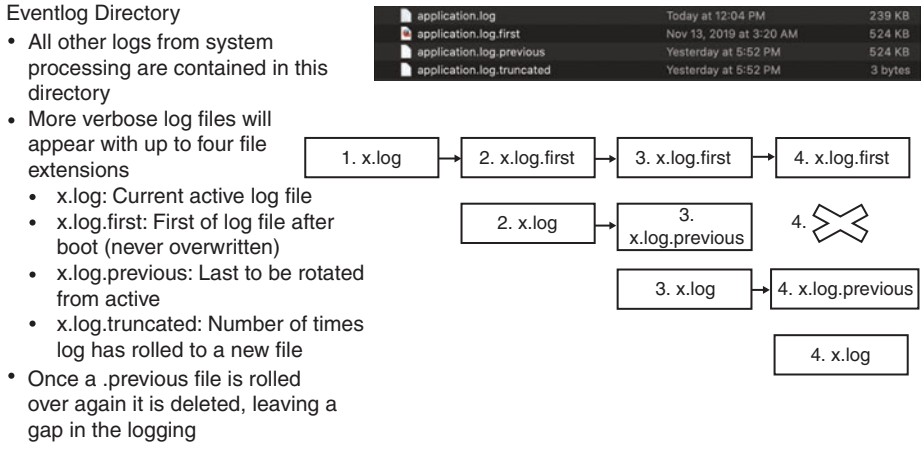

Figure 27-5 *Event Log File Extension Operation on CE Software-Based Endpoints*

1. The first x.log box represents an event log that is created when the endpoint first boots up, such as the application.log.
2. Once this log fills up, the data is saved under a new filename, such as application.log.first. The data in log.first logs will never be overwritten or deleted. Once this original content has been renamed, new logged information can be captured in the application.log file (second row).
3. When this second application.log file fills up, it too needs to be renamed, so application.log.previous will be used. Once this second set of content has been renamed, new logged information can be captured in the application.log file (third row). The log.first file does not change from when it was originally created.
4. When this third application.log file fills up, it too needs to be renamed, and application.log.previous will be used again. Two files with the same name cannot exist, so the first application.log.previous file is deleted, creating a gap in the logging. Once this third set of content has been renamed, new logged information can be captured in the application.log file (fourth row).

You may have noticed from Figure 27-5 that there is another file extension called x.log.truncated. This is not a renamed log file as previously described. This file is a logged reference point that tracks how many times an **x** log has rolled over since the last reboot of the endpoint. Many logs are included in the Event Log folder, most of which will never be referenced by a system administrator. The following is a list of the most commonly used event logs, with a description of the information found in these logs:

Key Topic

- **All.log:**
  - All logs rolled into a single file
  - Quick glance at what's happening
  - Overwrites quickly
  - Specific log files referenced by tagging in all.log.
- **Application.log:**
  - Includes application-level debugging (including additional debugs enabled)
  - Includes Call setup and signaling
  - Includes Registration information
  - Includes Provisioning information
  - Includes Phonebook searches
  - Includes Configuration changes
- **Main.log:**
  - Process and system monitoring
  - Boot information, software version, option keys
  - Place to look for boot issues or crashes

- **Audit.log:**
  - Logs commands from all connected devices
  - Includes physically connected devices, such as touch and network connected
  - Includes any configuration changes from users or management devices

Historical bundles are similar to the current directory but contain slightly less diagnostic info. All of the same event logs are included in each historical log. Each time the system reboots, all the logs in the current log files are compressed into a single historical log file. The log.tar.gz is always a reference to the most recently created historical log. After the next reboot, the log.tar.gz file is renamed with a numeric value, such as log.x.tar.gz, where x is a number value (0–9). When all historic log slots are full, the oldest historic log will be deleted. This deletion is based on the time-stamp, not on the numeric value associated with the name. Although up to 11 historical log files can exist in the Cisco CE software-based endpoint, the log bundle will only contain the previous five historical log files. All historical log files are cleared on a factory reset.

Call Signaling and Quality

The web interface of Cisco Telepresence CE software-based endpoints provides three basic places to check for call info:

Key Topic

- Call Control Page
- Status
- Call Logs

These quick reference locations are perfect for a quick understanding of call negotiation or basic media statistics, or to augment additional troubleshooting. The Call Control and Status pages will display call statistic information only while a call is active. However, Call Logs pages are available for previous calls or call attempts.

Select the Call Control menu from the top of the screen to view the Call Control settings. If there is an active call, you will see the connected party's alias listed in the Participants section with three buttons beside the listing. The last two are used to place the call on hold or hang up the call. The first button, the information button, is an *i* with a circle around it. Figure 27-6 illustrates how the Call Control Details page will appear.

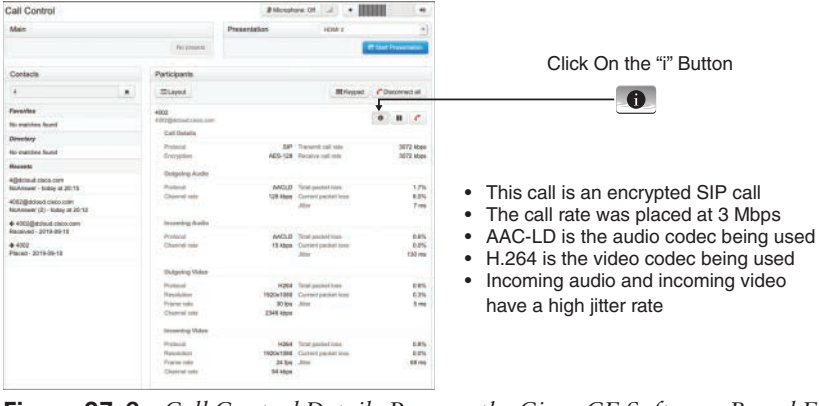

Figure 27-6 *Call Control Details Page on the Cisco CE Software-Based Endpoint*

The information button is used to display the call details. Information that can be observed from the call in Figure 27-6 includes that this call is an encrypted SIP call and that the call rate was placed at 3 Mbps. Also, AAC-LD is the audio codec being used, H.264 is the video codec being used, and incoming audio and incoming video have a high jitter rate.

From the menus at the top of the page, navigate to **Setup > Status**. A list of menus will appear in the column to the left. Select MediaChannels from the menu. This will display status information about any current calls that are connected at the immediate time. Status provides useful information, such as the UDP media ports used by both source and destination endpoints.

The call logs offer information after a call ends or for calls that never connected. Next, navigate to **Maintenance > Call Logs**. A list of all calls or call attempts since the last reboot will be listed in this log. From this page, you can access some basic information, such as the direction of the call, how long the call was connected, and the reason for the call ending. Select any of these call listings to view more detailed information. The following valuable information can be observed from this log:

Key Topic

- Remote number
- Call direction, protocol, and call rate
- Disconnect cause
- Disconnect cause code
- Media statistics averages (packet loss/jitter)

Signaling and Media Detail Capture

This chapter has already discussed how to enable extended logging, how to perform a packet capture, and how to pull this information from the endpoints. This section will introduce how to read the log information after it is captured. Debug log information is accessed in the eventlog.application.log file. There will be a lot of information to drill through to find what you need. The easiest way to find the call debug information is to use the search tool. The actual SDP information, which is available only in a debug, always begins with v=0. You can use this variable to find the specific call information you are looking for. Example 27-1 illustrates how some of the information captured in a debug will appear in the application.log file.

Example 27-1 *Debug Log from eventlog.application.log After Extended Logging Was Enabled and a Call Was Placed*

```
SipPacket I: SIP Msg: Outgoing => INVITE, CSeq: 100 INVITE, Remote:
14.49.23.20:5060, CallId: 88b866e5f94fdfb3e0214208bd6a34fe, SessionId: 0eaf314a5f335
bdbb6823e427b119680;remote=00000000000000000000000000000000
SipPacket[2]: INVITE sip:20001@tkratzke.local SIP/2.0
SipPacket[2]: Via: SIP/2.0/TCP 14.0.70.171:39293;branch=xxxx;rport
SipPacket[2]: Call-ID: 88b866e5f94fdfb3e0214208bd6a34fe
SipPacket[2]: CSeq: 100 INVITE...
SipPacket[2]: v=0
SipPacket[2]: m=video 21008 RTP/AVP 99 97 126 123
SipPacket[2]: b=TIAS:6000000
SipPacket[2]: a=rtpmap:99 H265/90000
```

```
SipPacket[2]: a=fmtp:99 level-id=90;max-lsr=125337600;max-lps=2088960;max-tr=22;max-
tc=20;max-fps=6000;x-cisco-hevc=529
SipPacket[2]: a=rtpmap:97 H264/90000
SipPacket[2]: a=fmtp:97 packetization-mode=0;profile-level-id=428016;max-
br=5000;max-mbps=490000;max-fs=8160;max-smbps=490000;max-fps=6000
! Output truncated for brevity
```

Key Topic

In Example 27-1, the shaded header information is all the information you will see in the application log file about each call when no debug is enabled. All the information below the header, which has been truncated for obvious reasons, is shown only when a debug is enabled. The m= represents the beginning of the SDP media capabilities exchange based on what codecs the endpoint supports. Audio will always be displayed first, followed by primary video, which is followed by content sharing. The example has omitted the audio SDP negotiation. Note that 21008 is the RTP port this endpoint wants to use to send the media, which is video in this case. A different port would be used for audio or content sharing, and different ports are used for media going in the other direction. The 99 97 126 123 numbers are SIP cause codes for codecs that this endpoint supports. You will always find a description for each number in the a= lines below, so you don't need to memorize what each one of these is. The preferred codec is always the first number in the list, and each number following is in sequential order of preference. Notice that 99 represents H265/90000 and 97 represents H264/90000. The 90000 is just the Hertz rate. The next line below each of these identifies the bandwidth rate requested with each codec.

Packet captures are tagged in a file called pcap. You will need special tools, such as Wireshark, to open and read this file. Reading a pcap file is not an intuitive process. You must know what you are looking for in order for the information to make any sense. Each line in a pcap file is a different packet that was sent. When one of these packet lines is selected, a new window is opened, usually at the bottom of the page, showing specific header information about the packet selected. Each of these new lines of information can be expanded so that engineers can drill down into the different OSI layers of information sent.

Because a pcap file contains a lot of information, an easy way to filter the information is to locate a distinguishing attribute you are trying to investigate, such as the UDP port used for video communications, and then filter the packets based on that port. There is another, and perhaps easier, way to filter information using Wireshark. Navigate to **Telephony > RTP > RTP Streams**. This tool will filter out the data for you and provide a summary of all the detected RTP streams in the capture. It will provide the source and destination address, identify the payload type, and identify packet loss and jitter values. This information might help in isolating where along the connection path a loss of data was detected.

Now that you know how to read these detailed log files, the following explanation will illustrate how these logs can be used to troubleshoot a call media issue that the Cisco Technical Assistance Center (TAC) gets lots of calls about. Figure 27-7 illustrates the logs used to troubleshoot this issue.

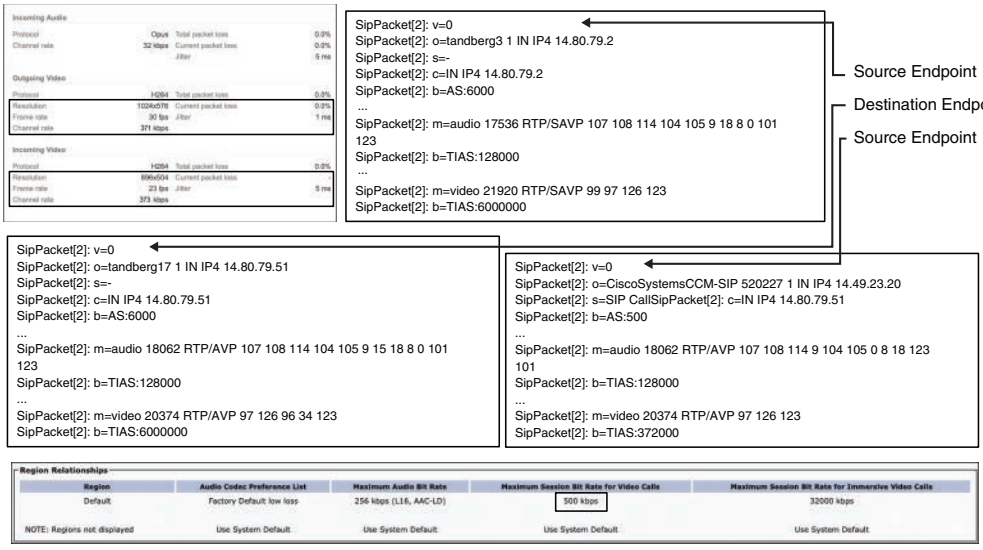

Figure 27-7 *Call Information Details Log and SIP Debug Logs Used for Troubleshooting Media Issue*

Imagine two endpoints in a basic call through the Cisco Unified Communications Manager. When they call each other, the audio quality is good; however, their video is pixelated and blurry. At first investigation, the Call Information Details from the Call Control menu on the endpoint reveal that video bandwidth is not getting the quantity requested. In Figure 27-7, the top-left screenshot reveals that only 371 kbps bandwidth has been allocated for outgoing video, and only 373 kbps has been allocated for incoming video. When extended logging is enabled, the SIP call setup log initially shows the source endpoint that initiated the call is requesting 6 Mbps for this call for video. The next part of the log, represented in the middle-left log, shows that the destination endpoint can also support 6 Mbps in this call, so the issue doesn't seem to be with the destination endpoint.

Further investigation in the log file finally reveals the issue, represented in the middle-right log. The source endpoint is actually requesting 500 kbps total bandwidth for this call, divided between audio at 128 kbps and video at 372 kbps. Something obviously renegotiated the bandwidth for the call. Since both endpoints are registered to the Cisco Unified Communications Manager, a quick check reveals that only 500 kbps has been provisioned for video calls, which is causing poor video quality during the call.

Common Registration Issues

Chapter 9, "Endpoint Registration," covered the registration process to the Cisco Unified Communications Manager. Understanding this registration process is important because any of these registration components could cause registration to fail. As a review, Cisco Unified IP phones and Cisco Telepresence endpoints require these elements to register successfully with Cisco Unified Communications Manager:

1. The phone must be on a correct voice and video VLAN that has network connectivity to Cisco Unified Communications Manager.
2. The phone must have a correct IP address, network mask, and default router, which the DHCP server can assign, or you can configure manually at the phone.

3. The phone must have a correct TFTP server address from which to download the configuration. The DHCP server assigns the option 150 TFTP IP address, or you can configure it manually at the phone.
4. There must be IP connectivity to the Cisco TFTP and Cisco Unified Communications Manager servers.
5. The phone must be able to exchange SCCP or SIP messages with Cisco Unified Communications Manager servers, without any filters applied.
6. The phone must be a known device to Cisco Unified Communications Manager. The phone must also be configured manually with the correct MAC address or must be able to use autoprovisioning.
7. If security is implemented in the Cisco Unified Communications network, the phone must have correct security elements applied. This configuration is performed on Cisco Unified Communications Manager, but the phone can cache old invalid information, such as trust lists.

Several factors can cause a "Phone not registered" message on a Cisco Unified IP phone. If a phone does not power up, the cause might be disabled PoE at the switch port or a switch that uses an incompatible PoE mechanism. Cisco switches support two types of PoE: Cisco proprietary and standards-based IEEE 802.3X. Cisco Unified IP phones support only the standards-based option. The IEEE standards for PoE are 802.3af and 802.3at. Some Cisco phones support the latter, which could cause issues. If a user were to connect an 802.3at device to an 802.3af supported switch, certain features on that phone would not work, which could include the phone not having enough power to fully boot up. The result is a cyclical reboot, where the phone begins to power on but shuts down during the boot process and starts the boot sequence all over again. Although PoE-related issues are not the most common issues found in a Cisco Collaboration network, they can occur. Figure 27-8 illustrates the more common issues that administrators will come across in a Cisco Collaboration network causing registration failures.

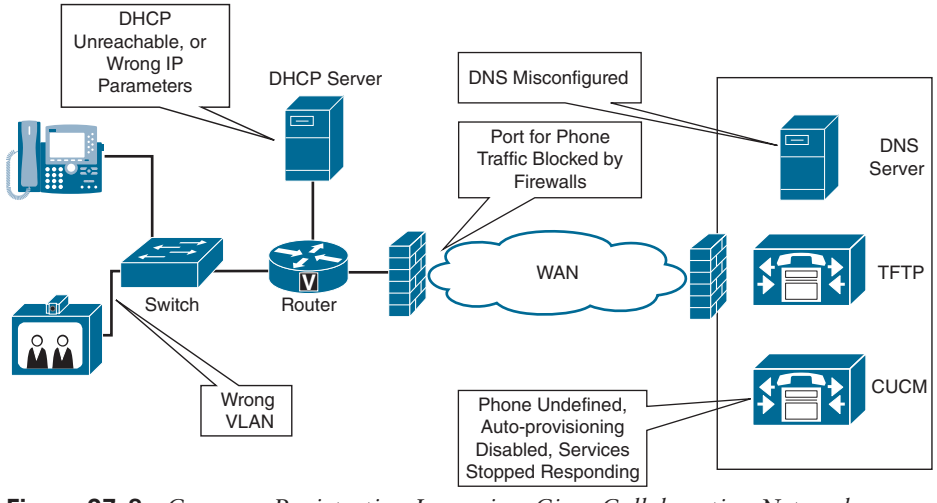

Figure 27-8 *Common Registration Issues in a Cisco Collaboration Network*

Key Topic

The phone might get an incorrect VLAN from the switch, so call control servers might be unreachable. You can configure each port on a Cisco Catalyst switch to which the Cisco Unified IP phones connect with two independent VLANs. Although one of the VLANs is called a Voice VLAN for historical reasons, it is also used to carry video traffic between the phone and the switch. This auxiliary VLAN carries tagged traffic using an IEEE 802.1Q header with the configured VVID. The Data VLAN stretches from the switch through the phone to the computer, which connects at the PC port of the phone. This VLAN is untagged. Although two VLANs are the recommended setup, the switch does not need to use separate VLANs for voice and video traffic and data traffic. One VLAN can be used for everything. The biggest advantage of using two VLANs is that the phone separates traffic types that have very different QoS requirements. Two VLANs also perform traffic classification that can be reused later downstream.

DHCP servers can be unreachable because a router is placed between the phone and the DHCP server path. Remember that DHCP requests are broadcasts that spread out on the local IP subnet only. The router must be configured to relay these broadcasts on the IP subnet on which the DHCP server connects. The DHCP server can run out of IP addresses and therefore has none to assign. The DHCP server can also be misconfigured and assign incorrect IP parameters. A Cisco Unified IP phone can obtain IP parameters from three general types of DHCP servers:

Key Topic

- DHCP server running as a service at Cisco Unified Communications Manager
- DHCP server that is implemented at a Cisco IOS router
- Third-party DHCP server—for example, a Microsoft Windows server

If the phone has incorrect IP parameters, you can erase them and wait until the phone obtains new parameters from the DHCP server. To erase network settings on the Cisco Unified IP Phone 7800 and 8800 series, choose **Applications > Administrator Settings > Reset Settings > Network Settings** and confirm that the network settings have been erased. If IP parameters are not being received from the DHCP server, consider network-connectivity problems. Verify Layer 2 connectivity to the switch first. Another potential issue might be that the DHCP server has assigned all its available addresses to other endpoints and does not have an address to assign. The DHCP server might use an incorrect configuration and assign incorrect IP parameters to the phone. For example, an incorrect default router or network mask can cause incorrect routing outbound from the phone. Also remember that option 150 must be configured on the DHCP server for the phones to obtain configuration settings from the TFTP server. Example 27-2 shows a typical router-based DHCP server configuration.

Example 27-2 *Sample Router-Based DHCP Server Configuration*

```
ip dhcp pool video-phones
   network 10.250.20.0 255.255.255.0
    dns-server 10.192.126.10 10.192.126.11
    default-router 10.250.20.2
    domain-name cisco.com
    option 150 ip 10.192.5.97
```

After the phone obtains correct IP parameters, it sends registration packets that follow the IP routed path. If the route does not exist in either direction, the traffic does not reach the servers that are required to complete the registration process. The Cisco Unified IP phone requires IP connectivity to the Cisco TFTP server and Cisco Unified Communications Manager. In small deployments, these two servers can be collocated on the same physical platform, but for larger deployments, they need to be separate for scalability reasons. You can verify IP connectivity by using two options:

- Ping the servers from a switch, by using an extended ping with the source IP address from the voice and video VLAN. This VLAN is the same VLAN that the phones use to communicate with the servers.
- Ping the servers from a computer in the voice and video VLAN that receives IP parameters from the DHCP server for the same IP subnet as the phone. If using this option, ensure that you do not connect the computer to the PC port of the phone because from there it might be assigned to a completely different VLAN.

If **ping** fails, use **traceroute** to verify the routed path. Be aware, however, that company security policies might intentionally block ping and traceroute, whereas other traffic types can pass. You can also verify IP routing tables at the downstream routers. Do not forget about the opposite direction. If ping succeeds, there might still be issues at firewalls. Firewalls can let the ping pass but block the ports that are required for successful registration, such as TFTP, SCCP, SIP, or instantiated RTP/RTCP ports.

If you suspect that broken IP connectivity to the Cisco Unified Communications Manager is causing the registration issue from a Cisco Telepresence endpoint, you can verify the IP connectivity directly from the CE software-based endpoints. Log in to the CLI as admin and display the network settings that the endpoint uses. Verify that the IP settings and VLAN are correct. Try to ping the call control server from the CLI by using the **systemtools network ping** *IP_address* command. If the ping shows lost packets, you can verify the IP routed path by using the **systemtools network traceroute** *IP_address* command. You can also use the **systemtools network netstat** command to verify that the network processes are running and listening on the correct protocol ports.

DNS, if used, can also contribute to failed communication with the call control servers resulting in a failed registration attempt. DNS can be unavailable or unreachable in the IP network. DNS can also be misconfigured, meaning that it resolves names to incorrect IP addresses. If a DNS server is used to address Cisco Unified Communications servers, you should verify that the DNS server resolves names to correct IP addresses. To verify name resolution, use the **ns lookup** command followed by the DNS name of the Cisco Unified Communications Manager and Cisco TFTP server, if a separate TFTP exists, instead of the IP address. Remember that this test must be performed from the same VLAN that is used for voice and video devices. Otherwise, the verification procedure does not check IP connectivity to the DNS server. If names do not resolve at all, ping the DNS server by its IP address. If no response is received, troubleshoot IP connectivity to the DNS server. If names resolve to incorrect IP addresses, check and correct the DNS database. If DNS names resolve to correct IP addresses but no response to the ping is received, troubleshoot IP connectivity to Cisco Unified Communications Manager and the Cisco TFTP server.

Firewalls might block the types of network protocols that are required during the registration process. Cisco Unified IP phones use TFTP to download configuration and use SCCP or SIP to complete the registration process. The firewalls in the path must open ports for these protocols:

Key Topic

- TFTP uses UDP port 69 and then uses ephemeral ports for the actual file download.
- SCCP uses TCP port 2000. Secure SCCP uses TLS ports 2443 and 2445.
- SIP uses TCP or UDP port 5060. The port that is used depends on how SIP is provisioned on Cisco Unified Communications Manager. Secure SIP uses TLS port 5061.

If no network connectivity issues exist but the Cisco Unified IP phone still cannot register, settings on the Cisco Unified Communications Manager could prevent registration from occurring. Administrators should verify that endpoint configuration exists and that the correct MAC address is configured in Cisco Unified Communications Manager. If autoprovisioning is used instead of manual configuration, verify that autoprovisioning is enabled and that enough directory numbers are left in the pool. Verify that the required services are running on the Cisco Unified Communications Manager and Cisco TFTP server. Also, verify that Cisco Unified Communications Manager has not reached the limit of maximum registered phones because of licensing or configured limitations.

The Cisco Unified IP phone and Telepresence endpoint supports digital certificates, device authentication, and encryption. Phones can use the IEEE 802.1X data link layer authentication protocol to connect to a Cisco Catalyst switch that implements this security mechanism. Phones that support 802.1X must be configured correctly; otherwise, the data link layer authentication will prevent the phone from registering. Phones that do not support 802.1X must be provided with some other means to access the voice network. Cisco Unified IP phones and Telepresence endpoints also provide security by default with these automatic security features:

- Signing of the phone configuration files
- Support for phone-configuration file encryption
- HTTPS with Cisco Tomcat and other web services (MIDlets)

These security features are provided by default, even without implemented security for signaling and media and without running the Cisco CTL client. If the Cisco Unified IP phone does not have an existing CTL file, it trusts the first ITL file automatically, as it trusts the CTL file. Subsequent ITL files must be signed by the same TFTP private key, or the Trust Verification Service must be able to return the certificate that corresponds to the signer. The ITL file contains the ITL. The ITL file has the same format as the CTL file and is basically a smaller, leaner version of the CTL file. These attributes apply to the ITL file:

- Unlike the CTL file, the system builds the ITL file automatically when you install the cluster, and the ITL file is updated automatically if the contents need to be changed.
- The ITL file does not require eTokens. It uses a soft eToken (the TFTP private key).
- The Cisco Unified IP phones download the ITL file at bootup time or during reset, just after downloading the CTL file (if present).

Key Topic

The CTL and ITL files can prevent a Cisco Unified IP phone from registering if the phone was moved to a different Cisco Unified Communications Manager cluster and the files that are stored at the phone still refer to the original cluster of the phone. In these circumstances, you must manually erase the CTL and ITL files from the phone to let the phone build or download new files as it boots.

Common Call Setup Issues

The components that are involved in call setup depend on the chosen deployment model. Cisco Unified Communications Manager can be deployed in single-site or multisite models, or as a single cluster or multiple clusters. The components also depend on whether the network interconnects with any other networks, such as video-capable ISDN (H.320) or traditional H.323 systems. Generally, two traffic types are used for each call:

- Call signaling traffic
- Call bearer traffic for audio and video

In most deployments, the components or mechanisms involved in call setup include dial plans, digit manipulation, and call privilege policies, which are implemented at the call control systems such as Cisco Unified Communications Manager and gateways. If any of these components are misconfigured, a call might fail to set up. Figure 27-9 illustrates the components involved with call setup in a Cisco Collaboration solution.

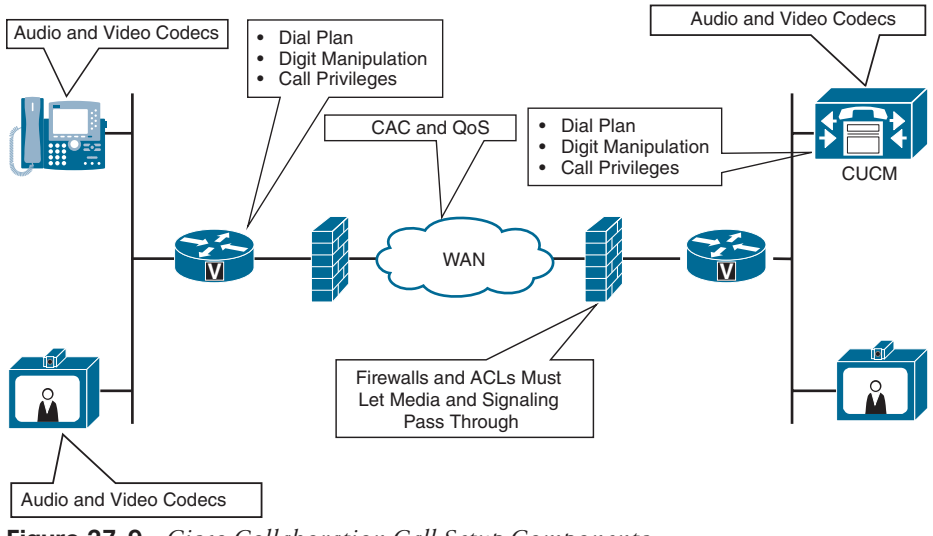

Figure 27-9 *Cisco Collaboration Call Setup Components*

If firewalls or access lists are implemented along the call path, they must allow call signaling and audio and video media to pass through to them for a call to set up. Media port numbers are dynamically negotiated during the call setup. Therefore, firewall adjustments must be implemented to track these negotiations and open proper ports through which the media will flow.

Key Topic

Some additional components can affect the call setup process. Audio and video codecs are implemented at endpoints and gateways. During the setup of a video call, the endpoints negotiate the codec that is used to encode and decode audio and video traffic. If two endpoints cannot agree on a common codec, the call setup will fail. Call Admission Control

(CAC), if implemented, can block the call as it is being set up because of the existence of too many calls between two sites across the WAN. CAC is used to guard network resources from oversubscription, which can affect the quality of existing calls. Gateways might implement media transcoding functions when passing the media from one network to another. Gateways require DSP resources to perform these processing-intensive functions. Starting with Cisco IOS Release 15.0(1)M, if a DSP is required but is not available, when the phone goes off-hook, the caller receives a fast-busy tone. Prior to Cisco IOS Release 15.0(1)M, the caller would hear dead air.

When troubleshooting call setup issues, you must make correct judgments concerning the potential cause at the beginning of the troubleshooting procedure. Many distinct components are involved in the call. Many common call-setup issues relate solely to call control components. Based on the type of issue, you can deduce what might be causing the problem. For example, getting a reorder tone during dialing is almost certainly the result of an incorrect dial plan. Table 27-2 lists the most common endpoint call setup issues and the probable causes to check related to these issues.

Key Topic

Table 27-2 Common Endpoint Call Setup Issues and Probable Causes

Call Setup Issue	Probable Causes
Reorder tone during and at the end of dialing	Misconfigured CUCM Components: ■ Misconfigured dial plan ■ Insufficient calling privileges ■ Misconfigured digit manipulation
	CAC can also cause this issue
No ring-back tone	IP reachability issues to CUCM
	Gateway may block or drop audio
	ISDN may not provide ring-back tone
Unexpected second dial tone	Misconfigured dial plan
Video is not set up, only audio	Video codec mismatch
	CAC (such as RSVP Video Desired Mode)
	Regions
Dead air is heard	Firewall or ACL blocking media
One-way audio or video	Firewall or ACL blocking media in one direction, or CAC
Call is dropped after dialed	Audio or video codec mismatch
Call is dropped in the middle of the call	Repeatable issues are usually due to network connectivity events, but could also be caused by CAC or permission on the CUCM

The first call-setup issue listed is the reorder tone heard during and at the end of dialing. It's important to remember that Cisco Unified IP phones and collaboration endpoints that register to the Cisco Unified Communications Manager are also controlled by it. Therefore, issues can often be diagnosed with only this call control element, which provides all intelligence and features to the endpoints.

Misconfiguring Cisco Unified Communications Manager components can cause a reorder tone during and at the end of dialing. These components could include misconfiguring the dial plan, insufficient calling privileges for the type of call, or misconfigured digit manipulation at the call control component. CAC can also block the call. Cisco Unified Communications Manager or a separate component in the voice and video network can implement CAC, such as an H.323 gatekeeper or an RSVP-enabled router. If CAC blocks the call, you usually see the message "Not Enough Bandwidth" on the IP phone display or Cisco Telepresence endpoint. If the caller does not have the appropriate privileges to dial the number, the call is blocked. Several companies have a policy to block international numbers or costly premium numbers, such as 900 numbers, for many employees.

IP reachability problems to Cisco Unified Communications Manager can cause no ring-back tone. Alternatively, a gateway, if it is used, might not cut through the audio, or another network, such as ISDN or H.323, simply might not provide the ring-back tone. You should verify that IP connectivity to Cisco Unified Communications Manager is stable. If the affected call was placed outbound to the public switched telephone network (PSTN), additional tools on the gateway or at the network edge may be needed to assess where the issue resides.

An unexpected second dial tone during dialing is typically caused by a misconfigured dial plan on a call control component. The Cisco Unified Communications Manager administrator can try to identify the issue by using the Cisco Unified Communications Manager Dialed Number Analyzer. You must work from the Cisco Unified Communications Manager to diagnose and solve this issue because it cannot be solved at the endpoint level.

If only audio was set up for the call and not video, this issue could indicate a mismatch of video codecs between the endpoints. It could also indicate that CAC has blocked more resource-demanding video. Cisco Unified Communications Manager might have CAC configured to set up only audio if the WAN runs out of bandwidth. This behavior is called RSVP video desired mode. Try to isolate the video codec cause by dialing the same destination from another endpoint model that might support a different set of codecs. To isolate the CAC cause, you can try dialing the destination later, when more bandwidth might be available.

Dead air after the call is set up indicates a media-blocking function in the routed path. An ACL entry might be blocking RTP audio traffic. You can use the ACL packet-logging feature, which will log the denied packets on the router console. Voice and video protocol inspection engines are required at the firewall for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the Cisco Adaptive Security Appliance (ASA) to perform a deep packet inspection instead of passing the packet through the fast path. Several common inspection engines are enabled on the Cisco ASA appliance by default, but you might need to enable others depending on your network.

One-way calling between the same two endpoints means that the call can be set up in one direction only. This issue is typically caused by calling privileges. This behavior might also be intentional and might not be an issue. CAC is a direction-oriented mechanism and can block a call in one direction while allowing media to flow in the opposite direction. You

should consider asymmetric bandwidth utilization in typical TCP/IP networks. A firewall or ACL that blocks a call setup in one direction can also cause the one-way calling issue. Use the packet-logging feature to identify the traffic that a firewall or access list denies. To isolate weak calling privileges, use the Cisco Unified Communications Manager Dialed Number Analyzer, or place the same call from the IP phone of a manager, who is expected to have higher calling privileges.

A call dropping immediately after it is dialed is typically caused by a mismatch of audio codecs; however, a mismatch of video codecs can also cause this issue. The behavior depends on endpoint algorithms if the call is dropped because of mismatched video codecs or is degraded to audio only. Try to place a test call from a different endpoint model that might support a different set of codecs. The Cisco Unified Communications Manager administrator should also check the codecs that are enabled for the affected endpoint or use the trace utility to see the sequence of events that led to the dropped call behavior.

Calls that are dropped in the middle of a conversation are usually caused by a network connectivity event. Make sure that the event is repeatable and not transient before you start troubleshooting. Check the IP connectivity between the endpoints in both directions. You can power-cycle the affected endpoints to initialize their software and then try to make a call again.

Key Topic

Calling issues on Cisco Telepresence CE software-based endpoints are like the potential issues of standard Cisco Unified IP phones. Because of the complex setup, a great variety of problems can be experienced when setting up a call with a Cisco Telepresence CE software-based endpoint. Cisco Telepresence CE software-based endpoints can register to the Cisco Unified Communications Manager or to the Cisco Expressway. Depending on the provisioning model, you should verify that the dial plan is correct and that calls can be successfully routed through that call control server. Also, related to the dial plan at the Cisco Unified Communications Manager, the caller can be blocked because of insufficient calling privileges. IP connectivity problems also contribute to failed call setup. Issues with IP routing, firewall configuration, and access lists that filter out audio or video in the routed path can all prevent a call from being successfully set up. CAC is used to maintain a good quality of media. In essence, CAC sets up a bandwidth limit that can accommodate a certain number of calls. If this bandwidth is consumed, no additional calls are admitted, and the calls fail to set up.

Common Media Issues

From the source of video information until the video is displayed at the remote endpoint, the video media might be subject to several factors that influence its final quality. Video call setup negotiates video parameters end to end, so the local endpoint adjusts its behavior according to the far-end endpoint capabilities, and vice versa. At the local endpoint that generates the video stream, a camera might not be properly tuned to local environmental conditions. As a result, the camera might not be able to focus on an object, such as because of a dark room or narrow contrast range. Figure 27-10 illustrates common issues that can affect media quality within a Cisco Collaboration solution.

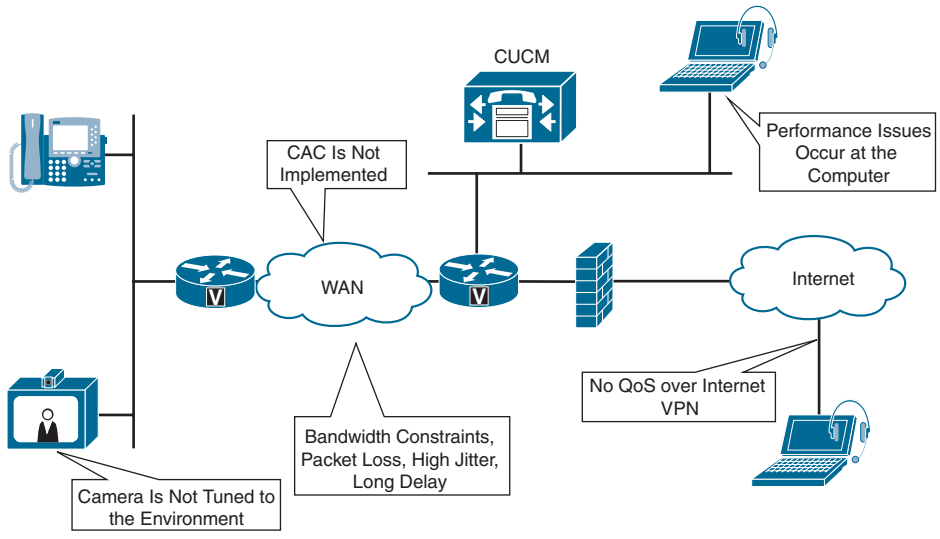

Figure 27-10 *Media Quality Issues Within a Cisco Collaboration Solution*

The result of each video-encoding algorithm is the video bit rate. Because the video bit rate is dynamic and depends on many factors, it might be calculated and provisioned incorrectly across the WAN. When the video stream is transferred over the WAN, bandwidth reservation might be insufficient. High packet loss and high jitter can cause the far-end video decoder to miss video information, and the resulting video can be jerky. High end-to-end delay makes an interactive video call a bad experience.

If CAC is not implemented in the network, there is no control over the bandwidth that is available for video. All video calls might struggle because of lack of bandwidth per call, when too many calls exist. In addition to CAC, voice and video have specific traffic behavior and QoS requirements. Voice generally requires the network to meet the following characteristics in each direction:

Key Topic

- Jitter, or the variation in delay, should be less than 30 ms.
- One-way end-to-end delay for an interactive call must be less than 150 ms.
- Packet loss should generally be less than 1 percent. However, the speech quality of some audio codecs decreases rapidly when reaching a 0.5-percent packet loss, such as G.729A or G.729AB.
- Voice traffic should be marked with a DSCP EF to benefit from specific queuing behavior while transferred.

Video has requirements that are similar to voice. The only considerable difference is that video demands much more bandwidth, and its bandwidth requirements vary greatly. Video traffic should be marked with DSCP AF41 to benefit from bandwidth guarantees while transferred end to end. Many deployments mark voice and video the same. Packet loss results in a poor user experience and can usually be seen by artifacts on the screen. You can view call statistics to determine whether the network is experiencing the packet loss.

Key Topic

Call statistics can be checked on Cisco CE software-based endpoints using the web interface, as previously discussed, using the on-screen display, and using the CLI. Call statistics include packet loss, jitter, and delay statistics for incoming and outgoing audio and video channels. Jitter is the variation of delay between packets received by the endpoint. To check for packet loss in a call, check the system information screen or use the **xstatus diagnostics** command from the CLI. The main areas of packet loss will probably be on the video channels but can also be on the audio. Dropped packets are packets that arrive too late to be used or that were lost and never received. You can also see the jitter statistics and the channel rate. The video channel rate can change dynamically, depending on how much change is in the picture. If there is excessive packet loss or high jitter, you need to look at the network to determine the cause.

If the CE software-based endpoint monitor displays an image outside of the monitor frame, be aware that most monitors will overscan TV resolutions such as 720p and 1080p. To fix this issue, you must enter the monitor menu system and find the appropriate setting. Different monitors use different terms, but look for something like Just Scan, Pixel by Pixel, or Underscan. Some monitors have this option for the native resolution of the display only. If you cannot find the option, go to the Endpoint Administrator Settings menu, change the video-output resolution, and try again.

If you cannot get any audio from the CE software-based endpoint when it is connected to a monitor by using HDMI, check whether the video output resolution for HDMI 1 is set to 800 × 600 or 1900 × 1200. Because of an issue with these two resolutions, they are run in DVI mode (DVI over HDMI), which does not support audio. To check or change the video-output resolution, navigate to **System > Configuration > Video > Output > HDMI 1 > Resolution** and choose another video output resolution for HDMI 1, or use the audio from Audio Line Out 1 and 2. If the audio on the dual stream is out of sync with the dual-stream video, be aware that Cisco CE software-based endpoints do not support lip synchronization between audio and video over SIP.

When you install but cannot control a second Cisco PrecisionHD camera on your CE software-based endpoint codec, you need to use a Video System Control Architecture (VISCA) cascading cable to connect the cameras. The VISCA cascading cable connects the first and second cameras. You need to configure the video-input source to set which camera you should control when this particular video-input source is active. To configure the second camera, which is connected to video input source 2, execute the CLI command **xConfiguration Video Input Source 2 CameraControl CameraId: 2**.

If audio is distorted, one of these issues might be the cause:

- Echo cancellation is not working and might be disabled.
- HDMI is being used for audio, but the audio is delayed by processing on the monitor. Modern TVs are often used and have picture processing that can delay the audio, and the codecs cannot echo-cancel this situation. Determine whether the monitor has a game mode or other no-processing mode to stop any video and audio processing. You can test whether the monitor is the issue by attaching some active speakers to the codec output and determining whether there are still echo cancellation issues. Echo cancellation on Cisco CE software-based endpoints supports a maximum of 340 ms on all the audio bandwidth.

Cisco ClearPath is an innovative technology that can improve the quality of video while minimizing the effect of packet loss in networks with uncontrolled quality of data transmission. This technology is a set of algorithms for compensation of network losses in the communication channel and can work with any image quality, up to full high definition. Table 27-3 summarizes the major elements that influence media quality in a Cisco Collaboration solution.

Key Topic

Table 27-3 Major Elements Influencing Media Quality in a Cisco Collaboration Solution

Element	Description
Input video peripherals	Video input peripheral quality, lens, exposure range, focus capabilities, lighting conditions
Video codec	Selection of video codec, performance, and capabilities of video codec
Output video peripherals	Video output peripheral quality: screen, image-enhancing capabilities
Amount of video information	Video resolution, frames per second, object-moving behavior, background complexity, resulting video bit rate
Network QoS	Network, packet loss, jitter and delay characteristics, CAC, differentiated services in converged network
Correct bandwidth provisioned	Bandwidth overhead calculation and provisioning
CPU utilization	Computer hosting Cisco Jabber running many applications

Input-video peripheral quality has a direct influence on the quality of the video information that is produced. The following are major characteristics of the input-video peripheral:

- The quality of the lens that the camera uses
- The exposure range that the camera can process
- The focus capabilities
- The recommended lighting conditions for which the camera is designed

Selection of the video codec must consider the performance requirements, codec capabilities, and limitations in terms of its susceptibility to declined QoS, such as higher jitter. H.264 standards do not cover precisely how to encode and decode a video stream or how to recover from issues, such as from missed frames. Some of these functions are left to the vendor-endpoint implementations, which Cisco has revolutionized with technologies like ClearPath.

An output-video peripheral can have a strong impact on how video information is presented at the far end. Screens that are not tuned to contrast and brightness according to surrounding conditions can greatly affect how the video is experienced. Vendors also implement various video-enhancing capabilities, such as sharpening and noise reduction, that can improve the final media content perceived by the far-end participants.

The amount of video information has an impact on how much bandwidth the associated video stream requires in terms of video bit rate. Parameters that influence the video bit rate are video resolution, number of frames per second, and object moving behavior or complexity of the background behind the object.

The quality of video information can decline most rapidly with decreasing QoS in the video-transport network. High packet loss and high jitter can have a strong influence on video decoding behavior and can cause rapid degradation in video quality. High end-to-end delay is a problem for human conversation and makes it difficult to continue as normal. If the network carries multiple traffic types, differentiated services must be implemented to provide each traffic type with its own QoS requirements.

When calculating bandwidth requirements for an expected video stream, bandwidth overhead is seldom considered. Failure to consider bandwidth overhead leads to insufficient bandwidth reservation and packet loss. In a Cisco Unified Communications network, some additional bandwidth overhead above the expected video-stream bit rate must be considered. The network designer must address this situation.

The computer that hosts the client requires available processor time to encode and decode a video stream. If the computer runs too many applications and processor utilization is high, there might be no processor time left for the software-based video codec within the client. Processing power is also reduced when the computer runs on batteries. Endpoint-to-endpoint interoperability is also important when it comes to video quality. Video that is coded on an endpoint can appear inadequate on one endpoint but look perfect on another endpoint.

Troubleshooting Cisco Jabber

A few Cisco Jabber tools can be used to help troubleshoot issues. You can display message notifications that are collected during the registration process, whenever you change the operational mode, or whenever you choose a different deskphone device. Navigate to **Help > Show Error Notifications** to see whether any error or warning messages have been logged.

If a user encounters a problem with Cisco Jabber, that user can create a problem report. The user can enter a description of the problem, and the description is included in an autogenerated report. The report contains logs from the user's computer, and the report is saved to the desktop. The user can then send this file to the system administrator to help analyze the problem. The problem report can be generated from the Cisco Jabber interface, or it can be created from outside the application by navigating to **Start > All Programs > Cisco Jabber > Cisco Jabber Problem Report**. Use the following steps to create the problem report from Cisco Jabber:

Step 1. On the Cisco Jabber client, navigate to **Help > Report a Problem**.

Step 2. Navigate through the three separate windows to create the report. In the first window, read through the welcome message, which has an important notice, and click **Continue**.

Step 3. In the second window, provide a single-sentence explanation of the issue, choose the problem category from the drop-down list, and optionally enter a more detailed problem description. Click **Continue**.

Step 4. In the third window, you can optionally attach a file to the report and click **Generate**. A copy of the report is saved to your computer desktop.

Cisco Jabber saves the problem report to the computer desktop as a .zip file. The file contains all the logs that are collected at the computer. You can unzip these files and assess the issue yourself, or send this .zip file to Cisco TAC for further analysis. Figure 27-11 illustrates the procedure to generate a problem report from Cisco Jabber.

Figure 27-11 *Problem Report from Cisco Jabber*

When Cisco Jabber is set to Automatic Registration, it is hard-coded to always try registering to the cloud first. Then it will try to register with Cisco Unified Communications Manager IM and Presence Service or the Cisco Unified Communications Manager. Last, it will try to send the registration communication to the Expressway Edge. The Cisco Jabber client can also be configured to manually search for any one of these services.

The client registers with different services of Cisco Unified Communications Manager, depending on whether the client operates in deskphone or softphone mode. Cisco Jabber shares all IP parameters with the host computer. These parameters must be set properly at the computer, which must have IP connectivity to Cisco Unified Communications Manager IM and Presence Service, Cisco Unified Communications Manager, and the Cisco TFTP

server, if the client operates as a softphone. Cisco Unified Communications Manager CTI Manager is the service that is usually run with other services on Cisco Unified Communications Manager. For larger deployments, Cisco Unified Communications Manager CTI Manager can be run on a separate server. Cisco Jabber must be able to exchange traffic with the servers for the protocols that are used during the registration process. Most communication is based on HTTPS, such as SOAP and Cisco CallManager. Cisco IP Phone Services. XMPP is also used, as are TFTP, SIP, and CTIQBE, depending on the mode of the client. You should make sure that the appropriate port numbers are open at the firewalls that are installed in the routed path:

Key Topic

- SOAP uses server ports TCP 8191 and TCP 8443.
- XMPP uses server port TCP 5222.
- Cisco CallManager Cisco IP Phone Services use server port TCP 8443.
- TFTP uses server port UDP 69.
- SIP registration uses server ports UDP 5060, depending on how SIP was provisioned at the servers. Secure SIP uses 5061.
- CTIQBE uses server port TCP 2748.

The registration process for Cisco Jabber starts with entering a user ID, password, and the Cisco Unified Communications Manager IM and Presence Service address. The client automatically receives from the servers all other parameters that are required during the registration in softphone or deskphone mode. If the Cisco Unified Communications Manager IM and Presence Service address was entered as an IP address rather than by name, the address will change to the Cisco Unified Communications Manager IM and Presence Service host name automatically during the SOAP configuration exchange. The host computer must be able to resolve this host name, or the address field must be rewritten manually to the IP address during each new user login. This issue is frequently the reason that the client does not register.

Cisco Jabber shows different error messages when the login fails to aid in troubleshooting. The error message “Invalid user ID or password. Please try again” indicates an issue with the user credentials. The credentials are authenticated against the local Cisco Unified Communications Manager database or the LDAP server. The credentials might have expired at either server, or the LDAP server might not be reachable by the Cisco Unified Communications Manager. The error message “Unable to connect to network. Please check your network connection” indicates network-connectivity issues, DNS-resolution issues, issues at Cisco Unified Communications Manager IM and Presence Service, or the user might not hold the license for the client.

Also, Cisco Jabber is incompatible with NAT. To traverse NAT, the client must be behind a VPN connection. NAT and firewalls can cause a range of registration issues that must be resolved with the help of a security engineer. When deploying a Cisco Adaptive Security Appliance and Cisco AnyConnect Secure Mobility Client for Cisco Jabber for Windows, you should configure NAT rules to support the Cisco AnyConnect Secure Mobility Client. If you do not configure NAT rules, the Cisco AnyConnect Secure Mobility Client cannot communicate with the Cisco ASA.

If Cisco Jabber logs in and you want to make sure that the client has registered properly, use the Server Status and Notifications window to determine whether any components failed to register. If so, you can generate a problem report that contains the logs to provide more information to diagnose the problem. Cisco Jabber is hosted on a computer with other applications; therefore, it can face very specific issues that relate to the host computer environment. Table 27-4 identifies potential issues that can arise from this specific computer environment for Cisco Jabber.

Key Topic

Table 27-4 Common Call Setup Issues on Cisco Jabber

Issue	Possible Causes
No calls possible	Softphone is not registered, or deskphone CTI control does not work. Network connectivity issues or misconfigured servers exist.
No audio (softphone mode)	Required network ports are not open on the computer that hosts the application.
One-way audio or video (softphone mode)	Computer audio or video device on either side does not work. If connected over a Cisco VPN client (for Windows), ensure that the stateful firewall is disabled.
Poor audio quality	If echo or feedback is heard, be sure to use a proper headset and not computer speakers.
Incoming video is black (permanent)	Required network ports are not open on the computer that hosts the application.
Incoming video is black (transient)	Local or far-end computer experiences lack resources to encode or decode the video signal. Camera could be muted.

When a user logs in to the client but no calls are possible, consider the operational mode that Cisco Jabber uses:

- If Cisco Jabber is in softphone mode, the client might not be successfully registered as being in softphone mode with Cisco Unified Communications Manager.
- If Cisco Jabber is in deskphone mode, the CTI control component might not work because of an unregistered hardware Cisco Unified IP phone. Alternatively, CTI communication issues might exist, or the CTI gateway might be misconfigured in Cisco Unified Communications Manager or the Cisco IM and Presence Server. Also, the called target might not be registered with the call control server.

To verify Cisco Jabber connectivity, navigate to **Help > Show Connection Status** or **Help > Show Error Notifications**. If no audio is received in softphone mode, verify that the correct RTP ports (UDP 16384 to 32766) are open on the computer firewall.

One-way audio or video in softphone mode means that voice is heard or video is seen in one direction only. If the client is connected over a Cisco VPN client (for Windows), make sure that the stateful firewall setting is disabled on the host computer. The microphone, as an audio device peripheral, or the camera also might not work at the far end. When experiencing black incoming video as a permanent issue, and when only audio works, you must also consider the firewall setup on the host computer. If black incoming video is only a transient issue, the far-end system that encodes the video or the computer that decodes the video might be experiencing a temporary lack of computing resources.

Some major elements influence video quality in a Cisco Unified Communications environment for Cisco Jabber as well. These elements are the same for Cisco Jabber as they are for any other Cisco Unified IP phone or Telepresence endpoints. Refer to the media quality discussion under the "Common Media Issues" section for a review of these elements.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 27-5 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 27-5 Key Topics for Chapter 27

Key Topic Element	Description	Page Number
Paragraph	Status Message Types on IP Phones	637
List	Log Archive Download Options	640
List	Files Included in the Current Log Bundle Folder	641
List	Files Included in the Eventlog Folder	643
List	Call Info Logs on CE Endpoints	644
List	Information Found in Call Logs on CE Endpoints	645
Paragraph	How to Read a Detailed Debug Log	646
Paragraph	How VLANs Impact Registration	649
List	Types of DHCP Servers	649
List	Tools for Testing IP Connectivity	650
List	TFTP SIP and SCCP Ports Used	651
Paragraph	How Registration Is Impacted from CTL and ITL Files	652
Paragraph	Call Setup Issues	652
Table 27-2	Common Endpoint Call Setup Issues and Probable Causes	653
Paragraph	Call Setup Issues in CE Endpoints	655
List	QoS Characteristics	656
Paragraph	Checking Call Statistics on CE Endpoints	657
Table 27-3	Major Elements Influencing Media Quality in a Cisco Collaboration Solution	658
List	Ports Used by Cisco Jabber	661
Table 27-4	Common Call Setup Issues on Cisco Jabber	662

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

802.1x, CAC, Codec, CTL, Current Logging, DHCP, DSP, Eventlog, Extended Logging, Historical Logs, ITL, Option 150, TVS, VLAN, VVID

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the three types of log archives that can be downloaded with their descriptions.
2. What are the four major characteristics of input-video peripherals that can influence media quality?
3. List the four steps to create a problem report on Cisco Jabber.

This page intentionally left blank

CHAPTER 28

Cisco Unified Communications Manager Reports

This chapter covers the following topics:

Dialed Number Analyzer: This topic will explain how to enable, access, and use the Dialed Number Analyzer tool on the Cisco Unified Communications Manager.

CAR Tool: This topic will explain how to enable, access, and use the Cisco CDR Analysis and Reporting tool, also known as the CAR tool, on the Cisco Unified Communications Manager.

CDR and CMR Logs on CUCM: This topic will explain the information contained within and the use for call detail records (CDRs) and call management records (CMRs). These records include user reports, system reports, and device reports.

Many tools can be used to troubleshoot issues and monitor analytical data from the Cisco Unified Communications Manager. Because this book is designed to introduce administrators to the basic functions of a Cisco Collaboration solution, this chapter will cover only two of these tools. Topics discussed in this chapter include the following:

- Dialed Number Analyzer
- CAR Tool
- CDR and CMR Logs on CUCM
  - User Reports on CUCM
  - System Reports on CUCM
  - Device Reports on CUCM

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.4 Troubleshoot these network components in a Cisco Collaboration solution
  - 1.4.c LDAP integration on Cisco Unified Communications Manager

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 28-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 28-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Dialed Number Analyzer	1–2
CAR Tool	3–4
CDR and CMR Logs on CUCM	5–6

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. How many services must be enabled on the CUCM before the Dialed Number Analyzer tool can be used?
   - **a.** None
   - **b.** One
   - **c.** Two
   - **d.** Four
2. How does the CAR tool process CDR files from the CDR Repository?
   - **a.** Compressed tar.gz files
   - **b.** Flat files
   - **c.** Direct from the source, not from the repository
   - **d.** Both flat files and compressed tar.gz files
3. Which of the following details are included in a CMR?
   - **a.** The date and time the call was started
   - **b.** The amount of data sent and received during the call
   - **c.** The time the call connected
   - **d.** The cause for the termination of a call
4. Which of the following user reports is available to users, managers, and administrators?
   - **a.** Top N
   - **b.** Cisco IP Phone Service
   - **c.** Cisco Unified Communications Manager Assistant (IPMA)
   - **d.** Bills
5. Which of the following system reports are available to managers and administrators?
   - **a.** QoS Summary
   - **b.** Traffic Summary
   - **c.** System Overview
   - **d.** CDR error

6. Which of the following is a device report available through the CAR tool on the CUCM?
   a. Endpoint
   b. Trunk
   c. Jabber
   d. Gateway

Foundation Topics

Dialed Number Analyzer

Many tools are available through the Cisco Unified Communications Manager. This vast ocean of tools exists for the many different call scenarios that can be processed by the CUCM. Although this chapter will not cover all the different tools available, a few tools are worth mentioning.

The Cisco Dialed Number Analyzer (DNA) is a calling simulator that will determine whether a call between a specific source and destination alias is possible through the current configuration within the Collaboration network. When DNA is used, a call is not actually placed. The returned data shows whether the call would be possible and indicates all of the call control mechanisms that will affect this call. Therefore, when calls are not possible, DNA will show what component in the Cisco Unified Communications Manager is preventing the call from connecting, if any apply. You can use the DNA tool to troubleshoot dial plan issues and issues that relate to calling privileges in Cisco Unified Communications Manager.

The Dialed Number Analyzer installs as a feature service along with Cisco Unified Communications Manager. The DNA tool can be used to analyze dial plans after they are deployed. You can use the tool to test a dial plan by providing dialed digits as input. This tool also analyzes the dialed digits and shows details of the call. You can use these results to diagnose the dial plan and identify any problems.

The Dialed Number Analyzer allows analysis of inbound and outbound calls in a Cisco Unified Communications Manager dial plan. It analyzes the calls and provides results that show complete details, including call patterns and calling- and called-party transformations that are applied to the dialed digits.

Before you can use the Dialed Number Analyzer, a Cisco Unified Communications Manager administrator must activate and implement the service. To activate the Dialed Number Analyzer services, follow these steps:

Step 1. From the Cisco Unified Serviceability page, navigate to **Tools > Service Activation**.

Step 2. Select the CUCM publisher from the drop-down list and click **Go**.

Step 3. Under the CM Services section, check both of the following DNA services:

- **Cisco Dialed Number Analyzer Server:** If you have more than one node in the CUCM cluster, activate this service on the node that will be dedicated to the Dialed Number Analyzer service.

- **Cisco Dialed Number Analyzer:** If you are planning to use the CUCM Dialed Number Analyzer, activate this service. This service may consume a lot of resources, so activate this service only on the CUCM node with the least amount of call-processing activity or during off-peak hours.

Step 4. Click **Save** to enable these services. Figure 28-1 illustrates how the services for the Dialed Number Analyzer can be activated.

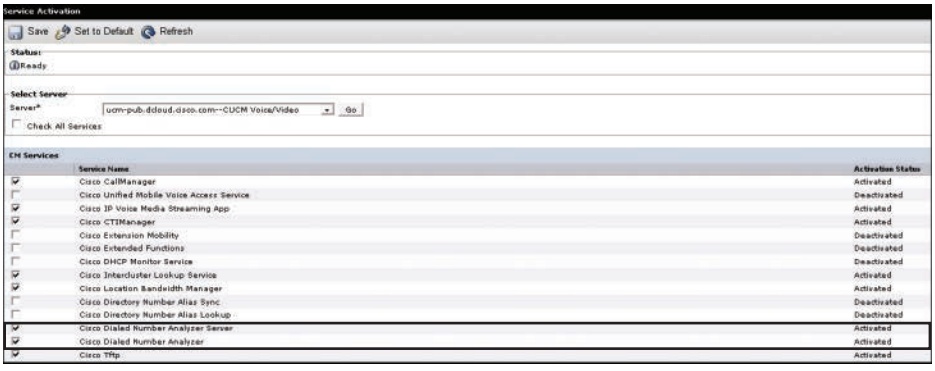

Figure 28-1 *Dialed Number Analyzer Services on the CUCM*

After activating DNA, you can use it to analyze call connection attempts. Although the DNA web page can be accessed from the Cisco Unified Serviceability menus, another way to access this tool may be a little bit easier. Follow these steps to access the Cisco Unified Communications Manager Dialed Number Analyzer:

Key Topic

Step 1. Navigate to **https://<***CUCM IP Address***>/dna**.

Step 2. Choose **Analysis > Phones**.

Step 3. Leave the text box next to the Find button empty and click **Find** to view a list of all available devices. Click the Device Name of the endpoint from which you want to simulate a call attempt.

Step 4. In the window that appears, do the following:

- From the Association Information section, choose **Line (1)**.
- In the Dialed Digits Settings section, choose the **Dialed Digits** radio button, and then enter the directory number of another device.
- Check the **SIP Analysis** check box.

Step 5. Click the **Do Analysis** button.

A new window will pop up with data from the analysis. You can use this data to analyze the results that appear in the results window. Figure 28-2 shows output from the Cisco Unified Communications Manager Dialed Number Analyzer.

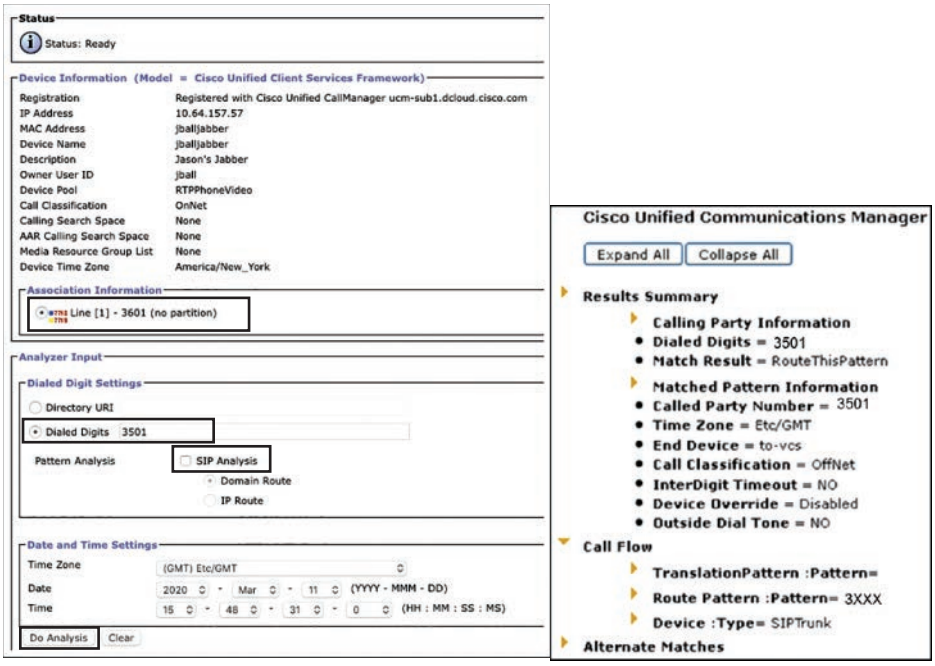

Figure 28-2 *Dialed Number Analyzer Tool and Output Results*

This output was generated from the Analyzer window, where the calling phone, number 3601, simulates dialing to the number 3501. The output window on the right side of the figure starts with a Results Summary and ends with Alternate Matches. The content of the analysis output depends on the dial plan configuration and the Analyzer window that was used to generate that output. This particular analysis shows that the called party, number 3501, is reachable by using a SIP trunk that is associated with route pattern 3XXX.

CAR Tool

All calls that the Cisco CallManager service processes can be logged by the Cisco Unified Communications Manager. This data includes different call details and call quality information. The Cisco Unified Communications Manager CDR Analysis and Reporting (CAR) tool uses this data to provide system-generated reports to users, managers, and administrators. These reports can help monitor QoS issues, device utilization, or call statistics.

As its primary function, CAR generates reports about the users of CUCM and generates reports on the system status with respect to call processing. CAR also performs CAR database management activities. Administrators can perform these tasks in one of the following ways:

- Automatically configure the required tasks to take place
- Manually perform the tasks by using the web interface

All CAR reports use CDR data. CAR processes the CDRs from flat files that the CDR Repository service places in the CDR Repository folder structure. CAR then processes

CDRs at a scheduled time and frequency. By default, CDR data loads continuously 24 hours per day and seven days per week. However, the administrator can set the loading time, interval, and duration as needed. In addition, the default setting loads only CDR records. CMR records are not loaded by default. Figure 28-3 illustrates the CAR tool characteristics previously outlined.

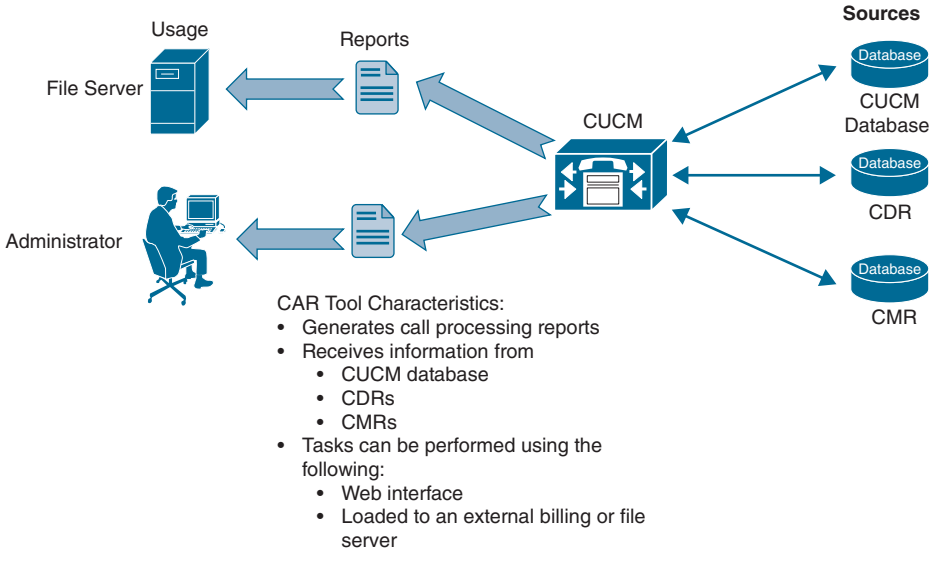

Figure 28-3 *CAR Tool Characteristics*

CAR is composed of a group of complementary services that administrators can activate in the Service Activation window of the Cisco Unified Serviceability page. Before CAR can be launched from the Tools menu in Cisco Unified Serviceability, you must activate the CAR services. In Cisco Unified Serviceability, navigate to **Tools > Service Activation**. Check the boxes next to the following CDR services:

- Cisco CAR Web Service
- Cisco SOAP-CDROnDemand Service (optional). (If using a third-party billing application that accesses CDR data via an HTTPS/Simple Object Access Protocol [HTTPS/SOAP] interface, activate this service.)

Certain CUCM service parameters must also be enabled to ensure that the CDR records are generated. The administrator can configure these parameters on the Service Parameters Configuration window from the Cisco Unified CM Administration page. To access the Service Parameters Configuration window, navigate to **System > Service Parameters**. Select the publisher CUCM from the first drop-down list, and then select the Cisco Call Manager service from the second drop-down list. Click the Advanced button to display the complete list of service parameters available. Table 28-2 identifies the list of service parameters that can affect CDR and CMR records.

Key Topic

Table 28-2 CDR and CMR Service Parameters on the CUCM

Service Parameter	Description
CDR Enabled Flag	This parameter determines whether CDRs are generated. Valid values specify True (CDRs are generated) or False (CDRs are not generated). For this required field, the default value specifies False. Enable this parameter on all servers.
CDR Log Calls with Zero Duration Flag	This parameter enables or disables the logging of CDRs for calls that never connected or that lasted less than one second. Cisco CallManager logs unsuccessful calls (calls that result in reorder, such as might occur due to a forwarding directive failure or calls that attempt to go through a busy trunk) regardless of this flag. This is a required field with a default value of False.
Call Diagnostics Enabled	Three settings can be configured under this menu option: ■ **Enabled Only When CDR Enable Flag is True:** Generates CMRs only when the CDR Enabled Flag service parameter is set to True. ■ **Enabled Regardless of CDR Enabled Flag:** Generates CMRs without regard to the setting in the CDR Enabled Flag service parameter. This parameter represents a required field. ■ The default value specifies **Disabled**, which will not generate CMRs.
Display FAC in CDR	This parameter determines whether the forced authorization codes (FAC) associated with the call display in the CDR. Valid values specify True (display authorization code in CDRs) or False (do not display authorization code in CDRs) for this required field. The default value specifies False.
Show Line Group Member DN in finalCalledPartyNumber CDR Field	This parameter determines whether the finalCalledPartyNumber field in the CDRs shows the directory number of the line group member who answers the call or the hunt pilot directory number. Valid values specify True (the finalCalledPartyNumber in CDRs will show the directory number of the phone that answered the call) or False (the finalCalledPartyNumber in CDRs will show the hunt pilot directory number). This parameter applies only to basic calls that are routed through a hunt list without future interactions, such as transfers, conference, and call park. If a feature is involved in the call, the hunt pilot directory number will show in the finalCalledPartyNumber field regardless of the setting in this parameter. The default value for this required field specifies False.

Service Parameter	Description
Add Incoming Number Prefix to CRD	This parameter determines whether CUCM adds the incoming prefix (as specified in the National Number Prefix, International Number Prefix, Subscriber Number Prefix, and Unknown Number Prefix Service Parameters) to the calling party number in the CDRs for that call. If the prefix is applied on the inbound side of the call, it is always added to the calling party number in the CDRs for that call. This occurs even if this parameter is set to False. If the prefix is applied on the outbound side, the prefix is added to the calling party number in the CDR or CDRs for that call, only if the parameter is set to True. If the Destination of the call is a gateway, CUCM will not add the prefix to the CDRs even if this parameter is enabled. This parameter is applied on a clusterwide basis. The default value for this required field specifies False.

The CAR tool provides different levels of user rights. In fact, the CAR tool provides reporting capabilities for three levels of users:

- Administrators are allowed to use all the features of CAR so that they can generate system reports to view system performance, verify load balancing, and troubleshoot.
- Managers can generate reports for users, departments, and QoS to help with call monitoring for budgeting or security purposes. Reports can also be used for determining the voice quality of the calls (for example, to ensure compliance with service-level agreements).
- End users can generate a billing report for their calls.

Any application or end user can act as a CAR administrator. Users who have been identified as CAR administrators have complete control over the CAR system. The administrator can modify all parameters that relate to the system and the reports. Cisco Unified Communications Manager CAR requires a minimum of one administrator.

CDR and CMR Logs on CUCM

The call detail records (CDR) and call management records (CMR) architectures include different ways in which CDR and CMR files can be loaded to the CAR tool. CDRs provide details about the following:

- The called number
- The calling number
- The date and time that the call was started
- The time that the call connected
- The time that the call ended
- The cause for the termination of a call

CMRs include jitter, lost packets, the amount of data that was sent and received during the call, and latency. CDR data comprises CDRs and CMRs, collectively. A single call can result in the generation of several CDRs and CMRs. The CUCM records information regarding each call in CDRs and CMRs. CDRs and CMRs, known collectively as CDR data, serve as the basic information source for Cisco CAR.

The Cisco CDR Agent service transfers CDR and CMR files that CUCM generates from the local host to the CDR Repository node, where the CDR Repository Manager service runs over an SFTP connection. If the SFTP connection fails, the Cisco CDR Agent service continues to make connection attempts to the CDR Repository node until a connection is made. The Cisco CDR Agent service sends any accumulated CDR files when the connection to the CDR Repository node resumes. The CDR Repository Manager service maintains the CDR and CMR files, allocates the amount of disk space for use by CMRs and CDRs, sends the files to up to three configured destinations, and tracks the delivery result for each destination. The Cisco CAR tool accesses the CDR and CMR files in the directory structure that the CDR Repository Manager service creates. By default, the CAR tool will retain data for 60 days and up to 80 percent for a low watermark and 90 percent for a high watermark of the database. High and low watermarks can be modified for data retention in CAR, and the max age of data can be set up to 180 days. To modify these settings, navigate to **System > Database > Configure Automatic Purge**. Figure 28-4 illustrates the CDR and CMR structure within the Cisco Unified Communications Manager.

CDR and CMR Architecture:

- Call processing nodes create the CDR and CMR flat files for registered devices
- The CDR agent service transfers these files to the CDR repository node via SFTP
- The CDR repository manager runs on the CDR repository node and maintains CDR and CMR files
- The CAR tool accesses the CDR and CMR files from the directory structure of the CDR repository node

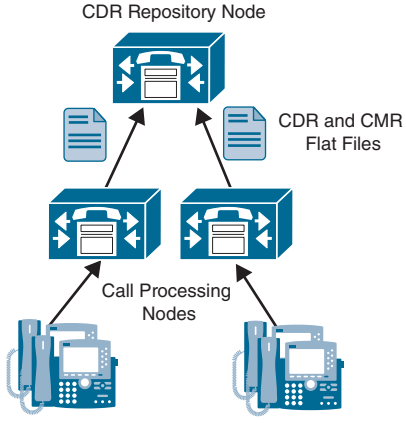

Figure 28-4 *CDR and CMR Structure on the CUCM*

The CAR tool allows CDR and CMR dump information to be exported to a location on the local computer. The CDR and CMR dump is in CSV format and can be used for external billing tools. The following procedure shows how to export CDR and CMR data:

Key Topic

Step 1. From the Cisco Unified CM CDR Analysis and Reporting page, navigate to **CDR > Export CDR/CMR**. The Export CDR/CMR Records window displays.

Step 2. In the From Date and To Date drop-down lists, choose a date range for the CDR and CMR dump.

Step 3. In the Select Records section, check the **CDR Records** box, or the **CMR Records** box, or both.

Step 4. Click **Export to File**.

Step 5. In the new window that appears, right-click **CDR Dump** or **CMR Dump** and choose **Save As** to download the files. Choose the folder to which the files will be saved on the local disk.

Step 6. If you check the Delete File box, the files will be deleted from the CUCM as soon as you click the Back or Close button.

The CDR Search feature can be used to display details about a phone call, such as to show QoS details. Follow these steps to search for a special call:

Step 1. Navigate to **CDR > Search > By User/Phone Number/SIP URL**. The configuration window for the CDR search appears.

Step 2. In the configuration window for CDR Search, you can specify a phone number or range. You can use the Search Internal Phone Number/SIP URL Based User link to display a search window to find configured end users in CUCM. In this window, search for a first name or last name and click **Select** to search for this user. Then, in the From Date and To Date fields, enter a date and time range. When ready, click **OK** to perform the search.

A table with the results of the CDR search will display in a new window. The result includes all call details. To send the report via email, click the Send Report button. To display the media information, click the Others link to open the details in a new window. The media information includes QoS details, such as jitter or latency values for each call. Figure 28-5 illustrates the search results generated in the table.

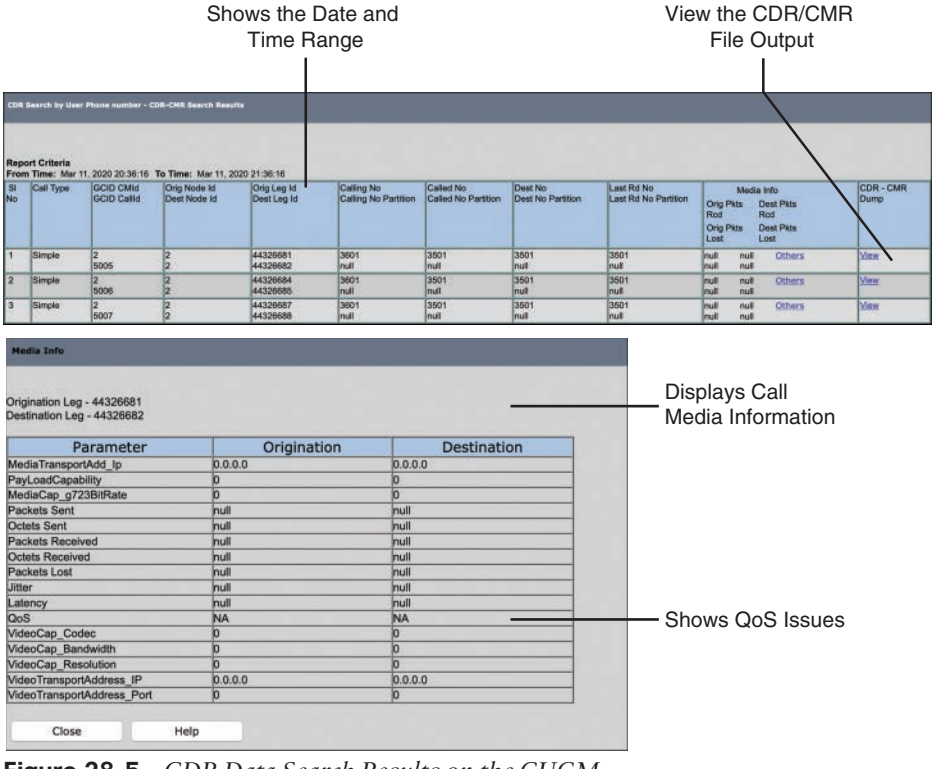

Figure 28-5 *CDR Data Search Results on the CUCM*

User Reports on CUCM

The CAR tool can be used to generate CDR reports. Users, managers, and CDR administrators can generate user reports. The CAR tool includes user reports for Bills, Top N, Cisco Unified Communications Manager Assistant (IPMA), and Cisco IP Phone Service.

Bills user reports can be generated for an individual or based on an entire department. Individual user reports are available for users, managers, and CAR administrators. Individual bills provide call information for the date range that is specified. This report allows you to generate, view, or email summary or detailed information about individual phone bills. CAR administrators who are application users cannot generate this report.

Department user reports are available for managers and CAR administrators. Department bills provide call information and QoS ratings. A manager can generate a summary or detailed report of the calls that all users make who report to the manager or only those users whom the manager selects. A CAR administrator can generate a summary or detailed report of the calls that some or all users in the system make.

Top N user reports provide data on users based on the top charges, longest duration, or most consumed calls. By Charge user reports are available for managers and CAR administrators. The Top N by Charge reports list the top number of users who incurred a maximum charge for calls during a period that the manager or CAR administrator specifies. Reports that are generated by destinations show the destinations that incurred the maximum charges. Reports that are generated by all calls list the calls that incurred the maximum charges. By Number of Calls user reports are available for managers and CAR administrators as well. The Top N by Number of Calls report lists the users who incurred the maximum number of calls. Reports that are generated using extensions list the extensions that placed or received the greatest number of calls during the period that is specified.

The Cisco Unified Communications Manager Assistant (IPMA) user reports are available only for administrators, and they are based on Manager Call Usage or Assistant Call Usage reports. For Manager Call Usage, the Cisco Unified Communications Manager Assistant Summary and Detail reports provide call completion usage details for Cisco Unified Communications Manager assistant managers. The manager reports can include only calls that managers manage for themselves, or only calls that assistants manage for managers, or calls that both managers and assistants handle for managers. For Assistant Call Usage, the Cisco Unified Communications Manager Assistant Summary and Detail reports provide call completion usage details for Cisco Unified Communications Manager Assistants. The Assistant reports can include only calls that assistants manage for themselves, or only calls that assistants manage for managers, or calls that assistants manage for themselves and for managers.

Cisco IP Phone Service user reports are also available for CAR administrators exclusively. The Cisco IP Phone Services report shows selected Cisco IP Phone Services, the number of users who are subscribed to each of the selected services, and the utilization percentage for each of the selected services. Table 28-3 provides a summary of all the different user reports available on the Cisco Unified Communications Manager.

Key Topic

Table 28-3 User Reports on the CUCM

User Report	Method of Application	User Access Allowed
Bills	Individual	Users, Managers, or Administrators
	Department	Managers or Administrators
Top N	By Charge	Managers or Administrators
	Duration	Managers or Administrators
	By Number of Calls	Managers or Administrators
Cisco Unified Communications Manager Assistant (IPMA)	Manager Call Usage	Administrators
	Assistant Call Usage	Administrators
Cisco IP Phone Service	Shows Cisco IP Phone Services, the number of users who are subscribed to each of the selected services, and the utilization percentage for each of the selected services	Administrators

System Reports on CUCM

In addition to user reports, the CAR tool can also generate different system reports. CAR provides system reports for managers and CAR administrators. Both managers and CAR administrators can access the QoS summary report. Only CAR administrators can access all of the other reports.

Four different QoS reports can be generated using the CAR tool:

- **Detail** reports are available for CAR administrators only. The QoS Detail report provides the QoS ratings that are attributed to inbound and outbound calls on the CUCM network for the period that was specified. You can use this report to help monitor the voice quality of all calls on a user-level basis for the entire system.
- **Summary** reports are available for managers and CAR administrators. This report provides a pie chart that shows the distribution of QoS grades that are achieved for the specified call classifications and period. The report also provides a table that summarizes the calls for each QoS grade.
- **By Gateway** reports are available for CAR administrators only. This report shows the percentage of the calls for each of the chosen gateways that meet the QoS criteria that the user chooses. Reports can be generated on an hourly, daily, or weekly basis.
- **By Call Types** reports are available only for CAR administrators. This report shows the percentage of the calls for each chosen call type that meets the QoS criteria that the user chooses. Reports can be generated on an hourly, daily, or weekly basis.

Two different types of Traffic reports are available for system reporting from the CAR tool, and all Traffic reports are available only to CAR administrators. Summary reports provide information about the call volume for a period that the administrator specified. They include only those call types and QoS voice-quality categories that were chosen. You can use this

report to determine the number of calls that are being made on an hourly, weekly, or daily basis. Summary by Extension reports provide information about the call volume for a period and set of extensions that the administrator specified. It includes only those call types and extensions that were chosen.

The next system report available to CAR administrators only is the Forced Authorization Code/Client Matter Code (FAC/CMC) reports. Three different types of reports can be generated here:

- **Client Matter Code** reports allow administrators to view the following information:
  - Originating and destination numbers
  - The date and time that the call originated
  - The call duration in seconds
  - The call classification for calls that relate to each chosen client matter code
- **Authorization Code Name** reports allow administrators to view the following information:
  - Originating and destination numbers
  - The date and time that the call originated
  - The call duration in seconds
  - The call classification
  - The authorization level for calls that relate to each chosen authorization code name
- **Authorization Level** reports allow administrators to view the following information:
  - Originating and destination numbers
  - The date and time that the call originated
  - The call duration in seconds
  - The authorization code name
  - The call classification for calls that relate to each chosen authorization level

Some other system reports also are useful to CAR administrators. The Cisco Unified Communications Manager MCID service tracks malicious calls. The Malicious Call Details report displays the details of malicious calls for a given date range. The Cisco Unified Communications Manager Call Precedence service allows authenticated users to preempt lower-priority phone calls. The PDF version of the CAR Precedence Call Summary report displays the call summary for the precedence values in the form of a bar chart, on an hour of the day, day of week, or day of month basis, for each of the precedence levels that were chosen. You can use the System Overview report to see a high-level picture of the CUCM network. The CDR Error report provides statistics for the number of error records in the CAR Billing Error table and the reason for the errors. You can use this report to determine whether CAR incurred any errors with CDR data while the CDR data was loaded. Table 28-4 summarizes the system reports available through CAR as outlined in this section.

Key Topic

Table 28-4 System Reports on the CUCM

System Report	Method of Application	User Access Allowed
QoS	Detail	Administrators
	Summary	Managers and Administrators
	By Gateway	Administrators
	By Call Types	Administrators
Traffic	Summary	Administrators
	Summary by Extension	Administrators
Forced Authorization Code/Client Matter Code (FAC/CMC)	Client Matter Code	Administrators
	Authorization Code Name	Administrators
	Authorization Level	Administrators
Malicious Call Details	CUCM MCID service	Administrators
Precedence Call Summary	CUCM Call Precedence service	Administrators
System Overview	High-level picture of the CUCM network	Administrators
CDR Error	Error records in the CAR Billing_Air table	Administrators

Device Reports on CUCM

Device reports help CAR administrators track the load and performance of Cisco Unified Communications Manager–related devices, such as gateways or conference bridges. The following device reports are available only for CAR administrators.

Three Gateway reports are available to administrators on the Cisco Unified Communications Manager. You can use the Gateway Detail report to track issues with specific gateways. This report provides a list of calls that used the specified gateways. This report can be used to review detailed information about the chosen gateways. The administrator can specify gateways by type. The Gateway Summary report provides a summary of all the calls that went through the gateways. It also provides the total number of calls and the duration for each of the categories. The categories are Incoming, Tandem, and Outgoing (Long-Distance, Local, International, Others, OnNet). The report also provides the total calls for each QoS value for each gateway in the system. The Gateway Utilization report provides an estimate of the utilization percentage for the gateways.

The Route Pattern/Hunt Pilot reports can be generated to accommodate five different levels of information. The Route and Line Group Utilization report provides an estimated utilization percentage of the chosen route and line groups. The administrator can examine the usage based on each hour of a day or by a specified number of days of the week or month. Reports are generated for each chosen route and line group. You can use the report to analyze whether the route and line group capacity is sufficient to meet the usage requirements. The Route/Hunt List Utilization report provides an estimated utilization percentage of the chosen route and hunt list. Reports are generated for each chosen route and hunt list. You can use the report to analyze whether the route and hunt list capacity is sufficient to meet the usage requirements. The Route Pattern/Hunt Pilot Utilization report provides an estimated utilization percentage of the chosen route patterns and hunt pilots. The CDR Hunt

Pilot Call Summary report displays the call details for the specified hunt pilot. This report displays only an overview of the calls for the hunt pilots; hunt member information is not included. Finally, the Hunt Pilot Detail report displays call details for a hunt pilot number or a hunt member directory number.

Two different Conference Bridge reports are available through the CAR tool. The Conference Call Details report allows the administrator to generate and view details about conference calls and conference bridges. The summary report displays the summary information for conference calls within a chosen date and time range, but it does not contain information about each individual conference participant call leg. The Conference Bridge Utilization report provides an estimate of the utilization percentage of the conference bridges. The administrator can examine the usage based on each hour of a day or by a specified number of days of the week or month.

The Voice Messaging Utilization report provides an estimate of the utilization percentage of the voice-messaging devices. Table 28-5 summarizes all the different device reports available on the CAR tool through the Cisco Unified Communications Manager.

Key Topic

Table 28-5 Device Reports on the CUCM

Device Report	Method of Application	User Access Allowed
Gateway	Detail	Administrator
	Summary	Administrator
	Utilization	Administrator
Route Pattern/Hunt Pilot	Route and Line Group Utilization	Administrator
	Route/Hunt List Utilization	Administrator
	Route Pattern/Hunt Pilot Utilization	Administrator
	Hunt Pilot Summary	Administrator
	Hunt Pilot Detail	Administrator
Conference Bridge	Conference Call Details	Administrator
	Conference Bridge Utilization	Administrator
Voice Messaging Utilization	Utilization percentage of the voice-messaging devices	Administrator

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 28-6 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 28-6 Key Topics for Chapter 28

Key Topic Element	Description	Page Number
Steps	Using DNA to Analyze Call Behavior	669
Table 28-2	CDR and CMR Service Parameters on the CUCM	672
List	Levels of CAR Users	673
List	Details Provided in CDRs	673
Steps	Steps to Export CDR/CMR Data	674
Table 28-3	User Reports on the CUCM	677
Table 28-4	System Reports on the CUCM	679
Table 28-5	Device Reports on the CUCM	680

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

CAR Tool, CDR, CDR Agent Service, CDR Repository Manager Service, CDR Repository Node, CMC, CMR, DNA, FAC

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. What are the steps to use the Cisco DNA tool?
2. List the three types of users who can access the CAR tool and the files they can access.
3. List all the CDR and CMR attributes that are included in the log file details.

CHAPTER 30

Understanding the Disaster Recovery System

This chapter covers the following topics:

Disaster Recovery System Overview: This topic will describe the features of the Disaster Recovery System. This system provides complete data backup and restore capabilities for all servers in a Cisco Unified Communications Manager cluster, Cisco Unity Connection, or Cisco Unified Communications Manager IM and Presence server.

Backup Cisco Unified Communications Solutions: This topic will describe the backup process for Cisco Unified Communications solutions.

Restore Cisco Unified Communications Solutions: This topic will describe the steps for performing a restore using the Restore Wizard.

One of the most important administrative functions is the backup and restore procedure on Cisco Unified Communications Manager, Cisco Unity Connection, or Cisco Unified Communications Manager IM and Presence server. Organizations depend on the competency of administrators to keep the IP telephony network functioning properly and to restore operations when an outage occurs for any reason. Therefore, knowing how to perform backups and setting a proper schedule and timing for the backup are crucial. Topics discussed in this chapter include the following:

- Disaster Recovery System Overview
- Backup Cisco Unified Communications Solutions
- Restore Cisco Unified Communications Solutions

This chapter covers the following objectives from the Cisco Collaboration Core Technologies v1.0 (CLCOR 350-801) exam:

- 1.4 Troubleshoot these network components in a Cisco Collaboration solution

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 30-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes."

Table 30-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Disaster Recovery System Overview	1–2
Backup Cisco Unified Communications Solutions	3–4
Restore Cisco Unified Communications Solutions	5–6

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is a capability of the Disaster Recovery System?
   - **a.** Backups can only be performed manually.
   - **b.** Backups can only be scheduled.
   - **c.** Backups require a remote SFTP server.
   - **d.** A user interface for performing backup tasks is available from any Cisco UC server.
2. Which of the following components is used to perform the actual backup and restore of the UC services?
   - **a.** Master Agent
   - **b.** Local Agent
   - **c.** Backup/Restore Agent
   - **d.** Backup/Restore Service
3. What is the first step to backing up the CUCM using the Disaster Recovery System?
   - **a.** Configure the backup scheduler.
   - **b.** Initiate a manual backup of the system.
   - **c.** Create a backup device.
   - **d.** Enable the backup service.
4. Which of the following features is used to back up licensing within the UC solutions?
   - **a.** CDR_CAR
   - **b.** UCM
   - **c.** ELM
   - **d.** SELFCARE
5. When you are restoring settings from an SFTP server, what should be the next step after you choose the Type?
   - **a.** Select the features to be restored.
   - **b.** Select the file integrity check.
   - **c.** Select the node to restore settings to.
   - **d.** Choose the SFTP server to restore settings from.

6. How many jobs does History display for previous backup or restore processes?
   a. 20
   b. 10
   c. 40
   d. 30

Foundation Topics

Disaster Recovery System Overview

The Disaster Recovery System (DRS) allows administrators to perform regularly scheduled automatic or user-invoked manual data backups. The DRS performs a cluster-level backup. In a cluster-level backup, the system collects backups for all servers in a cluster to a central location and archives the backup data to a physical storage device. The DRS restores its own settings, such as the backup device settings and schedule settings, as part of the platform backup or restore. The DRS backs up and restores drfDevice.xml and drfSchedule.xml files. When the server is restored with these files, the administrator does not need to reconfigure the DRS backup device and schedule. When performing a system data restoration, the administrator can choose which nodes should be restored. The DRS includes the following capabilities:

- A user interface for performing backup and restore tasks
- A distributed system architecture for performing backup and restore functions, including monitoring the current backup status and providing a history log
- Scheduled or manual backups
- Backups archived to a physical drive or remote SFTP server

The DRS cannot be used for migration between different Cisco Unified Communications Manager releases. Before restoring a server, the administrator should ensure that the version that is installed on the server matches the version of the backup file that should be restored. The DRS supports only matching versions of the application for restore procedures.

The DRS uses two components: the Master Agent and the Local Agent. They provide the features for the various DRS tasks. The system automatically starts the Master Agent service on each node of the cluster, but the Master Agent is functional only on the first node. The Master Agent on the subsequent nodes does not perform any functions. The Master Agent performs the following duties: stores systemwide component registration information, maintains a complete set of scheduled tasks in an XML file, and updates this file when it receives updates of schedules from the user interface. The Master Agent sends executable tasks to the applicable Local Agents, as scheduled. The Local Agents execute immediate backup tasks without delay. The Master Agent stores backup data on a local attached drive or at a remote network location. The administrator accesses the Master Agent through the DRS user interface to perform activities such as the following:

- Configuring backup devices
- Scheduling backups by adding new backup schedules

- Viewing or updating an existing schedule
- Displaying the status of executed schedules
- Performing system restoration

The server has a Local Agent to perform backup and restore functions as well. Each server in a cluster, including the server that contains the Master Agent, must have its own Local Agent to perform backup and restore functions for its server. The Local Agent runs backup and restore scripts on the server. In a cluster, the Local Agent runs backup and restore scripts on each node in the cluster. By default, a Local Agent is automatically activated on each node of the cluster.

The DRS web interface is separated into two menus: one for backup tasks and one for restore tasks. To access the DRS, select Disaster Recovery System from the Navigation drop-down list in the upper-right corner of the screen on the Cisco Unified Communications Manager, Cisco Unified IM and Presence server, or Cisco Unity Connection web interface window. You can also directly access this Disaster Recovery web interface by entering the address **https://<***server_IP***>/DRF/** in the address bar of a web browser. The DRS feature is also available in Cisco Unified Contact Center Express. You can log in to the DRS by using the same administrator username and password used for the Cisco Unified CM Operating System Administration web interface. Figure 30-1 illustrates the menu options available using the DRS.

In the Navigator drop-down menu, select Disaster Recovery System.
OR
In the Address bar enter https://<IP_Address>/drf.

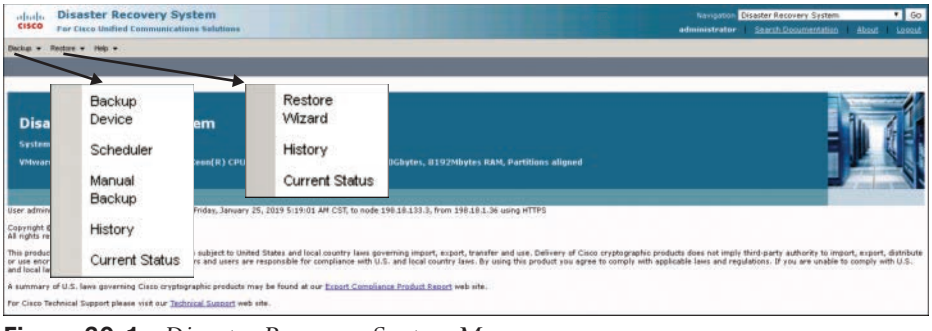

Figure 30-1 *Disaster Recovery System Menus*

For UC services, the platform, traces, syslogs, and license information will be backed up or restored. Additionally, the MOH, BAT BPS, CCM Preferences, TFTP Phone Device Profiles, SNMP Syslog Components, Cluster Management, and CEF are all backed up. The other components that are backed up or restored depend on the product type, as summarized in Table 30-2. This table identifies the differences and similarities between the backup and restore processes for each Cisco UC application.

Key Topic

Table 30-2 Disaster Recovery System Components

Cisco Unified Communications Manager	Cisco IM and Presence Server	Cisco Unity Connection
Platform	Platform	Platform
Cisco License Manager	Cisco License Manager	Cisco License Manager
Trace Collection Tool	Trace Collection Tool	Trace Collection Tool
Syslog	Syslog	Syslog
Cisco Unified CM DB	Cisco Unified Communications Manager IM and Presence Service DB	Cisco Unity Connection DB
TFTP/MOH Files	XCP Data	Mailbox Store
CDR/CAR Data	CUP Data	Greetings

Backup Cisco Unified Communications Solutions

Key Topic

The first step in creating a backup is to select a *backup device*. The backup device is a compilation of settings used to discover the database to which the backup information will be saved. Because no backup devices are created at the time of installation, you will need to create a new one. To create a new backup device from the DRS menus, navigate to **Backup > Backup Device** and click the Add New button. The DRSs use SFTP to store and retrieve these backup files, whether from a local or remote storage system. All files will be stored in the .tar format. The backup device configuration window allows you to determine the backup location. When you're using Network Directory as the desired destination for a backup, SFTP will be used. You define a Backup Device Name and enter the Host Name/IP Address, Path Name, User Name, and Password settings for the SFTP server. The Path name is the root of the SFTP software folder plus a subfolder denoted by a backslash (\). When ready, click Save. After the backup devices are created, click the Back button to view all backup devices. Figure 30-2 illustrates how to create a backup device.

The *scheduler* allows administrators to perform automatic backups in specific time frames. The backup process is resource intensive and can take longer for a larger database. It is advisable that you schedule backups during off hours or during a maintenance window. To create a scheduled backup task from the DRS menus, navigate to **Backup > Scheduler**. Then click Add New to create a new schedule. Define a Schedule Name for the scheduled backup and select the Device Name that should be used to save the backup files. Next, select the features for the backup scheduler. In Figure 30-3, CDR_CAR, IM_AND_PRESENCE, UCM, SELFCARE, and PLM are the features that can be selected in a Cisco Unified Communications Manager cluster. For Cisco Unity Connection, you can choose the following features:

Key Topic

- CONNECTION_DATABASE
- CONNECTION_GREETINGS_VOICENAMES
- CONNECTION_MESSAGES_UNITYMDXDB1
- CUC

After the features have been selected, define the time when the backup process should start. The schedule can be created to run Once or on a Daily, Weekly, or Monthly basis at a

specific time. When you're finished configuring the schedule, click Save. Once the schedule has been saved, click the Back button to go back to the schedule list. Locate the newly created backup schedule and check the box beside it. Click the Enable Selected Schedules button to activate the schedule. Figure 30-3 illustrates the settings used to schedule a backup through the DRS.

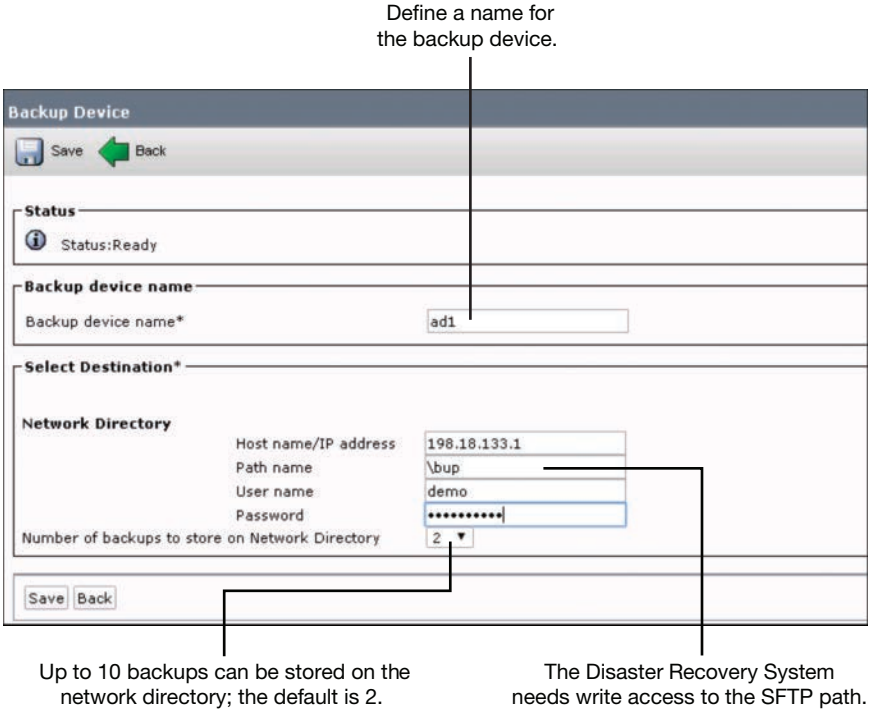

Figure 30-2 *Creating a Backup Device*

A *manual backup* is a one that starts immediately. To immediately run a manual backup, navigate to **Backup > Manual Backup**. Choose the Backup Device that should be used for the manual backup from the Device Name drop-down list. Select the features, which include the same features as a scheduled backup, and then click Start Backup.

Backup Status is a table that displays the status of a backup as each component is being backed up. When you're starting a manual backup, the status is automatically displayed, but you can navigate to **Backup > Current Status** to view the status of the manual backup if the page gets lost. As the backup procedure is being completed, you see the status of each component that is being backed up and its progress. Information about the status of each component can be viewed in the log file in the lower-right portion of the screen. Figure 30-4 illustrates how to start a manual backup and view the current status of that backup.

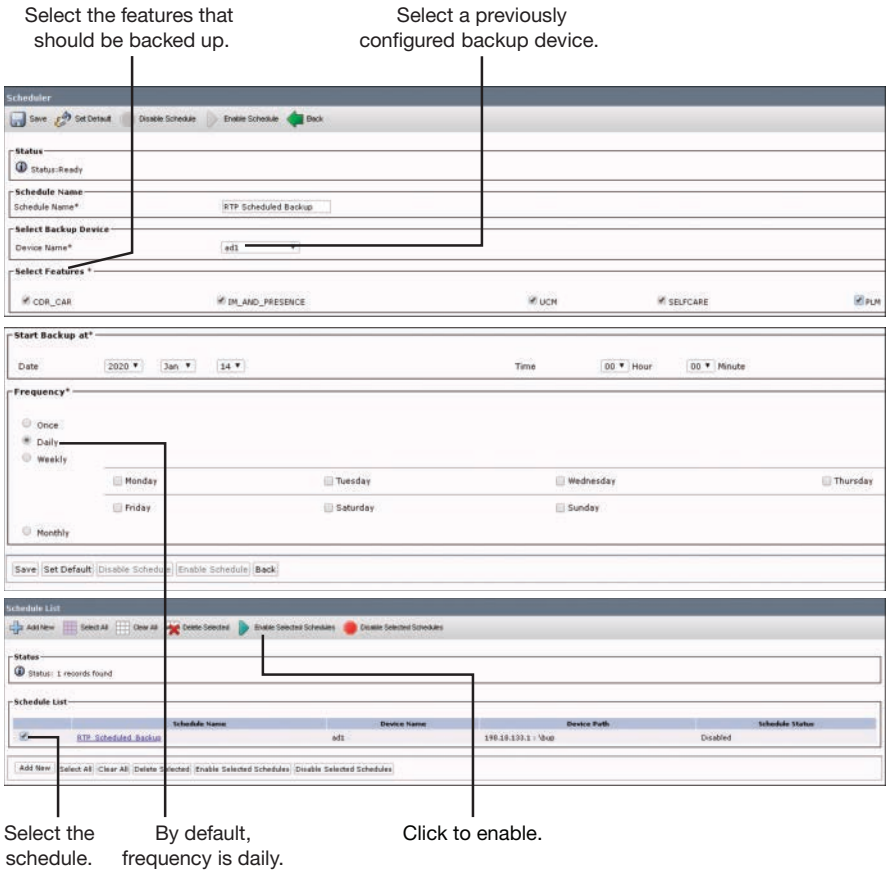

Figure 30-3 *Schedule a Backup Through the Disaster Recovery System*

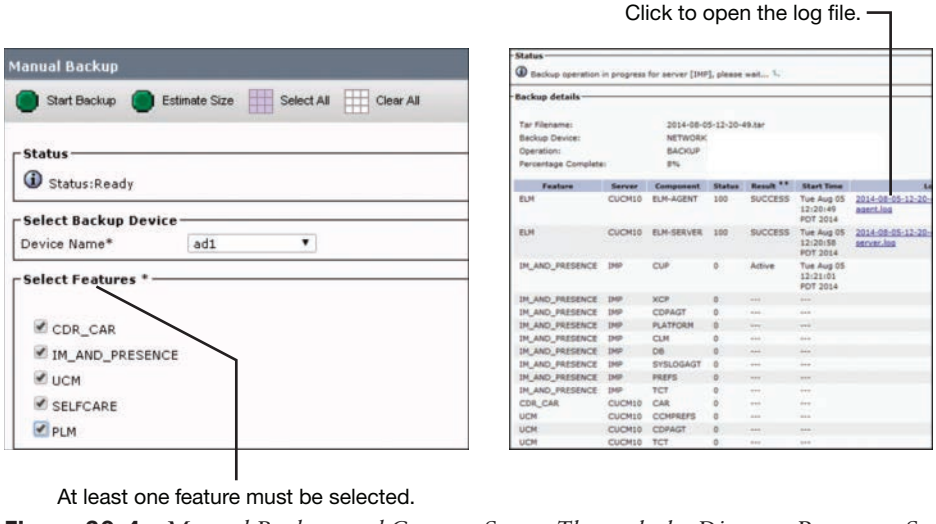

Figure 30-4 *Manual Backup and Current Status Through the Disaster Recovery System*

Restore Cisco Unified Communications Solutions

The Restore Wizard is embedded in the Backup and Restore System. It is used if a recovery from a server failure is needed. Before restoring Cisco Unified Communications Manager, the administrator must ensure that the hostname, IP address, version, and deployment type of the server being restored matches the hostname, IP address, version, and deployment type of the backup file that should be restored. To access the Backup and Restore System, navigate to **Restore > Restore Wizard** and use the following steps:

Key Topic

Step 1. Choose the Backup Device that should be used for the restore process, and then click **Next**.

Step 2. The Restore Wizard will check for valid backup files on the backup device. Select the desired backup file from the drop-down list and click **Next**.

Step 3. Once the restore device and backup file have been selected, choose the Type of restore. The type of restore is simply the features that should be restored. This selection depends on the features that are contained in the backup file. For example, if the backup file contains only the CCM feature, you cannot restore the CDR_CAR feature from this file. Once the Type has been selected, click **Next**.

Step 4. Choose file integrity check. The file integrity check is optional and is required only in the case of SFTP backups. Be aware that the file integrity check process consumes a significant amount of CPU and network bandwidth, which slows down the restore process considerably. Next, choose the server node to restore for each feature. If you choose the first node to restore the data, the DRS automatically restores the database on the subsequent nodes. Also, after you choose the node to which the data is to be restored, any existing data on that server is overwritten. When ready, click **Restore** to start the restore process. Figure 30-5 illustrates the steps required to run the Restore Wizard from the DRS.

When the restore process starts, the restore status will be displayed. As the restore procedure for a component is being completed, you can view the status as each component of each server is being restored. Information on the status of each component can be viewed in the log file in the lower-right portion of the screen. After the restore procedure has successfully completed, a reboot of the restored server is required for all changes to take effect. Even if restoring only the first node, you must restart all nodes in the cluster. Restart the subsequent nodes before restarting the first node. Figure 30-6 illustrates the restore status screen that displays after a restore is initiated from the DRS.

The history can be used to show the previous backup and restores. From the Backup History window and Restore History window, you can view the backups or restores that have been performed, including the filename, storage location, completion date, result, and features that were backed up or restored. Navigate to **Backup > History** to display the Backup History window or navigate to **Restore > History** to display the Restore History window. Each history displays the latest 20 jobs.

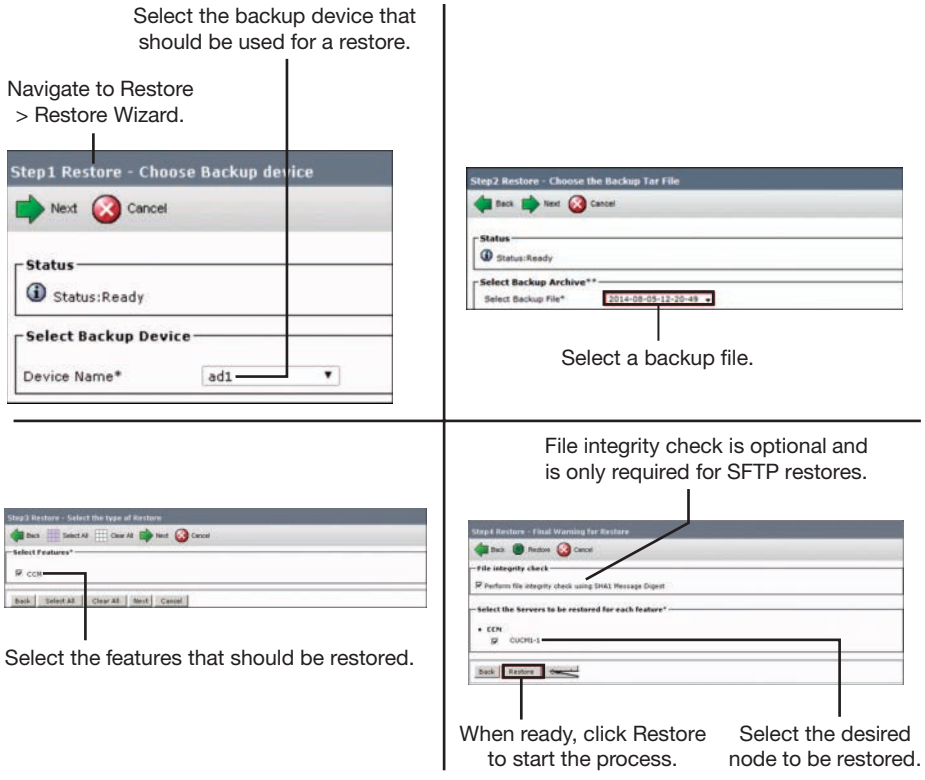

Figure 30-5 *Restore Wizard Steps on the Disaster Recovery System*

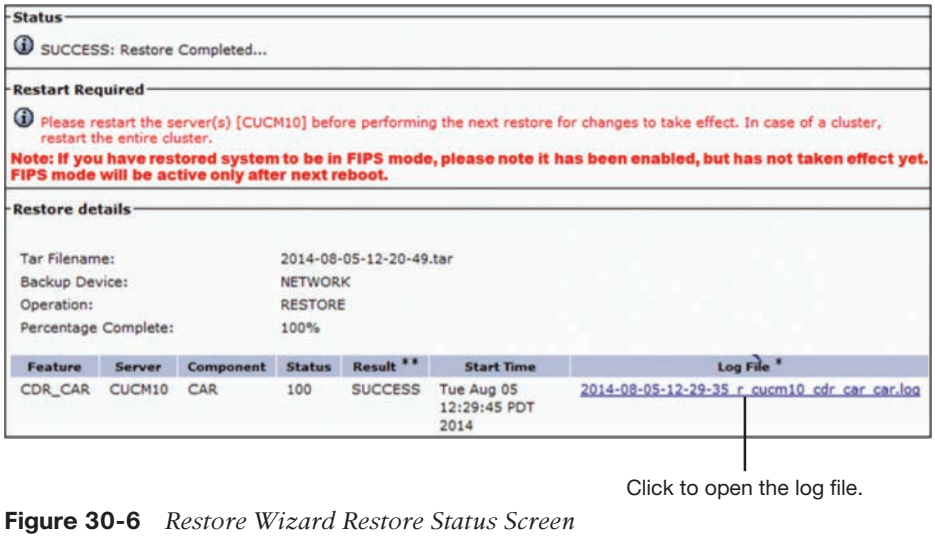

Figure 30-6 *Restore Wizard Restore Status Screen*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 32, "Final Preparation," and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 30-3 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Table 30-3 Key Topics for Chapter 30

Key Topic Element	Description	Page Number
Paragraph	Master Agent Explained	702
Table 30-2	Disaster Recovery System Components	704
Paragraph	Backup Device	704
List	Features That Can Be Backed Up	704
Steps	Steps to Run the Restore Wizard	707

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Backup Device, Backup Status, Cluster-Level Backup, Local Agent, Manual Backup, Master Agent, Restore Wizard, Scheduler

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. List the activities that can be performed in the DRS through the Master Agent.
2. List all the features that can be backed up from the CUCM.