



Campus Infrastructure Considerations

To ensure successful implementation of Cisco IP Telephony Solutions, you must first consider your LAN infrastructure. Before adding voice to your network, your data network must be configured properly.

You can use these concepts and implementation techniques regardless of whether you have a headquarters with tens of thousands of users or a small branch with fewer than a hundred users. However, the size of the network determines the actual components and platforms you will select and the details that determine the scalability, availability, and functionality of your network.

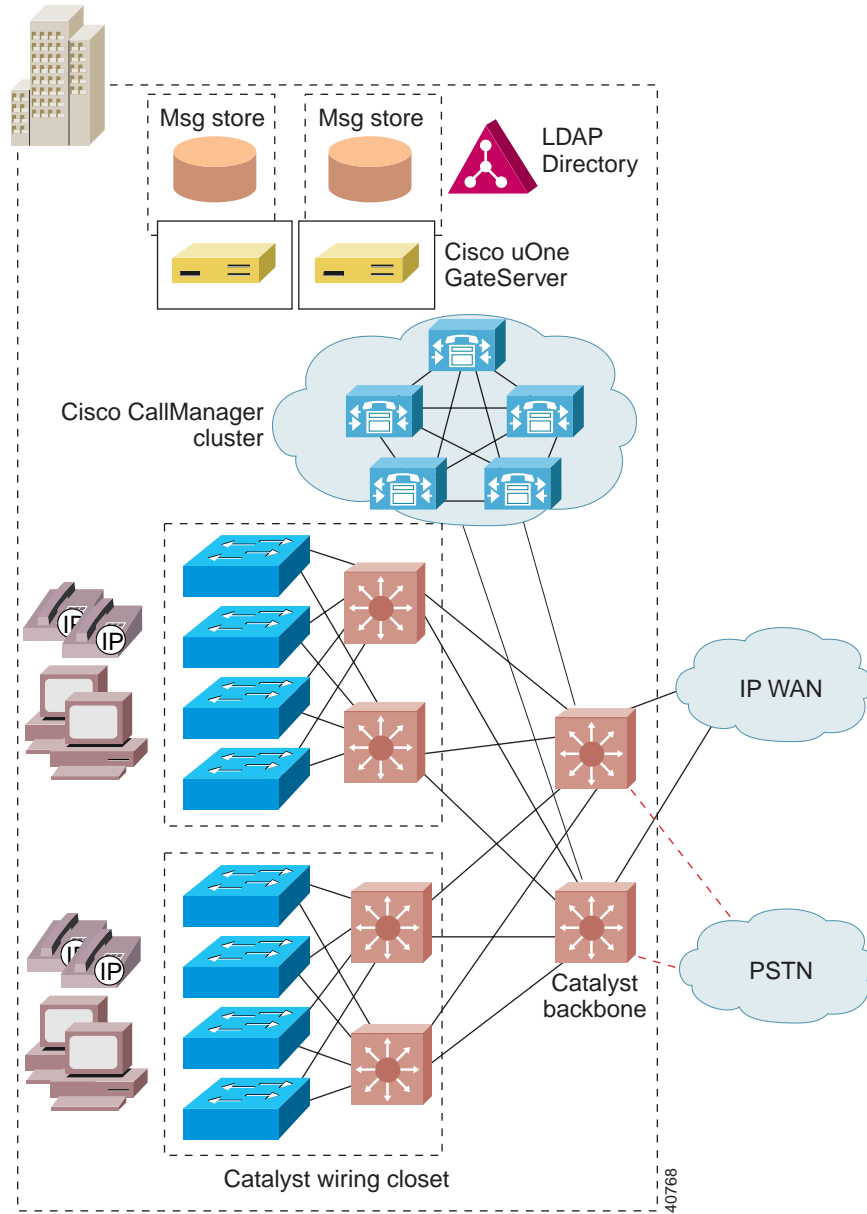
This chapter contains these sections:

- Overview, page 2-2
- Power Protection Strategies, page 2-4
- Network Infrastructure, page 2-5
- High Availability, page 2-7
- Physical Connectivity Options, page 2-9
- Power to IP Phones, page 2-10
- IP Addressing and Management, page 2-21
- Quality of Service, page 2-28

Overview

Cisco IP Telephony Solutions rely on the stable foundation of Cisco multiprotocol routers and Catalyst multilayer LAN switches, which are the building blocks in enterprise networks. Figure 2-1 illustrates a general model of a Cisco IP telephony network using these components.

Figure 2-1 Cisco IP Telephony General Deployment Model



Power Protection Strategies

Reliable power is vital to IP telephony. An uninterruptible power supply (UPS) can be used to ensure a reliable and highly available infrastructure by protecting it from power failures. Each UPS has some amount of battery that will keep the equipment running for a certain period of time. The UPS can be configured with the appropriate amount of battery for desired results.

**Caution**

Cisco strongly recommends that you provide some type of backup power for your IP telephony network. Cisco AVVID products do not ordinarily come with a backup power supply.

Here are some common strategies for using UPS:

- Back up the wiring closet switches and downstream data center using UPS. While this strategy ensures that power is maintained to the phones, wall powered devices such as PCs can still go down.
- Back up the whole building using UPS. This protects all devices and equipment from power failures. Protecting PCs in this fashion is useful because of the new breed of highly available data applications.
- Provide a separate generator for power (besides the feed from the utility company) and use it as backup. In this case you might still need to add UPS because it usually takes a few minutes for the generator to ramp up. The advantage of this strategy is that less battery time is needed for each UPS.

In addition, UPS can be configured with options such as Simple Network Management Protocol (SNMP) management, remote monitoring, alarm reporting, and so on.

Further Information

For more information on power protection, see the “Additional Information” section on page xvii.

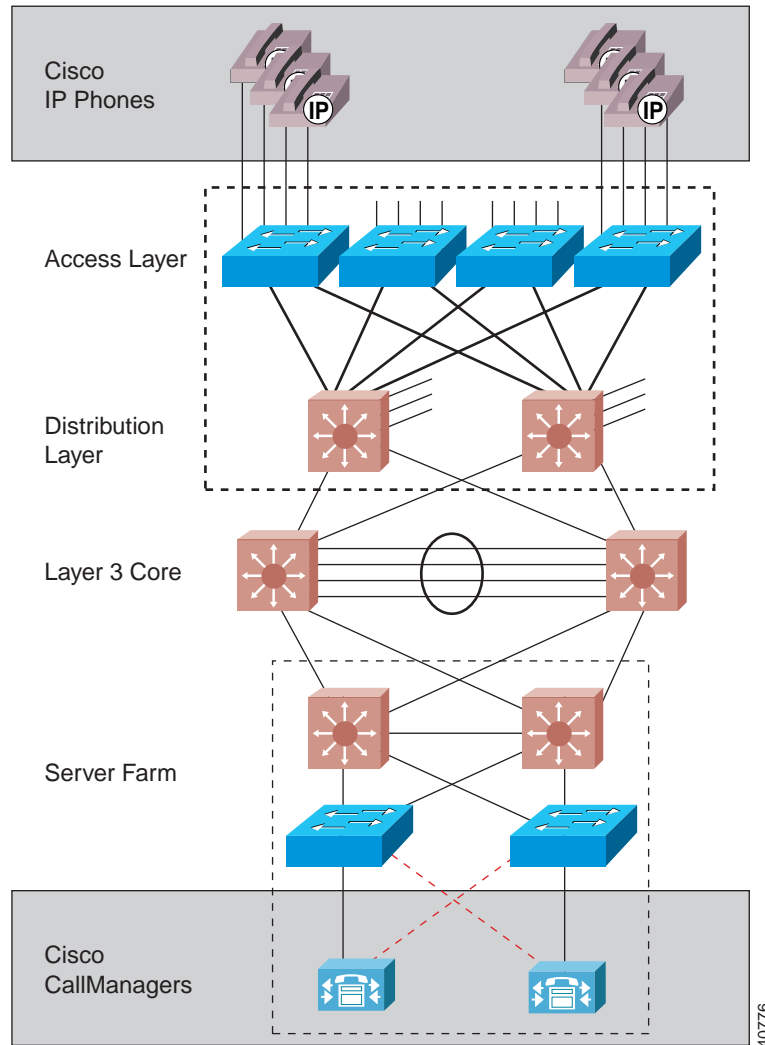
Network Infrastructure

Building an end-to-end IP telephony system requires an IP infrastructure based on Layer 2 and Layer 3 switches and routers, with switched connections to the desktop. Network designers must ensure that the endpoints are connected using switched 10/100 Ethernet ports, as illustrated in Figure 2-2.

**Caution**

Cisco does not support using hubs for shared connectivity to the switches because they can interfere with correct operation of the IP telephony system.

Figure 2-2 Switched 10/100 Ethernet Network Infrastructure



Cisco IP Phones, which are connected to the switch port, also provide connectivity for an attached computer. The phone electronics, which include a three-port switch, preserve the switched connectivity model for the computer and ensure quality of service for both the IP phone and the downstream computer.




Note

The three-port switch has two external ports and one internal port.

Figure 2-3 shows the two basic parts of the IP phone—phone circuitry and switching electronics—housed in the same package. There are two switched connections available as RJ-45 jacks: one goes to the switch in the wiring closet using a straight-through cable, and the other connects a PC or workstation. Two additional non-Ethernet connectors can be used for attaching a headset and for debugging purposes.

Figure 2-3 Cisco IP Phone Internals



High Availability

The distributed architecture of a Cisco IP telephony solution provides the inherent availability that is a prerequisite for voice networking. Cisco IP telephony solutions are also inherently scalable, allowing seamless provisioning of additional capacity for infrastructure, services, and applications.

In the world of converged networking, in contrast to the world of the PBX, availability is designed into a distributed system rather than into a box. Redundancy is available in the individual hardware components for services such as power and supervisor modules. Network redundancy, however, is achieved with a combination of hardware, software, and intelligent network design practices.

Network redundancy is achieved at many levels (see Figure 2-2). Physical connections exist from the edge devices where IP phones and computers are attached to two spatially diverse aggregation devices. In the event that an aggregation device fails, or connectivity is lost for any reason (such as a broken fiber or a power outage), failover of traffic to the other device is possible. By provisioning clusters of Cisco CallManagers to provide resilient call control, other servers can pick up the load if any device within the cluster fails.

Advanced Layer 3 protocols such as Hot Standby Router Protocol (HSRP) or fast converging routing protocols, such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP), can be used to provide optimum network layer convergence around failures.

Advanced tools are also available for the MAC layer (Layer 2). Tunable spanning-tree parameters and the ability to supply a spanning tree per virtual LAN (VLAN) allow fast convergence. Value-added features such as uplink-fast and backbone-fast allow intelligently designed networks to further optimize network convergence.

High availability of the underlying network plays a major role in ensuring a successful deployment. This translates into redundancy, resiliency, and fast convergence.


Further Information

For more information on high availability, see the “Additional Information” section on page xvii.

Physical Connectivity Options

This section describes the various ways in which IP phones and computers can be connected to the network (see Figure 2-4).

Figure 2-4 Network Connectivity Options



The first option shown in Figure 2-4 is to connect the IP phone to the switch and to connect the data device (computer or workstation) to the switched Ethernet port on the IP phone, as described in the “Network Infrastructure” section on page 2-5. This is the most common connectivity option and aids in rapid deployment with minimal modifications to the existing environment. This arrangement has the advantage of using a single port on the switch to provide connectivity to both devices. Also, no changes to the cabling plant are required if the phone is line powered (see the “Power to IP Phones” section on page 2-10). The disadvantage is that, if the IP phone goes down, the computer also loses connectivity.

The second option shown in Figure 2-4 is to connect the IP phone and the computer using different switch ports. Although this option doubles the switch port count for every user, it provides a level of redundancy for the user. If the phone goes down, the PC is not affected, and vice versa. Also, you can connect

the phone and PC to ports on different modules, thus achieving another layer of redundancy by providing protection for one of the devices if either module goes down.

The third option shown in Figure 2-4 differs from the others in that the phone is not a hardware device, but is a JTAPI application running on a computer. This option, the Cisco IP SoftPhone, could be particularly useful in environments where the need for a separate handset is minimal.

Power to IP Phones

Cisco IP Phones support a variety of power options. This section discusses each of the three available power schemes:

- Inline Power, page 2-10
- External Patch Panel Power, page 2-17
- Wall Power, page 2-20

Inline Power

The advantage of inline power is that it does not require a local power outlet. It also permits centralization of power management facilities.

With the inline power method, pairs 2 and 3 (pins 1, 2, 3, and 6) of the four pairs in a Category 5 cable are used to transmit power (6.3W) from the switch. This method of supplying power is sometimes called phantom power because the power signals travel over the same two pairs used to transmit Ethernet signals. The power signals are completely transparent to the Ethernet signals and do not interfere with their operation.

The inline method of supplying power requires the new power-enabled line card for the switch. This mechanism is currently available in the following Cisco Catalyst systems:

- Catalyst 6000 Family Switches with minimum Cisco CatOS Release 5.5 or later.
- Catalyst 4000 Family Switches (Catalyst 4006 with Power Entry Module and Auxiliary Power Shelf. Require minimum of two power supplies to power 240 ports.) Minimum Cisco CatOS Release 6.1 or higher.
- Catalyst 3524-PWR (standalone 24-port 10/100 two gigabit uplinks). Minimum Cisco IOS Release 12.0(5).XU or higher.

Figure 2-5) illustrates the new Catalyst 6000 power-enabled line card.

Figure 2-5 Catalyst 6000 Power-Enabled Line Card



Before the Catalyst switch applies power, it first tests for the presence of an IP phone. By first testing for the unique characteristics of the Cisco IP Phone and then applying power, using a low current limit and for a limited time, the Catalyst switch avoids damage to other types of 10/100 Ethernet terminating devices.

Establishing Power to the IP Phone

To establish power to the IP phone, the power-enabled Catalyst switch performs the following steps:

1. The switch performs phone discovery by sending specific tones down the wire to the IP phone. In its unpowered state, the IP phone loops these tones back to the switch.

When the switch receives this tone, it knows that the device connected is a Cisco IP Phone and it is safe to deliver power to the device. This behavior is exhibited only by Cisco IP Phones, so that other devices connected to the switch port are safe from receiving current. This hardware polling is done by the system at fixed intervals on a port-by-port basis until a LINK signal is seen or the system has been configured not to apply inline power to that port.

2. When the switch finds an IP phone by using phone discovery, it applies power to the device. The Cisco IP Phone powers up, energizing the relay and removing the loopback (normally closed relay becomes open) between transmit and receive pairs. It also sends a LINK packet to the switch. From this point, the IP phone functions as a normal 10/100 Ethernet device.

If the LINK packet is received within five seconds, the Catalyst switch concludes that the attached device is a Cisco IP Phone, and it maintains the power feed. Otherwise power is removed and the discovery process is restarted.

3. Once the Cisco IP Phone is powered and responding, the phone discovery mechanism enters a steady state. If the phone is removed or the link is interrupted, the discovery mechanism starts again. The port is checked every five seconds for a LINK packet and, in its absence, the test tone is generated.

The advantage of this mechanism is that power is supplied to the phone by the switch just as it is in a traditional telephony environment. In some installations, it is entirely possible that only two pairs have been terminated out of the four available for the data run between the wiring closet and the desktop location. In such cases the inline power method can allow customers to deploy IP telephony by using the existing cable plant without any modification.

Inline Power Configuration

The inline power method requires Catalyst software Release 5.5 for Catalyst 6000, Cisco CatOS 6.1 or higher for Catalyst 4000, and Cisco IOS Release 12.0(5)XU or later for Catalyst 3524-PWR. These software releases support all the necessary commands to enable the switch to deliver power through the power-enabled line card. You also have the option of explicitly not providing power through the line card, but the auto detection feature has the capability of determining whether an attached phone requires power or not.

Configuring the Inline Power Mode

The inline power mode can be configured on each port on the switch using the one of the following commands.

For Cisco CatOS:

```
set port inlinpower mod/port {auto | off}
```

For native Cisco IOS:

```
Switch(config-if)# power inline {auto | never}
```

The two modes are defined as follows:

- **auto**—The supervisor engine tells the port to supply power to the phone only if it has discovered the phone using the phone discovery mechanism, as described in the “Establishing Power to the IP Phone” section on page 2-12. This is the default behavior.
- **off**—The supervisor engine instructs the port not to apply power, even if it can and if it knows that there is a connected Cisco IP Phone device.

If the **set port inlinpower** command executes successfully, the system displays a message similar to

```
Inline power for port 7/1 set to auto
```

If the **set port inlinpower** command does not execute successfully, the system prints a message similar to

```
Failed to set the inline power for port 7/1
```

**Note**

The remainder of this chapter uses the Cisco CatOS command syntax. For native Cisco IOS commands, refer to the specific product documentation for the switches and line cards.

Configuring the Default Power Allocation

You can configure the default power allocation using the following command:

```
set inlinepower defaultallocation value
```

This command specifies how much power, in watts, to apply on a per-port basis. The default value of 10W is good for any currently available or planned Cisco IP Phone model. The phone has the intelligence to report to the switch how much power it actually needs (using Cisco Discovery Protocol), and the switch can adjust the delivered power accordingly, but under some circumstances you might want to reconfigure the default allocation. For example, if the switch has only 7W of available remaining power and you attach a new phone, the switch will refuse power to the phone because it initially needs to send the default 10W (even though the phone itself only requires 6.3W). In this case, you could reconfigure the default power allocation to 7W, and the switch would provide power.

If the **set inlinepower defaultallocation** command executes successfully, the system displays a message similar to

```
Default Inline Power allocation per port: 10.0 Watts (0.24 amps @42V)
```

If the **set inlinepower defaultallocation** command does not execute successfully, the system displays the following error message:

```
Default port inline power should be in the range of 2000..12500 (mW)
```

Displaying the Inline Power Status

You can display the details on the actual power consumed by using the following command:

```
show port inlinepower {mod | mod/port}
```

Here is an example display from the **show port inlinepower** command:

```

Default Inline Power allocation per port: 12.500 Watts (0.29 Amps
@42V)
Total inline power drawn by module 7: 37.80 Watts (0.90 Amps @42V)y
module 5: 37.80 Watts ( 0.90
Port      InlinePowered      PowerAllocated
      Admin Oper   Detected mWatt   mA @42V
-----
7/1 auto  off   no      0       0
7/2 auto  on    yes    12600   300
7/3 auto  faulty yes    12600   300
7/4 auto  deny  yes     0       0
7/5 on    deny  yes     0       0
7/6 on    off   no      0       0
7/7 off   off   no      0       0

```

Other Inline Power Considerations

This section briefly discusses miscellaneous issues related to inline power supply.

Power Consumption

Cisco IP Phone model 7960 consumes 6.3W. Depending upon the number of phones attached or planned, the system should be equipped with a 1300W power supply or the new power supply capable of delivering 2500W.



Note

The new power supply for the Cisco Catalyst 6000 family switches needs 220V to deliver 2500W of power. When powered with 110V, it delivers only 1300W. In addition, the power supply needs 20A regardless of whether it is plugged into 110V or 220V.

Error and Status Messages

You can configure the system to send syslog messages that indicate any deviations from the norm. These messages include the following deviations:

- Not enough power available

```
5SYS-3-PORT_NOPOWERAVAL:Device on port 5/12 will remain unpowered
```

- Link did not come up after powering up the port

```
%SYS-3-PORT_DEVICENOLINK:Device on port 5/26 powered but no link up
```

- Faulty port power

```
%SYS-6-PORT_INLINEPWRFLTY:Port 5/7 reporting inline power as faulty
```

Power status can also be displayed on a per-port basis using the **show port status** command. The command displays the following values:

- On—Power is being supplied by the port.
- Off—Power is not being supplied by the port.
- Power-deny—System does not have enough power, so the port does not supply power.

Dual Supervisors

When the system is using dual supervisors, power management per port and phone status are synchronized between the active and standby supervisor. This is done on an ongoing basis and is triggered with any change to the power allocation or phone status. The usefulness and functioning of the high availability features are unaffected by the use of inline power.

Power Protection

Cisco recommends that backup power be used for a higher degree of redundancy and availability. See the “Power Protection Strategies” section on page 2-4.

Ports and Power Supplies

Table 2-1 shows the number of IP phones that can be supported with the 1050W, 1300W, and 2500W power-enabled line cards on a Cisco Catalyst 6509 with the Policy Feature Card (PFC).

Table 2-1 IP Phones Supported with Power-Enabled Line Cards

Power Supply	IP Phones Supported at 6.3W per Phone
1050W	60 IP phones
1300W	96 IP phones (2 modules)
2500W	240 IP phones (5 modules)

External Patch Panel Power

If the switch does not have a power-enabled line card, or one is not available for the switch being used, then a Cisco power patch panel (Figure 2-6) can be used. The power patch panel can be inserted in the wiring closet between the Ethernet switch and the Cisco IP Phone.

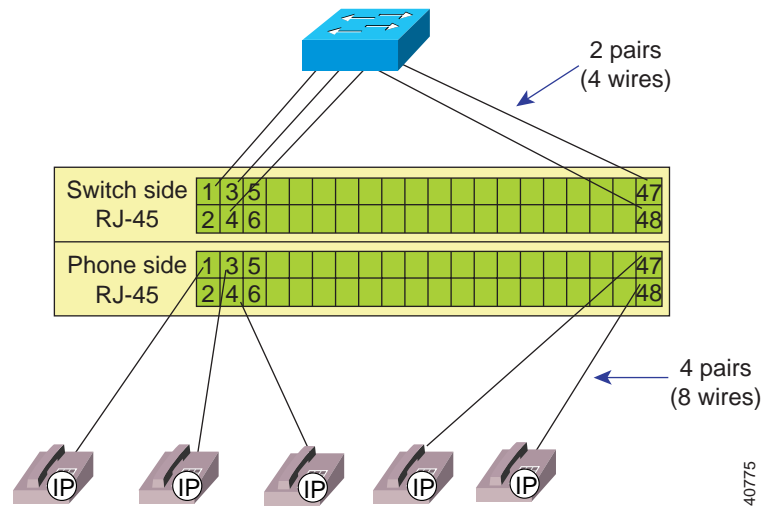
Figure 2-6 Cisco Power Patch Panel



The patch panel has a 250W power supply and draws its power from a 110 VAC source. It can accommodate 48 ports and is capable of supplying power to each of the 48 ports at 6.3W per Cisco IP Phone model 7960. We recommend an uninterruptible power supply (UPS) for backup in the event of a power failure.

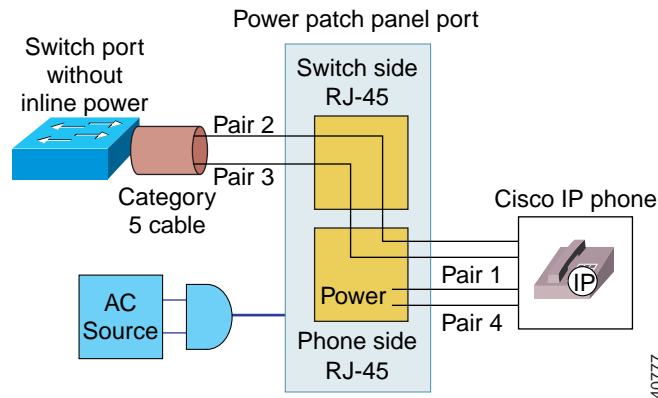
As shown in Figure 2-7, the patch panel has two ports per connection: one port on the switch side and one port on the phone side.

Figure 2-7 Power Patch Panel Connectivity to Cisco IP Phone



This arrangement of applying power to the phone uses all four pairs in the Category 5 cable. Unlike the inline method, Ethernet pairs do not carry power signals. Rather, the remaining pairs of Category 5 cable are used for delivering power from the patch panel (see Figure 2-8).

Figure 2-8 External Power Through the Power Patch Panel




As shown in Figure 2-8, pairs 2 and 3 from the switch are patched straight through to pairs 2 and 3 coming from the phone. Pairs 1 and 4 from the phone terminate at the patch panel (Ethernet does not use pairs 1 and 4) and power is applied across them to power the phone. The actual conductors used are pins 4 and 5 (pair 1) and pins 7 and 8 (pair 4) for power and ground return. This means that all four pairs in the Category 5 cable need to be terminated at the user's desk and in the wiring closet.

The Cisco power patch panel operates in discovery mode. In discovery mode, the patch panel tries to verify if the device connected to it is a Cisco IP Phone. It does this by using the phone discovery mechanism used in the inline power method, except that here the patch panel, rather than the switch, generates the test tone. Everything else about the process is identical to that described in the "Establishing Power to the IP Phone" section on page 2-12.

Wall Power

The last option is to power the Cisco IP Phone from a local transformer module plugged into a nearby outlet (maximum of 3 meters), as illustrated in Figure 2-9.

Figure 2-9 Wall Powered Cisco IP Phone



A combination of these power options can provide redundant power to the Cisco IP Phone. Internally, these three sources are combined through protection diodes, so that whatever combination is used, the phone shares the power.

Summary of Recommendations

You can purchase line cards that are capable of applying power to the IP phone. If you want to deploy IP phones with existing switches, you can either buy new line cards capable of applying power or use the external Cisco power patch panel to power the phones if powered line cards are not available for the switch. As a final option, you can use wall power to provide power to the IP phones.

IP Addressing and Management

Each IP phone requires an IP address, along with associated information such as subnet mask, default gateway, and so on. Essentially, this means that your organization's need for IP addresses doubles as you assign IP phones to users.

This information can be configured statically on the IP phone, or it can be provided by the Dynamic Host Configuration Protocol (DHCP).

The following sections describe various ways that you can meet these IP addressing requirements:

- Assigning IP Addresses Using Same Subnet as Data Devices
- Modifying the IP Addressing Plan
- Creating a Separate IP Subnet for IP Phones

Assigning IP Addresses Using Same Subnet as Data Devices

You might want to provide IP addresses to the IP phones using the same subnet as data devices. This might be a straightforward solution in your situation. However, many sites have IP subnets with more than 50% of subnet addresses already allocated. If your network fits this description, this is not the best solution for your needs.

Modifying the IP Addressing Plan

You could assign addresses for IP phones out of the existing subnets, but you must renumber the IP addressing plan. This may not always be feasible.

Creating a Separate IP Subnet for IP Phones

You can put the IP phones on a separate IP subnet. The new subnet could be in a registered address space or in a private address space, such as network 10.0.0.0. Using this scheme, the PC would be on a subnet reserved for data devices and the phone would be on a subnet reserved for voice. Configuration on the IP phone can be minimized by having the phone learn as much information dynamically as possible. Therefore, when the IP phone powers up it should get its voice subnet automatically, then send a DHCP request on that subnet for an IP address.

The automated mechanism by which the IP phone gets its voice subnet is provided through enhancements to the Cisco Discovery Protocol (CDP).

CDP Enhancements

Cisco Discovery Protocol (CDP) is a device discovery protocol that runs on all Cisco equipment. With CDP, each device sends periodic messages to a multicast address and in turn listens to the periodic messages sent by other devices. This allows devices on the network to discover one another and learn information such as protocols used, protocol addresses, native VLAN of interconnected ports, and so on. CDP is also used to send some Layer 2 and Layer 3 messages.

Cisco IP Phones use CDP to interact with the switch so that the switch knows that an IP phone is connected to it. To provide this level of support, three new fields have been added to CDP:

- Voice VLAN ID (VVID) for communicating the voice subnet to the IP phone
- Trigger field for soliciting a response from the connected device
- Power requirement field for getting the exact power requirement from the phone

VVID Field

A VLAN (Layer 2) maps to a subnet (Layer 3) as a broadcast domain, such that a VLAN is equivalent to a subnet. The VVID was introduced with release 5.5 of the Catalyst software. This is the voice VLAN that the switch assigns to the IP phone inside the CDP message. It allows the IP phone to get its VLAN ID automatically when it is plugged into the switch if a VLAN is configured for the phone (see the “Voice VLAN Configuration” section on page 2-24). If no VLAN is configured for the IP phone, the phone resides in the native VLAN (data subnet) of the switch.

Trigger Field

The trigger field is used to force a response from the connected device. Under normal circumstances, a device sends CDP update messages at a configured interval (default is one minute). If an IP phone is connected between CDP messages, it cannot receive its VVID. In this case, the IP phone issues a trigger in the CDP message it sends to the switch, forcing the switch to respond with a VVID.

Power Requirement Field

When the switch provides inline power to an IP phone, it has no way of knowing how much power the phone needs (this varies by model). Initially, the switch allocates 10W, then adjusts the delivered power according to the requirements sent by the IP phone in the CDP message.


Auxiliary VLANs and Data VLANs

The new voice VLAN is called an *auxiliary VLAN* in the Catalyst software command-line interface (CLI). In the traditional switched world, data devices reside in a data VLAN. The new auxiliary VLAN is used to represent other types of devices collectively. Today those devices are IP phones (hence the notion of a voice VLAN), but, in the future, other types of non-data devices will also be part of the auxiliary VLAN. Just as data devices come up and reside in the native VLAN (default VLAN), IP phones come up and reside in the auxiliary VLAN, if one has been configured on the switch.

When the IP phone powers up, it communicates with the switch using CDP. The switch then provides the phone with its configured VLAN ID (voice subnet), also known as the *voice VLAN ID* or *VVID*. Meanwhile, data devices continue to reside in the native VLAN (or default VLAN) of the switch. A data device VLAN (data subnet) is referred to as a *port VLAN ID* or *PVID*.

Figure 2-10 shows an IP phone and a PC in their respective VLANs.

Figure 2-10 Voice VLAN ID and Port VLAN ID



Voice VLAN Configuration

To configure the VVID from the Catalyst software CLI, use the **set port auxiliaryvlan** command. You can use this command to set the VVID on a single port, on a range of ports, or for an entire module. The following example shows how to display the command syntax:

```
Console> (enable) set port auxiliaryvlan help
Usage: set port auxiliaryvlan <mod/port>
      <vlan|untagged|dot1p|none>
      (vlan + 1..1000)
```

In the following example, the VVID is set to 222 for ports 2/1 through 2/3. When the phone powers up, the switch instructs it to register with VLAN 222.

```
Console> (enable) set port auxiliaryvlan 2/1-3 222
Auxiliaryvlan 222 configuration successful.
```

The following examples show how to display which ports are in which auxiliary VLAN:

```
Console> show port auxiliaryvlan 222
AuxiliaryVlan auxVlanStatus Mod/Ports
-----
222          222          1/2,2/1-3
Console> show port 2/1
Port AuxiliaryVlan AuxVlan-Status
-----
2.1 222          active
```

The following is an example of VVID configuration on Catalyst switches running Cisco IOS at the interface level (for example, Catalyst 3524-PWR and 2900XL):

```
interface FastEthernet0/1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan <PVID>
  switchport mode trunk
  switchport voice vlan <VVID>
  spanning-tree portfast
  switchport mode trust
```

Connecting to the Network

The following steps outline the process that takes place when an IP phone is powered up and plugged into the network:

1. The IP phone begins a CDP exchange with the switch. The phone issues a trigger CDP to force a response from the switch. That response contains the VVID for the phone.
2. If the IP phone is configured to use DHCP (the default), it issues a DHCP request on the voice subnet it got from the switch. This is the recommended mode of operation. Static addressing can be used, but it prevents mobility.
3. The IP phone gets a response from the DHCP server in the network. Along with the DHCP response, which provides the IP address to the telephone, it is also possible to supply the address of the TFTP server from which the phone gets its configuration. This is done by configuring option 150 on the DHCP server and specifying the address of the TFTP server; Cisco DHCP server supports this feature. Again, it is possible to specify the TFTP server address manually, but this would limit adds, moves, and changes, as well as remove some other benefits.
4. The IP phone contacts the TFTP server and receives a list of addresses of Cisco CallManagers. Up to three Cisco CallManagers can be specified in the list. This provides redundancy in case the first Cisco CallManager in the list is not available.
5. The IP phone now contacts the Cisco CallManager and registers itself, receiving in return a configuration file and runtime code necessary for the phone to operate. For each configuration, the IP phone receives a directory number (DN) from the Cisco CallManager to be used for calling that particular IP phone.
6. The IP phone is ready to make and receive calls.

**Note**

This process takes about 90 seconds. To speed it up, turn on portfast and turn off port channeling and trunking. This reduces the time to about 30 seconds.

Sample Addressing Plan and Recommendations

Figure 2-11 shows examples of preferred IP addressing for connecting IP phones and PCs.

Figure 2-11 Preferred IP Addressing Plans

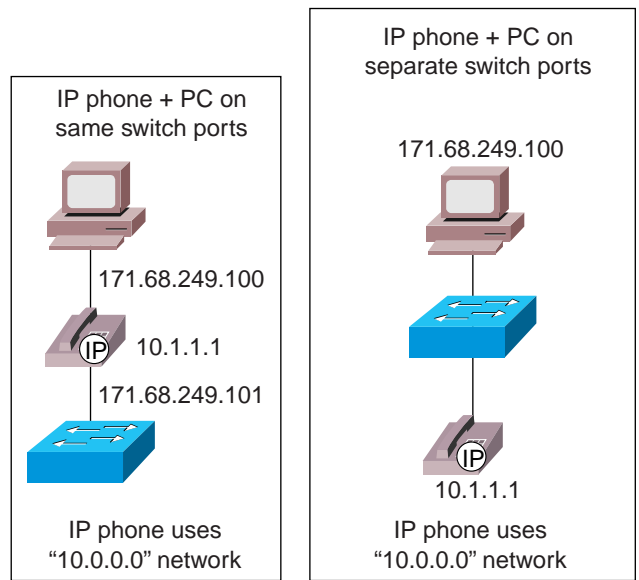



Figure 2-12 shows examples of preferred IP addressing for connecting IP phones, PCs, and Cisco IP SoftPhones.

Figure 2-12 *Optional IP Addressing Plans*



Here are some summary recommendations for IP addressing:

- Continue to use existing addressing for data devices.
- Add IP phones with DHCP as the mechanism for getting addresses.
- Use a unique range of IP addresses (for example, RFC 1918).
- Use the auxiliary VLAN feature where possible. This requires a switch capable of handling 802.1Q with enhanced software.

Quality of Service

In a converged environment, all types of traffic travel over a single transport infrastructure. Yet all traffic types are not the same. Data is bursty, loss intolerant, and not latency sensitive. Voice, on the other hand, is nonbursty and has some tolerance to loss but is latency sensitive. The challenge is in providing the required level of service for each of these traffic types.

Running both voice and data on a common network requires the proper quality of service (QoS) tools to ensure that the delay and loss parameters of voice traffic are satisfied. These tools are available as features in IP phones, switches, and routers.

See Chapter 8, “Quality of Service,” for information on WAN QoS.

Traffic Classification Types

The goal of protecting voice traffic from being run over by data traffic is accomplished by classifying voice traffic as high priority and then allowing it to travel in the network before low priority traffic. Classification can be done at Layer 2 or at Layer 3 as follows:

- At Layer 2 using the three bits in the 802.1p field (referred to as class of service, or CoS), which is part of the 802.1Q tag.
- At Layer 3 using the three bits of the differentiated services code point (DSCP) field in the type of service (ToS) byte of the IP header.

Classification is the first step toward achieving quality of service. Ideally, this step should be done as close to the source as possible, usually at the access layer of the network.

Trust Boundaries

The concept of trust is an important and integral one to implementing QoS. Once the end devices have a set class of service (CoS) or type of service (ToS), the switch has the option of trusting them or not. If the switch trusts the settings, it does not need to do any reclassification; if it does not trust the settings, then it must perform reclassification for appropriate QoS.

The notion of trusting or not trusting forms the basis for the trust boundary. Ideally, classification should be done as close to the source as possible. If the end device is capable of performing this function, then the trust boundary for the network is at the access layer in the wiring closet. If the device is not capable of performing this function, or the wiring closet switch does not trust the classification done by the end device, the trust boundary may shift. How this shift happens, depends on the capabilities of the switch in the wiring closet. If the switch can reclassify the packets, then the trust boundary remains in the wiring closet. If the switch cannot perform this function, then the task falls to other devices in the network going toward the backbone. In this case, the rule of thumb is to perform reclassification at the distribution layer. This means that the trust boundary has shifted to the distribution layer. It is more than likely that there is a high-end switch in the distribution layer with features to support this function. If possible, try to avoid performing this function in the core of the network.

In summary, try to maintain the trust boundary in the wiring closet. If necessary, move it down to the distribution layer on a case-by-case basis, but avoid moving it down to the core of the network. This advice conforms with the general guidelines to keep the trust boundary as close to the source as possible.


**Note**

This discussion assumes a three-tier network model, which has proven to be a scalable architecture. If the network is small, and the logical functions of the distribution layer and core layer happen to be in the same device, then the trust boundary can reside in the core layer if it has to move from the wiring closet.

Traffic Classification at Layer 2

Cisco IP Phones can mark voice packets as high priority using CoS as well as ToS. By default, the phone sends 802.1Q tagged packets with the CoS and ToS set to a value of 5. Figure 2-13 shows packets from the IP phone being sent as tagged frames with the 802.1p fields set to 5 and frames from the PC being sent untagged.


Figure 2-13 Frame Tagging with PVID and VVID



Because most PCs do not have an 802.1Q capable network interface card (NIC), they send the packets untagged. This means that the frames do not have a 802.1p field. Also, unless the applications running on the PC send packets with a specific CoS value, this field is zero. A special case is where the TCP/IP stack in the PC has been modified to send all packets with a ToS value other than zero. Typically this does not happen, and the ToS value is zero.

Even if the PC is sending tagged frames with a specific CoS value, Cisco IP Phones can zero out this value before sending the frames to the switch. This is the default behavior and is illustrated in Figure 2-14. Frames coming from the phone have a CoS of 5 and frames coming from the PC have a CoS of 0. When the switch receives these frames, it can take into account these values for further processing based on its capabilities.

Figure 2-14 PC Is Not Trusted



The switch uses its queues (available on a per-port basis) to buffer incoming frames before sending them to the switching engine. (It is important to remember that input queuing comes into play only when there is congestion.) The switch uses the CoS value(s) to put the frames in appropriate queues. The switch can also employ mechanisms such as weighted random early detection (WRED) to make intelligent drops within a queue (also known as congestion avoidance) and weighted round-robin (WRR) to provide more bandwidth to some queues than to others (also known as congestion management).

Example Scenario for the Catalyst 6000

Each port on the Catalyst 6000 family switches has one receive queue and two transmit queues. On the receive side, all packets go into a regular queue. Tail drop is used on this regular queue for congestion avoidance, but this mechanism comes into play *only* if there is congestion on the receive side. This is unlikely in most cases, because a frame coming in from a 10/100 Ethernet or Gigabit Ethernet port onto a 32-Gbps bus will not experience congestion.

On the transmit side, CoS values 0, 1, 2, and 3 go into the low regular queue and CoS values 4, 5, 6, and 7 go into the high regular queue. In addition, within each queue WRED can be used to make intelligent drops based on the CoS value and the percentage of buffers that are full. Finally, the high regular queue and low regular queue are serviced based on the WRR configuration. These queues are configurable; for example, they could be configured to be serviced in a 25 to 75 ratio.

**Note**

All the values for WRED, WRR, and queue size are configurable.

Cisco Catalyst 6000 family switches also support the notion of trusted and untrusted QoS on a per-port basis. This parameter is configured with the following command:

```
set port qos mod/ports.. trust { untrusted | trust-cos | trust-ipprec |
  trust-dscp }
```

This command allows you to configure the trust state as well as specify to trust CoS or ToS (**trust-ipprec**) or DSCP (**trust-dscp**), which is an emerging Layer 3 standard under the Differentiated Services working group of the Internet Engineering Task Force (IETF).

So far, this discussion has centered around the case depicted in Figure 2-14, where voice traffic comes in as CoS 5 and PC traffic is zeroed out if there is any tag. There may be times, however, when it is desirable to trust the PC CoS (if sending tagged packets) or assign a value other than zero. This can be achieved on Catalyst switches as well.

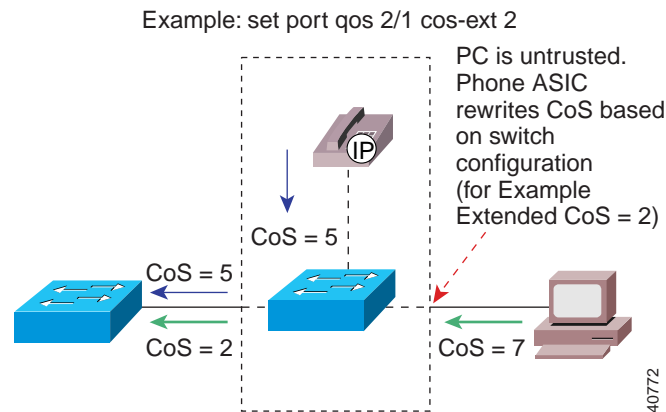
Figure 2-15 shows the case where the PC is trusted completely, and whatever CoS it presents is honored.

Figure 2-15 PC Is Trusted



Figure 2-16 shows a different case in which the PC is not trusted completely, yet it gets a level of service higher than it would with CoS=0. This is achieved by extending a specific CoS value to the PC traffic.

Figure 2-16 PC Is Not Trusted but Gets a Non-Zero CoS



Note

All of the previously discussed configurations can be used on any Catalyst switch that runs Cisco CatOS or native Cisco IOS software (for example, Catalyst 3524XL).

QoS Commands

Three commands are available for specifying classification and trust boundary:

- **set port qos *mod/ports* trust {untrusted | trust-cos | trust-ipprec | trust-dsep}**
Defines the trust boundary.
- **set port qos *mod/ports* {trust-ext | trust-cos}**
Extends the trust boundary to the PC.
- **set port qos *mod/ports* cos-ext *value***
Sets a defined CoS to the traffic from the PC.

Traffic Classification at Layer 3

Using the 802.1p bits within the 802.1Q tag provides the desired QoS results at Layer 2. When traffic has to cross a Layer 3 boundary, however, it becomes imperative to implement these mechanisms using Layer 3 parameters, such as the 3 IP precedence bits (commonly referred to as ToS) or the new DSCP parameter, which uses the six most significant bits within the ToS byte of the IP header. Traffic crosses a Layer 3 boundary when packets are routed between subnets by Layer 3 switches or routers. Traffic also crosses a Layer 3 boundary when packets need to go out of the campus network onto the WAN through edge routers. When this happens, Layer 2 classification does not help. Layer 3 classification is needed for achieving the desired level of QoS. All of the QoS techniques employed by the routers (including the very important WAN QoS) rely on Layer 3 classification.

Layer 3 classification can be achieved by using the appropriate platforms in the campus. Beginning with the IP phones, packets are already presented to the switch with CoS = ToS = 5. This Layer 3 classification is preserved even if the packets travel all the way through to the WAN edge router where the Layer 2 header is removed. So, if the trust boundary is at the source (IP phone), voice traffic has the ToS bits set to 5 and is presented to the network devices for appropriate treatment. WAN routers can use this classification to employ any of the queuing techniques. If the trust boundary is not at the source and packets need to be reclassified, then the device performing this function should be capable of doing it at Layer 3 before it can cross a Layer 3 boundary.

Layer 3 Traffic Classification on the Cisco Catalyst 6000

Cisco Catalyst 6000 family switches equipped with the Policy Feature Card (PFC) perform Layer 3 traffic classification by default when the port is trusted. Thus if a packet comes into a trusted port with CoS = 5, the switch takes this value and resets the ToS bits to 5 as well. No additional configuration is needed. If the port is untrusted, the packet gets a default CoS at the input port.

Then you can configure a QoS access control list (ACL) on the switch and rewrite the ToS to a desired value based on some matching criteria. For example, the following command sets a ToS of 5 for all packets coming from subnet 10.1.1.0 and destined to any address.

```
Console> (enable) set qos acl ip TEST dscp 40 10.1.1.0 0.0.0.255 any
```

QoS ACLs can also include Layer 4 information for classifying individual applications. Cisco Catalyst 6000 family switches are also capable of policing traffic based on Layer 3 addresses and Layer 4 port numbers. For example, you can police individual HTTP flows to 1 Mbps and aggregate all HTTP flows to 25 Mbps.

The following are important points in regard to QoS functionality on the Cisco Catalyst 6000 family switches:

- By default, QoS is not enabled. Use **set qos enable** to enable QoS on the switch.
- By default, ports are not trusted. Use the following command to enable trust on a port:

```
set port qos mod/ports.. trust { untrusted | trust-cos | trust-ipprec |  
trust-dscp }
```

- QoS configurations can be applied on a per-port basis or on a per-VLAN basis. This works very well for IP telephony implementations where phones are on a separate VLAN, as described in the “IP Addressing and Management” section on page 2-21.
- By default, Cisco Catalyst 6000 family switches map CoS to ToS when the port is trusted or by using QoS ACLs.

**Tips**

If the trust boundary happens to be on a wiring closet switch that is not capable of reclassifying at Layer 3, you can shrink the trust boundary to the distribution layer where a Layer 3 capable device is more likely to be present.

Summary of Capabilities and Recommendations

Table 2-2 briefly summarizes the capabilities within the Cisco Catalyst switch families.

Table 2-2 Summary of QoS Capabilities on the Cisco Catalyst Switch Family

Platform	Ability to Trust	Reclassify CoS	Reclassify ToS	Congestion Avoidance (WRED)	Priority Queue	Multiple Queues	Congestion Management (WRR)	Policing
Catalyst 6000	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
Catalyst 5000	No	Yes	Yes ¹	Yes	No	No	No	No
Catalyst 4000	No	Yes	No	No	No	Yes	No	No
Catalyst 3500	Yes	Yes	No	No	Yes	Yes	No ²	No

1. With additional configuration
2. Round robin only



Note

Currently the only Cisco LAN switches that support a minimum of two queues and that can guarantee voice quality are the Cisco Catalyst 8500, Catalyst 6XXX family, Catalyst 4XXX family, Catalyst 3500XL, and Catalyst 2900XL.

Here are some summary recommendations for QoS implementation:

- Create a trust boundary at the network edge in the wiring closet. Make ports trusted on the wiring closet switch where IP phones are attached.
- Reclassify ToS at the edge if devices cannot be trusted.
- Shrink the trust boundary to the distribution layer and reclassify ToS there if reclassification is not possible at the edge.
- Use a priority queue if possible for delay-sensitive traffic.

- Use QoS ACLs for more granular classification of packets using Layer 4 information.
- Use policing if necessary to limit traffic for individual flows as well as aggregate flows.
- Have traffic going to the WAN edge classified at Layer 3 so that the router can use it for advanced WAN queuing mechanisms.
- Use a WAN edge router as the classifier for very small remote site networks where a Layer 3 capable switch is not available.

