White Paper

Cisco PacketCable Primer

Introduction

This white paper describes the architecture of PacketCable 1.1 and Cisco Systems implementation of this architecture. The paper's intent is to give customers insight into the PacketCable architecture and how Cisco® can help them achieve the benefits of a PacketCable network.

Overview

Many cable operators today are deploying VoIP services using their hybrid fiber-coaxial (HFC) cable infrastructures that support high-speed (HSD) data services. Costing less to operate than traditional circuit-switched telephony solutions, VoIP over cable allows cable service providers to offer value-added services that take advantage of the integration of voice and data networks. Data over Cable Service Interface Specification (DOCSIS) 1.1 certifications of multiple cable modem termination systems (CMTSs) and multimedia terminal adapters (MTAs), solidification of PacketCable specifications, and recent qualifications—combined with top-tier multiple system operators (MSOs) announcing various VoIP trials—all underscore the fact that technology to deliver VoIP is ready to be deployed

What Is PacketCable?

For VoIP over cable to be widely adopted, it needs to meet or exceed the quality, security, and features of the traditional public switched telephone network (PSTN). To that end, the PacketCable architecture was developed.

The PacketCable specifications define a comprehensive approach to cable telephony. Led by CableLabs, PacketCable is an industry initiative to develop specifications for delivering advanced, real-time multimedia services over a two-way cable plant. This architecture addresses issues such as device interoperability, product compliance with international standards, and the business and high-level technical requirements for the PacketCable architecture.

Taken as a whole, the PacketCable specifications define a complete architecture for providing voice services over a DOCSIS 1.1 network.

PacketCable Architecture

The PacketCable working group of CableLabs has created specifications that address certain areas of a full-featured, residential VoIP architecture. This document covers the following major areas of functionality:

- Event messaging and billing
- Dynamic quality of service (DQoS)
- PSTN interconnect
- Provisioning
- Lawful intercept
- Security
- Reliability and redundancy

Figure 1 shows a high-level overview of the PacketCable architecture. The system uses IP telephony technology to provide high-quality transport of voice over an IP network. Residential gateways in the form of MTAs embedded in a cable modem (embedded MTA [EMTA]) provide access at the customer premises. By plugging a standard analog telephone into the MTA device, a user can make phone calls to another MSO's customer directly across the IP network or to anyone outside the network through a media gateway.

Figure 1
Overview of Cisco's Implementation of the PacketCable Architecture

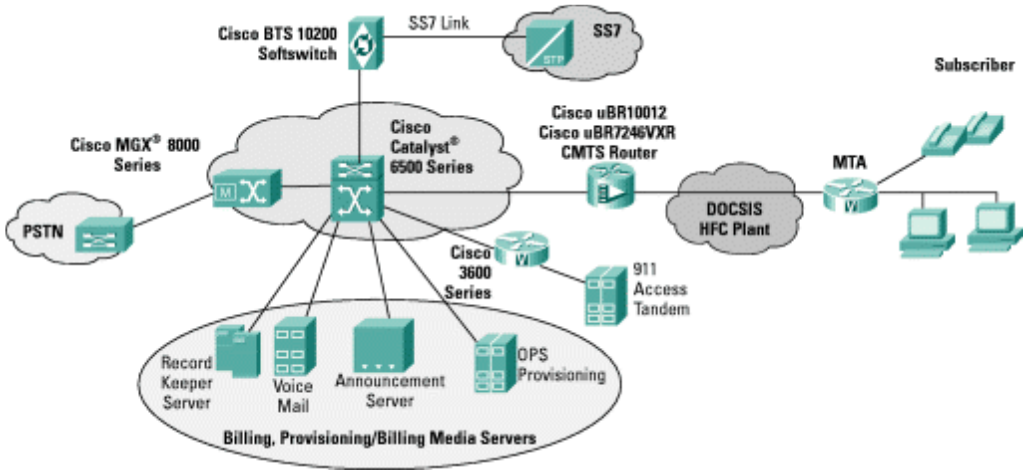

Call management servers (CMSs) and media gateway controllers (MGCs) provide centralized call-control processing by passing control information and setting up connections between residential MTAs. When these connections are established, voice passes directly between gateway endpoints in the form of Real-Time Transport Protocol (RTP) packet streams, as shown in Figure 2. Most connections with the PSTN are through ISDN User Part (ISUP) trunks with a media gateway providing the bearer connections and a signaling gateway providing the signaling connection into the Signaling System 7 (SS7) network. For some specialized requirements, for example, operator devices, MultiFrequency (MF) or channel associated signaling (CAS) trunks are provided.

Figure 2
PacketCable Signaling Architecture

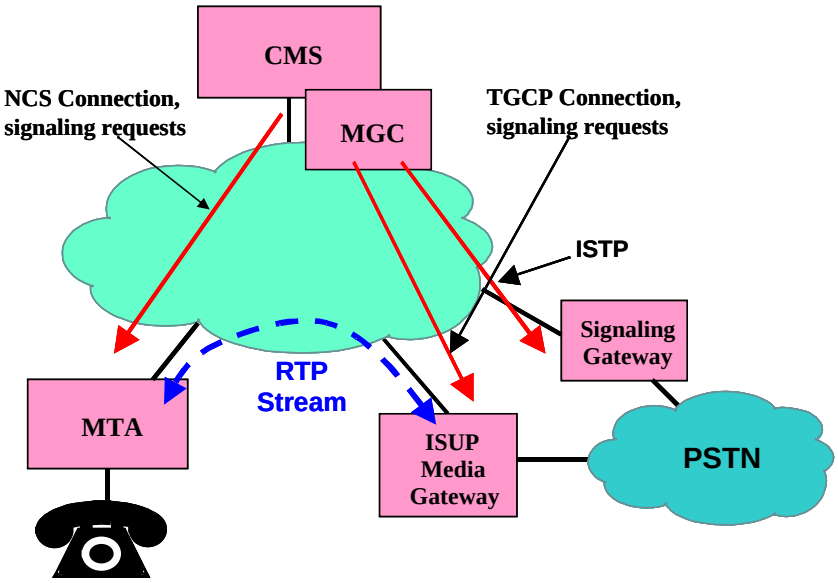

The PacketCable Network-Based Call Signaling (NCS) [1] protocol is used to communicate with the MTA endpoints. The PacketCable PSTN Gateway Call Signaling Protocol (TGCP) [2] is used to communicate with media gateways. NCS and TGCP are profiles of the Media Gateway Control Protocol (MGCP) [3], which belongs to the xGCP suite of protocols. These protocols allow a central call-control mechanism to control customer-premises-equipment (CPE) devices for voice services. The Simple Gateway Control Protocol (SGCP) [4] is another protocol included in this suite.

Functional Components

The following sections describe the functional components of the PacketCable architecture. These functional components generally apply to one or more PacketCable specification(s

Dynamic Quality of Service (DQoS)

Quality of service (QoS) is one of the most important aspects of a residential voice service. To ensure that a service offering is acceptable to an MSO's customers, the quality must be as good as or better than the existing voice service that is, the PSTN.

For this purpose, PacketCable has defined dynamic QoS (DQoS). The DQoS architecture focuses on the access portion of the network, specifically between the MTA and the CMTS.

The major concepts of the DQoS specification follow:

- PacketCable DQoS gates
- A *gate* is a logical entity that resides on a CMTS. A gate consists of a packet classifier, a traffic policy, and an interface to an entity that gathers statistics and events (the record keeping server [RKS]). A gate can ensure that only those sessions that have been authorized by the service provider receive high-quality service. This is accomplished by using the gate ID as an authorization token when requests are made to reserve or commit the flow.
- Billing based on the use of resources
- PacketCable use of DOCSIS QoS is accounted for on a per-session basis.
- Two-phase (reserve-commit) QoS activation model
- In the two-phase model, the application reserves the resource and then later commits it. As in the DOCSIS model, resources that are reserved but not yet committed are available for temporary assignment to other (for example, best effort) service flows. Billing is started when the resources are actually committed.
- Prevention of (minimizing) abusive QoS usage
- Two types of abusive QoS usage are identified: that which is accurately billed but leads to denying service to others, and that which is not accurately billed and leads to theft of service.
- Provisioning of admission control mechanisms for both upstream and downstream directions in the DOCSIS network
- Both upstream and downstream QoS should be subject to per-session admission control.

DQoS uses several protocols and components to establish a guaranteed service. These protocols include DOCSIS 1.1, the Resource Reservation Protocol (RSVP) (not necessary for EMTA), the Common Open Policy Services (COPS), and PacketCable Event Messaging.

Components involved in this architecture follow:

1. Cable modem—The cable modem is responsible for classifying, policing, and marking packets onto the proper service flows after these flows are established by appropriate signaling protocols.
2. CMTS—The CMTS is responsible for allocating and scheduling upstream and downstream bandwidth in accordance with MTA requests and QoS authorizations established by the CMS. The CMTS acts as a policy enforcement point (PEP) as defined by the COPS [5].

3. CMS—The CMS performs services that permit MTAs to establish multimedia sessions. The term *gate controller* refers to the portion of the CMS that performs the QoS-related functions. In the DQoS architecture, the gate controller controls the operation of the gates implemented on a CMTS. The gate controller acts as a policy decision point (PDP) as defined by the COPS [5].
4. RKS—The RKS coordinates the event messages, specifically DQoS-related messages, as described in the section "Event Messaging."

The early stages of the PacketCable architecture focuses on embedded MTAs only, and the following describes the relevant messages for this scenario (refer to Figure 3).

Figure 3
PacketCable DQoS Model

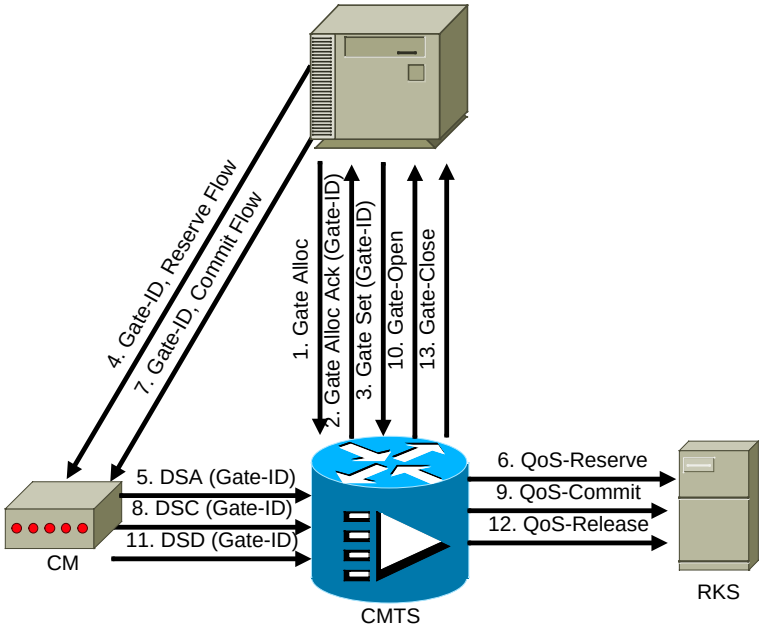

1. After the appropriate voice signaling, the gate controller portion of the CMS requests a gate from the CMTS, via the Gate-Alloc message. This Gate-Alloc message is transported using COPS.
2. If the EMTA is using less than the maximum number of gates allowed for the subscriber, the CMTS replies with a Gate-Alloc-Ack message, which contains the gate ID of the gate to be used for this session.
3. If the proper resources can be reserved on both the terminating CMTS and the backbone, the CMS issues a Gate-Set message to the CMTS to authorize the gate.
4. At this point, the CMS sends the Gate-ID message to the EMTA via voice signaling (NCS) and instructs the EMTA to reserve the flow.
5. The EMTA in turn sends a DOCSIS 1.1 Dynamic Service Add (DSA) request message to the CMTS that includes the gate ID.

6. The CMTS sends the QoS-Reserve event message to the RKS in order to audit the start of resource reservation. Event messages are discussed further in the section "Event Messaging." Event messages are transported using Remote Authentication Dial-In User Service (RADIUS).
7. After the terminating end picks up the phone, the CMS sends a signaling message to the EMTA to inform the EMTA that the flow should be committed.
8. The EMTA sends a DOCSIS 1.1 Dynamic Service Change (DSC) message to actually commit the resources.
9. At this point, the CMTS sends the QoS-Commit event message to the RKS in order to audit the start of resource commitment.
10. The CMTS also sends a Gate-Open message to the CMS to inform the CMS of the gate state change.
11. When the call is terminated, the EMTA sends a Dynamic Service Delete (DSD) message to remove the DOCSIS service flow.
12. The CMTS sends a QoS-Release event message to the RKS, and the auditing is terminated.
13. The CMTS also sends a Gate-Close message to the CMS to inform the CMS of the gate state change.
14. For more details on the PacketCable DQoS model, refer to the PacketCable Dynamic Quality-of-Service Specification [6].

Event Messaging

Event-based messaging is used in the PacketCable architecture as the means by which to convey network usage and activities. These messages, when correlated by a RKS, comprise a call detail record (CDR). By sending this generated CDR to a back-office billing server, MSOs can bill the subscriber appropriately for using network resources.

Event-based messaging has been designed as part of the PacketCable architecture because several different devices are involved in granting network resources to the subscriber, thus the call-state and network-usage information does not reside in only one device. Using event messages enables interoperation of devices from different vendors with RKSs from additional vendors. An added benefit of using event-based messaging is that network resources can be monitored in real time, because event messages are sent in real time. Network resource use does not have to complete before the usage is reported. Event messages are designed to be flexible enough that services other than voice can be monitored and billed for; however, PacketCable currently has defined only event messaging that pertains to voice services.

Figure 4
PacketCable Event Messaging Devices

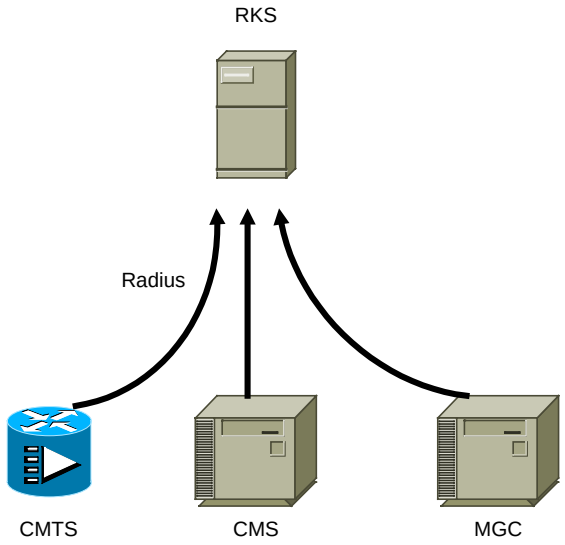

Figure 4 shows the interaction between event-message-producing devices and the RKS. Devices that create event messages include the CMTS, the CMS, and the MGC. These devices send the event messages via RADIUS. The RKS coordinates the event messages it receives, and sends them to the appropriate billing server using the appropriate protocol. This protocol is dependent upon the billing server being used, but the File Transfer Protocol (FTP) is a common protocol that has been used in this fashion.

The most commonly used event messages include the following:

Event	Message
Signaling_Star	The CMS or MGC sends this event message when call signaling starts. Technically this event message is sent after the digits are collected, to prevent event messages from being sent when the user has not actually made a call; for example, when the user picks up the phone and then puts it back on hook, without dialing any digits.
QoS_Reserve	The CMTS sends this event message when it reserves bandwidth on the access network
Call_Answer	The CMS or MGC sends this event message as an indicator that the media connection is open because the call has been answered
QoS_Commit	The CMTS sends this event message when it commits bandwidth on the access network.
Interconnect_Start	The MGC sends this event message when the start of network interconnect occurs. This is typically when a call is routed to or from the PSTN network.
Call_Disconnect	The CMS or MGC sends this event message when the media connection is closed because either the calling party or the called party terminated the call.
Signaling_Stop	The CMS or MGC sends this event message when call signaling is complete.

Event	Message
Interconnect_Stop	The MGC sends this event message when the bandwidth between the PacketCable network and the adjacent network (usually PSTN) is terminated
QoS_Release	The CMTS sends this event message when it has released its bandwidth commitment on the access network.

For more information about event messaging as it is used in the PacketCable architecture, refer to the PacketCable Event Messages Specification [7].

PSTN Interconnect

PSTN connectivity is extremely important to any residential VoIP service. In the early stages of any deployment, most calls are made to people who do not reside on the MSO's network. The number of calls to the PSTN decreases as the MSO's presence (and service subscription rates) increases, but for the foreseeable future, there will always be the need to connect to endpoints on the PSTN.

The PacketCable architecture has distributed the PSTN interconnect architecture among four distinct components, the CMS, the MGC, the signaling gateway, and the media gateway. Each component has a specific role:

- The CMS routes calls that originate from a subscriber on the MSO's network to the MGC. If a provider has multiple MGCs, the CMS must have the routing logic to know which MGC to send the call to. In the case where the MGC routes a call from a PSTN subscriber to a subscriber on the CMS, the CMS has to continue routing the call to the final destination.
- The MGC is in complete control of the media gateways. The MGC knows the state of each trunk in each trunk group. The MGC can route calls more efficiently by knowing which trunks are most actively used, which trunks are least actively used, and the charges for different providers. The MGC can even load share calls across different providers if interconnection costs are equivalent. It can also help track service-level agreements (SLAs) with providers and begin routing to other providers when the SLA for a particular provider is close to being fulfilled. The benefit of the MGC is that it alleviates the need for the CMS to store any of this dynamic information.
- The signaling gateway does the out-of-band call signaling to and from the PSTN network. What is convenient about separating this device from the other components is that usually the out-of-band signaling protocol is SS7, and this protocol has many different variants, usually depending on the country in which it is being used. This setup allows for country-specific signaling gateways to be used when necessary, without disruption to the rest of the components.
- The media gateway bridges the IP network and the analog PSTN network. Media gateways generally have banks of digital signal processors (DSPs) that do the conversion from digital to analog on the bearer paths.

Figure 5 shows the protocols used between the different devices.

Figure 5
PacketCable PSTN Interconnect Architecture

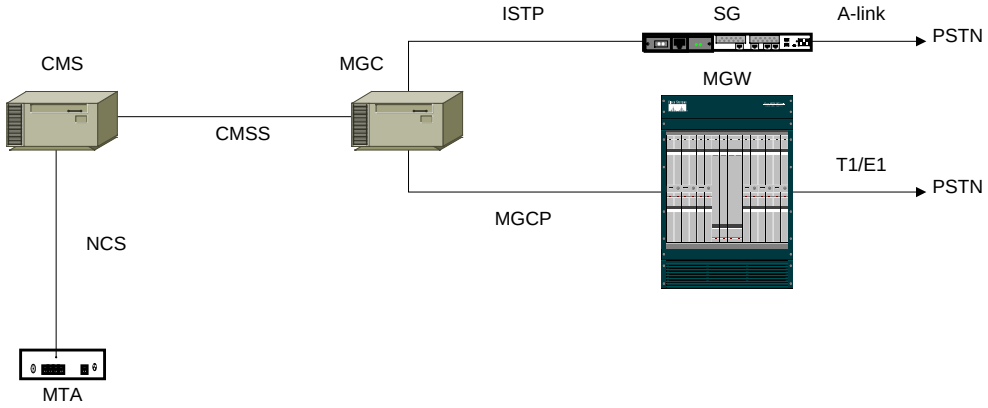

In a residential VoIP network, calls fall into two classifications, on net and off net. An on-net call is a call that remains on the MSO's network, between two MSO subscribers. An off-net call is a call that must interconnect with the PSTN to reach a particular subscriber.

In the case of an on-net to off-net call, it is the responsibility of the CMS to take an incoming NCS call from MTA endpoints and route those calls, via the PacketCable CMS to CMS Signaling Protocol (CMSS) to the MGC. This protocol is essentially the Session Initiation Protocol (SIP) [8] with some extensions. For more information about these extensions, refer to the CMSS specification [9]. The MGC is then responsible for routing those calls to the PSTN by using the PacketCable Internet Signaling Transport Protocol (ISTP) [10] to communicate with the signaling gateway, while controlling the media gateway via TGCP to open the appropriate bearer circuit. In the case of an off-net to on-net call, it is the responsibility of the signaling gateway to accept incoming SS7 traffic and communicate the incoming call to the MGC. The MGC controls the correct circuit on the media gateway and signals an incoming call to the CMS. It is then the responsibility of the CMS to control the destination MTA via NCS to participate in the incoming call.

Provisioning

In order for any solution to scale above 50–100 users, a solid provisioning system needs to be established. PacketCable has defined the provisioning flows for two different areas, one for provisioning MTA end-user devices, and the other for provisioning the CMS. However, only the MTA provision specification is currently required by PacketCable.

MTA Provisioning

MTA provisioning configures the MTA with the information necessary (for example, CMS fully qualified domain name (FQDN) and CMS User Datagram Protocol (UDP) port) to successfully operate with the PacketCable network, in addition to performing normal DOCSIS provisioning, that is, Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), or time of day (TOD), for the cable modem. The provisioning process is probably the most complex sequence of events in PacketCable. Figure 6 shows a sample flow for an EMTA.

Figure 6
Simplified PacketCable MTA Provisioning Flow

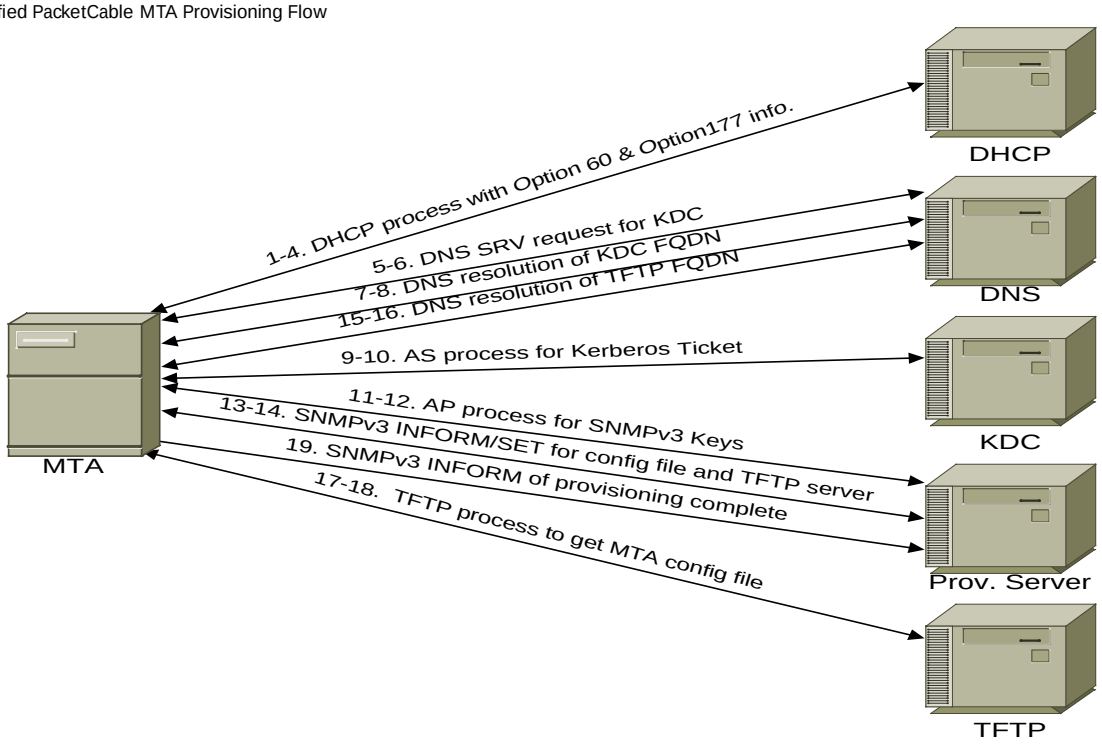

When the EMTA is powered on, the cable modem first performs normal DOCSIS 1.1 provisioning as specified in [11]. There is one difference, however: the cable modem requests and receives option 122 in the DHCP process. This option includes the address of the DHCP primary and backup servers used for the MTA. These are option 122 suboptions 1 and 2, respectively. The MTA uses this information to decide if DHCP responses it receives are from the correct DHCP server. This allows for the high-speed data and telephony providers to be different entities. Figure 6 shows the sequence of events for the subsequent MTA initialization. Note that the servers in Figure 6 can all be running on one host if desired; they are separated in the figure only for clarity.

1. The MTA sends a DHCP DISCOVER, and in the DHCP DISCOVER the MTA has populated DHCP option 60, which informs the DHCP server that it is a PacketCable MTA.
2. The MTA receives a DHCP OFFERs from one or more servers. The MTA checks the responses against the option 122 information that it received from the cable modem. The MTA uses the OFFER (s) from only the DHCP server(s) in that list. Also, for a DHCP OFFER to be valid, it must contain option 122 suboption 3, the provisioning server address.
3. When the MTA determines that it has a valid offer, it sends a DHCP REQUEST to the appropriate DHCP server.
4. 4The DHCP server completes the DHCP process by sending a DHCP ACK. During this process, the MTA has been provided option 122 suboptions 3, 4, and optionally 5, 6, and 7. These are the provisioning server, primary Domain Name System (DNS), secondary DNS, Kerberos realm, and the ticket-granting-ticket (TGT) flag, respectively. Suboption 7 (the TGT flag) is rarely used, so it is not explained here.

5. The MTA then uses option 122 suboption 6, the Kerberos realm, to perform a DNS SRV request to the DNS server(s) specified in option 122 suboptions 4 and 5.
6. The DNS server returns the FQDN of the key distribution center (KDC) used for the specified Kerberos realm.
7. The MTA then issues a DNS request to resolve the FQDN of the KDC.
8. The DNS server returns the IP address of the KDC in question.
9. At this point, the MTA performs an AS-REQ request to the KDC server. In the AS-REQ request, the MTA places its manufactures certificate. This certificate is unique to the MTA. The purpose of this request is to obtain a Kerberos ticket. The MTA needs this ticket later in the provisioning process.
10. Before responding to this request, the KDC must authenticate this MTA. It does so by checking the MAC address in the MTA certificate, and then verifying the FQDN used by the MTA against this MAC address. The KDC does this by interacting with the provisioning server, because the provisioning server is a trusted element, and it has the FQDN-to-MAC address mappings. If this authentication succeeds, the KDC responds with an AS-REP reply, which includes the Kerberos ticket the MTA was requesting.
11. The MTA uses the Kerberos ticket when sending an AP-REQ request to the provisioning server specified in option 122 suboption 3. This AP-REQ request is a request by the MTA for the keying information needed for the Simple Network Management Protocol Version 3 (SNMPv3) process, which is discussed later in this paper.
12. If the provisioning server correctly verifies this Kerberos ticket, it replies with the necessary keying information in an AP-REP reply message to the MTA.
13. The MTA now contacts the provisioning server via an SNMPv3 INFORM message that it is now ready to receive its configuration file. The SNMP process is encrypted using the keys provided in t he AP-REP response.
14. The provision server responds with an SNMPv3 SET with the configuration file name and TFTP server FQDN with which the MTA will request its configure file. The provisioning server also includes a secure hash algorithm 1 (SHA-1) hash of the configuration file. This is to ensure the integrity of the configuration file (that is, that it has not been altered in any way). Finally, the provision server includes the encryption key to be used to decrypt the configuration file, because the file is encrypted when sent. This prevents any other device from using this configuration file. For more details on the security aspects, refer to [15].
15. The MTA issues a DNS query to get the IP address of the TFTP server FQDN.
16. The DNS server responds with the IP address of the TFTP server.
17. The MTA issues a TFTP request from the TFTP server for the configuration file.
18. The TFTP server responds to the MTA with the configuration file requested, which is encrypted.
19. After decrypting the configuration file, the MTA completes the provisioning process by issuing an SNMPv3 INFORM to inform the provisioning server that provisioning is complete.

For more information about PacketCable provisioning, refer to the PacketCable MTA Device Provisioning Specification [12].

CMS Provisioning

CMS provisioning is just as important as MTA provisioning. In order to scale to high take rates and deal with churn, MSOs need a standard way of updating the CMS when a subscriber is added, deleted, modified, etc. This setup allows an MSO to use CMSs from different vendors, while keeping their provisioning workflow systems intact.

Because the related PacketCable specification is in the early stages of implementation, it is not described in detail here. In summary, however, this specification describes a schema to be used by provisioning servers when provisioning CMSs. Attributes such as subscriber name, phone number, etc. are provided in this schema. The specification recommends that the provisioning server use Extensible Markup Language (XML) to interact with the CMS. For more information about this specification, refer to the PacketCable CMS Subscriber Provisioning Specification [13].

Lawful Intercept

When establishing a residential VoIP service, local regulatory issues must be considered. For example, in the United States, the FCC mandates that all residential telephony systems have the capability to be "tapped;" that is, a law-enforcement agency (LEA) must be able to listen in on a suspect call. There is still debate about exactly what is needed and what timeframes need to be met, but it will be required at some point in the very near future.

The PacketCable Electronic Surveillance Specification [14] fully describes the necessary architecture for this capability.

Figure 7
Logical Lawful Intercept Call Flow

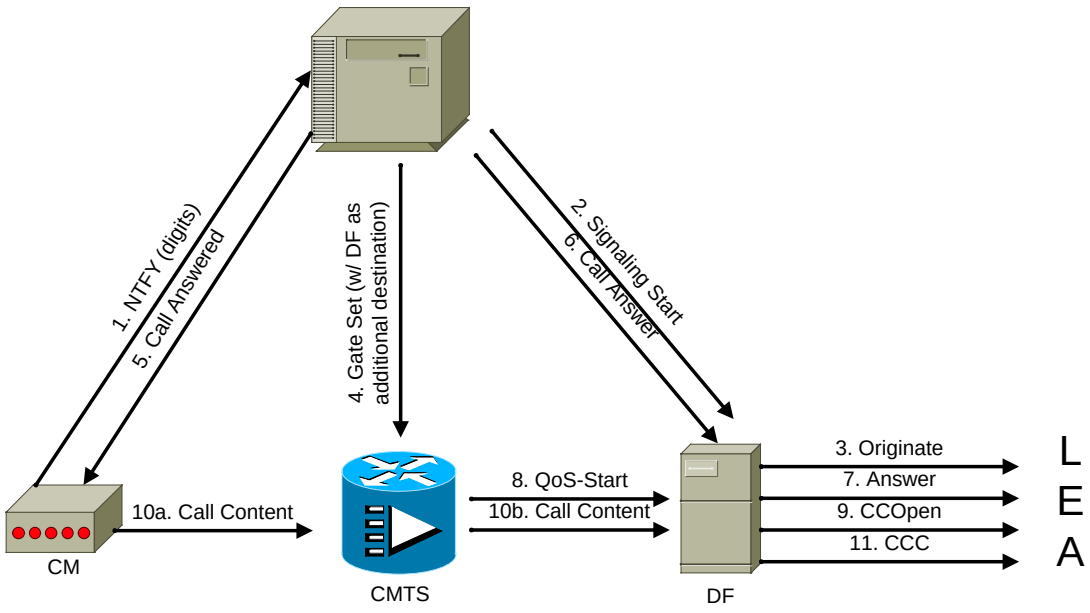

As Figure 7 shows, lawful intercept is performed in the following manner (some details are left out for clarity):

0. A suspect subscriber's phone number is provisioned in the CMS to be "tapped."
1. The suspect makes a call, and digits are sent to the CMS.

2. When the digits are received at the CMS, the CMS informs the distribution-function server with a Signaling_Start message.
3. The distribution function then sends an Originate message to the collection function at the appropriate LEA.
4. The CMS then sends a Gate_Set message with the original destination and adds the distribution-function server as a second destination.
5. The suspect's destination answers the phone.
6. When the call is answered, the CMS informs the distribution function via a Call_Answer message.
7. The distribution function then sends an Call_Answer message to the collection function at the LEA.
8. When the bandwidth is fully committed, the CMTS sends a QoS_Start message to the distribution function.
9. The distribution function then sends a CCOpen message to the collection function at the LEA. This message informs the collection function of the beginning of delivery of call content information.

10a The call is in progress and content is delivered to the CMTS.

10b The CMTS in turn duplicates the call content and delivers it to both the intended destination and the distribution function.

11. The distribution function forwards the call content to the collection function at the LEA over the call content channel (CCC).

Note that the FCC considers a solution that is in compliance with the PacketCable Electronic Surveillance Specification a "safe harbor." This means that if you implement this architecture, you are in accordance with the FCC Communications Assistance for Law Enforcement Act (CALEA) regulations.

For more details on the PacketCable lawful intercept functionality, refer to [14].

Security

The purpose of any security technology is to protect items of value, whether a revenue stream, purchasable information asset, or the delivery infrastructure. All types of VoIP architectures offer security provisions for individual network components. The PacketCable security specification spans the entire PacketCable network, not just individual components. This comprehensive, standards-based specification is designed to create a network that is as secure—or in many cases more secure—than the PSTN.

PSTN security largely depends on each telephone being connected to a dedicated line. Providing the same level of privacy and resistance to denial-of-service attacks when a PacketCable network is used for voice communications requires cryptography-based security mechanisms. Defined in the PacketCable security specification, these mechanisms secure both voice and signaling data transmitted over the shared HFC network and shared IP backbone.

Any fee-charging service that is offered to residential customers across a large, open network is susceptible to security intrusions. These intrusions can cost the MSO large amounts of money, and in some cases, can cause grief for customers (that is, if their MTA is compromised). The security aspects of the PacketCable solution are extremely complex, and an in-depth discussion is beyond the scope of this white paper. Consequently, this white paper does not discuss details of the security architecture, but rather focuses on the two main concepts: signaling security and media security.

Signaling Security

Any time the MTA has to interact with a trusted device (any device controlled by the MSO and on the MSO premises), it must have a security association with that device. A security association is a set of provisioned security elements (for example, security keys) on both the MTA and the trusted element. With a set of security associations, the trusted device can authenticate the MTA (because only one MTA can have the correct security association) when interacting with it. The interaction that takes place between the MTA and the trusted element is all encrypted using IP Security (IPSec), with the preprovisioned (preshared) keys. This setup ensures that all traffic between the two devices is from known sources and is encrypted.

An example of this would be communication between a MTA and a CMS. The CMS and MTA must have security associations with each other, and when the MTA contacts the CMS, the CMS knows that the MTA in question is a valid MTA because the security associations match. Also, this setup prevents intruders from assuming control of the MTA, because the MTA knows that CMS interactions can come only from the CMS with the matching security association.

Media Security

In order to protect customer privacy, customers' conversations must be kept private. In order to do this, the media streams generated by their conversations must be encrypted. The PacketCable security specification has provided a means by which to do this. Essentially, the endpoints in the conversation (MTA and MTA or MTA and media gateway) negotiate a set of ciphersuites (type of authentication and encryption to be used), and then encrypt all their traffic using the negotiated method. Currently the Advanced Encryption Standard (AES) is the only encryption method required by the PacketCable security specification. The authentication algorithms required for PacketCable are the multilinear modular hash (MMH) algorithm for RTP media and the SHA-1 for RTP Control Protocol (RTCP) media. These algorithms are described in detail in the PacketCable security specification.

For more information about the security architecture for PacketCable, refer to the PacketCable Security Specification [14].

Overview

In order to help MSOs deploy a primary-line residential VoIP service, Cisco has designed a solution based on the PacketCable architecture. This solution includes both Cisco and third-party components. This section focuses on the different components that make up this solution, and discusses an area that the solution addresses but PacketCable does not such as, redundancy. Figure 8 provides a detailed diagram of the Cisco PacketCable solution.

Figure 8
Cisco PacketCable Solution Architecture

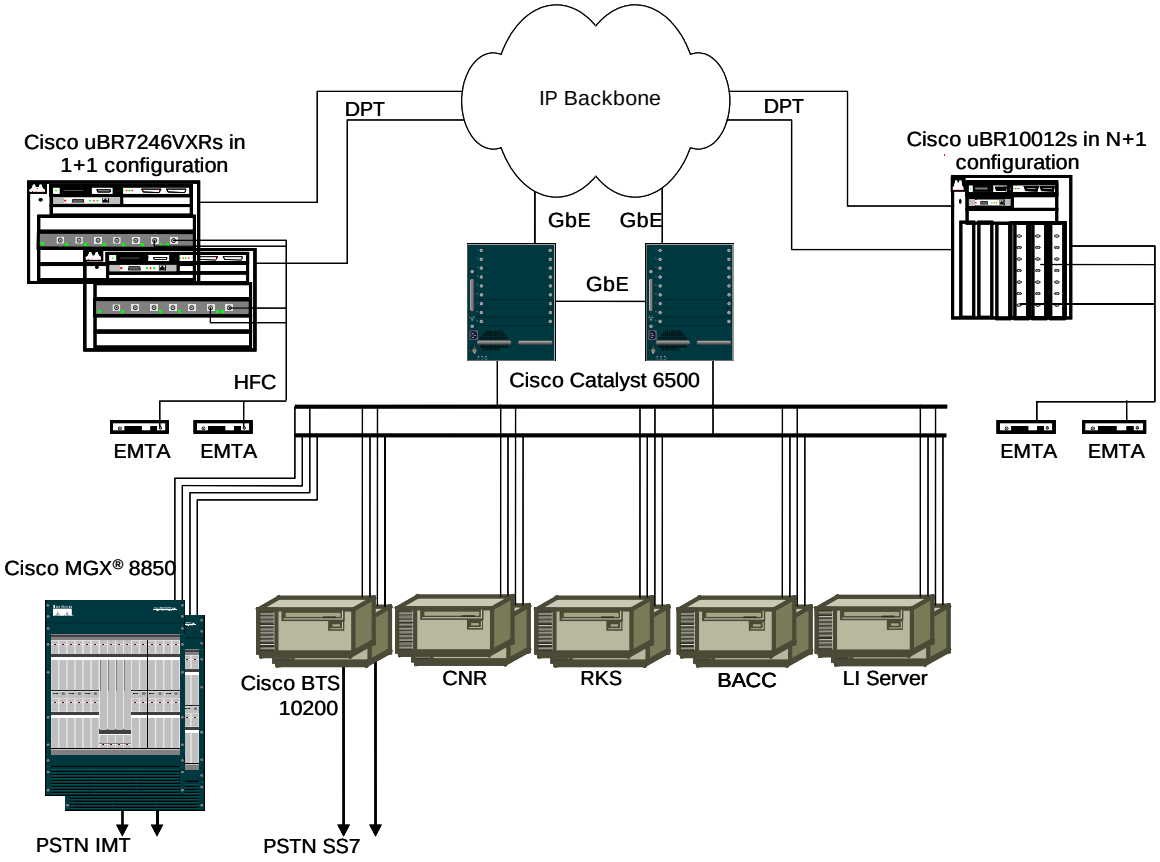

Product Components

Table 1 lists the major components of the Cisco PacketCable solution.

Table 1 Table 1 Components for Cisco PacketCable Solution Architecture

Component	Model
CMTS Cisco uBR7246VXR and uBR10012 universal broadband routers	
Media gateway*	Cisco MGX® 8850 with Voice Networking Service Module (VISM)
Provisioning servers	Cisco Network Registrar® system with Cisco Broadband Access Center for Cable (BACC)
Aggregation switch*	Cisco Catalyst® 6500 Series
CPE	Customer's choice; Cisco recommends either Thomson/RCA DHG445 or Toshiba PCX3000
CMS/MGC	Cisco BTS Version 3.5.1
RKS	BayPackets RSI Mediate 6000
Layer I server	SS8 Xcipio or Verint RELIANT

Switch and trunking gateway models may account for customer sizing.

CMTS – Cisco uBR7246 VXR and uBR10012

In brief, the Cisco uBR7246VXR and uBR10012 serve as the aggregation point of cable modems (with and without voice capabilities); they are connected to the IP network via a wide array of interface options. The Cisco CMTSs are both DOCSIS 1.1 and PacketCable 1.0 qualified.

Cisco uBR7246VXR Component Features

- Four slots for modem line cards
- Support for the G1 network processor
- DOCSIS 1.0 and 1.1 support
- Two port adapter slots for WAN uplink port adapters (optional Fast Ethernet interface on input/output (I/O) controller board for management network connectivity)
- Support for redundant power supplies
- Support for N + 1 redundancy
- National clock synchronization capability
- Support for up to 2 x 8 modem cards

Cisco uBR10012 Component Features

- Eight slots for modem line cards
- Support for performance routing engine (PRE)-1
- DOCSIS 1.0 and 1.1 support
- Four slots for WAN uplink port adapters
- Support for redundant PREs
- Support for redundant power supplies

- Support for N + 1 redundancy
- National clock synchronization capability
- Support for up to 5 x 20 modem cards

Media Gateway – Cisco MGX 8850 with VISM

The Cisco MGX 8850 Multiservice Switch with the VISM is the gateway between the IP and PSTN worlds for the bearer traffic. The Cisco VISM provides toll-quality voice, fax, and modem transmission and efficient utilization of wide-area bandwidth through industry-standard implementations of echo-cancellation, voice-compression, and silence-suppression techniques.

All other major Cisco media gateways can be used with the Cisco PacketCable architecture. However, because the Cisco MGX 8850 will potentially be the most widely used, component features are listed here for that media gateway only. Information about other Cisco media gateways can be found at http://www.cisco.com.

Component Features

- Support for up to eight E1/T1 ports per VISM, and 24 VISMs per Cisco MGX switch, for a total of 5952 or 4608 DS0s
- Support for T3 when using the service resource module (SRM)
- SRM support for up to three T3s, two active SRMs per chassis, for a total of six T3 interfaces
- Support for G.711 a-law and mu-law, G.726-16/24/32/40, and g.729a/b
- Support for two process switch modules, one active and one hot standby
- Carrier-quality echo cancellation: programmable 16-, 32-, 64-, and 128-ms tail lengths
- Tone detection and generation, which detects multifrequency tones and generates multifrequency and call-processing tones (for example, dial tone, call-waiting tone, etc.)
- From 1.2 to 45 Gbps per single chassis
- Network uplink interfaces from Ethernet to OC-48

Provisioning Server – Cisco Broadband Access Center for Cable and Cisco CNS Network Registrar

The Cisco CNS Network Registrar dynamic IP address management system assigns IP addresses to cable modems and other devices on the network based on a predefined set of policies. In addition, the Cisco CNS Network Registrar system is a full-featured DNS server. The Cisco Broadband Access Center for Cable system adds the PacketCable components to the provisioning system, in addition to dynamic configuration file creation and class-of-service grouping.

Component Features

- Full PacketCable 1.1 compliance
- Dynamic creation of both DOCSIS and PacketCable configuration files
- Dynamic DNS
- DHCP safe failover mode
- Primary and secondary DNS support
- DNS round-robin reply support

- Multiple-address reply support
- Client-class support, enabling the provisioning of modems for different services

Aggregation Switch – Cisco Catalyst 6509

The Cisco Catalyst 6509 delivers high-performance, multilayer switching solutions for enterprise and service provider networks.

- Component Features
- 256-Gbps switching fabric
- Support for Supervisor Engine 2
- Gigabit Ethernet switching modules
- Support for Inter-Switch Link (ISL) trunking
- Cisco Express Forwarding on Supervisor Engine 2 and Gigabit Ethernet modules
- Multiprotocol Label Switching (MPLS) provider edge support
- Hot Standby Router Protocol (HSRP) support between supervisor engines

Customer Premises Equipment

Although this architecture document does not specify the particular CPE device to be used, the Toshiba and Thomson modems listed in Table 1 have been successfully tested with the solution.

Component Features

CPE devices should have at least the following features:

- Voice coder-decoder (codec) support, including at least G.711a-law and mu-law, G.729e, and G.728
- Two to four voice ports
- DOCSIS 1.1 compatibility
- PacketCable 1.0 certification

CMS – Cisco BTS 10200

The CMS is the center of call control for the Cisco solution. The CMS has full control over the residential gateways in the network, and routes calls to the appropriate destinations, on net or off net.

Component Features

- Full PacketCable 1.0 or PacketCable 1.1 compliance
- Full fault tolerance
- Ability to handle 70+ calls per second with DQoS

MGC and Signaling Gateway – Cisco BTS 10200

In the Cisco solution, the MGC and signaling-gateway capabilities are built into the Cisco BTS 10200.

Component Features

- Media gateway control via MGCP 1.0
- Compatibility with major Cisco media gateways
- Support for ANSI ISUP

Record Keeping Server

The RKS is a trusted network element component that receives PacketCable event messages via the RADIUS protocol from other trusted PacketCable network elements such as the CMS, CMTS, and MGC. As part of its solution testing, Cisco has tested with the BayPackets RSI Mediate 6000. The BayPackets RKS assembles the event messages into CDRs, which are then made available to back-office systems such as billing, fraud detection, and other applications.

Component Features

- Bellcore standard GR-1303 is built into the solution to interface with billing applications.
- The RKS supports Automatic Message Accounting Billing AMA Format (AMA-BAF) CDR and IPDR standard formats, as well as Automatic Messaging Accounting (AMA) call codes 110, 001, 006, 008, and 009. The flexible rules engine can generate real-time custom data records based on the customer's unique needs.

Redundancy and Reliability

When deploying fee-charging services, the robustness of the solution must be considered. Interruptions in service annoy customers, and can cause them to cancel the service. In a residential VoIP service, it could mean the difference between life and death, for example, if a 911 call cannot get through because of a service outage. The "VoIP Availability and Reliability Model for the PacketCable Architecture Technical Report" [16] details the recommended reliability numbers for the different components of the network. This section describes how the Cisco solution attempts to meet those numbers.

The solution architecture is designed in such a way that there will be no single component whose failure could disrupt the service. Losing a single call-processing or IP network component does not affect established calls. Calls are lost only if the trunk that they use is lost, and this situation can occur in the PSTN today. Future versions of the solution will have stateful failover in the media gateway so that calls will not be lost even if a media gateway fails.

Where physically separate components are specified, it may make sense for the MSO to locate such components in geographically different areas, for example, different buildings, in case of damaging natural events.

CMS

The CMS needs to be composed of at least two physically separate call-processing servers, in order to continue to set up calls if one of the servers goes down. The BTS was designed from the beginning with this in mind. The BTS is composed of four servers, two servers for operations, administration, maintenance, and provisioning (OAM&P) and two call-control servers. If any one of these servers goes down, the system continues to operate, and the failure is transparent to the user. Actually, one OAM&P server *and* one call-control server can both fail, and the system will still continue to operate.

CMTS

The CMTS is one of the most important components in the network. Loss of a single CMTS affects all users connected to that CMTS. The Cisco architecture takes advantage of the N + 1 method to prevent service outages due to CMTS failures.

Figure 9 shows how the Cisco ubr7246VXRs are physically connected in an N + 1 setup.

Figure 9
Cisco uBR7246VXR N + 1 Redundancy Setup

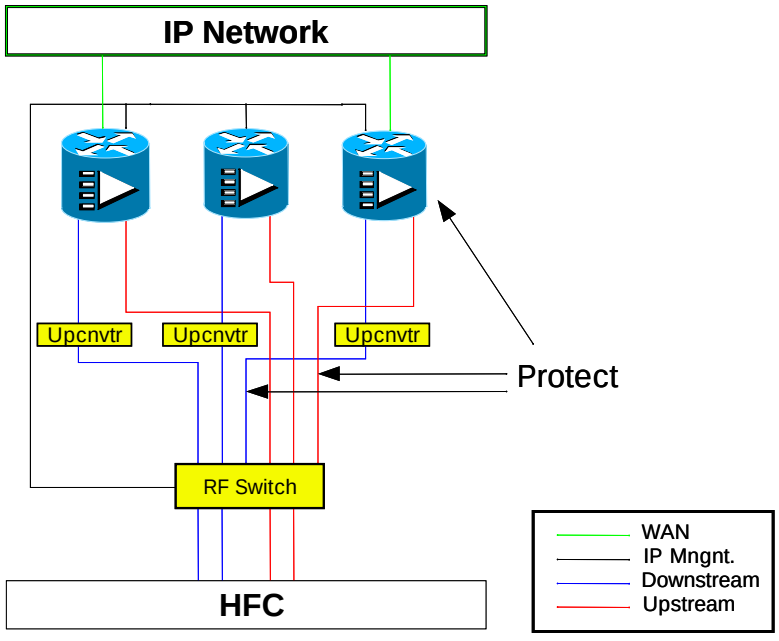

N + 1 redundancy is available for both the Cisco uBR7246VXR and the Cisco uBR10012 chassis. Both chassis have redundant power supply options, but the Cisco uBR10012 has additional intrachassis redundancy. This added redundancy includes dual PREs and intrachassis N + 1 redundancy. Figure 10 shows how N + 1 is implemented with the Cisco uBR10012.

Figure 10
Cisco uBR10012 N + 1 Redundancy Setup

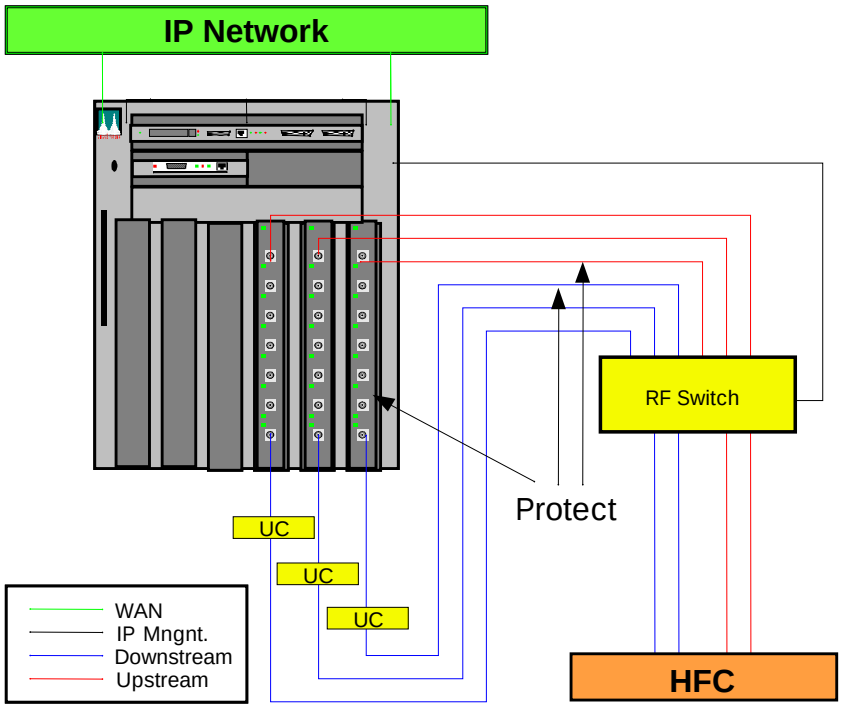

Media Gateway

The media gateway is the one component in the network whose failure causes a loss in active calls. The impact of a media gateway failure, however, can be minimized in numerous ways. Options are described as follows:

The use of a loopback interface on the media gateway protects against single IP interface failures. Loopback interfaces should be added to the list of networks that are advertised via the routing protocol running on the media gateway. This allows incoming traffic to be routed through the remaining operational interfaces. In addition to using the loopback interface, Cisco IOS® gateways have an added feature such that gateways with multiple, equivalent interfaces can have one interface configured as active and the other one as standby. For example, by using the backup interface command, a Cisco SN 5400 Storage Router with two Fast Ethernet interfaces can have one interface configured as the active interface and the second configured as backup. The backup interface uses the same IP and MAC as the original interface, so ARP delays will not be experienced.

Higher availability can be assured by using at least two gateways to handle the trunk group between the IP system and the PSTN. If one of the digital interfaces or the entire gateway handling a certain trunk group fails, the ISUP protocol in coordination with the MGC ensures that the corresponding trunk circuits are taken out of service, and that new calls are directed to the remaining gateways.

Additionally the use of continuity testing (COT), allows the MGC to test circuits before routing calls across them. If the circuit fails, the MGC notes this, and routes the call over the next operational circuit. An alarm is generated, and a manual intervention can occur to fix the failure and configure the MGC to start using that circuit again. COT is usually performed on a certain percentage of calls going through the system. It does increase call setup time, but not dramatically. The percentage of calls that should use COT can be configured in the MGC.

Provisioning Server

The Cisco Broadband Access Center for Cable (BACC) consists of three components. The regional data unit (RDU) is the heart of the system, responsible for the central data store, configuration engine, and interaction with administrators via the graphical user interface (GUI). The device provisioning engine (DPE) interacts with both the RDU and Cisco CNS Network Registrar to supply a device with the proper information through DHCP. The DPE is also a caching TFTP server. This reduces the number of interactions that must take place between the RDU and DPE, because the DPE already has a copy of the configuration in its cache. The third component is the Cisco CNS Network Registrar, which is the DHCP and DNS server.

The provisioning server ensures high availability by using multiple hosts for mission-critical applications. The most mission-critical aspects of the system are the DPE and the Cisco CNS Network Registrar because they interact directly with user devices. Each of these applications has failover mechanisms in place such that if one host fails, the other host takes over. Failover of either application is transparent to the end user. The only network configuration requirement to enable this failover capability is to configure the CMTS with multiple cable helper addresses. This setup allows the CMTS to forward DHCP broadcast to both hosts, thus ensuring that one server gets the request if the other is down.

Regional Data Center

Two Cisco Catalyst 6500 Series switches aggregate all nodes in the regional data center. Each switch has redundant power supplies and supervisor engines. In the rare event that an entire switch fails, the switches are configured such that the remaining switch takes over instantly, without any interruption in active calls or calls being established. The following paragraphs explain how this is done.

In most circumstances, hosts used for different service offerings are placed in different virtual LANs (VLANs) on the Cisco Catalyst 6500 Series switches. This offers the added benefits of manageability and potential reduction in the number of collisions on the VLAN, and helps when using a network sniffer for troubleshooting. The Cisco Catalyst 6500 Series assigns an interface to each VLAN on the switch. To the routing protocol, this interface acts as a routing interface, essentially the same as any interface on any Cisco router. The Cisco Catalyst 6500 Series switches run the HSRP between all VLAN interfaces (one VLAN interface is configured on each Catalyst 6500 Series). Each VLAN interface is given a priority, and the highest-priority interface is the active one. If that engine fails, the second-highest-priority interface takes over. HSRP allows all interfaces to serve the same IP address (although the actual assigned VLAN IP addresses must be different), because only one card is active at any instant. This setup allows all nodes being aggregated by the switches to use the HSRP address as the default gateway IP address, and if a failure occurs, no reconfiguration needs to be done on the nodes.

Figure 11 illustrates how the switches are physically connected to the backbone network. The goal of this connectivity is to allow traffic to flow from the hosts to the backbone, and vice versa, even if an entire switch or backbone router fails.

Figure 11

Regional Data Center High Availability

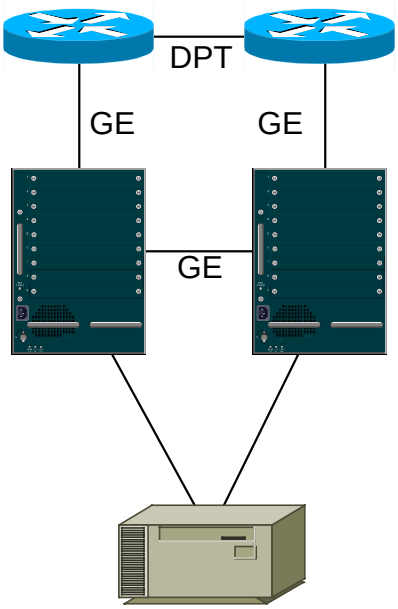

As shown in Figure 11, Gigabit Ethernet links connect each switch with one of the backbone routers. The switches are also connected using a Gigabit Ethernet link. The VLAN used for these hosts is trunked across this Gigabit Ethernet link, so intra-VLAN traffic can traverse the switches transparently. Each host has an interface into each Cisco Catalyst 6500 Series to protect against failure of an entire switch. If one switch fails, the host still has a route out the other switch.

Conclusion

The PacketCable working group of CableLabs has designed and specified an architecture capable of supporting voice, video, and data across a single HFC network. The needs of these services, however, require that the architecture allow for:

- Event messaging and billing
- DQoS
- PSTN interconnect
- Provisioning
- Lawful intercept
- Security
- Reliability and redundancy

With all these requirements, it is very difficult to design a simple solution. Therefore, the PacketCable solution looks, and is, very complex. This white paper was written to help clarify some of this complexity, and offer insight into how Cisco is implementing this architecture.

Reference Documents

[1] **PacketCable Network-Based Call Signaling Protocol Specification, Version I06, Nov. 27, 2002.**

[2] **PacketCable PSTN Gateway Call Signaling Protocol, Version I06, Nov. 27, 2002.**

[3] **Media Gateway Control Protocol, Version 1.0 – RFC 3435.**

[4] **Arango, M., and Huitema, C. Simple Gateway Control Protocol (SGCP), Version 1.1 Draft, July 30, 1998.**

[5] **Boyle, J., et al. The COPS (Common Open Policy Service) Protocol, RFC 2748, Jan. 2000.**

[6] **PacketCable Dynamic Quality-of-Service Specification, Version I05, Nov. 27, 2002.**

[7] **PacketCable Event Messages Specification, Version I05, Nov. 27, 2002.**

[8] **SIP: Session Initiation Protocol – RFC 3261.**

[9] **PacketCable CMS to CMS Signaling Protocol, Version I02, Dec. 5, 2002.**

[10] **PacketCable Internet Signaling Transport Protocol (ITSP) Specification, Version I02, Dec. 21, 2001.**

[11] **Data-over-Cable Service Interface Specifications, Radio Frequency Interface Specification, SPRFIv1.1- I09-020830, Cable Television Laboratories, Inc., Aug. 30, 2002.**

[12] **PacketCable MTA Device Provisioning Specification, Version I05, Nov. 27, 2002.**

[13] **PacketCable CMS Subscriber Provisioning Specification, Version I01, Dec. 5, 2002.**

[14] **PacketCable Electronic Surveillance Specification, Version I01, Dec. 29, 1999.**

[15] **PacketCable Security Specification, Version I07, Nov. 27, 2002.**

[16] **VoIP Availability and Reliability Model for the PacketCable Architecture Technical Report, Version 1, Nov. 28, 2000.**

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

RDA4671 04/03