Volume 06 Issue 04 Published, November 15, 2002 ISSN 1535-766X

Intel® Technology Journal

Interoperable Home Infrastructure

This issue of Intel Technology Journal (Vol. 6, Issue 4, 2002) explores exciting new infrastructures for the interoperable home and discusses how future homes can become more responsive to our changing life styles by combining new technologies in software and electronics.

Inside you'll find the following papers:

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

Volume 06 Issue 04 Published, November 15, 2002 ISSN 1535-766X

Intel® Technology Journal

Interoperable Home Infrastructure

Articles

Digital Content Distribution in the Home
Lin Chao, Publisher

In the not-so-distant future, the many new developments in consumer electronics, human-machine interface, wireless, digital content (such as photos and MP3 music files) and communication technologies will find a new home—and it may be your home.

The new "digital home" will connect all sorts of devices—from PCs to phones to TVs to microwave ovens to as-yet-to-be-created gadgets—and allow them to communicate easily with one another.

Imagine this: A parent could use a PDA (personal digital assistant) or a telephone keypad to tell the microwave to cook the dinner for 15 more minutes. At the same time, that PDA could tell the PC to record a special TV show. After eating the dinner, the family could watch the special TV show. These new possibilities for a wide range of products and services are under development today. The digital home is possible if consumer-electronics manufacturers and computer companies start integrating networking technologies into products.

We at Intel are especially interested in these new possibilities. We are working with consumer electronics and computer companies to develop a UPnP* standard that defines how digital content is distributed and consumed within a home network. The goal of the UPnP standard is to help connect new and existing devices within the home and make them easy to use.

This issue of Intel Technology Journal (Vol. 6, Issue 4, 2002) explores exciting new infrastructures for the interoperable home. The first paper explains the framework for the interoperable home. The next four papers examine the details of connecting together the parts, including media distribution, remote I/O, network security and content protection. The final two papers address networking with multi-hop wireless networks and IPv6 internetworking.

These seven papers show how future homes can become more responsive to our changing life styles by combining new technologies in software and electronics. UPnP technologies will lay the "foundation" in the new-era home.

Interoperable Home Infrastructure Foreword

Abel Weinrib, Director & GM, Network Architecture Lab
Gerald Holzhammer, Director & GM, Desktop Platform Architecture

In the last few years, consumers have seen incredible innovations in technology that they can apply to their everyday lives—and major shifts in their behavior when they have chosen to apply them. One such shift is about to happen in the emerging "digital" home, namely interoperability of PC and consumer electronic devices.

The vision for consumer electronic devices and personal computers to interoperate within the home is not new. It is fair to ask what has changed and why will it happen now?

- Powerful but affordable home networks built on robust wired or wireless Ethernet technology are quickly becoming a reality in many homes.
- Consumer devices are rapidly converting from analog to digital platforms to support new digital media types such as MP3, MPEG 2, DV or HTML.
- Lots of the digital content stored on PCs, such as music, images and increasingly video, is best experienced in other areas of the home.

While these trends are key enablers, they are not sufficient by themselves to make the digital home happen. Ease of installation and operation is another fundamental requirement for broad deployment. Intel has been working with the industry to simplify configuring home-networking and to enable new approaches for device interoperability so that consumers will be able to quickly, effortlessly, and securely install new peripherals.

To realize this seamless interoperability, each device will have to adhere to a well-defined set of protocols, formats and agreements – the digital home interoperability stack (see figure).

Central to the interoperability stack are a common transport protocol (Internet Protocol - IP) and a general device discovery and management protocol called Universal Plug and Play (UPnP*). Ethernet, both in its wired or wireless forms, is the physical network transport of choice. This basic set of protocols goes a long way to ensure coexistence of devices on the home network. But our vision goes beyond coexistence. Enabling off-the-shelf interoperability of devices from different manufacturers will require agreement on basic media formats and transports and a way to securely exchange protected high-value content (such as commercial movies) between devices.

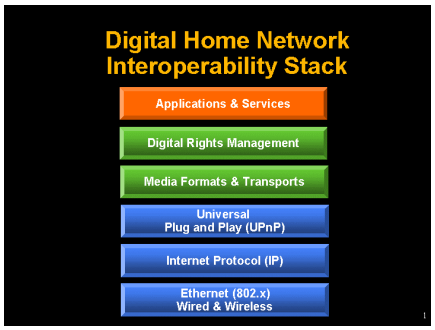

Intel is working across industries to drive a common digital home vision and the standards to support it. Collaboration between the computer and consumer electronics industries is required to make this vision a reality. As a result, devices in the home will be able to seamlessly work together through widely accepted open standards such as IEEE, IP networking protocols, and UPnP device protocols.

In this issue of Intel Technology Journal, we refer frequently to one of the key enablers of the digital home framework—the UPnP protocol. Intel is a founding member of the UPnP Forum, an association of more than 500 companies working together to develop interoperable specifications and standards for easy home networking. UPnP technologies address device discovery and control, and far more, including security, quality of service, and even protecting commercial content across the network. Another technology area covered in this issue that is crucial to the digital home is the network itself. Multi-hop, wireless networks are one way to find maximum bandwidth with lower power requirements. Finally, the transition to IPv6 in the home will restore transparency that Network Address Translation (NAT) for IPv4 has removed.

* UPnP is a certification mark of the UPnP Implementers Corp.

Home Interoperability Framework for the Digital Home

Yasser Rasheed, Corporate Technology Group, Intel Corporation
Jim Edwards, Desktop Platform Group, Intel Corporation
Charlie Tai, Corporate Technology Group, Intel Corporation

Index words: UPnP Technology, Home Interoperability, Digital Home, PC/CE Collaboration

ABSTRACT

The transition from analog to digital media, coupled with recent advancements in broadband and home networking technologies, is fostering a Digital Home environment where rich multimedia content can be delivered to and distributed throughout the home seamlessly. This promises mutually beneficial opportunities for collaboration between the Personal Computer (PC) and Consumer Electronics (CE) industries, and it paves the way for new usage scenarios and significantly improved consumer experiences.

A key element for a successful PC/CE collaboration in the Digital Home is interoperability between devices operating on a common home network. We define interoperability as the ability for devices in the home to discover, configure, and control the capabilities of peer devices and to negotiate common protocols and media formats for proper multimedia content distribution.

This paper presents a number of typical usage scenarios in the Digital Home and then uses these scenarios to extract interoperability requirements. A standards-based Home Interoperability Framework is then described that enables vendor interoperability. Finally, a case study is used to illustrate the Home Interoperability Framework along with its building blocks in action.

We believe that the interoperability of PC and CE devices will (1) reinforce consumers' expectations about anytime, anywhere access to their content; (2) create new device categories; (3) add functionality to existing devices; and (4) offer numerous opportunities for content providers, manufacturers, retailers, and service providers to profit from advances in technology. To achieve the necessary interoperability and provide the consumer with the ultimate digital entertainment experience, it is imperative for both the PC and CE industries to embrace a common framework.

INTRODUCTION

In recent years, the world has witnessed tremendous advancements in computer platforms, consumer electronics, and communication networks. The phenomenal growth of the Internet and the insatiable demand for bandwidth have resulted in a growing demand by consumers for broadband access. Meanwhile, the explosion of mobile devices has conditioned consumers to expect access to their information and content *anytime* and *anywhere*. The deployment of home networks promises to help fulfill this consumer need in the Digital Home and be the catalyst for a new wave of digital devices. Technological advancements will enable high-speed Internet services and rich multimedia content delivery to the home, paving the way for significant enhancements to consumers' entertainment experiences.

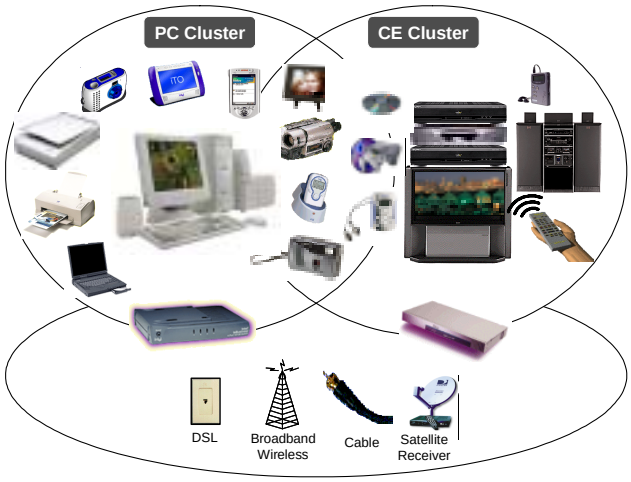

Figure 1: Clustered home entertainment environment

As illustrated in Figure 1, the typical home has two clusters of equipment used for entertainment purposes: the Consumer Electronics (CE) cluster and the Personal Computer (PC) cluster.

The CE cluster has traditionally been considered the main entertainment environment for home users. It normally consists of TVs, stereos, video cassette recorders (VCRs), DVD players, and personal video recorders (PVRs). Its content comes from terrestrial broadcast, cable, or satellite programming received through dedicated set-top boxes (STBs), and from removable media such as video tapes, CDs, and DVDs.

The PC cluster, consisting of one or more PCs and a number of attached peripheral devices, is emerging as an alternative, complementary entertainment environment. The content for the PC cluster is typically rich multimedia content acquired from the Internet over a broadband access pipe, or from removable media such as CD-ROMs and DVD-ROMs.

In the future, these two clusters will increasingly communicate with each other and form the basis for what we refer to as the Digital Home. This trend is exemplified by PC and CE manufacturers' integration of new digital features and functions into their products. For example, more and more CE devices are capable of communicating with the PC using Ethernet and IEEE 1394 connectivity. As a result, a 'network effect' is being unleashed in the home, resulting in a multitude of new devices and appliances that share content and resources, expanded revenue opportunities, and increased consumer satisfaction.

The new home devices that emerge will combine the power of the PC, especially with respect to management, storage, and processing capabilities, with the convenience and ease-of-use of CE devices. The rich multimedia content available in PC and CE clusters will be easily exchanged between a large number of devices, allowing the consumer to experience it *anytime* and *anywhere* in the home. This will result in a number of innovative usage scenarios that can greatly enhance consumers' entertainment experiences.

USAGE SCENARIOS

The following scenarios highlight the key drivers for the home network and the associated impact on device interoperability requirements.

Watching TV and Movies

On a Friday evening after a long day at work, Jim and his wife Pam sit on the living room couch and start to watch a movie on their TV. While they are watching, Jim is able to pause the live TV to accommodate an interruption by his kids, rewind to catch up on missed segments, and resume watching. After the movie is over, Pam remembers the preview she saw earlier and, pushing a few buttons on the remote control, sets the PVR to record the show.

On Saturday afternoon, Jim's kids take over the living room to watch their favorite shows recorded over the course of the week. When Jim sees his favorite chair occupied by his 11-year-old, he decides to watch TV in the kitchen. After glancing at the program guide and not finding anything he wants to watch on live TV, he pushes a button on the remote control to see what programs are stored on all the devices on his home network. A guide pops up and shows all the movies, TV shows, and music stored on his two PVRs and three PCs. He finds an old war movie (Jim loves the new service he signed up for that downloads movies based on his preferences over his broadband connection), makes himself a sandwich, and sits at the counter to watch the movie. After Jim finishes his sandwich, he realizes he prefers to watch the movie in a more comfortable position. He pauses the movie with the kitchen remote, turns off the TV, and heads up to his bedroom. Grabbing the bedroom remote, he lays down on his bed, locates the movie in the program guide, hits Play, and the movie resumes from where he paused it.

At the same time, Pam is doing the family finances on a PC in the den while listening to songs on the PC's jukebox application. She decides she needs a break and clicks on an icon on the desktop. A screen pops up (the same screen Jim saw on the kitchen TV) showing all the programs recorded on the home network. Pam sits back and starts watching her favorite documentary.

Benefits for Consumers

The entire family can easily and conveniently watch TV and movies on any display, whether it is a TV or a PC, anywhere in the home, all at the same time. The program guide shows a common integrated listing of all movies, TV programs, pictures, music, and other content from all the devices on the home network—at the touch of a button.

Listening to Music

Waking up late on Saturday morning, Jane reaches over to her bedside table for the remote control. She pushes the play button and the bedroom clock radio resumes playing her "weekend mix" playlist stored on her PC. Jane decides to stop the "weekend mix" and instead play all her songs from the U2 music group in shuffle mode. She grabs her wireless tablet, re-orders the songs by artist, and saves the resulting playlist to the PC. After a few more minutes of lying in bed, Jane goes to the kitchen to eat breakfast. She pushes a button on a remote control in the kitchen and the music starts playing on the kitchen speakers, without missing a beat.

Benefits for Consumers
Music can be conveniently stored on a PC for personalized access from anywhere on a home network and can be easily controlled by multiple devices like remote controls, tablets, and other PCs.

Capturing and Sharing Life's Moments in Pictures

Returning from their honeymoon in Hawaii, Bob and Alice are excited about sharing their pictures with friends and family. At the touch of a button, Bob wirelessly downloads the photos and video clips from his new digital camera to a home PC. Then Alice, feeling inspired while everything is still fresh in her mind, edits the photos and video clips into one file, adding captions and music to create a multimedia slideshow of their honeymoon.

At their party the next weekend, Bob and Alice invite their guests to join them in the living room to watch their honeymoon pictures. Bob pushes a button on the remote control and a display of available pictures and personal videos stored on the PC appears on the TV. Bob quickly locates the honeymoon slideshow and selects it. Everyone loves the pictures and clips, and compliments Alice on the captions she added.

After lunch the following Monday, Bob meets his friend Paul at a coffee shop and decides to show him his honeymoon slideshow. From his personal digital assistant (PDA), Bob connects to the 802.11 hotspot, logs on to his home network, and brings up the slideshow to show his friend.

Benefits for Consumers
Wireless technologies like Bluetooth* can be used to conveniently download pictures from a mobile device like a digital camera to a PC. 802.11 wireless home networks can be used to access content stored in one place (in this example, a PC) to view on a display elsewhere in the home. Broadband-enabled homes allow consumers to easily and securely access content stored in the home and from other locations outside the home.

Turning the Usage Scenarios into Reality

These scenarios demonstrate the range of new and innovative usage models that can be supported by a Digital Home network, where devices collaborate to provide the best user experience. Clearly, these represent only a few of the many possible usage models.

*Other brands and names are the property of their respective owners.

Interoperability between PCs and CE devices provided by different manufacturers is critical for these usage models to work seamlessly and without explicit user intervention. Interoperability must be realized at different levels for these devices to communicate with each other and exchange useful information. For example, interoperability must address networking and connectivity technologies, media formats, streaming protocols, and configuration and control mechanisms.

It is imperative that device manufacturers and application developers address interoperability as a design-level requirement, rather than as an optional feature.

HOME INTEROPERABILITY REQUIREMENTS

Vendor interoperability is accomplished between devices when they are capable of collaborating on a particular desired service that they provide to the consumer. Typically, this includes the capability for these devices to communicate with each other and exchange relevant information.

The list below identifies requirements at various layers of functionality to achieve vendor interoperability:

- *End-to-end connectivity between devices inside and outside the home.* This includes networking compatibility at the link layer (layer 2) and network layer (layer 3), such that devices are able to establish communication with each other.
- *Unified framework for device discovery, configuration, and control.* This refers to the ability of a device on the home network to discover the presence of other devices and services on the network, and identify their functionality and associated capabilities. It also includes the ability to configure these devices and services, and control their operation with the appropriate ease-of-use.
- *Common media formats and streaming protocols.* Once devices can communicate with each other, they need to agree on a common streaming protocol in order to establish a media streaming session. These devices also need to agree on a common media format to ensure that the media can be shared and consumed.
- *Common media management and control framework.* Media management and control refer to the ability to organize, package, browse, search, and select media items to be processed and the ability to control the operation of media streaming sessions. A media management and control framework must be established across all devices in the home to enable the proper exchange of information between devices provided by different vendors.

- *Flexible Quality of Service and policy management mechanisms.* Quality of Service (QoS) networking is essential for transporting multiple high-bitrate media streams in the home. In addition, users may want to apply certain usage rules that govern how their home network is used by household members. The key for both QoS and policy-based network management mechanisms is flexibility and ease-of-use. For this to work, all devices must agree on a common framework and associated mechanisms to implement these functions.
- *Compatible authentication and authorization mechanisms for users and devices.* A number of authentication and authorization mechanisms are being considered by device manufacturers and application developers to provide appropriate security for access and control. It is imperative to settle on a compatible framework to enable devices to request and/or grant access to particular devices and services in the home.
- *Common commercial content protection and Digital Rights Management framework.* A flexible and extensible content protection framework must be adopted by the industry that enables the chaining of content protection solutions to provide comprehensive, end-to-end protection of commercial entertainment content as it is delivered, managed, stored, and consumed on a wide range of devices within the Digital Home. A consistent set of technical and licensing mechanisms must be agreed upon to ensure that content is protected in an effective and efficient manner throughout the home.
- *Standard mechanisms for user interfaces at a distance.* In the future, devices will be created to specifically handle user interaction tasks. Applications running on other devices in the home can drive these "I/O devices" to deliver the desired end-user experiences. As a result, users will be free to interact with home networking applications from anywhere in the home, rather than being restricted by the location of the resources. The traditional fixed location, fixed user interaction style, such as a "PC-in-the-den" will be augmented with new kinds of user interaction, where a range of activities in the home will be supported by user interfaces tailored to the needs of the user.

HOME INTEROPERABILITY FRAMEWORK

This section defines a Home Interoperability Framework—one based on industry-wide standards and building blocks—that enables vendor interoperability. The functionality that results can be built upon, and enables the creation of compelling, interoperable devices and applications for today's market and future markets.

The Home Interoperability Framework (HIF) consists of a number of building blocks that work in harmony to ensure vendor interoperability between various types of devices in the home.

The HIF allows flexible interconnection between these devices, each with potentially different networking media technologies, using an internetworking layer based on Internet protocols.

The framework is divided into four layers, as shown in Figure 2.

- Device Connectivity Layer: The various connectivity options for home devices.
- Internetworking Layer: The routing and internetworking components based on the Internet Protocol (IP).
- Platform Middleware Layer: The collection of platform middleware building blocks used by various home devices.
- Applications and Services Layer: The applications and services running on a particular device. This layer includes the Content Protection and Digital Rights Management (DRM) building blocks for the management and distribution of premium content.

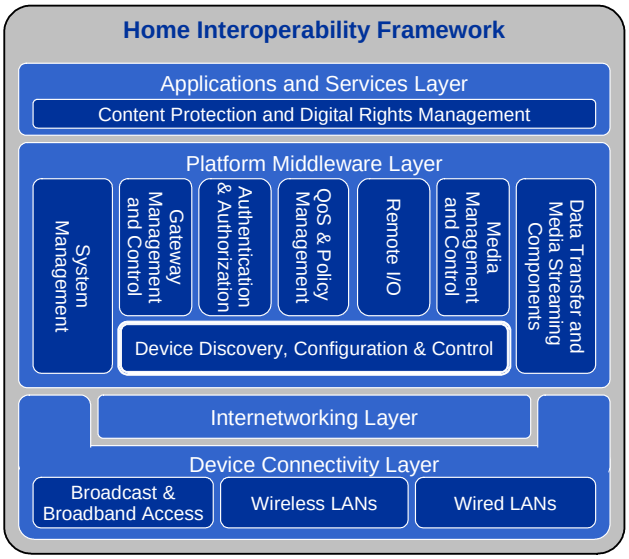

Figure 2: Home Interoperability Framework

The following sections explore the detail for each of the four layers of the Home Interoperability Framework.

DEVICE CONNECTIVITY LAYER

A fundamental requirement for devices to participate in a networked Digital Home is connectivity. Devices may include one or more connectivity options depending on their functionality, as illustrated at the bottom layer of the

Home Interoperability Framework (HIF) in Figure 2. The device connectivity options can be classified as follows:

- Connectivity to the outside world. This includes both broadcast and broadband access. Broadband access represents Internet connectivity to the Wide Area Network (WAN) through Cable, Digital Subscriber Line (DSL), Cable, or Wireless Local Loop (WLL). Broadcast access represents connectivity to content sources such as Terrestrial broadcast, Cable TV, and Satellite programming, all of which are typically non-IP based, in contrast to content received through broadband access connections.
- Connectivity to the home network. This includes both wireless Local Area Network (WLAN) options such as 802.11 technologies and wired LAN options such as Ethernet. In cases where 1394[1] is used as an entertainment cluster connect, further protocol translation may be needed at higher layers to make 1394 content and devices available for access on the Digital Home network.

INTERNETWORKING LAYER

The Internet Protocol (IP) is a family of protocols that provide the underlying network communications for devices on the Internet. IP also provides a suitable internetworking technology to facilitate pervasive connectivity for devices inside and outside the home, regardless of their physical connectivity technology. IP is based on industry-standard specifications implemented and supported in a wide range of devices with more than two decades of deployment in government, academic, and commercial environments.

IP encapsulates data from upper-layer protocols with address information. The address information is used to route the data traversing the network. Demand for IPv4 addresses for the burgeoning personal computer, consumer electronics, and mobile handheld markets is predicted to exhaust the supply of available IPv4 addresses, particularly in rapidly expanding market segments in the Asia-Pacific region. To conserve the limited supply of routable IPv4 addresses, many residential gateway vendors support "private" IPv4 addresses through Network Address Translation (NAT) technology.

Many interactive multimedia person-to-person communication applications, such as video conferencing and online interactive games, make use of upper-layer protocols that may break with NAT. This forces residential gateway vendors to add workarounds to NAT to repair the communications channel. Since each workaround is specific to an upper-layer protocol, many such fixes are necessary over time, and they limit the design of new upper-layer protocols. While NAT is a reasonable short-term answer for a limited supply of IPv4 addresses, a better solution is the new version of IP known as IPv6.

IPv6 is essentially IPv4 with built-in auto-configuration, enhanced support for mobility and security, and a much larger network address space that allows more devices to be transparently interconnected. The Internet Engineering Task Force (IETF) is actively pursuing a range of transition techniques for a smooth migration from IPv4 to IPv6, many of which are applicable to home devices and residential gateways. IPv6 is gaining traction in the PC, CE, and mobile industries as the long-term solution to the shortage of IPv4 addresses while maintaining end-to-end transparency.

The HIF supports both IPv4 and IPv6, with an emphasis on IPv4 in the short term and IPv6 in the longer term.

PLATFORM MIDDLEWARE LAYER

The Platform Middleware Layer in the Home Interoperability Framework (HIF) provides the basic platform building blocks that may be needed on various home devices. The cornerstone of this layer is the Device Discovery, Configuration, and Control building block, which are based on UPnP* technology.

The Platform Middleware Layer also includes platform building blocks such as Data Transfer and Media Streaming Components, Media Management and Control, Remote I/O, Quality of Service (QoS) and Policy Management, Authentication and Authorization, Gateway Management and Control, and System Management.

Discovery, Configuration and Control

The UPnP device architecture enables peer-to-peer network connectivity of PCs and CE devices of different form factors, intelligent appliances, and wireless devices. It is a distributed, open networking architecture that leverages Internet and Web technologies, such as Hyper-Text Transport Protocol (HTTP), Simple Object Access Protocol (SOAP), and eXtended Mark-up Language (XML), to set up flexible data communication between any two devices under the command of any controlling device on the network. UPnP specifications are defined

[1] This refers to the physical and link layers of the 1394 specifications

* Other brands and names are the property of their respective owners.

by the UPnP Forum [1], an industry consortium with more than 500 member companies including various PC and CE manufacturers.

The UPnP Forum working committees are responsible for establishing standard Device Control Protocols (DCPs). UPnP DCPs define the syntax and semantics for devices and services that implement a specific class of functions. The term *device* is used in the UPnP architecture specification to define a logical container of other devices and services, where *services* are logical entities providing a specific service to the UPnP device network. Services are controlled by control points, which are in turn defined as logical entities that can control specific services. A physical UPnP device may combine multiple services and/or control points. Several examples of UPnP services are discussed in the following sections.

The operation of UPnP devices is divided into the following phases:

- *Addressing:* In the addressing phase, devices obtain an IP address through Auto-IP or Dynamic Host Configuration Protocol (DHCP) mechanisms.
- *Discovery:* In this phase, control points search for devices and services, and devices advertise their services.
- *Description:* Once a control point finds a device or service of interest, it requests a description document. Devices and services respond by sending XML description documents that define the actions and attributes they support.
- *Control*: In this phase, control points invoke the actions described in the XML description documents associated with the services they control. These actions are executed by the services and typically cause changes in the service states and attributes. The syntax and semantics of these control actions are defined in the UPnP DCPs associated with the device class.
- *Eventing*: Control points subscribe to event servers hosted by the services, which allows them to receive events from a specific service they are interested in. Similar to control actions, events are defined in the corresponding DCPs.
- *Presentation*: Finally, devices may choose to host an HTML document that provides a user interface for the device.

UPnP technology forms the basis for the device discovery, configuration, and control building block in HIF. It defines a horizontal abstraction for distributed networking across PCs and CE devices in the home, independent of any particular operating system, programming language, or physical medium.

Data Transfer and Media Streaming

Media streaming refers to the ability to transfer real-time content between multiple, networked devices and between different software modules on a device (see Figure 3). Since these software modules usually perform media processing functions on the data being transferred, such as encoding and decoding, compression and decompression, and packetization and de-packetization, they are referred to as media streaming components.

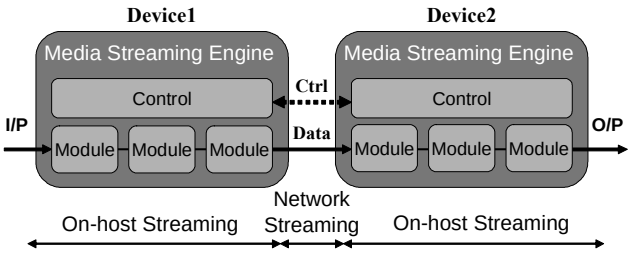

Figure 3: High-level architecture of media streaming

A media streaming engine is responsible for providing a connection (a virtual data channel) between the software modules to enable efficient data transfer between them. The media streaming engine is responsible for the setup, configuration, and maintenance of a network connection for streaming between devices.

A media streaming engine may provide the following features:

- dynamic (run-time) support for multiple media formats, streaming, and stream control protocols
- support for multiple, simultaneous streams
- audio/video synchronization and buffer management

A number of media streaming engines exist today. DirectShow*, which is available on Microsoft Windows* operating systems, is arguably the most prominent. Regardless of the diversity of the streaming engines on various devices in the Digital Home, it is necessary for all of the devices to agree on compatible components to allow the exchange of media content.

Media Management and Control

UPnP AV specifications define the interaction model between UPnP AV devices and associated control points. UPnP AV devices include TVs, VCRs, CD/DVD players, set-top boxes, stereo systems, still-image cameras, electronic picture frames, and PCs. The UPnP AV architecture allows devices to support various types of entertainment content such as MPEG2 and MPEG4 for

* Other brands and names are the property of their respective owners.

video, JPEG for pictures, and MP3 for audio. It also allows various types of transfer protocols such as HTTP and Real-time Transport Protocol (RTP).

UPnP AV specifications define two types of logical devices on the home network: Media Servers and Media Renderers (see Figure 4). They also define four "services" hosted by servers and renderers:

- The *Content Directory Service* enumerates the available "content" (videos, music, pictures, and so forth).
- The *Connection Manager Service* determines how the content can be transferred from the Media Server to the Media Renderer devices.
- The *AV Transport Service* controls the flow of the content (play, stop, pause, seek, etc.).
- The *Rendering Control Service* controls how the content is played (volume/mute, brightness, etc.).

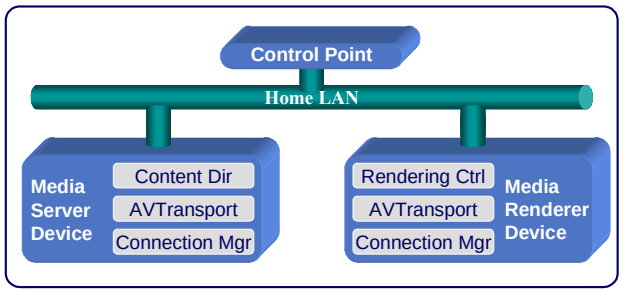

Figure 4: UPnP AV device architecture

The UPnP AV specifications are defined by the UPnP AV Working Committee (WC) as chartered by the UPnP Forum* steering committee. UPnP AV version 1.0 specifications have been defined with active support from many CE vendors and have been recently approved and published by the UPnP Forum.

The HIF uses UPnP AV technology for media management and control, allowing various PCs and CE devices to discover content on the home network and control media streaming sessions between these devices.

Remote I/O

Remote I/O is a technology under development based on the UPnP device architecture. It moves the point of user interaction away from an application running on a specific device, such as a PC or a CE device, to one or more remote I/O devices. The remote I/O device supplies input and output services such as mouse, keyboard, and display, that together comprise the user interface. Applications run on a host elsewhere on the home network and are matched with compatible I/O devices. Applications may take on different user interface characteristics depending on the I/O devices being used. Furthermore, the application user interface can migrate across I/O devices as the user moves about the home.

Remote I/O supports PC/CE collaboration by introducing *location independence* and *tailored I/O devices* to the home network. Applications can connect to wireless I/O devices anywhere in the home, freeing the user interaction from the location of the application. For example, a user might prefer to read the news in a comfortable chair in the family room instead of using the desktop PC in the den. In addition, I/O devices can be tailored to user activity. For example, a display for reading the news should be handheld, comfortable to hold, and allow the user to adjust visual settings such as colors and contrast. These two properties of remote I/O devices will improve the quality of user experiences in the home.

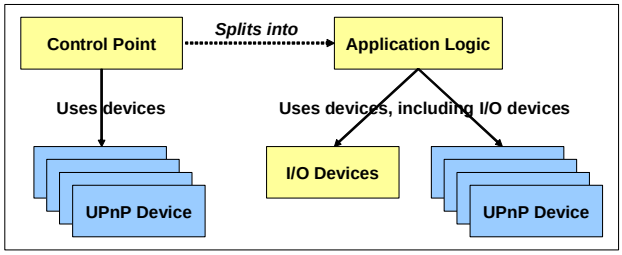

Figure 5: Control point splits into application logic and I/O device(s)

Remote I/O separates user interaction from application logic. In effect, a UPnP control point is being split in two. One part retains the controlling application logic, while the other part consists of the user interaction components and their related hardware. The I/O part becomes a standard UPnP device that is controlled along with other devices required by the application (see Figure 5).

Quality of Service and Policy Management

Streaming of rich multimedia content in the home typically requires using QoS networking and policy-based network management techniques to ensure the best consumer viewing experience possible. The QoS and Policy Management building block provides the necessary mechanisms for application developers to add support for QoS and policy management to their media streaming applications. The overall architecture of the QoS and policy management building block is shown in Figure 6.

The capabilities supported by the QoS and policy management building block may be divided into two categories:

* Other brands and names are the property of their respective owners.

- *Traffic Control:* Traffic Control (TC) refers to the ability of a device to perform functions such as packet classification, policing, shaping, queuing, and scheduling. We define two types of TC functions: 1) Traffic Shaping for end-systems such as a media source/sink device; and 2) Traffic Enforcement for network infrastructure devices such as residential gateways and wireless access points.
- *Policy Management:* Policy management refers to the ability to create, modify, and delete usage policies. A policy manager application residing on a device in the home would typically be able to evaluate policy requests against the policies it maintains, and provide decisions to a policy client upon request.

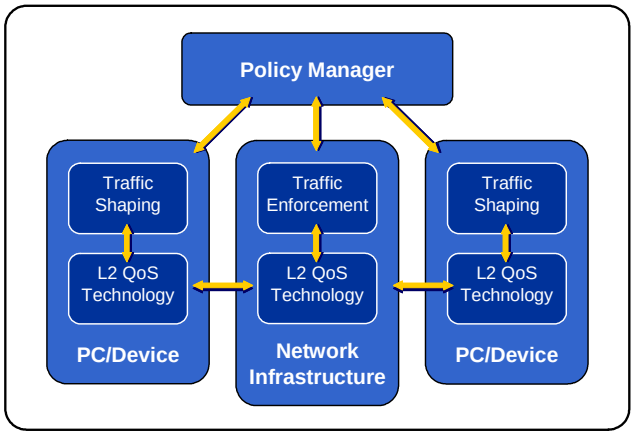

Figure 6: System architecture for QoS networking

Authentication and Authorization

An implicit assumption in the Digital Home is that devices and control points on the network can be trusted. It is possible, however, for unknown and potentially hostile entities to connect to a UPnP device network and listen to messages or control and query UPnP devices. It is important in such cases that devices and control points employ security mechanisms to verify the identities of each other and provide basic levels of access control for resources on the network. The solution for these devices is to use UPnP Security. UPnP Security protects UPnP messages from tampering and/or disclosure to untrusted parties, and it protects UPnP devices from unauthorized control operations. The primary focus of UPnP Security is access control, which is the ability to specify which control points can perform secured operations or receive sensitive information from specific devices. UPnP Security also permits fine-grained access control so that some control points can be granted limited access.

UPnP Security uses public-key cryptography to identify principals (entities to whom access is granted). However, unlike secure Web server public-key certificates, the public keys used in UPnP Security do not need to be issued or certified by a centralized authority. The decision to trust a public key is made by the device owner during device and/or control point installation.

Another important aspect of UPnP access control is the set of permissions that can be granted. The approach taken by UPnP Security is to define access control in terms of device-specific sets of abstract permissions. Each permission potentially enables one or more actions to be performed by a control point authorized by an entry in the device's Access Control List (ACL). Device manufacturers are free to define their own permissions and use whatever mechanism they want to map incoming control requests onto required permissions.

Gateway Management and Control

The term "Gateway" refers to an Internet access device that provides Wide Area Network (WAN) connectivity to the home LAN. In addition to addressing and naming services, it typically provides Network Address Translation (NAT) and firewall support. Ease-of-use for home networks can be greatly enhanced if applications can programmatically manage and query the gateway properties.

The gateway management and control building block enable applications on a home LAN to select and configure available WAN connections. It provides users with diagnostic information by maintaining the connection state and sending event notifications when appropriate. A LAN client can also configure LAN interface parameters, such as 802.11-related configurations. Although having a NAT on the gateway is commonplace and convenient, it breaks some important peer-to-peer applications. The Gateway Management building block provides a service for applications to dynamically configure automatic NAT traversal through port mapping for both inside-out and outside-in access. Applications can also query the gateway for the effective bandwidth on the WAN and LAN link and adjust their behavior dynamically.

It is important that a common management mechanism be used across gateway vendors and applications. Most of these capabilities are currently supported through the Internet Gateway Device (IGD) v1.0 specification defined recently in the UPnP Forum. This specification has been adopted by major gateway vendors including Linksys, Microsoft, and D-Link. Additional functionality, such as security bootstrapping for 802.11 access points, is being addressed by the IGDv2 working committee.

System Management

System management refers to the ability to manage the operational state of devices across the home network. To

enable this capability, all networked devices must support a common set of operational services. These system management services can be classified into the following categories:

- *Development, Test, and Manufacturing* services that assist device vendors during the design, test, and manufacturing stages of development.
- *Administration and Management* services that exist to ensure the correct run-time operation of a device, allowing for run-time diagnosis and correction of problems. This category also includes firmware and software upgrade operations.

APPLICATIONS AND SERVICES LAYER

The Applications and Services layer provides a number of functions that mostly deal with the applications and services running on a Digital Home device. This includes coordination between the underlying platform middleware building blocks and interaction with the user through a User Interface (UI). Most importantly, the applications and services layer is responsible for performing content protection and digital rights management for multimedia content on a device.

Content Protection and Digital Rights Management

Commercial content providers require technical and legal protection for their digital content to prevent unauthorized access and copying. PC and CE manufacturers are working with the content industry to develop content protection solutions, which typically consist of a combination of technical and legal mechanisms.

Technical protection mechanisms are effective at preventing unsophisticated attempts to circumvent a particular content protection solution. Licensing and other legal mechanisms are much more effective against business entities with assets, employees, and distribution channels than they are with individuals. Accordingly, technical mechanisms provide the basis for content protection, and an effective licensing structure provides for enforcement.

BENEFITS OF THE HOME INTEROPERABILITY FRAMEWORK

The Home Interoperability Framework (HIF) provides a set of building blocks that allow devices to perform desired functions in an interoperable fashion. Clearly, not all devices need to implement all HIF building blocks. For example, a gateway device may need to support both a LAN and a WAN connection. In contrast, an audio player device need only support a LAN connection. Similarly, selection of other HIF building blocks depends on the overall functionality of a particular device in the home, whether it be a PC, a gateway, or an embedded device that performs a limited set of functions.

The HIF provides various benefits for the constituents of the Digital Home entertainment value chain, including consumers, content providers, service providers, CE device manufacturers, and application developers.

- *Benefits for Consumers***:** The HIF allows devices in the home to collaborate and provide consumers with compelling and innovative entertainment experiences. The vendor interoperability features of HIF provide consumers with greater flexibility in selecting their entertainment equipment. This is a clear benefit for consumers, and has the potential to drive consumer demand for more content, devices, and services.
- *Benefits for Content Providers:* For content providers, HIF provides support for content protection mechanisms that help protect their high-value assets. The vendor interoperability features of HIF also provide content providers with assurances that their content can be delivered to various devices in the home in a protected fashion.
- *Benefits for Service Providers:* The HIF building blocks provide interoperable technical solutions that eliminate barriers for secure end-to-end connectivity and high-quality media streaming. Therefore, the HIF allows content and services to be delivered to more end-points in the home, increasing revenue opportunities for service providers.
- *Benefits for CE Device Manufacturers and Application Developers:* The HIF allows CE device manufacturers and application developers to build interoperability into their devices and applications. The resulting interoperability features will fuel more demand for devices and applications in the home. Clearly, this has the potential to create numerous business opportunities for both PC and CE vendors. The framework enables device vendors to provide new product features and customized interfaces for market differentiation.

CASE STUDY

The following case study uses one of the usage scenarios described earlier. It examines the scenario in which Jim decides to watch a movie on the TV and at the same time Pam decides to watch a documentary on the PC in the den. It is assumed that Jim's movie is stored on the Personal Video Recorder (PVR) in the living room and Pam's documentary is stored on another PC in the bedroom. The overall theory of operation is reviewed to better illustrate the effectiveness of the Home Interoperability Framework (HIF) in providing vendor interoperability.

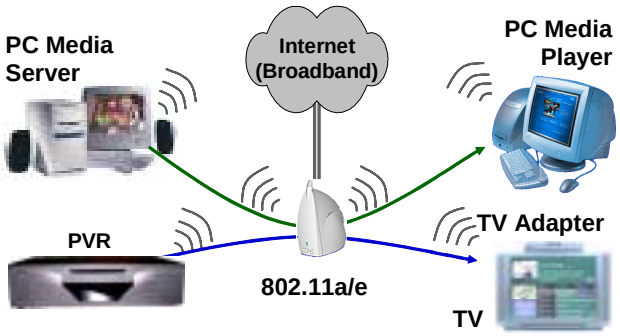

Figure 7: Equipment configuration for the case study

This case study involves a PVR acting as the media server for Jim's movie and a TV adapter that connects the legacy TV to the home network. It also involves the bedroom PC acting as a media server for Pam's documentary and the PC in the den acting as the media player. The scenario, as illustrated in Figure 7, also includes a wireless gateway acting as the access device to the Internet and as an 802.11a wireless access point with 802.11e Quality of Service (QoS) support.

From a functionality point of view, the PVR and the TV adapter perform similar functions for Jim as the PC server and the PC media player do for Pam. To avoid redundancy, we focus on Jim's movie session and describe the interaction between the PVR, the wireless gateway, and the TV adapter.

The Scenario

Jim finds his favorite chair in the living room occupied by his 11-year-old. He heads to the kitchen, browses the movie collection stored on all the devices in the home and decides to watch a war movie, *Gladiator*. He makes himself a sandwich and pushes the Play button on the kitchen remote.

Jim and Pam subscribe to a premium video-on-demand service that allows them to watch three movies per week for a flat rate, and any additional movies for an extra charge. The service allows movies to be downloaded overnight over the broadband Internet connection and stored on either the PVR or one of the PCs in encrypted format with appropriate Digital Rights Management

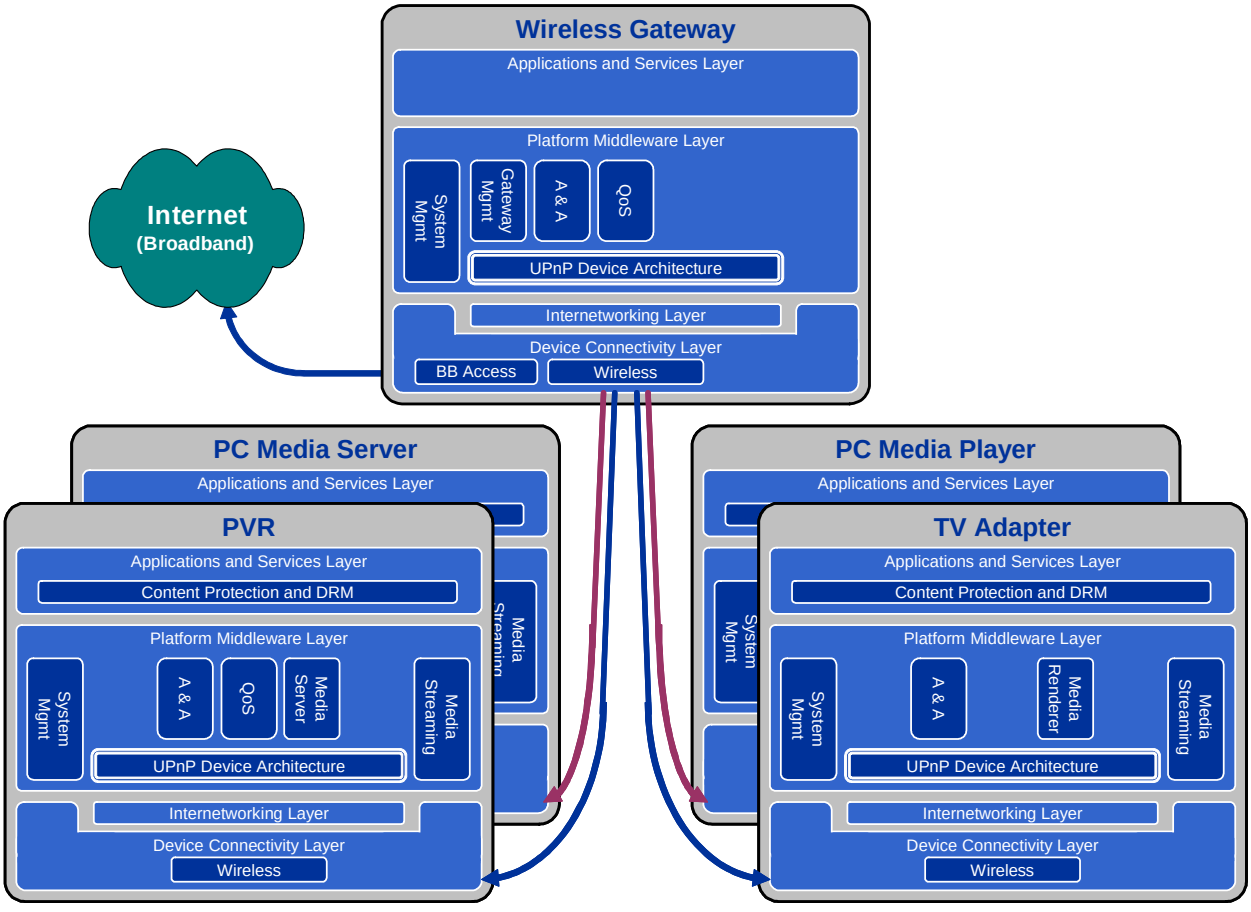

Figure 8: HIF Device Configurations

DRM) credentials. Billing is not activated until a movie is decrypted at playback time.

Figure 8 shows the HIF tailored to the desired functionality. The PVR, shown in the bottom left corner of the figure, acts as a UPnP* AV media server, with 802.11e QoS support over the wireless 802.11a connection; the QoS server is responsible for tagging all outgoing packets with appropriate priorities.

The wireless gateway, shown at the top of Figure 8 is an 802.11a access point with 802.11e QoS support and an additional DSL WAN interface supplying the connection to the broadband Internet. The wireless gateway serves as a UPnP Internet Gateway Device (IGD) for management and control of gateway services (NAT, DNS, Firewall, and so forth) to allow secure and transparent end-to-end connectivity to the WAN. The wireless gateway also supports the policy management service for the home LAN.

The TV adapter shown in the bottom right corner of Figure 8 has several functions. It connects the legacy TV to the wireless home LAN and is responsible for receiving, decoding, and playing back the movie over its video output connected to the TV. It acts as an UPnP AV media renderer, performing other functions such as content protection (DRM and media decryption), and it provides the user (Jim) with an electronic program guide for browsing content on all devices in the home. All three devices participating in this service support UPnP Security for device authentication and authorization.

Finally, in this usage scenario it is assumed that the TV adapter application includes a UPnP control point that is used to both browse the content on the home network and to control streaming sessions.

Theory of Operation

This section describes in detail the sequence of steps the UPnP control point implements and how the user interacts with it. The operation is divided into three phases: discovery, browsing and selecting content, and playback.

Discovery Phase

When the TV adapter is activated, its control point searches and discovers all UPnP devices and services on the home network. As a result, it discovers the PVR media server with support for QoS and security. It also discovers the wireless gateway as a UPnP IGD device with support for UPnP Security and policy management services.

* Other brands and names are the property of their respective owners.

Browsing and Selecting Phase

As soon as Jim points the kitchen remote at the TV and starts interacting with it, the TV adapter queries the UPnP AV content directory service running on all media servers in the home—including the living room PVR—and requests a list of movies that are available for viewing. The TV adapter presents an electronic program guide on the TV that allows Jim to select one of the movies to watch. As Jim navigates through the list of interesting titles, he finds the movie *Gladiator* and selects it for viewing. Jim may want to learn more about the movie by pressing the Info button on his remote, which then triggers the TV adapter to display a synopsis of the movie plot, the number of chapters, and languages supported by the soundtrack. Jim can also see if the movie is available for free viewing or whether he will be charged an extra fee for watching it.

During the process of finding movie content on the PVR, the TV adapter also determines certain technical details about the movie, such as information about the digital format used to encode the movie. Jim never sees this information, but it is used later by the TV adapter and the PVR media server to set up the correct network connection and playback configuration when the movie actually starts streaming.

Jim, unaware of the extra work being done by the TV adapter on his behalf, decides to start watching *Gladiator* from chapter 15 and settles back in his chair. The navigation and title information are not encrypted, so Jim is able to browse through his movie list without interacting with the DRM even though *Gladiator* is commercial content.

Playback Phase

When Jim presses the Play button on the remote to begin playing the movie on his TV, a number of activities are triggered. On the server, the Content Protection/DRM module makes an entry in a billing log to identify the movie being watched. Later, this billing log will be sent to a central server where Jim's credit card will be charged for the appropriate amount. The control point invokes the connection manager services on both the TV adapter and the PVR to set up the connection. It requests a list of streaming protocols and the supported media formats. After processing the response, the control point selects HTTP streaming and the MPEG2 format.

On the media server, the DRM retrieves the key used to unlock *Gladiator* from its database and securely sends that key to the TV adapter. Now the application can begin to stream the encrypted movie to the TV adapter. When the TV adapter receives the encrypted buffers, they can be decrypted and then processed. The control point also

invokes the AV Transport service on the TV adapter and sends it a "play" command, which triggers the media streaming engine to start decoding and playback.

Variation of the Above Scenario

Instead of the TV adapter acting as a control point, it is possible for the PVR Media Server to perform this role and provide the user interface that is rendered by the TV adapter. Remote I/O supports this capability. Using Remote I/O, user input is collected from the remote control by the TV adapter and forwarded to the PVR. The PVR provides appropriate user interface screens for display by the TV adapter.

CONCLUSION

Personal Computer (PC) and Consumer Electronic (CE) collaboration creates significant new business opportunities for both the PC and CE industries. Both industries are actively working to capitalize on these opportunities. An important key to success for these new opportunities is vendor interoperability.

This paper covered the basic interoperability requirements for the Digital Home and proposed a Home Interoperability Framework (HIF) for creating innovative and interoperable solutions that can greatly enhance the consumer's digital entertainment experience. The HIF is based on well-established standards for home networking such as 802.11 wireless LANs, IP-based internetworking, and the UPnP* architecture.

ACKNOWLEDGMENTS

This paper is the product of the collective thinking of many architects inside Intel Labs and Product divisions, too many to list here. These individuals deserve much credit for continuing to push the Digital Home vision for Intel and for the industry at large.

REFERENCES

[1] UPnP Forum, http://www.upnp.org

[2] UPnP Technology, Intel Research and Development, http://www.intel.com/labs/connectivity/upnp/

AUTHORS' BIOGRAPHIES

Yasser Rasheed is a Network Architect in the Corporate Technology Group, Intel Corporation. He has been with Intel for two years, and has been involved in a number of projects related to home networking and advanced media streaming over wired and wireless networks. His interests include UPnP technology, Quality of Service, and policy-based networking for real-time applications. Yasser holds a B.Sc. degree in Electrical Engineering from Cairo University, Egypt, and M.Sc. and Ph.D. degrees in Electrical and Computer Engineering from the University of Toronto, Canada. Yasser's e-mail is yasser.rasheed@intel.com

Jim Edwards is a lead architect and a technology development manager with Intel Corporation's Solutions Architecture lab in Hillsboro, Oregon, working on the architecture and technologies for the Digital Home. Solutions Architecture is a part of Desktop Platform Architecture, which is an architecture organization within the Desktop Platform Group responsible for defining and developing platforms of the future. His interests include home networking, digital media, and sports cars. Jim holds a B.S. and M.S. degree in Electrical Engineering from Cornell University. His e-mail is jim.edwards@intel.com

Charlie Tai is a Principal Engineer in the Corporate Technology Group, Intel Corporation, and has focused on communications and networking technologies and strategies. His interests include home networking, ad-hoc networking, broadband network, Quality of Service and traffic management. Charlie has represented Intel in standards and industry organizations such as the Winsock 2 Forum, ATM Forum, UAWG, DSL Forum, USB Forum, IETF, and the UPnP Forum. He has been on the UPnP Forum Steering Committee and on the Board of Directors of the UPnP Implementers Corporation. He was instrumental in the establishment of the Internet Gateway, Audio-Video, and Security working committees in the UPnP Forum. Charlie received his B.S. degree in Computer Science from National Taiwan University and an M.S. and Ph.D. degree in Computer Science from UCLA. His e-mail is charlie.tai@intel.com.

* Other brands and names are the property of their respective owners.

High-Quality Media Distribution in the Digital Home

Yasser Rasheed, Corporate Technology Group, Intel Corporation
John Ritchie, Desktop Platform Group, Intel Corporation

Index words: UPnP technology, UPnP AV, Media distribution, QoS, Digital Home, Interoperability

ABSTRACT

The proliferation of digital technologies within the home is paving the way for significant enhancements to the user's entertainment experience. A fundamental aspect of these new experiences is the ability to enjoy rich multimedia content in any location throughout the home, regardless of where the content is physically stored. Today's home networking technologies provide a solid foundation for distributing content throughout the home and achieving this type of "anytime, anywhere" access to content. In order to satisfy users' quality expectations, the ensuing demand for Quality of Service (QoS) networking, especially for wireless Local Area Network (WLAN) technologies, is becoming an area of high interest for content and service providers who are aggressively seeking opportunities to capitalize on the new paradigm. In reality however, media distribution and QoS networking technologies may not be ready yet for mass-market adoption due to the complexity associated with the installation, configuration, and management. This paper identifies a high-quality media distribution solution that provides consumers with a simple out-of-the-box installation and configuration experience coupled with the ability to manage the limited network bandwidth in order to achieve predictable results. We first analyze a typical Digital Home environment and discuss possible usage scenarios for media distribution in the home. We also provide an overview of Digital Home technologies related to high-quality media distribution, including UPnP* and UPnP Audio Visual (AV) technologies, and existing QoS technologies in both wired and wireless LANs. We then describe a novel QoS networking framework that integrates existing QoS technologies and provides application developers with simple interfaces to add QoS support to their applications based on UPnP technology. The result is a comprehensive solution that allows end-users to experience high-quality media distribution in a vendor interoperable fashion.

INTRODUCTION

As consumers continue to purchase digital multimedia products for their home, they increasingly gain access to a wide range of rich multimedia content from a variety of sources. These sources include Personal Computers (PCs), the Internet, digital terrestrial broadcast receivers, satellite receivers, and Consumer Electronics (CE) devices including CD/DVD players, camcorders, digital still cameras, portable audio players, etc. The physical location of these content sources is often not the user's preferred viewing location. For example, users may prefer to listen to their MP3 music collection on their living room stereo instead of having to sit in the den where the PC is typically located.

Although today's home networking technologies provide a solid foundation for distributing content throughout the home, some key barriers remain, which prevent the wide-spread adoption of high-quality media distribution solutions for consumers masses. The two primary barriers include: 1) the installation and configuration complexities and/or cost, and 2) the complexities associated with managing the relatively limited network bandwidth in order to reliably meet the user's quality expectations.

A Typical Digital Home Environment

Inside the home, a growing number of devices are getting connected together, forming a "Digital Home (DH) environment." The DH environment typically consists of multiple networked devices including PCs, CE devices, such as TVs and stereos. Legacy (analog) devices may also be connected to the Digital Home environment by special-purpose adapters.

Figure 1 shows a typical Digital Home environment that includes two TVs, each with a Set-top Box (STB)

* Other brands and names are the property of their respective owners

connected to the home LAN, two desktop PCs and a laptop. The STB connected to TV1 (STB1) is also connected to a satellite dish receiver, acting as a source for digital broadcast content. The home network in this example is a wireless LAN consisting of a wireless access point that acts as a Residential Gateway (RG) that connects the home network to the Internet through a broadband access pipe.

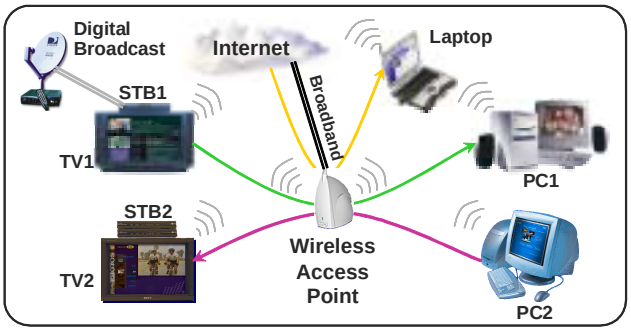

Figure 1: Digital Home environment

With the proliferation of Digital Home devices, we expect the users to explore and make the best use of their Digital Home environment. As a result, consumers will attempt to run various types of applications, simultaneously, over the home network. While many applications may be able to co-exist over the home LAN, many others might suffer from temporary shortage of bandwidth capacity. As a result, consumers might experience unpredictable and unsatisfactory behavior from those applications.

We present below usage scenarios that describe situations where consumers might run into difficulties. These scenarios illustrate the need for QoS networking to achieve high-quality media distribution in the home.

Multiple Simultaneous Video Streams

The configuration in the DH environment shown in Figure 1 allows consumers to experience simultaneous transport of multiple video streams over the wireless home LAN. For example Mom may be recording a TV documentary from TV1 to PC1, while the kids use TV2 to watch a pre-recorded sports program stored on PC2. During the same time, Dad may be working from home and viewing a live corporate video on his laptop that is connected to the corporate network using a Virtual Private Network (VPN) connection.

Technical Challenges

For most consumers, the Digital Home environment is built one device at a time over a number of years. When new devices are added, the challenge is for existing devices to recognize that a new device has appeared on the network, and to determine how it can interoperate with it. Similarly, the new device may need to discover existing devices on the network, and determine how to interoperate with them.

For example, when PC2 was originally added to the network, the two TVs may have needed to know then that PC2 became available to serve new content. Additionally, the TVs may also have needed to detect the new content available on PC2. This clearly highlights the need for a standard framework for automatic discovery, configuration and control.

Consumers expect the quality of home media streaming applications to be at the same level as they experience with CE devices, such as their TVs/VCRs, DVD players, etc. The multitude of applications running simultaneously over the wireless LAN in Figure 1 raises several challenging questions:

- Can the wireless LAN support all the potential video streams at the same time, without causing any application to break?
- Will the perceived quality of the video streams remain consistent in the presence of bandwidth contention created by other video streams?
- Some streams may be more important than others, or have stricter real-time transport requirements than others, and some may be more suitable for graceful degradation. Is it possible then to manage the bandwidth in such a way that certain streams are treated more favorably than others?
- Before they start an application, users may prefer to get assurances that the home environment can (and will) support it. Is it possible then to allow these applications to detect whether the capacity on the network is available to start a new stream?

These challenging questions highlight the need for bandwidth management and QoS-based networking. They also highlight the need for a home policy server that maintains and enforces policies such as giving certain traffic flows higher priority than others.

These scenarios also show the importance of ease-of-use to hide the sophistication involved in the setup and configuration of PCs, STBs and the wireless access point from the user. With minimal interaction, the user should be provided with the best Digital Home experience.

In the following sections, we describe a Media Distribution Architecture that allows devices to independently discover and interoperate with each other in order to stream content from one device to another, based on UPnP and UPnP AV technologies.

We then present an overview of wired and wireless home networking link-layer technologies, and existing QoS support at various layers of the network stack.

Subsequently, we describe a QoS framework that integrates the existing QoS technologies and provides a complete solution based on UPnP technology for high-quality media distribution in the home.

Finally, we use a case study based on the home environment in Figure 1 to illustrate the theory of operation of the overall system that can deliver a high-quality entertainment experience for consumers.

THE UPNP* TECHNOLOGY

Most media distribution systems are based on custom, end-to-end solutions that use proprietary technologies and require a trained specialist to set up and configure the system. These factors prevent such solutions from reaching price-points and convenience levels that are needed for mass-market adoption of these systems.

The key factor affecting both price-point and convenience is the involvement of a professional installer. In order to enable wide-spread adoption, these systems need to be installable and configurable by the end-user without help from a service representative just like today's VCR. This type of 'out-of-box' installation experience requires individual devices to be self-configurable. Additionally, devices need to discover other devices on the network that it can interact with in an autonomous manner.

This type of self-configuring system requires the adoption of a single interoperability technology that is pervasive throughout the industry (i.e., adopted by a large number of device manufacturers and supported by a wide range of devices). There have been many efforts in the past to define such interoperability technologies. Although these efforts adequately solve many of the technical issues, their lack of critical mass within the industry limits their success in the mass-market.

A survey of the current status of various interoperability technologies reveals that the UPnP specification offers the greatest promise for an interoperability technology that will actually be adopted and deployed by the key market-making leaders within the industry.

The UPnP device architecture specification [2] defines general interoperability mechanisms that enable self-configuring devices to create an ad-hoc, self-discovering system of interoperable network devices. This specification defines mechanisms for automatic address configuration, device discovery, command/control, and eventing. The UPnP specification also defines "Presentation Pages" that allow devices to expose dedicated Web pages for user-initiated interaction with a specific device.

UPnP Fundamentals

UPnP technology uses existing Internet standards including TCP/IP, HTTP, SSDP, SOAP, GENA, XML, etc. These open standards provide the communication infrastructure of the UPnP architecture. Although the UPnP architecture consists of a peer-to-peer network, nodes on the network communicate with each other in a client-server manner. Clients are called *Control Points* (CP) and typically provide a User Interface (UI) for end-users. Servers are called *Controlled Devices* (henceforth, called devices) and by definition expose a well-defined set of functions called *actions*. In all cases, Control Points invoke actions, and devices respond to actions that are received.

Within the UPnP architecture, device functionality is exposed using a set of *services*, each of which corresponds to a functional component of the device. Each service defines a set of 'state variables' and 'actions' that allow Control Points to obtain the current state of the device and to control the device's operation. Invoking an action usually causes a change in the internal state of the device that would affect the value of certain state variables.

In order to enable autonomous device interoperability, members of the UPnP Forum [1] constructed a set of device and service definitions (a.k.a. templates) which can be used to model various common devices. Since the behavior of these device and service templates is well defined, Control Points can interoperate with any device that implements the services that are supported by the Control Point. In this manner, Control Points and devices can be built independently by different manufacturers with the assurance that they will interoperate according to the functionality defined by the corresponding UPnP device/service templates.

Network Addressing

Since the UPnP architecture is built on top of the Internet Protocol IP, each node in the network requires a unique IP address. This address is assigned either via a Dynamic Host Configuration Protocol (DHCP) server or via the 'Auto-IP' protocol if a DHCP server is not available. When a DHCP service becomes available, all nodes are required to obtain an address from it. Once a device or a Control Point has been assigned an address, it is considered "added" to the network.

Discovery

When a Control Point is added to the network, it needs to discover (i.e., locate) the devices in the network that it is

* Other brands and names are the property of their respective owners

capable of controlling. This is accomplished via the Simple Service Discovery Protocol (SSDP) by broadcasting a discovery request that identifies the functional capabilities that the Control Point wants to control. Any device that exposes those capabilities responds to the request by identifying itself to the Control Point.

The device response contains the URL of the "XML device description document," which identifies the services that the device implements, as well as the specific actions and state variables that are supported by each service. By parsing this information, the Control Point is able to determine the exact capabilities of each device. This allows a Control Point to determine if it wants to interact with and control a particular device.

When a new device is added to the network, the device may broadcast an identification notification to the network. This notification informs existing Control Points that a new device has been added to the network and is available to be controlled. The notification information includes the URL of the new device's description document, as described above.

Command/Control

Once a Control Point has determined that it wants to control a particular device, the Control Point uses a Simple Object Access Protocol (SOAP) to invoke any of the actions exposed by the device's services. The behavior of each action is well defined by the service template document.

Eventing

As the internal state of a device changes, either in response to an action or via some internal condition, the device can inform one or more Control Points of the state change using Generic Event Notification Architecture (GENA). With this protocol, Control Points that desire to be informed of state changes within a particular device must register with that device to receive event notifications. A given device may be monitored by multiple Control Points. When an internal state change occurs, the device sends an event notification to each Control Point that has registered with the device. This event notification includes an identification of the state variable that has changed, along with its new value. The set of state variables that are evented by the device is defined in each of the service templates that are supported by the device. Additionally, each evented state variable may be moderated such that rapid changes in that state variable do not cause excessive network traffic.

THE UPnP* AV SPECIFICATIONS

The UPnP AV specifications [3] define a set of UPnP device and service templates that specifically target Consumer Electronics (CE) devices such as TVs, VCRs, DVD players, stereo systems, MP3 players, and so forth. In this context, a CE device refers to any device that interacts with entertainment content (e.g., movies, audio, and still images) which includes the PC.

In today's non-networked CE environment, CE devices interoperate with each other using dedicated cables. The UPnP AV specification enables CE devices to use the network instead of these dedicated cables to interoperate with each other. This network-wide interoperability allows CE devices to distribute entertainment content throughout the home network.

UPnP AV Fundamentals

The UPnP AV architecture shown in Figure 2 defines three main logical entities: a Media Server, a Media Renderer, and a UPnP AV Control Point. The Media Server has access to entertainment content and can send that content to another UPnP AV device via the network. A Media Renderer is able to receive external content from the network and render it on its local hardware. An AV Control Point coordinates the operation of the Media Server and Media Renderer in order to accomplish the desires of the end-user.

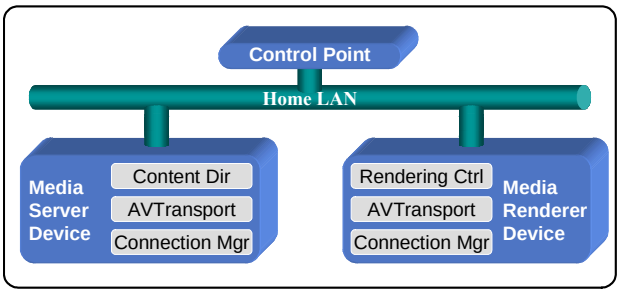

Figure 2: UPnP AV architecture

As described later in this section, Media Servers and Media Renderers implement a set of UPnP AV services. These services provide command and control functions that allow a Control Point to set up and configure the Server and Renderer for transferring the desired content from the Server to the Renderer.

The Control Point is involved only in command and control operations. It is not involved in the actual transfer of the content. Therefore, the Control Point, and hence

* Other brands and names are the property of their respective owners

the entire AV architecture, is not dependent on any particular transfer protocol and/or content data format.

Since the AV architecture can accommodate various transfer protocols and content formats, servers and renderers can transfer the desired content using any transfer protocol and data format that they *both* support. As part of its set-up and configuration responsibilities, the Control Point must identify and select which protocol and format is to be used. However, since the Control Point is not involved in the actual transfer, it does not need to implement the selected transfer protocol or data formats.

Although the AV architecture defines these three logical entities, a physical device may contain any combination of them. For example, many Renderers are likely to include an embedded Control Point so that the user may control the operation from the same location where the content is rendered.

UPnP AV Control Points

As described above, AV Control Points control the operation of the Media Servers and Media Renderers so that the user can render specific content on a particular rendering device. In most end-user scenarios, the Control Point uses a variation of the following algorithm:

- Locate the existing Server/Renderer devices in the network, i.e., discovery
- Enumerate the available content for the user to choose from, i.e., content enumeration
- Query the Server and Renderer to find a common transfer protocol and data format for the selected content, i.e., protocol/format negotiation
- Configure the Server and Renderer with the desired content and selected protocol/format, i.e., Server/Renderer setup
- Initiate the transfer of the content according to the desires of the users, such as Play, Pause, Seek, and so forth, i.e., control content flow
- Adjust how the content is rendered by the Renderer, such as Volume, Brightness, and so forth, i.e., control rendering characteristics

The Control Point accomplishes this general algorithm by invoking various actions on UPnP AV services exposed by the Server and Renderer. In this manner, the Control Point can perform the content distribution tasks that are desired by the user.

Media Server

A Media Server is a device that has access to entertainment content and can send that content to another device for rendering. Media Servers include familiar devices such as VCRs, set-top boxes (cable, satellite, digital broadcast, etc.), camcorders, CD/DVD players/jukeboxes, radio tuners, TV tuners, still-image cameras, etc.

Media Servers expose the Content Directory, Connection Manager, and (optionally) the AV Transport services. The Content Directory service allows a Control Point to discover and enumerate all of the content that is accessible by the Server. The Connection Manager service allows the Control Point to negotiate and select the common transfer protocol and data format that will be used by the Server and Renderer to transfer the desired content. The (optional) AV Transport service is used to control the flow of the content (e.g., play, stop, pause, etc.).

Media Renderer

A Media Renderer is a device that can receive content from another device and render it using some local hardware. This includes familiar devices such as a TV, a stereo system, a set of speakers, an Electronic Picture Frame (EPF), etc. Innovative Renderers can use any type of output hardware that can be controlled by the incoming content. For example, a "Music Fountain" can generate dancing streams of water based on the content of a song.

Each Renderer exposes the Rendering Control, Connection Manager, and (optionally) the AV Transport services. The Rendering Control service controls how the content is rendered (e.g., Volume, Brightness, etc.) As with the Media Server, the Connection Manager service is used to negotiate a common protocol/format, and the (optional) AV Transport service is used to control the flow of the content.

Content Directory Service

The Content Directory Service (CDS) allows Control Points to discover and enumerate content that is accessible by a Media Server. CDS "content" objects include individual "content items," which represent individual pieces of content such as a song, video clip, or a photo; and "content containers," which represent collections of items such as a playlist, CD, or a photo album. Each CDS object, either an item or container, includes meta-data that describe various attributes of the object, such as title, artist, duration, and so forth.

CDS provides both *Browse* and *Search* capabilities. Control Points that browse a CDS begin at the root of the CDS hierarchy and iteratively examine the structure, container by container, until the desired content item is found. This is similar to how a file system is used to locate a file that is nested several layers down from the root directory. Control Points typically use this

mechanism when the user does not immediately have a particular content item in mind.

Alternatively, a Control Point can use the CDS's Search capability to locate all of the items/containers that possess certain attributes (i.e., certain meta-data values such as "creator=Disney").

Part of the meta-data for each object is a list of transfer protocol and data format combinations that are supported for that piece of content. This information is used by the Control Point in conjunction with the Connection Manager service on the target Media Renderer to determine which protocols and formats can be used to transfer the content to the Renderer. Each protocol/format combination is identified by a unique Universal Resource Identifier (URI). This URI is used by the Control Point to identify the content, protocol, and format that are to be used during the transfer.

The data structures defined by CDS dictate the over-the-wire representation of content items/containers and associated meta-data. CDS does not define the Media Server's internal storage mechanisms or structures. Media Servers can store CDS information using any appropriate mechanism. For example, some servers may use a full-featured database system to store its CDS content hierarchy and to provide a rich set of meta-data for each object. Other servers may maintain their CDS information using only a directory/file hierarchy of their internal file system. In this case, the breadth of the meta-data for each object will be limited to the information that is stored by the file system for each directory/file.

As CDS requests are made by the Control Point, the server converts those requests into a set of operations carried out by the underlying database used to store the CDS hierarchy and meta-data. Depending on the Server's implementation, this may be a set of relational data operations or a set of system calls to the local file system.

Rendering Control Service

The Rendering Control Service (RCS) is implemented on a Media Renderer in order to provide the Control Point with a mechanism to control how the content is rendered (e.g., volume, brightness, contrast, etc.). These functions are directly related to the capabilities of the output hardware on the Renderer.

The internal logic of RCS is fairly simplistic. As RCS actions are invoked, the RCS simply converts the requested adjustment to the corresponding hardware request as needed.

Connection Manager Service

The Connection Manager (CM) service is implemented both on the Server and Renderer. The primary purpose of the CM service is to allow the Control Point to identify and select the common protocol/format that will be used to transfer the desired content from the Server to the Renderer.

The actions defined by the CM service provide a standardized interface to the Server's/Renderer's internal network and media codec subsystems. An implementation of CM must be able to enumerate and configure the transfer protocols that are supported by the device's network subsystem and the data formats that are supported by the device's media codecs.

In order to enumerate the list of supported protocols/formats, the CM service on most fixed-function devices (e.g., traditional CE equipment) is fairly simple since the set of supported protocols/formats is fixed (e.g., it is known what the device is designed for). For other general-purpose devices, such as the PC, the network and codec subsystems typically provide internal interfaces that allow a CM service to discover the network protocols and codec models that have been dynamically installed on the device.

When preparing to transfer a piece of content from the server to the renderer, the CM service must be able to set up and configure its network and codec subsystems according to the requested mechanism. In many cases, this may involve constructing a data path from the device's network subsystem, through the appropriate codec, to the device's output hardware. In many implementations, the CM service uses a set of pluggable modules that provide a data path from a source module (e.g., the network interface card), through zero or more intermediate modules (e.g., a codec), and finally to a sink module (e.g., the device's output hardware). A popular example of a pluggable media streaming engine is Microsoft's DirectX*. With this technology, individual filters correspond to particular media streaming functions (e.g., capture a data stream from the network, decode and/or transform the stream, etc.) Individual filters are plugged together to form a complete data path called a filter graph.

AV Transport Service

The AV Transport (AVT) service provides a number of actions that allow a Control Point to control the flow of the content. This includes many of the familiar operations typically associated with the mechanical "tape transport"

* Other brands and names are the property of their respective owners.

mechanism implemented on most VCRs, such as Play, Stop, Pause, Seek, FF/REW, and so forth.

AVT also provides the mechanism which the Control Point uses to identify the content that is to be played. This is done by passing in the unique Universal Resource Identifier (URI), which was obtained from the Content Directory Service for the desired content and the selected protocol and format.

Depending on which transfer protocol is used to transfer the content, either the server or the renderer may provide an instance of the AVT service. If the selected protocol is a "pull" model (e.g., HTTP GET), then the renderer is required to provide an instance of AVT to control the flow of the content (e.g., play, pause, seek). If the selected protocol is a "push" model, then the server must provide an instance of AVT.

The internal implementation of AVT must hook into the device's media streaming subsystem in order to configure it to access and stream the desired content and to control that content stream, as directed by the Control Point (i.e., the end-user). In most cases, the internal logic of the AVT service is fairly straightforward. When an AVT action is invoked, AVT invokes the corresponding operation(s) on the internal media streaming subsystem. When the device is using a pluggable media streaming technology like DirectX, AVT simply invokes the appropriate method(s) on the appropriate filter graph.

UPnP AV Control Point Algorithm

As described above, the UPnP AV architecture defines the external interfaces of the media server and media renderer so that a Control Point (CP) can manage the distribution of entertainment content as desired by the end-user. However, the AV architecture does not define any of the internal structure of the server/renderer. This is left entirely to the implementer. Nevertheless, in practice, there are some general implementation models that will be commonplace. We outline below some of these models.

When a Control Point joins the network, it locates all of the media servers and media renderers in the network. It does this using Simple Service Discovery Protocol (SSDP). In order to locate Media Servers in the network, the Control Point issues an SSDP IP-multicast request packet for any UPnP device that implements the UPnP AV Media Server device template. All devices that implement the media server template must respond to the request with the URL of its description document. Media renderers are located in a similar manner.

Once servers and renderers are located, the Control Point obtains and parses each device's XML description document to determine the device's exact capabilities (i.e., its UPnP services, actions, and state variables.) If the device implements the desired capabilities, the Control Point continues to interact with it as described below.

At some point after the Control Point initializes itself, a Control Point may display an initial User Interface (UI) so that the end-user can interact with the Control Point. The contents and layout of the UI is device-dependent to provide room for innovation and product differentiation.

For each media server that is found, the Control Point uses the server's Content Directory Service (CDS) to enumerate the content that is available from that server. Control Points often collect CDS information from multiple servers and aggregate it into a single "whole home" view of all of the content that is available from within the home, regardless of which service it is on.

Depending on the Control Point's UI, the CP will either browse through the CDS information, perform searches on it, or a combination of the above. Once the CP has received and processed the returned data, the Control Point updates its UI.

After the user has selected the desired content, the Control Point determines which transfer protocols and data formats are supported for that particular piece of content. This is done by examining the CDS meta-data for the selected item. Using the Connection Manager service on each renderer, the Control Point can obtain the set of protocols/formats that are supported. The Control Point then compares the protocol/format information from the server's CDS and the renderer's CM to determine which renderer(s) is capable of rendering the desired content.

After a common protocol/format has been identified, the Control Point invokes the CM services on both the server and renderer to inform each device of the target protocol/format. In response, the CM would set up and configure its internal network and media streaming subsystems based on the common protocol/format that has been chosen.

As a result of configuring each device, either the server or renderer will return an instance of the AV Transport (AVT) service that is associated with the media stream that has just been set up. The Control Point uses the returned AVT to specify the content that is to be transferred from the server to the renderer.

When the user indicates the desired operation that is to be performed on the current content (e.g., Play, Seek, etc.), the AVT service is invoked accordingly. After the content has begun to play, the user may select other operations (Stop, Pause, Seek, etc.), any time during the streaming session.

As the content is being rendered, the Control Point may provide a set of UI components that allow the user to control how the content is rendered. This includes various

rendering characteristics such as loudness of the volume, brightness of the video/image, etc. As the user adjusts various rendering characteristics, the Control Point invokes the appropriate action on the Rendering Control Service (RCS) as appropriate.

The core capabilities needed to distribute and render the selected content on a particular rendering device is provided by the UPnP AV architecture. The Control Point developer can concentrate on providing innovative and compelling media distribution UIs for the end-user.

Summary of the UPnP AV Architecture

The Media Distribution Architecture described above enables device manufacturers to develop and deploy interoperable multimedia products that are trivial for mass-market consumers to self-install and self-configure. This allows consumers to distribute their digital entertainment content throughout the home network. However, this architecture does not ensure high-quality distribution. In order to achieve a reliable high-quality experience, additional mechanisms must accompany the general media distribution architecture.

HIGH-QUALITY MEDIA DISTRIBUTION INFRASTRUCTURE

The presence of a home networking infrastructure is essential for the realization of media distribution in the home. Fundamentally, the home network must provide the capacity to distribute multimedia traffic over the network at the quality level that consumers expect.

Figure 3 shows the building blocks that enable high-quality media distribution. We divide the architecture into three functional layers: 1) A network subsystem layer that includes wired/wireless link-layer and IP internetworking technologies; 2) A middleware layer that includes the UPnP device architecture, UPnP* AV, media streaming components and a Quality of Service (QoS) networking building block; and 3) an application layer which includes the media distribution application.

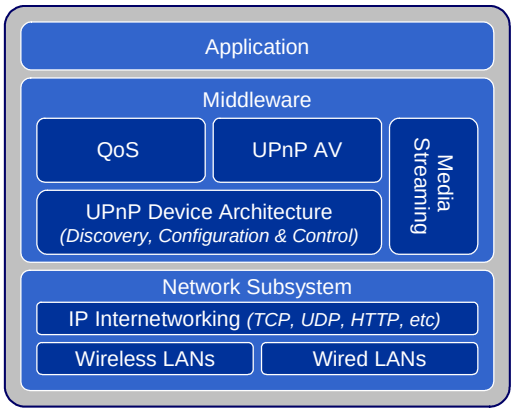

Figure 3: High-quality media distribution architecture

In the following, we start with background information on different wired and wireless home networking solutions. We then provide an overview of existing QoS technologies in the home at different layers of the network stack. In the next section, we present a framework for QoS networking based on UPnP technology. This framework integrates existing QoS technology building blocks and uses UPnP technology to provide automatic discovery, configuration, and control features.

Home Networking Link-Layer Technologies

A variety of home networking technologies exists today, including both wired and wireless solutions, which typically vary in terms of bandwidth capacity, wiring requirements, availability in a certain home environment, and above all, price. We describe below a number of these wired and wireless link-layer technologies for the home.

Wired LANs

Wired home networking solutions that are most suitable for home networking include Ethernet, phoneline, and powerline networking.

Ethernet

Ethernet comes in many variants starting from shared 10Mbps to the popular switched 10/100Mbps, and more recently 1Gbps. While the Medium Access Control (MAC) mechanism differs in these variants, the frame format has remained mostly unchanged. The vast majority of today's homes are not wired for Ethernet, and retrofitting them may represent an inconvenient alternative for home users. However, there exists a trend in some newly built homes to be wired with Ethernet cables.

Phoneline Networking

The Home Phoneline Networking Alliance (HomePNA) technology implements a 10BASE-T-class network using existing telephone wiring as a shared access medium.

* Other brands and names are the property of their respective owners

HomePNA devices use an RJ-11 jack to plug into telephone "wall jacks" using a standard phone wire, just like a telephone.

Powerline Networking

Powerline networking technologies use powerline cables as the physical transmission medium. Due to the availability of powerline cables in every home, these technologies are gaining momentum in the industry.

Wireless LANs

Wireless home networking solutions are mostly based on IEEE802.11 specifications, which include 802.11b and the higher data rate 802.11a.

802.11b Wireless LANs

IEEE 802.11b operates in the 2.4 GHz frequency band and offers theoretical throughputs of up to 11Mbps, with effective throughput of up to 6Mbps. 802.11b networks usually employ a wireless LAN access point that network devices communicate with, at a distance of 100-300 feet in home environments.

802.11a Wireless LANs

IEEE 802.11a operates in the 5 GHz frequency band and offers higher theoretical throughputs of up to 54Mbps, with effective throughput of up to 22Mbps. 802.11a networks also employ a wireless LAN access point for devices to connect to.

Existing QoS Technologies

In order to deploy a complete end-to-end QoS networking solution in a Digital Home, various system and network components have to work in concert in order to achieve the final result. Figure 4 shows the different elements of a QoS networking solution. We describe these QoS elements in detail below, which include link-layer QoS support, network-layer QoS, and policy-based network management.

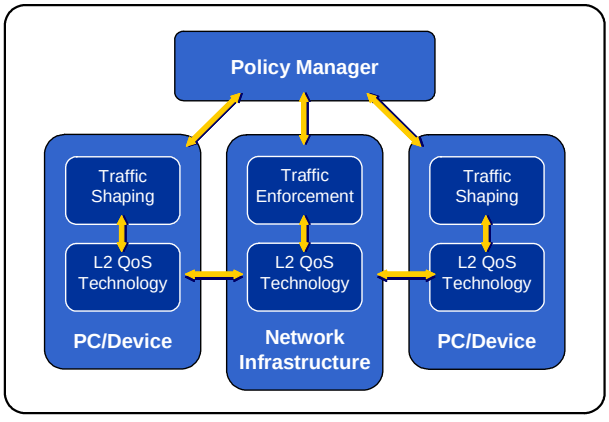

Figure 4: System overview for QoS networking

Before discussing these QoS technologies, we start first with a definition of Quality of Service.

Definition of Quality of Service

For the purpose of this work, we define Quality of Service as the ability to provide priority-based preferential treatment for certain traffic flows over the others, based on flexible usage policies.

Link-Layer QoS Technology

QoS support at the link layer (a.k.a. Layer 2, or L2) may take various forms depending on the specific link-layer technology being used. For example, wired Ethernet includes support for Quality of Service (QoS) in the form of 802.1p packet tagging based on the IEEE 802.1D specification, which defines the addition of four bytes to the legacy Ethernet frame format. The defined priority tagging mechanism is known as IEEE 802.1p priority tagging, and it allows for eight levels of priority.

Similarly, the HomePNA MAC layer uses Distributed Fair Priority Queuing (DFPQ), which enables an 8-level priority mechanism that is compatible with the IEEE 802.1p packet priority scheme used in switched Ethernet. Likewise, the HomePlug technology uses a MAC layer modeled after IEEE 802.11 (CSMA/CA) with QoS extensions that provide five priority levels with strict priority enforcement, based on a variation of IEEE 802.1p.

For Wireless 802.11 LANs, the IEEE 802.11e Task Group (TGe) is enhancing the current 802.11 MAC to add QoS support to both 802.11a and 802.11b. 802.11e QoS support will include a priority tagging mechanism based on the IEEE 802.1p definition as well.

Network-Layer QoS

At the network layer, flows are identified by IP header information such as source and destination addresses and protocol numbers, and in some cases transport layer information such as TCP/UDP port numbers. In order to provide QoS, the IP packets are first classified into flows, which are then policed and/or shaped according to the assigned flow rates, and finally queued and scheduled according to the QoS needed. This process of packet classification, policing, shaping, queuing, and scheduling is collectively called Traffic Control. Traffic Control may also mark or tag the packets with information to be used by other devices en-route to provide QoS.

Traffic control may be further classified into two categories: 1) traffic shaping; and 2) traffic enforcement. This classification is based on whether it is implemented

on end-systems such as source and sink PCs and devices, or on intermediate network nodes.

Traffic Shaping

On end system devices, traffic control is typically referred to as Traffic Shaping. Traffic Shaping modules usually interact with the Layer2 QoS module to mark or tag the packets accordingly.

Traffic Enforcement

In contrast to traffic shaping, traffic control on intermediate network nodes, such as wireless access points or residential gateways, is referred to in this work as Traffic Enforcement. Traffic Enforcement is similar to Traffic Shaping in terms of the functions being implemented; however, the actions taken on non-compliant traffic streams may vary from preferential scheduling to strict packet re-tagging with a lower priority, or eventually, packet dropping.

Policy-Based Management

Policy-based network management is used to dynamically manage and control the network behavior based on rules and actions [5]. In policy-based network management, policy clients get relevant policies from a policy server. A typical policy may state "Between 5pm and 10pm, allow 500kbits/sec bandwidth to a video flow from a PC media server to the TV."

The policy manager is typically responsible for creating, modifying, and deleting usage policies. It has the ability to evaluate a policy request against the policies it maintains and the dynamics of the network resources and provide a decision to the policy client upon request. Typically, policy servers provide a policy console that can be used to create new policies and/or modify existing policies. Policy-based network management can be effectively used to provide centralized and dynamic control over the home network.

Lack of a Complete Widely Adopted QoS Solution

As discussed above, various elements of QoS home networking exist today. In reality however, there is a lack of a complete widely adopted solution. A missing component for a complete QoS solution is the ability to discover QoS capabilities supported by devices on the home network.

In the following section, we describe a framework for QoS in the home based on UPnP technology. This framework uses UPnP capabilities such as automatic discovery, configuration, and control to provide a complete QoS solution in the home.

ADVANCED QUALITY OF SERVICE FRAMEWORK

Fundamentally, QoS networking provides capabilities to enable intelligent network resource management, based on their availability, as well as the demand for these resources represented by the number and networking requirements of multimedia streams.

We describe below a novel framework for QoS networking in the Digital Home environment based on UPnP* technology, taking into consideration the need for PCs and devices to discover and control QoS capabilities remotely.

Various applications may need to detect QoS capabilities on the network, such as whether a media server and/or an intermediate network switch supports packet tagging, or whether a particular wireless link has available capacity to support a media streaming session. Moreover, certain applications, such as when Dad decided to view a live corporate video (see Figure 1 in the Introduction section), may need to configure the residential gateway to reserve a portion of its bandwidth for the incoming video stream.

Figure 5 shows a block diagram of the QoS framework. This framework is currently under development and its main goal is to integrate the various QoS technology elements described earlier into a cohesive framework that can provide the necessary policy-based dynamic bandwidth management features needed to enhance the consumer's entertainment experience. In addition to the existing QoS elements, this framework also provides a means to discover, configure, and control QoS capabilities remotely over the home LAN.

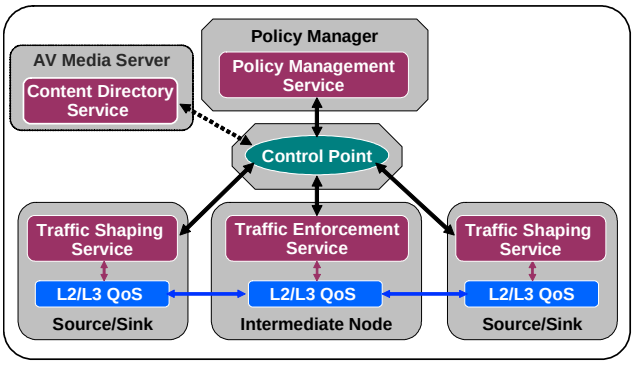

Figure 5: QoS framework based on UPnP technology

* Other brands and names are the property of their respective owners

In this framework, we introduce a number of UPnP services that work in concert to provide the necessary QoS networking functionality, namely a Policy Management Service, a Traffic Shaping service and a Traffic Enforcement service.

Traffic Shaping Service

The main role of the UPnP Traffic Shaping service is to enumerate the traffic control capabilities on end-systems (PCs and CE devices), expose them to the rest of the UPnP network, and allow control points to configure and control its QoS capabilities remotely. In other words, it provides a UPnP-based interface to access its traffic control functions such as packet classification, tagging, and scheduling.

Traffic Enforcement Service

Similar to the UPnP Traffic Shaping service, the UPnP Traffic Enforcement service provides a UPnP-based interface to access the underlying traffic enforcement (policing) capabilities. A Traffic Enforcement service is designed to allow applications to request network resources. It also allows a policy management application to enforce a particular policy, by pushing rules down to the traffic enforcement device.

Policy Management Service

The UPnP Policy Management service is responsible for exposing the capabilities of the Policy Server to the rest of the UPnP network, and as such, it listens to UPnP policy requests and relays them to the actual Policy Manger Application. When a change occurs at the Policy Manager, such as when a new policy comes into effect, the Policy Management service may generate an event that triggers other devices and control points to retrieve the new policy information.

Content Directory Service Extensions

The UPnP AV Content Directory service enumerates content available through the associated media server device. In addition to the traditional information stored in accordance with each content item, the QoS framework defines further QoS-related metadata extensions to be added for each item, such as the bit rate, packet size, and so forth. These extensions allow a control point to identify the QoS requirements for each stream.

Role of the Control Point

The UPnP Control Point plays a pivotal role in the overall QoS framework. In addition to the regular functionality it plays on the UPnP network such as discovery of devices and services, the control point in the QoS framework acts as a relay between the services defined. For example, the control point may obtain QoS requirements for a specific content item from the content directory, use the obtained information to send a resource request to the policy management service, obtain a policy decision and a priority setting from the policy manager, then send a command to the traffic shaping service on the media server to instruct it to tag and shape the packets according to the decision originating from the policy manager.

Summary of the QoS Framework

To summarize the capabilities of the QoS framework described above, it provides the following features:

- **End-to-end QoS support in the home:** The QoS model in this framework includes discovery, configuration, and control of QoS capabilities on end systems such as source (server) and sink (renderer) devices, as well as intermediate network equipment such as wireless access points and Ethernet switches.
- **Priority-based QoS as the baseline:** There are in general two broad categories of QoS mechanisms: priority-based and reservation-based mechanisms. This framework sets priority-based QoS with dynamic priority assignment as the baseline due to the availability of link-layer priority-based mechanisms.
- **Independence of link-layer technologies:** The framework is designed to provide a common interface for application developers, regardless of the underlying link-layer technology.

REALIZING HIGH-QUALITY MEDIA DISTRIBUTION IN THE DIGITAL HOME

In the previous sections, we presented an overview of existing infrastructure components for media distribution in the home. We have also described a QoS framework that integrates these infrastructure components and uses UPnP* technology to provide necessary features such as automatic discovery, configuration, and control for QoS networking capabilities.

We now look at a usage scenario based on Figure 1, to illustrate the QoS framework in action, and describe the associated theory of operation.

Media Distribution Scenario

Mom and Dad subscribe to a movie-on-demand service that entitles them to download a movie per week for a flat

* Other brands and names are the property of their respective owners

monthly fee. A typical 2-hour movie encoded in MPEG2 at 3Mbps takes about 2.7GB of disk space on one of their PCs, PC1 in this example. Mom usually selects the movie to download over the broadband Internet connection ahead of the weekend, and then invites their friends for a movie night on the weekend.

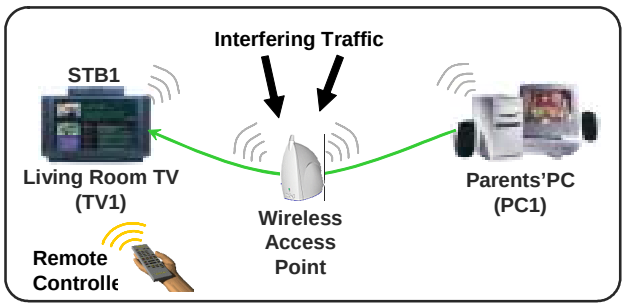

Figure 6: QoS-based movie playback

When friends arrive, Dad uses the remote control to set up the movie to be displayed on TV1 in the living room (through STB1). He also configures it such that the movie transfer from PC1 to TV1 is uninterrupted. In the meantime, their son decides to use the bedroom's TV (TV2) to watch his favorite basketball game that he had pre-recorded on PC2, and their daughter uses PC2 to download a large document from the Internet for a term paper she is preparing. The network traffic generated by their kids' applications is considered interfering traffic for their movie playback session. Figure 6 shows the scenario described above.

Theory of Operation

We now describe the theory of operation for the QoS framework based on the components described in the previous sections. We use the usage scenario described in Figure 6 to illustrate the sequence of events and the role of each component in the framework. We assume that the policy server is installed on PC1 and is already configured to provide the highest priority to the movie playback session during playback, and treat all other flows (basketball game video and file download) as Best Effort.

In the following, we describe the steps involved in the discovery of devices and services and the session setup, QoS configuration, and control. We then discuss the events generated during the session, as well as the effect of the interfering traffic and potential policy changes.

Devices and Services Discovery

In the usage scenario, PC1 advertises itself on the UPnP network as a device that contains a media server device, and two additional UPnP services: 1) a Policy Management service; and 2) a Traffic Shaping service. The media server device also includes a UPnP AV Content Directory service that provides a list of media contents stored on PC1, as well as a UPnP AV Connection Manager service that enumerates the media formats and streaming protocols it supports. Meanwhile, STB1 advertises itself as a media-rendering device that contains a UPnP AV Connection Manager service enumerating media formats and streaming protocols it supports, in addition to a UPnP AV Rendering Control service that enumerates rendering controls, such as volume, brightness, etc. The remote control acts as a control point for all services advertised. Once the control point discovers the services, it now knows it can create a QoS-based movie playback session from PC1 to STB1/TV1.

Session Setup, Configuration, and Control

When Dad picks up the remote control and points to STB1, all movies stored on PC1 are listed on the electronic program guide displayed on TV1. Dad then selects the movie that Mom had downloaded and presses the "play on TV" button. This action triggers the control point to request the description documents for the UPnP services on PC1 and STB1, which provide detailed information on the attributes and actions they support.

Capability Matching

The control point then evaluates the returned information and attempts to match the capabilities on both sides, such as the media format (MPEG2) and the streaming protocol (RTP over UDP). It also notes that both PC1 and STB1 support priority-based QoS and records the QoS parameters associated with the movie (bitrate, packet size, and delay requirements).

Policy Admission

When the control point learns the QoS requirements of the video stream that the media player is about to start, it sends a request to the policy manager in order to validate and approve the request. The policy manager then sends a response based on the policies that apply to the request.

QoS-based TC Setup and Configuration

Once the policy manager approves the QoS request, the control point sends a traffic shaping request to PC1, which in turn configures its Traffic Control module to tag all packets pertaining to the movie with a certain priority value. Finally, the control point issues UPnP requests to the media server on PC1 and the media renderer on STB1 to start playback.

Eventing

During the playback session, the control point receives events generated by each service. For example, the media server generates periodic events exposing the attributes and state of the server, such as "streaming started,"

"streaming complete," etc. Similarly, the media renderer on STB1 generates events about its playback state and any related statistics such as frame rate, packets lost, etc. The Traffic Shaping service on PC1 generates events that include QoS statistics such as number of packets tagged, etc. Finally, the Policy Management service generates events for the number of ongoing sessions, the priority assignments, and any changes in the stored policies.

Effect of Interfering Traffic

The usage scenario described above includes two background sessions, basketball game and file download, each generating interfering traffic that may threaten to degrade the quality of the movie playback. No QoS requests are associated with the interfering traffic, and as such, the wireless access point treats it as Best Effort. Since the movie playback traffic is configured for higher priority, it is tagged as higher priority so that the wireless access point ensures that it is uninterrupted by interfering traffic. In general, the policy manager is configured to allocate a certain percentage of the total bandwidth to Best-Effort traffic. This allocation regulates how QoS requests are admitted or rejected, thus ensuring adequate bandwidth sharing between Best Effort and QoS flows.

Effect of Policy Changes

Policy changes can happen due to either administrative or dynamic changes. Administrative changes consist of changes by an administrator to the actual policies at the policy server, such as which applications get higher priority bandwidth, amount of bandwidth to allow for an application, etc. Dynamic changes consist of changes resulting from external events. These could be based on time-based conditions or changes caused by other policy activity (such as a higher priority traffic needing the bandwidth from lower priority traffic). When a policy change does happen, all the impacted policy clients (the control point) must be notified. If the policy change results in traffic control change, the control point receives the Policy Management event and may issue a request to the Traffic Enforcement service on the wireless access point to handle the policy change.

CONCLUSION

In this paper, we described a high-quality media distribution architecture based on UPnP AV and a QoS framework based on UPnP technology that is currently under development. This architecture uses the automatic configuration, setup and control capabilities of UPnP* technology to provide devices and applications with adequate QoS support, without requiring these applications to be fully QoS-aware. We believe this approach enables easier deployment of high-quality media distribution solutions, which may in turn allow service providers to capitalize on the emergence of broadband access and home networking technologies in the home.

* Other brands and names are the property of their respective owners

REFERENCES

[1] UPnP Forum, http://www.upnp.org

[2] UPnP Device Architecture, Microsoft Corporation, http://www.upnp.org/download/UPnPDA10_20000613.htm

[3] UPnP AV 1.0 Specifications, UPnP Forum, http://www.upnp.org/standardizeddcps/mediaserver.asp

[4] UPnP Technology, Intel Research and Development, http://www.intel.com/labs/connectivity/upnp/

[5] A Framework for Policy-based Admission Control - IETF RFC2753.

AUTHORS' BIOGRAPHIES

Yasser Rasheed is a Network Architect in the Corporate Technology Group, Intel Corporation. He has been with Intel for two years, and has been involved in a number of projects related to home networking and advanced media streaming over wired and wireless networks. His interests include UPnP technology, Quality of Service, and policy-based networking for real-time applications. Yasser holds a B.Sc. degree in Electrical Engineering from Cairo University, Egypt, and M.Sc. and Ph.D. degrees in Electrical and Computer Engineering from the University of Toronto, Canada. Yasser's e-mail is yasser.rasheed@intel.com

John Ritchie joined Intel Corporation in 1986 and is a Staff Engineer with the Solutions Architecture Lab within the Desktop Platform Architecture organization. John is currently the co-chair of the UPnP AV working committee and was a principal architect and author of the UPnP AV specifications [3]. John has been involved with other interoperability technologies since 1994 including CE-Bus, Home PnP, Home API, and HAVi. John earned his B.S. degree in Information and Computer Science from the University of California, Irvine in 1983. His e-mail address is john.g.ritchie@intel.com

Remote I/O: Freeing the Experience from the Platform with UPnP* Architecture

Mark R. Walker, Desktop Platform Group, Intel Corporation
Jim Edwards, Desktop Platform Group, Intel Corporation
Michael Jeronimo, Desktop Platform Group, Intel Corporation.
John G. Ritchie, Desktop Platform Group, Intel Corporation
Ylian Saint-Hilaire, Desktop Platform Group, Intel Corporation

Index words: UPnP technology, remote, interface, discovery, protocol

ABSTRACT

Increasing adoption rates of home networking solutions raises the possibility of using the Personal Computer (PC) to remotely extend its power and capabilities into all areas of the home. A new specification to be developed under the auspices of the Universal Plug and Play (UPnP) Forum will allow home applications to be controlled and experienced with remote, network-connected devices. When adopted, this standard will enable the development of whole new classes of software applications that power new user experiences within the digital home. In this paper, we present two usage scenarios that illustrate new capabilities that would be enabled with this proposed standard. We discuss the foundational architecture of UPnP technology for home networking solutions and highlight its efficacy as a basis for constructing new, remote Input/Output (I/O) software services standards. Finally, we describe the technical challenges of the proposed standard and contrast this standard with existing and comparatively incomplete schemes for remote desktop and similar functionality.

INTRODUCTION

Recent reductions in the installation cost and advances in the average bandwidth of wired and wireless networking solutions have raised the possibility of employing the home network to remotely extend the considerable processing power of a fixed-location home PC into all regions of the home. As a result, a new standardization effort recently approved by the Universal Plug and Play (UPnP*) Forum [1] will focus on developing new software services for remote I/O across the home network.

Remote I/O is a technology under development based on the UPnP device architecture. It moves the point of user interaction away from an application running on a specific device, such as a PC or a CE (Consumer Electronic) device, to one or more remote I/O devices. The remote I/O device supplies input and output services such as mouse, keyboard, and display that together comprise the user interface. Applications run on a host elsewhere on the home network and are matched with compatible I/O devices. Applications may take on different user interface characteristics depending on the I/O devices being used. Furthermore, the application user interface can migrate across I/O devices as the user moves about the home.

Remote I/O supports interactions between PC and Œ devices by introducing location independence and tailored I/O devices to the home network. Applications can connect to wireless I/O devices anywhere in the home, freeing the user interaction from the location of the application. For example, a user might prefer to read the news in a comfortable chair in the family room instead of using the desktop PC in the den. In addition, I/O devices can be tailored to user activity. For example, a display for reading the news should be handheld, comfortable to hold, and allow the user to adjust visual settings such as colors and contrast. These two properties of remote I/O devices will improve the quality of user experiences in the digital home.

When in place, this proposed standard will effectively establish a UPnP remote I/O framework, allowing a fixed-location PC to serve user interfaces to a multitude of wired and wireless remote devices in the home. Under this new framework, remote User Interface (UI) devices will possess the "out-of-the-box" ability to discover and execute applications that can be run remotely. Correspondingly, these applications running on the PC would possess the out-of-the-box ability to discover remote UI devices and

project user interfaces onto them. By using this new middleware standard to drive remote UI devices, the PC-application user experience will be "freed" from the home PC platform.

User Scenarios

The following scenarios illustrate specific technical concepts of the proposed UPnP remote I/O architecture. It is important to note that there are many other compelling scenarios, especially around entertainment and communications, which could also benefit from remote I/O.

Scenario 1: Remote Control of Home Automation Functions

This is an illustration of the remote I/O standard-endowed ability of remote devices to discover and execute remote-enabled home applications.

Dave and Kathy's house has a home automation system. The system consists of some application software running on the home PC and a wire termination panel installed in the pantry. The termination panel is in turn connected to a large collection of interior light switches and dimmers, the thermostat, switches for outdoor lighting, the security system, and the lawn watering system. The PC implements the UPnP remote I/O standard on the home network. The termination panel is also connected to the home network, but communicates only with the PC and is not a UPnP device.

Kathy picks up the newly purchased universal wireless home remote from the kitchen table and sees an icon corresponding to the home-automation system in the activity display that she hadn't noticed previously. Touching the corresponding softkey, Kathy is surprised to see icons corresponding to all of the home-automation subsystems in the display. She selects the interior lighting system and watches the display change into a group of icons representing switches and dimmers for the entire house. She touches the on/off softkey corresponding to the light switch for the den and hears muffled, annoyed shouting coming from the den at the other end of the house. Dave is apparently still in the den using the PC.

Scenario 2: Integration of Remote User Interface Devices with a Home Security System

This is an illustration of the remote I/O standard-endowed ability of remote-enabled applications to discover and project user interfaces onto remote devices.

Continuing the previous scenario, Dave and Kathy's home-automation system is also connected to their security system. Dave configured the security system to use his wireless PDA and the connected TVs in the house as alarm enunciator and response devices, in addition to the alarm bell connected directly to the wire termination panel. The PDA and the TVs implement the UPnP remote I/O standard.

Late one evening, the family is suddenly awakened by the alarm. Dave and Kathy jump out of bed and immediately direct their attention to the TV. The TV displays text informing them that there is a smoke alarm active (in the garage), the best evacuation route from the master bedroom (down the front staircase and out the front door), and where to meet (near the oak tree in the front yard). As they dash from the room, Dave grabs his wireless PDA on the nightstand. The PDA screen displays some of the same text that appeared on the TV, but also has a menu displaying a choice of emergency responses. Dave uses the stylus to touch "Call Fire Department" as he reaches the front door. Their teenage son Rick meets Dave and Kathy on the lawn. The net-connected TV in Rick's room had instructed him to take the back staircase to the patio door in order to avoid the most likely location of the fire. In less than three minutes, a firefighter team arrives on the scene and determines that the smoke is emanating from the breaker panel in the garage. One firefighter uses a handheld extinguisher on the source of the smoke. The entire incident is over in less than twenty minutes. The damage to the house is minimal, and all family members are safe and sound.

General Requirements of the Remote I/O Framework

The above two scenarios would be possible only with some type of a preexisting framework for remote I/O. General requirements of the hypothetical framework underlying the above usage scenario descriptions include the following:

1. *Out-of-the-box device and service discovery capability*. The UPnP remote I/O framework must allow applications to discover remote-capable devices and initiate remote UI sessions. Correspondingly, remote devices must also have a way to discover and initiate UI sessions with remote-enabled applications.
2. *Scaling to match device capabilities*. Successfully remoting UIs to very thin devices would be made possible by endowing the application server with the ability, under the UPnP remote I/O framework, to adapt the UI and data transmission scheme according to the capabilities of the targeted remote device.
3. *Simplicity*. For remoting to very thin UI devices, it is important to keep the UPnP remote I/O framework as simple as possible. Ideally, the data transmission

protocol should require very little state on the client other than a frame buffer. The application server should additionally be responsible for maintaining its state independent of the remote device.

HIGH-LEVEL REMOTE I/O ARCHITECTURE

UPnP* 1.0 Architecture Overview

The UPnP Device Architecture V1.0 specification [2] provides a solid foundation for building UPnP remote I/O capabilities into local area networks. UPnP technology effectively establishes a middleware standard for out-of-the-box, cross-vendor discovery, description, control, eventing, and presentation.

UPnP Device Architecture V1.0 defines protocols for communication between *control points* and *devices*. Control points are essentially software applications and are the active components of UPnP architecture. Devices are physical or logical entities, enumerated via simple XML descriptions and containing APIs referred to as *services*. Physical devices may host multiple logical devices, and each device may host multiple services.

Messages are transported over UPnP networks via HTTP over UDP/IP or TCP/IP. The supported message formats are Simple Service Discovery Protocol (SSDP), General Event Notification Architecture (GENA), and Simple Object Access Protocol (SOAP). IP addresses are assigned by a DHCP sever, or auto-assigned when no server is present. There are four basic types of UPnP network activity:

1. *Discovery*. When a device is added to the network, the device is allowed to "advertise" its presence on the network using SSDP. When a control point is added, the control point is similarly allowed to generate a multicast search for devices. In either case, the message exchange consists of a brief description of the device that includes the UPnP device type, the device ID, and a URL to the full device description.
2. *Description*. A control point obtains more information about a specific device by retrieving the full description from the URL with HTTP GET. The full description is composed of a device description and a service description. The device and service descriptions are XML documents and are constructed by the device vendor with the aid of the device and service template schemas. The service description contains details of the hosted API commands, called *actions,* along with parameters, called *arguments*.
3. *Control*. A control point accomplishes device control by invoking actions on the service control URL and by polling for service state variables. UPnP architecture employs the SOAP remote procedure call scheme to deliver control messages and return results. Services keep state tables updated so that control points can obtain meaningful values. When state variables change, events are broadcast over the home IP network to all interested control points.
4. *Eventing*. Control points may receive notification of specific state variable changes by forwarding a subscription message containing a delivery URL to the selected service. The service maintains a list of URLs corresponding to subscribing control points. Events are expressed in XML and forwarded to subscribers via GENA-extended HTTP messages.

Overview of Proposed Architecture of Remote I/O

Like previous UPnP working-committee-developed standards, the UPnP remote I/O standard specification will build upon the basic architecture behaviors listed above and will specify new devices, services, control actions, and event types necessary to realize the usage scenarios.

Features of the proposed UPnP remote I/O framework needed to expose remote-capable *applications* include the following:

An application registry service that lists applications that implement remote I/O support would be required. The registry would be enumerated by an UPnP remote I/O application server device and visible to all other UPnP devices on the home network. At least one UPnP application server device would be associated with each physical application server on the network.

Each application listed in the registry would expose a list of input and output attributes in standard form. Typical output attributes of a single application would include video output, static display output in unformatted text or XHTML form, and display update events. Typical input attributes would include named UI events.

Features of the proposed UPnP remote I/O framework needed to describe capabilities of remote *devices* include the following:

Device input and output services enumerated by UPnP remote I/O devices and visible to all other devices on the home network would be required. At

least one UPnP remote device would be associated with each physical remote UI device on the network.

Each UPnP remote I/O device would expose input and output services along with device and service-specific properties. Typical input services would include video frame buffers, as well as text and bitmap display. Device-specific properties of frame buffers enumerated by UPnP remote I/O would include frame size, pixel resolution, and supported data transmission protocols. Typical output services would include events bound to physical device actions like mouse & key clicks.

Other than providing support for named events, the lower-level details of the session-specific data transmission protocol selected for use between the application and the remote UI device are out-of-scope with respect to the UPnP remote I/O framework. While UPnP remote I/O provides a framework for enumerating vendor-preferred protocols that may be supported by a given remote UI device, it will not require specific data transmission schemes for standard compliance, nor will it attempt to define new ones.

Analysis of Usage Scenarios

We can now analyze the usage scenarios against the proposed remote I/O framework described in the previous section.

Scenario 1 is an example of a device initiating the discovery of and establishing a remote I/O session with a remote-capable application. The wireless remote in this case hosts a logical remote I/O device with input and output services, as well as control point functions. Figure 1 contains a graphic illustration of the following sequence:

1. *Discovery*. When Kathy powers up the wireless remote, the local control point logic on the remote queries the network for all application-server devices with an SSDP multicast (Figure 1). The application-server PC hosting the home automation application responds to the remote with a URL for retrieving the root of the XML application-server device description.

2. *Description*. Using HTTP GET, the remote device acquires the detailed XML description of the application server device. Another URL pointing to the XML description of the specific remote I/O services hosted by the application server is contained in the device description XML.

   The wireless remote device automatically employs the URL pointing to the service description with HTTP GET to acquire the detailed XML listing of the services hosted by the application server. The control logic on the remote parses the service XML and finds both remote I/O input and output services. The input and output service XML document lists API-like "actions" that can be called on the

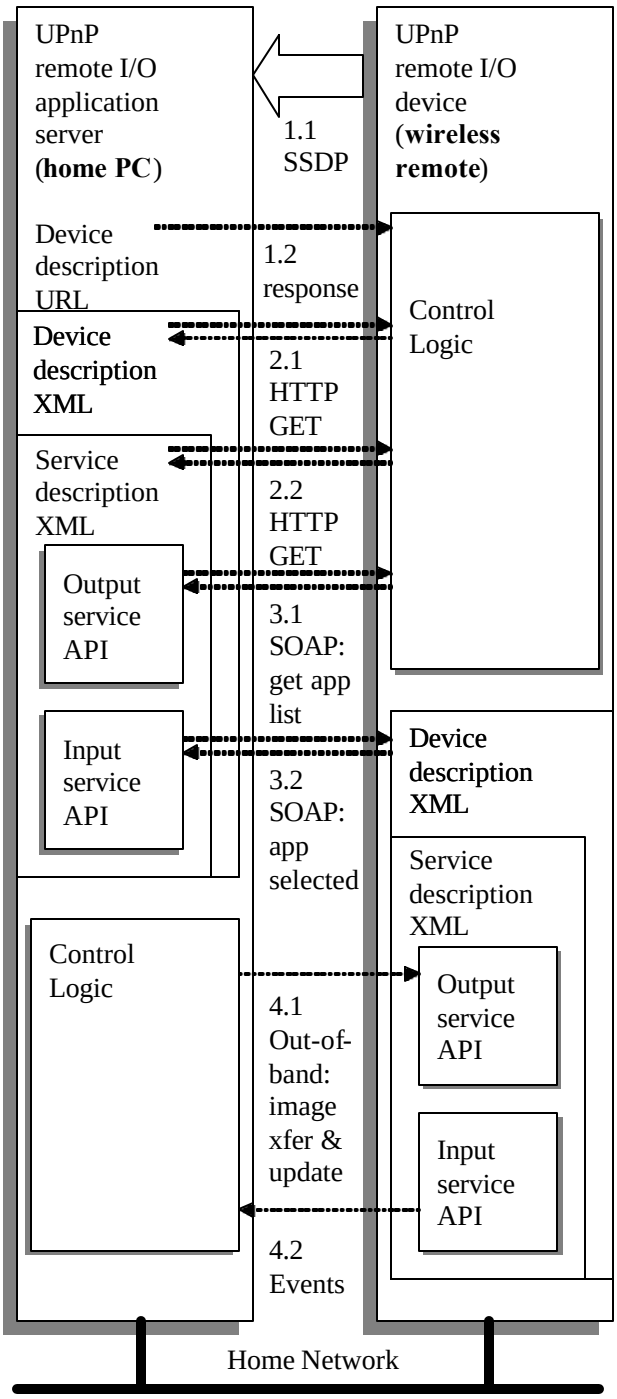

Figure 1: UPnP remote I/O communications sequence between application server and wireless remote UI device for Scenario 1

application server by the remote device, including an action that returns a listing of all remote-enabled applications accessible through this application server.

3. *Control*. The wireless remote automatically issues a SOAP packet to the application server containing the API action that returns the remote-enabled application list. The XML list is processed for UI display. Kathy then selects the lighting control application causing the wireless remote to issue another SOAP packet containing another requested action on the application server. The parameter list of this particular action contains not only the selected lighting control application, but also the XML UPnP remote I/O device and UPnP remote I/O service description of the wireless remote itself.

   The application server receives the SOAP packet containing the requested application along with the remote device description. At this point, control transfers from the remote device to the application server, as the application server establishes a UPnP remote I/O session with the wireless remote. The control logic on the application server matches a group of named events required by the application UPnP remote I/O input service with the softkey description contained in the XML UPnP remote I/O description of the wireless device.

4. *Out-of-band data transfer protocol and eventing*. An appropriately scaled image (possibly a bitmap) is forwarded from the application to the remote frame buffer. As Kathy touches the individual softkeys, corresponding command events are relayed to the subscribing input service on the application server and interpreted by the lighting control application.

Scenario 2 is an example of a remote-capable application initiating the discovery of compatible remote UI devices. The application server in this example hosts a logical application server device and registry service along with some control point functions.

1. *Discovery*. At installation time, using a discovery procedure identical to the one described in Scenario 1, the UPnP remote I/O application server hosting the home security system discovered all UPnP remote I/O devices and corresponding services. The security application established a table of remote I/O devices and their physical locations within the home.

2. *Description*. Prior to the alarm, the UPnP remote I/O application server hosting the home security system acquired the detailed device and service descriptions from each of the UPnP remote I/O devices responding to the discovery multicast.

3. *Control*. The above information was used to pre-configure the control action requests generated by the home security application requesting to set up sessions with each of the remote devices, as well as to pre-scale displays appropriately for each of the UPnP remote I/O devices.

4. *Out-of-band data transfer protocol and eventing*. When the smoke alarm activates, all pre-configured remote I/O sessions are instantly activated. Bitmaps with explanation texts are forwarded to each display device using the data transmission protocol appropriate for each device. The wireless PDA that Dave grabs from the nightstand implements input as well as displays functions. When Dave uses his stylus to touch the screen location corresponding to "call fire department," a command event is generated by the remote I/O output service on the wireless PDA. Dave's intent is interpreted by the home security application and an alert signal is sent to Dave and Kathy's security monitoring service, which in turn alert the local emergency services.

EXISTING, RELATED TECHNOLOGIES AND STANDARDS

Existing remote I/O schemes fall into two broad categories:

High-level methods for PC "desktop" remote User Interface (UI)

Low-level data transport protocols for remote devices

Remote desktop methods are generally designed to allow other interface devices (especially portable tablets) to transparently share desktop application functionality with a specific host machine.

HTML Presentation Pages

Under UPnP Architecture V1.0, if a UPnP device has a URL for presentation, then the control point can retrieve an HTML-based UI for controlling and/or viewing device status from this URL, load the page into a browser, and depending on the capabilities of the page, allow a user to control the device and/or view the device status [2]. The page is delivered via HTTP over TCP over IP. To retrieve a presentation page, the control point issues an HTTP GET request to the presentation URL, and the device returns a presentation page.

Unlike the UPnP Device and Service Templates, and standard device and service types, the capabilities of the presentation page are completely specified by the vendor, not the UPnP Forum.

Microsoft's RDP*

Microsoft's Remote Desktop Protocol (RDP) [3] is designed to provide remote display and input capabilities over network connections for Windows*-based applications running on a server. RDP is a protocol that allows for up to 64,000 separate virtual channels carrying device communication and presentation data from the server, as well as encrypted client mouse and keyboard data. RDP provides an extensible base from which to build many more capabilities. Other features of Microsoft RDP include the following:

bandwidth-reduction measures

support for roaming disconnect

support for print redirection

support for multipoint transmission

Sun SLIM*

Sun's SLIM [4] protocol is an experimental remote interface protocol that takes advantage of the fact that the display tends to respond to human input, which is quite slow. Refreshing the display from a local frame buffer and transmitting only pixel updates enable bandwidth savings between server and client. The SLIM client/console is essentially just a frame buffer. The server maintains the full, persistent contents of the frame buffer.

EIA-775, 775.1

EIA-775 [5] is the DTV 1394 Interface Specification developed by CEA's R4.8 1394 Interface Committee. The low-level EIA-775 protocol provides a two-way bus type connection for high-end audio/video products. The bus architecture allows several devices to send and receive audio, video, and control information to all other devices on the bus. For example, when going from a TV to a DVD, a consumer needs only to insert the disk into the player and hit play. The player will send the right messages to the audio and video displays to set up for stereo, or multi-channel and standard, or wide screen—whatever the case may be. EIA 775.1, the WEB-Enhanced DTV 1394 Interface specification, additionally includes Web browser and other Internet protocols. These allow a source of MPEG service (such as a cable or terrestrial set-top box, digital VCR or DTV) to utilize the MPEG decoding and display capabilities.

HAVI DDI

The Home Audio/Video Interoperability (HAVI) Data-Driven Interaction (DDI) protocol [6] defines an API-based scheme for allowing soft entities designated as DDI Controllers to manipulate DDI Target objects. The DDI Controller employs a description of the UI (DDI Data) to be presented to the user, which is obtained from the DDI Target. The HAVI DDI specification describes APIs and objects that are used to establish control sessions between Controllers and Targets.

Comparisons to UPnP Remote I/O

The essential component seen in virtually all of the technologies listed above is the type of specific, over-the-wire, data transmission protocol for exchanging UI information between an application and a remote UI device. Although these types of low-level communication protocol descriptions are indeed an essential requirement for remoting an application UI, they are insufficient for enabling general interoperability between applications and remote UI devices in a multi-vendor environment. In general, the piece that is missing in each of the above cases is a widely adopted framework for the discovery and enumeration of interoperable applications and UI devices.

CONCLUSION

The goal of the proposed standard is the establishment of a UPnP remote I/O framework, allowing a fixed-location PC to serve user interfaces to a multitude of wired and wireless remote devices in the home. It is anticipated that the use of this new middleware standard to drive remote User Interface (UI) devices will effectively "free" the application user experience from the home PC platform.

REFERENCES

[1] Mark R. Walker, John G. Ritchie, Michael Jeronimo, "Remote I/O: A UPnP Forum Working Committee Proposal," October 2002.

[2] *Universal Plug and Play Device Architecture* V1.0, June 2000. http://www.upnp.org/download/UPnPDA10_20000613.htm.

[3] Remote Desktop Protocol (RDP) Features and Performance.

Other brands and names are the property of their respective owners.

http://www.microsoft.com/TechNet/prodtechnol/win2kts/evaluate/featfunc/rdpfperf.asp

[4] B.K. Schmidt, M.S. Lam, J.D. Northcutt, "The Interactive Performance of SLIM: A Stateless, Thin-Client Architecture," *Operating Systems Review*, vol. 34, 5, December 1999, pp 32-47.

[5] DTV 1394 Interface Specification.

[6] The HAVi Specification: "Specification of the Home Audio/Video Interoperability (HAVi) Architecture," V1.0, January 2000.

AUTHORS' BIOGRAPHIES

Mark R. Walker has worked on numerous projects including PC-based video conferencing, wideband audio compression, and synthetic speech since joining Intel Architecture Labs in 1992. Mark received his Ph.D. degree from Arizona State University in 1991 and is now a member of the Intel Corporation's Solutions Architecture lab in Hillsboro, Oregon. Solutions Architecture is a part of Desktop Platform Architecture, which is an architecture organization within the Desktop Platform Group responsible for defining and developing platforms of the future. His e-mail is mark.r.walker@intel.com.

Michael Jeronimo is a software architect with Intel Corporation's Solutions Architecture lab in Hillsboro, Oregon, working on the architecture and technologies for the Digital Home. His interests include software architecture, UML, design patterns, and Internet security. Michael holds a B.A. degree in Computer Science from the University of California at Santa Cruz. His e-mail address is Michael.Jeronimo@intel.com.

Jim Edwards is a lead architect and a technology development manager with Intel Corporation's Solutions Architecture lab in Hillsboro, Oregon, working on the architecture and technologies for the Digital Home. His interests include home networking, digital media, and sports cars. Jim holds a B.S. and M.S. degree in Electrical Engineering from Cornell University. His e-mail is jim.edwards@intel.com.

John G. Ritchie joined Intel Corporation in 1986 and is a Staff Engineer with the Intel Corporation's Solutions Architecture lab in Hillsboro, Oregon. John is currently the co-chair of the UPnP AV working committee and was a principal architect and author of the UPnP AV specifications. John has been involved with other interoperability technologies since 1994 including CE-Bus, Home PnP, Home API, and HAVi. John earned his B.S. degree in Information and Computer Science from the University of California, Irvine in 1983. His e-mail address is john.g.ritchie@intel.com.

Ylian Saint-Hilaire is a software architect/developer working on UPnP/AV solutions. During his career, Ylian has been responsible for the development of network infrastructure and technology in a variety of areas including Intel's IPsec key negotiation technology and network mobility technologies. Ylian is currently a member of the Intel Corporation's Solutions Architecture lab in Hillsboro, Oregon. His e-mail address is ylian.saint-hilaire@intel.com.

Home Network Security

Carl M. Ellison, Corporate Technology Group, Intel Corporation

Index words: firewall, UPnP, 802.11, wireless, VPN, security, home networking

ABSTRACT

Home computers that are connected to the Internet are under attack and need to be secured. That process is relatively well understood, even though we do not have perfect solutions today and probably never will. Meanwhile, however, the home computing environment is evolving into a home network of multiple devices, which will also need to be secured. We have little experience with these new home networks and much research needs to be done in this area. This paper gives a view of the requirements and some of the techniques available for securing home networks.

INTRODUCTION

First, there was a single Personal Computer (PC) in a few homes with no connection to the outside world. Now, we have computers in most homes and most have Internet connections to the outside world. The next step, already happening, is not one computer but rather a large network of devices in a home. Some of these are mobile devices, which will be brought into the home by guests, friends, hired employees, maintenance personnel employed by service providers, and other strangers.

As these changes happen, the security needs of the home user also change. In the days of the disconnected single PC, the primary security threat was from virus contamination on floppy disks. With continuous connectivity to the Internet, many new attack channels have been opened (e-mail attachments, executable code or scripts fetched from Web pages, active penetrations at lower networking levels, etc.), while floppies have all but disappeared, closing that older channel. To the extent that these existing threats are understood, there are products available to help home users defend themselves against them.

However, the future home will have not one computer connected to the Internet but rather a network of many devices within the home, and that network might be connected to the Internet. In such an environment, the potential for attacks is greatly increased. Since this is still in the future, there are no products to counter these attacks. This is therefore an area ripe for research and product development. This paper primarily addresses researchers and product developers considering this new environment.

We briefly address the present state of affairs regarding the security of home computers. Present security measures will continue to be valuable in the future and will continue to evolve. Security solutions are always evolving, as no solution remains adequate for long.

The bulk of this paper, however, discusses the new home environment, in which there are threats not only from outside but also from inside. Those threats are characterized, and security mechanisms that can be built into products to secure the home user against these threats are described.

In our conclusion we describe how security mechanisms built for the corporate environment have serious flaws when used in the home environment. We discuss Universal Plug and Play (UPnP), developed in response to the unique needs of the home environment.

SECURING THE EXISTING HOME NET

Any home computer connected to the Internet is in danger of being attacked. A broadband connection leads to probes preparatory to an attack every few minutes. A dial-up connection, behind the firewall of an Internet Service Provider (ISP), leads to attacks from machines that are behind the same firewall. In the author's experience with one ISP, probes came once or twice a week.

There exist many papers, both academic and practical, on how to use existing products to secure current home computers from attacks via the Internet. It is not the

purpose of this paper to reiterate that advice, but to summarize it:

1. Computer owners should have a firewall and allow no responses to any attempts to connect into the home from outside. A firewall must have external administration disabled, and any passwords with which it was shipped need to be changed to very secure, hard to guess, passwords. These passwords can be written down, because they are defending against network attackers rather than in-home attackers.
2. A computer should have a modern virus scanner, which is enabled to scan all inputs to the computer, as well as automatic updating of virus signature files, at least daily.
3. Computer owners should update operating systems and applications with the latest security patches and scan for new patches daily. These patches must be digitally signed, and therefore authenticated, as having come from the software vendor and not an attacker.
4. Security settings should be set to maximum on both browsers and e-mail agents.
    a. E-mail agents should not allow incoming mail in HTML to be displayed if it accesses anything on the Internet.
    b. Neither application should allow any executable code or scripts to be accepted from the Internet and run.
5. If one uses wireless networking at home, the wireless access point must be placed outside the home firewall, rather than inside. Unfortunately, all current bundled firewall/access point products place the access point inside the firewall. Therefore, if one wants network security and wireless networking, and chooses a bundled product, then one must install a personal firewall on every machine in the house and allow no incoming connections on any of them.
6. For each operating system, there are numerous settings that must be made properly to maximize security. The documents describing such settings run to dozens of pages and need to be produced for each different home operating system.

These well-known security measures are both inadequate and burdensome. They are inadequate because any attack code that manages to penetrate a computer on the home network has free run within that computer. Solving this problem requires new operating system architectures–extremely long-term work. They are burdensome because with these measures in place, a computer user cannot view many modern Web pages because they require JavaScript; cannot read incoming e-mail transmitted in Hypertext Markup Language (HTML) so that the formatting will be as the sender intended; and cannot offer any Web services to friends out on the Internet.

There is a great deal of work yet to do before we have a good solution for the case of the single home computer connected to the Internet. Meanwhile, we as an industry are actively enhancing the home network. Few people today have real networks at home. Rather they have a single computer with a network connection, either dial-up or broadband. In the future, we anticipate home networks with hundreds of nodes. This future home network brings with it additional security problems that are not addressed by the products available today to secure the home computer and not completely addressed by projected modifications to operating systems that are needed to isolate hostile code from valuable resources within the home computer. This paper deals with those additional issues.

ELEMENTS OF SECURITY

It is a popular misconception that "security" is synonymous with "encryption." In many cases, confidentiality via encryption is the least important element of a security solution. Network security involves a number of different elements:

1. data origin authentication
2. command authorization
3. message integrity protection
4. message replay prevention
5. data confidentiality
6. key distribution
7. trust versus trustworthiness

Data Origin Authentication

Authentication is often tied in modern systems to integrity protection. To authenticate a message, one needs to establish that it came from a particular source. This can be established by physical point-to-point wiring, but can also be established by the use of cryptography, in which the sender of the message has a secret value and uses that secret value plus the message to compute a check value. The receiver/verifier checks the message origin (and integrity) by verifying that the check value could only have been produced by an entity in possession of the secret value. If public-key methods,

which are known as **digital signatures,** are used, then only the sender needs a copy of that secret value in order to get maximum security. If symmetric cryptography, via what is called a Message Authentication Code (MAC), is used, then the receiver also needs a copy of the secret value. Because there are two or more copies of that value in the system when we use a MAC, there is more opportunity for it to be compromised and therefore it is less secure. However, we still use MACs because symmetric methods are typically much faster than public-key methods. A hybrid scheme is often used, in which public-key methods are used to establish symmetric keys that are used for a short period of time.

Command Authorization

Establishing who sent a message, by authentication, is essential, but it is not enough. For example, there might be an incoming message commanding a home alarm system to turn itself off or a message to a home PC asking for a copy of a sensitive file to be sent to the requester.

An incoming message might be characterized as "Hi. I'm X. Do Y for me." Authentication verifies that the sender was X. Command authorization establishes whether X is allowed to do Y. Until you have established both authentication and authorization, you cannot make a security decision (namely, whether or not to do Y in response to this message).

Message Integrity Protection

It is essential to establish the integrity of incoming messages. This process is usually tied to authentication.

If the attacker could get a copy of a message saying "Hi, I'm X, do Y" and turn it into a message saying "Hi, I'm X, do Z," then if that new message passed the authentication verification process, the attacker could achieve a result that the legitimate parties did not desire. Normal authentication methods (digital signatures or MACs) include the entire message in the authentication and verification computation, so that any change to the body of the message would invalidate the authentication.

Message Replay Prevention

The attacker might capture a copy of a legitimate message, "Hi, I'm X. Turn off the home alarm system." That attacker could then re-use that message without any modification to it at all, except that it was sent at a time of the attacker's choosing. This is called a "replay attack." To prevent it, one must design network protocols that have unique, verifiable information (often called "freshness data") included among the data authenticated and verified in each message. This freshness data is often a sequence number or a time value. However, for home network use, especially when there are VCRs blinking 12:00 because the homeowner chooses not to set the clock, it is preferable not to rely on clock values being correct.

Data Confidentiality

Confidentiality could be achieved by dedicated, private network wiring but cryptographically it is achieved by encrypting the contents of the message. As with authentication, there are both symmetric- and public-key methods for doing this. In public-key systems, the receiver has the secret (called a private key); therefore, only the receiver is capable of reading a message encrypted for its key. In symmetric-key methods, the sender also needs a copy of the secret (the symmetric key) and as a result it is less secure. As with authentication, a hybrid method is often used: public-key methods are used to establish symmetric keys that are used for a short period of time or for a single message.

Key Distribution

Both authentication and confidentiality require the two communicating parties to have certain cryptographic keys. If public-key methods are used, the key distribution problem is a little simpler, but it is not trivial. It must be designed very carefully. Flaws or shortcuts in key distribution can completely invalidate the security benefit of the mechanism used.

Unfortunately for home networking, key distribution is considered an onerous task, and shortcuts are often employed to save the homeowner from having to do "geeky" things. So, for example, wireless network devices often come with built-in default keys that homeowners are allowed to just use. Use of such keys makes the security mechanism worthless, but the 802.11 devices don't know they are using worthless keys, so they spend the same amount of processing time (reducing network bandwidth) as they would with valid keys. Similarly, firewalls often control access by password and come with a default password (e.g., "admin"). Users who leave that password unchanged have completely invalidated the security mechanism.

How keys are distributed varies from one security tool to another and is discussed in more detail in a later section.

Trust Versus Trustworthiness

People sometimes use the words "trusted" and "trustworthy" as if they were synonyms. In fact, they are practically antonyms.

If a thing is trustworthy, then if you trust it you are not exposing yourself to risk. However, a thing is often called "trusted" not because it is trustworthy but because you are forced to trust it. In that case, you are exposed to risk. As a rule of thumb, it is good to have trustworthy things and bad to be required to trust things.

Unfortunately, we have no sure means of establishing trustworthiness when it comes to security. Therefore, it is standard practice to assume an entity is untrustworthy until proved otherwise. This is counter to standard social practice and calls for care on the part of the product designer. A homeowner should not have to rely on trust when it comes to friends or family using devices within a home. Rather, a product needs to be designed where rights can easily be granted to friends, the minimum rights necessary to do the job. Total access should generally not be granted to anyone except the homeowner regardless of how trustworthy the person is.

HOME NETWORK SECURITY REQUIREMENTS

The requirements for security in a home network depend on how "home" is defined. It also depends on what is envisioned as the network within that home.

If the network is just a link from a cable modem to a single PC, then one length of network cable would accomplish all the network security that the homeowner needs. However, we think ahead to a time in the not-too-distant future when a home contains dozens, if not hundreds of networked devices, some belonging to the entire household and some belonging to individuals within the home.

We summarize the security definitions of the previous section in two categories: authorization and confidentiality. For each device in the home network, we need to concern ourselves with two questions:

1. *Authorization:* Which things are authorized to do what actions or access what data on each device?
2. *Confidentiality*: Which things are allowed to read the messages being transferred to a given device from somewhere else?

The "things" referred to here could be networked devices or could be applications on a networked computer being operated by a particular person. Universal Plug and Play (UPnP) calls these things "Control Points" (CP). These CPs might all be within the home, but they might also be remote from the home, connecting into the home from the Internet.

Let us look at the definition of "home" more carefully, since people often use radically different definitions for the term without examining those definitions.

Single-Person Homes

The most basic home environment is a dwelling with only one person living in it. All the devices within the home belong to that one person. It is easy to provide a secure home network in such a home, assuming it is not connected to the Internet. Any device within the home can do anything with any other device within the home. One can, for example, use only a wired network and have no other security. If such a home network uses wireless networking, one can make sure that link encryption is used to enforce the policy that only home network devices are allowed to connect to wireless access points within the home.

This most basic home is of little interest, but it is the model that many security designers assume.

When the home network is connected to the Internet, the domain under consideration is no longer the home. It has many people, some to be kept out at all costs and some to be allowed access, but only to carefully selected resources.

Couples With Small Children

The task of securing the network in the home of a couple with small children might be as easy as that of a single person, provided the two adults agree on the security policy.

Families With Teenagers

Life becomes more complex with teenagers. Most teenagers are trying to establish some degree of independence. This might include ownership of personal networked devices and probably would include inviting friends into the house. What if those friends want to plug their own networked components into the home network? The establishment then of a security policy becomes much more complex than it was in the single person's household.

How much autonomy does the teenage child need? How much autonomy must the child's guests be allowed? How much does the head of the household have to trust either the child or the child's friends?

Adult Guests and Roommates

Adult guests and roommates are presumably more trustworthy than the guests of teenage children, but by the *Principle of Least Privilege* (that no person should be granted more access than he or she needs to do his or her job), the same questions apply to adults as to teens.

SECURITY DOMAINS AND POLICY

For the purposes of this paper, let us define a security domain as a set of objects that are allowed to interact with each other. A person is yet another object, according to this definition, although a person is usually represented on the network as an application that has access to a particular private key and can be operated only by a particular person.

A security policy is the specification of how objects in a security domain are allowed to interact.

The objects in these domains are all networked and computerized, for our purposes, but they are not all network components. For example, a networked home alarm system might be in a security domain with one particular control application on the family PC that can only be accessed by one user (let's say the head of the household). Other persons or other applications on that PC would not be allowed access. Another example might be in a single-occupant home where the PC has a directory of financial files that can be accessed by the homeowner *and* by the homeowner's tax accountant (on an office computer, connected by the Internet into the home). In this case, a security domain that includes that specific directory, the accountant's application, and some of the homeowner's applications might need to be defined.

The actual specification of security domains is up to the owner of the resource(s) being protected. What product designers and researchers need to be aware of is that these domains will contain objects that are much finer grained than network nodes, and that a resource owner might define as many security domains as he or she today defines file folders. In other words, some people would define only one while others would define hundreds.

UpnP , described below, was designed to take this into account. The "object" could be as fine grained as one action performed by one process in one PC running logged in as one particular user, or it could be as large as an entire networked device (e.g., a printer or scanner).

Interacting with these objects are what UPnP calls "Control Points" (CP). The user can define an arbitrary number of security domains in this structure. The user can also define named groups of devices and CPs. In the simplest form, there would be one policy statement: "my Control Points can do everything with my devices." In the most complex form, each pair-wise association would be defined carefully and intentionally.

KEY DISTRIBUTION MECHANISMS

It is not possible to say that one element of a security solution is more important than another, with the implication that you can do just the important parts. Doing 80% of a security solution is like closing 80% of a submarine's hatches and diving. [3]

That said, key distribution is the first and arguably the most important part of a security solution. Included under the term "key distribution" are the following:

1. passwords
2. DES, AES or WEP keys
3. public keys
4. PKI

Passwords

Typed passwords are typically converted by algorithm to cryptographic keys. When they are not converted to keys, they are used for authentication, just as a key is. Therefore, we consider passwords to be in the category of keys. Passwords can be distributed by being set by the manufacturer and printed for the user to read, but they are more secure if the user chooses a password and uses that. A fundamental problem with passwords though is that for security reasons, they should never be written down, but in reality, they are often written down. Most people cannot keep passwords in their memory unless they are very simple. This makes passwords a weak form of security: if they can be memorized, they are probably too simple and so can be guessed; if they are too complex, they are written down and are therefore available to a passerby.

DES, AES or WEP keys

There are symmetric encryption algorithms, such as Data Encryption Standard (DES), Advanced Encryption Standard (AES), and protocols using symmetric algorithms, such as WEP (Wired Equivalent Privacy). The keys for these algorithms and protocols are like passwords, in that both ends of a communication need to know the key. These are typically expressed in HEX digits and have the advantage that they can carry more

entropy (information) than a typical password, but have the disadvantage that they are not memorable and must therefore be written down. That makes them potentially available to someone other than the user although not to attackers on the Internet.

Public Keys

Public-Key (PK) cryptography differs from symmetric-key cryptography in that one encrypts with a different key from the one used to decrypt. It is also a characteristic of PK systems that one key, called the private key, can easily be used to generate the other key, called the public key, but the reverse is not true. One cannot easily use the public key to generate the private key. This allows the public key to be published.

When one encrypts with the private key, one gets what is called a "signature." Anyone with access to the public key can verify that this encrypted quantity was encrypted by the private key that corresponds to the public key that verifies the message.

When one encrypts with a public key, one gets confidentiality. Only the holder of the private key can decrypt the message thus encrypted.

For key exchange, PK systems have an advantage in that the public key can be transmitted without any need for privacy. In particular, it can be transmitted over the network. A public key can also be stored wherever it is needed without any efforts to keep it secret, although it must be protected from being replaced with an attacker's public key.

Although a public key can be transmitted without special protection, the machine receiving it needs to decide whether to use that key for a particular purpose. It is that decision, entirely within the receiving machine, that constitutes the security of a public-key distribution mechanism. The keys may flow freely, but there is a security decision to be made in any acceptance of such a key.

Public keys are simpler for a user than are passwords, DES, or WEP keys with which the user needs to enter the actual password or key over a secure channel. Typically this is done by typing. Public keys can be sent over the network and the user need only say "yes" or "no" when the machine that received the key asks if it should accept that offered key. The user can make that determination by comparing keys, without having to type any values. Typically the user will compare some function of the keys, such as a SHA-1 (Secure Hash Algorithm revision 1) hash.

PKI

A traditional Public-Key Infrastructure (PKI) is a mapping from names to public keys, with that mapping created by some trusted third party (usually called a Certification Authority or CA). It sounds good at first, but turns out to have severe problems.

Humans use names. They prefer to deal with them, especially over nonsensical things like keys or hash values. The issue is where those names come from.

With a traditional PKI, the CA must come up with the name to bind to the device's public key and must do that without knowing anything about the person who will eventually look at that name and try to make sense of it. These names must also be unique among all keys being certified. So, for example, a CA might create a name like ***Acme MP3 server, model 5489023-M, serial number 20020115-598003***. The user of the name, on the other hand, needs only to distinguish this device from other devices the user owns or otherwise has to deal with, so for the user a name like ***MP3s*** might make more sense. If the user has two MP3 servers, the second one might be named, by a CA, ***Acme MP3 server, model 5489023-M, serial number 20020115-598083.*** The user, however, if selecting his or her own name for the device, might call it ***bedroom MP3***.

The UPnP* Security Key Management Choice

For UPnP Security, we looked at the methods of key distribution and decided to use public keys and also to name keys personally. In other words, a user would acquire a new device and learn from that device the SHA-1 hash of its public key. That public-key hash is reported to an application the user runs, called the Security Console, and the user gets to compare what was reported over the network to what was learned from the new device (e.g., printed on a card shipped with the device). After a satisfactory comparison, the user then names the key with some name meaningful to the user. From then on, the user refers to the device by that name.

AUTHORIZATION MECHANISMS

Once a key for a given device or component or user has been learned, that entity can be authenticated, but a security decision cannot be made based only on authentication. A device must know what each authenticated entity is allowed to do. Devices cannot be

---

Other brands and names are the property of their respective owners.

manufactured with that knowledge built in, so it is the job of the device owner to implant that information.

There are many mechanisms available for this, but the three predominant ones are an Access Control List (ACL), an Authorization Server, and an Authorization Certificate.

Access Control List (ACL)

An ACL is a protected table residing in memory in the same device as the resource whose access is being protected. It is an array of entries, and each entry contains the following:

1. *subject*: an identifier of the entity being granted access
2. *authorization*: an indicator of the rights being granted that subject
3. *delegation*: a flag, indicating whether the subject may further delegate these rights
4. *validity*: optional conditions on validity of the entry, such as a "not-after" date and time

Some ACL entries contain fewer than all four of these fields, but these are enough to cover any home network authorization decision we have encountered.

A device can control access by an ACL alone. This makes programming easier and also allows an access entry to be deleted with ease, assuming one can access the device holding the ACL. It has the disadvantage of requiring a great deal of ACL editing if there are a large number of ACLs or a large number of subjects. It also could require a large amount of ACL storage. Since ACLs must survive power failures, this memory must be non-volatile.

For example, a traditional time-sharing file system ACL would contain a username (or group name) as the subject and some set of file permissions as the authorization (e.g., {read, append)). It would typically not allow delegation or have expiration dates.

The application SSH (Secure Shell) uses a file .ssh/authorized_keys which is an ACL whose entries contain only subject entries. Each subject is a public key. The authorizations are all the same (the ability to log in on that account and to do SCP (Secure Copy) commands to it). There is no delegation or validity interval.

In Universal Plug and Play (UPnP) Security, an ACL can have all four fields. The subject is either the hash of a public key, a name of a group of keys, or the reserved element "<any/>." The authorization is an XML (Extensible Markup Language) element with sub-elements listing individual permissions being granted. Since the subjects are public keys, the subject is able to delegate rights via authorization certificates, so there is a delegation field, with the default being permission to delegate. Validity fields are available if desired.

Authorization Server

If one has an environment (e.g., in a corporation) that contains a large number of devices all of which need the same ACL and if that ACL is very large, because there would be a large number of subjects, and if network costs are low, then it might make sense to move the ACL from each local machine to a server, often called an *authorization server*. This solution does not apply to the home environment, but products are typically developed for both environments at once, by people trained in the corporate environment, so an authorization server might be considered for home use.

However, this does not eliminate the need for an ACL in each device. When one uses an authorization server, the device would generate a message of the form "May X do Y?" send that message to the server and get back an answer, 'yes' or 'no.' That message from the server back to the device needs to be secured in all the ways described above under the definition of security. Each reply from the server needs to be protected from modification, replay, or imposture. Therefore, it needs to be authenticated and authorized. Since this is the message from the authorization server, one cannot use an authorization server to authenticate and authorize this message. Therefore, the device needs an ACL listing the authorization server. That ACL, in effect, grants all access rights to the server and allows it to delegate rights to others.

Even though each device needs an ACL, there might still be advantages to using a server. The ACL in each device is very small (one entry) and should rarely have to change.

For home use, however, an authorization server probably makes little sense. It complicates the network and adds cost in an environment where there is likely to be very little duplication of devices and therefore little benefit from the consolidation of ACL entries in one server.

Authorization Certificate

Another way to administer authorization without requiring each device ACL to list each subject and its access rights is to allow delegation by way of authorization certificates [1]. An authorization certificate is a digitally signed ACL entry.

A subject listed in the device ACL might be given the right to delegate some set of permissions. That subject can delegate permissions on to a second subject, where what gets delegated to the second subject is the intersection of the rights granted the first subject and the rights delegated on to the second.

With delegation of rights, the burden of administering security is spread out. One could also spread this out by allowing multiple entities to edit the ACL itself, but in that case, one entity could remove rights added by another entity. The entity empowered to edit the ACL also gets complete access to the device. With delegation by authorization certificate, the entity to whom rights have been delegated does not get total rights to the device, cannot further delegate any more than the rights it has been given to delegate, and cannot remove rights of others.

UPnP Security supports authorization certificates although their implementation is at the discretion of the device manufacturer.

Group Definition Certificates

Another way to spread out the administration of authorization is to have ACL entries (or authorization certificates) that grant rights to named groups.

A name in this context is not just a text string. There is no source of globally unique text string names for arbitrary objects nor is there likely ever to be. DNS (Domain Name System) is a working global name space, but the political attacks mounted on the Internet Corporation for Assigned Names and Numbers (ICANN) shows that even DNS is under siege. However, we need globally unique names to avoid ambiguity that can be exploited by an attacker.

For the purposes of this paper (and for UPnP Security) we define a name to have two fields:

1. hash of a public key
2. text name (as defined by the holder of the private key corresponding to the public key)

The pair is globally unique because the hash of the public key is globally unique.

We allow named groups to be defined in the Simple Distributed Security Infrastructure (SDSI) style [4], by a name definition certificate containing the following:

1. *issuer key hash:* the hash of the public key of the entity defining the name
2. *name:* the text of the name being defined
3. *subject:* the specification of the group member, either the hash of a key or a name of a subgroup
4. *validity:* a possible limitation of the lifetime of this group membership, e.g. via not-after dates

This name definition is then digitally signed by the issuer key and stands for the statement that "the subject is an element (or subgroup) of the specified named group."

With named groups, one can share the administrative load of granting access, but with more limitation than with authorization certificates. Because the name definition certificate contains no authorization field, every entity in the group gets the same access grant as every other entity in the group, that being the access granted that named group in an ACL entry or authorization certificate.

UPnP Security allows named groups, but their use is at the discretion of the device manufacturer.

SECURITY PRODUCTS

The products available in 2002 tend to support the hardened perimeter model of security. This is appropriate to the most basic concept of home (with only one user and no interactions with the outside world) but not to the more complex forms of home environment.

These products also tend to have been designed based on requirements of industry rather than of the home, making their administration difficult and sometimes assuming the existence of both physical security and a group of on-call support professionals.

Universal Plug and Play (UPnP) Security, described below, is a new standard designed for home use, but is too new to have any products for sale as of the fall of 2002.

Firewalls/Gateways

An Internet gateway or firewall secures an internal network from the Internet, to the extent that it blocks

---

unsolicited traffic from the outside. As long as there is a single security domain inside the home (a single-person home or a couple with small children, for example), the home can be secured by a single firewall. However, if there is more than one security domain inside the home (e.g., roommates or guests) then a single firewall would not help guard the interests of one internal security domain from other internal nodes. One might create a separate wired network for each security domain and give each of those networks its own firewall. However, that solution gets expensive as the number of domains increases.

Even in homes in which there are multiple security domains whose security is defined through mechanisms other than firewalls, one will probably want a firewall to protect the collection of domains from hostile outside entities.

Wireless Security

Wireless networking is becoming popular at home. It relieves the homeowner of the work of running network wires through and within finished walls. It can also reduce the clutter of wires within a room.

However, with this benefit comes a security drawback. By relieving the homeowner of the work of individually running network wires to each device in the home, wireless networking prevents the homeowner from selecting which devices should connect to a given network as might be accomplished by running wires. Instead, with wireless networks, cryptographic keys need to be used to individually choose which devices should be connected to a network. Devices allowed onto a network would be given the key to use that network.

The choice of wide area coverage networking, as with wireless or power-line networking, might also restrict the number of networks the homeowner could define. With individual wires, the homeowner can set up separate networks for only the cost of some hubs and wires. With 802.11, each separate network would require a separate Access Point and separate channels. Since there are fewer than seven 802.11 channels that can operate in the same area without getting in each other's way, this limits the number of networks that can be declared in a small space like a home and implemented by 802.11.

WEP

Wire Equivalent Privacy (WEP) was the original security measure for 802.11. It has been shown to have a flaw in key usage that allows an attacker to recover the key used after eavesdropping on a few thousand messages. Therefore, for real security, WEP is not useful. It can be an annoyance for a casual attacker, but not for a determined attacker.

802.11i and 802.1x

There is an on-going standardization effort in the IEEE under the titles of 802.11i and 802.1x to define security mechanisms to replace WEP. Although these definitions will presumably be cryptographically correct, they retain the problem of being wire-like (therefore unable to secure things more fine-grained than whole devices) and being limited by channel assignment. For a home network with a single occupant and therefore a single security domain, this might be a good solution. For a more complex home network, security must be achieved in other ways.

VPN

There are various Virtual Private Network (VPN) products that permit one to associate devices together in virtual networks, each created cryptographically. Most modern VPNs use the IPSEC (Internet Protocol Security) protocol.

With this technology, one can individually connect pairs of machines and build arbitrary security domains, provided the elements of those domains all have IP addresses (are full devices).

One potential problem with IPSEC or other VPN solutions is that if you have a network node (e.g., the homeowner's PC) that is in multiple security domains, a device in one domain might be able to link to a device in another domain by routing traffic through that PC. Preventing this linkage requires proper network administration (e.g., routing tables) within the PC.

UNIVERSAL PLUG AND PLAY

Universal Plug and Play (UPnP) [2] is an industry initiative designed to make home networking easy. It does not include security in the basic protocol. One can secure UPnP networks by wiring, if there is a single home domain and no wireless or power-line networking. However, in more general cases, one will need UPnP Security.

UPnP Security defines a service to be added to each secured device that allows its security to be managed. It also defines a service and control point behavior for an application called a Security Console, which edits the Access Control List (ACL) of a secured UPnP device and controls other security functions of that Device.

Overview

The basic architecture of UPnP V.1 is client-server, with the client called a "Control Point" (CP) and the server called a "Device." There are three protocols used to let components interact with each other:

1. SOAP (Simple Object Access Protocol): for remote procedure calls from a CP to a Device.
2. SSDP (Simple Service Discovery Protocol): for discovery of Devices.
3. GENA (General Event Notification Architecture): for subscribing to event reports and publication of those events.

SOAP carries the bulk of the work and the security of SOAP is described in detail below. In brief, SOAP is secured by allowing only authorized Control Points to invoke any secured action within a Device. This is accomplished by an ACL in each secured Device, each of the entries of which lists a Control Point (CP) unique ID, a name of a group of CPs or the universal group, <any/>, and what that CP or group is allowed to do on that Device.

SSDP is difficult to secure. It is vital for the authorized user to discover other Devices on the home network, but it is desirable that an attacker not be able to take an inventory of the home network's equipment. Therefore, to secure SSDP, the existence of Devices is announced, but they are announced as generic Devices, and are not described in any detail except in response to requests from authorized Control Points. It is still possible for a determined attacker to take an inventory of a home network, even if all traffic was completely encrypted (e.g., via IPSEC), just based on timing and length of messages. Therefore, securing SSDP is considered low priority until there are significant new research results in anonymity protection.

For security of GENA, we define a normal event variable that anyone can read, named "EncryptedEvent" but it contains any event variables that are to be made available only to authorized Control Points. An EncryptedEvent is sent to a Control Point encrypted in a key known only to the Device and that Control Point. Control Points not allowed to see such an event, are not allowed to subscribe to those events, and are not able to see the events by eavesdropping on the network.

Security Console

UPnP Security has defined a combination Device and Control Point called the Security Console (SC). This can be a separate component or part of some other component, but for the sake of definition, it is treated as separate. Its purpose is to take security ownership of Devices and then to authorize Control Points (or other Security Consoles) to have access to Devices over which the SC has control.

An SC can own a Device, meaning that it has the right to edit that Device's ACL and can do anything on that Device, or it can have been given some subset of rights to the Device and have the privilege to grant all or some of those rights to another SC or CP (by way of an authorization certificate).

The SC also has the job of defining names of individual Control Points and of groups of Control Points.

Discovery and Component Naming

The first security requirement is to discover your own network components. UPnP uses a broadcast protocol, SSDP, for physical discovery of Devices but nothing for discovery of Control Points. Therefore, the SC, acting as a Device, offers an action by which security-capable CPs can announce themselves to the SC. In that way, it learns of local Control Points while it uses SSDP to learn of local Devices.

Selection of one's own components can be done in a variety of ways, and there is room for much creativity here. UPnP Security offers a generic method, using a Security ID for each component. The Security ID is the SHA-1 hash of a component's public key, expressed in BASE-32 (using only upper case letters and six of the ten digits). It looks like a product registration ID:

```
GM3GK-RTMOI-4GYK2-ZK5FC-WMTRK
```

This ID is unique among all components in the world. By comparing the Security ID of a physical component to the ID announced to the Security Console, one can determine precisely which Devices are his or hers, even if the local network contains hundreds of Devices of the same kind (e.g., in a college dorm).

One might also select one's own components physically. For example, if the SC was a handheld unit with a physical link to the component or a very short-range radio link, then selection could be physical.

Once a Device is selected, whether physically or via comparison of Security IDs or otherwise, the user needs a better way to refer to the Device. The manufacturer could supply a name for the Device, but we give the user the chance to set his or her own name for the component. For example, where the manufacturer might specify a name like:

```
Media Store Model 5328-I-71
```

a user who thinks of this media store only as a place for MP3 files might choose a name like:

```
tunes
```

That name needs to be meaningful only to that user, so the choice of name is entirely up to the user. The name is used by the Security Console user interface, while over the network the hash of the public key (from which the Security ID is generated) is used as the component's unique ID. The key hash is also the ID by which a Control Point is known in a certificate or Device ACL.

Security Ownership

A Control Point does not need to be exclusive about which Security Consoles it advertises itself to. It is the beneficiary of grants of authority and all decision making is done by the Security Console in that case.

The situation is reversed for Devices. A Device has the resources (SOAP Actions) to which access must be restricted. The Security Console, by editing the Device's ACL, tells the Device which Control Points to obey. Therefore, the Device needs to be very exclusive in choosing which Security Console to associate with. This process is called "security ownership," in UPnP Security.

By the generic ownership protocol defined by UPnP Security, an SC can take ownership of a Device only if the SC knows the Device's secret password and the Device is not already owned. Once a device is owned, an SC that owns it can grant co-ownership to another SC or revoke it, but more importantly, an SC that owns a Device can completely re-write the Device's ACL (or do any other ACL editing operation).

Authorization and Permissions

What an SC does by editing a Device ACL is grant authorization to a Control Point or some other SC to exercise certain permissions. These permissions are arbitrary names, chosen by the manufacturer, to correspond to SOAP actions within that device. For example, one manufacturer might choose to have permissions with the same names as the actions that Device offers, with each permission allowing a caller to invoke just that one action. More likely, there will be far fewer permissions than actions and their names will be intuitive to the average user.

In addition, UPnP Security allows permissions to be parameterized, if the manufacturer desires that. For example, the demo of a UPnP Secure Media Server, as presented at the February 2002 Intel Developer Forum, had only one permission "Play" but it was parameterized with one or more file names, so that one could grant permission to play one or more individual MP3 files.

Delegation and Named Groups

With just ACLs, one can control authorization. However, as the complexity of the home network increases (e.g., when it scales to a college dorm or when it includes older teens and the mobile computers of their visiting friends), maintaining an ACL on each device might prove onerous. If a task becomes onerous, it tends not to be done, hence security is weakened.

The task of maintaining very large ACLs can be made more manageable by the use of certificates. UPnP Security uses both name and authorization certificates, as described previously in this paper.

With both kinds of certificates, the ACL ends up smaller. A possibly large part of the detail of what would have been a large ACL is expressed in certificates instead. When those certificates are issued by someone other than the operator of the SC, this reduces the workload of that operator, spreading it out among others. It also allows a measure of autonomy for teenage children or guests.

Delegation for Constrained ACLs

Some Devices might have very constrained local persistent memories. Since an ACL must survive power cycles, it must be held in persistent memory (such as flash) and a large ACL might overflow the Device's memory.

One might use delegation via name or authorization certificates, not to reduce the manual workload of administering authorization but rather to offload the Device's memory. In the most extreme case, one can have an ACL with only one entry, allowing a particular SC to have and delegate all rights on the device. That SC would then issue authorization certificates to express the same thing that might be expressed in a larger ACL. The user interface for this operation would probably look the same as ACL editing, so that the operator remains unaware of the difference.

OTHER PRODUCT LINES

Universal Plug and Play (UPnP) might not be used for all home security. A UPnP interface consists of SOAP messages and some products might not use SOAP as the preferred protocol.

However, one can use UPnP key distribution and authorization mechanisms (or the equivalent) in all of these product lines. It is key distribution and authorization that provide most of the security and those represent all of the user-visible work.

If a new product development can improve on the UPnP key distribution and administration, then that would be beneficial. Meanwhile, the UPnP forum is working very hard to improve on those areas.

New product developers must carefully consider the following:

> the whole range of environments (definitions of "home")
>
> the range of security policies a user might want to establish (e.g., matrices of what component may do what with what other component)
>
> delegation of rights, via named groups or authorization certificates
>
> the interaction between home devices and those in the global Internet

The body of this paper should give all the details needed to at least make a checklist for that process. Much more detailed discussion of this process can be found at the author's personal distributed authorization Web page [5].

CONCLUSION

One conclusion of this paper is that with proper security against the insider threats in the home environment, the security of the home network against threats from outside is increased.

Securing the home network is not the easy job some people would like to believe. A home network security policy can be much more complex than a corporate security policy. The homeowner would have to implement via network security policy controls what the corporation implements via door guards.

Most network security thinking to date has assumed that network access is binary: that one would allow access to the network or not. The idea of controlling access to individual components (or parts of those components, such as individual SOAP actions) is relatively new to network security design. While we adjust our product design process, this will produce a period of gradually increasing security and there will be a gradually lessening tension between the desire for ubiquitous computing and connectivity on the one hand and the desire for real security on the other.

ACKNOWLEDGMENTS

I thank Jesse Walker for keeping me up to date on 802.11i and 802.1x and my fellow members of the UPnP Security team, both within Intel and outside, for the collaborative effort that produced the UPnP Security design and implementation. I also thank Vic Lortz for being the driving force in the UPnP Security effort, chairing the Working Committee and bringing me into that work as security architect.

REFERENCES

[1] Ellison, et al., "SPKI Certificate Theory," RFC2693.

[2] Universal Plug and Play Forum, www.upnp.org

[3] Jesse Walker, private communication.

[4] Rivest and Lampson, "Simple Distributed Security Infrastructure." http://theory.lcs.mit.edu/~cis/sdsi.html

[5] SPKI and related resources. http://world.std.com/~cme/html/spki.html

AUTHOR'S BIOGRAPHY

Carl M. Ellison is a Senior Security Architect in the Network Architecture Lab of the Corporate Technology Group of Intel Corporation. His current research is devoted to distributed, public-key authorization that can be delegated, as in UPnP Security. He also has a special interest in self-organizing networks. His concentration on security has been a side-effect of a more general career focus on distributed and fault tolerant systems. His e-mail address is carl.m.ellison@intel.com.

Content Protection in the Digital Home

Michael Ripley, Corporate Technology Group, Intel Corporation
C. Brendan S. Traw, Corporate Technology Group, Intel Corporation
Steve Balogh, Corporate Technology Group, Intel Corporation
Michael Reed, Corporate Technology Group, Intel Corporation

Index words: content, protection, transmission, storage, DTCP, CPRM

ABSTRACT

This paper describes a flexible and extensible framework for distributing digital content in a protected and interoperable manner within the home. Within this framework a "chain" of solutions provide comprehensive, end-to-end protection of digital entertainment content as it is delivered to the home and managed, stored, and consumed on a wide range of devices. These solutions employ a consistent set of technical and licensing mechanisms to ensure that the content is protected in an effective and efficient manner throughout the domain. This leads to a framework in which content protection solutions from a variety of licensors can be selected and combined in a flexible and interoperable manner, based on market forces.

A number of content protection solutions that operate within this framework have been developed and are incorporated into products on the market today. This paper describes two such solutions:

- Content Protection for Recordable Media (CPRM)
- Digital Transmission Content Protection (DTCP)

INTRODUCTION

As more content enters the digital domain, the desire to protect that content grows. With consumers increasingly eager to move content between devices such as personal computers, DVD players and recorders, set-top boxes, and digital televisions, a variety of content protection technologies have been developed. These point solutions come together to form an overall "chain" of content protection technologies. Figure 1 depicts an illustrative example of such a chain, where for the sake of clarity analog connections are not shown. This illustration

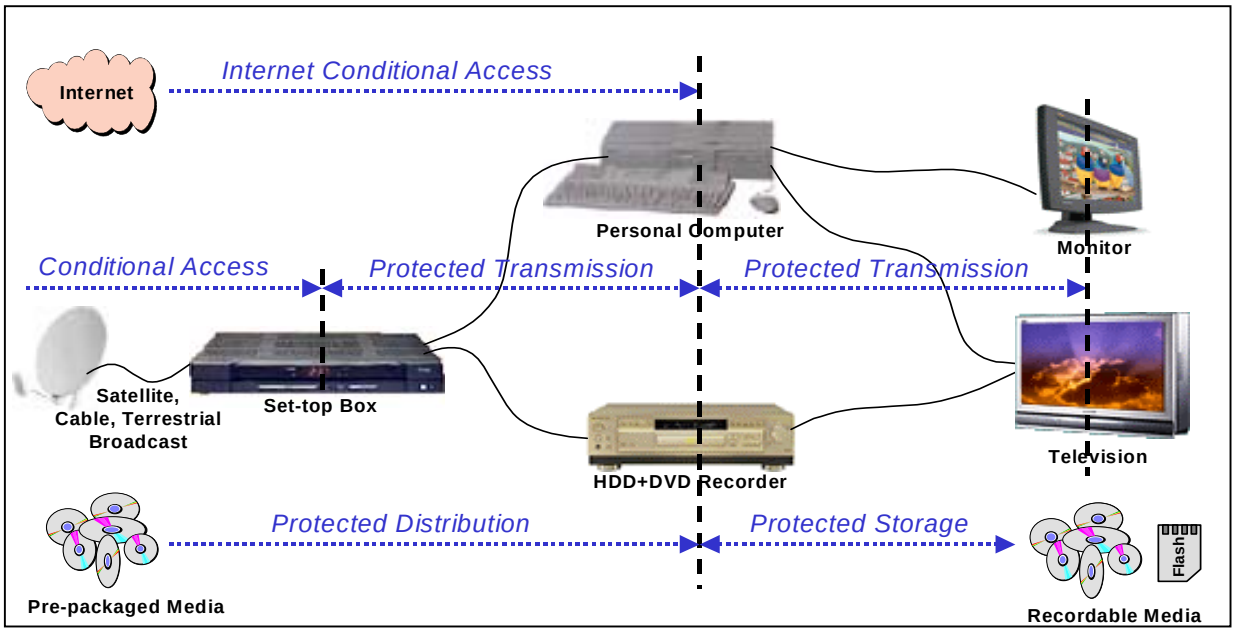

Figure 1: Digital content protection chain

should not be considered definitive.

The strength and completeness of the chain depends on more than just the individually developed links. The rising number of content protection technologies makes clear that an overall system architecture is needed to ensure that the individual pieces form a coherent, interoperable whole.

Without such a unifying architecture, inconsistencies, gaps, and even conflicts can occur between the various technologies, reducing the effectiveness of an overall content protection solution. The lack of a unifying architecture also leads to redundant and costly development efforts.

What is needed is an architecture that defines a set of overall principles that content owners and product developers can apply to ensure that content is protected in an efficient and effective way as it passes from one technology to another within the content protection system. Such an architecture can strengthen the overall content protection system, and ease implementation burdens on developers. By promoting the development of a comprehensive, compatible content protection system, this architecture stands to benefit content owners, content providers, device manufacturers, and above all consumers.

BASIC CONTENT PROTECTION STRUCTURE

Content protection solutions use a combination of technical and legal mechanisms to protect content against use that is inconsistent with the terms under which it was obtained from the content owner. The technical mechanisms take the form of a cryptographic protocol through which content is distributed or stored in an encrypted form. Access to the cryptographic keys and other intellectual property necessary to decrypt the protected content is subject to a license. This license is a legal tool to enforce the conditions under which such access is provided, including rules governing robust implementation and continued protection of content that is received subject to the license and subsequently stored or output.

The combination of technology and licensing provides a potent solution for preventing circumvention of content protection systems, while accommodating consumer expectations. Content protection technologies are effective at preventing unsophisticated attempts to circumvent a particular content protection solution. At the opposite extreme, well-financed professional pirates have routinely demonstrated an ability to defeat content protection technologies that have been incorporated within the economic constraints of consumer products. Conversely, licensing and other legal mechanisms are much more effective against business entities with assets, employees, and distribution channels.

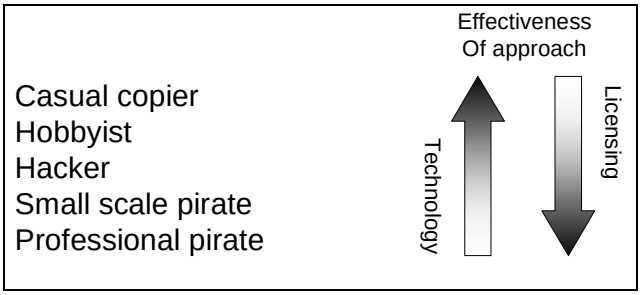

Figure 2: Effectiveness of technology and licensing

Thus, cryptographic technology implementations provide the basis for content protection, and an effective licensing structure provides for enforcement. The following subsections describe technical and licensing elements in further detail.

TECHNICAL ELEMENTS

Cryptography represents the technical foundation for content protection. This foundation can be broken down into four key elements:

- **Authentication.** To ensure that only licensed products have access to the protected content, technical means are needed to verify that a product is authorized. This is accomplished by the licensor of a content protection technology providing secret values that are only available to licensed products, which can be explicitly or implicitly verified as part of the process to gain access to the content. An example of implicit verification is two products independently calculating and using a secret value that can only be calculated by compliant licensed products. Note that a licensed product can be not only a stand-alone physical device such as a television or set-top box, but also a software application running on a personal computer.
- **Encryption.** Encryption is used to prevent unauthorized access to protected content. Decryption should only be possible by products whose compliance with the licensing conditions is verified via authentication.
- **Usage States.** State is typically associated with content that governs how the content may be used. This information is stored and communicated in a manner that ensures its integrity is maintained. The information can be carried along with the content, with its integrity cryptographically protected.

Additionally, it can be embedded within the content using a watermarking technology or via online mechanisms that obtain a "license." ("License" in this context means a file that is downloaded from a trusted source that conveys how the content may be used.) Products that license a content protection solution can be compelled through that license to respond to such watermark Usage State information as a means of extending the application of Usage States into the unprotected (e.g., analog) domain.

- **Renewability.** Renewability is used to extend the viability of a content protection solution. This is accomplished through a variety of mechanisms including updating a Digital Rights Management (DRM) system via online mechanisms and/or utilizing revocation. Either technique ensures a system's integrity is preserved in the event that a licensed product's authentication and/or encryption-related secrets are compromised and distributed. When such a compromise is detected and verified, the licensor of the content protection solution (through a process described below) may request an update or revocation of the compromised information. Revocable information should be assigned as finely as possible, ideally with each licensed product receiving unique information, thereby localizing the effect of the revocation as narrowly as possible.

LICENSE ELEMENTS

Each content protection solution within the aforementioned framework comprises an adopter's license, which is a technology license between the purveyor of the given content protection technology and manufacturers who want to include support for the technology within their products. Requiring adopters to enter a license as a condition of implementing the technology is important to maintaining the overall integrity and effectiveness of the content protection solution. Note that the adopter's license is complemented by additional licenses between the technology licensor and content providers and others, such as a "content participant agreement."

Apart from language covering confidentiality, intellectual property mutual non-assertions, payment of fees, disclaimers, liability, termination, remedies, etc., an adopter's license contains special provisions in the following areas:

- *License grant***.** A license to patents, trade secrets, and copyrights associated with the technology is granted to adopters only for the purpose of implementing the technology in a manner consistent with the specification and the other terms of the license, which includes the robustness and compliance rules. Designing and manufacturing a circumvention device under the license is strictly prohibited.

- *Specification changes*. The licensor may make certain changes to the specifications and license documents at the request of adopters, content providers, or on its own initiative. Before doing so, consideration of factors such as the integrity, security, and commercial viability of the content protection solution, and whether such changes would impose additional substantial obligations on adopters, is undertaken. Notice and comment opportunities are provided for content companies participating in the particular technology to allow any concerns about possible adverse effects on content protection to be known and, if the concern remains at the time a change is actually to be made, to allow a neutral arbitrator to determine whether a particular change would actually have a negative effect on content protection.

- *Third-party enforcement.* Content industry companies that have entered into a license agreement with the licensor have a strong vested interest in ensuring the integrity and viability of the content protection solution. These companies are typically given the ability to seek injunctive relief from adopters who are in violation of provisions of the license that directly relate to the effectiveness of the technology.

- *Revocation.* Renewability of the content protection technology is important to maintain its effectiveness. To ensure that devices are only revoked under the specified circumstances described above, a process is provided for consulting with the manufacturer whose device is proposed for revocation. If there is any disagreement about whether the conditions for revocation have been met, arbitration is used to establish the facts and resolve the dispute according to the established process.

- *Compliance rules.* Compliance rules are technical documents embedded within the license that specify when and how subsequent protection technologies are to be applied to protected content that is output or stored. Compliance rules may also include requirements such as the encoding and carriage of Usage States, response to watermarks at unprotected inputs to the device, and limitations on the number or quality of permitted copies of protected content.

- *Robustness rules.* Like the Compliance rules, the Robustness rules are a technical description of how products must be designed and manufactured to make reverse engineering or other modification difficult.

Typically there are separate rules for software and hardware implementation styles. For instance, software requires integrity verification mechanisms, such as signed code. Mechanisms to hide or obscure the operation of code and any secrets contained within may also be required. Hardware rules may require that implementations hide secret values and algorithms, by, for example, embedding them in silicon. Furthermore, products should be designed such that attempts to modify or otherwise tamper with them will likely result in the device being rendered unable to access protected content. Regardless of the implementation style, there are clear prohibitions against incorporating "defeating functions" that would enable a consumer to trivially disable any of the protection for the content. As well there are requirements to keep confidential information secret and to prevent access to the protected content when it traverses busses within the device. For all of these requirements, the levels of threat that must be resisted are described in terms of the types of tools and the skill level of the hacker.

CONTENT PROTECTION CHAIN

The license structure described above enables the formation of a "chain" of solutions for protecting content on an end-to-end basis as it is transferred, stored, managed, and consumed on entertainment devices within the home. Various technologies may be available for any particular part of the "chain" and may be selected according to device capability and application. Figure 3 shows an example, using representative content protection solutions, two of which are described in this paper.

Here, content is delivered to the home in an encrypted form via a conditional access technology. Products such as a set-top box that are designed and equipped to access this content must be licensed in order to receive the keying material necessary to decrypt the protected content. One of the terms of the conditional access license is that subsequent output of content that is subject to that license must be protected by an approved technology. In this example, the set-top box uses Digital Transmission Content Protection (DTCP) on IEEE 1394 interconnects.

In like manner, a recorder receiving content via the IEEE 1394 interconnect must be licensed if it is to decrypt DTCP protected content. One of the terms of the DTCP license is that exchangeable copies of copy-one-generation content that is subject to that license must be protected by an approved technology. In this example, the recorder uses Content Protection for Recordable Media (CPRM) to protect such copies on writeable DVD formats. Note that the DTCP license also includes provisions for making copies of content subject to that license that are uniquely associated with the licensed product (e.g., time-shifting on a local hard disk drive). Such localized copies must be protected to prevent further copying or playback on another device, using a method that meets the robustness rules of the DTCP license. One suitable method for such local protection may involve encrypting the content using an unpredictable key that is unique to the device.

For a product to play back the exchangeable DVD copy previously mentioned, the product must be licensed in order to decrypt the CPRM-protected content. In this example, that licensed product is a software application running on a personal computer equipped with a DVD drive. Thus, the chaining process is iterated until the content is ultimately consumed at the final device in the chain.

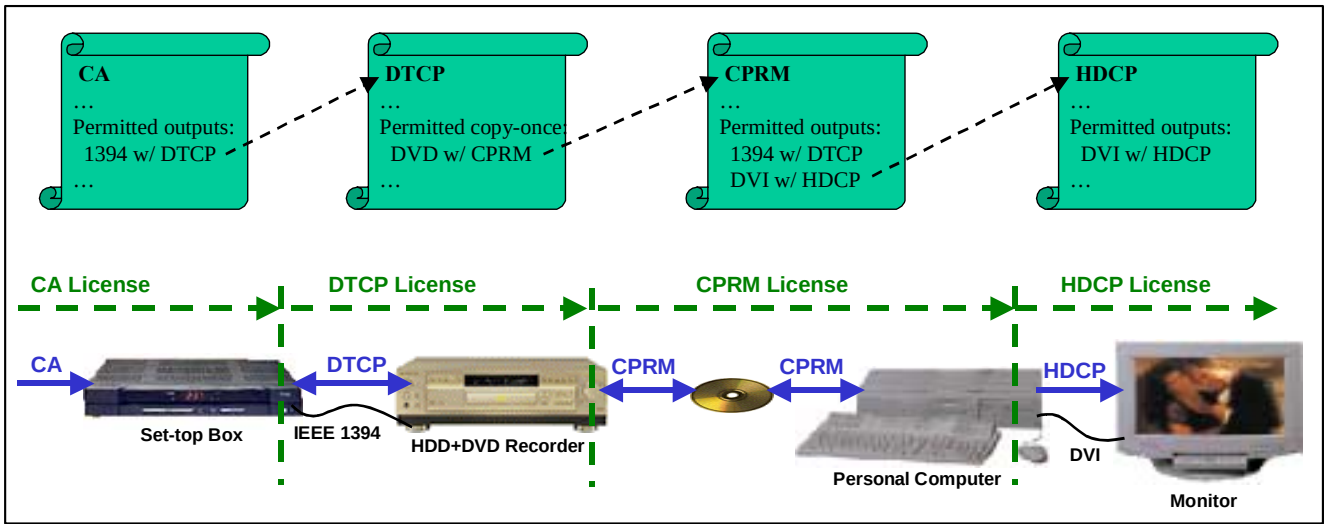

Figure 3: Content protection chain example

EXISTING SOLUTIONS

This section describes two existing solutions that operate within the framework described above:

- Content Protection for Recordable Media (CPRM)
- Digital Transmission Content Protection (DTCP)

CPRM

CPRM was developed by IBM, Intel, Matsushita and Toshiba (collectively, the "4C"). It provides a robust and renewable method for protected exchange of content via storage on portable/removable recording media. To date, CPRM has been defined and licensed for use in protecting content in a number of formats stored on a number of physical media types. These include DVD formats, SD Memory Cards, and Secure CompactFlash.

There are two primary technical components of CPRM: the C2 cipher and the Media Key Block.

C2 is a 10-round Feistal network block cipher with a 64-bit block size and a 56-bit key. The 4C companies designed and adopted C2, despite general cryptographic design principles which encourage use of well-known and well-evaluated ciphers, since no "well-known" alternatives had been identified that provided the necessary balance between suitability of hardware and software implementation, minimal licensing fees, and the ability to exclusively license C2 for use in 4C content protection solutions. This last attribute is particularly important, as circumvention of the 4C technologies will likely require use of the C2 cipher algorithm, which must be licensed from 4C. The C2 cipher is used to both encrypt and decrypt content and also as the basis of one-way and hash functions. Also, a license is now available for certain "stand-alone" uses of the C2 cipher (see the C2 License at http://www.4Centity.com for details).

Media Key Blocks (MKBs) are tables of cryptographic values that implement a form of broadcast key distribution, and they provide for renewability in 4C content protection solutions. MKBs are generated by the 4C Entity, LLC, and enable compliant licensed products to calculate a common "media key." Each licensed product is given a set of "device keys" when manufactured (also provided by the 4C Entity, LLC), which are used to process the MKB to calculate the media key. Device key sets may either be unique per device, or used commonly by multiple devices (4C licenses describe the details and requirements associated with these two alternatives). If a set of device keys is compromised in a way that threatens the integrity of the system, updated MKBs can be released that cause the compromised set of keys to calculate a different media key than is computed by the remaining compliant devices. In this way, the compromised device keys are "revoked" by new MKBs. In existing 4C solutions, MKBs are carried on compliant portable storage media, and devices use the corresponding medium's key as the basis for encrypting and decrypting protected content stored on that medium.

Figure 4 provides an illustrative example of CPRM being used to protect video content stored on writable DVD media. Note that this figure provides a simplified overview; for complete details see the CPRM technical

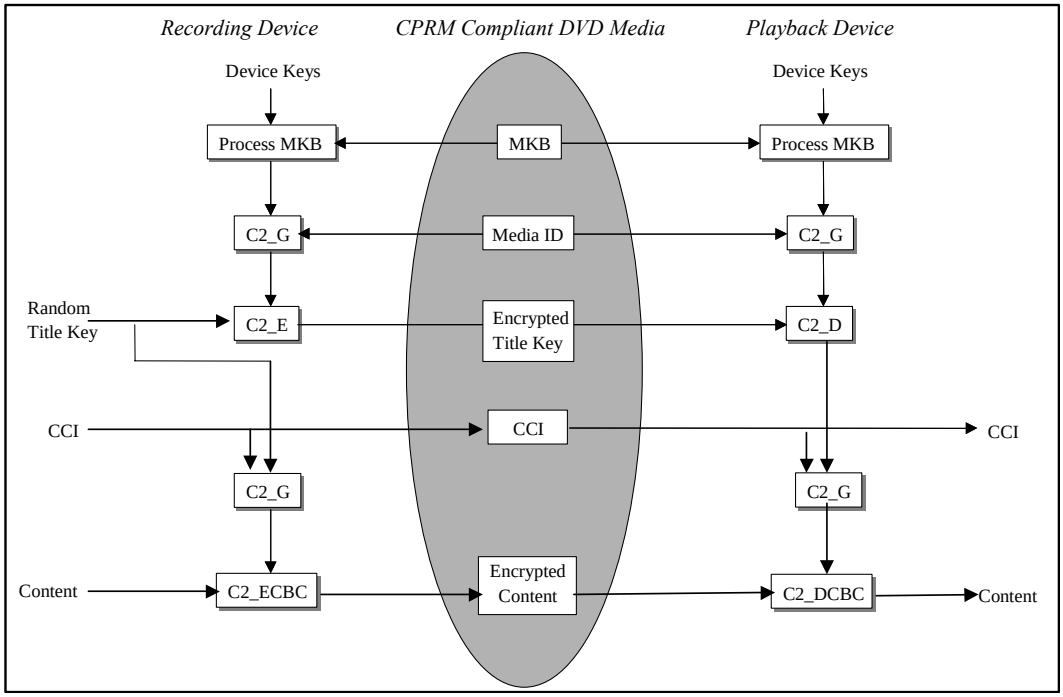

Figure 4: CPRM operation

specifications *www.4centity.com*

Each writable DVD disc that supports CPRM carries a Media Key Block (MKB) as read-only data stored in the lead-in area and a unique media identifier (Media ID) stored in the burst-cutting area, a region that is uniquely written on each disc in a manner that cannot be recorded or modified with consumer DVD equipment. A DVD recorder (e.g., stand-alone recorder, or software application running on a PC) that is equipped and licensed to use CPRM makes a permitted copy onto such a disc in the following manner. The recorder reads the MKB from the disc, and uses its secret Device Keys to process the MKB and calculate the Media Key. The Media Key is then combined with the Media ID, using the C2 One-way Function (C2_G), to form a Media Unique Key. The Media Unique Key is used to encrypt a randomly generated Title Key using the C2 cipher encryption function C2_E, and the encrypted value is stored on the data area of the disc. The Title Key is also combined with the content's Copy Control Information (CCI) using the C2 One-way Function (C2_G), and the result is used as a key to encrypt the content using the C2 cipher in cipher block chaining mode, a function denoted as C2_ECBC.

Thus, the content is encrypted in a manner that cryptographically "binds" it to that particular disc, through the use of the unique Media ID. The protected content can be played back from that disc by any compliant player that is equipped and licensed to use CPRM. Such a player uses its Device Keys and the relevant C2 decryption functions to carry out the corresponding playback process, as shown above.

The following subsections describe various physical media and content formats for which CPRM is currently defined and licensed. It is anticipated that CPRM will be applied to additional formats in the future.

DTCP

DTCP was developed by Hitachi, Intel, Matsushita, Sony, and Toshiba (collectively, the "5C"). DTCP provides system renewability as well as an encrypted exchange of content and CCI between authenticated devices. To date, DTCP has been defined and licensed for use in protecting the transmission of content via the IEEE 1394 serial bus (1394), the Universal Serial Bus (USB), and the Media Oriented Systems Transport (MOST), which is used in the automotive sector.

CCI is carried embedded in the content stream according to the content format (e.g., MPEG). For instance an MPEG-2 transport stream descriptor has been defined to carry this CCI. In addition, the CCI is mapped into an Encryption Mode Indicator (EMI) that provides protected, yet easily accessible, access to the CCI.

Content is encrypted using the M6 block cipher that is used in converted cipher block chaining mode with 56-bit keys.

Two authentication and key exchange (AKE) procedures based on challenge/response procedures are defined to enable manufacturers to trade off implementation complexity versus value of content to be handled:

Full Authentication (for all content) is based on Digital Signatures and Diffie-Hellman Key Exchange using a 160-bit elliptic curve public-key cryptosystem compatible with IEEE P1363.

Restricted Authentication (acceptable for "copy_once" and "copy_no_more" content only) is based on shared-secret techniques.

System renewability is provided through device certificate revocation. The license administrator can, under a rigorously specified set of conditions, exclude individual, compromised devices from participating in the protection system with devices supporting Full Authentication. Revocation lists are carried in System Renewability Messages that are distributed with content and between compliant devices.

Figure 5 shows an overview of the operation of the content protection system. The device that is the source of protected content has been instructed to transmit the content via the IEEE 1394 serial bus isochronous transport.

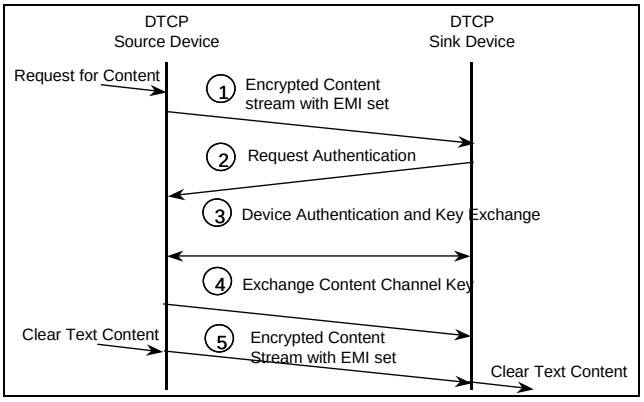

Figure 5: DTCP operation

Step 1: The source device is requested to initiate the transmission of a stream of protected content. The embedded CCI of the content is examined to determine the appropriate EMI value (e.g., "copy_once, "copy_never," or "copy_no_more") to associate with the encrypted content stream. The source device may choose to transmit an empty content stream until at least

one device has completed the appropriate authentication procedure.

Step 2: Upon receiving the content stream, the sink device inspects the EMI to determine the copy protection status of the content. If the content is marked "copy_never," the sink device requests that the source device initiate Full AKE. If the content is marked "copy_once" or "copy_no_more" the sink device can request Restricted AKE if Full Authentication is not available. If the sink device has previously performed the appropriate authentication, it can immediately proceed to Step 4.

Step 3: When the source device receives the authentication request, it proceeds with the type of authentication requested by the sink device, ensuring that Full AKE is performed if the content is marked "copy_never."

Step 4: Once the devices have completed AKE, the keys required to access the encrypted content stream are exchanged between the devices.

Step 5: Encrypted content flows between the devices.

The application of DTCP is not limited to 1394, USB, and MOST. It is suitable for use with any digital interconnect that supports bi-directional communications. With the emergence of wired and wireless IP networking technologies within the home environment, including Ethernet and 802.11 wireless solutions, efforts are currently underway to add DTCP support to that infrastructure. The principal challenges associated with this mapping are to constrain DTCP protected content to the home and personal network space and prevent anonymous, "hot-spot"-based sharing of content.

CONCLUSION

Figure 6 shows an example of end-to-end protection within the digital home, using representative existing content protection solutions.

The strength and completeness of the end-to-end protection stems from an overall system architecture that ensures the individual technologies form a protected whole. Intel is promoting the development of a comprehensive, compatible content protection system, because it will benefit content owners, content providers, device manufacturers, and above all, consumers. Consumers are demonstrating a growing desire to move digital content among devices such as personal computers, DVD players, set-top boxes, and digital televisions. A variety of content protection technologies have been developed to protect digital content within each of these devices. Chaining these point solutions together will enable digital content usage throughout the home, maintain the integrity of the content, and minimize implementation burdens on developers.

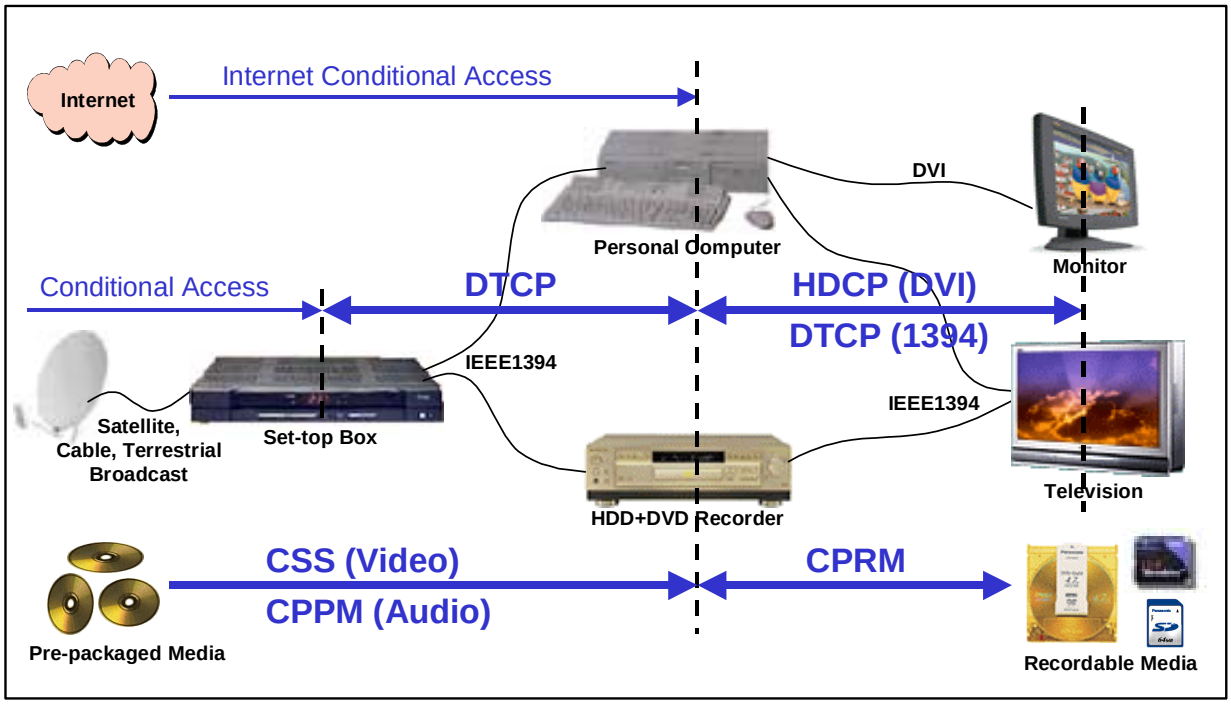

Figure 6: Example of end-to-end content protection

REFERENCES

[1] C. Brendan Traw, "Protecting Digital Content within the Home," *IEEE Magazine*, October 2001, pp. 42-47.

[2] Bloom, Cox, Kalker, Linnartz, Miller, and Traw, "Copy Protection for DVD Video," in *Proceedings of IEEE*, Vol. 87, No. 7, July 1999, pp. 1267-1276.

[3] Marks and Turnbull, "Technical Protection Measures: The Intersection of Technology, Law, and Commercial Licenses," *Workshop on Implementation Issues of the WIPO Copyright Treaty and the WIPO Performances and Phonograms Treaty*, World Intellectual Property Organization, Geneva, Switzerland, 1999.

[4] Lotspiech, Nusser, and Pestoni, "Broadcast Encryption's Bright Future," *IEEE Computer*, Vol. 35, No. 8, August 2002, pp. 57-63.

AUTHORS' BIOGRAPHIES

Michael Ripley is a Staff Engineer within Intel's Corporate Technology Group. He joined Intel in 1995, and during the past several years has been the technical lead on a number of Intel's content protection efforts, including development of CPPM for DVD-Audio and CPRM for several formats. Mike is the recipient of various awards, including the Intel Achievement Award. His e-mail address is michael.ripley@intel.com.

C. Brendan S. Traw is a Senior Principal Engineer within Intel's Corporate Technology Group. His current focus is on the management and protection of entertainment content within emerging digital environments. He has lead the development of industry-leading content protection solutions including DTCP for digital networks, HDCP for DVI, CPPM for prerecorded DVD Audio, and CPRM for various recordable media formats. Brendan is the author of over a dozen papers and book chapters and has received eight US patents. He received a Ph.D. degree in Computer Science from the University of Pennsylvania. His e-mail address is brendan.traw@intel.com.

Stephen Balogh is a Business Development/Marketing Manager within Intel's Corporate Technology Group. He has an extensive background in establishing, growing, and managing content protection technology licensing and key generation infrastructures. Steve has held leading positions in CP technology diffusion efforts including President of the Digital Content Protection LLC, President of the Digital Transmission Licensing Administrator LLC, Manager of the DTCP Key generation facility, and he is presently Intel's Steering Committee member for the DVD Forum. His e-mail address is stephen.p.balogh@intel.com.

Michael Reed is a Marketing Director within Intel's Corporate Technology Group. He has 16 years of marketing experience within the technology industry. Michael joined Intel in 2000. Prior to that he was the Vice President of Marketing for Diamond Multimedia's Rio Division, makers of the industry-leading Rio digital audio player. Michael holds a B.S. degree in Computer Science from Bowling Green State University and an MBA degree from the University of Oregon. His e-mail address is michael.j.reed@intel.com.

Meeting the Demands of the Digital Home with High-Speed Multi-Hop Wireless Networks

Lakshman Krishnamurthy, Network Architecture Lab, Intel Corporation
Steven Conner, Network Architecture Lab, Intel Corporation
Mark Yarvis, Network Architecture Lab, Intel Corporation
Jasmeet Chhabra, Network Architecture Lab, Intel Corporation
Carl Ellison, Network Architecture Lab, Intel Corporation
Chuck Brabenac, Communications and Interconnect Technology Lab, Intel Corporation
Ernest Tsui, Communications and Interconnect Technology Lab, Intel Corporation

ABSTRACT

In the near future, homes will be equipped with wireless networks that bridge data and consumer electronics networks, interconnecting desktop PCs, mobile laptops and handhelds, High-Definition TVs (HDTVs), DVD players, camcorders, and other multimedia devices. This environment introduces new wireless network requirements, including high and dependable bandwidth, low latency, and coverage throughout the home. Multi-hop wireless technology offers unique benefits for creating a high-speed, robust home wireless network. However, to support these demanding usage models, significant wireless networking innovations are required across the physical, MAC, and routing layers, and solutions need to be found for higher level issues such as Quality of Service (QoS) guarantees, device discovery, and security. In addition, user acceptance of multi-hop wireless networks will require ease of installation. Intel R&D is currently researching self-organizing multi-hop wireless networks for home environments. This paper introduces the technologies and tradeoffs needed to create a multi-hop wireless home network, identifying benefits and limitations. In particular, we describe usage scenarios and assumptions that drive the requirements. Finally, we provide an outline of the key technology problems that must be solved and recommend the necessary next steps to make this vision a reality.

INTRODUCTION

Wireless Local Area Network (WLAN) technologies are beginning to gain a foothold in the home computer market. Wireless networks allow the home user to share data and Internet access without the inconvenience and cost of pulling cables through walls or under floors and without unsightly network jacks. Pulling cables and installing new network jacks are particularly challenging in existing homes and apartments. In addition, wireless LANs provide the convenience of untethered computing for laptops and handhelds from anywhere in or around a house.

The benefits of wireless need not be limited to computer networking. As the bandwidth of wireless networks increases, audio/video home entertainment will be the next target, replacing device-to-device cabling as well as providing distribution throughout the home. Rather than maintaining separate networks for different types of devices, as is common with wireless LANs and cordless telephones, a unified technology is desirable to reduce the cost and complexity of installation and to allow cross-device communication and functionality. For instance, a consumer should be able to insert a DVD into a player in the living room and watch it on a TV in the bedroom or on a laptop on the back porch. Such a network will need to span the entire home, allowing any two devices to communicate. Figure 1 below shows typical devices in a home: entertainment devices, HDTVs, DVD players, game consoles, PCs, and laptops. A multi-hop network that interconnects these clusters is also shown.

Figure 1: Devices in a typical home clustered in various rooms

In a multi-hop network, a node transmits with low power to reach nearby neighboring nodes (routers), which will forward the data toward the intended destination. Using a multi-hop network provides significant benefits assuming a limited channel capacity. The alternative is a single-hop network, where each device transmits with enough power to be received by any other device in the home. When two devices transmit simultaneously, the resulting channel contention can limit capacity in a high-bandwidth environment. The typical solution to this problem is to divide the channel into subchannels by some combination of frequency, time, or coding. Dividing the channel into N subchannels allows N devices to transmit without contention. However, each of these subchannels will have a capacity of at most $1/N$, which must be sufficient to carry the desired traffic.

Multi-hop networking, on the other hand, increases the aggregate capacity of the network by using lower transmit power to only reach nearby neighboring nodes. This allows channel re-use, thereby improving spatial capacity. By using lower transmit power, devices at different locations can transmit simultaneously without interference. Thus, despite a limitation of N channels, more than N devices can potentially transmit simultaneously without contention.

In addition to preserving spatial capacity, the low-transmission power requirements of multi-hop networks allow them to support higher bandwidth despite Federal Communications Commission (FCC) regulations that limit maximum transmission power. At a given transmission power level, the number of reception errors increases as transmission distance increases, due to a diminishing signal-to-noise ratio. To allow transmission over greater distances, wireless devices use variable forward error correction encoding schemes or step-down to simpler modulation schemes [4]. These schemes result in a decrease in channel capacity as the transmission distance increases (Table 1). Multi-hop networking avoids this problem by transmitting data over several short hops rather than one large hop.

	802.11a	802.11b	802.11 g	UWB
5m	54	11	54	660
10m	48	11	54	188
20m	36	11	48	20
30m	24	11	36	5
40m	18	5	24	2
50m	12	5	18	1
60m	9	2	12	0.5

Table 1: Raw channel capacity (in Mbps) for several wireless technologies at various communication ranges. Higher bandwidths are available at shorter distances irrespective of the technology or standard.

Multi-hop communication also provides greater redundancy. When the network is dense, each device can have many neighbors within communication range, potentially creating multiple paths between two communicating devices. In the presence of localized interference or attenuation, such as a person standing in a room, a multi-hop network can route data along an alternate path. In a single-hop network, it is not possible to route around interferences or degradation between two devices.

While multi-hop networking solves many problems, many challenges remain before it can become a reality. In particular, a home multi-hop networking environment must be self-administered and cannot require the user to be technologically savvy. Currently, the installation of even a single-hop wireless LAN is a challenge [8] [10]. Identifying the optimal location for access points, selecting communication channels, setting the transmit power, and enabling security all require a high level of technical expertise and engineering. This paper discusses approaches to these and other issues in the context of multi-hop networking. While we do not offer definitive solutions we do suggest a number of research directions that will help to make high-speed multi-hop wireless networks for the home a reality.

USAGE SCENARIOS AND REQUIREMENTS

The first step in the process of creating deployable multi-hop networks is to identify the requirements of the home network based on network usage scenarios. Home networks may be deployed in a variety of domains:

- a small house, well separated from other houses
- an apartment in an apartment complex
- a large suburban house, with computer-savvy children
- a townhouse in a row of townhouses
- a college dorm or other similar facilities

Each of these domains may be viewed as a collection of interconnected clusters of devices. We need to answer three questions in connection with this network. First, how does a user install a network in a diverse home environment and maintain interoperability? Second, what are the traffic characteristics of devices and applications in this network? Third, how do multiple

networks co-exist in each domain? The following subsections explore each of these questions.

Interoperability Requirements

The interoperability issue is not unique to multi-hop wireless networks. However, any multi-hop solutions must account for the diversity of devices in the home network. In the future, several classes of wireless networking devices will be purchased by consumers for installation in the home. Consumer electronics and computing devices such as DVD players, televisions, remote-control devices, and handheld computers will come with radios pre-installed (e.g., Radio Free Intel [11]).

In the home environment, the wide range of types of equipment means that we cannot use the same radio everywhere. In many cases, wireless interfaces included in home products should be optimized for a particular task, such as short-range wire replacement [1] [2]. But even these interfaces are diverse. For example, using an expensive high-bandwidth radio on a wireless keyboard is wasteful and will result in a needless increase in cost. On the other hand, a high-bandwidth radio might be appropriate for a DVD player as the increased cost will be small relative to the total cost of the device, and it would support the high-bandwidth usage requirements of the player.

In order to support this rich diversity of devices, the multi-hop network must interface with each kind of device.

Installation Requirements

To enable longer-range multi-hop communication between devices distributed throughout the home, we envision the use of specialized low-cost router devices. Such routers might be packaged in compact form-factors that can be conveniently installed by plugging them into power outlets throughout the home. External add-on radios will convert legacy devices such as home appliances and older audio and video equipment into wireless-enabled devices.

For a non-technical home or apartment dweller to install and configure a home network, it is imperative that multi-hop wireless networking not increase the installation complexity of consumer electronics equipment. Given the large variation and unpredictability of Radio Frequency (RF) propagation in different deployment environments [6] and given the lack of technical expertise by typical installers (homeowners), tools to aid in correct deployment will be very important to ensure good connectivity between devices.

Traffic Characterization

Traffic in this network may be generated by a variety of applications ranging from Internet browsing, data backup, and telephony, to entertainment and gaming. These applications generate a range of traffic patterns.

Interactive traffic: PCs, laptops, and handheld devices will require regular Internet access over a home broadband connection. For applications such as Web surfing, digital photos, and e-mail, bottlenecks are likely to remain in the Internet or on the broadband connection. This traffic has more stringent latency requirements but less stringent bandwidth requirements than entertainment applications. Traffic generated by devices such as wireless mice and keyboards require very low latencies, but these devices are likely to require only single-hop communication over a short distance.

Bulk traffic: Applications such as data back-up, network file storage, and printing require higher bandwidth and, unless metered, may even saturate the available bandwidth. Such traffic would require less priority than interactive applications.

Audio applications: Cordless telephones are common today. Extending these devices to support Internet telephony is a logical next step. This class of applications would require low latency and uninterrupted connectivity while roaming throughout the home.

Entertainment-quality traffic: Audio Visual (A/V) home entertainment devices such as High-Definition TVs (HDTVs), DVDs, stereos, camcorders, multi-media PCs, and musical instruments require high bandwidth. Many of the usage scenarios call for communication in close proximities, but the ability of users to store legal digital content in Personal Video Recorders (PVRs) for distribution around the home introduces a traffic pattern beyond device clusters.

This discussion demonstrates the need for different traffic classification and prioritization within the nodes and the network. We have demonstrated the need to support low-latency and high-bandwidth real-time applications. For multi-hop networking to be useful, it must be able to provide sufficient bandwidth in order to differentiate itself from single-hop solutions.

Coexistence

When multiple networks exist within radio range of one another, these networks must be able to coexist. This situation typically occurs in high-density housing (apartments, college dorms, townhouses). The coexistence question is not unique to multi-hop networks [23]; these networks make the problem harder to solve by requiring a solution at every node in the network.

Coexistence imposes specific requirements for channelization, routing, Quality of Service (QoS), and security.

Channelization

One possible approach to coexistence is the *possessive* one: "I use my network for my devices and you use your network for your devices." Possessiveness may lead to disaster when it comes to channelization. If two neighbors choose channels for their networks independently, nothing will stop them from choosing the same channels and thus interfering with each other [13].

Channel selection must be cooperative and, to minimize the human administrative effort, should be achieved automatically by the network devices without requiring the network owners to meet and make plans. In a dense enough apartment complex, a network owner might not even know which neighbor is running the competing wireless network.

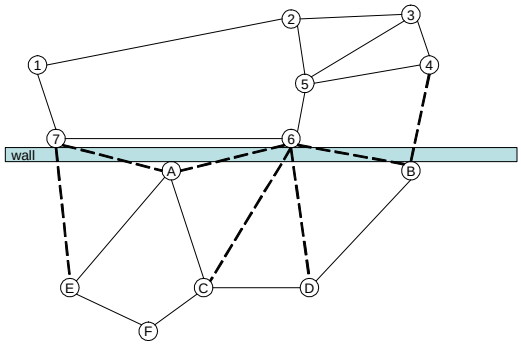

Figure 2: Two neighboring networks where Nodes 1 through 7 belong to one apartment and Nodes A through F belong to the adjacent apartment

Routing

If the two networks in Figure 2 must remain isolated, the dotted links will not be used. As a result, each network will offer lower performance to its owner than if the two networks were to cooperate in routing. For example, consider the network shown in Figure 2 and assume that the wall shown does not attenuate the signal any more than the internal walls in each apartment (not shown in the figure). Also assume that message transmission time is proportional to the square of the distance covered. Using a simple simulation, we get a throughput improvement between nodes {1, 7} and {2, 3, 4, 5, 6} of a factor of 3.6, while end-to-end latency improves by a factor between 1.5 and 1.9. For link A-B, the throughput improves by a factor of 1.5, and the latency improves by a factor of 1.6. While this result is based on a simplistic simulation, it illustrates the fact that cooperation between coexisting networks could give enough performance improvement to warrant the effort.

Quality of Service

Even an isolated network requires QoS routing to support streaming audio and video [19]. To achieve the desired QoS in a multi-hop network, every node needs to cooperate with its neighbors whether it is owned by the same user or not. Using low-power radios reduces the severity of the problem by limiting the range, and therefore reducing the number of cross-wall radio neighbors. If we allow neighboring networks to cooperate in order to gain a performance advantage, then we must ensure fairness of network usage. This, too, is a QoS problem that depends first on defining "fairness."

Security

In a multi-hop network, security is required to enable four major protection functions:

End-device and router introduction: When new end-point devices (e.g., DVD players and music jukeboxes) or router nodes are added, their introduction must be authenticated. This essentially determines the notion of who owns a particular device. This problem is not unique to a multi-hop network, but solutions to the introduction problem must work across the network.

User data integrity and secrecy: Link-level encryption has been proposed for the protection of both data and access in single-hop networks [4]. As discussed in a later section, end-to-end encryption is more appropriate for protecting user data than link-level encryption.

Device control and authentication: Commands sent to devices in a network must also be authenticated. In a wireless network, it is particularly important for end-devices to authenticate users before granting access and control permissions. For example, nodes in a neighboring network should not be allowed to control the television next door, nor should they be able to access personal home movies stored on a neighbor's media server. These issues must be solved in the context of a multi-hop network.

Network authentication and coexistence: Solving the network authentication problem is especially important with respect to the coexistence problem. A hostile neighbor or intruder could introduce inaccurate routing information or inject an unauthorized traffic load. Packets containing routing updates or QoS-protected streams must be authenticated. Implicit authentication by encryption is a poor substitute for real authentication. Moreover, relying on link encryption is a poor choice in multi-hop networks as end-devices lose access to origin authentication information.

SOLUTION SPACE

In the previous section, we identified the requirements for interoperability, installation, supporting home multi-media traffic, and coexistence. This section looks at the solution choices and tradeoffs that meet these requirements.

Interoperability

In the requirements section, we noted that a typical home could have a variety of devices using a mixture of Physical layers (PHY) and Medium Access Control (MAC) layers, not necessarily directly interoperable. In this diverse home environment, we envision a multi-hop network, using dedicated homogenous router devices. These routers use the same PHY/MAC layer for inter-router communication.

Using the same PHY/MAC in the multi-hop backbone provides an opportunity to better control the network, which makes it easier to provide entertainment-quality connectivity throughout the home. However, the multi-hop network must still interoperate with the wide variety of home devices. As shown in Figure 3, the leaf nodes of the multi-hop network must support every possible PHY/MAC standard that may be used in the home.

Figure 3: A multi-hop backbone must interface with multiple PHY/MAC layers in the home

This issue leads to an important cost-complexity tradeoff. Providing every possible PHY/MAC hardware implementation in each router node is simply too expensive. Creating many different kinds of router devices, each with a particular PHY/MAC implementation increases the complexity of the installation and limits support for mobility. An alternative is to create one type of device (or a small number of them) combining a limited number of PHY/MAC choices based on the likely set of home devices. The cost of this approach needs to be traded off with a third approach that makes use of reconfigurable radios.

Reconfigurable radio, or Software-Defined Radio (SDR) technology, allows a single piece of silicon to be reconfigured to implement many different PHYs and MACs. Such a flexible radio can allow interoperability with a larger set of end-point devices at a lower cost than including multiple radios in the same device. More information on reconfigurable radios and network configuration protocols needed to support self-configuration in the network may be found in Appendix A; these issues also apply to single-hop networks.

Multi-Hop Network Installation

Using software-defined reconfigurable radios will address the issues of legacy equipment and non-interoperable wireless standards to some degree. However, as described in the requirements section, multi-hop wireless networks must also support ease of installation and placement of nodes in a multi-hop network.

A rule of thumb based on typical deployment scenarios could be supplied to users as a starting point. An example of such a rule could be that wireless routers should typically be deployed every 10 feet. However, in a real home, Radio Frequency (RF) shadows are likely to exist due to home furnishings, household items, people, metal, and other attenuators built into the building structure, thereby greatly limiting the usefulness of such rules of thumb.

Given such unpredictability, users will need help deciding where specifically to deploy their nodes. One possibility is to provide a feedback mechanism, perhaps through Light-Emitting Diodes (LEDs), indicating the signal strength on each node. Such an approach would be particularly helpful for one-hop networks, where the user must simply make sure each device is close enough to another device. However, in the multi-hop router case, each router must be strategically placed to provide sufficient connectivity among multiple nodes (often in more than one direction). One technique to verify signal strength between specific pairs of nodes is to install a switch on each node allowing the user to select a single pair of nodes at a time to verify their connectivity. Using this technique to deploy even a few nodes in a home may become tedious.

An alternate technique would be to deploy the initial network using a rule of thumb, and then to connect a PC or other specialized network monitoring device to the network, which would collect signal strength and connectivity statistics from each node in the network and display a simplified summary of the results to the user.

Any nodes that are not discovered at the network monitoring location indicate a network partition. This information would let the user know that some routers must be moved, or that more routers must be installed in the area between the detected and missing nodes. Nodes that are expected to be in communication with one another, but which have low inter-communication signal strength, should be supplemented by placing a router between them. Such a network monitoring device would make it easy for users to gain a global view of how well they deployed their network.

SUPPORTING DIGITAL HOME TRAFFIC

Beyond the interoperability and installation issues, the Quality of Service (QoS) need of traffic described in the requirements section must be met. Towards this end, we look at the alternatives and tradeoffs in the area of routing, QoS support and channelization.

Routing and QoS Support

The topology of a multi-hop wireless network is a graph with an edge between each pair of nodes that can communicate directly. Even when all nodes are stationary, the network topology may be constantly changing due to variations in RF propagation and interference [6]. Thus, the network topology will likely be different when it rains than when it is sunny and different during a party than when a house is empty. A multi-hop network must adapt to the dynamically changing topology to allow nodes to communicate.

One approach to routing data in a multi-hop network is to have every node repeat every new packet received. This approach is advantageous in that it is simple, it routes between any two nodes, and it utilizes all redundant paths for greater reliability. However, because every node sends every packet once, this approach does not benefit from spatial re-use. Instead, this type of network "flooding" is typically used to identify a route through the network over which many data packets can then flow.

Network routes can be identified proactively or reactively. Proactive approaches maintain connectivity and resource availability information even when no traffic is present [20]. This approach reduces start-up latency, but wastes power and bandwidth when no routes are required and when the network changes frequently (in which case the information becomes stale). In contrast, a reactive routing approach identifies a route only after a packet transmission request or stream connection request is received [12] [21]. In networks with low utilization or highly dynamic topologies, a reactive approach is typically superior [5]. A hybrid approach is possible in which some paths are maintained proactively while others are identified reactively. A network might also switch between proactive and reactive routing in response to network load.

Multi-Hop Channelization

Channelization is often used to alleviate the problems of single-hop and two-hop interference in a wireless network. The Request to Send (RTS)/Clear to Send (CTS) scheme, part of the Distributed Coordination Function (DCF) in WLAN standards, such as 802.11a or 802.11e [4], is one technique that may be used for this purpose. This technique allows nodes to acquire the channel and suppress other nodes from contending for the same channel. This, however, does not take advantage of multiple channels available in most standards [1] [2] [4]. These schemes allow nodes separated by two hops to communicate at the same time, if they choose non-conflicting channels. In a multi-hop network, such a channelization scheme may allow nodes to take advantage of multiple access features.

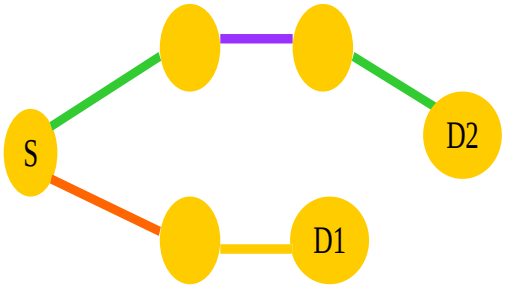

Figure 4: Colors representing channels (sub bands) are assigned in order to avoid interference

Solving this problem is akin to solving channelization based on the well-known graph coloring problem (Figure 4), which is known to be NP-complete[1]. However, heuristic solutions may be implemented in the network using static, centralized, or distributed techniques. The small scale of home networks makes brute-force approaches possible.

QoS Routing and Multi-Hop Channelization

Routing with QoS may be implemented assuming that the underlying network supports a contention-free environment, derived through channelization schemes. However, if multi-hop channelization is completed a priori, without considering the needs of traffic, there is no way to guarantee the needed QoS (even using

[1] "No polynomial time algorithm has been discovered for this class of problem" [7].

RTS/CTS schemes). To maintain QoS in a dynamic network with varying application and usage scenarios, two problems must be addressed: resource management across the network and resource management at the node/link level.

1. *Resource management across the network–routing and channel assignment:* For a given connectivity graph of nodes and set of flows with a known link capacity requirement, channel resources (time and frequency) need to be allocated in an efficient manner. A choice made by one node will affect all other nodes in the network.
2. *Resource management at a single node and link level:* For a given set of flows entering and leaving a node and a known link capacity (assuming routing is complete and an end-to-end path is set up), the link must service flows to meet the QoS needs of each flow. At the outset, this problem seems similar to its wired counterpart. But new evidence suggests that wireless channel characteristics require special attention [16].

While separation of routing and channel assignment from QoS simplifies the path assignment problem, it is at a cost to the effective system capacity, as the flows are not known a priori. On the other hand, if we increase the complexity of the QoS path set-up problem, we can solve the channel assignment problem along with the QoS constraints (Figure 5). Even though the allocation in Figure 5 uses more sub bands than the one shown in Figure 4, the sub bands are allocated based on the QoS constraints. This approach results in tighter QoS guarantees and better overall network utilization.

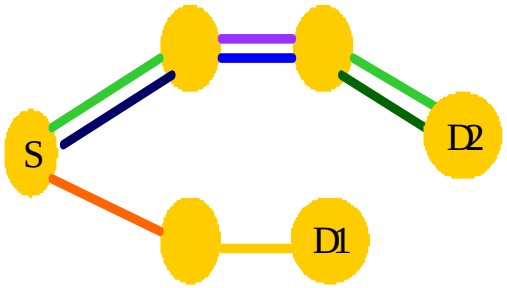

Figure 5: Colors representing channels are assigned based on QoS requirements

Further performance advantages may be achieved by enabling nodes from neighboring networks to cooperatively provide QoS (Figure 2). For example, we might predict the capacity that would be available if only using nodes from a given network and allow QoS reservations up to that capacity, but then allow messages to actually flow over any available nodes, thus taking advantage of extra performance. The remaining bandwidth would serve as leeway for the QoS algorithm or extra bandwidth for non-reserved uses (such as Web surfing).

These issues require the tradeoff of computation complexity to increase system utilization. This leads to the following questions:

1. What is the maximum size of the network to which the design of our algorithms should scale?
2. Should we jointly optimize the channel assignment and QoS path selection problems?
3. Should we assume all nodes have the same MAC and routing capabilities?
4. How do we allow new data traffic to preempt old data traffic?

Centralized Vs. Distributed Path Selection

Additionally, an important question in the choice of a QoS routing algorithm is whether to use a central controller or make decisions in a distributed manner. A decentralized QoS routing algorithm distributes the complexity across the network. This may increase the cost of nodes in the network, but distributing the decision making relaxes the need for network-wide synchronization and channel assignment, increasing the network scalability. A distributed approach also eliminates the need to communicate with a centralized controller, reducing the traffic overhead. Finally, while a decentralized approach can react more quickly to local changes, by eliminating the global-view of the network, it generally makes non-optimal decisions.

In a home networking environment, the relatively small network size and cost considerations are probably the most important factors to consider for making a decision between the two approaches. In such an environment, a centralized approach is more advantageous, as we can move most of the cost/complexity to the central nodes. The relatively small network size also implies that the overhead traffic flows are small, the reaction to the changes is quicker, and we may be able to make globally optimal decisions on the central controller.

Using a centralized method to solve problems of channelization, multi-hop routing, and path selection in a combined fashion will allow the system to meet the requirements of the Digital Home. Additionally, these algorithms need to comprehend security in order to protect the performance of the network. Not doing so will render any effort on their part useless in many of the deployment environments.

SECURITY

Security for home networking is a large topic covered more fully in "Home Network Security" also in this issue of the *Intel Technology Journal* [9]. This section mainly covers those security issues directly related to multi-hop networking. These issues were highlighted in the requirements section, namely, device introduction, user data/command integrity and secrecy, and network authentication.

Device and Router Introduction

In general, the introduction problem has been addressed by Universal Plug and Play (UPnP*) Security [9], and further work on simplifying introduction in the general case is being conducted as part of that R&D activity. This device introduction is a key component that enables solutions to the other security issues.

User Data and Command Security

User data security is not unique to multi-hop networks and is orthogonal to other security issues in these networks. The "Home Network Security" paper [9] in this issue of the *Intel Technology Journal* makes a case for end-to-end security. A summary of the motivation and its relationship to multi-hop networking is provided here.

Many homes require multiple security domains. For example, when the home contains older teenagers, adult roommates, boarders, or guests, we would expect each to have his or her own security domain, but these people would all share the same physical multi-hop network.

A security domain is a set of entities (devices or even processes) that are allowed to work together, excluding other entities. When traffic is encrypted for confidentiality, members of a domain can read one another's traffic. When user-command traffic (to control DVD players, jukeboxes, etc.) is strongly authenticated and authorized, one member of a domain is permitted to issue commands to another. When a home has multiple security domains, security cannot be implemented at the link layer between two directly communicating devices; rather, it must be implemented as a protocol at a higher layer between the end points in a communication that may traverse multiple devices in a multi-hop network.

One advantage of a higher-layer end-to-end security protocol is that the physical network carrying the user's traffic is not responsible for providing security for that traffic, and as a result, neighboring networks may be used to help carry a user's traffic without introducing security concerns.

Network Authentication

Even though user data security and integrity is provided end-to-end, devices still need to ensure network packets are received from authorized nodes in a multi-hop network. In particular, authentication must be supported for the one-hop origin of all packets, while the multi-hop originator must be authenticated for packets that access or control QoS or channelization on routers. A packet could be implicitly authenticated by encrypting it via a symmetric key known only to the origin and the verifier, but such a mechanism is more expensive than necessary and requires n^2 keys, where n is the number of network nodes. An alternative is to use a routing header in packets that can be authenticated with a lower computation overhead, resulting in a cost savings for routing nodes.

Routing and QoS Security

A routing algorithm makes decisions based on information such as latency and bandwidth between nodes. While the algorithm can acquire this information from neighboring network nodes, it is not possible to trust the routing information without verification. Without authentication, it is not possible to know if neighboring nodes belong to a possible adversary.

If a malicious node were used for routing, one way to interfere with the routing of messages is to insert an artificial delay in the message delivery path. It is not possible to prevent a malicious node from doing this. However, one can learn of this delay and route around it. This is therefore a normal routing problem rather than a routing security problem.

Another way to interfere with routing decisions is to advertise more bandwidth or lower latency for delivery of messages than can truly be achieved. One can test any advertised path for actual performance, provided that the destination node can authenticate a reply to a ping from the sending node. If we use public-key cryptography for this authentication, then we need a way to bind a node address to a public key. If the node address is the hash of the public key, then we achieve that binding without any additional cost.

Similarly, QoS establishment and maintenance requires nodes to trust information they learn from neighbors. QoS security cannot be free of administration. At the very least, a network node must learn which other nodes are part of its network. Clients will know their own QoS requirements and can make a reservation request from the nearest node belonging to the same owner. However, in a case of over-reservation, some human

* Other brands and names are the property of their respective owners.

administration will be required to establish priorities for different classes of use.

Solving the Coexistence Problem

Routing between neighboring networks (i.e., in adjacent apartments) may not be entirely independent and may require interference avoidance. Because security is provided end-to-end, the problem of coexistence between networks is only about sharing available bandwidth and not about maintaining privacy or data integrity. Three sharing strategies are possible:

1. Neighbors could choose to compete for channel access with no coordination between channel assignments. Such competition introduces channel contention. Channel contention reduces the overall channel capacity at high load and makes it difficult to predict the realizable channel capacity.
2. Neighbors could agree to statically split the available channels to avoid interference, thus introducing a constraint that must be reflected in QoS routing decisions. Channel assignment would be made independent of load, arbitrarily restricting the maximum bandwidth provided to each user wherever the networks overlap (typically along the common wall).
3. Finally, neighbors could agree to cooperate in channel assignment. While channel assignments could be initially made arbitrarily, one network could "borrow" a channel from the neighboring network (when not in use) or route traffic through a node on the neighboring network (i.e., nodes along a common wall) in response to high load. In this case, the network must be prepared to react if these resources are later reallocated by the neighboring network. One approach is to only allow neighboring resources to be borrowed for low-priority or non-QoS traffic. The cooperative approach decreases privacy, since neighboring networks must exchange load requirements to achieve QoS scheduling.

In each of these cases, the one-hop origin of packets containing routing updates or data from QoS-protected streams must be authenticated. We have described solutions for such authentication in the previous section.

RELATED WORK

The use of multi-hop wireless networking is gaining traction in everyday life. In fact, several companies already provide services using multi-hop wireless networks. MeshNetworks [17], for example, uses multi-hop routing between nodes installed on light poles, buildings, vehicles, and end-user devices such as laptops and handhelds to provide Internet access to subscribers in cities. Nokia supplies kits to enable multi-hop networking between nodes installed on rooftops [18] to provide broadband Internet access. Most of these services focus on extending the reach of Internet access beyond the range typically supported by access points.

Multi-hop wireless networks exhibit many unique problems, but they also overlap with wired home networks. Device discovery and auto-configuration protocols such as Universal Plug and Play (UPnP*) [3] that were originally designed for wired IP networks can easily be applied to wireless networks. Wired home network security issues [9] also must be dealt with in wireless networks. Finally, QoS routing [19] and preemption have been extensively explored for wired networks.

QoS routing in ad hoc wireless networks has only recently been investigated in simulations. Several schemes were originally developed using MAC-independent techniques based on existing ad hoc routing protocols [14] [22]. More recently, advances have been made by optimizing QoS in multi-hop wireless networks with Code Division Multiple Access (CDMA) [15] and Time Division Multiple Access (TDMA) [24].

CONCLUSION

Multi-hop wireless technology offers unique benefits for creating a high-speed, robust home wireless network. The benefits over traditional infrastructure wireless networks include extending coverage without requiring deployment of multiple wired base stations, increasing utilization of spatial capacity to realize higher throughput, and offering alternate communication paths to provide failure recovery and better throughput.

To support the demanding usage models for the digital home, wireless networking innovations are required across the physical, MAC, and routing layers. In addition, higher level issues such as Quality of Service (QoS) guarantees, device discovery, and security must be solved. We have outlined a roadmap to meet these challenges; solving them will usher in new opportunities for wireless networking in the digital home.

ACKNOWLEDGMENTS

We acknowledge contributions from Stephen Wood, Sumit Roy, Rahul Mangharam and Arjunan Rajeswaran. We also thank reviewers of this paper for their valuable comments and suggestions.

* Other brands and names are the property of their respective owners.

APPENDIX A

Reconfigurable Radios

Support for a wide variety of consumer wireless PHY/MAC layers will be possible by using low-cost reconfigurable radios at the "edge" of the network. Furthermore, within the network it is highly advantageous, due to the varying Quality of Service (QoS) and bandwidth requirements, to have features like variable modulation [4] in the PHY to optimize the throughput based on channel loading and propagation conditions (party-time vs. home alone). So an example radio might have a very flexible PHY/MAC to interface externally, and it might use a common technique like 802.11, which has variable modulations, to route within the network itself.

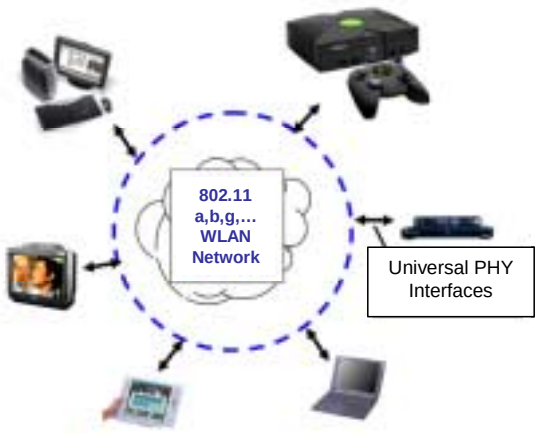

Figure 6: Reconfigurable radio interfaces allow multi-hop and other wireless networks to bridge multiple PHY technologies

Reconfigurable radios should then be configurable to handle 802.11 a,b,g, low-rate consumer devices (lower requirement than 802.11) [25], and perhaps future high-data-rate (~100-500 Mbps) wireless personal area networks [1] and USB standards.

Radios should also be "intelligent," assessing the environment and channel propagation around them, including interference. First, the radio should be able to broadcast and receive beacon signals introducing itself to the surrounding devices that are "reachable." Next, the radio should be able to assess the existing frequency channels, interference, and noise conditions on each channel, so that it can determine (perhaps in conjunction with central or distributed control algorithms) the best channel to support a given bandwidth and QoS requirement. Having measured the channel, the radio can use the minimal power (to save power to the power amplifier) to make a reliable connection. Thus, unlike conventional radios, a desirable radio requires a simple processor to help direct its various reconfigurable modes.

Finally, to lower cost, the reconfigurable radio (which by its very name implies higher cost) should take advantage of extensive silicon (Si) re-use. The baseband radio will have re-usable components such as various filters, digital mixers, etc., and the Radio Frequency (RF) portion will have a degree of agile frequency capability (for example, it will be able to jump to the 5.2 GHz band if there is microwave interference detected at the 2.4 GHz band). The MAC layer will also be flexible to allow download from the host of the typical extensive memory resources required to support the wide variety of anticipated MAC protocols: this will allow significant cost savings via memory re-use.

In summary, desirable home network radio architecture will require intelligence and reconfigurability to minimize power and keep QoS high while utilizing extensive Si re-use to keep costs down.

REFERENCES

1. "IEEE Wireless Personal Area Networks (WPAN) 802.15 High Rate (HR) Task Group (TG3)." http://grouper.ieee.org/groups/802/15/pub/TG3.html
2. "Specification of the Bluetooth System," Ver. 1.1, February 22, 2001. http://www.bluetooth.com
3. "Universal Plug and Play (UPnP) Device Architecture Document," Ver. 1.0, June 2000. http://www.upnp.org
4. "Wireless LAN Medium Access Control and Physical Layer Specifications," Aug. 1999, IEEE 802.11 Standard (IEEE Computer Society LAN MAN Standards Committee).
5. J. Broch, D. A. Maltz, D. B. Johnson, Y. C. Hu, and J. Jetcheva, "A Performance Comparison of Multi Hop Wireless Ad-Hoc Network Routing Protocols," in *Proc. of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM '98)*, Dallas, TX, October 1998.
6. D. Cheung and C. Prettie, "A Path Loss Comparison Between the 5 GHz UNII Band (802.11a) and the 2.4 GHz ISM Band (802.11b)," Intel Corp. Technical Report, January 2002. http://impulse.usc.edu/resources/802_11a-vs-b_report.pdf
7. T. H. Cormen, C. E. Leiserson, R. L Rivest, *Introduction to Algorithms*, McGraw-Hill, 1998.
8. S. Diaz, "Wireless networking: daunting but doable,"

San Jose Mercury News, May 29, 2002. http://www.siliconvalley.com/mld/siliconvalley/3363171.htm

9. C. Ellison, "Home Network Security," *Intel Technology Journal* Vol. 6, Issue 4, 2002.

10. J. Geier, "802.11a/b Site Survey: A Testimonial," *802.11 Planet*, October 10, 2002. http://www.80211-planet.com/columns/article.php/1479831 .

11. P. Gelsinger, Intel Developer Forum Keynote, Feb. 28, 2002. http://www.intel.com/pressroom/archive/speeches/gelsinger20020228.htm

12. D. B. Johnson, "Dynamic Source Routing in Ad Hoc Wireless Networks," in *Mobile Computing* (ed. T. Imielinski and H. Korth), 1996.

13. J. Kosseff and E. Hand, "Wireless channel use sets up turf battle," *The Oregonian*, August 19, 2002.

14. S. Lee and A. T. Campbell, "INSIGNIA: In-band signaling support for QoS in mobile ad hoc networks," in *Proc. of the 5th International Workshop on Mobile Multimedia Communications*, 1998.

15. C. R. Lin, "On-demand QoS routing in multihop mobile networks," in *Proc. of INFOCOM 2001*, Anchorage, AK, April 2001.

16. R. Mangharam, "HotSpot: Access Point Services," Technical Presentation, July 24, 2002.

17. Mesh Networks, Inc. http://www.meshnetworks.com/

18. Nokia Rooftop Solution. http://www.wbs.nokia.com/solution/

19. C. Parris, H. Zhang, D. Ferrari, "Dynamic Management of Guaranteed Performance Multimedia Connections," *ACM Springer-Verlag Multimedia Systems Journal*, Vol. 1, No. 6, 1994.

20. C. E. Perkins and P. Bhagwat, "Highly dynamic destination sequenced distance vector routing (DSDV) for mobile computers," in *Proc. of the SIGGCOMM'94 Conference on Communication Architectures, Protocols and Applications*, August 1994.

21. C. E. Perkins, E. M. Royer, and S. R. Das, "Ad-hoc On-Demand Distance Vector (AODV) Routing," IETF draft. draft-ietf-manet-aodv-06.txt.

22. E. M. Royer, C. Perkins, S. R. Das, "Quality of Service for Ad Hoc On-Demand Distance Vector Routing," IETF Draft, draft-ietf-manet-aodvqos-00.txt, Work in Progress, July 2000.

23. L. A. Rusch, "Indoor Wireless Communications: Capacity and Coexistence on the Unlicensed Bands," *Intel Technology Journal*, Q3 2001,

24. C. Zhu and M. S. Corson, "QoS routing for mobile ad hoc networks," in *Proc. of the Twenty First International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, NY, June 23-27, 2002.

25. ZigBee Alliance. http://www.zigbee.com/

AUTHORS' BIOGRAPHIES

Lakshman Krishnamurthy is a Senior Staff Engineer in the Network Architecture Lab, Intel R&D. His current research interests include high-speed wireless networking, wireless sensor networks, and software-defined radios. He received B.E. and Ph.D. degrees from the University of Mysore, India and the University of Kentucky, respectively. His e-mail is Lakshman.Krishnamurthy@intel.com.

Steven Conner is a Senior Network Software Engineer in the Network Architecture Lab, Intel R&D. His current research interests include wireless networking and network self-configuration. He received B.S. and M.S. degrees from the University of Arizona. His e-mail is w.steven.conner@intel.com.

Mark Yarvis is a Senior Researcher in the Network Architecture Lab, Intel R&D. His current research interests include techniques for leveraging network heterogeneity, sensor networks, and pervasive computing. He received B.S., M.S., and Ph.D. degrees from the University of California, Los Angeles. His e-mail is Mark.D.Yarvis@intel.com.

Jasmeet Chhabra is a Senior Network Software Engineer in the Network Architecture Lab, Intel R&D. His current research interests include wireless sensor networks and software-defined radios. He received a B.E. and M.S. degree from the University of Delhi, India and the University of Maryland, College Park, respectively. His e-mail is Jasmeet.Chhabra@intel.com.

Carl M. Ellison is a Senior Security Architect in the Network Architecture Lab of the Corporate Technology Group of Intel Corporation. His current research is devoted to delegatable, distributed, public-key authorization with a special emphasis on self-organizing networks. His concentration on security has been a side effect of a more general career focus on distributed and fault tolerant systems. His e-mail is cme@jf.intel.com.

Chuck Brabenac is a senior architect in the Communications and Interconnect Technology Lab, Intel R&D. He is currently involved in ultra wideband (UWB) wireless communications research, with focus on its associated MAC and application stacks. He is also active in UWB-related standards activity in IEEE 802.15.3a, where he will be serving as vice-chair of this task group that is chartered to develop a high data rate WPAN standard. His e-mail is chuck.brabenac@intel.com.

Ernest Tsui is a principal engineer and manager of the Wireless Architecture Research group, Communications and Interconnect Technology Lab, Intel R&D. His interests at Intel are in wireless (and wireline) communications algorithms and architectures. He is owner of the reconfigurable baseband PHY architecture for CTG. He received his B.S., M.S., and Ph.D. degrees from U.C. Berkeley and is a Senior Member of the IEEE. His e-mail is ernest.tsui@intel.com.

Internetworking Using IPv6 Technology Inside and Outside the Home

Venkat R Gokulrangan, Desktop Products Group, Intel Corporation

Index words: IPv6, IPv4, UPnP, Home Networking, NAT, Broadband

ABSTRACT

The lack of sufficient globally routable Internet Protocol (IP) addresses to be assigned to subscribers has led service providers to deploy techniques such as Network Address Translation (NAT). However, NAT breaks end-to-end connectivity in several applications resulting in a poor end-user experience. Internet Protocol version 6 (IPv6) is the next-generation network layer technology designed to provide globally unique IP addresses to every endpoint in the Internet for years to come, thereby resolving the address depletion problem. IPv6 technology has several features such as auto-configuration, security, and mobility, built into its design, thereby creating an easy, rich, plug and play networking experience for the end user.

The key to IPv6 deployment is a smooth transition from current IP version 4 (IPv4)-based networks to new IPv6-based networks. It is generally accepted that the transition will span several years, marked by the emergence of isolated IPv6 islands in the customer premises, co-existing with IPv4. The inter-island communication will be carried over tunnels created over existing IPv4 networks. Eventually, as the core Internet progressively deploys native IPv6 networks, the tunnels will be removed, leading to the completion of the transition.

This paper discusses the core architecture of IPv6, its key features related to end-user experience, and the common IPv4 to IPv6 transition models. We also present a simple solution implemented on an Intel® XScale™ core-based platform, which illustrates the easy adoption of IPv6 in a home network. Finally, we discuss the evolving support for IPv6 in the Universal Plug and Play (UPnP*) forum that facilitates the easy deployment of IPv6 in a home network.

INTRODUCTION

With the rapid growth of endpoints requiring Internet access across the globe, assigning global Internet Protocol (IP) addresses to the connecting endpoints is becoming an increasing problem for Service Providers (SP) due to the paucity of available global IPv4 addresses. The effect is compounded with new devices such as cellular and wireless hosts being enabled to access Internet content. With the available pool of global IPv4 addresses predicted to be exhausted in the near future [1] and in order to meet the growing demands of subscribers, SPs are starting to deploy techniques such as Network Address Translation (NAT) (RFC2663). While the deployment of NAT has not prevented endpoints from seamlessly using common Internet applications such as the World Wide Web, e-mail etc., it has resulted in the breaking of several existing peer-to-peer applications, often resulting in a poor end-user experience.

With broadband Internet deployment on the rise, home networks that connect multiple PCs and consumer Internet appliances that need global Internet connectivity are emerging. Normally, Residential Gateways (RG) are used to multiplex the global Internet connection by assigning private IP addresses to the devices needing Internet access. While the private addresses assigned to these devices are sufficient for in-home networking, they cannot be used by applications outside the home that need access to those devices and their services, due to the private nature of the addresses.

The above issues result in several customer support calls that place an undue burden on the SP's support network.

Another challenge in a home network is to make it easy for the average user to add, install, and configure new Internet appliances. The networking layer is key to facilitating this "plug and play" usage model. The absence of mandated support for auto-address

---

™Intel XScale is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

* Other brands and names are the property of their respective owners.

configuration in IPv4 technology remains a key barrier for plug and play networking, despite the emergence of temporary solutions [2].

IPv6 is the next-generation network layer technology that attempts to solve the problems outlined above. With its 128-bit address width, IPv6 provides a very large address space, enough to provide global Internet addresses to endpoints for years to come. Moreover, it makes routing more efficient by virtue of its hierarchical address architecture. This solves the critical problems of address paucity and ubiquitous access to devices and services from anywhere, anytime. With its built-in design and support for auto-configuration, security, and mobility, IPv6 facilitates easy-to-use, end-to-end secure networking.

While the benefits of migrating to IPv6 networks are clear, the transition itself cannot happen instantaneously due to the involvement of several network elements. The transition from IPv4 to IPv6 networks will span the next several years, during which time existing IPv4 networks will be used to transport IPv6 packets. It is expected that IPv6 technology will initially emerge in customer premises as islands. As such, it will be relatively easy to deploy IPv6-based networks that co-exist with IPv4-based networks both inside and outside the home without affecting existing operations.

In this paper, we first present the problems and pitfalls of deploying NAT in an IPv4 environment and propose IPv6 as the remedy at the network layer. We follow this with a brief discussion of the IPv6 core addressing architecture and its key features. We also discuss the migration from IPv4 to IPv6 networks—the transition architecture and the models that facilitate a smooth transition. Then, we discuss a simple solution that was implemented on an Intel® XScale™ core-based platform, to deploy IPv6 in home networks that co-existed with IPv4 networks. Finally, we discuss the current deployment status of IPv6 and related issues.

IN-HOME NETWORKING

Endpoint Addressing and IP Multiplexing

As the number of homes with multiple PCs increases, the necessity for those PCs to share the Internet connection becomes obvious. With Broadband Internet access becoming more common across the globe, the emergence of special-purpose Internet appliances and digital media devices also necessitates the sharing of the Internet access as the Internet becomes the means to deliver rich, dynamic multi-media content. Figure 1 shows a home network with two PCs, an Internet appliance, and a printer, all sharing a single Internet connection.

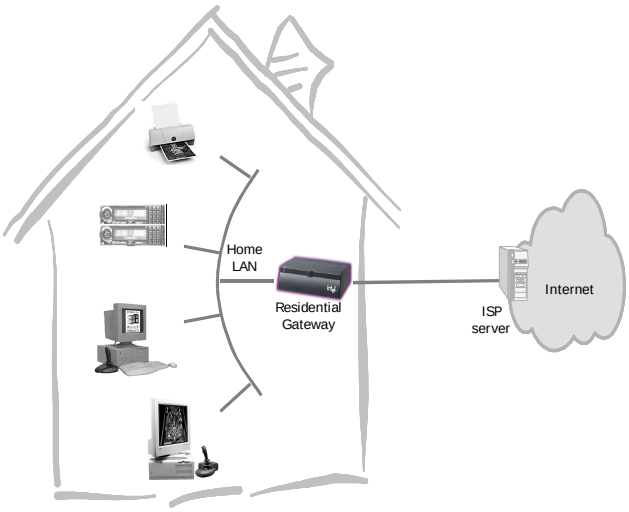

Figure 1: Two PCs, an Internet appliance, and a printer sharing the Internet connection

Since service providers charge more to provision global IPv4 addresses to devices inside the home, consumers often buy Residential Gateways (RGs) to multiplex the single Internet connection they have in their home. RGs assign private addresses to the various endpoints within the home and translate the private addresses to the global address as the packets fly through them. However, this fragments the Internet-enabled endpoints into private address space and public address space.

While such a Network Address Translation-based (NAT) scheme is relatively transparent to the end user, NATs with private addresses do break certain operational semantics that would have been preserved if the endpoints had obtained global IP addresses. Of particular interest are applications that communicate point-to-point by exchanging their IP addresses and port numbers in the IP datagram itself. As the exchanged addresses are private ones, the packet communication between the endpoints breaks down, due to NAT, since private addresses cannot be routed in the Internet. In general, any communication that takes place by the exchange of host addresses may be broken when a NAT is deployed in the path of the communication.

RGs use corrective measures such as Application Layer Gateways (ALGs) to make the private addresses and ports embedded in the packet payload appear to have emanated from the Internet connection of the RG. Since ALGs need to know the application protocol to fix the payload, it is difficult to dynamically create ALGs for new network applications that use unknown protocols.

™Intel XScale is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

NAT also prohibits standard Internet services from being hosted on more than one endpoint. Since services are often hosted on a combination of IP addresses and ports, simultaneous usage of a port by two applications or hosts is not possible as they share the Internet address. For example, it is not possible to host more than one game server when the game uses a fixed port. While specialized applications can be used to configure RGs to allow service forwarding, it is beyond the average end-user to effectively use such customized applications.

Therefore, even though it is evident that addresses for multiple endpoints inside the home are required, NAT-based private IP addressing is not the correct solution for the following reasons:

- Several peer-to-peer applications that exchange their IP address and port will not work.
- ALGs alleviate the problem by translating the payload of well-known protocols; however, they cannot be extended to new unknown application protocols.
- Hosting multiple identical services behind a NAT is not possible.
- Service forwarding requires knowledge and skills beyond the average user.

While the obvious solution is to assign global IPv4 addresses to every endpoint, the cost of IP addresses in conjunction with the very limited pool of available IPv4 addresses, where cost is not a concern, prevents this from being the solution.

IPv4 Protocol Issues

The lack of global IPv4 addresses now is often cited as the result of the less-efficient design of the IPv4 addressing architecture. In particular, the "class A" IPv4 addresses, which constitute half of the address space, were given to a small number of organizations. This resulted in an uneven distribution of addresses and underutilization of the overall IPv4 address space.

Further, the lack of inherent support for secure end-to-end communication and the sub-optimal support for mobility make IPv4 less suitable for modern communication patterns. These features were added later as quick-fix solutions. The support for implementing Quality of Service (QoS) was very coarse-grained and for all intents and purposes it resulted in the underutilization of the feature. Finally, communication anonymity could not be added into the protocol; instead it was added at a higher layer such as application tunnels.

IPV6 TECHNOLOGY

The next-generation network layer protocol, IP version 6 (IPv6) provides a remedy for the problems found in IPv4 as follows:

Global Internet Address: By provisioning a very large address space with 128 bits for every IP address, IPv6 can assign global addresses to every Internet endpoint for several decades.

Plug and Play Networking: By connecting to a network, an IPv6-enabled endpoint automatically acquires IPv6 addresses. Multiple addresses could be assigned to an interface in different realms—local addresses and global addresses. The scheme has also been made flexible enough to re-address endpoints quickly if necessary.

Better Quality of Service (QoS) Support: IPv6 provides better support for fine-grained QoS. This facilitates better delivery of multi-media data.

Mobility, Anonymity, and Security by packet encryption (IPSec) have also been built into the protocol itself. It has been mandated that security must be supported by every IPv6 implementation.

IPv6 Core Architecture

The IPv6 provisions (RFC2460) a small protocol header where the essential information is stored and allows dynamic extensions to it when necessary. This is shown in Figure 2 below.

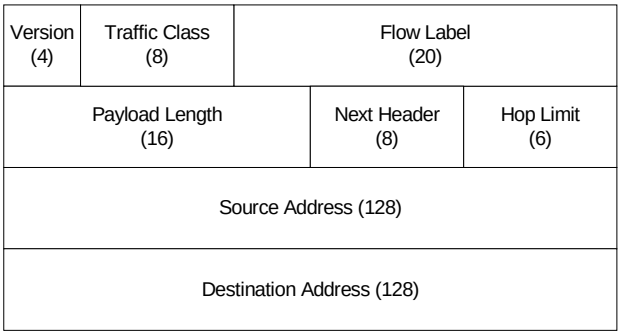

Figure 2: IPv6 core address architecture

The core header looks similar to that of IPv4 with the exception that the number of fields is reduced. Except for the flow label field, all fields functionally correspond to similar fields in IPv4. Additional header information can be dynamically added as extension headers by using an indirection field (Next Header) in the core header that points to a chain of extension headers.

IPv6 technology classifies every address into one of three types: Unicast, Multicast, and Anycast.

Unicast and Multicast addresses are comparable to their IPv4 counterparts. A Unicast IPv6 address is the address

assigned to an interface in a host. It differs from IPv4 in that an interface can have one or more IPv6 Unicast addresses assigned to it.

Multicast addresses are usually assigned to a set of interfaces, usually a set of hosts. Every packet destined for the Multicast address is received by every node/interface in the set. The delivery semantics are similar to that of IPv4.

The new class of addresses in IPv6, the Anycast address, is an address that is assigned to a set of interfaces/hosts. However, a packet destined for the Anycast address is delivered to at least one of the interface hosts as defined by a set of criteria (usually governed by the underlying routing protocol). We discuss the global Unicast address in detail as it is the most relevant address for an interoperable home.

IPv6 Global Unicast Address

The structure of a global IPv6 Unicast address is shown in Figure 3. The 128-bit native IPv6 address is hierarchically partitioned as follows:

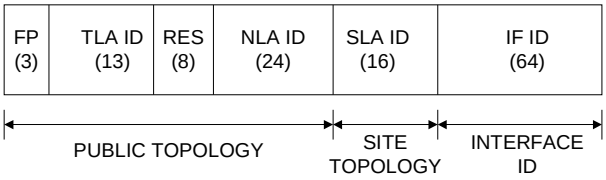

Figure 3: IPv6 Unicast address structure

FP – 3-bit Format Prefix field. This indicates the type of IPv6 address. At the time of this writing, all native global IPv6 Unicast addresses will have a value of 001.

Top-Level Aggregator (TLA). This is a 13-bit identifier used to denote the infrastructure provider.

RES. This is an 8-bit field reserved for IANA usage.

Next-Level Aggregator (NLA). This is used to denote the different entities the infrastructure provider services.

Site-Level Aggregator (SLA). This is used to identify one of the different partitions in a given site.

INTERFACE-ID This is the interface identifier used to identify a unique host. This is referred to as the Extended Unique Identifier (EUI), and normally an extended form of the MAC address is used.

In general, the 64 bits to the left are used to identify the logical subnet to which the interface belongs, while the 64 bits to the right are used to identify the interface in the subnet. Often the subnets are represented in a CIDR-like notation (RFC1519).

The key features of the IPv6 architecture are the large address space and the routing efficiency resulting from the hierarchical organization.

IPv6 Address Auto-Configuration

One of the mandated requirements of the IPv6 protocol is that it be able to automatically acquire an IPv6 address to instantly communicate in the network neighborhood. This is facilitated by making every host implement an interface with the special local IPv6 address as described in RFC2373. These are referred to as link-local addresses and can be used to communicate only in the local subnet. Additionally, a host can configure itself with addresses using any route advertisement appearing in the subnet to which it is connected. Usually, IPv6-enabled routers send such advertisements so that hosts can self-configure.

Transition Address Architecture

For existing IPv4 endpoints running an IPv6 protocol stack, several compatibility and transition architecture addresses have been provisioned. The important transition architecture address involves assigning the special TLA-ID of 2. Implied in almost every global Unicast IPv6 address that has this special TLA-ID is an IPv4 address in the next 32 bits (bits 16-47) and an arbitrary number in its SLA to form a unique subnet prefix **2002:<IPV4>::/48**. This prefix and the EUI combine to form a unique global IPv6 address that enables an IPv4 host to communicate via the Internet.

These are referred to as *6to4 addresses* (Figure 4), and the hosts that acquire 6to4 addresses based on their global IPv4 address are referred to as 6to4 hosts. Normally, such hosts require the implementation of both IPv4 and IPv6 networking protocol stack support in their operating system and are called *dual-stacked hosts*.

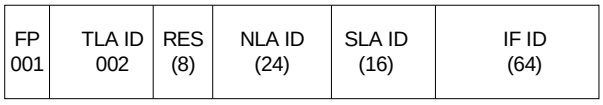

Figure 4: IPv6 6to4 address

THE APPROACH TO TRANSITIONING TO IPV6

The necessity to adopt and natively support IPv6 across the entire Internet involves the following:

- Every host that accesses the Internet needs to support IPv6 at every layer.
- The edge network, a collective term for both the home network and the provider network, needs to support IPv6, i.e., both intra-home and home-to-provider communication must support IPv6.

- The core backbone of the Internet connected to the provider network must support IPv6.

As you can see, the transition to an IPv6 network involves a great deal of effort in terms of cost, time, and planning. The migration to IPv6 will happen incrementally only over a period of time; it won't happen with a simple flip of a switch. During the transition, it is expected that all IPv6 communication will be transported as payload in existing IPv4 networks. This technique is called *tunneling,* or *tunneling IPv6 over IPv4*, as shown in Figure 5.

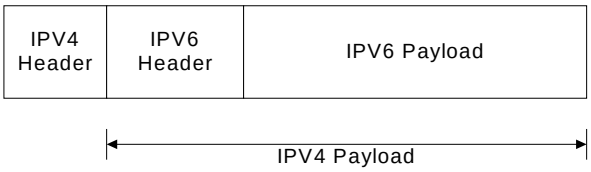

Figure 5: IPv6 packet tunneled in an IPv4 packet

We expect the transition to evolve and emerge initially in the home network where IPv6 deployment is relatively simple. This will result in the creation of islands of IPv6-enabled homes that are connected to the Internet over IPv4 networks. During this period, applications will be IPv6-enabled but will preserve their IPv4 connectivity for maximum Inter-networking. We also anticipate that host applications will automatically initiate IPv6 connectivity and use it if available. Otherwise, IPv4 will be used for connectivity. Such IPv6 islands will expect minimal assistance from the service providers to resolve host names to IPv6 addresses. We also expect communication between the IPv6 islands to happen over the IPv4 Internet by the use of tunneling.

As the transition progresses and more Internet content is delivered using IPv6 as the preferred protocol for connectivity, networked homes are expected to use IPv6 on a regular basis. Increased usage and transport of content using IPv6, although over IPv4 networks, coupled with the reduced support burden, will serve as incentives for the service providers to deploy IPv6 natively in their networks. Similarly, as more traffic gets transported over IPv6, the core network providers will start deploying IPv6 in their networks alongside IPv4 networks resulting in a dual-stacked core network.

The above two actions would help establish IPv6 as the mainstream network protocol, while IPv4 usage diminishes. It is important to note that, for IPv6 to become the mainstream networking protocol, host operating systems need to incorporate dual networking stacks, and most network applications need to be modified to make use of IPv6 connectivity if present.

IPV6 TRANSITION ARCHITECTURE FOR HOME NETWORKS

In general, the transition architecture is classified into two broad categories: tunneling and translation.

The tunneling models transport IPv6 datagrams over IPv4 as the protocol payload. The translation models involve translating the IPv4 headers into IPv6 headers and vice versa as the packets fly through a translation device.

During the initial stages of the transition, we expect the tunneling models to prevail over the translation models due to a lack of standards and the problems associated with translation. However, it is to be noted that wherever IPv6 support is absent or when heterogeneous communication has to happen, such as an IPv4-only host to an IPv6-only host, the network has to deploy translation techniques such as NAT-PT(RFC2766).

6to4 Tunneling

Several tunneling techniques have been proposed to transport IPv6 over IPv4 [1]. All the tunneling schemes work as long as the applications are IPv6-enabled and communicate end-to-end using IPv6. To support the applications, an endpoint must implement IPv4/IPv6 dual stack to make IPv4 the transport layer for the IPv6 communication.

Of the several tunneling techniques, the important ones are the (manually) configured tunnels and automatic tunnels. Automatic tunnels are the preferred choice as they leverage existing global IPv4 addresses to set up the tunnel. Of particular interest are the 6to4 automatic tunnels (RFC3056). Since the 6to4 tunneling scheme uses an existing IPv4 address to create IPv6 subnets, it is an important element of the transition architecture. In fact, the transition model that is important for IPv6 deployment during the transition in a home uses 6to4 transition techniques. We briefly discuss the 6to4 automatic tunneling scheme below and follow it up with our implementation.

The 6to4 model uses the ISP-assigned global IPv4 address (32 bit) to obtain the special prefix {**2002:IPv4 address}** and uses this global 48-bit prefix to advertise to all the devices in the home network. Any IPv6 device inside the home automatically acquires a global 6to4 IPv6 address, thus becoming a 6to4 host.

The communication between two 6to4 hosts in the Internet is relatively straightforward as each host can tunnel the IPv6 packets using their respective IPv4 connection. However, the communication between a 6to4 host and a native IPv6 host travels through a tunnel to a well-known router, called the *relay router*. The relay router in turn routes the IPv6 packet in the native IPv6 cloud. Usually

the relay routers are assigned a fixed Anycast address so hosts can route the native IPv6 payload to them over the IPv4 network. This is illustrated in Figure 6 below: In this figure, the RG itself is the 6to4 host communicating with other 6to4 hosts.

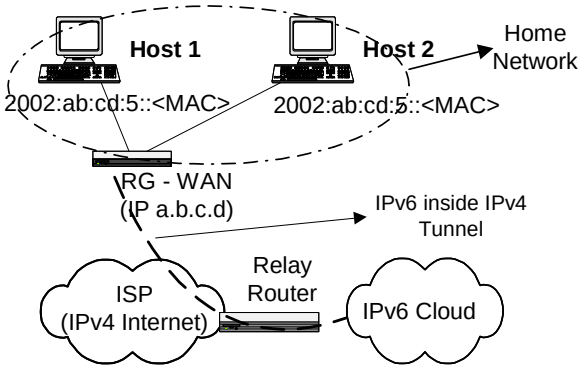

Figure 6: 6to4 tunneling

The communication back from the native IPv6 cloud to the 6to4 host is routed by the nearest dual-stacked router and transported over an IPv4 network back to the host and not necessarily via the relay router.

Translation Model

In the translation model, the IPv4 packets are translated into IPv6 packets (and vice versa) using specialized translation schemes [1]. Usually such translations are more specific to the protocol headers and are performed by network elements dedicated to such purposes. Thus, these translation schemes are transparent to network applications and usually do not require the application to change. Since translation schemes rewrite headers, they are susceptible to the same problems as ALGs. The disadvantages of the translation schemes outweigh the advantages, and we do not use them unless there is a compelling necessity to do so.

Sometimes, the translation schemes have to be used in conjunction with the tunneling schemes such as in the case of some legacy hosts that do not support dual IP stacks.

Requirements for Home Network Transition

The key to successful implementation of the above transition architecture is as follows:

- Hosts should be dual-stack-enabled to transport IPv6 packets over their IPv4 layer.
- ISPs should have the ability to support at least one global IPv4 address to the home.
- An always-on device, usually Residential Gateways (RGs), should be present to administer the model in the home network without additional intervention.
- The endpoints should initiate IPv6 tunnels such as Teredo [4] when an IPv6 supporting RG is not present or an ISP cannot provision a global IPv4 address.

OUR IMPLEMENTATION OF THE IPV6 TRANSITION

Internet appliances such as Residential Gateways (RGs) are important for the IPv6 transition inside the home as they always remain powered on and can make the route advertisement service available to the home network. Additionally, the RG is responsible for the following:

- It must automatically convert the global IPv4 address assigned by the ISP to an IPv6 global prefix and announce it in the home network.
- It must take the native IPv6 packets emanating in the home and tunnel them to the relay router or route them natively in the Internet.
- When heterogeneous communication happens (IPv4 to IPv6 networks or vice versa), the RG must translate the packets, if necessary.

Due to this important role of RGs, we decided to use an Intel® XScale™ core-based reference platform (IOP80310) and make it an RG with IPv6 routing capability. This platform was chosen because of its suitability in an embedded environment. The platform has an on-board Ethernet port and two expansion slots. We used the on-board Ethernet port to connect to the home network and one of the expansion slots to connect to the Broadband Internet.

Due to its immediate availability for the IOP80310 platform with source code, we chose Linux* as the operating system to implement our solution in this platform.

We used the IPv6 code base from the USAGI project [3] and ported it to the Intel XScale core-based platform. The resultant IPv6 Linux kernel was used along with some software tools, which were required to configure static IPv6 routes and default rules for routing in the kernel

™ Intel XScale is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

* Other brands and names are the property of their respective owners.

routing table. We also ported the route advertisement daemon to the IOP80310 platform to announce the 6to4 global prefix in the home network.

Home Network Setup

We used the IOP80310 reference platform as the residential gateway (RG). The on-board Ethernet port was connected to a four-port hub, so the home network computers and devices could connect to the RG. We used an Ethernet card on one of the expansion slots to connect and simulate a broadband Wide Area Network (WAN) connection.

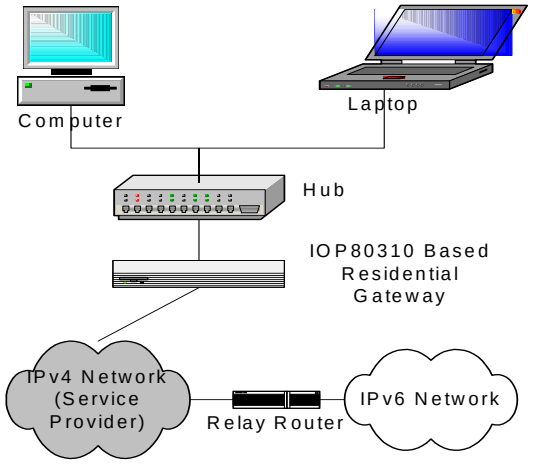

Figure 7: Home network setup of our implementation

We used a couple of computers running Windows XP* and Windows 2000* operating systems to connect to the hub and create a home network environment. Both computers were configured to support dual IPv4/IPv6 stack. This setup is illustrated in Figure 7.

Software Setup

We used the Linux kernel, ported to the IOP80310 platform and configured it to provision the functions usually found in RGs such as NAT etc. This enabled all clients to run all existing Internet applications by sharing the broadband connection. The RG was assigned a static IPv4 address on its WAN interface and connected to the Internet.

The IPv6 software stack was added as a loadable driver in the RG. The RG was configured to create the 6to4 global prefix from its static IPv4 address and announce it as the global subnet prefix in the home network with the RG as the router for both IPv4 and IPv6 packets.

Additionally, the RG's initialization software created a IPv4-based tunnel to the standard relay router so that all non-local IPv6 traffic can be routed to the relay router. Entries in the routing table were created to provision this. Additionally, specific routing table entries were also created so that all IPv6 traffic destined for the home network clients coming from the Internet are properly routed by the RG's IPv6 stack back to the home network. Our service provider had provisioned a DNS server that supported name resolution to both IPv4 and IPv6 records so applications could take advantage of it.

End-to-End IPv6 Connectivity

To test end-to-end IPv6 connectivity and internetworking over the (IPv4) Internet, we used an IPv6-enabled version of the popular PC game Quake* and used it to connect to an IPv6-enabled Quake server over the Internet. Even though the client and the server were communicating over IPv6, the packets were actually transported over an existing IPv4 network.

We were able to successfully connect to the Quake server and execute the game validating our end-to-end internetworking vision.

THE CHALLENGES OF TRANSITIONING TO IPV6 IN THE HOME

From an implementation standpoint, the key challenge faced, in our experience, was the inability to address endpoints using the Fully Qualified Domain Names (FQDN) of the endpoints. There were two reasons attributed to this: a) a lack of standards to dynamically discover the IPv6 Domain Name Server (DNS); and b) the inability to dynamically register the name with the DNS server so peers could communicate end-to-end without having to know the IPv6 address. Standards are still evolving in the Internet Engineering Task Force (IETF) to solve this barrier to IPv6 deployment.

The other key challenge we faced is very specific to the transition mode. By using 6to4, we compromised one of the distinguishing features of IPv6 over IPv4 i.e., communication anonymity. By virtue of using the global IPv4 address as part of the prefix, we simply ended up with the same problem existing in IPv4, i.e., the traffic from an endpoint could be tracked based on the IP address.

From a deployment and ease-of-use perspective, RGs play a crucial role as they function at the junction of both the home and the service provider networks. As RGs hold the key to enable IPv6 services inside the home, the clear challenge is to be able to configure them to adopt the transition model in the provider network. In the following

section, we discuss the role of Universal Plug and Play (UPnP*) technology in enabling IPv6 in-home networks.

IPV6 AND UNIVERSAL PLUG AND PLAY

Universal Plug and Play (UPnP*) technology provides a control protocol to easily install, configure and control devices and appliances used in small and home networks. In a home network with different kinds of digital devices present, UPnP technology holds the key to facilitating ease-of-use leading to rich end-user experience. In this section, we discuss the design aspects that influence the easy deployment of IPv6 using UPnP technology. To use IPv6 in conjunction with UPnP technology, there are two aspects that have to be considered.

1. UPnP technology currently uses IPv4 as the underlying network layer for all communication. To use IPv6 as the network layer, extensions and modifications to the UPnP protocol and specifications are needed. Currently, this is being reviewed by the UPnP Forum.
2. Residential Gateways (RGs) fall under the category of Internet Gateway Devices (IGDs). In order for IGDs to support IPv6, and administer IPv6-related functions used in the home network, standardized means of configuring the IGDs using the UPnP protocol are needed.

To successfully enable IPv6 addresses for clients inside the home, the IGD should have the ability to be configured using the UPnP protocol to support the following:

- A default configuration of the UPnP IGD should automatically detect all the IPv6 services present in the attached service provider network and relay those services locally to make them available to all the devices in the home network. In the absence of IPv6 services in the provider network, an RG should have the ability to be configured to support one of the transition models, with the preferred automatic tunneling model being 6to4. The RG should supplementally act as an IPv6 router so that IPv6 packets from/to the home are properly handled as they flow to and from the provider network.
- Endpoints such as PCs should detect the presence of the IPv6 capabilities of an UPnP IGD and use them. In the absence of IPv6 support in the network, the PCs should initiate techniques such as Teredo to enable their IPv6 stack.

* Other brands and names are the property of their respective owners.

THE STATUS OF IPV6 DEPLOYMENT IN THE HOME

At present, most modern operating systems including Windows* XP*, and Windows 2000* (clients and servers), all flavors of Unix* including Linux*, and several embedded operating systems support dual-stacked networking protocol stacks. While the Windows operating systems support them, they are not enabled for usage automatically. Integrated support for IPv6 in the operating system has been included in the test releases of future versions of the Windows* operating system such as Windows XP-SP1*.

Residential Gateway (RG) vendors are starting to show an interest in incorporating dual-stack operating systems in their access devices. From the perspective of UPnP technology, several vendors including Intel are participating in the UPnP forum to standardize IPv6 implementation in RGs. Initial Intel participation has resulted in a set of guidelines for IPv6 implementation in RGs. Standardization in the UPnP forum will further solidify support for IPv6 in the RG community.

Core network equipment vendors such as Hitachi have started to add IPv6 support in the operating platforms and environments of their infrastructure equipment, such as routers [1].

CONCLUSION

In a broadband-enabled digital home, several Internet appliances and digital devices are expected to be present. IPv4 Network Address Translations (NATs) would only result in limited end-user experiences when using those devices. Current solutions based on IPv4 are mere patch-work solutions. For easy networking inside and outside the home, IPv6 is emerging as the preferred next-generation protocol for communication in the Internet. The design of the entire gamut of Intel's products that participate in the digital home including PCs, Internet appliances, and digital home products should be enabled to use IPv6 as the protocol of choice.

ACKNOWLEDGMENTS

I thank my colleagues in the Corporate Technology Group, Intel, for exchanging and sharing ideas, views, and opinions while writing this paper. In particular, I thank Vijay Rao for supporting the IPv6 work on Intel® XScale™

* Other brands and names are the property of their respective owners.

™ Intel XScale is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

core-based platforms. Sincere thanks go to Christopher Lord and Ulhas Warrier for several fruitful discussions that laid the foundations for the arguments presented here. Many thanks go to Yasser Rasheed and Charlie Tai for their encouragement.

REFERENCES

[1] D.G. Waddington and F Chang, "Realizing the transition to IPv6," *IEEE Communications Magazine,* June 2002, pp. 139-144.

[2] Stuart Cheshire and Bernard Aboba, "Dynamic Configuration of IPv4 Link-local Addresses," *IETF Draft*, March 2001.

[3] USAGI Project. *http://www.linux-ipv6.org/*

[4] C Huitema, "Teredo: Tunneling IPv6 over UDP through NATs," *IETF Draft,* Sep 2002.

AUTHOR'S BIOGRAPHY

Venkat R Gokulrangan is a Senior Software Engineer in the Desktop Platform Group. While at Intel Labs, he initiated the work on IPv6 for Residential Gateways using Intel platforms. Venkat actively participates in standards activities for easy IPv6 adoption in residential networks. His work resulted in proposing standards to the IGD subcommittee in the UPnP forum. Venkat has designed and implemented data subsystems such as 802.1D bridges and voice subsystems such as VoDSL software. He has an M.S. degree in Computer Science and Engineering from the University of Michigan, Ann Arbor. His e-mail is venkat.r.gokulrangan@intel.com

For further information visit:

developer.intel.com/technology/itj/index.htm