



Intel[®] Technology Journal

Interoperable Home Infrastructure

High-Quality Media Distribution in the Digital Home

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.

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

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 widespread 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)

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.

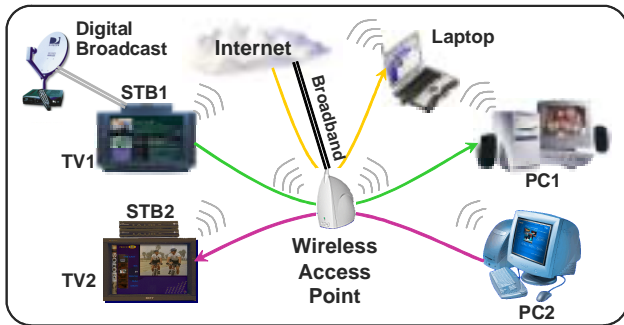


Figure 1: Digital Home environment

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.

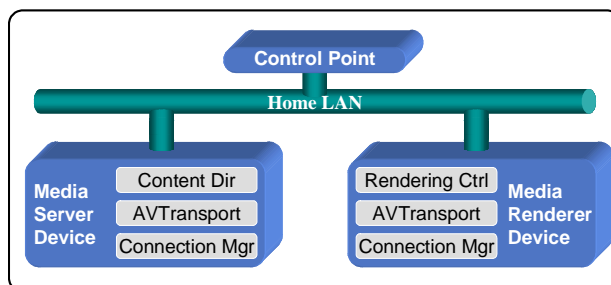


Figure 2: UPnP AV architecture

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.

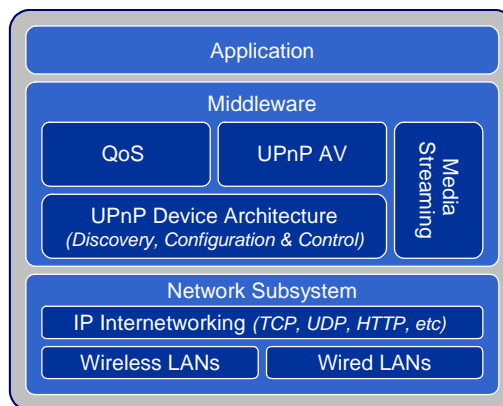


Figure 3: High-quality media distribution architecture

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, phonenumber, 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.

Phonenumber Networking

The Home Phonenumber 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

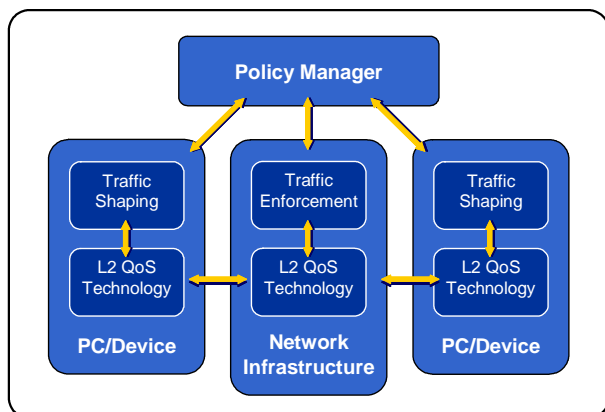


Figure 4: System overview for QoS networking

support, network-layer QoS, and policy-based network management.

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.

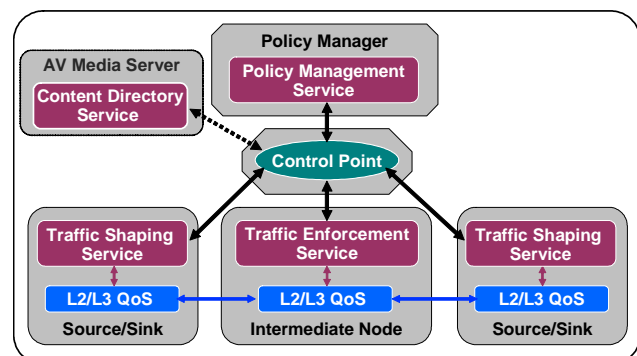


Figure 5: QoS framework based on UPnP technology

* 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*

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

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.

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/standardizeddcp/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

Copyright © Intel Corporation 2002. This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://www.intel.com/sites/corporate/tradmarx.htm>.

For further information visit:

developer.intel.com/technology/itj/index.htm