



UPnP™ Device Architecture 1.1

Document Revision Date: October 15, 2008

© 2008 Contributing Members of the UPnP Forum. All rights reserved. See <http://www.upnp.org/info/copyright.asp> for more information.

Authors*	Company
Alan Presser	AllegroSoft
Lee Farrell	Canon
Devon Kemp	Canon
William Lupton	Conexant
Shinichi Tsuruyama	Epson
Shivaun Albright	HP
Andrew Donoho	IBM
John Ritchie	Intel
Bryan Roe	Intel
Mark Walker	Intel
Toby Nixon	Microsoft
Colleen Evans	Microsoft
Henry Rawas	Microsoft
Trevor Freeman	Microsoft
Joonyoung Park	Motorola
Cathy Chan	Nokia
Franklin Reynolds	Nokia
Jose Costa-Requena	Nokia
Yinghua Ye	Nokia
Tom McGee	Philips
Geert Knapen	Philips
Maarten Bodlaender	Philips
Jarno Guidi	Philips
Lex Heerink	Philips
John Gildred	Pioneer
Alan Messer	Samsung
YoonSoo Kim	Samsung
Markus Wischy	Siemens
Andrew Fiddian-Green	Siemens
Bruce Fairman	Sony
Jonathan Tourzan	Sony
John Fuller	Sony

*Note: The UPnP Forum in no way guarantees the accuracy or completeness of this author list and in no way implies any rights for or support from those members listed. This list is not the specifications' contributor list that is kept on the UPnP Forum's website.

Table of Contents

List of Tables.....	v
List of Figures.....	vi
i Introduction	1
i.1 What is UPnP™ Technology?	1
i.2 UPnP™ Forum	1
i.3 In this document	2
i.4 Audience.....	5
i.5 Conformance terminology	5
i.6 Acronyms	6
i.7 Glossary.....	6
i.8 References and resources.....	7
0 Addressing	8
0.1 Determining whether to use Auto-IP	8
0.2 Choosing an address	8
0.3 Testing the address	9
0.4 Forwarding rules.....	10
0.5 Periodic checking for dynamic address availability.....	10
0.6 Device naming and DNS interaction	10
0.7 Name to IP address resolution	11
0.8 References	11
1 Discovery.....	12
1.1 SSDP message format	15
1.1.1 SSDP Start-line	16
1.1.2 SSDP message header fields	16
1.1.3 SSDP header field extensions.....	16
1.1.4 UUID format and RECOMMENDED generation algorithms.....	17
1.1.5 SSDP processing rules.....	17
1.2 Advertisement	17
1.2.1 Advertisement protocols and standards	18
1.2.2 Device available - NOTIFY with ssdp:alive	19
1.2.3 Device unavailable -- NOTIFY with ssdp:byebye.....	25
1.2.4 Device Update - NOTIFY with ssdp:update	27
1.3 Search.....	29
1.3.1 Search protocols and standards	29
1.3.2 Search request with M-SEARCH	30
1.3.3 Search response	33
1.4 References	36
2 Description	37
2.1 Generic requirements on HTTP usage.....	40
2.2 Generic requirements on XML usage	43
2.3 Device description	43

2.4	UPnP Device Template	48
2.5	Service description.....	48
2.5.1	Defining and processing extended data types	55
2.5.2	String equivalents of extended data types	57
2.5.3	Generic requirements	58
2.5.4	Ordering of Elements	58
2.5.5	Versioning	59
2.6	UPnP Service Template.....	59
2.7	Non-standard vendor extensions and limitations.....	60
2.7.1	Placement of Additional Elements and Attributes	61
2.8	UPnP Device Schema.....	61
2.9	UPnP Service Schema	62
2.10	UPnP Datatype Schema.....	62
2.11	Retrieving a description using HTTP	62
2.12	References	65
3	Control	67
3.1	Control protocols.....	69
3.1.1	SOAP Profile	69
3.2	Actions	73
3.2.1	Action invocation.....	73
3.2.2	Action Response	76
3.2.3	UPnP Action Schema	79
3.2.4	Recommendations and additional requirements	79
3.2.5	Action error response.....	79
3.2.6	UPnP Error Schema.....	82
3.3	Query for variable.....	83
3.4	References	83
4	Eventing.....	84
4.1	Unicast eventing.....	85
4.1.1	Subscription	86
4.1.2	SUBSCRIBE with NT and CALLBACK	88
4.1.3	Renewing a subscription with SUBSCRIBE with SID.....	91
4.1.4	Canceling a subscription with UNSUBSCRIBE	93
4.2	Multicast Eventing.....	95
4.3	Event messages	96
4.3.1	Error Cases	97
4.3.2	Unicast eventing: Event messages: NOTIFY	97
4.3.3	Multicast Eventing: Event messages: NOTIFY	101
4.4	UPnP Event Schema	104
4.5	Augmenting the UPnP Device and Service Schemas	104
4.6	References	105
5	Presentation.....	106
5.1	References	107
Appendix A. IP Version 6 Support.....		109

A.1	Introduction	109
A.2	General Principles	109
A.2.1	Device operation	110
A.2.2	Control point operation	110
A.3	Addressing	110
A.3.1	Summary of boot/startup process	111
A.3.2	Short overview of protocol specified by RFC 2462	111
A.4	Discovery	112
A.4.1	Advertisement	113
A.4.2	Advertisement: Device unavailable	114
A.4.3	Advertisement: Device update	114
A.4.4	Search	114
A.4.5	Search response	115
A.5	Description	115
A.6	Control	115
A.7	Eventing	115
A.8	Presentation	116
A.9	References	116
Appendix B.	Schemas	117
B.1	UPnP Device Schema	117
B.2	UPnP Service Schema	122
B.3	UPnP Control Schema	126
B.4	UPnP Error Schema	127
B.5	UPnP Event Schema	128
B.6	Schema references	129

List of Tables

Table i-1: Acronyms	6
Table 1-1: Root device discovery messages	20
Table 1-2: Embedded device discovery messages	20
Table 1-3: Service discovery messages	20
Table 2-1: Vendor extensions	60
Table 3-1: SOAP 1.1 UPnP Profile	70
Table 3-2: <code>mustUnderstand</code> attribute	71
Table 3-3: UPnP Defined Action error codes	82
Table 4-4: HTTP Status Codes indicating a Subscription Error	91
Table 4-5: HTTP Status Codes indicating a Resubscription Error	93
Table 4-6: HTTP Status Codes indicating a Cancel Subscription Error	94
Table 4-7: HTTP Status Codes indicating a Notify Error	101
Table 4-7: Multicast event levels	103

List of Figures

Figure i-1: Protocol stack	2
Figure 1-1: Discovery architecture.....	13
Figure 1-2: Advertisement protocol stack	19
Figure 1-3: Initial and repeat announcements, no announcement spreading.....	21
Figure 1-4: Initial and repeat announcements, message spreading of repeat announcements	21
Figure 1-5: Search protocol stack.....	30
Figure 2-1: Description architecture	37
Figure 2-2: Description retrieval protocol stack.....	63
Figure 3-1: Control architecture	67
Figure 3-2: Control protocol stack.....	69
Figure 4-1: Unicast eventing architecture.....	85
Figure 4-2: Unicast eventing protocol stack.....	86
Figure 4-3: Multicast eventing architecture.....	95
Figure 4-4: Multicast eventing protocol stack.....	95
Figure 5-1: Presentation architecture.....	106
Figure 5-2: Presentation protocol stack	107

i Introduction

[Informative]

i.1 What is UPnP™¹ Technology?

UPnP™ technology defines an architecture for pervasive peer-to-peer network connectivity of intelligent appliances, wireless devices, and PCs of all form factors. It is designed to bring easy-to-use, flexible, standards-based connectivity to ad-hoc or unmanaged networks whether in the home, in a small business, public spaces, or attached to the Internet. UPnP technology provides a distributed, open networking architecture that leverages TCP/IP and Web technologies to enable seamless proximity networking in addition to control and data transfer among networked devices.

The UPnP Device Architecture (UDA) is more than just a simple extension of the plug and play peripheral model. It is designed to support zero-configuration, "invisible" networking, and automatic discovery for a breadth of device categories from a wide range of vendors. This means a device can dynamically join a network, obtain an IP address, convey its capabilities, and learn about the presence and capabilities of other devices. Finally, a device can leave a network smoothly and automatically without leaving any unwanted state behind.

The technologies leveraged in the UPnP architecture include Internet protocols such as IP, TCP, UDP, HTTP, and XML. Like the Internet, contracts are based on wire protocols that are declarative, expressed in XML, and communicated via HTTP. Using Internet protocols is a strong choice for UDA because of its proven ability to span different physical media, to enable real world multiple-vendor interoperation, and to achieve synergy with the Internet and many home and office intranets. The UPnP architecture has been explicitly designed to accommodate these environments. Further, via bridging, UDA accommodates media running non-IP protocols when cost, technology, or legacy prevents the media or devices attached to it from running IP.

What is "universal" about UPnP technology? No device drivers; common protocols are used instead. UPnP networking is media independent. UPnP devices can be implemented using any programming language, and on any operating system. The UPnP architecture does not specify or constrain the design of an API for applications; OS vendors may create APIs that suit their customers' needs.

i.2 UPnP™ Forum

The UPnP Forum is an industry initiative designed to enable easy and robust connectivity among stand-alone devices and PCs from many different vendors. The UPnP Forum seeks to develop standards for describing device protocols and XML-based device schemas for the purpose of enabling device-to-device interoperability in a scalable, networked environment.

The UPnP Implementers Corporation (UIC) is comprised of UPnP Forum member companies across many industries that promote the adoption of uniform technical device interconnectivity standards and testing and certifying of these devices. The UIC

¹ UPnP™ is a certification mark of the UPnP™ Implementers Corporation.

develops and administers the testing and certification process, administers the UPnP logo program, and provides information to UIC members and other interested parties regarding the certification of UPnP devices. The UPnP device certification process is open to any vendor who is a member of the UPnP Forum and UIC, has paid the UIC dues, and has devices that support UPnP functionality. For more information, see <http://www.upnp-ic.org>.

The UPnP Forum has set up working committees in specific areas of domain expertise. These working committees are charged with creating proposed device standards, building sample implementations, and building appropriate test suites. This document indicates specific technical decisions that are the purview of UPnP Forum working committees.

UPnP vendors can build compliant devices with confidence of interoperability and benefits of shared intellectual property and the logo program. Separate from the logo program, vendors may also build devices that adhere to the UPnP Device Architecture defined herein without a formal standards procedure. If vendors build non-standard devices, they determine technical decisions that would otherwise be determined by a UPnP Forum working committee.

i.3 In this document

The UPnP Device Architecture (formerly known as the DCP Framework) contained herein defines the protocols for communication between controllers, or *control points*, and devices. For discovery, description, control, eventing, and presentation, the UPnP Device Architecture uses the following protocol stack (the indicated colors and type styles are used throughout this document to indicate where each protocol element is defined):

Figure i-1: Protocol stack

<i>UPnP vendor [purple-italic]</i>			
<i>UPnP Forum [red-italic]</i>			
UPnP Device Architecture [green-bold]			
SSDP [blue]	Multicast events [navy-bold]	SOAP [blue]	GENA [navy-bold]
		HTTP [black]	HTTP [black]
UDP [black]		TCP [black]	
IP [black]			

At the highest layer, messages logically contain only UPnP vendor-specific information about their devices. Moving down the stack, vendor content is supplemented by information defined by UPnP Forum working committees. Messages from the layers above are hosted in UPnP-specific protocols such as the Simple Service Discovery Protocol (SSDP), the General Event Notification Architecture (GENA) and the multicast event protocol defined in this document, and others that are referenced. SSDP is delivered via either multicast or unicast UDP. Multicast events are delivered via multicast UDP. [GENA](#) is delivered via HTTP. Ultimately, all messages above are delivered over IP. The remaining sections of this document describe the content and format for each of these protocol layers in detail. For reference, colors in [square brackets] above indicate which protocol defines specific message components throughout this document.

Two general classifications of devices are defined by the UPnP architecture: controlled devices (or simply “devices”), and control points. A controlled device functions in the role of a server, responding to requests from control points. Both control points and controlled devices can be implemented on a variety of platforms including personal computers and embedded systems. Multiple devices, control points, or both may be operational on the same network endpoint simultaneously.

Note: This document is oriented toward an IPv4 environment. Considerations for an IPv6 environment are expressed in Appendix A.

The foundation for UPnP networking is IP addressing. In an IPv4 environment, each device or control point must have a Dynamic Host Configuration Protocol (DHCP) client and search for a DHCP server when the device or control point is first connected to the network. If a DHCP server is available, i.e., the network is managed; the device or control point MUST use the IP address assigned to it. If no DHCP server is available, i.e., the network is unmanaged; the device or control point MUST use Auto IP to get an address. In brief, Auto IP defines how a device or control point intelligently chooses an IP address from a set of reserved addresses and is able to move easily between managed and unmanaged networks. If during the DHCP transaction, the device or control point obtains a domain name, e.g., through a DNS server or via DNS forwarding, the device or control point should use that name in subsequent network operations; otherwise, the device or control point should use its IP address.

Certain UPnP networks have more complex configurations such as multiple physical networks and/or multiple logical networks to accommodate multiple non-overlapping addressing schemes. Devices and control points may also have two or more network interfaces, and/or two or more IP addresses assigned to each interface. In such configurations, a single device or control point may be assigned multiple IP addresses from different logical networks in the same UPnP network, resulting in devices appearing to a control point multiple times in the network. Devices and control points that have multiple IP addresses on the same UPnP network are referred to as multi-homed. Throughout this document, the term “UPnP-enabled interface” is used to refer to an interface which is assigned an IP address belonging to the UPnP network. Additional behaviors specific to multi-homed devices and control points will be covered in applicable sections throughout the document. However, as a general principle, related interactions between control points and devices (e.g. action control request and response messages, event subscription and event messages) MUST occur using the same pair of outgoing and incoming UPnP-enabled interfaces.

Given an IP address, Step 1 in UPnP networking is *discovery*. When a device is added to the network, the UPnP discovery protocol allows that device to advertise its services to control points on the network. Similarly, when a control point is added to the network, the UPnP discovery protocol allows that control point to search for devices of interest on the network. The fundamental exchange in both cases is a discovery message containing a few essential specifics about the device or one of its services, e.g., its type, identifier, and a pointer to more detailed information. The section on Discovery below explains how devices advertise, how control points search, and contains details about the format of discovery messages.

Step 2 in UPnP networking is *description*. After a control point has discovered a device, the control point still knows very little about the device. For the control point to learn more about the device and its capabilities, or to interact with the device, the

control point must retrieve the device's description from the URL provided by the device in the discovery message. Devices may contain other logical devices, as well as functional units, or *services*. The UPnP description for a device is expressed in XML and includes vendor-specific manufacturer information like the model name and number, the serial number, the manufacturer name, URLs to vendor-specific Web sites, etc. The description also includes a list of any embedded devices or services, as well as URLs for control, eventing, and presentation. For each service, the description includes a list of the commands, or *actions*, to which the service responds, and parameters, or *arguments* for each action; the description for a service also includes a list of variables; these variables model the state of the service at run time, and are described in terms of their data type, range, and event characteristics. The section on Description below explains how devices are described and how control points retrieve those descriptions.

Step 3 in UPnP networking is *control*. After a control point has retrieved a description of the device, the control point can send actions to a device's services. To do this, a control point sends a suitable control message to the control URL for the service (provided in the device description). Control messages are also expressed in XML using the Simple Object Access Protocol (SOAP). Like function calls, in response to the control message, the service returns any action-specific values. The effects of the action, if any, are modeled by changes in the variables that describe the run-time state of the service. The section on Control below explains the description of actions, state variables, and the format of control messages.

Step 4 in UPnP networking is *eventing*. A UPnP description for a service includes a list of actions the service responds to and a list of variables that model the state of the service at run time. The service publishes updates when these variables change, and a control point may subscribe to receive this information. The service publishes updates by sending event messages. Event messages contain the names of one or more state variables and the current value of those variables. These messages are also expressed in XML. A special initial event message is sent when a control point first subscribes; this event message contains the names and values for all evented variables and allows the subscriber to initialize its model of the state of the service. To support scenarios with multiple control points, eventing is designed to keep all control points equally informed about the effects of any action. Therefore, all subscribers are sent all event messages, subscribers receive event messages for all evented variables that have changed, and event messages are sent no matter why the state variable changed (either in response to a requested action or because the state the service is modeling changed). Multicast eventing is a variant of Step 4 in UPnP networking. Through multicast eventing, control points can listen to state changes in services without subscription. This form of eventing is useful first when events which are not relevant to specific UPnP interactions should be delivered to control points to inform users, and second when multiple controlled devices want to inform multiple other control points. Multicast eventing is inherently unreliable since it is based on UDP. To increase the probability of successful transmission, the option to retransmit multicast event notifications is outlined. UPnP Working committees should define whether specific events are multicast events. The section on Eventing below explains unicast event subscription and the format of both unicast and multicast event messages.

Step 5 in UPnP networking is *presentation*. If a device has a URL for presentation, then the control point can retrieve a page 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 device status. The degree to which each of these can be accomplished depends on the specific capabilities of the presentation page and device. The section on Presentation below explains the protocol for retrieving a presentation page.

i.4 Audience

The audience for this document includes UPnP device and control point vendors, members of UPnP Forum working committees, and anyone else who has a need to understanding the technical details of UPnP protocols.

This document assumes the reader is familiar with the HTTP, TCP, UDP, IP family of protocols; this document makes no attempt to explain them. This document also assumes most readers will be new to XML, and while it is not an XML tutorial, XML-related issues are addressed in detail given the centrality of XML to the UPnP Device Architecture. This document makes no assumptions about the reader's understanding of various programming or scripting languages.

i.5 Conformance terminology

In this document, features are described as REQUIRED, RECOMMENDED, OPTIONAL or DEPRECATED as follows:

REQUIRED (or MUST or MANDATORY).

These basic features **MUST** be implemented to comply with UPnP Device Architecture. The phrases “**MUST NOT**”, and “**PROHIBITED**” indicate behavior that is prohibited, i.e. that if performed means the implementation is not in compliance.

RECOMMENDED (or SHOULD).

These features add functionality supported by UPnP Device Architecture and **SHOULD** be implemented. **RECOMMENDED** features take advantage of the capabilities UPnP Device Architecture, usually without imposing major cost increases. Notice that for compliance testing, if a **RECOMMENDED** feature is implemented, it **MUST** meet the specified requirements to be in compliance with these guidelines. Some **RECOMMENDED** features could become requirements in the future. The phrase “**SHOULD NOT**” indicates behavior that is permitted but **NOT RECOMMENDED**.

OPTIONAL (or MAY).

These features are neither **REQUIRED** nor **RECOMMENDED** by UPnP Device Architecture, but if the feature is implemented, it **MUST** meet the specified requirements to be in compliance with these guidelines. These features are not likely to become requirements in the future.

DEPRECATED.

Although these features are still described in this specification, they should not be implemented except for backward compatibility. The occurrence of a deprecated feature during operation of an implementation compliant with the current specification has no effect on the implementation's operation and does not produce any error conditions. Backward compatibility may require that a feature is implemented and functions as specified but it **MUST** never be used by implementations compliant with this specification.

i.6 Acronyms

Table i-1: Acronyms

Acronym	Meaning	Acronym	Meaning
ARP	Address Resolution Protocol	SOAP	Simple Object Access Protocol
CP	Control Point	SSDP	Simple Service Discovery Protocol
DCP	Device Control Protocol	UDA	UPnP Device Architecture
DDD	Device Description Document	UPC	Universal Product Code
DHCP	Dynamic Host Configuration Protocol	URI	Uniform Resource Identifier
DNS	Domain Name System	URL	Uniform Resource Locator
GENA	General Event Notification Architecture	URN	Uniform Resource Name
HTML	Hypertext Markup Language	UUID	Universally Unique Identifier
HTTP	Hypertext Transfer Protocol	XML	Extensible Markup Language
SCPD	Service Control Protocol Description		

i.7 Glossary

action

Command exposed by a service. Takes one or more input or output arguments. May have a return value. For more information, see section 2, "Description" and section 3, "Control".

argument

Parameter for action exposed by a service. May be in or out. For more information, see section 2, "Description" and section 3, "Control".

control point

Retrieves device and service descriptions, sends actions to services, polls for service state variables, and receives events from services.

device

Logical device. A container. May embed other logical devices. Embeds one or more services. Advertises its presence on network(s). For more information, see section 1, "Discovery" and section 2, "Description".

device description

Formal definition of a logical device, expressed in the UPnP Template Language. Written in XML syntax. Specified by a UPnP vendor by filling in the placeholders in a UPnP Device Template, including, e.g., manufacturer name, model name, model number, serial number, and URLs for control, eventing, and presentation. For more information, see section 2, "Description".

device type

Standard device types are denoted by urn:schemas-upnp-org:device: followed by a unique name assigned by a UPnP Forum working committee. One-to-one relationship with UPnP Device Templates. UPnP vendors may specify additional device types; these are denoted by urn:domain-name:device: followed by a unique name assigned by the vendor, where *domain-name* is a Vendor Domain Name. For more information, see section 2, "Description".

event

Notification of one or more changes in state variables exposed by a service. For more information, see section 4, "Eventing".

GENA

General Event Notification Architecture. The event subscription and notification protocol defined in section 4, "Eventing".

publisher

Source of event messages. Typically a device's service. For more information, see section 4, "Eventing".

root device

A logical device that is not embedded in any other logical device. For more information, see section 2, "Description".

service

Logical functional unit. Smallest units of control. Exposes actions and models the state of a physical device with state variables. For more information, see section 3, "Control".

service description

Formal definition of a logical service, expressed in the UPnP Template language. Written in XML syntax. Specified by a UPnP vendor by filling in any placeholders in a UPnP Service Template. (Was SCPD.) For more information, see section 2, "Description".

service type

Standard service types are denoted by urn:schemas-upnp-org:service: followed by a unique name assigned by a UPnP forum working committee, colon, and an integer version number. One-to-one relationship with UPnP Service Templates. UPnP vendors may specify additional services; these are denoted by urn:*domain-name*:service: followed by a unique name assigned by the vendor, colon, and a version number, where *domain-name* is a Vendor Domain Name. For more information, see section 2, "Description".

SOAP

Simple Object Access Protocol. A remote-procedure call mechanism based on XML that sends commands and receives values over HTTP. For more information, see section 3, "Control".

SSDP

Simple Service Discovery Protocol. A multicast discovery and search mechanism that uses a multicast variant of HTTP over UDP. Defined in section 1, "Discovery".

state variable

Single facet of a model of a physical service. Exposed by a service. Has a name, data type, optional default value, optional constraints values, and may trigger events when its value changes. For more information, see section 2, "Description" and section 3, "Control".

subscriber

Recipient of event messages. Typically a control point. For more information, see section 4, "Eventing".

UPnP Device Template

Template listing device type, required embedded devices (if any), and required services. Written in XML syntax and derived from the UPnP Device Schema. Defined by a UPnP Forum working committee. One-to-one relationship with standard device types. For more information, see section 2, "Description".

UPnP Service Template

Template listing action names, parameters for those actions, state variables, and properties of those state variables. Written in XML syntax and derived from the UPnP Service Schema. Defined by a UPnP Forum working committee. One-to-one relationship with standard service types. For more information, see section 2, "Description".

UPnP Device Schema

Defines the elements and attributes used in UPnP Device and Service Templates. Written in XML syntax and derived from XML Schema (Part 1: Structures, Part 2: Datatypes). Defined by the UPnP Device Architecture herein. For more information, see section 2, "Description".

Vendor Domain Name

A domain name that is supplied by a vendor. It is owned by the vendor, and MUST be registered with an ICANN accredited Registrar, such that it is unique. The vendor MUST keep the domain name registration up to date. A Vendor Domain Name length SHOULD be chosen to be compatible with the use of the domain name in the UDA.

i.8 References and resources

RFC 2710

Multicast Listener Discovery for IPv6. Available at: <http://www.ietf.org/rfc/rfc2710.txt>.

RFC 2616

HTTP: Hypertext Transfer Protocol 1.1. Available at: <http://www.ietf.org/rfc/rfc2616.txt>.

RFC 2279

UTF-8, a transformation format of ISO 10646 (character encoding). Available at: <http://www.ietf.org/rfc/rfc2279.txt>.

XML

Extensible Markup Language. W3C recommendation. Available at: <http://www.w3.org/XML/>.

Each section in this document contains additional information about resources for specific topics.

0 Addressing

[Normative]

Addressing is Step 0 of UPnP™ networking. Through addressing, devices and control points get a network address. Addressing enables discovery (Step 1) where control points find interesting device(s), description (Step 2) where control points learn about device capabilities, control (Step 3) where a control point sends commands to device(s), eventing (Step 4) where control points listen to state changes in device(s), and presentation (Step 5) where control points display a user interface for device(s).

The foundation for UPnP networking is IP addressing. A UPnP device or control point MAY support IP version 4-only, or both IP version 4 and IP version 6. This section, and the examples given throughout sections 1 through 5 of this document, assumes an IPv4 implementation. Annex A of this document describes IPv6 operation. Each UPnP device or control point which does not itself implement a DHCP server MUST have a Dynamic Host Configuration Protocol (DHCP) client and search for a DHCP server when the device or control point is first connected to the network (if the device or control point itself implements a DHCP server, it MAY allocate itself an address from the pool that it controls). If a DHCP server is available, i.e., the network is managed; the device or control point MUST use the IP address assigned to it. If no DHCP server is available, i.e., the network is unmanaged; the device or control point MUST use automatic IP addressing (Auto-IP) to obtain an address.

Auto-IP (defined in RFC 3927) defines how a device or control point: (a) determines if DHCP is unavailable, and (b) intelligently chooses an IP address from a set of link-local IP addresses. This method of address assignment enables a device or control point to easily move between managed and unmanaged networks.

This section provides an overview of the basic operation of Auto-IP. The operations described in this section are detailed and clarified in the reference documents listed below. Where conflicts between this document and the reference documents exist, the reference document always takes precedence.

0.1 Determining whether to use Auto-IP

A device or control point that supports Auto-IP and is configured for dynamic address assignment begins by requesting an IP address via DHCP by sending out a DHCPDISCOVER message. The amount of time this DHCP Client listens for DHCPOFFERS is implementation dependent. If a DHCPOFFER is received during this time, the device or control point MUST continue the process of dynamic address assignment. If no valid DHCPOFFERS are received, the device or control point MUST then auto-configure an IP address using Auto-IP.

0.2 Choosing an address

To auto-configure an IP address using Auto-IP, the device or control point uses an implementation dependent algorithm for choosing an address in the 169.254/16 range. The first and last 256 addresses in this range are reserved and MUST NOT be used.

The selected address MUST then be tested to determine if the address is already in use. If the address is in use by another device or control point, another address MUST be chosen and tested, up to an implementation dependent number of retries. The

address selection **MUST** be randomized to avoid collision when multiple devices or control points are attempting to allocate addresses. The device or control point chooses an address using a pseudo-random algorithm (distributed over the entire address range from 169.254.1.0 to 169.254.254.255) to minimize the likelihood that devices or control points that join the network at the same time will choose the same address and subsequently choose alternative addresses in the same sequence when collisions are detected. This pseudo-random algorithm **SHOULD** be seeded using the device's or control point's Ethernet hardware MAC address.

0.3 Testing the address

To test the chosen address, the device or control point **MUST** use an Address Resolution Protocol (ARP) probe. An ARP probe is an ARP request with the device or control point hardware address used as the sender's hardware address and the sender's IP address set to 0s. The device or control point **MUST** then listen for responses to the ARP probe, or other ARP probes for the same IP address. If either of these ARP packets is seen, the device or control point **MUST** consider the address in use and try a different address. The ARP probe **MAY** be repeated for greater certainty that the address is not already in use; it is **RECOMMENDED** that the probe be sent four times at two-second intervals.

After successfully configuring a link-local address, the device or control point **MUST** send two gratuitous ARPs, spaced two seconds apart, this time filling in the sender IP address. The purpose of these gratuitous ARPs is to make sure that other hosts on the net do not have stale ARP cache entries left over from some other host that may previously have been using the same address.

Devices and control points that are equipped with persistent storage **MAY** record the IP address they have selected and on the next boot use that address as their first candidate when probing, in order to increase the stability of addresses and reduce the need to resolve address conflicts.

Address collision detection is not limited to the address testing phase, when the device or control point is sending ARP probes and listening for replies. Address collision detection is an ongoing process that is in effect for as long as the device or control point is using a link-local address. At any time, if a device or control point receives an ARP packet with its own IP address given as the sender IP address, but a sender hardware address that does not match its own hardware address, then the device or control point **MUST** treat this as an address collision and **MUST** respond as described in either (a) or (b) below:

- (a) Immediately configure a new link-local IP address as described above; or,
- (b) If the device or control point currently has active TCP connections or other reasons to prefer to keep the same IP address, and has not seen any other conflicting ARP packets recently (e.g., within the last ten seconds) then it **MAY** elect to attempt to defend its address once, by recording the time that the conflicting ARP packet was received, and then broadcasting one single gratuitous ARP, giving its own IP and hardware addresses as the source addresses of the ARP. However, if another conflicting ARP packet is received within a short time after that (e.g., within ten seconds) then the device or control point **MUST** immediately configure a new Auto-IP address as described above.

The device or control point **MUST** respond to conflicting ARP packets as described in either (a) or (b) above; it **MUST NOT** ignore conflicting ARP packets. If a new address is selected, the device or control point **MUST** cancel previous advertisements and re-advertise with the new address.

After successfully configuring an Auto-IP address, all subsequent ARP packets (replies as well as requests) containing an Auto-IP source address **MUST** be sent using link-level *broadcast* instead of link-level *unicast*, in order to facilitate timely detection of duplicate addresses.

0.4 Forwarding rules

IP packets whose source or destination addresses are in the 169.254/16 range **MUST NOT** be sent to any router for forwarding. Instead, the senders **MUST** ARP for the destination address and then send the packets directly to the destination on the same link. IP datagrams with a multicast destination address and an Auto-IP source address **MUST NOT** be forwarded off the local link. Devices and control points **MAY** assume that all 169.254/16 destination addresses are on-link and directly reachable. The 169.254/16 address range **MUST NOT** be subnetted.

0.5 Periodic checking for dynamic address availability

A device or control point that has auto-configured an IP address **MUST** periodically check for the existence of a DHCP server. This is accomplished by sending DHCPDISCOVER messages. How often this check is made is implementation dependent, but checking every 5 minutes would maintain a balance between network bandwidth required and connectivity maintenance. If a DHCP OFFER is received, the device or control point **MUST** proceed with dynamic address allocation. Once a DHCP assigned address is in place, the device or control point **MAY** release the auto-configured address, but **MAY** also choose to maintain this address for a period of time (or indefinitely) to maintain connectivity.

To switch over from one IP address to a new one, the device **SHOULD**, if possible, cancel any outstanding advertisements made on the previous address and **MUST** issue new advertisements on the new address. The section on Discovery explains advertisements and their cancellations. In addition, any event subscriptions are deleted by the device (see section on Eventing).

For a multi-homed device with multiple IP addresses, to switch one of the IP addresses to a new one, the device **SHOULD** cancel any outstanding advertisements made on the previous IP address, and **MUST** issue new advertisements on the new IP addresses. Furthermore, it **MUST** also issue appropriate update advertisements on all unaffected IP addresses. The section on Discovery explains advertisements, their cancellations and updates. The section on Eventing explains the effect on event subscriptions.

0.6 Device naming and DNS interaction

Once a device has a valid IP address for the network, it can be located and referenced on that network through that address. There may be situations where the end user needs to locate and identify a device. In these situations, a friendly name for the device is much easier for a human to use than an IP address. If a device chooses to provide a host name to a DHCP server and register with a DNS server, the device **SHOULD** either ensure the requested host name is unique or provide a means for the user

to change the requested host name. Most often, devices do not provide a host name, but provide URLs using literal (numeric) IP addresses.

Moreover, names are much more static than IP addresses. Clients referring a device by name don't require any modification when the IP address of a device changes. Mapping of the device's DNS name to its IP address could be entered into the DNS database manually or dynamically according to RFC 2136. While devices supporting dynamic DNS updates can register their DNS records directly in the DNS, it is also possible to configure a DHCP server to register DNS records on behalf of these DHCP clients.

0.7 Name to IP address resolution

A device that needs to contact another device identified by a DNS name needs to discover its IP address. The device submits a DNS query according to RFC1034 and 1035 to the pre-configured DNS server(s) and receives a response from a DNS server containing the IP address of the target device. A device can be statically pre-configured with the list of DNS servers.

Alternatively a device could be configured with the list of DNS server through DHCP, or after the address assignment through a DHCPINFORM message.

0.8 References

RFC1034

Domain Names - Concepts and Facilities. Available at: <http://www.ietf.org/rfc/rfc1034.txt>.

RFC1035

Domain Names - Implementation and Specification. Available at: <http://www.ietf.org/rfc/rfc1035.txt>.

RFC 2131

Dynamic Host Configuration Protocol. Available at: <http://www.ietf.org/rfc/rfc2131.txt>.

RFC 2136

Dynamic Updates in the Domain Name System. Available at: <http://www.ietf.org/rfc/rfc2136.txt>.

RFC 3927

Dynamic Configuration of IPv4 Link-Local Addresses. Available at: <http://www.ietf.org/rfc/rfc3927.txt>.

1 Discovery


[Normative]

Discovery is Step 1 in UPnP™ networking. Discovery comes after addressing (Step 0) where devices get a network address.

Through discovery, control points find interesting device(s). Discovery enables description (Step 2) where control points learn about device capabilities, control (Step 3) where a control point sends commands to device(s), eventing (Step 4) where control points listen to state changes in device(s), and presentation (Step 5) where control points display a user interface for device(s).

Discovery is the first step in UPnP networking. When a device is added to the network, the UPnP discovery protocol allows that device to advertise its services to control points on the network. Similarly, when a control point is added to the network, the UPnP discovery protocol allows that control point to search for devices of interest on the network. The fundamental exchange in both cases is a discovery message containing a few, essential specifics about the device or one of its services, e.g., its type, universally unique identifier, a pointer to more detailed information and optionally parameters that identify the current state of the device.

Figure 1-1: Discovery architecture



When a device knows it is newly added to the network, it MUST multicast a number of discovery messages advertising itself, its embedded devices, and its services (initial announce). Any interested control point can listen to the standard multicast address for notifications that new capabilities are available. A multi-homed device MUST multicast the discovery messages on all UPnP-enabled interfaces. A multi-homed control point MAY listen to the standard multicast address on one, some or all of its UPnP-enabled interfaces.

When a new control point is added to the network, it MAY multicast a discovery message searching for interesting devices, services, or both. All devices MUST listen to the standard multicast address for these messages and MUST respond if any of their root devices, embedded devices or services matches the search criteria in the discovery message. In addition, a control point MAY unicast a discovery message to a specific IP address on port 1900 or on the port specified by the optional

SEARCHPORT.UPNP.ORG header field (which supersedes port 1900 for this use), searching for a UPnP device or service at that specific IP address. This action presumes the control point already knows the device at this IP address is a UPnP 1.1 device (which listens on the appropriate port). The control point can use unicast search for a number of applications. A unicast search can quickly confirm a specific device and provide the corresponding discovery information (e.g. UUID, URL) of this device. All devices MUST listen to incoming unicast search messages on port 1900 or, if provided, the port number specified in the SEARCHPORT.UPNP.ORG header field and MUST respond if any of their root devices, embedded devices or services matches the search criteria in the discovery message.

A multi-homed control point MAY multicast discovery messages on one, some or all of its UPnP-enabled interfaces. Multi-homed devices MUST listen to the standard multicast address on all UPnP-enabled interfaces for multicast discovery messages. Multi-homed devices MUST also listen to incoming unicast search messages on port 1900 or, if provided, the port number specified in the SEARCHPORT.UPNP.ORG header field. The devices MUST respond if any of their root devices, embedded devices or services matches the search criteria in the discovery message.

To reiterate, a control point MAY learn of a device of interest because that device sent discovery messages advertising itself or because the device responded to a discovery message searching for devices. In either case, if a control point is interested in a device and wants to learn more about it, the control point uses the information in the discovery message to send a *description* query message. The section on Description explains description messages in detail.

When a device is removed from the network, it SHOULD, if possible, multicast a number of discovery messages revoking its earlier announcements, effectively declaring that its root devices, embedded devices and services will no longer be available. When the IP address of a device is changed, it SHOULD revoke any earlier announcements and it MUST advertise using the new IP address.

When a multi-homed device becomes unavailable to the network on any of its UPnP-enabled interfaces, it SHOULD, if possible, multicast a number of discovery messages revoking its earlier announcements on the affected UPnP-enabled interfaces, effectively declaring that its root devices, embedded devices and services will no longer be available on those interfaces. If it remains available to the network on any of its other UPnP-enabled interfaces, it MUST NOT multicast such discovery messages on the unaffected UPnP-enabled interfaces.

When a multi-homed device becomes available to the network on a new UPnP-enabled interface (in addition to any existing UPnP-enabled interfaces), it MUST increase its BOOTID.UPNP.ORG field value (see section 1.2 "Advertisement"), and multicast a number of update messages on the existing UPnP-enabled interfaces to announce the new BOOTID.UPNP.ORG field value. After all the update messages have been sent, it MUST multicast a number of discovery messages on all (existing and new) UPnP-enabled interfaces with the new BOOTID.UPNP.ORG field value.

Similarly, when one of the IP addresses of a multi-homed device is changed, it SHOULD revoke any earlier announcements on the previous IP address. It MUST increase its BOOTID.UPNP.ORG field value (see section 1.2 "Advertisement"), and multicast a

number of update messages on the existing UPnP-enabled interfaces to announce the new BOOTID.UPNP.ORG field value. After all the update messages have been sent, it MUST multicast a number of discovery messages on all (existing and new) UPnP-enabled interfaces with the new BOOTID.UPNP.ORG field value.

Finally, if a multi-homed device loses connectivity on one of its UPnP-enabled interfaces and then regains connectivity, it MUST increase its BOOTID.UPNP.ORG field value (see section “1.2 Advertisement”), and multicast a number of update messages on the unaffected UPnP-enabled interfaces to announce the new BOOTID.UPNP.ORG field value. After all the update messages have been sent, it MUST multicast a number of discovery messages on all (affected and unaffected) UPnP-enabled interfaces with the new BOOTID.UPNP.ORG field value.

To limit network congestion, the time-to-live (TTL) of each IP packet for each multicast message SHOULD default to 2 and SHOULD be configurable. When the TTL is greater than 1, it is possible for multicast messages to traverse multiple routers; therefore control points and devices using non-AutoIP addresses MUST send an IGMP Join message so that routers will forward multicast messages to them (this is not necessary when using an Auto-IP address, since packets with Auto-IP addresses will not be forwarded by routers).

Versioning: Discovery plays an important role in the interoperability of devices and control points using different versions of UPnP networking. The UPnP Device Architecture (defined herein) is versioned with both a major and a minor version, usually written as *major.minor*, where both *major* and *minor* are integers (for example, version 2.10 [two dot ten] is *newer* than version 2.2 [two dot two]). Advances in minor versions MUST be a compatible superset of earlier minor versions of the same major version. Advances in major version are not required to be supersets of earlier versions and are not guaranteed to be backward compatible. Version information is communicated in discovery and description messages. Discovery messages include the version of UPnP networking that the devices and control points support (in the SERVER and USER-AGENT header fields); the version of device and service types supported is also included in relevant discovery messages. Additionally, description documents also include the same information. SERVER and USER-AGENT header fields are also used in control and eventing to communicate which version of UPnP networking the devices and control points support. This section explains the format of version information in discovery messages and specific requirements on discovery messages to maintain compatibility with advances in minor versions.

The remainder of this section explains the UPnP discovery protocol known as SSDP (Simple Service Discovery Protocol) in detail, enumerating how devices advertise and revoke their advertisements as well as how control points search and devices respond.

1.1 SSDP message format

SSDP uses part of the header field format of HTTP 1.1 as defined in RFC 2616. However, it is NOT based on full HTTP 1.1 as it uses UDP instead of TCP, and it has its own processing rules. This section defines the generic format of a SSDP message.

All SSDP messages MUST be formatted according to RFC 2616 section 4.1 “generic message”. SSDP messages MUST have a start-line and a list of message header fields. SSDP messages SHOULD NOT have a message body. If a SSDP message is received with a message body, the message body MAY be ignored.

1.1.1 SSDP Start-line

Each SSDP message MUST have exactly one start-line. See section 1.2, “Advertisement” and section 1.3, “Search” below for the definition of all possible SSDP messages. The start-line MUST be formatted either as defined in RFC 2616 section 5.1 or section 6.1. Furthermore, the start-line MUST be one of the following three:

[NOTIFY * HTTP/1.1\r\n](#)

[M-SEARCH * HTTP/1.1\r\n](#)

[HTTP/1.1 200 OK\r\n](#)

As a clarification, while the start-line MUST include “HTTP/1.1”, this does not signal that SSDP is fully based on HTTP 1.1; this start-line element is included for backward compatibility reasons only.

1.1.2 SSDP message header fields

The message header fields in a SSDP message MUST be formatted according to RFC 2616 section 4.2. This specifies that each message header field consist of a case-insensitive field name followed by a colon (":"), followed by the case-sensitive field value. SSDP restricts allowed field values.

Example SSDP header:

[HOST: 239.255.255.250:1900](#)

1.1.3 SSDP header field extensions

UPnP working committees and UPnP vendors are allowed to extend SSDP messages with additional SSDP header fields. Additional message header fields can also be defined by the UPnP Forum Technical committee (e.g. section 1.2, “Advertisement” defines BOOTID.UPNP.ORG, CONFIGID.UPNP.ORG, NEXTBOOTID.UPNP.ORG, and SEARCHPORT.UPNP.ORG header fields). To prevent name-clashes of header field definitions (two parties accidentally define the same header field name with different semantics), vendor-defined header field names MUST have the following format:

field-name = token "." domain-name

where the domain-name MUST be Vendor Domain Name, and in addition MUST satisfy the token format as defined in RFC 2616, section 2.2.

Example vendor-defined SSDP header fields:

[myheader.philips.com: "some value"](#)

[myheader.sony.com: "other value"](#)

1.1.4 UUID format and RECOMMENDED generation algorithms

UPnP 1.1 devices MUST format UUIDs according to the format specified below. However, UPnP 1.1 control points MUST also be able to accept UUIDs that have not been formatted according to the rules specified below, as formatting rules are not specified in UPnP 1.0 other than the requirement that a UUID is a string.

UUIDs are 128 bit numbers that MUST be formatted as specified by the following grammar (taken from [1]):

```
UUID = 4 * <hexOctet> "-" 2 * <hexOctet> "-" 2 * <hexOctet> "-" 2 * <hexOctet> "-" 6 * <hexOctet>
hexOctet = <hexDigit> <hexDigit>
hexDigit = "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9" | "a" | "b" | "c" | "d" | "e" | "f" | "A" | "B" | "C" | "D" | "E" | "F"
```

The following is an example of the string representation of a UUID:

```
"2fac1234-31f8-11b4-a222-08002b34c003"
```

UUIDs MAY be generated using *any* suitable generation algorithm² that satisfies the following requirements:

1. It is very unlikely to duplicate a UUID generated from some other resource.
2. It maps down to a 128-bit number.
3. UUIDs MUST remain fixed over time.

The following UUID generation algorithm is RECOMMENDED:

Time & MAC-based algorithm as specified in [1], where the UUID is generated once and stored in non-volatile memory if available.

1.1.5 SSDP processing rules

When an SSDP message is received that is not formatted according to section 1.1, "SSDP message format" (the sections above), receivers SHOULD silently discard the message. Receivers MAY try to parse such SSDP messages to try to interoperate.

When parsing SSDP header fields, receivers MUST parse all REQUIRED SSDP-defined header fields (see section 1.2, "Advertisement" and section 1.3, "Search" below) and MAY skip all other header fields. Receivers MUST be able to skip header fields they do not understand.

1.2 Advertisement

When a device is added to the network, the device advertises its services to control points. It does this by multicasting discovery messages to a standard address and port (239.255.255.250:1900). Control points listen to this port to detect when new capabilities are available on the network. To advertise the full extent of its capabilities, a device MUST multicast a number of discovery messages corresponding to each of its root devices, embedded devices and services. Each message contains information specific to the embedded device (or service) as well as information about its enclosing device. Messages MUST include duration


² The UUID generation algorithm specified in [1] is RECOMMENDED, but is not MANDATORY, other UUID generation algorithms may be used instead, as long as they satisfy the three requirements.

until the advertisements expire; if the device remains available, the advertisements MUST be re-sent (with new duration). If the device becomes unavailable, the device SHOULD explicitly cancel its advertisements, but if the device is unable to do this, the advertisements will expire on their own. If a multi-homed device becomes unavailable on some, but not all, of its UPnP-enabled interfaces, the device SHOULD explicitly cancel its advertisements on the affected UPnP-enabled interfaces (but NOT on the unaffected UPnP-enabled interfaces), but if the device is unable to do this, the advertisements on those interfaces or IP addresses will expire on their own. In addition, messages include the following header fields defined in this document: BOOTID.UPNP.ORG, NEXTBOOTID.UPNP.ORG, CONFIGID.UPNP.ORG, SEARCHPORT.UPNP.ORG. The field value of the BOOTID.UPNP.ORG header field MUST be increased each time a device (re)joins the network and sends an initial announce (a “reboot” in UPnP terms), or adds a UPnP-enabled interface. Unless the device explicitly announces a change in the BOOTID.UPNP.ORG field value using an SSDP message, as long as the device remains continuously available in the network, the same BOOTID.UPNP.ORG field value MUST be used in all repeat announcements, search responses, update messages and eventually bye-bye messages. Control points can parse this header field to detect whether the device has potentially lost its state (event subscriptions will have been lost, DCP specific state may have been changed) due to a “reboot”. Since a device cannot change IP addresses without changing the BOOTID.UPNP.ORG field value, the BOOTID.UPNP.ORG field value can also be used to distinguish multi-homed devices (in this case, a control point will see SSDP messages from different IP addresses with the same UUID, BOOTID.UPNP.ORG field value) from devices that changed IP addresses (in this case, the BOOTID.UPNP.ORG field value will be different). The field value of the NEXTBOOTID.UPNP.ORG header field indicates the field value of the BOOTID.UPNP.ORG header field that a multi-homed device intends to use in future announcements after adding a new UPnP-enabled interface. The field value of the CONFIGID.UPNP.ORG header field identifies the current set of device and service descriptions; control points can parse this header field to detect whether they need to send new *description* query messages. The field value of the SEARCHPORT.UPNP.ORG header field identifies the port at which the device listens to unicast M-SEARCH messages; control points can parse this header field to know to which port unicast M-SEARCH messages MUST be sent. These header fields are explained in detail below.

1.2.1 Advertisement protocols and standards

To send (and receive) advertisements, devices (and control points) use the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

Figure 1-2: Advertisement protocol stack



At the highest layer, discovery messages contain vendor-specific information, e.g., URL for the device description and device identifier. Moving down the stack, vendor content is supplemented by information from a UPnP Forum working committee, e.g., device type. Messages from the layers above are hosted in UPnP-specific protocols, defined in this document. In turn, the SSDP messages are delivered via UDP over IP. For reference, colors in [square brackets] above indicate which protocol defines specific header fields and field values in discovery messages listed below.

1.2.2 Device available - NOTIFY with `ssdp:alive`

When a device is added to the network, it MUST multicast discovery messages to advertise its root device, any embedded devices, and any services. Each discovery message MUST contain four major components:

1. A notification type (e.g., device type), sent in an NT (Notification Type) header field.
2. A composite identifier for the advertisement, sent in a USN (Unique Service Name) header field.
3. A URL for more information about the device (or enclosing device in the case of a service), sent in a LOCATION header field.
4. A duration for which the advertisement is valid, sent in a CACHE-CONTROL header field.

To advertise its capabilities, a device multicasts a number of discovery messages. Specifically, a root device MUST multicast:

- Three discovery messages for the root device.

Table 1-1: Root device discovery messages

NT		USN *
1	<u>upnp:rootdevice</u>	uuid: <i>device-UUID</i> : <u>upnp:rootdevice</u>
2	uuid: <i>device-UUID</i> **	uuid: <i>device-UUID</i> (for root device UUID)
3	urn: <u>schemas-upnp-org:device:deviceType:ver</u> or urn: <i>domain-name</i> : <u>device:deviceType:ver</u>	uuid: <i>device-UUID</i> : urn: <u>schemas-upnp-org:device:deviceType:ver</u> (of root device) or uuid: <i>device-UUID</i> : urn: <i>domain-name</i> : <u>device:deviceType:ver</u>

- Two discovery messages for each embedded device.

Table 1-2: Embedded device discovery messages

NT		USN *
1	uuid: <i>device-UUID</i> **	uuid: <i>device-UUID</i>
2	urn: <u>schemas-upnp-org:device:deviceType:ver</u> or urn: <i>domain-name</i> : <u>device:deviceType:ver</u>	uuid: <i>device-UUID</i> : urn: <u>schemas-upnp-org:device:deviceType:ver</u> or uuid: <i>device-UUID</i> : urn: <i>domain-name</i> : <u>device:deviceType:ver</u>

- Once for each service type in each device.

Table 1-3: Service discovery messages

NT		USN *
1	urn: <u>schemas-upnp-org:service:serviceType:ver</u> or urn: <i>domain-name</i> : <u>service:serviceType:ver</u>	uuid: <i>device-UUID</i> : urn: <u>schemas-upnp-org:service:serviceType:ver</u> or uuid: <i>device-UUID</i> : urn: <i>domain-name</i> : <u>service:serviceType:ver</u>

* Note that the prefix of the USN header field (before the double colon) MUST match the value of the UDN element in the device description. (Section 2, "Description" explains the UDN element.)

** Note that the field value of this NT header field MUST match the value of the UDN element in the device description.

If a root device has d embedded devices and s embedded services but only k distinct service types, this works out to $3+2d+k$ requests. If a particular device or embedded device contains multiple instances of a particular service type, it is only necessary to advertise the service type once (rather than once for each instance). Note that if two embedded devices contain a service of the same service type, these services MUST still be separately announced. This advertises the full extent of the device's capabilities to interested control points. These messages MUST be sent out as a series with roughly comparable expiration times; order is unimportant, but refreshing or canceling individual messages is PROHIBITED.

Updated UPnP device and service types are REQUIRED to be fully backward compatible with previous versions of the same type. Devices MUST advertise the highest supported version of each supported type. For example, if a device supports version 2 of the "Audio" service, it would advertise only version 2, even though it also supports version 1. It MUST NOT advertise additional supported versions. Control points that support a given version of a device or service are able to also interact with higher

versions because of this backward compatibility requirement, but only using the functionality that was defined in the lower version. For example, if a control point supports only version “1” of the “Audio” service, and a device advertises that it supports version “2” of the “Audio” service, the control point MUST recognize the device and be able to use it.

Choosing an appropriate duration for advertisements is a balance between minimizing network traffic and maximizing freshness of device status. Relatively short durations close to the minimum of 1800 seconds will ensure that control points have current device status at the expense of additional network traffic; longer durations, say on the order of a day, compromise freshness of device status but can significantly reduce network traffic. Generally, device vendors should choose a value that corresponds to expected device usage: short durations for devices that are expected to be part of the network for short periods of time, and significantly longer durations for devices expected to be long-term members of the network. Devices that frequently connect to and leave the network (such as mobile wireless devices) SHOULD use a shorter duration so that control points have a more accurate view of their availability. Advertisements in a set (both initial and subsequent) SHOULD have comparable durations. Advertisements in the initial set SHOULD be sent as quickly as possible. Subsequent refreshments of the advertisements MAY be spread over time rather than being sent as a group.

Spreading refreshments of advertisements over time rather than being sent as a group improves reliability in case there are network glitches: without increasing the total network load it increases the frequency of sending announcements from devices to control points. The two figures below show the announcement behavior without spreading and with spreading the messages over the entire interval. The figures show a timeline from the moment a device joins the network, sends its initial announcements (represented by vertical lines), and subsequently periodically sends repeat announcements. In the second figure, these repeat announcements are spread over the entire period rather than sent as a bunch.

Figure 1-3: Initial and repeat announcements, no announcement spreading





Figure 1-4: Initial and repeat announcements, message spreading of repeat announcements



Devices SHOULD wait a random interval (e.g. between 0 and 100milliseconds) before sending an initial set of advertisements in order to reduce the likelihood of network storms; this random interval SHOULD also be applied on occasions where the device obtains a new IP address or a new UPnP-enabled interface is installed.

Due to the unreliable nature of UDP, devices SHOULD send the entire set of discovery messages more than once with some delay between sets e.g. a few hundred milliseconds. To avoid network congestion discovery messages SHOULD NOT be sent more than three times. In addition, the device MUST re-send its advertisements periodically prior to expiration of the duration specified in

the CACHE-CONTROL header field; it is RECOMMENDED that such refreshing of advertisements be done at a randomly-distributed interval of less than one-half of the advertisement expiration time, so as to provide the opportunity for recovery from lost advertisements before the advertisement expires, and to distribute over time the advertisement refreshment of multiple devices on the network in order to avoid spikes in network traffic. Note that UDP packets are also bounded in length (perhaps as small as 512 Bytes in some implementations); each discovery message MUST fit entirely in a single UDP packet. There is no guarantee that the above 3+2d+k messages will arrive in a particular order.

A multi-homed device MUST perform the above announcement procedures on each of its UPnP-enabled interfaces. Advertisements sent on multiple UPnP-enabled interfaces MUST contain the same field values except for the HOST, CACHE-CONTROL and LOCATION header fields. The HOST field value of an advertisement MUST be the standard multicast address specified for the protocol (IPv4 or IPv6) used on the interface. The URL specified by the LOCATION header field MUST be reachable on the interface on which the advertisement is sent. Finally, advertisements sent on different interfaces MAY have different CACHE-CONTROL field values and MAY be sent with different frequencies.

When a device is added to the network, it MUST send a multicast message with method NOTIFY and ssdp:alive in the NTS header field in the following format. Values in *italics* are placeholders for actual values.

```
NOTIFY * HTTP/1.1
HOST: 239.255.255.250:1900
CACHE-CONTROL: max-age = seconds until advertisement expires
LOCATION: URL for UPnP description for root device
NT: notification type
NTS: ssdp:alive
SERVER: OS/version UPnP/1.1 product/version
USN: composite identifier for the advertisement
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or an update message
CONFIGID.UPNP.ORG: number used for caching description information
SEARCHPORT.UPNP.ORG: number identifies port on which device responds to unicast M-SEARCH
```

Note: No body is sent for messages with method NOTIFY, but note that the message MUST have a blank line following the last header field.

The TTL for the IP packet SHOULD default to 2 and SHOULD be configurable.

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive.

All field values are case sensitive except where noted.

Request line

Must be "NOTIFY * HTTP/1.1"

NOTIFY

Method for sending notifications and events.

*

Message applies generally and not to a specific resource. MUST be *.

HTTP/1.1

HTTP version.

Header fields

HOST

REQUIRED. Field value contains multicast address and port reserved for SSDP by Internet Assigned Numbers Authority (IANA). MUST be [239.255.255.250:1900](#). If the port number (":1900") is omitted, the receiver MUST assume the default SSDP port number of 1900.

CACHE-CONTROL

REQUIRED. Field value MUST have the max-age directive ("max-age=") followed by an integer that specifies the number of seconds the advertisement is valid. After this duration, control points SHOULD assume the device (or service) is no longer available; as long as a control point has received at least one advertisement that is still valid from a root device, any of its embedded devices or any of its services, then the control point can assume that all are available. The number of seconds SHOULD be greater than or equal to 1800 seconds (30 minutes), although exceptions are defined in the text above. Specified by UPnP vendor. Other directives MUST NOT be sent and MUST be ignored when received.

LOCATION

REQUIRED. Field value contains a URL to the UPnP description of the root device. Normally the host portion contains a literal IP address rather than a domain name in unmanaged networks. Specified by UPnP vendor. Single absolute URL (see RFC 3986).

NT

REQUIRED. Field value contains Notification Type. MUST be one of the following. (See Table 1-1, "Root device discovery messages", Table 1-2, "Embedded device discovery messages", and Table 1-3, "Service discovery messages" above.) Single URI.

[upnp:rootdevice](#)

Sent once for root device.

uuid:[device-UUID](#)

Sent once for each device, root or embedded, where [device-UUID](#) is specified by the UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

urn:[schemas-upnp-org:device:deviceType:ver](#)

Sent once for each device, root or embedded, where [deviceType](#) and [ver](#) are defined by UPnP Forum working committee, and [ver](#) specifies the version of the device type.

urn:[schemas-upnp-org:service:serviceType:ver](#)

Sent once for each service where [serviceType](#) and [ver](#) are defined by UPnP Forum working committee and [ver](#) specifies the version of the service type.

urn:[domain-name:device:deviceType:ver](#)

Sent once for each device, root or embedded, where [domain-name](#) is a Vendor Domain Name, [deviceType](#) and [ver](#) are defined by the UPnP vendor, and [ver](#) specifies the version of the device type. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

urn:[domain-name:service:serviceType:ver](#)

Sent once for each service where [domain-name](#) is a Vendor Domain Name, [serviceType](#) and [ver](#) are defined by UPnP vendor, and [ver](#) specifies the version of the service type. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

NTS

REQUIRED. Field value contains Notification Sub Type. MUST be [ssdp:alive](#). Single URI.

SERVER

REQUIRED. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form [OS name/OS version](#), the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form [product name/product version](#). For example, "SERVER: [unix/5.1 UPnP/1.1 MyProduct/1.0](#)". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

USN

REQUIRED. Field value contains Unique Service Name. Identifies a unique instance of a device or service. MUST be one of the following. (See Table 1-1, "Root device discovery messages", Table 1-2, "Embedded device discovery messages", and Table 1-3, "Service discovery messages" above.) The prefix (before the double colon) MUST match the value of the UDN element in the device description. (Section 2, "Description" explains the UDN element.) Single URI.

uuid:[device-UUID:upnp:rootdevice](#)

Sent once for root device where [device-UUID](#) is specified by UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

uuid: *device-UUID*

Sent once for every device, root or embedded, where *device-UUID* is specified by the UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

uuid: *device-UUID*:urn:*schemas-upnp-org*:*device*:*deviceType*:*ver*

Sent once for every device, root or embedded, where *device-UUID* is specified by the UPnP vendor, *deviceType* and *ver* are defined by UPnP Forum working committee and *ver* specifies version of the device type. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

uuid: *device-UUID*:urn:*schemas-upnp-org*:*service*:*serviceType*:*ver*

Sent once for every service where *device-UUID* is specified by the UPnP vendor, *serviceType* and *ver* are defined by UPnP Forum working committee and *ver* specifies version of the device type. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

uuid: *device-UUID*:urn:*domain-name*:*device*:*deviceType*:*ver*

Sent once for every device, root or embedded, where *device-UUID*, *domain-name* (a Vendor Domain Name), *deviceType* and *ver* are defined by the UPnP vendor and *ver* specifies the version of the device type. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format. Period characters in the Vendor Domain Name MUST be replaced by hyphens in accordance with RFC 2141.

uuid: *device-UUID*:urn:*domain-name*:*service*:*serviceType*:*ver*

Sent once for every service where *device-UUID*, *domain-name* (a Vendor Domain Name), *serviceType* and *ver* are defined by the UPnP vendor and *ver* specifies the version of the service type. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format. Period characters in the Vendor Domain Name MUST be replaced by hyphens in accordance with RFC 2141.

BOOTID.UPNP.ORG

REQUIRED. The BOOTID.UPNP.ORG header field represents the boot instance of the device expressed according to a monotonically increasing value. Its field value MUST be a non-negative 31-bit integer; ASCII encoded, decimal, without leading zeros (leading zeroes, if present, MUST be ignored by the recipient) that MUST be increased on each initial announce of the UPnP device or MUST be the same as the field value of the NEXTBOOTID.UPNP.ORG header field in the last sent SSDP update message. Its field value MUST remain the same on all periodically repeated announcements. A convenient mechanism is to set this field value to the time that the device sends its initial announcement, expressed as seconds elapsed since midnight January 1, 1970; for devices that have a notion of time, this will not require any additional state to remember or be "flushed". However, it is perfectly acceptable to use a simple boot counter that is incremented on every initial announcement as a field value of this header field. As such, control points MUST NOT view this header field as a timestamp. The BOOTID.UPNP.ORG header field MUST be included in *all* announcements of a root device, its embedded devices and its services. Unless the device explicitly updates its value by sending an SSDP update message, as long as the device remains available in the network, the same BOOTID.UPNP.ORG field value MUST be used in *all* announcements, search responses, update messages and eventually bye-bye messages.

Control points can use this header field to detect the case when a device leaves and rejoins the network ("reboots" in UPnP terms). It can be used by control points for a number of purposes such as re-establishing desired event subscriptions, checking for changes to the device state that were not evented since the device was off-line.

CONFIGID.UPNP.ORG

REQUIRED. The CONFIGID.UPNP.ORG field value MUST be a non-negative, 31-bit integer, ASCII encoded, decimal, without leading zeros (leading zeroes, if present, MUST be ignored by the recipient) that MUST represent the configuration number of a root device. UPnP 1.1 devices MAY freely assign configid numbers from 0 to 16777215 ($2^{24}-1$). Higher numbers are reserved for future use, and can be assigned by the Technical Committee. The **configuration** of a root device consists of the following information: the DDD of the root device and all its embedded devices, and the SCPDs of all the contained services. If any part of the configuration changes, the CONFIGID.UPNP.ORG field value MUST be changed. The CONFIGID.UPNP.ORG header field MUST be included in all announcements of a root device, its embedded devices and its services. The configuration number that is present in a CONFIGID.UPNP.ORG field value MUST satisfy the following rule:

- if a device sends out two messages with a CONFIGID.UPNP.ORG header field with the same field value K, the **configuration** MUST be the same at the moments that these messages were sent.

Whenever a control point receives a CONFIGID.UPNP.ORG header field with a field value K, and subsequently downloads the configuration information, this configuration information is associated with K. As an additional safeguard, the device MUST include a `configId` attribute with value K in the returned description (see section 2, "Description"). The following caching rules for control points supersede the caching rules that are defined in UPnP 1.0:

- Control points MAY ignore the CONFIGID.UPNP.ORG header field and use the caching rules that are based on advertisement expirations as defined in Chapter 2, Description: as long as at least one of the discovery advertisements from a root device, its embedded devices and its services have not expired, a control point MAY assume that the root device and all its embedded devices and all its services are available. The device and service descriptions MAY be retrieved at any point since the device and service descriptions are static as long as the device and its services are available.

- If no configuration number is included in a received SSDP message, control points SHOULD cache based on advertisement expirations as defined in Chapter 2 Description.
- If a CONFIGID.UPNP.ORG header field with field value K is included in a received SSDP message, and a control point has already cached information associated with field value K, the control point MAY use this cached information as the current configuration of the device. Otherwise, a control point SHOULD assume it has not cached the current configuration of the device and needs to send new *description* query messages.

The CONFIGID.UPNP.ORG header field reduces peak loads on UPnP devices during startup and during network hiccups. Only if a control point receives an announcement of an unknown configuration is downloading required.

SEARCHPORT.UPNP.ORG

OPTIONAL. If a device does not send the SEARCHPORT.UPNP.ORG header field, it MUST respond to unicast M-SEARCH messages on port 1900. Only if port 1900 is unavailable MAY a device select a different port to respond to unicast M-SEARCH messages. If a device sends the SEARCHPORT.UPNP.ORG header field, its field value MUST be an ASCII encoded integer, decimal, without leading zeros (leading zeroes, if present, MUST be ignored by the recipient), in the range 49152-65535 (RFC 4340). The device MUST respond to unicast M-SEARCH messages that are sent to the advertised port.

Note: No responses are sent for messages with method NOTIFY.

1.2.3 Device unavailable -- NOTIFY with ssdp:byebye

When a device and its services are going to be removed from the network, the device SHOULD multicast an `ssdp:byebye` message corresponding to *each* of the `ssdp:alive` messages it multicast that have not already expired. If the device is removed abruptly from the network, it might not be possible to multicast a message. As a fallback, discovery messages MUST include an expiration value in a CACHE-CONTROL field value (as explained above); if not re-advertised, the discovery message eventually expires on its own.

(Note: when a control point is about to be removed from the network, no discovery-related action is required.)

When a device is about to be removed from the network, it SHOULD explicitly revoke its discovery messages by sending one multicast message for *each* `ssdp:alive` message it sent. Each multicast message MUST have method NOTIFY and `ssdp:byebye` in the NTS header field in the following format. Values in *italics* are placeholders for actual values.

When a multi-homed device is about to be removed from the network on one or more of its UPnP-enabled interfaces, it SHOULD explicitly revoke its discovery messages by sending one multicast message for *each* `ssdp:alive` message it has previously sent on those interfaces and IP addresses. It MUST NOT send such multicast messages to any of the UPnP-enabled interfaces that remain available.

When `ssdp:byebye` messages are sent on multiple UPnP-enabled interfaces, the messages MUST contain identical field values except for the HOST field value. The HOST field value of an advertisement MUST be the standard multicast address specified for the protocol (IPv4 or IPv6) used on the interface.

```

NOTIFY * HTTP/1.1
HOST: 239.255.255.250:1900
NT: notification type
NTS: ssdp:byebye
USN: composite identifier for the advertisement
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or an update
message
CONFIGID.UPNP.ORG: number used for caching description information

```


Note: No body is present for messages with method NOTIFY, but note that the message MUST have a blank line following the last header field.

The TTL for the IP packet SHOULD default to 2 and SHOULD be configurable.

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Request line

Must be "NOTIFY * HTTP/1.1"

NOTIFY

Method for sending notifications and events.

*

Message applies generally and not to a specific resource. MUST be *.

HTTP/1.1

HTTP version.

Header fields

HOST

REQUIRED. Field value contains multicast address and port reserved for SSDP by Internet Assigned Numbers Authority (IANA). MUST be [239.255.255.250:1900](#). If the port number (":1900") is omitted, the receiver MUST assume the default SSDP port number of 1900.

NT

REQUIRED. Field value contains Notification Type. (See list of required field values for the NT header field in NOTIFY with `ssdp:alive` above.) Single URI.

NTS

REQUIRED. Field value contains Notification Sub Type. MUST be [ssdp:byebye](#). Single URI.

USN

REQUIRED. Field value contains Unique Service Name. (See list of required field values for the USN header field in NOTIFY with `ssdp:alive` above.) Single URI.

BOOTID.UPNP.ORG

REQUIRED. As defined in section 1.2, and 1.2.2.

CONFIGID.UPNP.ORG

REQUIRED. As defined in section 1.2, and 1.2.2.

Note: No responses are sent for messages with method NOTIFY.

If a control point has received *at least one* `ssdp:byebye` message of a root device, any of its embedded devices or any of its services then the control point can assume that all are no longer available. As a fallback, if a control point fails to receive notification that a root device, its embedded devices and its services are unavailable, the original discovery messages will eventually expire yielding the same effect. Only when *all* original advertisements of a root device, its embedded devices and its services have expired can a control point assume that they are no longer available.

If a multi-homed control point has received *at least one* `ssdp:byebye` message of a root device, any of its embedded devices or any of its services on one of its UPnP-enabled interfaces then the control point can assume that all are no longer available on that UPnP-enabled interface. However, the control point MUST NOT assume that the device is also no longer available on all of its other UPnP-enabled interfaces. As a fallback, if a control point fails to receive notification that a root device, its embedded devices and its services are unavailable on a particular UPnP-enabled interface, the original discovery messages will eventually expire yielding the same effect. Only when *all* original advertisements of a root device, its embedded devices and its services received on a UPnP-enabled interface have expired can a control point assume that they are no longer available on that interface or IP address.

1.2.4 Device Update – NOTIFY with `ssdp:update`

When a new UPnP-enabled interface is added to a multi-homed device, the device MUST increase its `BOOTID.UPNP.ORG` field value, multicast an `ssdp:update` message for each of the root devices, embedded devices and embedded services to all of the existing UPnP-enabled interfaces to announce a change in the `BOOTID.UPNP.ORG` field value, and re-advertise itself on all (existing and new) UPnP-enabled interfaces with the new `BOOTID.UPNP.ORG` field value. Similarly, if a multi-homed device loses connectivity on a UPnP-enabled interface and regains connectivity, or if the IP address on one of the UPnP-enabled interfaces changes, the device MUST increase the `BOOTID.UPNP.ORG` field value, multicast an `ssdp:update` message for each of the root devices, embedded devices and embedded services to all the unaffected UPnP-enabled interfaces to announce a change in the `BOOTID.UPNP.ORG` field value, and re-advertise itself on all (affected and unaffected) UPnP-enabled interfaces with the new `BOOTID.UPNP.ORG` field value. In all cases, the `ssdp:update` message for the root devices MUST be sent as soon as possible. Other `ssdp:update` messages SHOULD be spread over time. However, all `ssdp:update` messages MUST be sent before any announcement messages with the new `BOOTID.UPNP.ORG` field value can be sent.

When `ssdp:update` messages are sent on multiple UPnP-enabled interfaces, the messages MUST contain identical field values except for the `HOST` and `LOCATION` field values. The `HOST` field value of an advertisement MUST be the standard multicast address specified for the protocol (IPv4 or IPv6) used on the interface. The URL specified in the `LOCATION` field value MUST be reachable on the interface on which the advertisement is sent.

```
NOTIFY * HTTP/1.1
HOST: 239.255.255.250:1900
LOCATION: URL for UPnP description for root device
NT: notification type
NTS: ssdp:update
USN: composite identifier for the advertisement
BOOTID.UPNP.ORG: BOOTID value that the device has used in its previous announcements
CONFIGID.UPNP.ORG: number used for caching description information
NEXTBOOTID.UPNP.ORG: new BOOTID value that the device will use in subsequent announcements
SEARCHPORT.UPNP.ORG: number identifies port on which device responds to unicast M-SEARCH
```

Note: No body is present for messages with method `NOTIFY`, but note that the message MUST have a blank line following the last header field.

The TTL for the IP packet SHOULD default to 2 and SHOULD be configurable.

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive.

All field values are case sensitive except where noted.

Request line

Must be "NOTIFY * HTTP/1.1"

NOTIFY

Method for sending notifications and events.

*

Message applies generally and not to a specific resource. MUST be *.

HTTP/1.1

HTTP version.

Header fields

HOST

REQUIRED. Field value contains multicast address and port reserved for SSDP by Internet Assigned Numbers Authority (IANA). MUST be [239.255.255.250:1900](#). If the port number (":1900") is omitted, the receiver MUST assume the default SSDP port number of 1900.

LOCATION

REQUIRED. Field value MUST be the same as the LOCATION field value that has been sent in previous SSDP messages. Single absolute URL (see RFC 3986).

NT

REQUIRED. Field value contains Notification Type. (See list of required field values for the NT header field in NOTIFY with `ssdp:alive` above.) Single URI.

NTS

REQUIRED. Field value contains Notification Sub Type. MUST be [ssdp:update](#). Single URI.

USN

REQUIRED. Field value contains Unique Service Name. (See list of required field values for the USN header field in NOTIFY with `ssdp:alive` above.) Single URI.

BOOTID.UPNP.ORG

REQUIRED. As defined in section 1.2, and 1.2.2, Field value MUST be the same as the BOOTID.UPNP.ORG field value that has been sent in previous SSDP messages.

CONFIGID.UPNP.ORG

REQUIRED. As defined in section 1.2, and 1.2.2.

NEXTBOOTID.UPNP.ORG

REQUIRED. Field value contains the new BOOTID.UPNP.ORG field value that the device intends to use in the subsequent device and service announcement messages. Its field value MUST be a non-negative 31-bit integer; ASCII encoded, decimal, without leading zeros (leading zeroes, if present, MUST be ignored by the recipient) and MUST be greater than the field value of the BOOTID.UPNP.ORG header field.

SEARCHPORT.UPNP.ORG

OPTIONAL. As defined in section 1.2, and 1.2.2.

Note: No responses are sent for messages with method NOTIFY.

If a control point with a single UPnP-enabled interface receives an `ssdp:update` message, the NEXTBOOTID.UPNP.ORG field value replaces the BOOTID.UPNP.ORG field value that the control point has previously recorded for the device. It can expect future announcements, search responses, update messages and eventually bye-bye messages from the device to contain the "new" BOOTID.UPNP.ORG field value (that is: the field value of the NEXTBOOTID.UPNP.ORG header field in the received `ssdp:update`

message). The field value in the NEXTBOOTID.UPNP.ORG header field MUST be recorded as the current BOOTID.UPNP.ORG field value of the device which is to be expected on all subsequent SSDP messages.

If a multi-homed control point receives an `ssdp:update` message on its UPnP-enabled interface(s), and the message arrives on the interface(s) that it uses for UPnP communications with the device (such as event subscriptions), it can assume that the device has remained continuously available (including all device state), and that the NEXTBOOTID.UPNP.ORG field value replaces the BOOTID.UPNP.ORG field value that the control point has previously recorded for the device. It can expect future announcements, search responses, update messages and eventually bye-bye messages from the device to contain the “new” BOOTID.UPNP.ORG field value (that is: the field value of the NEXTBOOTID.UPNP.ORG header field in the received `ssdp:update` message). The field value in the NEXTBOOTID.UPNP.ORG header field MUST be recorded as the current BOOTID.UPNP.ORG field value of the device which is to be expected on all subsequent SSDP messages.

If a control point receives an SSDP message with a BOOTID.UPNP.ORG field value different (either higher or lower) from the value that the control point has previously recorded for the device, it can assume that the device has become temporarily unavailable on that interface and has become available again, and any stored state information about the device has become invalid. It MUST treat the device as a newly discovered device.


1.3 Search

When a control point is added to the network, the UPnP discovery protocol allows that control point to search for devices of interest on the network. It does this by multicasting on the reserved address and port (239.255.255.250:1900) a search message with a pattern, or target, equal to a type or identifier for a device or service. Responses from devices contain discovery messages essentially identical to those advertised by newly connected devices; the former are unicast while the latter are multicast. Control points can also send a unicast search message to a known IP address and port 1900 or the port indicated by SEARCHPORT.UPNP.ORG, to verify the existence of UPnP device(s) and service(s) at the IP address. For example, a unicast search may be used to quickly check whether a known UPnP device or service is still available on the network. Multi-homed control points MAY choose to send discovery messages on any, some or all of its UPnP-enabled interfaces.

1.3.1 Search protocols and standards

To search for devices (and be discovered by control points), control points (and devices) use the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

Figure 1-5: Search protocol stack



At the highest layer, search messages contain vendor-specific information, e.g., the control point, device, and service identifiers. Moving down the stack, vendor content is supplemented by information from a UPnP Forum working committee, e.g., device or service types. Messages from the layers above are hosted in UPnP-specific protocols, defined in this document. In turn, search requests are delivered via multicast and unicast SSDP messages defined in this document. Search responses are delivered via a unicast SSDP messages defined in this document. Both kinds of messages are delivered via UDP over IP. For reference, colors in [square brackets] above indicate which protocol defines specific header fields and field values in discovery messages listed below.

1.3.2 Search request with M-SEARCH

When a control point desires to search the network for devices, it MUST send a multicast request with method M-SEARCH in the following format. Control points that know the address of a specific device MAY also use a similar format to send unicast requests with method M-SEARCH.

For multicast M-SEARCH, the message format is defined below. Values in *italics* are placeholders for actual values.

```
M-SEARCH * HTTP/1.1
HOST: 239.255.255.250:1900
MAN: "ssdp:discover"
MX: seconds to delay response
ST: search target
USER-AGENT: OS/version UPnP/1.1 product/version
```

Note: No body is present in requests with method M-SEARCH, but note that the message MUST have a blank line following the last header field.

Note: The TTL for the IP packet SHOULD default to 2 and SHOULD be configurable.

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive.

All field values are case sensitive except where noted.

Request line

Must be "M-SEARCH * HTTP/1.1"

M-SEARCH

Method for search requests.

*
Request applies generally and not to a specific resource. MUST be *.

HTTP/1.1
HTTP version.

Header fields

HOST

REQUIRED. Field value contains the multicast address and port reserved for SSDP by Internet Assigned Numbers Authority (IANA). MUST be [239.255.255.250:1900](#).

MAN

REQUIRED by HTTP Extension Framework. Unlike the NTS and ST field values, the field value of the MAN header field is enclosed in double quotes; it defines the scope (namespace) of the extension. MUST be "[ssdp:discover](#)".

MX

REQUIRED. Field value contains maximum wait time in seconds. MUST be greater than or equal to 1 and SHOULD be less than 5 inclusive. Device responses SHOULD be delayed a random duration between 0 and this many seconds to balance load for the control point when it processes responses. This value MAY be increased if a large number of devices are expected to respond. The MX field value SHOULD NOT be increased to accommodate network characteristics such as latency or propagation delay (for more details, see the explanation below). Specified by UPnP vendor. Integer.

ST

REQUIRED. Field value contains Search Target. MUST be one of the following. (See NT header field in NOTIFY with [ssdp:alive](#) above.) Single URI.

[ssdp:all](#)

Search for all devices and services.

[upnp:rootdevice](#)

Search for root devices only.

uuid:*device-UUID*

Search for a particular device. *device-UUID* specified by UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

urn:[schemas-upnp-org:device:deviceType:ver](#)

Search for any device of this type where *deviceType* and *ver* are defined by the UPnP Forum working committee.

urn:[schemas-upnp-org:service:serviceType:ver](#)

Search for any service of this type where *serviceType* and *ver* are defined by the UPnP Forum working committee.

urn:*domain-name*:[device:deviceType:ver](#)

Search for any device of this type where *domain-name* (a Vendor Domain Name), *deviceType* and *ver* are defined by the UPnP vendor and *ver* specifies the highest supported version of the device type. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

urn:*domain-name*:[service:serviceType:ver](#)

Search for any service of this type. Where *domain-name* (a Vendor Domain Name), *serviceType* and *ver* are defined by the UPnP vendor and *ver* specifies the highest supported version of the service type. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

USER-AGENT

OPTIONAL. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, "USER-AGENT: *unix/5.1* [UPnP/1.1](#) *MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

For unicast M-SEARCH, the message format is defined below. Values in *italics* are placeholders for actual values.

```
M-SEARCH * HTTP/1.1
HOST: hostname:portNumber
MAN: "ssdp:discover"
ST: search target
USER-AGENT: OS/version UPnP/1.1 product/version
```

Note: No body is present in requests with method M-SEARCH, but note that the message MUST have a blank line following the last header field.

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive.

All field values are case sensitive except where noted.

Request line

Must be "[M-SEARCH * HTTP/1.1](#)"

M-SEARCH

Method for search requests.

*

Request applies generally and not to a specific resource. MUST be *.

HTTP/1.1

HTTP version.

Header fields

HOST

REQUIRED. For unicast requests, the field value MUST be the domain name or IP address of the target device and either port 1900 or the SEARCHPORT provided by the target device.

MAN

REQUIRED by HTTP Extension Framework. Unlike the NTS and ST field values, the field value of the MAN header field is enclosed in double quotes; it defines the scope (namespace) of the extension. MUST be "[ssdp:discover](#)".

ST

REQUIRED. Field value contains Search Target. MUST be one of the following. (See NT header field in NOTIFY with [ssdp:alive](#) above.) Single URI.

[ssdp:all](#)

Search for all devices and services.

[upnp:rootdevice](#)

Search for root devices only.

uuid:*device-UUID*

Search for a particular device. *device-UUID* specified by UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

urn:[schemas-upnp-org:device:deviceType:ver](#)

Search for any device of this type where *deviceType* and *ver* are defined by the UPnP Forum working committee.

urn:[schemas-upnp-org:service:serviceType:ver](#)

Search for any service of this type where *serviceType* and *ver* are defined by the UPnP Forum working committee.

urn:[domain-name:device:deviceType:ver](#)

Search for any device of this type where *domain-name* (a Vendor Domain Name), *deviceType* and *ver* are defined by the UPnP vendor and *ver* specifies the highest supported version of the device type. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

urn:[domain-name:service:serviceType:ver](#)

Search for any service of this type where *domain-name* (a Vendor Domain Name), *serviceType* and *ver* are defined by the UPnP vendor and *ver* specifies the highest supported version of the service type. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

USER-AGENT

OPTIONAL. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, "USER-AGENT: *unix/5.1 UPnP/1.1 MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

Due to the unreliable nature of UDP, control points SHOULD send each M-SEARCH message more than once. As a fallback, to guard against the possibility that a device might not receive the M-SEARCH message from a control point, a device SHOULD re-send its advertisements periodically (see `CACHE-CONTROL` header field in `NOTIFY` with `ssdp:alive` above).

For a multicast request, the control point SHOULD wait at least the amount of time specified in the `MX` header field for responses to arrive from devices. The random distribution of responses over the `MX` interval means that a responder MAY send a response at `MX` seconds after receiving the M-SEARCH request. The `MX` field value MAY be adjusted by heuristics at the requester based on, for example, observed number of responders. Network characteristics affecting the propagation of traffic cannot be addressed by increasing the `MX` field value because of the reason cited above. A requester MAY adapt to network characteristics with heuristics based on observed network behavior (the exact heuristics are out of scope). The net effect is that the M-SEARCH request persists at the requester for a period of time exceeding `MX` such that the characteristics of the network are properly accommodated to minimize lost responses.

When a device receives a unicast M-SEARCH, it SHOULD respond within 1 second and it MAY respond sooner. The sender of the unicast request SHOULD wait at least 1 second for the response.

Updated versions of device and service types are REQUIRED to be fully backward compatible with previous versions. Devices MUST respond to M-SEARCH requests for any supported version. For example, if a device implements "`urn:schemas-upnp-org:service:xyz:2`", it MUST respond to search requests for both that type and "`urn:schemas-upnp-org:service:xyz:1`". The response MUST specify the same version as was contained in the search request. If a control point searches for a device or service of a particular version and receives no responses (presumably because no device present on the network supports the specified version), but is willing to operate using a lower version, it MAY repeat the search request specifying the lower version.

1.3.3 Search response

To be found by a network search, a device MUST send a unicast UDP response to the source IP address and port that sent the request to the multicast address. Devices respond if the `ST` header field of the M-SEARCH request is "`ssdp:all`", "`upnp:rootdevice`", "`uuid:`" followed by a UUID that exactly matches the one advertised by the device, or if the M-SEARCH request matches a device type or service type supported by the device. Multi-homed devices MUST send the search response using the same UPnP-enabled interface on which the search request was received. The URL specified in the `LOCATION` field value MUST specify an address that is reachable on that interface.

Devices responding to a multicast M-SEARCH SHOULD wait a random period of time between 0 seconds and the number of seconds specified in the `MX` field value of the search request before responding, in order to avoid flooding the requesting control point with search responses from multiple devices. If the search request results in the need for a multiple part response from the device, those multiple part responses SHOULD be spread at random intervals through the time period from 0 to the number of seconds specified in the `MX` header field. Devices MAY assume an `MX` field value less than that specified in the `MX` header field. If

the [MX](#) header field specifies a field value greater than 5, the device SHOULD assume that it contained the value 5 or less. Devices MUST NOT stop responding to other requests while waiting the random delay before sending a response.

For multicast M-SEARCH requests, if the search request does not contain an [MX](#) header field, the device MUST silently discard and ignore the search request. If the [MX](#) header field specifies a field value greater than 5, the device SHOULD assume that it contained the value 5 or less.

Any device responding to a unicast M-SEARCH SHOULD respond within 1 second.

The URL specified in the LOCATION header field of the [M-SEARCH](#) response MUST be reachable by the control point to which the response is directed.

Responses to [M-SEARCH](#) requests are intentionally parallel to advertisements, and as such, follow the same pattern as listed for [NOTIFY](#) with [ssdp:alive](#) (above) except that instead of the [NT](#) header field there is an [ST](#) header field here. The response MUST be sent in the following format. Values in *italics* are placeholders for actual values.

```
HTTP/1.1 200 OK
CACHE-CONTROL: max-age = seconds until advertisement expires
DATE: when response was generated
EXT:
LOCATION: URL for UPnP description for root device
SERVER: OS/version UPnP/1.1 product/version
ST: search target
USN: composite identifier for the advertisement
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or an update message
CONFIGID.UPNP.ORG: number used for caching description information
SEARCHPORT.UPNP.ORG: number identifies port on which device responds to unicast M-SEARCH
```

Note: No body is present in a response to a request with method [M-SEARCH](#), but note that the message MUST have a blank line following the last header field.

(Note: No need to limit TTL for the IP packet in response to a search request.)

Listed below are details for the header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Response line

Must be "HTTP/1.1 200 OK"

Header fields

CACHE-CONTROL

REQUIRED. Field value MUST have the max-age directive ("max-age=") followed by an integer that specifies the number of seconds the advertisement is valid. After this duration, control points SHOULD assume the device (or service) is no longer available; as long as a control point has received at least one advertisement that is still valid from a root device, any of its embedded devices or any of its services, then the control point can assume that all are available. The number of seconds SHOULD be greater than or equal to 1800 seconds (30 minutes), although exceptions are defined in the text above. Specified by UPnP vendor. Other directives MUST NOT be sent and MUST be ignored when received.

DATE

RECOMMENDED. Field value contains date when response was generated. "rfc1123-date" as defined in RFC 2616.

EXT

REQUIRED for backwards compatibility with UPnP 1.0. (Header field name only; no field value.)

LOCATION

REQUIRED. Field value contains a URL to the UPnP description of the root device. Normally the host portion contains a literal IP address rather than a domain name in unmanaged networks. Specified by UPnP vendor. Single absolute URL (see RFC 3986).

SERVER

REQUIRED. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, "SERVER: *unix/5.1 UPnP/1.1 MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

ST

REQUIRED. Field value contains Search Target. Single URI. The response sent by the device depends on the field value of the ST header field that was sent in the request. In some cases, the device MUST send multiple response messages as follows. If the received ST field value was:

[ssdp:all](#)

Respond $3+2d+k$ times for a root device with d embedded devices and s embedded services but only k distinct service types (see section 1.1.2, "SSDP message header fields" for a definition of each message to be sent). Field value for ST header field MUST be the same as for the NT header field in NOTIFY messages with [ssdp:alive](#). (See above.)

[upnp:rootdevice](#)

Respond once for root device. Must be [upnp:rootdevice](#).

uuid:*device-UUID*

Respond once for each matching device, root or embedded. Must be uuid:*device-UUID* where *device-UUID* is specified by the UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

urn:[schemas-upnp-org;device:deviceType:ver](#)

Respond once for each matching device, root or embedded. MUST be urn:[schemas-upnp-org;device:deviceType:ver](#) where *deviceType* and *ver* are defined by UPnP Forum working committee and *ver* MUST contain the version of the device type contained in the M-SEARCH request.

urn:[schemas-upnp-org;service:serviceType:ver](#)

Respond once for each matching service type. MUST be urn:[schemas-upnp-org;service:serviceType:ver](#) where *serviceType* and *ver* are defined by the UPnP Forum working committee and *ver* MUST contain the version of the service type contained in the M-SEARCH request.

urn:*domain-name*:[device:deviceType:ver](#)

Respond once for each matching device, root or embedded. MUST be urn:*domain-name*:[device:deviceType:ver](#) where *domain-name* (a Vendor Domain Name), *deviceType* and *ver* are defined by the UPnP vendor and *ver* MUST contain the version of the device type from the M-SEARCH request. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

urn:*domain-name*:[service:serviceType:ver](#)

Respond once for each matching service type. MUST be urn:*domain-name*:[service:serviceType:ver](#) where *domain-name* (a Vendor Domain Name), *serviceType* and *ver* are defined by the UPnP vendor and *ver* MUST contain the version of the service type from the M-SEARCH request. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

USN

REQUIRED. Field value contains Unique Service Name. (See list of required field values for the USN header field in NOTIFY with [ssdp:alive](#) above.) Single URI.

BOOTID.UPNP.ORG

REQUIRED. As defined in section 1.2, and 1.2.2.

CONFIGID.UPNP.ORG

OPTIONAL. As defined in section 1.2, and 1.2.2.

SEARCHPORT.UPNP.ORG

OPTIONAL. As defined in section 1.2, and 1.2.2.

If there is an error with the search request (such as an invalid field value in the MAN header field, a missing MX header field, or other malformed content), the device MUST silently discard and ignore the search request; sending of error responses is PROHIBITED due to the possibility of packet storms if many devices send an error response to the same request.

1.4 References

RFC 2141

URN Syntax. Available at: <http://www.ietf.org/rfc/rfc2141.txt>.

RFC 2616

HTTP: Hypertext Transfer Protocol 1.1. Available at: <http://www.ietf.org/rfc/rfc2616.txt>.

RFC 2774

HTTP Extension Framework. Available at: <http://www.ietf.org/rfc/rfc2774.txt>.

RFC 3986

Uniform Resource Identifiers (URI): Generic Syntax. Available at: <http://www.ietf.org/rfc/rfc3986.txt>.

RFC 4340

Datagram Congestion Control Protocol (DCCP). Available at: <http://www.ietf.org/rfc/rfc4340.txt>.

[1]

DCE variant of Universal Unique Identifiers (UUIDs), The Open group, 1997, Available at: <http://www.opengroup.org/onlinepubs/9629399/apdx.htm>.

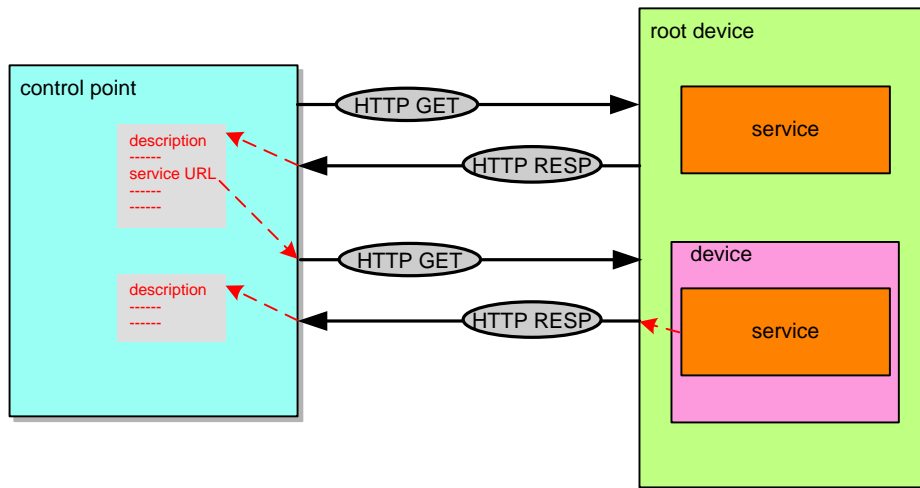
2 Description

[Normative]

Description is Step 2 in UPnP™ networking. Description comes after addressing (Step 0) where devices get a network address, and after discovery (Step 1) where control points find interesting device(s). Description enables control (Step 3) where control points send commands to device(s), eventing (Step 4) where control points listen to state changes in device(s), and presentation (Step 5) where control points may display an html user interface for device(s).

After a control point has discovered a device, the control point still knows very little about the device -- only the information that was in the discovery message, i.e., the device's (or service's) UPnP type, the device's universally-unique identifier, and a URL to the device's UPnP description. For the control point to learn more about the device and its capabilities, or to interact with the device, the control point MUST retrieve a description of the device and its capabilities from the URL provided by the device in the discovery message.

Figure 2-1: Description architecture



The UPnP description for a device is partitioned into two logical parts: a *device description* describing the physical and logical containers, and *service descriptions* describing the capabilities exposed by the device. A UPnP device description includes vendor-specific manufacturer information like the model name and number, serial number, manufacturer name, URLs to vendor-specific Web sites, etc. (details below). For each service included in the device, the device description lists the service type, service name, a URL for a service description, a URL for control, and a URL for eventing. A device description also includes a description of all embedded devices and a URL for presentation of the aggregate. This section explains UPnP device descriptions, and the sections on Control, Eventing, and Presentation explain how URLs for control, eventing, and presentation are used respectively.

Note that a single physical device MAY include multiple logical devices. Multiple logical devices can be modeled as a single root device with embedded devices (and services) or as multiple root devices (perhaps with no embedded devices). In the former case,

there is one UPnP device description for the root device, and that device description contains a description for all embedded devices. In the latter case, there are multiple UPnP device descriptions, one for each root device.

A UPnP device description is written by a UPnP vendor. The description is in XML syntax and is usually based on a standard UPnP Device Template. A UPnP Device Template is produced by a UPnP Forum working committee; they derive the template from the UPnP Device Schema, which was derived from standard constructions in XML. This section explains the format for a UPnP device description, UPnP Device Templates, and the part of the UPnP Device Schema that covers devices.

A UPnP service description includes a list of commands, or *actions*, to which the service responds, and parameters, or *arguments* for each action. A service description also includes a list of variables. These variables model the state of the service at run time, and are described in terms of their data type, range, and event characteristics. This section explains the description of actions, arguments, state variables, and the properties of those variables. The section on Eventing explains event characteristics.

Like a UPnP device description, a UPnP service description is written by a UPnP vendor. The description is in XML syntax and is usually based on a standard UPnP Service Template. A UPnP Service Template is produced by a UPnP Forum working committee; they derived the template from the UPnP Service Schema, augmenting it with human language where necessary. The UPnP Service Schema is derived using the conventions of XML Schema. This section explains the format for a UPnP service description, UPnP Service Templates, typical augmentations in human language, and the part of the UPnP Service Schema that covers services.

UPnP vendors can differentiate their devices by extending services (see section 2.7, "Non-standard vendor extensions and limitations"), including additional UPnP services, or embedding additional devices. When a control point retrieves a particular device's description, these added features are exposed to the control point for control and eventing. The device and service descriptions authoritatively document the implementation of the device.

Retrieving a UPnP device description is simple: the control point issues an HTTP GET request on the URL in the discovery message, and the device returns the device description. Retrieving a UPnP service description is a similar process that uses a URL within the device description. The protocol stack, method, header fields, and body for the response and request are explained in detail below. Description documents MUST be sent using the same IP address on which the HTTP GET request was received.

As long as *at least one* of the discovery advertisements from a root device, any of its embedded devices or any of its services have not expired and *none* of the advertisements have been cancelled, a control point MAY assume that the root device and all its embedded devices and all its services are available. The device and service descriptions MAY be retrieved at any point since the device and service descriptions are static as long as the device and its services are available. If a device cancels *at least one* of its advertisements or if *all* the advertisements expire, a control point SHOULD assume the device and its services are no longer available. If a device needs to change one of these descriptions, it MUST cancel its outstanding advertisements and re-advertise. Consequently, control points SHOULD NOT assume that device and service descriptions are unchanged if a device re-appears on the network, but they can detect whether descriptions changed if a changed CONFIGID.UPNP.ORG field value is present in the announcements.

Like discovery, description plays an important role in the interoperability of devices and control points using different versions of UPnP networking. As explained in section 1, "Discovery", the UPnP Device Architecture is versioned with both a major and a minor version. The major version and minor version are separate integer numbers; they are not to be interpreted or compared as though they were a single decimal number, even though they may appear as such in print. Advances in minor versions MUST be a compatible superset of earlier minor versions of the same major version; therefore device vendors are free to implement standardized devices and services on versions of the architecture with a higher minor version number. Advances in major version are NOT REQUIRED to be supersets of earlier versions and are not guaranteed to be backward compatible. The architecture version of a root device, all its embedded devices and all its services MUST be the same. Version information is communicated in description messages as a backup to the information communicated in discovery messages. This section explains the format of version information in description messages.

Device and service types standardized by UPnP Forum working committees or created by vendors have an integer version. Every later version of a device or service MUST be a fully backwardly compatible superset of the previous version, i.e., compared to earlier versions of the device, it MUST include all mandatory embedded devices and services of the same or later version. The UPnP device or service type remains the same across all versions of a device whereas the device or service version MUST be larger for later versions. Versions of device and service templates MAY have non-integer versions (such as "0.9") during development in the working committee, but this MUST become an integer upon standardization. Devices and services MAY have a version number greater than the major version number of the architecture they are designed for (e.g., "Power:2" MAY be designed to work on UDA version 1.0); there is no direct correlation between the version of a device or service template and the architecture version with which it is designed to work. If a non-backward-compatible version of a device or service is defined, it MUST have a different device or service name to indicate that it is not backwardly compatible (and version numbers of the new type MUST restart at 1).

UPnP device and service types are "building blocks" that MAY be assembled in various combinations. Both standard and vendor-defined device types MAY be embedded in standard device types. Both standard and vendor-defined device types MAY be embedded in vendor-defined device types. Likewise, both standard and vendor-defined service types MAY be embedded in both standard and vendor-defined device types. A control point that is capable of operating with a particular device or service type MUST at least recognize that device or service type even when it is embedded within another device type (standard or vendor-defined) that it does not recognize. For example, if a standard service type "Print:1" is defined, and a standard device type "Printer:1" is defined that contains the "Print:1" service, a control point that wishes to use the "Print:1" service must find and use it whether the service is embedded within a "urn:schemas-upnp-org:device:Printer:1" device or embedded within a vendor-defined "urn:acme-com:device:Printer:1" or "urn:acme-com:device:AcmeMultifunctionPrinter:1" device.

The remainder of this section first explains how devices are described, explaining details of vendor-specific information, embedded devices, and URLs for control, eventing, and presentation. Second, it explains UPnP Device Templates. Third, it explains how services are described, explaining details of actions, arguments, state variables, and properties of those variables.

Then it explains UPnP Service Templates, and the UPnP Service Schema. Finally, this section explains in detail how a control point retrieves device and service descriptions from a device.

2.1 Generic requirements on HTTP usage

This section defines generic requirements on HTTP usage in UPnP Version 1.1. HTTP is the underlying transport for:

- Description (see section 2, “Description”)
- Control (see section 3, “Control”)
- Eventing (see section 4, “Eventing”)
- Presentation (section 5, “Presentation”)

The baseline transport for all devices and control points is RECOMMENDED to be HTTP/1.1 compliant (as defined in RFC 2616) but at least MUST be HTTP/1.0 compliant (as defined in RFC 1945). Vendors are free to implement and Working Committees are free to require for new device classes implementations of more recent versions of HTTP that are backwards compatible with HTTP version 1.0, such as HTTP version 1.1 as defined in RFC 2616. However whatever version is implemented, all REQUIRED components defined by the specified HTTP version MUST be implemented.

If a control point uses an HTTP/1.0 binding on a SOAP request without setting the KeepAlive token, the device MUST close the socket after responding. If a control point uses an HTTP/1.1 binding on a SOAP request, and sets the “Connection:CLOSE” token, the device MUST close the socket after responding.

USER-AGENT header field

Control points can add the USER-AGENT header field to any UPnP-related HTTP request to signal that they support UPnP 1.1. Working Committees MAY require presence of this header on description retrieval, action invocations and event subscriptions for newly defined services.

```
USER-AGENT: OS/version UPnP/1.1 product/version
```

USER-AGENT

OPTIONAL. Specified by UPnP vendor. String. Field value MUST begin with the following “product tokens” (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be UPnP/1.1, and the third token identifies the product using the form *product name/product version*. For example, “USER-AGENT: *unix/5.1 UPnP/1.1 MyProduct/1.0*”. Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

Vendor-defined or working committee-defined HTTP Header fields

HTTP field names defined by vendors or working committees MUST have the following format:

field-name = token “.” domain-name

where the domain-name MUST be a Vendor Domain Name or MUST be “UPNP.ORG” (for working committee defined field names), and in addition MUST satisfy the token format as defined in RFC 2616 section 2.2. Field names are case-insensitive.

HTTP/1.0 Persistent connections

Some implementations of HTTP/1.0 defined what is known as persistent connections. There are many practical uses for this functionality, as it may reduce overhead for a given device by allowing resources to be used more efficiently. However, this functionality for HTTP/1.0 is not officially defined in the specification and classified as experimental. Further, the way it has been experimentally defined is flawed in such a way that it may cause sessions to hang in certain scenarios. This functionality MUST NOT be implemented by any UPnP devices or control points that implement HTTP version 1.0.

HTTP/1.0 HEAD request

Some implementations utilize the HEAD request to try to predetermine the amount of memory required to process a GET request. Some servers may not know that size of the content because it may be dynamic. In such cases, the responses will not contain a CONTENT-LENGTH header field. As such, control points MUST NOT rely on the CONTENT-LENGTH header field being specified for a HEAD response.

HTTP/1.1 General

When a device or control point implements HTTP/1.1, all requirements of HTTP/1.0 MUST be maintained, with the exception of the CONTENT-LENGTH header field, which MUST NOT be specified when doing chunked transfers.

HTTP status codes

Servers MUST return appropriate HTTP status codes for invalid requests. A device or control point MUST use a 4xx HTTP status code for responses that indicate a problem with the format of a request or response. For example, if an HTTP client makes a PUT request to a server that does not implement the PUT method, the server SHOULD return a "405 Method not Allowed" HTTP status code and MUST return a 4xx series HTTP status code. Another example is if an HTTP client makes a request to a server that is malformed HTTP or not well formed XML, the server SHOULD return a "400 Bad Request" HTTP status code and MUST return a 4xx series HTTP status code. While clients are not REQUIRED to understand specific status codes, they MUST understand classes of status codes. For example, a 4xx series HTTP status code signifies an improper request, whereas a 5xx series HTTP status code signifies a processing error for a valid request.

HTTP/1.1 and HTTP/1.0 compatibility

Devices and control points that implement HTTP/1.1 MUST be able to interoperate with HTTP/1.0 control points and devices. Care MUST be taken when devices and control points process requests, such that the response generated is compatible with the HTTP version specified in the request. For example, if an HTTP/1.0 request is made, the device or control point MUST NOT return an HTTP/1.1 chunked response.

HTTP/1.1 HOST header field and use of the HOST header field with HTTP/1.0

The 'HOST' header field MUST be specified in all requests, because HTTP/1.1 allows support for virtual domains, which rely on this header field to determine the target destination.

The HOST header field MUST also be included in HTTP/1.0 requests, for backwards compatibility with UPnP 1.0, which REQUIRES the HOST header field to be present without explicitly mentioning a HTTP version.

HTTP/1.1 EXPECT: 100-Continue

Servers MAY send a "100-Continue" HTTP status code to let the client know that the header fields received have been processed. If a client will rely on this status response before sending the body, it MUST send the "EXPECT: 100-Continue" header field in the request. If a server received this header field in the request, it MUST NOT wait for the request body before sending the continue response. However, a client MUST be prepared to handle cases when the "EXPECT: 100-Continue" header field is not sent, but a "100-Continue" HTTP status code is still received from the server.

HTTP/1.1 Chunked Encoding

Devices and control points that advertise support for HTTP/1.1 MUST have support for decoding chunked encoded messages. Chunked encoded messages MAY contain Chunk-Extensions, which are delineated with a ';'. Extensions that are not recognized MUST be ignored, which includes the absence of an extension, but the presence of the delineator.

Chunked encoding also allows responses and requests to include trailer fields, which are header fields that follow the body. Devices and control points MUST only send trailer fields if the request contained the 'TE' header field (indicates trailer processing is supported), or if the trailer fields in the response only contain OPTIONAL metadata that can be safely ignored.

Before a control point uses chunked encoding to make a request to a device, it MUST check to ensure that the device is an HTTP/1.1 device. Devices MAY use different HTTP engines (that support different versions) for description, control, eventing and presentation. Therefore, to correctly identify which HTTP version is used for processing control requests, a HEAD request MAY be issued to the corresponding control URL.

HTTP/1.1 Persistent Connections

Persistent connections is the default behavior defined by HTTP/1.1. It is strongly RECOMMENDED that this behavior be maintained, as it may be beneficial in many scenarios, as it allows for resources to be utilized more efficiently. Support for Pipelined request handling is also RECOMMENDED if persistent connections are supported.

If a server responds with a "CONNECTION: close" header line, it MUST close the session after responding. Similarly if a client specifies "CONNECTION: close" in the request, the server MUST also close the session after responding.

When Requests are pipelined to a server, the server MUST answer the requests in the order that they are received. Clients MUST also be prepared to retry connections if pipelining fails, for example, if the server does not support them.

HTTP/1.1 Redirect restrictions

HTTP/1.1 defines OPTIONAL support for redirecting an HTTP request. UPnP 1.1 devices MAY redirect a request, although this is NOT RECOMMENDED. If a UPnP 1.1 device redirects a request, it MUST respond with a "307 Temporary Redirect" HTTP status code (see also RFC 2616). UPnP 1.1 devices MUST NOT return any other HTTP/1.1 redirect options. Control points MUST implement HTTP/1.1 redirect and SHOULD redirect the request upon receiving a "307 Temporary Redirect" HTTP status code (see also RFC 2616).

2.2 Generic requirements on XML usage

XML namespace prefixes do not have to be the specific strings that are used in the examples in this specification. They can be any value that obeys the rules of the general XML namespace mechanism as outlined in the *Namespaces in XML* specification.

Devices MUST accept requests that use other legal XML namespace prefixes.

If an XML element has no value (i.e. it contains the empty string), it is valid to combine the opening and closing XML tags (e.g., "<actionname/>" instead of "<actionname></actionname>").

2.3 Device description

The UPnP description for a device contains several pieces of vendor-specific information, definitions of all embedded devices, URL for presentation of the device, and listings for all services, including URLs for control and eventing. In addition to defining non-standard devices (which MAY contain both vendor-defined and standard embedded devices and services), UPnP vendors MAY add embedded devices and services to standard devices. To illustrate these, below is a listing with placeholders (in *italics*) for actual elements and values. Some of these placeholders would be specified by a UPnP Forum working committee (colored *red*) or by a UPnP vendor (colored *purple*). For a non-standard device, all of these placeholders would be specified by a UPnP vendor. Elements defined by the UPnP Device Architecture are colored *green*. Immediately following the listing is a detailed explanation of the elements, attributes, and values.

```
<?xml version="1.0"?>
<root xmlns="urn:schemas-upnp-org:device-1-0"
  <u>configId</u>="configuration number">
  <u>specVersion</u>
    <u>major</u>1</u>major>
    <u>minor</u>1</u>minor>
  </u>specVersion>
  <u>device</u>
    <u>deviceType</u>urn:schemas-upnp-org:device:deviceType:v</u>deviceType>
    <u>friendlyName</u>short user-friendly title</u>friendlyName>
    <u>manufacturer</u>manufacturer name</u>manufacturer>
    <u>manufacturerURL</u>URL to manufacturer site</u>manufacturerURL>
    <u>modelDescription</u>long user-friendly title</u>modelDescription>
    <u>modelName</u>model name</u>modelName>
    <u>modelNumber</u>model number</u>modelNumber>
    <u>modelURL</u>URL to model site</u>modelURL>
    <u>serialNumber</u>manufacturer's serial number</u>serialNumber>
    <u>UDN</u>uuid:UUID</u>UDN>
    <u>UPC</u>Universal Product Code</u>UPC>
    <u>iconList</u>
      <u>icon</u>
        <u>mimetype</u>image/format</u>mimetype>
        <u>width</u>horizontal pixels</u>width>
        <u>height</u>vertical pixels</u>height>
        <u>depth</u>color depth</u>depth>
        <u>url</u>URL to icon</u>url>
      </u>icon>
      <!-- XML to declare other icons, if any, go here -->
    </u>iconList>
    <u>serviceList</u>
      <u>service</u>
        <u>serviceType</u>urn:schemas-upnp-org:service:serviceType:v</u>serviceType>
        <u>serviceId</u>urn:upnp-org:serviceId:serviceID</u>serviceID>
        <u>SCPURL</u>URL to service description</u>SCPURL>
        <u>controlURL</u>URL for control</u>controlURL>
        <u>eventSubURL</u>URL for eventing</u>eventSubURL>
```

```

</service>
<!-- Declarations for other services defined by a UPnP Forum working committee
      (if any) go here -->
<!-- Declarations for other services added by UPnP vendor (if any) go here -->
</serviceList>
<deviceList>
<!-- Description of embedded devices defined by a UPnP Forum working committee
      (if any) go here -->
<!-- Description of embedded devices added by UPnP vendor (if any) go here -->
</deviceList>
<presentationURL>URL for presentation</presentationURL>
</device>
</root>

```

Listed below are details for each of the elements, attributes, and values appearing in the listing above. All elements and attributes are case sensitive; HTTP specifies case sensitivity for URLs; other values are not case sensitive except where noted. The order of elements is significant. Except where noted: REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once. Note that some implementations MAY strictly enforce the length limits for various elements noted below, and therefore working committees are advised to heed all limits specified.

<?xml>

REQUIRED for all XML documents. Case sensitive.

<root>

REQUIRED. MUST have "urn:[schemas-upnp-org:device-1-0](#)" as the value for the xmlns attribute; this references the UPnP Device Schema (described below). Case sensitive. Has the following attribute:

configId

REQUIRED. Specifies the configuration number to which the device description belongs. See section 1, "Discovery" for further definition and usage of the configuration number.

Contains all other elements describing the root device, i.e., contains the following child elements:

<specVersion>

REQUIRED. In device templates, defines the lowest version of the architecture on which the device can be implemented. In actual UPnP devices, defines the architecture on which the device is implemented. Contains the following sub elements:

<major>

REQUIRED. Major version of the UPnP Device Architecture. MUST be 1 for devices implemented on a UPnP 1.1 architecture.

<minor>

REQUIRED. Minor version of the UPnP Device Architecture. MUST be 1 for devices implemented on a UPnP 1.1 architecture. MUST accurately reflect the version number of the UPnP Device Architecture supported by the device. Control points MUST be prepared to accept a higher version number than the control point itself implements.

<URLBase>

Use of URLBase is deprecated in UPnP 1.1; UPnP 1.1 devices MUST NOT include URLBase in their description documents. For full definition of URLBase, see the UPnP 1.0 specification.

<device>

REQUIRED. Contains the following sub elements:

<deviceType>

REQUIRED. UPnP device type. Single URI.

- For standard devices defined by a UPnP Forum working committee, MUST begin with "urn:[schemas-upnp-org:device](#):" followed by the standardized device type suffix, a colon, and an integer device version i.e. urn:[schemas-upnp-org:device:deviceType:ver](#)". The highest supported version of the device type MUST be specified.
- For non-standard devices specified by UPnP vendors, MUST begin with "urn:", followed by a Vendor Domain Name, followed by ":", followed by a device type suffix, colon, and an integer version, i.e., "urn:[domain-name:device:deviceType:ver](#)". Period characters in the Vendor Domain

Name MUST be replaced with hyphens in accordance with RFC 2141. The highest supported version of the device type MUST be specified.

The device type suffix defined by a UPnP Forum working committee or specified by a UPnP vendor MUST be <= 64 chars, not counting the version suffix and separating colon.

<friendlyName>
REQUIRED. Short description for end user. MAY be localized (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. String. SHOULD be < 64 characters.

<manufacturer>
REQUIRED. Manufacturer's name. MAY be localized (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. String. SHOULD be < 64 characters.

<manufacturerURL>
OPTIONAL. Web site for Manufacturer. MAY have a different value depending on language requested (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. Single URL.

<modelDescription>
RECOMMENDED. Long description for end user. MAY be localized (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. String. SHOULD be < 128 characters.

<modelName>
REQUIRED. Model name. MAY be localized (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. String. SHOULD be < 32 characters.

<modelNameNumber>
RECOMMENDED. Model number. MAY be localized (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. String. SHOULD be < 32 characters.

<modelURL>
OPTIONAL. Web site for model. MAY have a different value depending on language requested (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. Single URL.

<serialNumber>
RECOMMENDED. Serial number. MAY be localized (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Specified by UPnP vendor. String. SHOULD be < 64 characters.

<UDN>
REQUIRED. Unique Device Name. Universally-unique identifier for the device, whether root or embedded. MUST be the same over time for a specific device instance (i.e., MUST survive reboots). MUST match the field value of the NT header field in device discovery messages. MUST match the prefix of the USN header field in all discovery messages. (Section 1, "Discovery" explains the NT and USN header fields.) MUST begin with "uuid:" followed by a UUID suffix specified by a UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

<UPC>
OPTIONAL. Universal Product Code. 12-digit, all-numeric code that identifies the consumer package. Managed by the Uniform Code Council. Specified by UPnP vendor. Single UPC.

<iconList>
REQUIRED if and only if device has one or more icons. Specified by UPnP vendor. Contains the following sub elements:

- <icon>
RECOMMENDED. Icon to depict device in a control point UI. MAY have a different value depending on language requested (see ACCEPT-LANGUAGE and CONTENT-LANGUAGE header fields). Icon sizes to support are vendor-specific. Contains the following sub elements:
 - <mimetype>
REQUIRED. Icon's MIME type (see RFC 2045, 2046, and 2387). Single MIME image type. At least one icon SHOULD be of type "image/png" (Portable Network Graphics, see IETF RFC 2083).
 - <width>
REQUIRED. Horizontal dimension of icon in pixels. Integer.
 - <height>
REQUIRED. Vertical dimension of icon in pixels. Integer.
 - <depth>
REQUIRED. Number of color bits per pixel. Integer.
 - <url>
REQUIRED. Pointer to icon image. (XML does not support direct embedding of binary data. See note below.) Retrieved via HTTP. MUST be relative to the URL at which the device description is located in accordance with section 5 of RFC 3986. Specified by UPnP vendor. Single URL.

<serviceList>

OPTIONAL. Contains the following sub elements:

<service>

OPTIONAL. Repeated once for each service defined by a UPnP Forum working committee. If UPnP vendor differentiates device by adding additional, standard UPnP services, repeated once for each additional service. Contains the following sub elements:

<serviceType>

Required. UPnP service type. MUST NOT contain a hash character (#, 23 Hex in UTF-8). Single URI.

- For standard service types defined by a UPnP Forum working committee, MUST begin with "urn:schemas-upnp-org:service:" followed by the standardized service type suffix, colon, and an integer service version i.e. urn:schemas-upnp-org:device:serviceType:ver. The highest supported version of the service type MUST be specified.
- For non-standard service types specified by UPnP vendors, MUST begin with "urn:", followed by a Vendor Domain Name, followed by ":service:", followed by a service type suffix, colon, and an integer service version, i.e., "urn:domain-name:service:serviceType:ver". Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141. The highest supported version of the service type MUST be specified.

The service type suffix defined by a UPnP Forum working committee or specified by a UPnP vendor MUST be <= 64 characters, not counting the version suffix and separating colon.

<serviceId>

REQUIRED. Service identifier. MUST be unique within this device description. Single URI.

- For standard services defined by a UPnP Forum working committee, MUST begin with "urn:upnp-org:serviceId:" followed by a service ID suffix i.e. urn:upnp-org:serviceId:serviceID. If this instance of the specified service type (i.e. the <serviceType> element above) corresponds to one of the services defined by the specified device type (i.e. the <deviceType> element above), then the value of the service ID suffix MUST be the service ID defined by the device type for this instance of the service. Otherwise, the value of the service ID suffix is vendor defined. (Note that upnp-org is used instead of schemas-upnp-org in this case because an XML schema is not defined for each service ID.)
- For non-standard services specified by UPnP vendors, MUST begin with "urn:", followed by a Vendor Domain Name, followed by ":serviceId:", followed by a service ID suffix, i.e., "urn:domain-name:serviceId:serviceID". If this instance of the specified service type (i.e. the <serviceType> element above) corresponds to one of the services defined by the specified device type (i.e. the <deviceType> element above), then the value of the service ID suffix MUST be the service ID defined by the device type for this instance of the service. Period characters in the Vendor Domain Name MUST be replaced with hyphens in accordance with RFC 2141.

The service ID suffix defined by a UPnP Forum working committee or specified by a UPnP vendor MUST be <= 64 characters.

<SCPDURL>

REQUIRED. URL for service description. (See section 2.5, "Service description" below.) MUST be relative to the URL at which the device description is located in accordance with section 5 of RFC 3986. Specified by UPnP vendor. Single URL.

<controlURL>

REQUIRED. URL for control (see section 3, "Control"). MUST be relative to the URL at which the device description is located in accordance with section 5 of RFC 3986. Specified by UPnP vendor. Single URL.

<eventSubURL>

REQUIRED. URL for eventing (see section 4, "Eventing"). MUST be relative to the URL at which the device description is located in accordance with section 5 of RFC 3986. MUST be unique within the device; any two services MUST NOT have the same URL for eventing. If the service has no evented variables, this element MUST be

present but MUST be empty (i.e., <eventSubURL></eventSubURL>.) Specified by UPnP vendor. Single URL.

<deviceList>

REQUIRED if and only if root device has embedded devices. Contains the following sub elements:

<device>

REQUIRED. Repeat once for each embedded device defined by a UPnP Forum working committee. If UPnP vendor differentiates device by embedding additional UPnP devices, repeat once for each embedded device. Contains sub elements as defined above for root sub element device.

<presentationURL>

RECOMMENDED. URL to presentation for device (see section 5, "Presentation"). MUST be relative to the URL at which the device description is located in accordance with the rules specified in section 5 of RFC 3986. Specified by UPnP vendor. Single URL.

Control points SHOULD recognize and interoperate with services using serviceld values other than the value defined by the device type. If multiple instances of a service exist, control points SHOULD by default (unless directed otherwise by user action) use the service instance associated with the serviceld value defined by the device type. If none of the instances of the service have the serviceld value defined by the device type, the control point may use any service instance. When only one instance of the service exists, control points SHOULD use that instance even if the serviceld value does not match that defined by the device type.

For future extensibility and according to the requirements in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values.

Subject to the constraints defined in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", control points and devices MUST ignore any XML comments or XML processing instructions embedded in UPnP device and service descriptions that they do not understand. UPnP device descriptions MUST be encoded using UTF-8.

When the value of any text element or attribute contains one or more characters reserved as markup (such as ampersand ("&") or less than ("<")), the text MUST be escaped in accordance with the provisions of section 2.4 of the XML specification and each such character replaced with the equivalent numeric representation or string (such as "&" or "<"). Such characters appearing in URLs MAY also be percent-encoded in accordance with the URL percent-encoding rules specified in sections 2.1 and 2.4 of RFC 3986.

XML does not support directly embedding binary data, e.g., icons in UPnP device descriptions. Binary data MAY be converted into text (and thereby embedded into XML) using an XML data type of either bin.base64 (a MIME-style base 64 encoding for binary data) or bin.hex (hexadecimal digits represent octets). Alternatively, the data can be passed indirectly, as it were, by embedding a URL in the XML and transferring the data in response to a separate HTTP request; the icon(s) in UPnP device descriptions are transferred in this latter manner.

If any icons are included, at least one SHOULD be in the Portable Network Graphics (PNG) format defined in RFC 2083, indicated by the MIME type "image/png", and not use progressive encoding. NO specific icon sizes are RECOMMENDED due to the wide variety preferred by various control points; control point vendors are encouraged to publish implementation guidelines.

The use of URLBase element is deprecated by this specification. UPnP 1.1 devices MUST NOT include URLBase in their description documents. To ensure interoperability with UPnP 1.0 devices, control points MUST be able to process URLBase if it is specified and use it for resolving relative URLs that appear elsewhere in the description. If relative URLs are included in the device description, control points MUST resolve them into absolute URLs in accordance with section 5 of RFC 3986, using either URLBase (if specified) or the location from which the device description was retrieved as the base URL, before using these URLs for their respective purposes.

Note that in version 1.0 of the UPnP Device Architecture, the serviceList element was REQUIRED, and it was REQUIRED to contain at least one service element. These requirements were subsequently rescinded to accommodate the InternetGatewayDevice:1 and Basic:1 device types. If the device has no services, the serviceList element MAY be omitted entirely, or it MAY be present but contain no service elements.

2.4 UPnP Device Template

The listing above also illustrates the relationship between a UPnP device description and a UPnP Device Template. As explained above, the UPnP device description is written by a UPnP vendor, in XML, following a UPnP Device Template. A UPnP Device Template is produced by a UPnP Forum working committee as a means to standardize devices.

By appropriate specification of placeholders, the listing above can be either a UPnP Device Template or a UPnP device description. Recall that some placeholders would be defined by a UPnP Forum working committee (colored *red*), i.e., the UPnP device type identifier, REQUIRED UPnP services, and REQUIRED UPnP embedded devices (if any). If these were defined, the listing would be a UPnP Device Template, codifying the standard for this type of device. UPnP Device Templates are one of the key deliverables from UPnP Forum working committees.

Taking this another step further, the remaining placeholders in the listing above would be specified by a UPnP vendor (colored *purple*), i.e., vendor-specific information. If these placeholders were specified (as well as the others), the listing would be a UPnP device description, suitable to be delivered to a control point to enable control, eventing, and presentation.

Put another way, the UPnP Device Template defines the overall type of device, and each UPnP device description instantiates that template with vendor-specific information. The first is created by a UPnP Forum working committee; the latter, by a UPnP vendor.

2.5 Service description

The UPnP description for a service defines actions and their arguments, and state variables and their data type, range, and event characteristics.

Each service MUST have zero or more actions. Each action MUST have zero or more arguments. Each argument is designated as either an input or an output argument. Input arguments MUST be listed first. If an action has one or more output arguments, the

first output argument MAY be marked as a return value. Each argument MUST correspond to one of the <stateVariable> elements in the <serviceStateTable> in the SCPD.

Each service MUST have one or more state variables.

In addition to defining non-standard services, UPnP vendors MAY add actions and services to standard devices, and MAY embed standard services and devices in non-standard devices.

To illustrate these points, below is a listing with placeholders (in *italics*) for actual elements and values. For a standard UPnP service, some of these placeholders would be defined by a UPnP Forum working committee (colored *red*) or specified by a UPnP vendor (*purple*). For a non-standard service, all of these placeholders would be specified by a UPnP vendor. Elements defined by the UPnP Device Architecture are colored *green*. Immediately following the listing is a detailed explanation of the elements, attributes, and values.

```
<?xml version="1.0"?>
<scpd
  xmlns="urn:schemas-upnp-org:service-1-0"
  xmlns:dt1="urn:domain-name:more-datatypes"
  <!-- Declarations for other namespaces added by UPnP Forum working committee (if any) go here -->
  <!-- The value of the attribute must remain as defined by the UPnP Forum working committee. -->
  xmlns:dt2="urn:domain-name:vendor-datatypes"
  <!-- Declarations for other namespaces added by UPnP vendor (if any) go here -->
  <!-- Vendors must change the URN's domain-name to a Vendor Domain Name -->
  <!-- Vendors must change vendor-datatypes to reference a vendor-defined namespace -->
  configId="configuration number">
  <specVersion>
    <major>1</major>
    <minor>1</minor>
  </specVersion>
  <actionList>
    <action>
      <name>actionName</name>
      <argumentList>
        <argument>
          <name>argumentNameIn1</name>
          <direction>in</direction>
          <relatedStateVariable>stateVariableName</relatedStateVariable>
        </argument>
        <!-- Declarations for other IN arguments defined by UPnP Forum working Committee (if any) go here -->
        <argument>
          <name>argumentNameOut1</name>
          <direction>out</direction>
          <retval/>
          <relatedStateVariable>stateVariableName</relatedStateVariable>
        </argument>
        <argument>
          <name>argumentNameOut2</name>
          <direction>out</direction>
          <relatedStateVariable>stateVariableName</relatedStateVariable>
        </argument>
        <!-- Declarations for other OUT arguments defined by UPnP Forum working committee (if any) go here -->
      </argumentList>
    </action>
    <!-- Declarations for other actions defined by UPnP Forum working committee (if any) go here -->
```



```

    <!-- Declarations for other actions added by UPnP vendor (if any) go here -->
  </actionList>
  <serviceStateTable>
    <stateVariable sendEvents="yes"|"no" multicast="yes"|"no">
      <name>variableName</name>
      <dataType>basic data type</dataType>
      <defaultValue>default value</defaultValue>
      <allowedValueRange>
        <minimum>minimum value</minimum>
        <maximum>maximum value</maximum>
        <step>increment value</step>
      </allowedValueRange>
    </stateVariable>
    <stateVariable sendEvents="yes"|"no" multicast="yes"|"no">
      <name>variableName</name>
      <dataType type="dt1:variable data type">string</dataType>
      <defaultValue>default value</defaultValue>
      <allowedValueList>
        <allowedValue>enumerated value</allowedValue>
        <!-- Other allowed values defined by UPnP Forum working committee
             (if any) go here -->
        <!-- Other allowed values defined by vendor (if any) go here -->
      </allowedValueList>
    </stateVariable>
    <stateVariable sendEvents="yes"|"no" multicast="yes"|"no">
      <name>variableName</name>
      <dataType type="dt2:vendor data type">string</dataType>
      <defaultValue>default value</defaultValue>
    </stateVariable>
    <!-- Declarations for other state variables defined by UPnP Forum working committee
         (if any) go here -->
    <!-- Declarations for other state variables added by UPnP vendor (if any) go here -->
  </serviceStateTable>
</scpd>

```

Listed below are details for each of the elements, attributes, and values appearing in the listing above. All elements and attributes (including action, argument, and state variable names) are case sensitive; values are not case sensitive except where noted. Except where noted, REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once.

<?xml>

REQUIRED for all XML documents. Case sensitive.

<scpd>

REQUIRED. MUST have "urn:schemas-upnp-org:service-1-0" as the value for the xmlns attribute; this references the UPnP Service Schema (explained below). Case sensitive. Has the following attribute:

configId

REQUIRED. Specifies the configuration number to which the service description belongs. See section 1, "Discovery" for further definition and usage of the configuration number.

Contains all other elements describing the service, i.e., contains the following sub elements:

<specVersion>

REQUIRED. In service templates, defines the lowest version of the architecture on which the service can be implemented. In actual UPnP services, defines the architecture on which the service is implemented. Contains the following sub elements:

<major>

REQUIRED. Major version of the UPnP Device Architecture. MUST be 1 for services implemented on a UPnP 1.1 architecture.

<minor>

REQUIRED. Minor version of the UPnP Device Architecture. MUST be 1 for services implemented on a UPnP 1.1 architecture. MUST accurately reflect the version number of the UPnP Device Architecture supported by the device. Control points MUST be prepared to accept a higher version number than the control point itself implements.

<actionList>
 REQUIRED if and only if the service has actions. (Each service MAY have >= 0 actions.) Contains the following sub element(s):

<action>
 REQUIRED. Repeat once for each action defined by a UPnP Forum working committee. If UPnP vendor differentiates service by adding additional actions, repeat once for each additional action. Contains the following sub elements:

<name>
 REQUIRED. Name of action. MUST NOT contain a hyphen character ("-"), 2D Hex in UTF-8) nor a hash character ("#", 23 Hex in UTF-8). Case sensitive. First character MUST be a USASCII letter ("A"- "Z", "a"- "z"), USASCII digit ("0"- "9"), an underscore ("_"), or a non-experimental Unicode letter or digit greater than U+007F. Succeeding characters MUST be a USASCII letter ("A"- "Z", "a"- "z"), USASCII digit ("0"- "9"), an underscore ("_"), a period ("."), a Unicode combiningchar, an extender, or a non-experimental Unicode letter or digit greater than U+007F. The first three letters MUST NOT be "XML" in any combination of case.

- For standard actions defined by a UPnP Forum working committee, MUST NOT begin with "X" nor "A".
- For non-standard actions specified by a UPnP vendor and added to a standard service, MUST begin with "X", followed by a Vendor Domain Name, followed by the underscore character ("_"), followed by the vendor-assigned action name. The vendor-assigned action name must comply with the syntax rules defined above.

String. SHOULD be < 32 characters.

<argumentList>
 REQUIRED if and only if parameters are defined for action. (Each action MAY have >= 0 parameters.) Contains the following sub element(s):

<argument>
 REQUIRED. Repeat once for each parameter. UPnP vendors MUST NOT add vendor-defined arguments to actions defined by a UPnP Forum working committees. Contains the following sub elements:

<name>
 REQUIRED. Name of formal parameter. The name SHOULD be chosen to reflect the semantic use of the argument. MUST NOT contain a hyphen character ("-"), 2D Hex in UTF-8). First character MUST be a USASCII letter ("A"- "Z", "a"- "z"), USASCII digit ("0"- "9"), an underscore ("_"), or a non-experimental Unicode letter or digit greater than U+007F. Succeeding characters MUST be a USASCII letter ("A"- "Z", "a"- "z"), USASCII digit ("0"- "9"), an underscore ("_"), a period ("."), a Unicode combiningchar, an extender, or a non-experimental Unicode letter or digit greater than U+007F. The first three letters MUST NOT be "XML" in any combination of case. String. Case sensitive. SHOULD be < 32 characters.

<direction>
 REQUIRED. Defines whether argument is an input or output parameter. MUST be either "in" or "out" and not both. All input arguments MUST be listed before any output arguments.

<retval>
 OPTIONAL. Identifies at most one output argument as the return value. If included, MUST be included as a subelement of the first output argument. (Element only; no value.)

<relatedStateVariable>
 REQUIRED. MUST be the name of a state variable. Case Sensitive. Defines the type of the argument; see further explanation below in this section.

<serviceStateTable>
 REQUIRED. (Each service MUST have => 1 state variables.) Contains the following sub element(s):

<stateVariable>
 REQUIRED. Repeat once for each state variable defined by a UPnP Forum working committee. If UPnP vendor differentiates service by adding additional state variables, repeat once for each additional state variable. Has the following attributes:

sendEvents

OPTIONAL. Defines whether event messages will be generated when the value of this state variable changes. Default value is **yes**. Non-evented state variables MUST set this attribute to **no**.

- For standard state variables defined by a UPnP Forum working committee, the working committee decides whether the variable is evented and the value of the `sendEvents` attribute MUST NOT be altered by a vendor.
- For non-standard state variables specified by a UPnP vendor and added to a standard service, the vendor MAY decide whether the non-standard state variable will be evented or not.

multicast

OPTIONAL. Defines whether event messages will be delivered using multicast eventing. Default value is **no**. If the multicast is set to **yes**, then all events sent for this state variable MUST be unicast AND multicast.

- For standard state variables defined by a UPnP Forum working committee, the working committee decides whether the state variable is multicast and the value of the `multicast` attribute MUST NOT be altered by a vendor.
- For non-standard variables specified by a UPnP vendor and added to a standard service, the vendor may decide whether the non-standard variable will be delivered using multicast eventing.

The `<stateVariable>` element contains the following sub elements:

<name>

REQUIRED. Name of state variable. MUST NOT contain a hyphen character (“-”, 2D Hex in UTF-8). First character MUST be a USASCII letter (“A”-“Z”, “a”-“z”), USASCII digit (“0”-“9”), an underscore (“_”), or a non-experimental Unicode letter or digit greater than U+007F. Succeeding characters MUST be a USASCII letter (“A”-“Z”, “a”-“z”), USASCII digit (“0”-“9”), an underscore (“_”), a period (“.”), a Unicode combiningchar, an extender, or a non-experimental Unicode letter or digit greater than U+007F. The first three letters MUST NOT be “XML” in any combination of case. Case sensitive.

- For standard state variables defined by a UPnP Forum working committee, MUST NOT begin with **X_** nor **A_**.
- For non-standard state variables specified by a UPnP vendor and added to a standard service, MUST begin with **X_**, followed by a Vendor Domain Name, followed by the underscore character (“_”), followed by the vendor-assigned state variable name. The vendor-assigned state variable name must comply with the syntax rules defined above.

String. SHOULD be < 32 characters.

<dataType>

REQUIRED. Same as data types defined by XML Schema, Part 2: Datatypes. Defined by a UPnP Forum working committee for standard state variables; specified by UPnP vendor for extensions. Has an OPTIONAL `type` attribute:

type

OPTIONAL. If the `type` attribute is present, the value of the `<dataType>` element MUST be “string”. The value of the `type` attribute overrides the “string” value; it defines the data type using a fully qualified data type name according to the conventions of XML schema and can refer to XML Schema simple types, service-local complex types and service-local extended simple types. Service-local data types are defined in a corresponding UPnP Service Template or they MAY be vendor-specific.

See also section 2.5.1, “Defining and processing extended data types” and section 2.5.2, “String equivalents of extended data types”.

For example: `<dataType type="xsd:byte">string</dataType>`

For a state variable using an extended data type via the `type` attribute and containing complex data, the `<defaultValue>`, `<allowedValueList>` and `<allowedValueRange>` elements MUST NOT be present. In such case the restrictions for the data type MUST be described in the data type schema provided in the service template document.

The <dataType> element MUST have one of the following values:

- ui1
Unsigned 1 Byte int. Same format as int without leading sign.
- ui2
Unsigned 2 Byte int. Same format as int without leading sign.
- ui4
Unsigned 4 Byte int. Same format as int without leading sign.
- i1
1 Byte int. Same format as int.
- i2
2 Byte int. Same format as int.
- i4
4 Byte int. Same format as int. MUST be between -2147483648 and 2147483647.
- int
Fixed point, integer number. MAY have leading sign. MAY have leading zeros, which SHOULD be ignored by the recipient. (No currency symbol.) (No grouping of digits to the left of the decimal, e.g., no commas.)
- r4
4 Byte float. Same format as float. MUST be between 3.40282347E+38 to 1.17549435E-38.
- r8
8 Byte float. Same format as float. MUST be between -1.79769313486232E308 and -4.94065645841247E-324 for negative values, and between 4.94065645841247E-324 and 1.79769313486232E308 for positive values, i.e., IEEE 64-bit (8-Byte) double.
- number
Same as r8.
- fixed.14.4
Same as r8 but no more than 14 digits to the left of the decimal point and no more than 4 to the right.
- float
Floating point number. Mantissa (left of the decimal) and/or exponent MAY have a leading sign. Mantissa and/or exponent MAY have leading zeros, which SHOULD be ignored by the recipient. Decimal character in mantissa is a period, i.e., whole digits in mantissa separated from fractional digits by period ("."). Mantissa separated from exponent by "E". (No currency symbol.) (No grouping of digits in the mantissa, e.g., no commas.)
- char
Unicode string. One character long.
- string
Unicode string. No limit on length.
- date
Date in a subset of ISO 8601 format without time data.
- dateTime
Date in ISO 8601 format with OPTIONAL time but no time zone.
- dateTime.tz
Date in ISO 8601 format with OPTIONAL time and OPTIONAL time zone.
- time
Time in a subset of ISO 8601 format with no date and no time zone.
- time.tz
Time in a subset of ISO 8601 format with OPTIONAL time zone but no date.
- boolean
"0" for false or "1" for true. The values "true", "yes", "false", or "no" are deprecated and MUST NOT be sent but MUST be accepted when received. When received, the values "true" and "yes" MUST be interpreted as true and the values "false" and "no" MUST be interpreted as false.

bin.base64

MIME-style Base64 encoded binary BLOB. Takes 3 Bytes, splits them into 4 parts, and maps each 6 bit piece to an octet. (3 octets are encoded as 4.) No limit on size.

bin.hex

Hexadecimal digits representing octets. Treats each nibble as a hex digit and encodes as a separate Byte. (1 octet is encoded as 2.) No limit on size.

uri

Universal Resource Identifier.

uuid

Universally Unique ID. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

<defaultValue>

RECOMMENDED. Expected, initial value. Defined by a UPnP Forum working committee or delegated to UPnP vendor. MUST match data type. MUST satisfy <allowedValueList> or <allowedValueRange> constraints. For a state variable using an extended data type via the type attribute and containing complex data, the <defaultValue> element MUST NOT be present.

<allowedValueList>

RECOMMENDED. Enumerates legal string values. PROHIBITED for data types other than string. At most one of the <allowedValueRange> or <allowedValueList> elements MAY be specified. Sub elements are ordered. For a state variable using an extended data type via the type attribute and containing complex data, the <allowedValueList> element MUST NOT be present. Contains the following sub elements:

<allowedValue>

REQUIRED. A legal value for a string variable. Defined by a UPnP Forum working committee for standard state variables; if the UPnP Forum working committee permits it, UPnP vendors MAY add vendor-specific allowed values to standard state variables. Specified by UPnP vendor for extensions. String. MUST be < 32 characters.

<allowedValueRange>

RECOMMENDED. Defines bounds for legal numeric values; defines resolution for numeric values. Defined only for numeric data types (i.e. integers and floats). At most one of the <allowedValueRange> or <allowedValueList> elements MAY be specified. For a state variable using an extended data type via the type attribute and containing complex data, the <allowedValueRange> element MUST NOT be present. Contains the following sub elements which MUST have the same type as the state variable:

<minimum>

REQUIRED. Inclusive lower bound. Defined by a UPnP Forum working committee or delegated to UPnP vendor. Single numeric value. The value of the <minimum> element MUST be less than the value of the <maximum> element. If a working committee has assigned an explicit value for this element, then vendors MUST use that value. Otherwise, vendors MUST choose their own value, but always within the allowed range for the data type of this state variable. If the working committee defines an allowed range for this state variable, then the value MUST be within that allowed range as defined by the <step> value (See below).

<maximum>

REQUIRED. Inclusive upper bound. Defined by a UPnP Forum working committee or delegated to UPnP vendor. Single numeric value. The value of the <maximum> element MUST be greater than the value of the <minimum> element. If a working committee has assigned an explicit value for this element, then vendors MUST use that value. Otherwise, vendors MUST choose their own value, but always within the allowed range for the data type of this state variable. If the working committee defines an allowed range for this state variable, then the value MUST be within that allowed range as defined by the <step> value (See below).

<step>

RECOMMENDED. Defines the set of allowed values permitted for the state variable between the <minimum> and <maximum>. The value of the <step> element divides the inclusive range from <minimum> value to <maximum> value into an integral number of equal parts. Additionally, <maximum> value = <minimum> value + n * <step> value, where n is a positive integer. Defined by a UPnP Forum working committee or delegated to UPnP vendor. If a working committee has assigned an explicit value for this element, then vendors MUST use that value. Otherwise, vendors MAY choose their own value. When the <step> element is omitted and the

data type of the state variable is an integer, the default value of step is 1; otherwise, when step is omitted, the state variable MAY be set to any value within the inclusive range of <minimum> value to <maximum> value. Single numeric value. Note: care must be taken when dealing with floating point values so that conversions and/or rounding errors do not cause inaccurate comparison operations.

The <relatedStateVariable> element of an <argument> element definition MUST be the name of a state variable defined in the same service description. The <relatedStateVariable> element defines the *data type* of the argument; there is not necessarily any semantic relationship between an argument and the related state variable used to define its type. The <relatedStateVariable> element MUST specify the name of a state variable in the service state table which has the same data type, allowed value list, or allowed value range as the argument. If no state variable exists with an appropriate definition, the working committee (or vendor) MUST define an additional state variable for that purpose; state variables which are defined solely for the purpose of describing the type of an argument MUST have a name that includes the prefix "A_ARG_TYPE".

The <allowedValueList> and <allowedValueRange> elements MAY be used to indicate optional device capabilities. Working committees MAY REQUIRE all values in the list or range to be supported by all vendors (no options), REQUIRE a minimum subset with additional values being OPTIONAL, or allow vendors to entirely decide which portions of the list or range to support. Vendors MAY add additional, vendor-specific values to the allowed value list by using the "X" prefix on the vendor-defined allowed values, if permitted by working committees. However, it should be noted that greater flexibility in OPTIONAL capabilities reduces the number of values that control points can depend on to be present, with corresponding impacts on interoperability. If device capabilities are expected to change during device operation, working committees SHOULD define evented state variables or separate actions to detect device capabilities rather than embedding capabilities information in the service description, because the latter requires cancellation of advertisements and readvertisement each time the service description document is changed. If the service description is used to convey capabilities information, the device MUST omit from the service description any OPTIONAL elements (actions, allowed values, etc.) that are not implemented.

For a state variable using an extended data type via the `type` attribute and containing complex data, the <defaultValue>, <allowedValueList> and <allowedValueRange> elements MUST NOT be present. In such case the restrictions for the data type MUST be described in the data type schema typically provided in the service template document.

2.5.1 Defining and processing extended data types

The optional `type` attribute of the <dataType> element as defined in section 2.5, "Service description" above allows a service description document (SCPD) to include extended data types (defined by the UPnP technical committee, a UPnP working committee or vendor-specific data types) that have more structure and expression than the existing XSD data types. As mentioned above, this `type` attribute can only be applied when the base data type is of type string. The value attached to the `type` attribute refers to a data type from a separate schema defined outside this document.

As a first RECOMMENDATION on the use of extended data types, if UPnP actions only have simple arguments, these SHOULD be declared using UPnP defined data types, instead of XML schema simple types. This enables use of such UPnP actions by UPnP 1.0 stacks that are not XML-schema enabled.

As a further RECOMMENDATION on extended data types that MAY be defined, arrays SHOULD be declared by using a sequence with an element type of which the number of occurrences is restricted. For example, if an array-type “myArrayType” of 50 long integers needs to be declared, this could be the corresponding schema fragment:

```
<xsd:complexType name="myArrayType">
  <xsd:sequence>
    <xsd:element name="x" type="xsd:long" minOccurs="50" maxOccurs="50"/>
  </xsd:sequence>
</xsd:complexType>
```

References to this type (as with any XML namespace) can be made in one of two ways. The first option is a fully qualified namespace reference in the `type` attribute alone. In this case the namespace reference in the `type` attribute MUST not only refer to the schema, but also to the type within that schema.

```
<scpd xmlns="urn:schemas-upnp-org:service-1-0"
  configId="configuration number">
  ...
  <dataType type="urn:domain-name:schema-name:datatype-name">
    string
  </dataType>
  ...
</scpd>
```

The second option is to define the namespace at the beginning of the SCPD document and then make a reference in the type definition. In this case, the `type` attribute contains a prefix that identifies the namespace, followed by the data type-named.

```
<scpd xmlns="urn:schemas-upnp-org:service-1-0"
  xmlns:dt="urn:domain-name:schema-name"
  configId="configuration number">
  ...
  <dataType type="dt:datatype-name">
    string
  </dataType>
  ...
</scpd>
```

Implementations MUST support both formats. The first format is potentially easier to parse, while the second format may result in shorter description files (i.e. when the same namespace is used multiple times in the same document).

These data types written in XSD Schema need not be processed at run-time. Instead, an implementer MAY use the referred schema as a standard description of the type to parse for that particular `type` attribute. To allow run-time schema processing of extended data types, an optional location of the extended data type schema MAY be expressed in the SCPD using the standard XSD `xsi:schemaLocation` and `xsi:noNamespaceSchemaLocation` attributes as defined in section 4.3.2 of XML Schema Part 1. These attributes can be used in the root SCPD element (essentially an instance document) where the extended data type is defined, as illustrated below:

```

<scpd xmlns="urn:schemas-upnp-org:service-1-0"
  xmlns:dt="urn:domain-name:schema-name"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:domain-name:schema-name
    http://some.company.com/dir/file.xsd"
  configId="configuration number">
  ...
  <dataType type="dt:datatype-name">
    string
  </dataType>
  ...
</scpd>

```

Alternatively, these attributes MAY be declared on use in the <dataType> element where the existing fully qualified namespace and type name for the extended data type are defined. An example for reference is given below:

```

<scpd xmlns="urn:schemas-upnp-org:service-1-0"
  configId="configuration number">
  ...
  <dataType type="urn:domain-name:schema-name:datatype-name"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:noNamespaceSchemaLocation="http://some.company.com/dir/file.xsd">
    string
  </dataType>
  ...
</scpd>

```

2.5.2 String equivalents of extended data types

A number of working committees have created services based on UPnP 1.0 (which does not support extended data types) that define their own encoding of information inside specific string-type state variables. To provide these working committees with an upgrade path to extended datatypes written in native XML³, a mechanism is defined that gives working committees the *option* to define the “string equivalent” of an extended data type (working committees can decide not to). If a string equivalent of an extended data type is defined, there are two valid ways to represent the value of the data type: either as an extended data type, written in native XML, or as a string, that encodes the datatype as specified by the working committee.

The mechanism uses the USER-AGENT header field in action invocations and event subscriptions (see also section 2.1, “Generic requirements on HTTP usage”). If a control point invokes an action without a USER-AGENT header field, or if the USER-AGENT header field does not specify UPnP version 1.1 or greater, the values of in-arguments and out-arguments MUST be encoded using the “string equivalent”.

If a control point invokes an action with a USER-AGENT header field that specifies UPnP version 1.1 or greater, the values of in-arguments and out-arguments MUST be encoded using the extended data type written in native XML.

³ In this text “native XML” refers to datatypes formatted according to XML-schema using the normal XML format, while “string-equivalent of an extended datatype” refers to encoding a complex data type inside a UPnP string, examples of which (escaped XML, comma-separated lists) can be found in the ContentDirectory:1 specification.

If a control point has subscribed to events without a USER-AGENT header field, or if the USER-AGENT header field specifies a UPnP version less than 1.1, all values of complex-type evented state variables that are sent to the control point MUST be encoded using the "string equivalent". If no "string equivalent" is defined for an evented state variable, subscription without the correct USER-AGENT header field MUST be refused.

If a control point has subscribed to events with a USER-AGENT header field that specifies UPnP version 1.1 or greater, all values of complex-type evented state variables that are sent to the control point MUST be encoded using the extended data type written in native XML.

2.5.3 Generic requirements

For future extensibility and according to the requirements in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values.

Subject to the constraints defined in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", control points and devices MUST ignore any XML comments or XML processing instructions embedded in UPnP device and service descriptions that they do not understand. UPnP service descriptions MUST be encoded using UTF-8.

When the value of any text element or attribute contains one or more characters reserved as markup (such as ampersand ("&") or less than (" $<$ ")), the text MUST be escaped in accordance with the provisions of section 2.4 of the XML specification and each such character replaced with the equivalent numeric representation or string (such as "&" or "<"). Such characters appearing in URLs MAY also be percent-encoded in accordance with the URL percent-encoding rules specified in sections 2.1 and 2.4 of RFC 3986. Note that it is logically possible for a service to have no actions but have state variables and eventing; though unlikely, such a service would be an autonomous information source. However, a service with no state variables is PROHIBITED.

Unlike device descriptions, service descriptions and associated values MUST NOT use locale-specific values; this includes service descriptions, values of action arguments, and values of state variables. Instead, most action arguments and state variables MUST use values that are expressed in a locale-independent manner; control points MAY convert and/or format the information from a standard form into the correct language and/or format for the locale. For example, dates are represented in a locale-independent format (ISO 8601), and integers are represented without locale-specific formatting (e.g., no currency symbol, no grouping of digits). String values MUST be represented in a locale-independent manner. Variables with an allowedValueList MUST use token values in the language of UPnP standards and not reflect strings intended to be displayed in a localized user interface.

2.5.4 Ordering of Elements

The order of XML elements in device and service description documents MUST adhere to the order as defined in the corresponding specification as defined by the working committee for that device or service type. Furthermore, the order of

elements (e.g. arguments) in control messages and in their responses **MUST** adhere to the order defined in the device's service description document.

Note: UPnP 1.0 does **NOT REQUIRE** that the order of XML elements in device and service description documents adheres to the order as defined in the corresponding schema (as defined by the working committee) for that device or service type. However, it *does REQUIRE* that control messages and responses are ordered according to the corresponding device's service description, a **REQUIREMENT** that is sometimes overlooked. Therefore, when receiving messages from UPnP 1.0 services, control points **SHOULD** be able to process out-of-order elements; and when transmitting messages to UPnP 1.0 services, control points **MUST** send elements in the order defined by that particular device's service description.

2.5.5 Versioning

Services standardized by UPnP Forum working committees have an integer version. Every later version of a service **MUST** be a superset of the previous version, i.e., it **MUST** include all actions and state variables exactly as they are defined by earlier versions of the service. The UPnP service type remains the same across all versions of a service whereas the service version **MUST** be larger for later versions. Versions of device and service templates **MAY** have non-integer versions (such as "0.9") during development in the working committee, but this **MUST** become an integer upon standardization. Devices and services **MAY** have a version number greater than the major version number of the architecture they are designed for (e.g., "Power:2" may be designed to work on UDA version 1.0).

2.6 UPnP Service Template

The listing above also illustrates the relationship between a UPnP service description and a UPnP Service Template. As explained above, the UPnP description for a service is written by a UPnP vendor, in XML, following a UPnP Service Template. A UPnP Service Template is produced by a UPnP Forum working committee as a means to standardize devices.

By appropriate specification of placeholders, the listing above can be either a UPnP Service Template or a UPnP service description. Recall that some placeholders would be defined by a UPnP Forum working committee (colored *red*), i.e., actions and their parameters, and states and their data type, range, and event characteristics. If these were specified, the listing above would be a UPnP Service Template, codifying the standard for this type of service. Along with UPnP Device Templates (see section 2, "Description"), UPnP Service Templates are one of the key deliverables from UPnP Forum working committees.

Taking this another step further, the remaining placeholders in the listing above would be specified by a UPnP vendor (colored *purple*), i.e., additional, vendor-specified actions and state variables. If these placeholders were specified (as well as the others), the listing would be a UPnP service description, suitable for effective control of the service within a device.

Put another way, the UPnP Service Template defines the overall type of service, and each UPnP service description instantiates that template with vendor-specific additions. The first is created by a UPnP Forum working committee; the latter by a UPnP vendor.

2.7 Non-standard vendor extensions and limitations

As explained above, UPnP vendors MAY differentiate their devices and extend a standard device by including additional services and embedded devices. Similarly, UPnP vendors MAY extend a standard service by including additional actions, state variables or allowed values. Naming conventions and conditions for each of these are listed in the table below and explained in detail above.

Table 2-1: Vendor extensions

| Type of extension | Standard | Non-Standard |
|--|--|---|
| device type | urn: schemas-upnp-org:device:deviceType:v | urn: domain-name:device:deviceType:v |
| service type | urn: schemas-upnp-org:service:serviceType:v | urn: domain-name:service:serviceType:v |
| service ID | urn: upnp-org:serviceId:serviceID | urn: domain-name:serviceId:serviceID |
| action name | MUST comply with the syntax rules of the standardized action name as defined in section 2.5, "Service description". | MUST comply with the syntax rules of the non-standardized action name as defined in section 2.5, "Service description". |
| stateVariable name | MUST comply with the syntax rules of the standardized stateVariable name as defined in section 2.5, "Service description". | MUST comply with the syntax rules of the non-standardized stateVariable name as defined in section 2.5, "Service description". |
| allowedValue value | MUST be a legal value for a string variable. Only values explicitly defined by a working committee are allowed. | Permitted only if allowed by the working committee. MUST begin with a Vendor Domain Name, followed by the underscore character (" _"), followed by a legal value for a string variable. |
| XML elements and their attributes in device or service description | Defined by the UPnP Device and Service Schemas. | Arbitrary XML, scoped by one or more XML namespaces owned by the vendor. MUST be enclosed in an element that begins with " X ". |
| XML attributes of standard elements in device or service description | Defined by the UPnP Device and Service Schemas. | Arbitrary attributes, scoped by one or more XML namespaces, owned by the vendor. MUST begin with " X ". |

As the last two rows of the table above indicate, UPnP vendors MAY also add non-standard XML to a device or service description. Each addition MUST be scoped by a vendor-owned XML namespace. Arbitrary XML MUST be enclosed in an element that begins with "[X](#)," and this element MUST be a sub element of a standard complex type. Non-standard attributes MAY be added to standard elements provided these attributes are scoped by a vendor-owned XML namespace and begin with "[X](#)".

To illustrate this, below are listings with placeholders (in *italics*) for actual elements and values. Some of these placeholders would be specified by a UPnP vendor (*purple*) and some are defined by the UPnP Device Architecture (*green*).

```

<RootStandardElement
  xmlns="urn:schemas-upnp-org:device-1-0"
  xmlns:n="domain-name:schema-name">
  <!-- other XML -->
  <AnyStandardElement n:XVendorAttribute="arbitrary string value">
    <!-- other XML -->
  </AnyStandardElement>
  <!-- other XML -->
</RootStandardElement>

```

<RootStandardElement>

A standard root element. `xmlns` attribute defines namespaces, in this case, a standard UPnP namespace and a non-standard namespace with the prefix *n*. (Note: *n* is just a placeholder. A vendor can specify any prefix to identify the namespace that is valid according to the *Namespaces in XML* specification.)

- For device descriptions, MUST be <root>.
- For service descriptions, MUST be <scpd>.

<AnyStandardElement>

Any standard element, root or otherwise, content of text or element only. MUST already be included as part of the standard device or service description. `X_VendorAttribute` MUST begin with "X". (Prefix "A" is reserved.) MAY have an arbitrary string value.

```
<StandardComplexType n:X_VendorAttribute="vendor value">
  <n:X_VendorElement xmlns:n="domain-name:schema-name">
    <!-- arbitrary XML -->
  </n:X_VendorElement>
</StandardComplexType>
```

<StandardComplexType>

Element of complex type. MUST already be included as part of the standard device or service description.

- For device descriptions, MUST be one of: <root>, <specVersion>, <device>, <iconList>, <icon>, <serviceList>, <service>, or <deviceList>.

For service descriptions, MUST be one of: <scpd>, <actionList>, <action>, <argumentList>, <argument>, <serviceStateTable>, <stateVariable>, <allowedValueList>, or <allowedValueRange>.

<X_VendorElement>

MUST begin with "X". (Prefix "A" is reserved.) MUST have a value for the `xmlns` attribute. MAY contain arbitrary XML.

2.7.1 Placement of Additional Elements and Attributes

Instances of any UPnP schema, including device and service descriptions, control actions, errors and event notifications, MAY include additional XML elements (other than those defined by the UPnP Forum) *only* at the end of an ordered sequence of elements corresponding to a given complex type. Additionally, instances of any UPnP schema MAY include additional attributes with any element.

Exception for UPnP 1.0 devices:

UPnP 1.0 allows the inclusion of additional elements anywhere within device and service descriptions, control actions, errors and event notifications, provided that the XML is well-formed. Therefore, when receiving messages from UPnP 1.0 devices, control points MUST handle unknown elements and attributes found anywhere within the message.

2.8 UPnP Device Schema

The paragraphs above explain UPnP device descriptions and illustrate how one would be instantiated from a UPnP Device Template. As explained, UPnP Device Templates are produced by UPnP Forum working committees, and these templates are based upon the UPnP Device Schema. This schema defines the structures and data types used to create UPnP Device Templates. Appendix B.1, "UPnP Device Schema" contains the schema; below is an explanation of this schema.

The UPnP Device Schema is written in XML and according to the conventions of XML Schema (Part 1: Structures, Part 2: Datatypes). XML Schema provides a method of describing the structure of an XML document. The XML Schema description language itself is based upon XML. The language is very robust; it specifies which elements are REQUIRED vs. OPTIONAL, element nesting, data types for values (as well as other properties not of interest here) and much more. The UPnP Device Schema uses these XML Schema constructions to define elements like `<specVersion>`, `<URLBase>`, `<deviceType>`, et al listed in detail above. Because the UPnP Device Schema is constructed using a precise description language, it is unambiguous. As the UPnP Device Schema, UPnP Device Templates, and UPnP device descriptions are all machine-readable, software tools MAY be devised to validate the latter two, checking that they contain all the REQUIRED elements, are correctly nested, and have values of the correct data types.

2.9 UPnP Service Schema

The paragraphs above explain UPnP service descriptions and illustrate how one would be instantiated from a UPnP Service Template. Like UPnP Device Templates, UPnP Service Templates are produced by UPnP Forum working committees, and these templates are based upon the UPnP Service Schema. This schema defines the structure and data types used to create UPnP Service Templates. As explained above, the UPnP Service Schema is written in XML according to the conventions of XML Schema (Part 1: Structures, Part 2: Datatypes). Appendix B.2, "UPnP Service Schema" contains a listing of this schema

2.10 UPnP Datatype Schema


The UPnP basic data types for state variables are defined in section 2.5, "Service description". For any extended data types for state variables used by a service template, the service template MUST include either a reference to all relevant schemas for the extended data types or include a new service specific datatype schema with a corresponding unique target namespace. If any extended data types are used for state variables within an SCPD, the corresponding namespace for each extended data type MUST be referenced within the SCPD according to the "Namespaces in XML" specification. Section 2.5, "Service description" contains an example SCPD with namespace declarations.

2.11 Retrieving a description using HTTP

As explained above, after a control point has discovered a device, it still knows very little about the device. To learn more about the device and its capabilities, the control point MUST retrieve the UPnP description for the device using the URL provided by the device in the discovery message. Then, the control point MUST retrieve one or more service descriptions using the URL(s) provided in the device description. This is a simple HTTP-based process and uses the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

A multi-homed device MUST send description documents using the UPnP-enabled interface on which the HTTP GET request was received. To retrieve the UPnP description using a particular interface, a multi-homed control point MUST use the URL provided in the discovery message which arrived on that interface.

Figure 2-2: Description retrieval protocol stack



At the highest layer, description messages contain vendor-specific information, e.g., device type, service type, and required services. Moving down the stack, vendor content is supplemented by information from a UPnP Forum working committee, e.g., model name, model number, and specific URLs. Messages from the layers above are hosted in UPnP-specific protocols, defined in this document. In turn, the above messages are delivered via HTTP over TCP over IP. For reference, colors in [square brackets] above indicate which protocol defines specific header fields and body elements in the description messages listed below.

When a control point discovers a device on the network, it MAY wish to retrieve the Device Description document and Service Description Documents. Retrieving the UPnP device description is simple: the control point issues an HTTP GET request to the URL in the discovery message, and the device returns its description in the body of an HTTP response. Similarly, to retrieve a UPnP service description, the control point issues an HTTP GET request to the corresponding URL in the device description, and the device returns the description in the body of an HTTP response. The header fields and body for the response and request are explained in detail below.

First, a control point MUST send a request with method GET in the following format. Values in *italics* are placeholders for actual values.

```
GET /descriptionPath HTTP/1.1
HOST: hostname:portNumber
ACCEPT-LANGUAGE: language preferred by control point
```

(No body for request to retrieve a description, but note that the message MUST have a blank line following the last HTTP header field.)

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted. See RFC 2616 and RFC 1945 for further requirements on encoding of values of these fields.

Request line

GET

Method defined by HTTP. Can be GET or HEAD.

descriptionPath

Path component of device description URL (LOCATION header field in discovery message) or of the fully qualified service description URL. (If the SCPDURL sub element of the service element in the device description is an absolute URL, the fully qualified service description URL is the SCPDURL sub element. Otherwise (the SCPDURL sub element is a relative URL), the fully qualified service description URL is the URL resolved from the SCPDURL sub element in accordance with section 5 of

RFC 3986, using either the URLBase element, if specified, or the URL from which the device description was retrieved as the base URL.) Single, absolute path (see also RFC 2616).

HTTP/1.1

The version supported by the control point. (Note: the control point MUST implement all MANDATORY components of the version specified). MAY be any HTTP version that is backwards compatible to HTTP/1.0 (like HTTP/1.1).

Header fields

HOST

REQUIRED. Field value contains domain name or IP address and optional port components of device description URL (LOCATION header field in discovery message) or of the fully qualified service description URL. (If the SCPDURL sub element of the service element in the device description is an absolute URL, the fully qualified service description URL is the SCPDURL sub element. Otherwise (the SCPDURL sub element is a relative URL), the fully qualified service description URL is the URL resolved from the SCPDURL sub element in accordance with section 5 of RFC 3986, using either the URLBase element, if specified, or the URL from which the device description was retrieved as the base URL.) If the port is empty or not given, port 80 is assumed.

ACCEPT-LANGUAGE

OPTIONAL. RECOMMENDED for retrieving device descriptions. Field value contains preferred language(s) for description. If no description is available in this language, device MAY return a description in a default language. See RFC 1766 language tag(s).

After a control point sends a request, the device takes the second step and responds with a copy of its description. Including expected transmission time, a device MUST respond within 30 seconds. If it fails to respond within this time, the control point SHOULD re-send the request. A device MUST send a response in the following format and in accordance with section 2.1, "Generic requirements on HTTP usage". Two example responses are provided below: one that uses the CONTENT-LENGTH header field, and one that uses chunked encoding (with 2 chunks). Values in *italics* are placeholders for actual values.

Note: XML namespace prefixes do not have to be the specific examples shown below (e.g., "s" or "u"); they can be any value that obeys the rules of the general XML namespace mechanism; a device MUST accept action invocations that use other legal XML namespace prefixes.

Response using CONTENT-LENGTH header field – Valid with HTTP/1.0 or HTTP/1.1

```
HTTP/1.1 200 OK
CONTENT-LANGUAGE: language used in description
CONTENT-LENGTH: bytes in body
CONTENT-TYPE: text/xml; charset="utf-8"
DATE: when responded

Body
```

Response using chunked encoding – Valid with HTTP/1.1 only

```
HTTP/1.1 200 OK
TRANSFER-ENCODING: chunked
CONTENT-TYPE: text/xml; charset="utf-8"
CONTENT-LANGUAGE: language used in description
DATE: when responded

Length of chunk 1 in hexadecimal notation
Chunk 1
Length of chunk 2 in hexadecimal notation
Chunk 2

0
```

The body of this response is a UPnP device or service description as explained in detail above. Listed below are details for the header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Status Line

HTTP/1.1

The highest version supported by the origin server that is compatible with the control point that issued the request. For example, if the control point specified support for HTTP/1.0 in the request, the response MUST contain HTTP/1.0.

200 OK

HTTP defined status code indicating that no HTTP errors have occurred.

Header fields

CONTENT-LANGUAGE

REQUIRED if and only if request included an ACCEPT-LANGUAGE header field. Field value contains language of description. RFC 1766 language tag(s).

CONTENT-LENGTH

REQUIRED if Origin Server does not close the session after sending the response AND Origin Server does not send the response using chunked encoding.
PROHIBITED if Origin Server sends the response using chunked encoding. OPTIONAL otherwise.
Field value specifies the length of the body in bytes. Integer.

TRANSFER-ENCODING

OPTIONAL for HTTP/1.1 and above. Field value specifies whether the response is chunked encoded by having field value "chunked". MUST NOT be specified if CONTENT-LENGTH header field is present.

CONTENT-TYPE

REQUIRED. Field value MUST be "text/xml; charset="utf-8" " for description documents.

DATE

RECOMMENDED according to RFC 2616, section 14.18. Field value contains date when response was generated. "rfc1123-date" as defined in RFC 2616.

SERVER

(No SERVER header field is required for description messages.)

Note that because HTTP 1.1 allows use of chunked encoding, some devices MAY send the description using chunked encoding if the GET request specifies HTTP 1.1. Therefore all implementations that include HTTP 1.1 client support MUST support receiving chunked encoding.

2.12 References

ISO 8601

ISO (International Organization for Standardization). Representations of dates and times, 1988-06-15. Available at: <http://www.w3.org/TR/1998/NOTE-datetime-19980827>.

RFC 822

Standard for the format of ARPA Internet text messages. Available at: <http://www.ietf.org/rfc/rfc822.txt>.

RFC 1123

Includes format for dates, for, e.g., HTTP DATE header field. Available at: <http://www.ietf.org/rfc/rfc1123.txt>.

RFC 1766

Format for language tag for, e.g., HTTP ACCEPT-LANGUAGE header field. Available at: <http://www.ietf.org/rfc/rfc1766.txt>. See also <http://www.loc.gov/standards/iso639-2> for language codes.

RFC 2045

Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies. Available at: <http://www.ietf.org/rfc/rfc2045.txt>.

RFC 2046

Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types. Available at: <http://www.ietf.org/rfc/rfc2046.txt>.

RFC 2083
PNG (Portable Network Graphics) Specification Version 1.0. Available at: <http://www.ietf.org/rfc/rfc2083.txt>. See also <http://www.w3.org/TR/REC-png.html>.

RFC 2387
Format for representing content type, e.g., mimetype element for an icon. Available at: <http://www.ietf.org/rfc/rfc2387.txt>.

RFC 2616
HTTP: Hypertext Transfer Protocol 1.1. Available at: <http://www.ietf.org/rfc/rfc2616.txt>.

RFC 3986
Uniform Resource Identifiers (URI): Generic Syntax. Available at: <http://www.ietf.org/rfc/rfc3986.txt>.

UPC
Universal Product Code. 12-digit, all-numeric code that identifies the consumer package. Managed by the Uniform Code Council. Available at: http://www.uc-council.org/main/ID_Numbers_and_Bar_Codes.html.

XML
Extensible Markup Language. Available at: <http://www.w3.org/TR/2000/REC-xml-20001006>.

XML Schema (Part 1: Structures, Part 2: Datatypes). Available at: <http://www.w3.org/TR/xmlschema-1>, <http://www.w3.org/TR/xmlschema-2>.

Namespaces in XML
Available at: <http://www.w3.org/TR/REC-xml-names/>.

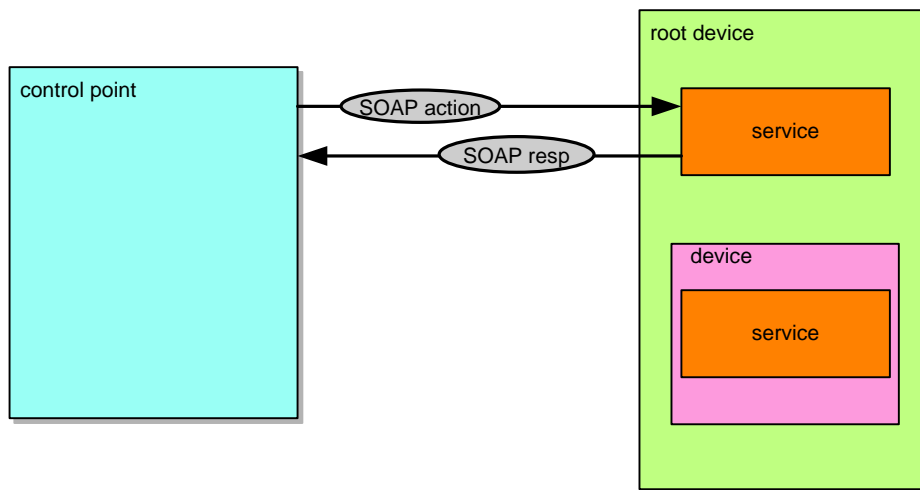
3 Control

[Normative]

Control is Step 3 in UPnP™ networking. Control comes after addressing (Step 0) where devices get a network address, after discovery (Step 1) where control points find interesting device(s), and after description (Step 2) where control points learn about device capabilities. Control is independent of eventing (Step 4) where control points listen to state changes in device(s). Through control, control points invoke actions on devices and poll for values. Control and eventing are complementary to presentation (Step 5) where control points display a user interface provided by device(s).

Given knowledge of a device and its services, a control point can ask those services to invoke actions and receive responses indicating the result of the action. Invoking actions is a kind of remote procedure call; a control point sends the action to the device's service, and when the action has completed (or failed), the service returns any results or errors.

Figure 3-1: Control architecture



To control a device, a control point invokes an action on the device's service. To do this, a control point sends a suitable control message to the fully qualified control URL for the service obtained from the controlURL sub element of the service element of the device description. If the controlURL sub element is an absolute URL, the fully qualified control URL for the service is the controlURL sub element. Otherwise (the controlURL sub element is a relative URL), the fully qualified control URL for the service is the URL resolved from the controlURL sub element in accordance with section 5 of RFC 3986, using either the URLBase element of the device description, if specified, or the URL from which the device description was retrieved as the base URL. A multi-homed control point that sends the control message on a particular interface MUST use the fully qualified control URL from the description document received on that interface. In response, the service returns any results or errors from the action. The effects of the action, if any, MAY also be modeled by changes in the variables that describe the run-time state of the service. When these state variables change, events are published to all interested control points. This section explains the protocol stack for, and format of, control messages. Section 4, "Eventing" explains event publication.

Working committees and vendors MAY define actions to allow control points to determine the current value of one or more state variables. Similar to invoking an action, a control point sends the defined query message to the control URL for the service. In response, the service provides the value of the variable or variables; each service is responsible for keeping its state table consistent so control points can poll and receive meaningful values for those state variables for which query actions are defined. Section 4, "Eventing" explains automatic notification of variable values.

As long as one of the discovery advertisements from a device have not expired, a control point MAY assume that the device and its services are available. If a device cancels at least one of its advertisements, a control point MUST assume the device and its services are no longer available.

Control points and devices MUST use UTF-8 for all UPnP control messages and responses.

While UDA does define a means to invoke actions and poll for values, UDA does not specify or constrain the design of an API for applications running on control points; OS vendors MAY create APIs that suit their customers' needs.

If a large amount of data must be sent in association with an action (particularly if the amount of data is not known in advance), it is NOT RECOMMENDED to send the data as part of a SOAP argument or as a MIME attachment to the SOAP message. Instead, it is RECOMMENDED that out-of-band transfer be used. For example, a URL could be sent as an argument value, and an HTTP GET, PUT, or POST be used to transfer the data. HTTP chunked encoding can be used when the amount of data is not known in advance.

Responses to SOAP messages during the Control phase MUST be sent to the same IP address from which the request was received. Any fully-qualified URLs contained in an action or response that refer to a resource on the device itself MUST have the HOST portion of the URL set appropriately so that the resource will be reachable by the control point that requested the action. This might be accomplished by using the field value specified in the HTTP HOST header field of the control request.


Services that use complex datatype arguments MUST follow the requirements in section 2.5, "Service description"

The remainder of this section explains in detail how control messages are formatted and sent to devices.

3.1 Control protocols

To invoke actions and poll for values, control points (and devices) use the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

Figure 3-2: Control protocol stack



At the highest layer, control messages contain vendor-specific information, e.g., argument values. Moving down the stack, vendor content is supplemented by information from a UPnP Forum working committee, e.g., action names, argument names, variable names. Messages from the layers above are hosted in UPnP-specific protocols, defined in this document. In turn, the above messages are formatted using a Simple Object Access Protocol (SOAP) header and body elements, and the messages are delivered via HTTP over TCP over IP. For reference, colors in [square brackets] above indicate which protocol defines specific header field elements in the subscription messages listed below.

The generic requirements on HTTP usage in UPnP 1.1 (as defined in section 2.1, “Generic requirements on HTTP usage” of this document) MUST be followed by devices and control points that implement Control.

3.1.1 SOAP Profile

UPnP profiles SOAP 1.1, NOT REQUIRING that all devices support all OPTIONAL features of SOAP 1.1, but devices and control points MUST support all MANDATORY features of SOAP 1.1. The following table summarizes the UPnP profiling of SOAP.

Table 3-1: SOAP 1.1 UPnP Profile

UPnP Control Request	Mandatory Optional Prohibited	Comment
<Envelope> element	M	
encodingStyle attribute of <Envelope>	O	If present, must be " http://schemas.xmlsoap.org/soap/encoding/ "
<Header> element (child element of <Envelope>)	O	
actor attribute of <Header>	O	
mustUnderstand attribute of <Header>	O	Only allowed if defined by the service to which it is directed
encodingStyle attribute of <Header>	O	If present, must be " http://schemas.xmlsoap.org/soap/encoding/ "
<Body> element (child element of <Envelope>)	M	Exactly one <Body> child element allowed
encodingStyle attribute of <Body> element	P	
root Attribute of <Body> child element	O	SHOULD NOT be used
UPnP Control Response		
<Envelope> element	M	
encodingStyle attribute of <Envelope>	O	If present, must be " http://schemas.xmlsoap.org/soap/encoding/ "
<Header> element (child element of <Envelope>)	O	
actor attribute of <Header>	O	
mustUnderstand attribute of <Header>	O	Only allowed if defined by the service to which it is directed
encodingStyle Attribute of <Header>	O	If present, must be " http://schemas.xmlsoap.org/soap/encoding/ "
<Body> element (child element of <Envelope>)	M	Exactly one <Body> child element allowed
encodingStyle attribute of <Body> element	P	
root attribute of <Body> child element	O	SHOULD NOT be used
UPnP Control Error Response		
<Envelope> element	M	
encodingStyle attribute of <Envelope>	O	If present, must be " http://schemas.xmlsoap.org/soap/encoding/ "
<Body> element (child element of <Envelope>)	M	Exactly one <Body> child element allowed containing exactly one <Fault> child element
<Fault> child element of <Body>	M	
<faultcode> child element of <Fault>	M	
<faultstring> child element of <Fault>	M	
<detail> child element of <Fault>	M	

SOAP 1.1 allows the use of footers, which are disallowed in SOAP 1.2. A UPnP message MUST NOT have any child elements of the <Envelope> element following the <Body> element.

SOAP <Header> element

UPnP working committees and the UPnP technical committee MAY define OPTIONAL or MANDATORY SOAP header entries that are included in the SOAP <Header> element of UPnP action and UPnP action response messages. In addition, vendors MAY include other SOAP header entries in the SOAP <Header> element of UPnP action and UPnP action response messages. If there are no SOAP header entries in a message, the SOAP <Header> element can be omitted.

SOAP `mustUnderstand` Attribute of <Header> element

The `mustUnderstand` attribute MUST NOT be added to SOAP <Header> element targeted at (see also `actor` attribute below) standardized UPnP services or targeted at control points that interact with standardized UPnP services, *unless* its use has been explicitly defined by the UPnP technical committee or a working committee (e.g. UPnP security).

The `mustUnderstand` header attribute MUST NOT be included on non-standard header entries that are targeted at (see also `actor` attribute) standardized services, as this breaks the basic interoperability of UPnP. `mustUnderstand` header entries MAY be included on non-standard header entries that are neither targeted at (see also `actor` attribute) standardized services (e.g. vendor defined services), nor targeted at control points interacting with standardized services.

Table 3-2: `mustUnderstand` attribute

SOAP Node Type	v1.0	v1.1
Transmitting Node targeting a standardized service or a control point that interacts with a standardized service.	The <code>mustUnderstand</code> attribute MUST NOT be added to SOAP header entries, <i>unless</i> the UPnP technical committee or a working committee has explicitly defined its use.	
Transmitting Node targeting a vendor specific service or a vendor specific SOAP node.	MUST target endpoint (see <code>actor</code> section below). The <code>mustUnderstand</code> attribute MAY be used at the discretion of the vendor.	MAY target intermediaries (see <code>actor</code> section below). The <code>mustUnderstand</code> attribute MAY be used at the discretion of the vendor.
Receiving Node	All unknown <Header> entries are ignored, except when explicitly defined differently by a working committee (UPnP Security).	All devices MUST honor the <code>actor</code> (see <code>actor</code> section below) and <code>mustUnderstand</code> attributes. If a header entry with <code>mustUnderstand="1"</code> is not understood, the whole message fails and a <Faultcode> element MUST be returned.

The SOAP `mustUnderstand` attribute has a restricted type of "xsd:boolean" that takes only "0" or "1" with "1" being true and "0" being false. A header entry with the `mustUnderstand` attribute set to a value of "1" MUST be processed by targeted nodes or message processing MUST fail. Such elements are considered "MANDATORY header entries". A SOAP node is considered to understand a SOAP header entry if that SOAP node understands the semantics specified for the XML expanded name of the outermost element information item of that header entry. MANDATORY SOAP header entries are presumed to modify the semantics of other SOAP header entries or SOAP <Body> elements and therefore MUST be understood for correct semantics. UPnP nodes receiving header entries flagged with the `mustUnderstand` attribute MUST process and understand MANDATORY header entries that are targeted at that node or the node MUST NOT process the SOAP message at all. If a node fails to process or understand a

MANDATORY entry, that node MUST generate a SOAP [Fault](#) with the [faultcode](#) element set to "MustUnderstand". Support for MANDATORY header entries assures that key message parts that are targeted at a particular SOAP node will not be erroneously ignored.

If a `<Header>` entry is a MANDATORY `<Header>` entry and contains entries not understood by the targeted SOAP node, the SOAP node MAY attempt processing without understanding the semantics of the extensions. MANDATORY extensions are not possible.

SOAP actor Attribute of `<Header>` element

The SOAP [actor](#) attribute is used in SOAP 1.1 to identify the URI of SOAP node that is to process the `<Header>` entry. All SOAP nodes play the role of "http://schemas.xmlsoap.org/soap/actor/next", which is the first node (device or control point) that processes the message. The lack of an [actor](#) attribute indicates that the entry is targeted at the destination. All UPnP defined `<Header>` elements MUST be targeted at the destination, unless explicitly defined otherwise by the UPnP technical committee or a working committee. Therefore, it is RECOMMENDED that the actor attribute is not included on UPnP `<Header>` entries.

`<Header>` entries within messages that are sent to UPnP 1.0 devices or control points MUST NOT be targeted at intermediaries (no actor attribute), since UPnP 1.0 devices and control points might ignore the actor attribute and parse a `<Header>` entry that is not intended for them.

If `<Header>` entries with an actor attribute are targeted at an intermediary and tagged [mustUnderstand="1"](#), the device or control point MUST NOT return a SOAP Fault containing the [faultcode](#) element set to "MustUnderstand" due to failure to process the relevant `<Header>` element targeted at another entity.

SOAP root Attribute

UPnP 1.1 REQUIRES that the first child element of the `<Body>` element MUST be the root of the RPC request. Since UPnP 1.1 defines an RPC-architecture, there can only be one root. The serialization root SHOULD NOT use the [root](#) attribute, but it is NOT PROHIBITED.

SOAP encodingStyle Attribute

UPnP 1.1 REQUIRES that devices and control points MUST be able to receive messages that do not contain the SOAP [encodingStyle](#) attribute, as well as messages that contain the SOAP [encodingStyle](#) attribute with value "[http://schemas.xmlsoap.org/soap/encoding/](#)". When [encodingStyle](#) is not included, the [encodingStyle](#) is "[http://schemas.xmlsoap.org/soap/encoding/](#)".

When communicating with UPnP 1.0 devices or control points, an [encodingStyle](#) attribute MUST be included on the SOAP `<Envelope>` element with value "[http://schemas.xmlsoap.org/soap/encoding/](#)". When communicating with UPnP 1.1 devices or control points, the [encodingStyle](#) attribute SHOULD be included and, if present, MUST have the value "[http://schemas.xmlsoap.org/soap/encoding/](#)".

If additional encodings are needed for application data, applications MAY use out of band data encoding for the relevant data.

SOAP <Body> element

UPnP 1.1 REQUIRES a <Body> element. It contains body entries for UPnP Actions and Responses. The actual entries are derived from the Service Description for the chosen Action. A response is either successful, in which case it contains output arguments, or unsuccessful, when it contains a <Fault> element as the only entry.

SOAP <Fault> element of <Body> element

UPnP REQUIRES the use of SOAP <Fault> elements when a failure response is returned. Please see Table 3-2, “mustUnderstand attribute” on usage of the [mustUnderstand](#) attribute for how the <detail> element MUST be constructed. When a <Header> element is encountered that is a MANDATORY <Header> element, the control point or device MUST either recognize the element or return the appropriate SOAP <Fault> element, containing the <faultcode> element set to “[MustUnderstand](#)”.

Backwards-compatible services MUST NOT use MANDATORY <Header> elements since previous UDAs allowed unknown <Header> elements to be ignored.

Acceptable SOAP Character Encodings

All messages MUST use UTF-8 serialization. The device or control point MUST indicate the content type for all control messages using the HTTP “charset” parameter.

3.2 Actions

Control points MAY invoke actions on a device's services and receive results or errors back. The action, results, and errors are encapsulated in SOAP, sent via HTTP requests, and received via HTTP responses.

3.2.1 Action invocation

The Simple Object Access Protocol (SOAP) defines the use of XML and HTTP for remote procedure calls. UPnP 1.1 uses HTTP to deliver SOAP 1.1 encoded control messages to devices and return results or errors back to control points. See section 2.1, “Generic requirements on HTTP usage” on use of HTTP in UPnP 1.1.

UPnP 1.1 deprecates the use of the HTTP Extension Framework (RFC 2774) for control. Specifically, control points MUST send a request with method POST and MUST NOT use the M-POST method. Devices MUST NOT reject POST methods with a “405 Method Not Allowed” HTTP status code since this causes UPnP 1.0 control points to issue a request using M-POST.

Below is a listing of a control message sent using the POST method followed by an explanation of the header fields and body. To invoke an action on a device's service, a control point MUST send a request with method POST in the following format. Two examples are provided: one using the CONTENT-LENGTH header and one using chunked encoding (with 2 chunks). Values in *italics* are placeholders for actual values.

Note: XML namespace prefixes do not have to be the specific examples shown below (e.g., “s” or “u”); they can be any value that obeys the rules of the general XML namespace mechanism; a device MUST accept action invocations that use other legal XML namespace prefixes.

Action invocation using the CONTENT-LENGTH header – Valid with HTTP/1.0 or HTTP/1.1

```
POST path control URL HTTP/1.0
HOST: hostname:portNumber
CONTENT-LENGTH: bytes in body
CONTENT-TYPE: text/xml; charset="utf-8"
USER-AGENT: OS/version UPnP/1.1 product/version
SOAPACTION: "urn:schemas-upnp-org:service:serviceType:v#actionName"

<?xml version="1.0"?>
<s:Envelope
  xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
  s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
  <s:Body>
    <u:actionName xmlns:u="urn:schemas-upnp-org:service:serviceType:v">
      <argumentName>in arg value</argumentName>
      <!-- other in args and their values go here, if any -->
    </u:actionName>
  </s:Body>
</s:Envelope>
```

Action invocation using chunked encoding – Valid with HTTP/1.1 only

```
POST path control URL HTTP/1.1
HOST: hostname:portNumber
TRANSFER-ENCODING: "chunked"
CONTENT-TYPE: text/xml; charset="utf-8"
USER-AGENT: OS/version UPnP/1.1 product/version
SOAPACTION: "urn:schemas-upnp-org:service:serviceType:v#actionName"

Length of first chunk in hexadecimal notation
<?xml version="1.0"?>
<s:Envelope
  xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
  s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
  <s:Body>
    <u:actionName xmlns:u="urn:schemas-upnp-org:service:serviceType:v">
      <argumentName>in arg value</argumentName>
      <!-- other in args and their value go here, if any -->
    </u:actionName>
  </s:Body>
Length of second chunk in hexadecimal notation
</s:Envelope>
0
```

Listed below are details for the request line, header fields, and body elements appearing in the listing above. Field names are not case sensitive. All HTTP field values and XML element names are case sensitive; XML values are not case sensitive except where noted. Except where noted, REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once.

Request line

POST

Method defined by HTTP.

path control URL

Path component of the fully qualified control URL for this service. Single, absolute path (see also RFC 2616, section 3.2.2).

HTTP/1.1

The version supported by the control point. (Note: the control point MUST implement all MANDATORY components of the version specified). MAY be any HTTP version that is backwards compatible to HTTP/1.0 (like HTTP/1.1).

Header fields

HOST

REQUIRED. Field value contains domain name or IP address and OPTIONAL port components of fully qualified control URL for this service. If the port is empty or not given, port 80 is assumed.

ACCEPT-LANGUAGE

PROHIBITED. The ACCEPT-LANGUAGE header field MUST NOT be used in control messages.

CONTENT-LENGTH

REQUIRED if Origin Server does not close the session after sending the action invocation AND Origin Server does not send the action invocation using chunked encoding.

PROHIBITED if Origin Server sends the action invocation using chunked encoding. OPTIONAL otherwise. Field value specifies the length of the body in bytes. Integer.

TRANSFER-ENCODING

OPTIONAL for HTTP/1.1 and above. Field value specifies whether the action invocation is chunked encoded by having field value "chunked". MUST NOT be specified if CONTENT-LENGTH header field is present.

CONTENT-TYPE

REQUIRED. Field value MUST be "text/xml; charset=utf-8".

USER-AGENT

OPTIONAL. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, "USER-AGENT: *unix/5.1 UPnP/1.1 MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1. See section 2.5, "Service description".

SOAPACTION

REQUIRED header field defined by SOAP. Field value MUST be the service type, hash mark, and name of action to be invoked, all enclosed in double quotes. The specified service version MUST indicate the version of the service that the control point wants to use while invoking the action. Its value may be any version of the service type in which the specified action was defined. When a control point invokes an action that has been defined in version K of a service, version number v MUST be equal or higher than K. For example, if an action has been defined in version 2 of a service, it MUST NOT be invoked using v=1. Furthermore; version v MUST be a version that is supported by the device. For example, for devices that support only version 1 of a service, v MUST be 1. Single URI.

Body

<Envelope>

REQUIRED element defined by SOAP. `xmlns` namespace attribute MUST be "<http://schemas.xmlsoap.org/soap/envelope/>". MUST include `encodingStyle` attribute with value "<http://schemas.xmlsoap.org/soap/encoding/>". A receiver MUST generate a fault if it encounters a message whose `<document>` element has a local name of "`Envelope`" but a namespace name that is not "<http://schemas.xmlsoap.org/soap/envelope/>". Contains the following sub elements:

<Body>

REQUIRED element defined by SOAP. MUST be qualified with SOAP namespace. Contains the following entry:

<actionName>

REQUIRED. Name of element is name of action to invoke. `xmlns` namespace attribute MUST be the service type enclosed in double quotes. The version specified MUST be the same version specified in the SOAPACTION header field. Case sensitive. MUST be the first child element of `<Body>`. Contains the following, ordered sub element(s):

<argumentName>

REQUIRED if and only if action has in arguments. Value to be passed to action. Repeat once for each in argument. (An element name is not qualified by a namespace; element nesting context is sufficient.) Case sensitive. Single data type as defined by UPnP service description. Every "in" argument in the definition of the action in the service description MUST be included, in the same order as specified in the service description (SCPD) that is available from the device.

If the CONTENT-TYPE header field specifies an unsupported field value (other than "text/xml") the device MUST return a "415 Unsupported Media Type" HTTP status code.

For future extensibility and according to the requirements in section 2.7, “Non-standard vendor extensions” and section 2.8, “UPnP Device Schema”, when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values.

Subject to the constraints defined in section 2.7, “Non-standard vendor extensions” and section 2.8, “UPnP Device Schema”, control points and devices MUST ignore any XML comments or XML processing instructions embedded in action requests that they do not understand.

When the value of any argument contains one or more characters reserved as markup (such as ampersand (“&”) or less than (“<”)), then the text MUST be escaped in accordance with the provisions of section 2.4 of the XML specification and each such character replaced with the equivalent numeric representation or string (such as “&” or “<”). Such characters appearing in URLs MAY also be percent-encoded in accordance with the URL percent-encoding rules specified in sections 2.1 and 2.4 of RFC 3986.

Note that because HTTP 1.1 allows use of chunked encoding, some control points MAY send the action request using chunked encoding if the POST method specifies HTTP 1.1. Device implementations that only support HTTP/1.0 and thus do not support receiving action requests using chunked encoding MUST return a “505 HTTP Version Not Supported” HTTP status code. Control points MUST NOT make HTTP 1.1 chunked POST requests to devices that are known to support only HTTP 1.0.

On a multi-homed control point, all fully qualified URLs contained in the action arguments that refer to resources on the control point MUST be reachable on the interface on which the action request is sent.

3.2.2 Action Response

The service MUST complete invoking the action and respond within 30 seconds, including expected transmission time (measured from the time the action message is transmitted until the time the associated response is received). Actions that take longer than this SHOULD be defined to return early and send an event when complete. If the service fails to respond within this time, what the control point SHOULD do is application-specific. A multi-homed device MUST send the response on the same UPnP-enabled interface on which the request was received. The service MUST send a successful completion response using the following format. The following two examples illustrate an action response using the CONTENT-LENGTH header and an action response using chunked encoding. The values in *italics* are placeholders for actual values.

Note; XML namespace prefixes do not have to be the specific examples shown below (e.g., “s” or “u”); they can be any value that obeys the rules of the general XML namespace mechanism; control points MUST accept action responses that use other legal XML namespace prefixes.

Action response using the CONTENT-LENGTH header – Valid with HTTP/1.0 or HTTP/1.1

```
HTTP/1.0 200 OK
CONTENT-TYPE: text/xml; charset="utf-8"
DATE: when response was generated
SERVER: OS/version UPnP/1.1 product/version
CONTENT-LENGTH: bytes in body
<?xml version="1.0"?>
<s:Envelope
  xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
  s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
  <s:Body>
    <u:actionNameResponse xmlns:u="urn:schemas-upnp-org:service:serviceType:v">
      <argumentName>out arg value</argumentName>
      <!-- other out args and their values go here, if any -->
    </u:actionNameResponse>
  </s:Body>
</s:Envelope>
```

Action response using chunked encoding – Valid with HTTP/1.1 only

```
HTTP/1.1 200 OK
TRANSFER-ENCODING: "chunked"
CONTENT-TYPE: text/xml; charset="utf-8"
DATE: when response was generated
SERVER: OS/version UPnP/1.1 product/version

Length of first chunk in hexadecimal notation
<?xml version="1.0"?>
<s:Envelope
  xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
  s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
  <s:Body>
    <u:actionNameResponse xmlns:u="urn:schemas-upnp-org:service:serviceType:v">
      <argumentName>out arg value</argumentName>
      <!-- other out args and their values go here, if any -->
    </u:actionNameResponse>
  </s:Body>
</s:Envelope>
0
```

Listed below are details for the response line, header fields, and body elements appearing in the listing above. Field names are not case sensitive. All HTTP field values and XML element names are case sensitive; XML values are not case sensitive except where noted. Except where noted, REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once.

Response line

HTTP/1.1

The highest version supported by the origin server that is compatible with the control point that issued the request. For example, if the control point specified support for HTTP/1.0 in the request, the response MUST contain HTTP/1.0.

200 OK

HTTP defined status code indicating that no HTTP errors were detected.

Header fields

CONTENT-LANGUAGE

PROHIBITED. The CONTENT-LANGUAGE header field MUST NOT be used in control messages.

CONTENT-LENGTH

REQUIRED if Origin Server does not close the session after sending the response AND Origin Server does not send the response using chunked encoding.

PROHIBITED if Origin Server sends the response using chunked encoding. OPTIONAL otherwise.

Field value specifies the length of the body in bytes. Integer.

TRANSFER-ENCODING

OPTIONAL for HTTP/1.1 and above. Field value specifies whether the response is chunked encoded by having field value "chunked" (in the example, the entire body is sent in a single chunk). MUST NOT be specified if CONTENT-LENGTH header field is present.

CONTENT-TYPE

REQUIRED. Field value MUST be "text/xml; charset=utf-8".

DATE

RECOMMENDED. Field value contains date when response was generated. "rfc1123-date" as defined in RFC 2616.

SERVER

REQUIRED. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, "SERVER: *unix/5.1 UPnP/1.1 MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

Body

<Envelope>

REQUIRED element defined by SOAP. xmlns namespace attribute MUST be "<http://schemas.xmlsoap.org/soap/envelope/>". MUST include [encodingStyle](#) attribute with value "<http://schemas.xmlsoap.org/soap/encoding/>". A receiver MUST generate a fault if it encounters a message whose document element has a local name of "[Envelope](#)" but a namespace name that is not "<http://schemas.xmlsoap.org/soap/envelope/>". Contains the following sub elements:

<Body>

REQUIRED element defined by SOAP. MUST be qualified with SOAP namespace. Contains the following entry:

<actionNameResponse>

REQUIRED. Name of element is action name prepended to [Response](#). xmlns namespace attribute MUST be service type enclosed in double quotes. Devices that support the same action in multiple namespaces MUST use the same namespace in the response as was used in the action invocation. For example, if an action was invoked using namespace:

urn:[schemas-upnp-org:service:ContentDirectory:2](#)

The response MUST also use namespace:

urn:[schemas-upnp-org:service:ContentDirectory:2](#)

Case sensitive. MUST be the first sub element of <[Body](#)>. Contains the following sub element:

<argumentName>

REQUIRED if and only if action has out arguments. Value returned from action. Repeat once for each out argument. If action has an argument marked with the <retVal/> element, this argument MUST be the first element. (An element name not qualified by a namespace; element nesting context is sufficient.) Case sensitive. Single data type as defined by UPnP service description. Every out argument in the definition of the action in the service description MUST be included, in the same order as specified in the service description (SCPD) available from the device.

For future extensibility and according to the requirements in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values.

Subject to the constraints defined in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", control points and devices MUST ignore any XML comments or XML processing instructions embedded in action responses that they do not understand.

On a multi-homed device, all fully qualified URLs contained in response arguments that refer to resources on the device MUST be reachable on the UPnP-enabled interface on which the response message is sent.

3.2.3 UPnP Action Schema

The UPnP Action Schema defines the structures and data types used in the body of UPnP actions and action responses. As explained with the UPnP Device and Service Schemas, the UPnP Action Schema is written in XML syntax according to the conventions of XML Schema (Part 1: Structures, Part 2: Datatypes). The UPnP Action Schema is defined within a UPnP service template; however, the schema MUST conform to the format as defined in appendix B.3, "UPnP Control Schema". The elements it defines are used in actions and action responses.

3.2.4 Recommendations and additional requirements

Control points and devices MUST ignore any XML comments or XML processing instructions they may receive that they do not understand.

XML namespace prefixes do not have to be the specific examples given above (e.g., "s" or "u"); they can be any value that obeys the rules of the general XML namespace mechanism; control points MUST accept responses that use other legal XML namespace prefixes.

If an action has no "out" arguments, it is valid to combine the opening and closing XML tags (e.g., "<actionNameResponse/>" instead of "<actionNameResponse></actionNameResponse>").

When the value of any argument contains one or more characters reserved as markup (such as ampersand ("&") or less than ("<")), the text MUST be escaped in accordance with the provisions of section 2.4 of the XML specification and each such character replaced with the equivalent numeric representation or string (such as "&" or "<"). Such characters appearing in URLs MAY also be percent-encoded in accordance with the URL percent-encoding rules specified in sections 2.1 and 2.4 of RFC 3986.

3.2.5 Action error response

Where the normal outcome of processing a SOAP message would have resulted in the transmission of a SOAP response, but rather a SOAP Fault is generated instead, a receiver MUST transmit a SOAP Fault message in place of the response. If the service encounters an error while invoking the action sent by a control point, the service MUST send a response within 30 seconds, including expected transmission time. Out arguments MUST only be used to return data and MUST NOT be used to convey error information. Error responses MUST be sent using the following format. The following two examples illustrate an error response using the CONTENT-LENGTH header and an error response using chunked encoding. Values in *italics* are placeholders for actual values.

Note: XML namespace prefixes do not have to be the specific examples shown below (e.g., "s" or "u"); they can be any value that obeys the rules of the general XML namespace mechanism; control points MUST error responses that use other legal XML namespace prefixes.

Error response using the CONTENT-LENGTH header – Valid using HTTP/1.0 and HTTP/1.1

```
HTTP/1.0 500 Internal Server Error
CONTENT-TYPE: text/xml; charset="utf-8"
DATE: when response was generated
SERVER: OS/version UPnP/1.1 product/version
CONTENT-LENGTH: bytes in body
<?xml version="1.0"?>
<s:Envelope
  xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
  s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
  <s:Body>
    <s:Fault>
      <faultcode>s:Client</faultcode>
      <faultstring>UPnPError</faultstring>
      <detail>
        <UPnPError xmlns="urn:schemas-upnp-org:control-1-0">
          <errorCode>error code</errorCode>
          <errorDescription>error string</errorDescription>
        </UPnPError>
      </detail>
    </s:Fault>
  </s:Body>
</s:Envelope>
```

Error response using chunked encoding – Valid using HTTP/1.1 only

```
HTTP/1.1 500 Internal Server Error
TRANSFER-ENCODING: "chunked"
CONTENT-TYPE: text/xml; charset="utf-8"
DATE: when response was generated
SERVER: OS/version UPnP/1.1 product/version

Length of first chunk in hexadecimal notation
<?xml version="1.0"?>
<s:Envelope
  xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
  s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
  <s:Body>
    <s:Fault>
      <faultcode>s:Client</faultcode>
      <faultstring>UPnPError</faultstring>
      <detail>
        <UPnPError xmlns="urn:schemas-upnp-org:control-1-0">
          <errorCode>error code</errorCode>
          <errorDescription>error string</errorDescription>
        </UPnPError>
      </detail>
    </s:Fault>
  </s:Body>
Length of second chunk in hexadecimal notation
</s:Envelope>
0
```

Listed below are details for the response line, header fields, and body elements appearing in the listing above. HTTP field names are not case sensitive. All HTTP field values and XML element names are case sensitive; XML values are not case sensitive except where noted. Except where noted, REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once.

Response line

HTTP/1.1

The highest version supported by the origin server that is compatible with the control point that issued the request.

For example, if the control point specified support for HTTP/1.0 in the request, the response MUST contain HTTP/1.0.

500 Internal Server Error

HTTP defined status code indicating that an error has been detected.

Header fields

CONTENT-LANGUAGE

PROHIBITED. The CONTENT-LANGUAGE header field MUST NOT be used in control messages.

CONTENT-LENGTH

REQUIRED if Origin Server does not close the session after sending the response AND Origin Server does not send the response using chunked encoding.

PROHIBITED if Origin Server sends the response using chunked encoding. OPTIONAL otherwise.

Field value specifies the length of the body in bytes. Integer.

TRANSFER-ENCODING

OPTIONAL for HTTP/1.1 and above. Field value specifies whether the response is chunked encoded by having field value "chunked" (in the example above the body is sent in 2 chunks). MUST NOT be specified if CONTENT-LENGTH header field is present.

CONTENT-TYPE

REQUIRED. Field value MUST be "text/xml; charset="utf-8" ".

DATE

RECOMMENDED. Field value contains date when response was generated. "rfc1123-date" as defined in RFC 2616.

SERVER

REQUIRED. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, "SERVER: *unix/5.1 UPnP/1.1 MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

Body

<Envelope>

REQUIRED element defined by SOAP. xmlns namespace attribute MUST be "<http://schemas.xmlsoap.org/soap/envelope/>". MUST include [encodingStyle](#) attribute with value "<http://schemas.xmlsoap.org/soap/encoding/>". A receiver MUST generate a fault if it encounters a message whose document element has a local name of "Envelope" but a namespace name that is not "<http://schemas.xmlsoap.org/soap/envelope/>". Contains the following sub elements:

<Body>

REQUIRED element defined by SOAP. MUST be qualified with SOAP namespace. Contains the following sub element:

<Fault>

REQUIRED element defined by SOAP. Error encountered while invoking action. MUST be qualified with SOAP namespace. Contains the following sub elements:

<faultcode>

REQUIRED element defined by SOAP. Value MUST be qualified with the SOAP namespace. MUST be "[Client](#)" for DCP specific errors. When MANDATORY header XML elements within the SOAP header cannot be processed it MUST be the SOAP fault code "[MustUnderstand](#)".

<faultstring>

REQUIRED element defined by SOAP. MUST be "[UPnPError](#)" for DCP specific errors.

<detail>

REQUIRED element defined by SOAP. Contains the following subelement:

<UPnPError>

REQUIRED element for DCP specific errors. MAY be empty for other errors. Contains the following subelements:

<errorCode>

REQUIRED element defined by UDA. Code identifying what error was encountered. See Table 3-3, "UPnP Defined Action error codes" for values. Integer.

<errorDescription>
 RECOMMENDED element defined by UDA. Short description. See Table 3-3, "UPnP Defined Action error codes" for RECOMMENDED values; other values MAY be used by vendors. Human-readable string. RECOMMENDED < 256 characters.

The following table summarizes defined error types and the corresponding value for the <errorCode> and <errorDescription> elements.

Table 3-3: UPnP Defined Action error codes

ErrorCode	errorDescription	Description
401	Invalid Action	No action by that name at this service.
402	Invalid Args	Could be any of the following: not enough in args, args in the wrong order, one or more in args are of the wrong data type.
403	<i>(Do Not Use)</i>	<i>(This code has been deprecated.)</i>
501	Action Failed	MAY be returned if current state of service prevents invoking that action.
600	Argument Value Invalid	The argument value is invalid
601	Argument Value Out of Range	An argument value is less than the minimum or more than the maximum value of the allowed value range, or is not in the allowed value list.
602	Optional Action Not Implemented	The requested action is optional and is not implemented by the device.
603	Out of Memory	The device does not have sufficient memory available to complete the action. This MAY be a temporary condition; the control point MAY choose to retry the unmodified request again later and it MAY succeed if memory is available.
604	Human Intervention Required	The device has encountered an error condition which it cannot resolve itself and required human intervention such as a reset or power cycle. See the device display or documentation for further guidance.
605	String Argument Too Long	A string argument is too long for the device to handle properly.
606-612 ⁴	<i>Reserved</i>	These ErrorCodes are reserved for UPnP DeviceSecurity.
613-699	<i>TBD</i>	Common action errors. Defined by UPnP Forum Technical Committee.
700-799	<i>TBD</i>	Action-specific errors defined by UPnP Forum working committee.
800-899	<i>TBD</i>	Action-specific errors for non-standard actions. Defined by UPnP vendor.

3.2.6 UPnP Error Schema

The UPnP Error Schema defines the structures and data types used in the body of UPnP error messages. As with the UPnP Device and Service Schemas, the UPnP Error Schema is written in XML syntax and according to the conventions of XML Schema (Part 1: Structures, Part 2: Datatypes). Appendix B.4, "UPnP Error Schema" contains a listing of this schema. The elements it defines are used in error messages.

For future extensibility and according to the requirements in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values.

Subject to the constraints defined in section 2.7, “Non-standard vendor extensions” and section 2.8, “UPnP Device Schema”, control points and devices MUST ignore any XML comments or XML processing instructions embedded in UPnP device and service descriptions that they do not understand.

XML namespace prefixes do not have to be the specific examples given above (e.g., “s” or “u”); they can be any value that obeys the rules of the general XML namespace mechanism; control points MUST accept responses that use other legal XML namespace prefixes.

3.3 Query for variable

The QueryStateVariable action has been deprecated by the UPnP Forum and MUST NOT be used by control points except in limited testing scenarios. Working committees and vendors MUST explicitly define actions for querying of state variables for which this capability is desired. Such explicit query actions MAY include multiple state variables, if desired. For the full definition of QueryStateVariable see the UPnP 1.0 specification.

3.4 References

RFC 1123

Includes format for dates, for, e.g., HTTP DATE header field. Available at: <http://www.ietf.org/rfc/rfc1123.txt>.

RFC 2616

HTTP: Hypertext Transfer Protocol 1.1. Available at: <http://www.ietf.org/rfc/rfc2616.txt>.

RFC 2774

HTTP Extension Framework. Available at: <http://www.ietf.org/rfc/rfc2774.txt>.

RFC 3986

Uniform Resource Identifiers (URI): Generic Syntax. Available at: <http://www.ietf.org/rfc/rfc3986.txt>.

SOAP

Simple Object Access Protocol. Available at: <http://www.w3.org/TR/2000/NOTE-SOAP-20000508>.

XML

Extensible Markup Language. Available at: <http://www.w3.org/XML>.

XML Schema (Part 1: Structures, Part 2: Datatypes)

Available at: <http://www.w3.org/TR/xmlschema-1>, <http://www.w3.org/TR/xmlschema-2>.

4 Eventing

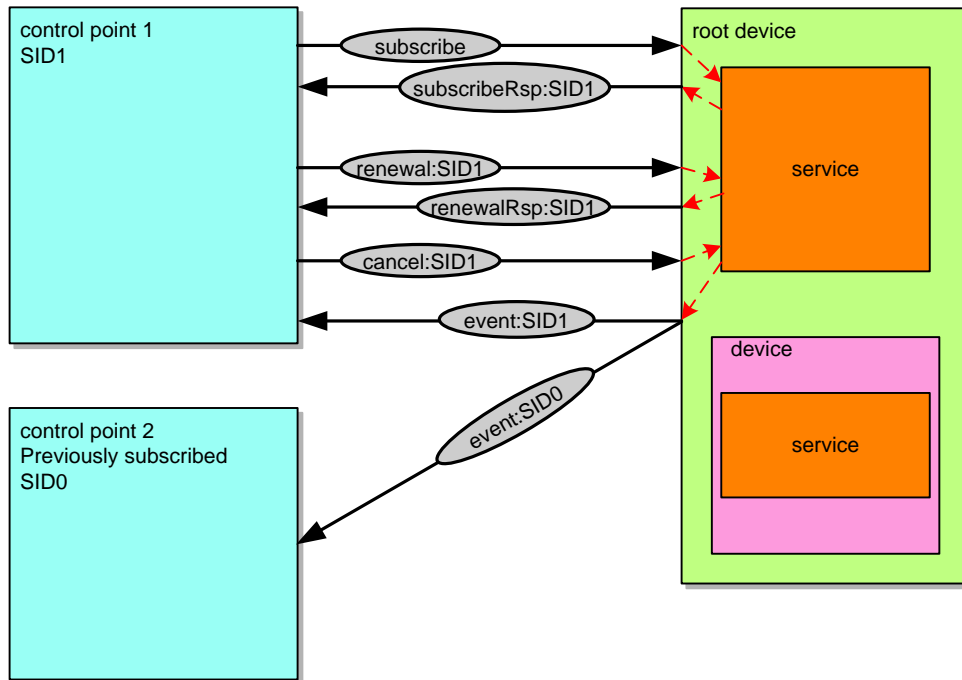
[Normative]

Eventing is Step 4 in UPnP™ networking. Eventing comes after addressing (Step 0) where devices get a network address, after discovery (Step 1) where control points find interesting device(s), and after description (Step 2) where control points learn about device capabilities. Eventing is intimately linked with control (Step 3) where control points send actions to devices. Through eventing, control points listen to state changes in device(s). Control and eventing are complementary to presentation (Step 5) where control points display a user interface provided by device(s).

After a control point has (1) discovered a device and (2) retrieved a description of the device and its services, the control point has the essentials for eventing. As section 2, "Description" explains, a UPnP service description includes a list of actions the service responds to and a list of variables that model the state of the service at run time. If one or more of these state variables are evented, then the service publishes updates when these variables change, and a control point MAY subscribe to receive this information. Two types of eventing are supported by this specification: unicast eventing as found in version 1.0 of the UPnP specification where a control point MAY subscribe to receive variable updates; and multicast eventing where variables MAY be defined as multicast events and can be additionally sent over UDP to any listening device on the multicast event address. This form of eventing is useful when events which are not relevant to a specific UPnP interaction SHOULD be delivered to *control points* to inform users, and when multiple *controlled devices* MAY want to inform multiple other *controlled devices*. Throughout this section, *publisher* refers to the source of the events (typically a device's service), *subscriber* refers to the destination of events (typically a control point), and the term *receiver* refers to the listener of multicast events (typically a control point, but MAY also be a controlled device).

4.1 Unicast eventing

Figure 4-1: Unicast eventing architecture



To subscribe to eventing, a subscriber sends a *subscription message*. If the subscription is accepted, the publisher responds with a duration for the subscription. To keep the subscription active, a subscriber **MUST** renew its subscription before the subscription expires. When a subscriber no longer needs eventing from a publisher, the subscriber **SHOULD** cancel its subscription. This section explains subscription, renewal, and cancellation messages in detail below.


The publisher notes changes to state variables by sending *event messages*. Event messages contain the names of one or more state variables and the current value of those variables, expressed in XML. A special *initial event message* is sent when a subscriber first subscribes; this event message contains the names and values for all evented variables and allows the subscriber to initialize its model of the state of the service. To support scenarios with multiple control points, eventing can be used to keep interested control points informed about the effects of actions performed by other control points or using other mechanisms for device control (such as front panel controls). All subscribers are sent all event messages, subscribers receive event messages for all evented variables (not just some), and event messages are sent no matter why the state variable changed (either in response to a requested action or because the state the service is modeling changed). This section explains the format of event messages in detail below.

Some state variables **MAY** change value too rapidly for eventing to be useful. One alternative is to filter, or moderate, the number of event messages sent due to changes in a variable's value. Some state variables may contain values too large for eventing to be useful; for this, or other reasons, a service **MAY** designate one or more state variables as *non evented* and never send event messages to subscribers. To determine the current value for such non-evented variables, control points **MUST** poll the

service explicitly, presuming that an action is provided to obtain the value of the state variable. This section explains how variable eventing is described within a service description.

To send and receive subscription and event messages, control points and services use the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

Figure 4-2: Unicast eventing protocol stack



At the highest layer, subscription and event messages contain vendor-specific information like URLs for subscription and duration of subscriptions or specific variable values. Moving down the stack, vendor content is supplemented by information from a UPnP Forum working committee, like service identifiers or variable names. Messages from the layers above are hosted in UPnP-specific protocols, defined in this document. In turn, the above messages are delivered via HTTP that has been extended using additional methods and header fields. The HTTP messages are delivered via TCP over IP. For reference, colors in [square brackets] above indicate which protocol defines specific header fields in the subscription messages listed below.

The remainder of this section first explains subscription, including details of subscription messages, renewal messages, and cancellation messages. Second, it explains in detail how event messages are formatted and sent to control points, and the initial event message. Finally, it explains the UPnP Device and Service Schemas as they pertain to eventing.

The generic requirements on HTTP usage in UPnP 1.1 (as defined in section 2.1, “Generic requirements on HTTP usage” of this document) **MUST** be followed by devices and control points that implement eventing.

Services that use evented complex datatypes **MUST** follow the requirements in section 2.5, “Service description”.

4.1.1 Subscription

A service has eventing if and only if one or more of the state variables are evented.

If a service has eventing, it publishes event messages to interested subscribers. The publisher maintains a list of subscribers, keeping for each subscriber the following information.

unique subscription identifier

REQUIRED. MUST be unique over the lifetime of the subscription, however long or short that may be. Generated by publisher in response to subscription message. **RECOMMEND** universally-unique identifiers to ensure uniqueness. Single URI.

delivery URL for event messages

REQUIRED. Provided by subscriber in subscription message. Single URL.

event key

REQUIRED. Key is 0 for initial event message. Key MUST be sequentially numbered for each subsequent event message; subscribers can verify that no event messages have been lost if the subscriber has received sequentially numbered event keys. MUST wrap from 4294967295 to 1 (32-bit unsigned decimal integer). Some implementations MAY include leading "0" characters in the event key, which MUST be ignored.

subscription duration

REQUIRED. Amount of time, or duration until subscription expires. Single integer, preceded in subscription messages by the keyword "**Second-**" (no spaces). UPnP 1.0 defines the use of the keyword **infinite** instead of an integer. This keyword is deprecated in UPnP 1.1 (it leads to problems if control points disappear without unsubscribing and is hardly used): UPnP 1.1 control points MUST NOT subscribe using keyword **infinite**, UPnP 1.1 devices MUST NOT set actual subscription durations to "infinite". The presence of **infinite** in a request MUST be silently ignored by a UPnP 1.1 device (the presence of infinite is handled by the device as if the TIMEOUT header field in a request was not present) . The keyword infinite MUST NOT be returned by a UPnP 1.1 device.

HTTP version supported by the subscriber

REQUIRED if the publisher supports chunked encoding of event notification messages, so that chunked messages are *not* sent to subscribers that do not support them.

A multi-homed publisher MUST also maintain information on the UPnP-enabled interface on which each subscription message was received. The same interface MUST be used when publishing event messages to the corresponding subscriber.

The publisher SHOULD accept as many subscriptions as it can reasonably maintain, taking into account that the number of event messages that need to be delivered per event, which increases linearly with the number of subscriptions.

The list of subscribers is updated via subscription, renewal, and cancellation messages explained immediately below and event messages explained later in this section.

To subscribe to eventing for a service, a subscriber sends a *subscription message* containing a URL for the publisher, a service identifier for the publisher, and a delivery URL for event messages. The subscription message MAY also include a requested duration for the subscription. The URL and service identifier for the publisher come from a description message. As section 2, "Description" explains, a description message contains a device description. A device description contains (among other things), for each service, an eventing URL (obtained from the eventSubURL element) and a service identifier (in the serviceId element); these correspond to the URL and service identifier for the publisher, respectively. If eventSubURL is an absolute URL, the fully qualified event subscription URL is the eventSubURL. If eventSubURL is a relative URL, the fully qualified event subscription URL is the URL resolved from eventSubURL in accordance with section 5 of RFC 3986, using either the URLBase element, if specified, or the URL from which the device description was retrieved as the base URL. If the eventSubURL is empty, no subscriptions are possible. The fully qualified event subscription URL for the publisher MUST be unique to a particular service within this device. A multi-homed control point that sends the subscription message on a particular UPnP-enabled interface MUST use the fully qualified eventing URL from the description document received on that UPnP-enabled interface. The delivery URL contained in the subscription message MUST be reachable on that interface.

The subscription message is a request to receive all event messages. No mechanism is provided to subscribe to event messages on a variable-by-variable basis. A subscriber is sent all event messages from the service. This is one factor to be considered when designing a service.

If the subscription is accepted, the publisher responds with a unique identifier for this subscription and a duration for this subscription. A duration SHOULD be chosen that matches assumptions about how frequently control points are removed from the network; if control points are removed every few minutes, then the duration SHOULD be similarly short, allowing a publisher to rapidly deprecate any expired subscribers; if control points are expected to be semi-permanent, then the duration SHOULD be very long, minimizing the processing and traffic associated with renewing subscriptions.

As soon as possible after the subscription is accepted, the publisher also sends the first, or *initial* event message to the subscriber. This message includes the names and current values for all evented variables. (The data type and range for each variable is described in a service description. Section 2, "Description" explains this in more detail.) This initial event message is always sent, even if the control point unsubscribes before it is delivered. The device MUST insure that the control point has received the response to the subscription request before sending the initial event message, to insure that the control point has received the SID (subscription ID) and can thereby correlate the event message to the subscription.

To keep the subscription active, a subscriber MUST renew its subscription before the subscription expires by sending a renewal message. The renewal message is sent to the same URL as the subscription message, but the renewal message does not include a delivery URL for event messages; instead the renewal message includes the subscription identifier. The response for a renewal message is the same as one for a subscription message.

If a subscription expires, the subscription identifier becomes invalid, and the publisher stops sending event messages to the subscriber and can clean up its list of subscribers. If the subscriber tries to send any message other than a subscription message, the publisher MUST reject the message because the subscription identifier is invalid.

When a subscriber no longer needs eventing from a particular service, the subscriber SHOULD cancel its subscription. Canceling a subscription generally reduces service, control point, and network load. If a subscriber is removed abruptly from the network, it might be impossible to send a cancellation message. As a fallback, the subscription will eventually expire on its own unless renewed.

It is strongly RECOMMENDED that subscribers monitor discovery messages from the publisher. If the publisher cancels its advertisements or if the value of the BOOTID.UPNP.ORG is increased without a prior `ssdp:update` message with a matching NEXTBOOTID.UPNP.ORG field value, subscribers MUST assume that their subscriptions have been cancelled.

Below is an explanation of the specific format of requests, responses, and errors for subscription, renewal, and cancellation messages.

4.1.2 SUBSCRIBE with NT and CALLBACK

For each service in a device, a description message contains an event subscription URL (obtained from the eventSubURL sub element of service element in the device description) and the UPnP service identifier (serviceId sub element in service element in device description). To subscribe to eventing for a particular service, a subscription message is sent to that service's fully

qualified event subscription URL. If eventSubURL is an absolute URL, the fully qualified event subscription URL is the eventSubURL. Otherwise, if eventSubURL is a relative URL, the fully qualified event subscription URL is the URL resolved from eventSubURL in accordance with section 5 of RFC 3986, using either the URLBase element, if specified, or the URL from which the device description was retrieved as the base URL. The message contains that service's identifier as well as a delivery URL for event messages. A multi-homed control point that sends the subscription message on a particular UPnP-enabled interface MUST use the fully qualified eventing URL from the description document received on that interface. The delivery URL contained in the subscription message MUST be reachable on that interface. A subscription message MAY also include a requested subscription duration.

To subscribe to eventing for a service, a subscriber MUST send a request with method SUBSCRIBE and NT and CALLBACK header fields in the following format. Values in *italics* are placeholders for actual values.

```
SUBSCRIBE publisher path HTTP/1.1
HOST: publisher host:publisher port
USER-AGENT: OS/version UPnP/1.1 product/version
CALLBACK: <delivery URL>
NT: upnp:event
TIMBOUT: Second-requested subscription duration
```

(No body for request with method SUBSCRIBE, but note that the message MUST have a blank line following the last HTTP header field.)

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Request line

SUBSCRIBE

Method to initiate or renew a subscription.

publisher path

Path component of the fully qualified event subscription URL. Single, absolute path (see also RFC 2616, section 3.2.2).

HTTP/1.1

The version supported by the control point. (Note: the control point MUST implement all MANDATORY components of the version specified). MAY be any HTTP version that is backwards compatible to HTTP/1.0 (like HTTP/1.1).

Header fields

HOST

REQUIRED. Field value contains domain name or IP address and optional port components of the fully qualified event subscription URL. If the port is missing or empty, port 80 is assumed.

USER-AGENT

OPTIONAL. Specified by UPnP vendor. String. Field value MUST begin with the following "product tokens" (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be UPnP/1.1, and the third token identifies the product using the form *product name/product version*. For example, "USER-AGENT: *unix/5.1 UPnP/1.1 MyProduct/1.0*". Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1. See section 2.5, "Service description".

CALLBACK

REQUIRED. Field value contains location to send event messages to. Defined by UPnP vendor. If there is more than one URL, when the service sends events, it will try these URLs in order until one succeeds. One or more URLs each enclosed by angle

brackets (“<” and “>”). Each URL MUST be an HTTP over TCP URL (prefixed by “http://”). The device MUST NOT truncate this URL in any way; if insufficient memory is available to store the entire CALLBACK URL, the device MUST reject the subscription. At least one of the delivery URLs MUST be reachable by the device.

NT

REQUIRED. Field value contains Notification Type. MUST be [upnp:event](#).

SID

(No SID header field is used to subscribe.)

TIMEOUT

RECOMMENDED. Field value contains requested duration until subscription expires. Consists of the keyword [Second-](#) followed (without an intervening space) by an integer. UPnP 1.0 defined that the integer can be replaced by the keyword [infinite](#). This has been deprecated in UPnP 1.1: UPnP 1.1 control points MUST NOT subscribe using keyword [infinite](#).

If there are enough resources to maintain the subscription, the publisher SHOULD accept it. To accept the subscription, the publisher assigns a unique identifier for the subscription, assigns a duration for the subscription, and sends an initial event message (explained in detail later in this section). To accept a subscription request, a publisher MUST send a response in the following format within 30 seconds, including expected transmission time. A multi-homed publisher MUST send the response on the same UPnP-enabled interface on which the subscription message was received. Values in *italics* are placeholders for actual values.

```
HTTP/1.1 200 OK
DATE: when response was generated
SERVER: OS/version UPnP/1.1 product/version
SID: uuid:subscription-UUID
CONTENT-LENGTH: 0
TIMEOUT: Second-actual subscription duration
```

(No body for response to a request with method SUBSCRIBE, but note that the message MUST have a blank line following the last HTTP header field.)

If the device sends the response over HTTP/1.0 without setting the KeepAlive token, or over HTTP/1.1 with the CONNECTION: close header field, the device MUST insure that the TCP FIN flag is sent BEFORE sending the initial event message. In all other cases, (unless the response is chunked), a CONTENT-LENGTH MUST be specified, (and set to 0), prior to sending the initial event.

Listed below are details for header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Response line

HTTP/1.1

The highest version supported by the origin server that is compatible with the control point that issued the request. For example, if the control point specified support for HTTP/1.0 in the request, the response MUST contain HTTP/1.0.

200 OK

HTTP defined status code indicating that no HTTP errors were detected..

Header fields

DATE

RECOMMENDED. Field value contains date when the response was generated. “rfc1123-date” as defined in RFC 2616.

SERVER

REQUIRED. Specified by UPnP vendor. String. Field value MUST begin with the following “product tokens” (defined by HTTP/1.1). The first product token identifies the operating system in the form *OS name/OS version*, the second token represents the UPnP version and MUST be [UPnP/1.1](#), and the third token identifies the product using the form *product name/product version*. For example, “SERVER: *unix/5.1 UPnP/1.1 MyProduct/1.0*”. Control points MUST be prepared to accept a higher minor version number of the UPnP version than the control point itself implements. For example, control points implementing UDA version 1.0 will be able to interoperate with devices implementing UDA version 1.1.

SID

REQUIRED. Field value contains Subscription Identifier. MUST be universally unique. MUST begin with `uuid:.`. Defined by UPnP vendor. See section 1.1.4, “UUID format and RECOMMENDED generation algorithms” for the MANDATORY UUID format.

TIMEOUT

REQUIRED. Field value contains actual duration until subscription expires. Keyword “**Second-**” followed by an integer (no space). SHOULD be greater than or equal to 1800 seconds (30 minutes).

CONTENT-LENGTH

REQUIRED if TCP FIN flag cannot be guaranteed to be sent before the initial event is sent. MUST have field value “0”.

If a publisher cannot accept the subscription, or if there is an error with the subscription request, the publisher MUST send a response with one of the following errors. The response MUST be sent within 30 seconds, including expected transmission time.

Table 4-4: HTTP Status Codes indicating a Subscription Error

ErrorCode	errorDescription	Description
400	Incompatible header fields	An SID header field and one of NT or CALLBACK header fields are present.
412	Precondition Failed	CALLBACK header field is missing or does not contain a valid HTTP URL; or the NT header field does not equal upnp:event .
5xx	Unable to accept renewal	If the publisher is unable to accept a renewal, it MUST respond with an appropriate 500-series HTTP status code.

Other errors, including other HTTP status codes, MAY be returned by layers in the protocol stack below the UPnP protocols.

Consult documentation on those protocols for details.

4.1.3 Renewing a subscription with SUBSCRIBE with SID

To renew a subscription to eventing for a particular service, a renewal message is sent to that service's fully qualified event subscription URL (See section 4.1.2, “SUBSCRIBE with NT and CALLBACK”). However, unlike an initial subscription message, a renewal message does not contain either the service's identifier nor a delivery URL for event messages. Instead, the message contains the *subscription* identifier assigned by the publisher, providing an unambiguous reference to the subscription to be renewed. Like a subscription message, a renewal message MAY also include a requested subscription duration. A multi-homed control point MUST send the renewal message using the same pair of UPnP-enabled interfaces used for the initial subscription.

The renewal message uses the same method as the subscription message, but the two messages use a disjoint set of header fields; renewal uses SID and subscription uses NT and CALLBACK. A message that includes SID and either of NT or CALLBACK header fields is an error.

To renew a subscription to eventing for a service, a subscriber MUST send a request with method SUBSCRIBE and SID header field in the following format. Values in *italics* are placeholders for actual values.

```
SUBSCRIBE publisher path HTTP/1.1
HOST: publisher host:publisher port
SID: uuid:subscription UUID
TIMEOUT: Second-requested subscription duration
```

(No body for method with request SUBSCRIBE, but note that the message MUST have a blank line following the last HTTP header field.)

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Request line

SUBSCRIBE

Method to initiate or renew a subscription.

publisher path

Path component of the fully qualified event subscription URL. Single, absolute path (see also RFC 2616, section 3.2.2).

HTTP/1.1

The version supported by the control point. (Note: the control point MUST implement all MANDATORY components of the version specified). MAY be any HTTP version that is backwards compatible to HTTP/1.0 (like HTTP/1.1)

Header fields

HOST

REQUIRED. Field value contains domain name or IP address and optional port components of fully qualified event subscription URL. If the port is missing or empty, port 80 is assumed.

CALLBACK

(No CALLBACK header field is used to renew an event subscription.)

NT

(No NT header field is used to renew an event subscription.)

SID

REQUIRED. Field value contains Subscription Identifier. MUST be the subscription identifier assigned by publisher in response to subscription request. MUST be universally unique. MUST begin with *uuid:*. Defined by UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms for the MANDATORY UUID format.

TIMEOUT

RECOMMENDED. Field value contains requested duration until subscription expires. Keyword **Second-** followed by an integer (no space). UPnP 1.0 defined that the integer can be replaced by the keyword **infinite**. This has been deprecated in UPnP 1.1: UPnP 1.1 control points MUST NOT subscribe using keyword **infinite**. See reference above.

To accept a renewal, the publisher reassigns a duration for the subscription and MUST send a response in the same format and with the same conditions as a response to a request for a new subscription, except that the initial event message is not sent again.

If a publisher cannot accept the renewal, or if there is an error with the renewal request, the publisher MUST send a response with one of the following errors. The response MUST be sent within 30 seconds, including expected transmission time.

Table 4-5: HTTP Status Codes indicating a Resubscription Error

ErrorCode	errorDescription	Description
400	Incompatible header fields	An SID header field and one of NT or CALLBACK header fields are present.
412	Precondition Failed	An SID does not correspond to a known, un-expired subscription; or the SID header field is missing or empty.
5xx	Unable to accept renewal	If the publisher is unable to accept a renewal, it MUST respond with an appropriate 500-series HTTP status code.

Other errors, including other HTTP status codes, MAY be returned by layers in the protocol stack below the UPnP protocols. Consult documentation on those protocols for details.

4.1.4 Canceling a subscription with UNSUBSCRIBE

When eventing is no longer needed from a particular service, a cancellation message SHOULD be sent to that service's fully qualified event subscription URL (see section 4.1.2, "SUBSCRIBE with NT and CALLBACK"). The message contains the subscription identifier. A multi-homed control point MUST send the cancellation message using the same pair of UPnP-enabled interfaces used for the initial subscription. Canceling a subscription generally reduces service, control point, and network load. If a control point is removed abruptly from the network, it might be impossible to send a cancellation message. As a fallback, the subscription will eventually expire on its own unless renewed.

To explicitly cancel a subscription to eventing for a service, a subscriber MUST send a request with method UNSUBSCRIBE in the following format. Values in *italics* are placeholders for actual values.

```
UNSUBSCRIBE publisher path HTTP/1.1
HOST: publisher host:publisher port
SID: uuid:subscription UUID
```

(No body for request with method UNSUBSCRIBE, but note that the message MUST have a blank line following the last HTTP header field.)

Listed below are details for the request line and header fields appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted.

Request line

UNSUBSCRIBE
Method to cancel a subscription.

publisher path
Path component of the fully qualified event subscription URL. Single, absolute path (see also RFC 2616, section 3.2.2).

HTTP/1.1
The version supported by the control point. (Note: the control point MUST implement all MANDATORY components of the version specified). MAY be any HTTP version that is backwards compatible to HTTP/1.0 (like HTTP/1.1)

Header fields

HOST

REQUIRED. Field value contains domain name or IP address and optional port components of fully qualified event subscription URL. If the port is missing or empty, port 80 is assumed.

CALLBACK

(No CALLBACK header field is used to cancel an event subscription.)

NT

(No NT header field is used to cancel an event subscription.)

SID

REQUIRED. Field value contains Subscription Identifier. MUST be the subscription identifier assigned by publisher in response to subscription request. Must be universally unique. Must begin with uuid:. Defined by UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

TIMEOUT

(No TIMEOUT header field is used to cancel an event subscription.)

To cancel a subscription, a publisher MUST send a response in the following format within 30 seconds, including expected transmission time.

```
HTTP/1.1 200 OK
```

Response line

HTTP/1.1

The highest version supported by the origin server that is compatible with the control point that issued the request. For example, if the control point specified support for HTTP/1.0 in the request, the response MUST contain HTTP/1.0.

200 OK

HTTP defined status code indicating that no HTTP errors were detected.

If there is an error with the cancellation request, the publisher MUST send a response with one of the following errors. The response MUST be sent within 30 seconds, including expected transmission time.


Table 4-6: HTTP Status Codes indicating a Cancel Subscription Error

ErrorCode	errorDescription	Description
400	Incompatible header fields	An SID header field and one of NT or CALLBACK header fields are present.
412	Precondition Failed	An SID does not correspond to a known, un-expired subscription; or the SID header field is missing or empty.

Other errors, including other HTTP status codes, MAY be returned by layers in the protocol stack below the UPnP protocols. Consult documentation on those protocols for details.

4.2 Multicast Eventing

Figure 4-3: Multicast eventing architecture



The publisher MAY note changes to state variables by sending *multicast event messages*. Multicast event messages contain the names of one or more state variables and the current value of those variables, expressed in XML. To send and receive multicast event messages, control points and services use the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

Figure 4-4: Multicast eventing protocol stack

<i>UPnP vendor [purple-italic]</i>
<i>UPnP Forum [red-italic]</i>
UPnP Device Architecture [green-bold]
Multicast Eventing [navy-bold]
UDP [black]
IP [black]

At the highest layer, multicast event messages contain vendor-specific information like vendor-specific state variable or specific variable values. Moving down the stack, vendor content is supplemented by information from a UPnP Forum working committee, such as service identifiers or variable names. Messages from the layers above are hosted in UPnP-specific protocols to transport events in a similar format to unicast UPnP events, but over a multicast address where subscriptionless eventing fits the desired usage. These messages are based on the HTTP protocol header and body format, but are not HTTP compliant because they are defined over UDP sockets. Throughout this section, the same formatting and extension rules for SSDP as set forth in section 1.1.2, “SSDP message header fields” and section 1.1.3, “SSDP header field extensions” are used to give HTTP-like header field formatting. In addition, services that use evented complex datatypes MUST follow the requirements in section 2.5, “Service description”. Lastly, like SSDP, to limit network congestion, the time-to-live (TTL) of each IP packet for each multicast message SHOULD default to 2 and SHOULD be configurable. This SHOULD be the same value as that used in SSDP. When the TTL is greater than 1, it is possible for multicast messages to traverse multiple routers; therefore control points and devices using non-AutoIP addresses MUST send an IGMP Join message so that routers will forward multicast messages to them (this is not necessary when using an Auto-IP address since packets with Auto-IP addresses will not be forwarded by routers).

Multicast eventing is inherently unreliable since it is based on UDP. In addition, there will be a greater possibility of message loss with greater packet size. To increase the probability of successful transmission, each message MAY be retransmitted one or more times. Therefore, UPnP working committees MUST specify the event size and event retransmission rules, based on their need for reliability.

4.3 Event messages

A service publishes changes to certain state variables by sending event messages. These messages contain the names of one or more state variables and the current value of those variables. Event messages MUST be sent in a timely manner so that subscribers are accurately informed about the state of the service and can provide a responsive user interface. If the value of more than one variable is changing at the same time, the publisher SHOULD bundle these changes into a single event message to reduce processing and network traffic.

As explained above, an initial event message is sent when a subscriber first subscribes; this event message contains the names and values for all evented variables and allows the subscriber to initialize its model of the state of the service. This message SHOULD be sent as soon as possible after the publisher accepts a subscription. This message MUST be sent, even if the control point unsubscribes before the message is delivered. Subscription does not cause multicast event messages.

Multicast event messages are constrained to being transported in a single UDP payload. This consideration is important when identifying variables that are to be multicast. If the cumulative size of the variables that are eligible for being sent by multicast exceeds the UDP packet’s capacity, it may be necessary to send multiple, distinct multicast events.

Both unicast and multicast event messages are tagged with an event key. In unicast eventing, a separate event key MUST be maintained by the publisher for each subscription to facilitate error detection (as explained below). The event key for a

subscription is initialized to 0 when the publisher sends the initial event message. For each subsequent event message, the publisher increments (by one) the event key for a subscription, and includes that updated key in the event message. The event key for multicast events is also initialized to 0 when the publisher sends the initial event message. For each subsequent multicast event message, the publisher increments (by one) the event key for the multicast events, and includes that updated key in the event message. Any implementation of event keys MUST handle overflow and wrap the event key from 4294967295 back to 1 (not 0). Unicast subscribers and multicast receivers MUST also handle this special case when the next event key is not an increment of the previous key. The key MUST be implemented as a 4 Byte (32 bit) unsigned integer.

All UPnP event messages MUST be encoded using UTF-8.

4.3.1 Error Cases

For unicast eventing, the publisher MUST send all event messages to the subscriber until the subscription expires even when the subscriber fails to respond. When a subscriber has missed one or more event messages, the subscriber MAY synchronize with the device's evented state by unsubscribing and re-subscribing. By doing so, the subscriber will get a new subscription identifier, a new initial event message, and a new event key.

For multicast eventing, since UDP is inherently unreliable, retransmission of a multicast event message (using the same SEQ field value) can increase the reliability. The receiver MUST interpret the same SEQ field value from separate multicast event messages from a same service (identified by USN field value) as being the exactly the same message sent multiple times and MUST therefore ignore such duplicates. Some state variables may change value too rapidly for some environments, for example enterprises. Working committees MUST specify traffic constraints for the DCP given these concerns and guidelines. Working committees SHOULD consider both the interval for transmission of multicast events per event type (LVL:) and the retransmission rules for particular event instances.

4.3.2 Unicast eventing: Event messages: NOTIFY

To send an event message, a publisher MUST send a request with method NOTIFY using the following format. The following two examples illustrate an event message using the CONTENT-LENGTH header and an event message using chunked encoding. Values in *italics* are placeholders for actual values.

Event messages sent to different subscribers that have the same sequence number MUST contain the same content except for the HOST header field. A multi-homed device MUST send the event message using the same pair of UPnP-enabled interfaces used for the initial subscription.

Note: XML namespace prefixes do not have to be the specific examples shown below (e.g., "s" or "u"); they can be any value that obeys the rules of the general XML namespace mechanism; control points MUST accept event messages that use other legal XML namespace prefixes.

Event message using the CONTENT-LENGTH header – Valid with HTTP/1.0 or HTTP/1.1

```
NOTIFY delivery path HTTP/1.0
HOST: delivery host:delivery port
CONTENT-TYPE: text/xml; charset="utf-8"
NT: upnp:event
NTS: upnp:propchange
SID: uuid:subscription-UUID
SEQ: event key
CONTENT-LENGTH: bytes in body
<?xml version="1.0"?>
<e:propertyset xmlns:e="urn:schemas-upnp-org:event-1-0">
  <e:property>
    <variableName>new value</variableName>
  </e:property>
  Other variable names and values (if any) go here.
</e:propertyset>
```

Event message using chunked encoding – Valid with HTTP 1.1 only

```
NOTIFY delivery path HTTP/1.1
HOST: delivery host:delivery port
CONTENT-TYPE: text/xml; charset="utf-8"
NT: upnp:event
NTS: upnp:propchange
SID: uuid:subscription-UUID
TRANSFER-ENCODING: "chunked"
SEQ: event key

Length of chunk 1 in hexadecimal notation
<?xml version="1.0"?>
<e:propertyset xmlns:e="urn:schemas-upnp-org:event-1-0">
  <e:property>
    <variableName>new value</variableName>
  </e:property>
  Other variable names and values (if any) go here.
</e:propertyset>
0
```

Listed below are details for the request line, header fields, and body elements appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted. All body elements and attributes are case sensitive; body values are not case sensitive except where noted. Except where noted, REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once. In particular, a single **propertyset** element MUST NOT include more than one **property** element that specifies the same *variableName* element; separate event notification messages MUST be used.

Request line

NOTIFY

Method to notify client about event.

delivery path

Path component of delivery URL (CALLBACK header field in subscription message). Destination for event message. Single, absolute path (see also RFC 2616). MUST be from one of the URLs contained in the CALLBACK header field, without truncation or modification.

HTTP/1.1

Highest HTTP version supported by the device. (Note: chunked encoding MUST NOT be used if the control point supports only HTTP 1.0).

Header fields

HOST

REQUIRED. Field value contains domain name or IP address and optional port components of delivery URL (CALLBACK header field in subscription message). If the port is missing or empty, port 80 is assumed.

ACCEPT-LANGUAGE

(No ACCEPT-LANGUAGE header field is used in event messages.)

CONTENT-LENGTH

REQUIRED if Origin Server does not close the session after sending the response AND Origin Server does not send the response using chunked encoding.

PROHIBITED if Origin Server sends the response using chunked encoding. OPTIONAL otherwise. Field value specifies the length of the body in bytes. Integer.

TRANSFER-ENCODING

OPTIONAL for HTTP/1.1 and above. Field value specifies whether the response is chunked encoded by having field value "chunked" (in the example above the body is sent in a single chunk). MUST NOT be specified if CONTENT-LENGTH header field is present.

CONTENT-TYPE

REQUIRED. Field value MUST be "text/xml; charset=utf-8"

NT

REQUIRED. Field value contains Notification Type. MUST be **upnp:event**.

NTS

REQUIRED. Field value contains Notification Sub Type. MUST be **upnp:propchange**.

SID

REQUIRED. Field value contains Subscription Identifier. MUST be universally unique. MUST begin with uuid:. Defined by UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

SEQ

REQUIRED. Field value contains Event Key. MUST be 0 for initial event message. MUST be incremented by 1 for each event message sent to a particular subscriber. To prevent overflow, MUST be wrapped from 4294967295 to 1. 32-bit unsigned value represented as a single decimal integer without leading zeroes (some implementations MAY include leading zeroes, which SHOULD be ignored by the recipient).

Body

<propertyset>

REQUIRED. xmlns namespace attribute MUST be urn:[schemas-upnp-org:event-1-0](#). Contains the following sub element:

<property>

REQUIRED. Repeat once for each variable name and value in the event message. MUST be qualified by the namespace prefix defined in the xmlns attribute of the <propertyset> element. Contains the following sub element:

<variableName>

REQUIRED. Element is name of a state variable that changed (<name> sub element of <stateVariable> element in service description). MUST NOT be qualified with any namespace. Value is the new value for this state variable. Case sensitive. Single data type as specified by UPnP service description.

For future extensibility and according to the requirements in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values. Note that when subscribing to eventing with a service that is of a higher version than what is supported by the control point, event notifications MAY be sent by the service to the control point containing state variable names that are not recognized by the control point. The control point SHOULD discard and ignore such unrecognized state variables within event notification messages.

When the new value of any variable contains one or more characters reserved as markup (such as ampersand ("&") or less than ("<>"), the text MUST be escaped in accordance with the provisions of section 2.4 of the XML specification and each such character replaced with the equivalent numeric representation or string (such as "&" or "<"). Such characters appearing

in URLs that appear as values MAY also be percent-encoded in accordance with the URL percent-encoding rules specified in sections 2.1 and 2.4 of RFC 3986.

On a multi-homed device, all fully-qualified URLs contained in event body that refer to resources on the device MUST be reachable on the UPnP-enabled interface on which the event message is sent.

Subject to the constraints defined in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", control points and devices MUST ignore any XML comments or XML processing instructions embedded in UPnP event messages that they do not understand. Note that because HTTP 1.1 allows use of chunked encoding, some devices MAY send the event notification using chunked encoding if the SUBSCRIBE request specified HTTP 1.1. It is therefore RECOMMENDED that all implementations that include HTTP 1.1 in the SUBSCRIBE request support receiving chunked encoding.

To acknowledge receipt of this event message, a subscriber MUST respond within 30 seconds, including expected transmission time. A multi-homed subscriber MUST send the response using the same pair of UPnP-enabled interfaces used for the event message. If a subscriber does not respond within 30 seconds, or if the publisher is unable to connect to the subscription URL, the publisher SHOULD abandon sending this message to the subscriber but MUST keep the subscription active and send future event messages to the subscriber until the subscription expires or is cancelled. The subscriber MUST send a response in the following format.

```
HTTP/1.1 200 OK
```

Response line

HTTP/1.1

Highest HTTP version supported by the control point that is compatible with the device that sent the event message.

200 OK

HTTP defined status code indicating that no HTTP errors were detected.

(No body for a request with method NOTIFY, but note that the message MUST have a blank line following the last HTTP header field.)

If a device sends an event to a control point using HTTP/1.0 without the KeepAlive token, the control point MUST close the socket after responding. If a device sends an event to a control point using HTTP/1.1 and sets the Connection:CLOSE token, the control point MUST close the socket after responding.

If there is an error with the event message, the subscriber MUST respond with one of the following errors. The response MUST be sent within 30 seconds, including expected transmission time.

Table 4-7: HTTP Status Codes indicating a Notify Error

ErrorCode	errorDescription	Description
400	Bad request	The NT or NTS header field is missing; or the request is malformed.
412	Precondition Failed	An SID does not correspond to a known, un-expired subscription; or the NT header field does not equal upnp:event ; or the NTS header field does not equal upnp:propchange ; or the SID header field is missing or empty.

Other errors, including other HTTP status codes, MAY be returned by layers in the protocol stack below the UPnP protocols. Consult documentation on those protocols for details.

4.3.3 Multicast Eventing: Event messages: NOTIFY

To send a multicast event message, a publisher MUST send a message with method NOTIFY using the following format. The following example illustrates an event message using the CONTENT-LENGTH header. Values in *italics* are placeholders for actual values.

A multi-homed publisher MUST multicast the event message on each of its UPnP-enabled interfaces. Event messages sent on different UPnP-enabled interfaces that have the same sequence number MUST contain the same content except for possibly the HOST header field and any fully-qualified URLs contained in the event body. The HOST header field of an advertisement MUST be the standard multicast eventing address specified for the protocol (IPv4 or IPv6) used on the interface. All fully-qualified URLs contained in the event body that refer to resources on the device MUST be reachable on the UPnP-enabled interface on which the event message is sent.

Note: XML namespace prefixes do not have to be the specific example shown below (e.g., "s" or "u"); they can be any value that obeys the rules of the general XML namespace mechanism; control points MUST accept event messages that use other legal XML namespace prefixes.

Multicast event message using the CONTENT-LENGTH header – Valid with HTTP/1.0 or HTTP/1.1

```
NOTIFY * HTTP/1.0
HOST: 239.255.255.246:7900 *** note the port number is different than SSDP ***
CONTENT-TYPE: text/xml; charset="utf-8"
USN: Unique Service Name for the publisher
SVCID: ServiceID from SCPD
NT: upnp:event
NTS: upnp:propchange
SEQ: monotonically increasing sequence count
LVL: event importance
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or update message
CONTENT-LENGTH: bytes in body

<?xml version="1.0"?>
<e:propertyset xmlns:e="urn:schemas-upnp-org:event-1-0">
  <e:property>
    <variableName>new value</variableName>
  </e:property>
  <!-- Other variable names and values (if any) go here. -->
</e:propertyset>
```

Listed below are details for the request line, header fields, and body elements appearing in the listing above. Field names are not case sensitive. All field values are case sensitive except where noted. All body elements and attributes are case sensitive; body values are not case sensitive except where noted. Except where noted, REQUIRED elements MUST occur exactly once (no duplicates), and RECOMMENDED or OPTIONAL elements MAY occur at most once. In particular, a single **propertyset** element MUST NOT include more than one **property** element that specifies the same *variableName* element; separate event notification messages MUST be used.

Request line

MUST be "NOTIFY *HTTP/1.1"

Header fields

HOST

REQUIRED. Field value MUST be 239.255.255.246:7900. Please note that port number 7900 is different from SSDP port number 1900.

CONTENT-LENGTH

REQUIRED. Field value specifies the length of the body in bytes. Integer. Chunked encoding MUST NOT be used for multicast event messages.

CONTENT-TYPE

REQUIRED. Field value MUST be "text/xml; charset="utf-8"".

USN

REQUIRED. Field value contains Unique Service Name for the publisher. Identifies a unique instance of a service in a unique instance of a device. It MUST be one of the following forms. The prefix (before the double colon) MUST match the value of the UDN element in the device description. (Section 2, "Description" explains the UDN element.) Single URI.

uuid: *device-UUID*::urn:schemas-upnp-org:*service*:*serviceType*:*ver*

where device-UUID is specified by the UPnP vendor; serviceType and ver are defined by the UPnP Forum working committee. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format.

uuid: *device-UUID*::urn:*domain-name*:*service*:*serviceType*:*ver*

where device-UUID, domain-name, serviceType and ver are defined by the UPnP vendor. See section 1.1.4, "UUID format and RECOMMENDED generation algorithms" for the MANDATORY UUID format. Period characters in the domain name MUST be replaced by hyphens in accordance with RFC 2141.

SVCID

REQUIRED. Field value contains ServiceID from the SCPD to uniquely identify which service generated the event. As defined in section 2.2, "Generic requirements on XML usage"

NT

REQUIRED. Field value contains Notification Type. MUST be **upnp:event**.

NTS

REQUIRED. Field value contains Notification Sub Type. MUST be **upnp:propchange**.

SEQ

REQUIRED. Field value contains Event Key. The numeric sequence count MUST be 0 for initial multicast event message. MUST be incremented by 1 for each multicast event message per a service; however, when a multicast message is retransmitted, it MUST be sent with its original event key. To prevent overflow, MUST be wrapped from 4294967295 to 1. 32-bit unsigned value represented as a single decimal integer without leading zeroes (leading zeroes, if present, MUST be ignored by the recipient).

LVL

REQUIRED. Field value MUST be a string in UTF-8. Event level allows the receiver to first level filter messages based on the value and is defined by the UPnP Technical Committee. See Table 4-8, "Multicast event levels" for the Event Levels defined with this version of the UPnP architecture. UPnP Working Committees MUST specify event level values when defining events that will be multicast.

The following table summarizes defined event levels and the expected meaning of those values. Event levels defined by the UPnP Forum Technical Committee start with the prefix "upnp:". Vendor and other extensions outside the UPnP Forum MUST be prefixed by the domain name of the defining organization. For example: "domain.org:/alerts/level/"

Table 4-8: Multicast event levels

Event Level	Description
upnp:/emergency	The event carries critical information that the device SHOULD act upon immediately.
upnp:/fault	The event carries information related to an error case
upnp:/warning	The event carries information that is a non-critical condition that the device MAY want to process or pass to the user
upnp:/info	The event carries information about the normal operation of the device that may be of interest to end-users. This information is simply informative and does not indicate any abnormal condition or status such as a warning or fault. Other event levels are defined for those purposes.
upnp:/debug	The event carries debug information typically used by programmers and test engineers to evaluate the internal operation of the device. This information is typically not displayed to end users.
upnp:/general	For events that fit into no other defined category
<code><domain>:/<level></code>	Example vendor extension. Domain is the ICANN domain name for the vendor and level is an arbitrary string defined by the vendor. E.g. domain.org:/alerts/type/

BOOTID.UPNP.ORG

REQUIRED. As defined in section 1.2, and 1.2.2.

Body

<propertyset>

REQUIRED. xmlns namespace MUST be "urn:schemas-upnp-org:event-1-0". Contains the following sub element:

<property>

REQUIRED. Repeat once for each variable name and value in the event message. MUST be qualified by the namespace prefix defined in the xmlns attribute of the <propertyset> element. Contains the following sub element:

<variableName>

REQUIRED. Element is name of a state variable that changed ([name](#) sub element of [stateVariable](#) element in service description). MUST NOT be qualified with any namespace. Value is the new value for this state variable. Case sensitive. Single data type as specified by UPnP service description.

Note that for simplicity many of the header fields for multicast eventing are the same as for unicast eventing. These include:

HOST, CONTENT-TYPE, USN, NT, NTS, and SEQ. In addition, the body of the message (propertyset) has the same format as unicast events.

For future extensibility and according to the requirements in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", when processing XML like the listing above, devices and control points MUST ignore: (a) any unknown elements and their sub elements or content, and (b) any unknown attributes and their values. Subject to the constraints defined in section 2.7, "Non-standard vendor extensions" and section 2.8, "UPnP Device Schema", control points and devices MUST ignore any XML comments or XML processing instructions embedded in UPnP device and service descriptions that they do not understand. The control point SHOULD discard and ignore unrecognized state variables within multicast event notification messages.

When the new value of any variable contains one or more characters reserved as markup (such as ampersand (“&”) or less than (“<”)), the text MUST be escaped in accordance with the provisions of section 2.4 of the XML specification and each such character replaced with the equivalent numeric representation or string (such as “&” or “<”). Such characters appearing in URLs that appear as values MAY also be percent-encoded in accordance with the URL percent-encoding rules specified in sections 2.1 and 2.4 of RFC 3986.

4.4 UPnP Event Schema

The UPnP Event Schema defines the structures and data types used in the body of UPnP event notifications. As explained with the UPnP Device and Service Schemas, the UPnP Event Schema is written in XML syntax according to the conventions of XML Schema (Part 1: Structures, Part 2: Datatypes). The UPnP Event Schema is defined within a UPnP service template; however, the schema MUST conform to the format as defined in appendix B.5, “UPnP Event Schema”. The elements it defines are used in event notifications.

As explained in section 2, “Description”, the UPnP Service Schema also specifies a `sendEvents` attribute for a state variable. The default value for this attribute is “[yes](#)”. To denote that a state variable is evented, the value of this attribute is “[yes](#)” (or the attribute is omitted) in a service description; to denote that a state variable is non-evented, the value is “[no](#)”. Note that if all of a service’s state variables are non-evented, the service has nothing to publish, and control points cannot subscribe and will not receive event messages from the service.

4.5 Augmenting the UPnP Device and Service Schemas

Some state variables may change value too rapidly for eventing to be useful. UPnP Forum Working Committees or UPnP vendors may document moderation in the number of event messages sent due to changes in a variable’s value. Event moderation may include limitation on the frequency in reporting change of a state variable value or a minimum degree of change that must occur before a change is reported.

Parameter	Description
maximumRate	Single numeric value (in seconds) of type integer or float. State variable <i>v</i> MUST NOT be part of an event message more often than every <i>n</i> seconds. If <i>v</i> is the only state variable changing, then an event message containing the state variable MUST NOT be generated more often than every <i>n</i> seconds. If <i>v</i> has changed sooner than <i>n</i> seconds from the last event message that contains <i>v</i> , then an event message containing the current value of <i>v</i> MUST be sent in a timely manner after <i>n</i> seconds from the previous event message containing <i>v</i> . If <i>v</i> has not changed within <i>n</i> seconds following the last event message that contains <i>v</i> , then when <i>v</i> does change an event message with the current value of <i>v</i> MUST be sent in a timely manner. Specifying a maximumRate value is useful for variables that model frequently changing state variables.
minimumDelta	Single numeric value (minimum change required) whose type MUST match the corresponding state variable. State variable <i>v</i> MUST NOT be part of an event message unless its value has changed (plus or minus) by at least minimumDelta since the last time an event message was sent that contains <i>v</i> . Only valid for state variables with a numeric (integer or float) data type. Specifying a minimumDelta value is useful for variables that model continuously changing state variables.

The publisher MAY send out any changed moderated variable when an event goes out. The publisher MUST meet moderation rules described above, but the publisher MAY flush recent changes when it sends out an event message.

Note that moderation affects events only and not state table updates. Specifically, control actions which return the value of state variables MAY return a more current value than published via eventing. Put another way, moderation means that not all state table changes result in events.

Decisions about which variables to event and any possible moderation is up to the appropriate UPnP Forum working committee (for standard services) or a UPnP vendor (for non-standard services).

4.6 References

RFC 2616

HTTP: Hypertext Transfer Protocol 1.1. Available at: <http://www.ietf.org/rfc/rfc2616.txt>.

RFC 3986

Uniform Resource Identifiers (URI): Generic Syntax. Available at: <http://www.ietf.org/rfc/rfc3986.txt>.

XML

Extensible Markup Language. Available at: <http://www.w3.org/TR/2000/REC-xml-20001006>.

XML Schema (Part 1: Structures, Part 2: Datatypes)

Available at: <http://www.w3.org/TR/xmlschema-1>, <http://www.w3.org/TR/xmlschema-2>.

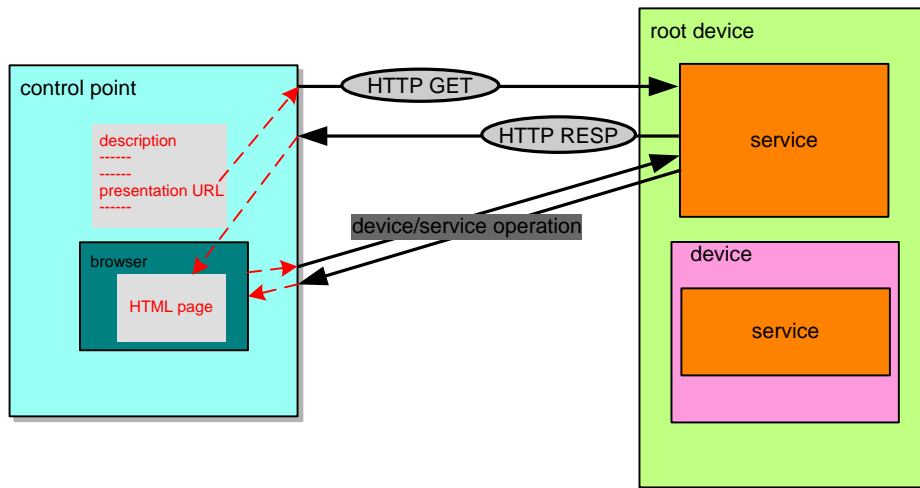
5 Presentation

[Normative]

Presentation is Step 5 in UPnP™ networking. Presentation comes after addressing (Step 0) where devices get network addresses, after discovery (Step 1) where control points find interesting device(s), and after description (Step 2) where control points learn about device capabilities. Presentation exposes an HTML-based user interface for controlling and/or viewing device status. Presentation is complementary to control (Step 3) where control points send actions to devices, and eventing (Step 4) where control points listen to state changes in device(s).

After a control point has (1) discovered a device and (2) retrieved a description of the device, the control point is ready to begin presentation. If a device has a URL for presentation, then the control point can retrieve a page 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 device status. The degree to which each of these can be accomplished depends on the specific capabilities of the presentation page and device.

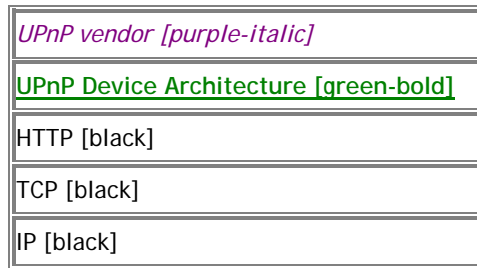
Figure 5-1: Presentation architecture



The URL for presentation is obtained from the presentationURL element in the device description. If presentationURL is an absolute URL, the fully qualified presentation URL is the presentationURL. Otherwise, if presentationURL is a relative URL, the fully qualified presentation URL is the URL resolved from presentationURL in accordance with section 5 of RFC 3986, using either the URLBase element, if specified, or the URL from which the device description was retrieved as the base URL. A multi-homed control point that attempts to retrieve a presentation page on a particular UPnP-enabled interface MUST use the fully qualified presentation URL from the description document received on that interface. The device description is delivered via a description message. Section 2, “Description” explains the device description and description messages in detail.

Retrieving a presentation page is a simple HTTP-based process and uses the following subset of the overall UPnP protocol stack. (The overall UPnP protocol stack is listed at the beginning of this document.)

Figure 5-2: Presentation protocol stack



At the highest layer, the presentation page is specified by a UPnP vendor. Moving down the stack, the UPnP Device Architecture specifies that this page be written in HTML. The page is delivered via HTTP over TCP over IP. For reference, colors in [square brackets] are included for consistency with other sections in this document.

To retrieve a presentation page, the control point issues an HTTP GET request to the presentation URL, and the device returns a presentation page. Responses to HTTP GET requests for presentation pages **MUST** be sent using the same address on the same interface on which the HTTP GET was received. The generic requirements on HTTP usage in UPnP 1.1 (as defined in section 2.1, “Generic requirements on HTTP usage” of this document) **MUST** be followed by devices and control points that implement presentation.

Unlike the UPnP Device and Service Templates, and standard device and service types, the capabilities of the presentation page are completely specified by the UPnP vendor. The page **MUST** be an HTML page; it is **RECOMMENDED** that the page be based upon XHTML-Basic. However, other design aspects are left to the vendor to specify. This includes, but is not limited to, all capabilities of the control point's browser, scripting language or browser plug-ins used, and means of interacting with the device. To implement a presentation page, a UPnP vendor **MAY** wish to use UPnP mechanisms for control and/or eventing, leveraging the device's existing capabilities but is not constrained to do so.

Presentation pages **SHOULD** use mechanisms provided by HTML for localization (e.g., META tag with charset attribute). Control points **SHOULD** use the ACCEPT-LANGUAGE and CONTENT-LANGUAGE feature of HTTP to try to retrieve a localized presentation page. Specifically, a control point **MAY** include a HTTP ACCEPT-LANGUAGE header field in the request for a presentation page; if an ACCEPT-LANGUAGE header field is present in the request, the response **MUST** include a CONTENT-LANGUAGE header field to identify the page's language.

It is **RECOMMENDED** that fully qualified URLs to resources on the device are not embedded in HTML presentation pages, but that relative URLs are used instead, so that the host portion of the embedded URLs does not need to be modified when sent on different UPnP-enabled interfaces.

5.1 References

RFC 3986

Uniform Resource Identifiers (URI): Generic Syntax. Available at: <http://www.ietf.org/rfc/rfc3986.txt>.

HTML

HyperText Markup Language. Available at: <http://www.w3.org/TR/html4>.

XHTML™ Basic
Available at: <http://www.w3.org/TR/xhtml-basic/>.

Appendix A. IP Version 6 Support

[Normative]

A.1 Introduction

Most of today's Internet uses IPv4, which is now nearly twenty years old. IPv4 has been remarkably resilient in spite of its age, but it is beginning to have problems. Most importantly, there is a growing shortage of IPv4 addresses, which are needed by all new machines added to the Internet. Deployment of large numbers of UPnP devices will only exacerbate the shortage.

IPv6 fixes a number of problems in IPv4, such as the limited number of available IPv4 addresses. It also adds many improvements to IPv4 in areas such as routing and network auto configuration. IPv6 is expected to gradually replace IPv4, with the two coexisting for a number of years during a transition period.

This Annex describes mechanisms by which devices and control points based on the UPnP Device Architecture MAY be used on IPv6 networks.

A.2 General Principles

Devices and control points MUST support IPv4-only, and MAY also support IPv6. IPv6-only devices and control points are NOT allowed in UPnP, since these cannot interoperate with IPv4-only control points and devices. All IPv6 devices and control points are therefore inherently multi-homed and must adhere to all multi-homed behaviors described in the rest of the document.

The following requirements apply to devices and control points using the UPnP Device Architecture over IPv6:

- Devices and control points MUST obtain a link-local address according to RFC 2462. Devices and control points SHOULD use the same address as previously used whenever possible. Appendix A.3.2, "Short overview of protocol specified by RFC 2462" presents a short overview of the theory of operation as specified by RFC 2462.
- Devices and control points MUST try to obtain (this may fail) a global address in each advertised prefix with the A bit set according to RFC 2462 (stateless autoconfiguration) and MAY try to obtain an IP address using DHCPv6, according to RFC 3315 (stateful autoconfiguration, DHCPv6). Devices and control points SHOULD use the same address as previously used whenever possible. A device MAY decide not to use its link-local addresses for UPnP if it supports UPnP on at least one global address on the same interface. This reduces overhead involved with announcing the device on all its addresses. A device MUST offer UPnP on one of its global addresses if it offers UPnP on a link-local address and it has successfully obtained a global address. This promotes interoperability and visibility of the UPnP device in larger networks.
- Devices and control points MUST listen for multicast messages on link local scope FF02::C, site local scope FF05::C and global scope FF0E::C (even if the device or control point does not have a global address, it MUST listen to all scopes). See also RFC 4291 for scope definitions.

- Devices **MUST** multicast SSDP messages for each of the UPnP-enabled interfaces. The scope of multicast SSDP messages **MUST** be link local FF02::C if the message is sent from a link local address. If the message is sent from a global address it **MUST** be multicast using either global scope FF0E::C or site local scope FF05::C. In networks with complex topologies and overlapping sites, use of global scope is **RECOMMENDED**.
- The hop limit of each IP packet for a global scope multicast message with a permanently assigned multicast address (e.g. FF0E::C for SSDP and FF0E::130 for multicast eventing) **SHOULD** be set to 254 and **SHOULD** be configurable.
- Devices and control points **MAY** select the interface or interfaces over which UPnP is enabled, when the device or control point supports multiple interfaces.
- Devices **MUST** listen for unicast SSDP traffic on all its UPnP-enabled addresses.

A.2.1 Device operation

A device supporting both IPv4 and IPv6 simultaneously **MUST** be advertised using the same USN on both IPv4 and IPv6 and **MUST** have identical device description documents and service description documents when accessed from both protocols. The device must also conform to other multi-homed descriptions in the respective sections of the document.

A.2.2 Control point operation

Control points can use the matching USNs of IPv4 and IPv6 announcements of dual stack IPv4/IPv6 devices to treat dual-stack devices as a single device. For example, a control point can subscribe to events on IPv6 and invoke actions on IPv4 based on the state information received on the IPv6 interface. In addition, the control point must also conform to other multi-homed descriptions in the respective sections of the document.

A.3 Addressing

RFC 2462 and RFC 3315 describe how a device or control point obtains an IPv6 address. Unlike when using IPv4, where each device or control point **MUST** have a DHCP client to try to obtain an address initially, DHCP is not always necessary in an IPv6 network. Addresses are automatically configured in new ways.

IPv6 multicast addresses are formed using a component of the address to determine the propagation of the message sent. The component is called scope and for UPnP 1.1, the scope **MUST** be one of: link local, site local or global. Further, since the multicast addresses used are permanently assigned, the scope is encompassing. That is, link local scope is contained in site local scope which is contained in global scope.

In IPv6 networks, a link-local address is determined per interface on the device or control point, and therefore IPv6 devices or control points will always have a link-local address. In addition, unicast addresses are also determined for each interface although a device or control point **MAY** or **MAY NOT** have a global address. Typical IPv6 devices or control points are multi-homed because they always have at least two addresses with which they can receive packets - a link-local address for local link traffic

and a routable global address. In some scenarios, devices or control points MAY only have a link-local address; reasons for this include device or control point sophistication (and thereby device or control point capability) and administrative policy. Link-local addresses are assigned immediately at the device or control point, without referring to an outside server such as a DHCP server. Global addresses are determined through the use of RA (Router Advertisement) messages in conversation with the local router.

This section describes the IPv6 autoconfiguration of unicast link-local addresses in more detail. In addition to the address assigned by this process, each device or control point, acting as a normal IPv6 host, listens for traffic on several multicast addresses: node-local scope all-nodes multicast address `FF01::1`; link-local scope all-nodes multicast address `FF02::1`; and multicast addresses of joined groups on each interface.

Note that in most implementations, the actions described are likely performed by the IPv6 stack and do not require any special coding by UPnP implementers.

A.3.1 Summary of boot/startup process

- For IPv4, Auto-IP addressing is performed as specified in section 0, "Addressing" of this document.
- Optionally, for IPv6, address assignment is performed as specified in RFC 2462 and RFC 3315. A short summary of the protocol is provided below.

Note that it is possible to obtain a global IPv4 address and only a link-local IPv6 address, or the other way around.

A.3.2 Short overview of protocol specified by RFC 2462

1. Hosts generate a link-local address for the interface. They test uniqueness using Neighbor Solicitation messages containing the tentative address.
2. Hosts obtain a Router Advertisement or determine that no routers are present
 - Hosts can send a Router Solicitation to obtain an advertisement quickly
3. Router Advertisements specify what sort of autoconfiguration a host SHOULD do, if no routers are present, stateful autoconfiguration MUST be invoked. Implementations MAY have an option to disable autoconfiguration, but the default SHOULD be that autoconfiguration is enabled.
 - A "ManagedFlag" flag indicates whether hosts SHOULD use stateful autoconfiguration to obtain addresses (false by default).

- An "other stateful configuration" flag (false by default) indicates whether hosts SHOULD use stateful autoconfiguration to obtain additional information (excluding addresses). By default, stateless autoconfiguration is used.
4. In the presence of Router Advertisements with ManagedFlag=false, Router Advertisements contain prefixes that identify the subnet(s) associated with a link. A (global) address is formed by combining a prefix with the identity of the interface (typically MAC address) on which IPv6 is used.
 5. In the presence of Router Advertisements with ManagedFlag=true, DHCPv6 MUST be used to obtain an address.
 6. Hosts will continually receive new advertisements, adding to and refreshing information received in previous advertisements.
 7. To speed the autoconfiguration process, a host MAY generate its link-local address (and verify its uniqueness) in parallel with waiting for a Router Advertisement.

A.4 Discovery

The UPnP discovery phase does not substantially change when used over IPv6. All definitions of section 1, "Discovery" of this document MUST be followed, except when a change is mentioned in this section.

IGMP is the protocol used by IPv4 to ensure that incoming multicast traffic is forwarded by a router to the network segment to which the router is attached. IGMP requires that the devices and control points attached to the network segment contact the router to notify it of their interest in certain multicast addresses. The protocol that provides this service in IPv6 is Multicast Listener Discovery protocol (MLD). Control points and devices MUST participate in the protocol for any IPv6 multicast scope that is not link local and is listened to by an interface.

IP Addresses embedded in UPnP messages and descriptions sent in response to requests received on IPv6 addresses will generally be literal addresses formatted according to RFC 2732 (including those in discovery messages, the URLBase element of the device description (if specified), and HTTP HOST header fields). Together with the UUID, the BOOTID.UPNP.ORG header field allows control points to recognize when a message received on a different protocol or address is referring to the same device that is multi-homed (in this case, the BOOTID.UPNP.ORG field value will be the same in all announcements), as opposed to being a new advertisement from a device which has changed from one protocol or address to another (in this case, the BOOTID.UPNP.ORG field value will differ between the old and the new announcement).

For backward compatibility with control points implementing UPnP over IPv6 according to the provisions of Annex A to UPnP Device Architecture version 1.0, devices SHOULD include an OPT header field and NLS header field in addition to the BOOTID.UPNP.ORG header field. The OPT header field is defined by the HTTP Extension Framework (RFC 2774); the OPT header field is used (rather than MAN) because it is possible for a control point to function without recognizing the NLS header field, although the user experience will be suboptimal (and IPv4-only control points may not recognize NLS). The NLS field value, as defined in Annex A to UPnP Device Architecture version 1.0, contains a string value which must change whenever the network

configuration of the device changes. It was recommended in that Annex that a GUID be used as the NLS field value, but other mechanisms for producing a unique value were permitted. Since under UPnP Device Architecture version 1.1 the BOOTID.UPNP.ORG field value is required to be unique on each device reboot or configuration change, the field value of the NLS header field can be set the same as the field value of the BOOTID.UPNP.ORG header field to simplify implementation.

A.4.1 Advertisement

For IPv6, a device advertises over IPv6 according to the following guidelines:

- SSDP announcements are sent to `[FF0X::C]:1900` (with "X" being set appropriately depending on the multicast scope upon which the announcement is being sent). Control points listen to these addresses and ports to detect when new devices are available on the network.
- As described in section 1.2.2 "Device available - NOTIFY with `ssdp:alive`", announcements sent over IPv6 MAY have a different CACHE-CONTROL field value and MAY be sent with a different frequency than announcements sent over IPv4. When all advertisements, both over IPv4 and IPv6, have expired, the control point MUST assume that the device (or service) is no longer available.
- The SSDP HOST field value contains an IPv6 address instead of an IPv4 address. The Internet Assigned Numbers Authority (IANA) has registered multicast address and port for SSDP: an address MUST be of the form `FF0X::C`. This is a variable scope multicast address where X is changed to represent the appropriate scope. For example, a device advertising on the local link would use a scope of 2 and address `FF02::C`. Port 1900 MUST be specified. For example, for a global scope broadcast the HOST header field is: `HOST: [FF0E::C]:1900`. In IPv6 announcements, the SSDP HOST header field will typically contain a literal IPv6 address, formatted according to RFC 2732, followed by a port number.
- The SSDP LOCATION field value contains the URL of the root device description document. Typically, a literal IPv6 address formatted according to RFC 2732 will be used. An IPv6 address MUST be contained within brackets if a port is specified. The host address in the URL MUST be valid within the current scope (the address or scope on which the announcement is being sent). Specifically, a device advertising over IPv6 MUST NOT use an IPv4 address in the SSDP LOCATION header field.
- The OPT and NLS header fields SHOULD be included.

The example below incorporates this syntax.

```
NOTIFY * HTTP/1.1
HOST: [FF02::C]:1900
CACHE-CONTROL: max-age = seconds until advertisement expires
LOCATION: URL for UPnP description of this device
OPT: "http://schemas.upnp.org/upnp/1/0/"; ns=01
01-NLS: same value as BOOTID field value
NT: notification type
NTS: ssdp:alive
```



```
SERVER: OS/version UPnP/1.1 product/version
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or update
message
CONFIGID.UPNP.ORG: number used for caching description information
USN: composite identifier for the advertisement
```

A.4.2 Advertisement: Device unavailable

All [ssdp:byebye](#) messages MUST be sent to the IPv6 multicast address as described in section A.4.1 “Advertisement”, and SHOULD contain the OPT and NLS header fields. Otherwise, the behavior is the same as IPv4. An example of an [ssdp:byebye](#) message has the following syntax.

```
NOTIFY * HTTP/1.1
HOST: \[FF02::C\]:1900
NT: notification type
OPT: "http://schemas.upnp.org/upnp/1/0/"; ns=01
01-NLS: same value as BOOTID field value
NTS: ssdp:byebye
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or update
message
CONFIGID.UPNP.ORG: number used for caching description information
USN: composite identifier for the advertisement
```

A.4.3 Advertisement: Device update

All [ssdp:update](#) messages MUST be sent to the IPv6 multicast address as described in section A.4.1 “Advertisement”, and SHOULD contain the OPT and NLS header fields. Otherwise, the behavior is the same as IPv4. An example of an [ssdp:update](#) message has the following syntax.

```
NOTIFY * HTTP/1.1
HOST: \[FF02::C\]:1900
LOCATION: URL for UPnP description for root device
NT: notification type
OPT: "http://schemas.upnp.org/upnp/1/0/"; ns=01
01-NLS: same value as BOOTID field value
NTS: ssdp:update
USN: composite identifier for the advertisement
BOOTID.UPNP.ORG: BOOTID value that the device has used in its previous announcements
CONFIGID.UPNP.ORG: number used for caching description information
NEXTBOOTID.UPNP.ORG: new BOOTID value that the device will use in subsequent announcements
SEARCHPORT.UPNP.ORG: number identifies port on which device responds to unicast M-SEARCH
```

A.4.4 Search

When a control point is added to the network, it MAY send multicast M-SEARCH requests on its IPv4 address(es), IPv6 address(es), or both. When searching over IPv6, a control point can freely choose the scope of the search (link local scope, site scope, administrative scope, organizational scope, global scope), allowing it to better direct its search. It should be noted that the source address of the M-SEARCH can influence how and where the M-SEARCH is routed. Aside from using an IPv6 multicast address and including an IPv6 address in the header fields, M-SEARCH messages are unchanged. An example of an M-SEARCH message has the following syntax. In addition, M-SEARCH messages MAY be unicast to IPv6 addresses of known devices, similar to IPv4 unicast M-SEARCH messages.

```
M-SEARCH * HTTP/1.1
HOST: [FF02::C]:1900
MAN: "ssdp:discover"
MX: seconds to delay response
ST: search target
```

A.4.5 Search response

To be found, a device MUST send a response to the source IP address and port that sent the request to the multicast address, and SHOULD include the OPT and NLS header fields in the message. An example of a search response message has the following syntax.

```
HTTP/1.1 200 OK
CACHE-CONTROL: max-age = seconds until advertisement expires
DATE: when response was generated
EXT:
LOCATION: URL for UPnP description of this device
SERVER: OS/version UPnP/1.1 product/version
OPT: "http://schemas.upnp.org/upnp/1/0/"; ns=01
01-NLS: same value as BOOTID field value
ST: search target
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or update message
CONFIGID.UPNP.ORG: number used for caching description information
USN: composite identifier for the advertisement
```

A.5 Description

Description documents MUST be sent using the same address on which the HTTP GET was received. Otherwise, behavior is the same as IPv4.

A.6 Control

Responses to SOAP messages during the Control phase MUST be sent on the same address on which the request was received. Otherwise, behavior is the same as IPv4.

A.7 Eventing

When subscribing to events over IPv6, the <deliveryURL> (or URLs) specified in the CALLBACK header field of the SUBSCRIBE message MUST be reachable by the device. This means, for example, when sending a SUBSCRIBE request to a device using a link-local IPv6 address, the <deliveryURL> MUST specify an IPv6 address on the same link, preferably a link-local address.

IPv4 addresses MUST NOT be included in the CALLBACK header field of a SUBSCRIBE message sent over IPv6. IPv6 addresses MUST NOT be included in the CALLBACK header field of a SUBSCRIBE message sent over IPv4.

IPv6 multicast event messages MUST be sent to [FF0X::130]:7900 (with "X" being equal to the address scope used in advertisement). To receive IPv6 multicast event messages, control points MUST listen to these addresses and ports.

To send a multicast event message, a publisher MUST send a message with method NOTIFY in the following format. Values in *italics* below are placeholders for actual values. Refer to section 4.3.3 Multicast Eventing: Event messages: NOTIFY for explanation of the elements. All IP addresses contained in the event MUST be IPv6 format and scoped as above.

```
NOTIFY * HTTP/1.1
HOST: [FF0X::130]:7900 *** note the address and the port number are different from SSDP ***
CONTENT-TYPE: text/xml; charset="utf-8"
USN: Unique Service Name for the publisher
SVCID: ServiceID from SCPD
NT: upnp:event
NTS: upnp:propchange
SEQ: monotonically increasing sequence count
LVL: event importance
BOOTID.UPNP.ORG: number increased each time device sends an initial announce or an update message

<?xml version="1.0"?>
<e:propertyset xmlns:e="urn:schemas-upnp-org:event-1-0">
  <e:property>
    <variableName>new value</variableName>
  </e:property>
  <!-- Other variable names and values (if any) go here. -->
</e:propertyset>
```

A.8 Presentation

Responses to HTTP GET requests for presentation pages MUST be sent using the same address on the same interface on which the HTTP GET was received.

Presentation pages retrieved over IPv6 MUST NOT contain IPv4 addresses. Presentation pages retrieved over IPv4 MUST NOT contain IPv6 addresses.

It is RECOMMENDED that fully qualified URLs to resources on the device are not embedded in HTML presentation pages, but that relative URLs are used instead, so that the host portion of the embedded URLs does not need to be modified to match the address on which the GET was received.

A.9 References

RFC 2732

Format for Literal IPv6 Addresses in URLs. Available at: <http://www.ietf.org/rfc/rfc2732.txt>.

RFC 2774

HTTP Extension Framework. Available at: <http://www.ietf.org/rfc/rfc2774.txt>.

RFC 2462

IPv6 Stateless Address Autoconfiguration. Available at: <http://www.ietf.org/rfc/rfc2462.txt>.

RFC 3315

Dynamic Host Configuration Protocol for IPv6 (DHCPv6). Available at: <http://www.ietf.org/rfc/rfc3315.txt>.

RFC 3986

Uniform Resource Identifiers (URI): Generic Syntax. Available at: <http://www.ietf.org/rfc/rfc3986.txt>.

RFC 4291

IP Version 6 Addressing Architecture. Available at: <http://www.ietf.org/rfc/rfc4291.txt>.

Appendix B. Schemas

[Informative]

B.1 UPnP Device Schema

Below is the UPnP Device Schema for devices (see also section 2.10, “UPnP Datatype Schema”). The elements it defines are used in UPnP Device Templates; they are colored [green](#) throughout this specification. Immediately following this is a brief explanation of the XML Schema elements, attributes, and values used. The reference to XML Schema at the end of the section has further details.

UPnP 1.0 specifies that the namespace of the device schema is “urn:schemas-upnp-org:device-1-0”. UPnP 1.1 does not change that namespace, but redefines it in a backwards-compatible way by restricting the order in which elements can be sent and REQUIRING the presence of the `configId` attribute. Therefore, the schema below specifies the syntax to which a UPnP 1.1 Device Description Document has to adhere. UPnP 1.1 control points also SHOULD expect Device Description Documents from UPnP 1.0 devices that can send elements in any order, and will not have the `configId` attribute.

```
<xsd:schema targetNamespace="urn:schemas-upnp-org:device-1-0"
xmlns="urn:schemas-upnp-org:device-1-0"
xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xsd:element name="root" type="rootType"/>
  <xsd:complexType name="deviceType">
    <xsd:sequence>
      <xsd:element name="deviceType">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:anyURI">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="friendlyName">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:string">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="manufacturer">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:string">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="manufacturerURL" minOccurs="0">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:anyURI">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
    </xsd:sequence>
  </xsd:complexType>
</xsd:schema>
```

```

    </xsd:complexType>
  </xsd:element>
  <xsd:element name="modelDescription" minOccurs="0">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:string">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="modelName">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:string">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="modelNumber" minOccurs="0">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:string">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="modelURL" minOccurs="0">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:anyURI">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="serialNumber" minOccurs="0">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:string">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="UDN">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:anyURI">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="UPC" minOccurs="0">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:string">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="iconList" minOccurs="0">
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="icon" maxOccurs="unbounded">
          <xsd:complexType>
            <xsd:sequence>

```

```

<xsd:element name="mimetype">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:string">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="width">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:int">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="height">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:int">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="depth">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:int">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="url">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:anyURI">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
  processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
  processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:element>
<xsd:element name="serviceList">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="service" maxOccurs="unbounded">
        <xsd:complexType>
          <xsd:sequence>
            <xsd:element name="serviceType">
              <xsd:complexType>
                <xsd:simpleContent>
                  <xsd:extension base="xsd:anyURI">
                    <xsd:anyAttribute namespace="##other" processContents="lax"/>
                  </xsd:extension>
                </xsd:simpleContent>
              </xsd:complexType>
            </xsd:element>
          </xsd:sequence>
        </xsd:complexType>
      </xsd:element>
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>

```

```

</xsd:element>
<xsd:element name="serviceId">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:anyURI">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="SCPDURL">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:anyURI">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="controlURL">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:anyURI">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="eventSubURL">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:anyURI">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
  processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
  processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:element>
<xsd:element name="deviceList" type="deviceListType" minOccurs="0"/>
<xsd:element name="presentationURL" minOccurs="0">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="xsd:anyURI">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
<xsd:complexType name="deviceListType">
  <xsd:sequence>
    <xsd:element name="device" type="deviceType" maxOccurs="unbounded"/>
    <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
  </xsd:sequence>
  <xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
<xsd:complexType name="rootType">

```

```

<xsd:sequence>
  <xsd:element name="specVersion">
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="major">
          <xsd:complexType>
            <xsd:simpleContent>
              <xsd:extension base="xsd:int">
                <xsd:anyAttribute namespace="##other" processContents="lax"/>
              </xsd:extension>
            </xsd:simpleContent>
          </xsd:complexType>
        </xsd:element>
        <xsd:element name="minor">
          <xsd:complexType>
            <xsd:simpleContent>
              <xsd:extension base="xsd:int">
                <xsd:anyAttribute namespace="##other" processContents="lax"/>
              </xsd:extension>
            </xsd:simpleContent>
          </xsd:complexType>
        </xsd:element>
        <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
          processContents="lax"/>
      </xsd:sequence>
      <xsd:anyAttribute namespace="##other" processContents="lax"/>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="URLBase" minOccurs="0">
    <xsd:complexType>
      <xsd:simpleContent>
        <xsd:extension base="xsd:anyURI">
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:extension>
      </xsd:simpleContent>
    </xsd:complexType>
  </xsd:element>
  <xsd:element name="device" type="deviceType"/>
  <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
</xsd:sequence>
<xsd:attribute name="configId" type="xsd:int"/>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:schema>

```

<element>

Defines a new element. `name` attribute defines element name. `type` attribute defines the data type for the content of element. `minOccurs` attribute defines minimum number of times the element MUST occur; default is `minOccurs="1"`; OPTIONAL elements have `minOccurs="0"`. `maxOccurs` attribute defines maximum number of times the element MUST occur; default is `maxOccurs="1"`; elements that can appear one or more times have `maxOccurs="unbounded"`. For a more detailed description, see the XML Schema specification.

<complexType>

Defines a new data type, often containing sub-elements. For a more detail description, see the XML Schema specification.

<attribute>

Defines a new attribute for the purpose of declaring in which elements it MAY appear. Like any XML element, the <attribute> element MAY have attributes of its own. The `use` attribute within this element indicates whether the attribute MUST be present; OPTIONAL attributes have `use="optional"`. For a more detail description, see the XML Schema specification.

B.2 UPnP Service Schema

Below is the UPnP Service Schema (see also section 2.9, “UPnP Service Schema”). The elements it defines are used in UPnP Service Templates; they are colored [green](#) throughout this specification. Immediately following this is a brief explanation of the XML Schema elements, attributes, and values used. The reference to XML Schema at the end of the section has further details.

UPnP 1.0 specifies that the namespace of the service schema is “urn:schemas-upnp-org:service-1-0”. UPnP 1.1 does not change that namespace, but redefines it in a backwards-compatible way by restricting the order in which elements can be sent and requiring the presence of the `configId` attribute. Therefore, the schema below specifies the syntax to which a UPnP 1.1 SCPD has to adhere. UPnP 1.1 control points also SHOULD expect SCPDs from UPnP 1.0 devices that can send elements in any order, and will not have the `configId` attribute.

```
<xsd:schema targetNamespace="urn:schemas-upnp-org:service-1-0"
  xmlns="urn:schemas-upnp-org:service-1-0"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xsd:element name="scpd" type="scpdType"/>
  <xsd:complexType name="directionType">
    <xsd:simpleContent>
      <xsd:extension base="xsd:string">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
  <xsd:complexType name="dataTypeType">
    <xsd:simpleContent>
      <xsd:extension base="xsd:string">
        <xsd:attribute name="type" type="xsd:string"/>
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
  <xsd:complexType name="scpdType">
    <xsd:sequence>
      <xsd:element name="specVersion">
        <xsd:complexType>
          <xsd:sequence>
            <xsd:element name="major">
              <xsd:complexType>
                <xsd:simpleContent>
                  <xsd:extension base="xsd:int">
                    <xsd:anyAttribute namespace="##other" processContents="lax"/>
                  </xsd:extension>
                </xsd:simpleContent>
              </xsd:complexType>
            </xsd:element>
            <xsd:element name="minor">
              <xsd:complexType>
                <xsd:simpleContent>
                  <xsd:extension base="xsd:int">
                    <xsd:anyAttribute namespace="##other" processContents="lax"/>
                  </xsd:extension>
                </xsd:simpleContent>
              </xsd:complexType>
            </xsd:element>
            <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
              processContents="lax"/>
          </xsd:sequence>
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="actionList" minOccurs="0">
```

```

<xsd:complexType>
  <xsd:sequence>
    <xsd:element name="action" maxOccurs="unbounded">
      <xsd:complexType>
        <xsd:sequence>
          <xsd:element name="name">
            <xsd:complexType>
              <xsd:simpleContent>
                <xsd:extension base="xsd:string">
                  <xsd:anyAttribute namespace="##other" processContents="lax"/>
                </xsd:extension>
              </xsd:simpleContent>
            </xsd:complexType>
          </xsd:element>
          <xsd:element name="argumentList" minOccurs="0">
            <xsd:complexType>
              <xsd:sequence>
                <xsd:element name="argument" maxOccurs="unbounded">
                  <xsd:complexType>
                    <xsd:sequence>
                      <xsd:element name="name">
                        <xsd:complexType>
                          <xsd:simpleContent>
                            <xsd:extension base="xsd:string">
                              <xsd:anyAttribute namespace="##other" processContents="lax"/>
                            </xsd:extension>
                          </xsd:simpleContent>
                        </xsd:complexType>
                      </xsd:element>
                      <xsd:element name="direction">
                        <xsd:complexType>
                          <xsd:simpleContent>
                            <xsd:restriction base="directionType">
                              <xsd:enumeration value="in"/>
                              <xsd:enumeration value="out"/>
                            </xsd:restriction>
                          </xsd:simpleContent>
                        </xsd:complexType>
                      </xsd:element>
                      <xsd:element name="retval" minOccurs="0">
                        <xsd:complexType>
                          <xsd:complexContent>
                            <xsd:restriction base="xsd:anyType">
                              <xsd:anyAttribute namespace="##other" processContents="lax"/>
                            </xsd:restriction>
                          </xsd:complexContent>
                        </xsd:complexType>
                      </xsd:element>
                      <xsd:element name="relatedStateVariable">
                        <xsd:complexType>
                          <xsd:simpleContent>
                            <xsd:extension base="xsd:string">
                              <xsd:anyAttribute namespace="##other" processContents="lax"/>
                            </xsd:extension>
                          </xsd:simpleContent>
                        </xsd:complexType>
                      </xsd:element>
                    <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
                      processContents="lax"/>
                    </xsd:sequence>
                    <xsd:anyAttribute namespace="##other" processContents="lax"/>
                  </xsd:complexType>
                </xsd:element>
                <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
                  processContents="lax"/>
                </xsd:sequence>
                <xsd:anyAttribute namespace="##other" processContents="lax"/>
              </xsd:complexType>
            </xsd:element>
          <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
            processContents="lax"/>
          </xsd:sequence>
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:complexType>
      </xsd:element>
    <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
      processContents="lax"/>
  </xsd:sequence>

```

```

        </xsd:sequence>
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
    </xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
    processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:element>
<xsd:element name="serviceStateTable">
    <xsd:complexType>
        <xsd:sequence>
            <xsd:element name="stateVariable" maxOccurs="unbounded">
                <xsd:complexType>
                    <xsd:sequence>
                        <xsd:element name="name">
                            <xsd:complexType>
                                <xsd:simpleContent>
                                    <xsd:extension base="xsd:string">
                                        <xsd:anyAttribute namespace="##other" processContents="lax"/>
                                    </xsd:extension>
                                </xsd:simpleContent>
                            </xsd:complexType>
                        </xsd:element>
                        <xsd:element name="dataType">
                            <xsd:complexType>
                                <xsd:simpleContent>
                                    <xsd:restriction base="dataTypeType">
                                        <xsd:enumeration value="ui1"/>
                                        <xsd:enumeration value="ui2"/>
                                        <xsd:enumeration value="ui4"/>
                                        <xsd:enumeration value="i1"/>
                                        <xsd:enumeration value="i2"/>
                                        <xsd:enumeration value="i4"/>
                                        <xsd:enumeration value="int"/>
                                        <xsd:enumeration value="r4"/>
                                        <xsd:enumeration value="r8"/>
                                        <xsd:enumeration value="number"/>
                                        <xsd:enumeration value="fixed.14.4"/>
                                        <xsd:enumeration value="float"/>
                                        <xsd:enumeration value="char"/>
                                        <xsd:enumeration value="string"/>
                                        <xsd:enumeration value="date"/>
                                        <xsd:enumeration value="dateTime"/>
                                        <xsd:enumeration value="dateTime.tz"/>
                                        <xsd:enumeration value="time"/>
                                        <xsd:enumeration value="time.tz"/>
                                        <xsd:enumeration value="boolean"/>
                                        <xsd:enumeration value="bin.base64"/>
                                        <xsd:enumeration value="bin.hex"/>
                                        <xsd:enumeration value="uri"/>
                                        <xsd:enumeration value="uuid"/>
                                    </xsd:restriction>
                                </xsd:simpleContent>
                            </xsd:complexType>
                        </xsd:element>
                        <xsd:element name="defaultValue" minOccurs="0">
                            <xsd:complexType>
                                <xsd:simpleContent>
                                    <xsd:extension base="xsd:string">
                                        <xsd:anyAttribute namespace="##other" processContents="lax"/>
                                    </xsd:extension>
                                </xsd:simpleContent>
                            </xsd:complexType>
                        </xsd:element>
                        <xsd:choice minOccurs="0">
                            <xsd:element name="allowedValueList">
                                <xsd:complexType>
                                    <xsd:sequence>
                                        <xsd:element name="allowedValue" maxOccurs="unbounded">

```

```

        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:string">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
        processContents="lax"/>
    </xsd:sequence>
    <xsd:anyAttribute namespace="##other" processContents="lax"/>
  </xsd:complexType>
</xsd:element>
<xsd:element name="allowedValueRange">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="minimum">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:double">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="maximum">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:double">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="step" minOccurs="0">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:double">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
        processContents="lax"/>
    </xsd:sequence>
    <xsd:anyAttribute namespace="##other" processContents="lax"/>
  </xsd:complexType>
</xsd:element>
</xsd:choice>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
  processContents="lax"/>
</xsd:sequence>
<xsd:attribute name="sendEvents" default="1">
  <xsd:simpleType>
    <xsd:restriction base="xsd:string">
      <xsd:enumeration value="1"/>
      <xsd:enumeration value="0"/>
    </xsd:restriction>
  </xsd:simpleType>
</xsd:attribute>
<xsd:attribute name="multicast" default="0">
  <xsd:simpleType>
    <xsd:restriction base="xsd:string">
      <xsd:enumeration value="1"/>
      <xsd:enumeration value="0"/>
    </xsd:restriction>
  </xsd:simpleType>
</xsd:attribute>
</xsd:complexType>

```

```

</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
  processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
</xsd:sequence>
<xsd:attribute name="configId" type="xsd:int"/>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:schema>

```

<element>

Defines a new element. `name` attribute defines element name. `type` attribute defines the data type for the content of element. `minOccurs` attribute defines minimum number of times the element MUST occur; default is `minOccurs="1"`; OPTIONAL elements have `minOccurs="0"`. `maxOccurs` attribute defines maximum number of times the element MUST occur; default is `maxOccurs="1"`; elements that can appear one or more times have `maxOccurs="unbounded"`. For a more detailed description, see the XML Schema specification.

<complexType>

Defines a new data type, often containing sub-elements. For a more detail description, see the XML Schema specification.

<attribute>

Defines a new attribute for the purpose of declaring in which elements it MAY appear. Like any XML element, the <attribute> element MAY have attributes of its own. The `use` attribute within this element indicates whether the attribute MUST be present; OPTIONAL attributes have `use="optional"`. For a more detail description, see the XML Schema specification.

B.3 UPnP Control Schema

Below is the template for UPnP Control Schemas (see also section 3.2.3, "UPnP Action Schema"). The elements it defines are used in actions and action responses; they are colored [green](#) throughout this specification. Immediately following this is a brief explanation of the XML Schema elements, attributes, and values used. The reference to XML Schema at the end of the section has further details.

```

<xsd:schema targetNamespace="urn:schemas-upnp-org:service:[serviceType:v]"
  xmlns="urn:schemas-upnp-org:service:[serviceType:v]"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xsd:element name="[actionName]" type="[actionName]Type"/>
  <xsd:element name="[actionName]Response" type="[actionName]ResponseType"/>
  <xsd:complexType name="[actionName]Type">
    <xsd:sequence>
      <!-- Use this for an argument of simple content. -->
      <xsd:element name="[argumentName]">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="[argumentType]">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <!-- Use this for an argument of complex content. -->
      <xsd:element name="[argumentName]" type="[argumentType]"/>
      <!-- Other arguments and their types go here, if any. -->
      <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
    </xsd:sequence>
    <xsd:anyAttribute namespace="##other" processContents="lax"/>
  </xsd:complexType>
  <xsd:complexType name="[actionName]ResponseType">
    <xsd:sequence>

```

```

<!-- Use this for an argument of simple content. -->
<xsd:element name="[argumentName]">
  <xsd:complexType>
    <xsd:simpleContent>
      <xsd:extension base="[argumentType]">
        <xsd:anyAttribute namespace="##other" processContents="lax"/>
      </xsd:extension>
    </xsd:simpleContent>
  </xsd:complexType>
</xsd:element>
<!-- Use this for an argument of complex content. -->
<xsd:element name="[argumentName]" type="[argumentType]"/>
<!-- Other arguments and their types go here, if any. -->
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:schema>

```

<element>

Defines a new element. `name` attribute defines element name. `type` attribute defines the data type for the content of element. `minOccurs` attribute defines minimum number of times the element MUST occur; default is `minOccurs="1"`; OPTIONAL elements have `minOccurs="0"`. `maxOccurs` attribute defines maximum number of times the element MUST occur; default is `maxOccurs="1"`; elements that can appear one or more times have `maxOccurs="unbounded"`. For a more detailed description, see the XML Schema specification.

<complexType>

Defines a new data type, often containing sub-elements. For a more detail description, see the XML Schema specification.

<attribute>

Defines a new attribute for the purpose of declaring in which elements it MAY appear. Like any XML element, the <attribute> element MAY have attributes of its own. The `use` attribute within this element indicates whether the attribute MUST be present; OPTIONAL attributes have `use="optional"`. For a more detail description, see the XML Schema specification.

B.4 UPnP Error Schema

Below is the template for UPnP Error Schemas (see also section 3.2.6, "UPnP Error Schema"). The elements it defines are used in error messages; they are colored [green](#) throughout this specification. Immediately following this is a brief explanation of the XML Schema elements, attributes, and values used. The reference to XML Schema at the end of the section has further details.

```

<xsd:schema targetNamespace="urn:schemas-upnp-org:control-1-0"
  xmlns="urn:schemas-upnp-org:control-1-0"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xsd:element name="UPnPError" type="UPnPErrorType"/>
  <xsd:complexType name="UPnPErrorType">
    <xsd:sequence>
      <xsd:element name="errorCode">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:int">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="errorDescription">
        <xsd:complexType>
          <xsd:simpleContent>
            <xsd:extension base="xsd:string">
              <xsd:anyAttribute namespace="##other" processContents="lax"/>
            </xsd:extension>
          </xsd:simpleContent>
        </xsd:complexType>
      </xsd:element>
    </xsd:sequence>
  </xsd:complexType>

```

```

</xsd:element>
<xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
</xsd:sequence>
<xsd:anyAttribute namespace="##other" processContents="lax"/>
</xsd:complexType>
</xsd:schema>

```

<element>

Defines a new element. `name` attribute defines element name. `type` attribute defines the data type for the content of element. `minOccurs` attribute defines minimum number of times the element MUST occur; default is `minOccurs="1"`; OPTIONAL elements have `minOccurs="0"`. `maxOccurs` attribute defines maximum number of times the element MUST occur; default is `maxOccurs="1"`; elements that can appear one or more times have `maxOccurs="unbounded"`. For a more detailed description, see the XML Schema specification.

<complexType>

Defines a new data type, often containing sub-elements. For a more detail description, see the XML Schema specification.

<attribute>

Defines a new attribute for the purpose of declaring in which elements it MAY appear. Like any XML element, the <attribute> element MAY have attributes of its own. The `use` attribute within this element indicates whether the attribute MUST be present; OPTIONAL attributes have `use="optional"`. For a more detail description, see the XML Schema specification.

B.5 UPnP Event Schema

Below is the template for the UPnP Event Schema (see also section 4.4, "UPnP Event Schema"). The elements it defines are used in event notifications; they are colored [green](#) throughout this specification. Immediately following this is a brief explanation of the XML Schema elements, attributes, and values used. The reference to XML Schema at the end of the section has further details.

```

<xsd:schema targetNamespace="urn:schemas-upnp-org:event-1-0"
  xmlns="urn:schemas-upnp-org:event-1-0"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xsd:element name="propertyset" type="propertysetType"/>
  <xsd:complexType name="propertysetType">
    <xsd:sequence>
      <xsd:element name="property" maxOccurs="unbounded">
        <xsd:complexType>
          <xsd:sequence>
            <!-- Use this for a stateVariable of simple content. -->
            <xsd:element name="[stateVariableName]" form="unqualified">
              <xsd:complexType>
                <xsd:simpleContent>
                  <xsd:extension base="[stateVariableType]">
                    <xsd:anyAttribute namespace="##other" processContents="lax"/>
                  </xsd:extension>
                </xsd:simpleContent>
              </xsd:complexType>
            </xsd:element>
            <!-- Use this for a stateVariable of complex content. -->
            <xsd:element name="[stateVariableName]" type="[stateVariableType]"
              form="unqualified"/>
            <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded"
              processContents="lax"/>
          </xsd:sequence>
          <xsd:anyAttribute namespace="##other" processContents="lax"/>
        </xsd:complexType>
      </xsd:element>
      <xsd:any namespace="##other" minOccurs="0" maxOccurs="unbounded" processContents="lax"/>
    </xsd:sequence>
    <xsd:anyAttribute namespace="##other" processContents="lax"/>
  </xsd:complexType>
</xsd:schema>

```

<element>

Defines a new element. `name` attribute defines element name. `type` attribute defines the data type for the content of element. `minOccurs` attribute defines minimum number of times the element MUST occur; default is `minOccurs="1"`; OPTIONAL elements have `minOccurs="0"`. `maxOccurs` attribute defines maximum number of times the element MUST occur; default is `maxOccurs="1"`; elements that can appear one or more times have `maxOccurs="unbounded"`. For a more detailed description, see the XML Schema specification.

<complexType>

Defines a new data type, often containing sub-elements. For a more detail description, see the XML Schema specification.

<attribute>

Defines a new attribute for the purpose of declaring in which elements it MAY appear. Like any XML element, the <attribute> element MAY have attributes of its own. The `use` attribute within this element indicates whether the attribute MUST be present; OPTIONAL attributes have `use="optional"`. For a more detail description, see the XML Schema specification.

B.6 Schema references

XML

Extensible Markup Language. Available at: <http://www.w3.org/XML>.

XML Schema (Part 1: Structures, Part 2: Datatypes)

Available at: <http://www.w3.org/TR/xmlschema-1>, <http://www.w3.org/TR/xmlschema-2>.

Namespaces in XML

Available at: <http://www.w3.org/TR/REC-xml-names/>.
