Universidade de Coimbra
Departamento de Engenharia Informática
Faculdade de Ciências e Tecnologia

A Management Framework for Residential Broadband Environments

Tiago José dos Santos Martins da Cruz

Coimbra
2011

UNIVERSIDADE DE COIMBRA
DEPARTAMENTO DE ENGENHARIA INFORMÁTICA

FACULDADE DE CIÊNCIAS E TECNOLOGIA

A MANAGEMENT FRAMEWORK FOR RESIDENTIAL BROADBAND ENVIRONMENTS

Tiago José dos Santos Martins da Cruz

Dissertação

COIMBRA
2011

APOIO:
FCT Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA CIÊNCIA, TECNOLOGIA E ENSINO SUPERIOR
BD REF. SFRH/BD/29118/2006

Tese realizada sob a orientação do

Prof. Doutor Paulo Alexandre Ferreira Simões

Professor Auxiliar do Departamento de Engenharia Informática da

Faculdade de Ciências e Tecnologia da Universidade de Coimbra

Com o apoio da FCT – Fundação para a Ciência e Tecnologia, ao abrigo da Bolsa de Doutoramento SFRH/BD/29118/2006.

Palavras-Chave

Gestão de serviços; Gestão de dispositivos; Redes de acesso de banda larga; *Gateways* domésticas; CWMP; Broadband Forum; UPnP; DLNA; Redes domésticas; Armazenamento distribuído; Segurança

Keywords

Service management; Device management; Broadband access networks; Residential gateways; CWMP; Broadband Forum; UPnP; DLNA; Residential networks; Distributed Storage; Security

Sumário

O crescimento do número de clientes servidos por redes de acesso de banda larga acarreta um novo conjunto de preocupações ao nível da gestão, serviços e segurança, com potenciais consequências para os operadores de telecomunicações, seus clientes e terceiros. O número elevado de clientes domésticos e tecnicamente impreparados que são atualmente servidos por conexões de elevado débito e natureza permanente constitui um cenário de risco para o qual o modelo tradicional de gestão e segurança dos operadores, centrado na sua infraestrutura interna, é incapaz de dar resposta.

Esta tese propõe abordar estas questões na perspetiva dos mecanismos de gestão do operador, propondo uma plataforma de gestão de equipamentos e serviços adaptada a um novo paradigma de operação que pressupõe o alargamento da esfera de intervenção à rede local do cliente, no sentido de colmatar a cada vez mais frequente incapacidade deste para gerir adequadamente os equipamentos e recursos ao seu dispor.

Neste prisma, vários aspetos relacionados com a temática dos ambientes de banda larga serão abordados de um modo integrado e ortogonal à plataforma de gestão proposta, através de um conjunto de cenários de aplicação, nomeadamente:

- **Gestão de dispositivos:** procurando dar resposta à necessidade crescente da integração do ecossistema heterogéneo de equipamentos e respetivos protocolos de gestão existentes na esfera da rede local dos clientes, abrindo as plataformas de gestão de operador à interação com estes. Deste modo, os operadores podem intervir de um modo ativo na rede dos clientes, tirando dos seus ombros a responsabilidade da configuração, diagnóstico e gestão dos equipamentos aí existentes, com benefícios evidentes em termos de operação e qualidade de serviço.
- **Exploração de novos paradigmas de serviços:** a exploração de novos paradigmas de serviços sobre redes de acesso de banda larga é outra das vertentes que esta tese propõe abordar, no sentido de investigar e validar um conjunto de serviços concebidos para tirar partido das características particulares destes ambientes. Estes serviços são propostos como complementos que adicionam mais-valia à oferta de conectividade e serviços básicos por parte dos operadores, ao mesmo tempo que oferecem um modelo de operação com gestão supervisionada.
- **Novos modelos de segurança:** a natureza específica das redes de acesso de banda larga, combinada com a sua proliferação, veio criar e/ou acentuar um conjunto de problemas de segurança que cada vez constituem uma séria ameaça a vários níveis, com várias consequências que podem ir da degradação de serviços ao comprometimento de informação pessoal. Nesta ótica é proposto um modelo distribuído de segurança com gestão partilhada que procura aliar operadores e clientes na deteção e combate às potenciais ameaças que afetam os utilizadores dos serviços de banda larga.

Estas três propostas procuram ir de encontro a um conjunto de necessidades e lacunas que se têm feito sentir ao nível das redes de acesso de banda larga, e que se revestem de particular importância quando contextualizadas no âmbito da introdução de serviços integrados baseados em IP (como é o caso do *triple-play*) e do crescimento considerável do número de clientes. Deste modo, o trabalho desenvolvido no âmbito desta tese procura ir ao encontro destas necessidades, contribuindo para reequilibrar as assimetrias identificadas através da proposta de modelos inovadores de segurança, serviços e gestão.

Abstract

The expansion of high-speed broadband access networks, with an increasing growth in the number of connected households has brought a new set of concerns related to aspects such as management, services and security, with potential consequences for communication operators, clients and third-parties. The considerable number of residential customers served by broadband networks that lack the necessary technical knowledge to manage their equipment and infrastructure, in a self-sufficient manner, together with the high bandwidth available for each permanent connection, contribute to a scenario that conventional centralized operator security and management models are unable to deal with.

This thesis addresses these issues in the perspective of the operator management infrastructure, by proposing a management framework for devices and services based on a different operation paradigm in which the operator is able to extend its influence to the customer premises LAN, instead of remaining confined to its own infrastructure. This has the benefit of relieving the users from the LAN configuration and management burden, while allowing operators to deliver a better service, by easing diagnostics and configuration procedures.

In this perspective, several related aspects will be addressed in the form of application scenarios, always in an integrated perspective orthogonal to the proposed management framework, namely:

- **Device management:** in order to integrate the heterogeneous device and management standards ecosystem of the residential network in the scope of the operator management infrastructure. By bridging both worlds, operators are able to extend their reach into the customers' premises networks, managing all sorts of devices and services while relieving users from such burden and improving service quality.
- **Exploration of new service paradigms:** another aspect which is addressed in the scope of this thesis has to do with researching and evaluating new service paradigms for leveraging the benefits of broadband environments. Those value-added proposals are conceived as complementary to the existing operators' connectivity and service portfolio, being proposed in the form of managed services.
- **New security models:** the specific nature of broadband network environments, together with its increasing household penetration ratio has contributed to create and/or increase a number of security issues which are growing to the point of becoming a serious threat, with repercussions at several levels, from service degradation to compromising personal information. In this perspective, a distributed security model based on the concept of shared security is proposed, bringing together operators and users in an effort to detect and fight the potential menaces which threaten modern broadband environments.

Not only these topics are of particular concern in the scope of broadband access networks, but they are also becoming increasingly relevant with the inclusion of other factors such as the introduction of integrated broadband services over IP (such as triple-play) and the expansion of the customer base. As such, his thesis proposes to contribute to this discussion by proposing innovative models for security, services and management in the context of broadband access networks.

Acknowledgments

I would like to express my deepest gratitude to my scientific advisor, Dr. Paulo Simões, for his guidance, understanding, encouragement and expertise. Every time I needed advice or guidance, he was there to help me.

To Professor Edmundo Monteiro, my *de facto* scientific co-advisor, for his advice and pertinent remarks.

To the people at PT Inovação, especially to Eng. Fernando Bastos and Eng. Alexandre Laranjeira for their valuable contributions and advice.

To Dr. José Marques, for his friendship and support, true to the motto: *a friend in need is a friend indeed.*

To my lab colleagues and co-researchers – the S3P team in all its incarnations: João Almeida, João Rodrigues, Patrício Batista, Rui Vilão, Thiago Leite, for their team spirit, enthusiasm and dedication.

To the *mighty G6.6 crew*: Vitor Bernardo, David Palma, João Gonçalves, Bruno Sousa and remaining colleagues, for the sympathetic ears, advice and true friendship.

To my friends, especially to Alexadre Leão, João Santos, Isidro Caramelo, Luís Cordeiro, Eduardo Costa, João Martins, José Manuel, José Martins, José Pião, Jorge Tavares, Nuno Oliveira, Ricardo Riquito e Sérgio Peres, for their support and help in the hour of need.

To my wife and best friend, Olga, without whose love and encouragement I would not have finished this thesis. You gave me the confidence I needed to sail across the oceans of uncertainty.

I would also like to thank my family for the support they provided me.

And least, but not last, to God, into Whom I deposited my faith and Who never left me alone.

Table of Contents

List of Figures

List of Tables

List of Acronyms

AAL5	ATM Adaptation Layer 5
ACS	Auto-Configuration Server
ADSL	Asynchronous Digital Subscriber Line
ALG	Application Level Gateway
API	Application Programming Interface
ATA	Analog Telephony Adaptor
ATM	Asynchronous Transfer Mode
AV	Audio Video
BIS	Boot Integrity Services
BRAS	Broadband Remote Access Server
BSS	Business Support Systems
CDC	Connected Device Configuration
CE	Consumer Electronics
CENELEC	Comité Européen de Normalisation Électrotechnique
CERT	Computer Emergency Response Team
CIFS/SMB	Common Internet File System/Server Message Block
CIM	Common Information Model
CLDC	Connected Limited Device Configuration
CLI	Command Line Interface
CLR	Common Language Runtime
CMIP	Common Management Information Protocol
CMOT	CMIP Operations over TCP/IP
COTS	Commercial of-the-Shelf
CPCS	Common Part Convergence Sub-layer
CPE	Customer Premises Equipment
CWMP	CPE Wan Management Protocol
DaaS	Desktop as a Service
DCOM	Distributed Component Object Model
DCP	Device Control Protocol
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DHF	Digital Home-network Forum
DIDL	Digital Item Declaration Language
DLNA	Digital Living Network Alliance
DMTF	Distributed Management Task Force
DNS	Domain Name System
DNS-SD	DNS Service Discovery
DOCSIS	Data Over Cable Service Interface Specification
DPWS	Devices Profile for Web Services
DSLAM	Digital Subscriber Line Access Multiplexer
DTCP-IP	Digital Transmission Content Protection over IP
DU	Deployment Units
DVB	Digital Video Broadcasting
DVD	Digital Versatile Disk
EE	Execution Environments
EU	Execution Units
FCC	Federal Communications Comission
FCS	Frame Check Sequence
FIGARO	Future Internet Gateway-based Architecture of Residential Networks

FM	Frequency Modulation
FTP	File Transfer Protocol
FTTB	Fiber to the Building
FTTH	Fiber to the Home
FTTP	Fiber to the Premises
GEM	Generic Encapsulation Method
GENA	General Event Notification Architecture
GPON	Gigabit Passive Optical Network
GSM	Global System for Mobile Communications
HES	Home Electronic System
HGI	Home Gateway Initiative
HIDS	Host-based Intrusion Detection System
HTTP	Hypertext Transfer Protocol
HTTPS	Secure HTTP
HVAC	Heating Ventilation and Air Conditioning
iAMT	Intel Active Management Techology
ICA	Independent Computing Architecture
IDMEF	Intrusion Detection Message Exchange Format
IDS	Intrusion Detection System
IEC	International Electrotechnical Comission
IEEE	Institute of Electric and Electronics Engineers
IETF	Internet Engineering Task Force
IGD	Internet Gateway Device
IGRS	Intelligent Grouping and Resource Sharing
IPS	Intrusion Prevention System
IPTV	IP Television
IPv4	Internet Protocol, version 4
IPv4LL	IPv4 Link-local Addresses
IPv6	Internet Protocol, version 6
ISC	Internet Systems Consortium
ISDN	Integrated Service Digital Networks
ISO	International Standards Organization
ISP	Internet Service Provider
ITSP	Internet Telephony Service Providers
ITU	International Telecommunications Union
IVR	Interactive Voice Response
JAR	Java Archive
JERI	Jini Extensible Remote Invocation
Jini	Java Intelligent Network Infrastructure
JMX	Java Management Extensions
JPEG	Joint Pictures Expert Group
JSR	Java Specification Request
JVM	Java Virtual Machine
L2	Layer 2
L3	Layer 3
LAN	Local Area Network
LDAP	Lightweight Directory Application Protocol
LLC	Logical Link Control
LLDP	Link Layer Discovery Protocol
LLMNR	Link-local Multicast Name Resolution
LLTD	Link Layer Topology Discovery
MAC	Media Access Control
mDNS	Multicast Domain Name System

MHP	Multimedia Home Platform
MIB	Managed Information Base
MOWS	Management Of Web Services
MPEG	Moving Pictures Expert Group
MPLS	Multiprotocol Layer Switching
MTP	Media Transfer Protocol
MTU	Maximum Transfer Unit
MUSE	Multi-Service Access Anywhere
MUWS	Management Using Web Services
NAPT	Network Address and Port Translation
NAS	Network Attached Storage
NAT	Network Address Translation
NBP	Network Boot Program
NGN	Next Generation Network
NIDS	Network Intrusion Detection System
OA&M	Operation, Administration, Maintenance and Provisioning
OASIS	Organization for the Advancement of Structured Information Standards
OCAP	OpenCable Application Platform
OGMP	Open Gateway Management Protocol
OLT	Optical Line Termination
OMA	Open Mobile Alliance
OMA-DM	Open Mobile Alliance-Device Management
OS	Operating System
OSGi	Open Services Gateway Initiative
OSS	Operation Support Systems
OUI	Organizational Unique Identifier
P2P	Peer-to-Peer
PBC	Push-button Configuration
PBX	Private Branch Exchange
PC	Personal Computer
PCM	Pulse Code Modulation
PCoIP	PC over IP
PDA	Personal Digital Assistant
PDU	Protocol Data Unit
PD-XXX	Proposed Draft XXX
PIN	Personal Identification Number
PLC	PowerLine Communications
PnP-X	Plug and Play Extensions
POTS	Plain Old Telephone System
PPP	Point-to-Point Protocol
PPPoE	PPP over Ethernet
PSTN	Public Switched Telephone Network
PXE	Preboot Execution Environment
QoS	Quality of Service
RAID	Redundant Array of Inexpensive Disks
RDP	Remote Desktop Protocol
RFCXXXX	Request for Comments XXXX
RGW	Residential Gateway
RPC	Remote Procedure Call
RTP	Real-time Transport Protocol
SaaS	Software as a Service
SIP	Session Initiation Protocol
SLA	Service Level Agreement

SLP	Service Location Protocol
SNMP	Simple Network Management Protocol
SOA	Service Oriented Architectures
SOAP	Simple Object Access Protocol
SOHO	Small Office Home Office
SSDP	Simple Service Discovery Protocol
SSH	Secure Shell
SSL	Secure Sockets Layer
STB	Set-top Box
STUN	Session Traversal Utilities for NAT
TAHI	The Application Home Initiative
TCO	Total Cost of Ownership
TCP	Transmission Control Protocol
TEAHA	The European Application home Alliance
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
TR-XXX	Technical Report XXX
UDN	Unique Device Name
UDP	User Datagram Protocol
UNDI	Universal Network Device Interface
UPnP	Universal Plug and Play
URL	Universal Resource Locator
USB	Universal Serial Bus
USN	Unique Service Name
UUID	Universal Unique Identifier
VCI	Vendor Class Identifier
VLAN	Virtual LAN
VM	Virtual Machine
VNC	Virtual Network Computing
VoD	Video on Demand
VoIP	Voice Over IP
VPN	Virtual Private Network
W3C	World Wide Web Consortium
WAN	Wide Area Network
WBEM	Web Based Enterprise Management
WDSL	Web Service Description Language
WT-XXX	Working Text XXX
WinRM	Windows Remote Management
WMI	Windows Management Instrumentation
WMM	Wi-Fi Multimedia
WPA2	Wi-Fi Protected Access 2
WQL	WMI Query Language
WSD	Web Services on Devices
WSDM	Web Services Distributed Management
WWW	World Wide Web
XBMC	Xbox Media Center
X-CWMP	Extended CWMP
xDSL	Digital Subscriber Line
XML	Extensible Markup Language

Part I:

Management on broadband access networks

1. Introduction

In this chapter we will introduce the theme of this thesis, the motivation behind it and the topics which constitute the basis for the associated research work. In the same way, it will be identified and discussed the approaches which were adopted to pursue the established objectives, also enumerated in the form of expected scientific contributions and impact. This chapter will end with a description of the structure of this thesis.

1.1 Context

Contemporary broadband access network environments are the natural outcome of a process that started several years ago with the first POTS (Plain Old Telephone System). That was the first access network, which played an important role by connecting communication carriers and service providers with the individuals and companies they served – by then the original access network was formed by the ensemble of wires, cables and equipment[1] deployed between a subscriber telephone termination point and the local telephone exchange.

Meanwhile, the access network has constantly developed, expanding its reach and evolving to support new services – dial-up, which was the earlier and dominant form of public Internet access in the 80s and 90s was gradually displaced by broadband networks since the beginning of the 21st century. Nowadays, it is common to find that the same lines which formed part of the original telephone and (later) cable television access network carry high-speed broadband services such as xDSL (Digital Subscriber Line) or DOCSIS (Data Over Cable Service Interface Specification), in addition to telephone or cable television signals.

Having reached a perceived critical status similar to that of water or electricity supply networks, broadband access networks are gradually being included in the USO (Universal Service Obligation) regulations of several countries, through the definition of a minimum capacity in terms of broadband internet access as a universal right, accessible for every and each citizen. What once was regarded as a luxury is now becoming a basic right – in fact, broadband internet access coverage is a development indicator used to measure not only the technological level of the communication infrastructure of a specific country, but also its competitiveness.

Contemporary broadband access networks are a result of an evolution process, mainly driven by an unbalanced mix of both technological advances and customer demands, with one outreaching the other in frequent occasions throughout their historical timeline.

If the birth of the World Wide Web in 1990 [Berners-Lee1990] paved the way for a more accessible and user-friendly Internet experience, what came next was a rapid succession of technologies and standards that went to make of it a diversified experience of content and form. Static web pages rapidly gave place to a new generation of web sites of dynamic nature with personalized user experiences and full of multimedia contents. With the era of the Web 2.0, P2P (Peer to Peer), Social Networks, Messaging, VoIP (Voice over IP), IPTV (IP Television) and VoD (Video on Demand), the rich internet has arrived.

Such diversity of content and offer had a “price”: bandwidth. Dial-up connections were unable to cope with such pace of evolution and, as a consequence, the rich internet pushed its way through the access networks. Users demanded more bandwidth, lower bills and better quality connections to take advantage of everything Internet had to offer and, as a result, access networks had to evolve in order to answer users’ demands. This situation, together with the introduction of flat rates, was the reason why finally broadband service adoption took off after a slow start back when the first services were launched to the public in the year 2000 [Anderson2002].

When fixed network operators planned broadband access services, they were concerned about ways of expanding their subscriber base and increasing revenue – apart from high-speed internet access, the rapid evolution of broadband technologies lacked some substance without the equal development of other complementary services. With the increased availability of permanent

[1] Typically little more than the telephone fixed-line loop – a twisted copper or aluminium pair, used to support voice communications, which was terminated in a circuit switch at the edge of the carrier service provider network.

broadband connections to a wider number of households, new service paradigms emerged, as it was the case of converged *Triple Play* (aggregation of voice, television and data services over IP using broadband as the common communication medium) – not only increasing operators' revenue and customers' loyalty but also reducing logistical complexity.

As a result, broadband networks have reshaped computing usage paradigms as a whole, being somewhat analogous to Moore's Law [Moore1965] in terms of impact, even overshadowing it. If all that mattered in the past was processing capacity and storage, nowadays the consumer wants bandwidth and convenience, sometimes up to the point of willing to trade up CPU power and storage capacity for them, for a reasonable price. As a result, a new generation of thin computing platforms, like smartphones and netbooks (both frequently bundled in mobile broadband contracts) were born to fulfill the need for small, lightweight and cheap networked computing devices with modest processing capabilities and good battery life. The computer is no longer the sole broadband consumer device.

Widespread broadband residential access, together with the introduction of novel services to domestic users – such as VoIP, IPTV, VoD and video surveillance – is also gradually transforming the shape of the domestic LANs (which are now commonplace among broadband customers), spawning a growing ecosystem of heterogeneous IP-enabled devices and network technologies (Wi-Fi, Powerline, Ethernet). Those devices share the same broadband connection with the help of a residential gateway (RGW)[2] which provides the frontier between two worlds: the access network and the customer LAN.

Until recently, Internet Service Providers (ISPs) defined their management scope to end at their borderline equipment, with the domestic customer being responsible for his own frontier equipment (such as home/residential gateways) and everything beyond that. This state of affairs was convenient for both customers and ISPs, due to ethical, practical and legal reasons.

But the situation is gradually changing, since the need for ISPs to be able to manage equipment inside the customer LAN is increasing as customers expect new IP-based services to have equal or better performance and reliability, when compared to their conventional counterparts. Since most customers lack the willingness or technical skills to properly manage those devices, this implies that operators must be able to remotely manage the devices (configuration management, monitoring, etc.) and the path between them and the access network (i.e., at least a segment of the customer LAN), in order to maintain adequate service levels.

In fact, some of the key services now provided by operators (VoIP, IPTV, VoD, femtocell-based applications, etc.) heavily depend on equipment placed on the customer LAN but intended to be managed by the ISP (e.g. set-top boxes). Despite legitimate customer privacy concerns, the truth is that this trend has already implicitly started – most triple play customers already have ISP-provided devices on their LAN with customized configurations and/or firmware which they cannot control – and seems unavoidable in general terms (of course, there will always be a minority of users with the time and the skills to circumvent this scenario).

As a result, the idea of allowing operators to manage equipment inside the customer's premises gains increased acceptance by both sides, creating the need for adequate management mechanisms capable of addressing the necessities of service providers while safeguarding customer's autonomy and privacy concerns. This is the main scope of the research work hereby presented, which explores the possibility of developing a management framework designed

[2] The terms *Residential Gateway* (RGW), *Home Gateway* (HGW), *Broadband Gateway* (BGW) or *Domestic Gateway* will be used interchangeably along this document, referring to the same meaning.

from the ground up to address the device and service management needs of broadband access network environments, while also exploring new managed service paradigms.

1.2 Research topics and objectives

The following section details the specific research topics which will be dealt with in this document and the expected results. Apart from the research topics, other aspects such as research methodologies, contributions and document structure will be also discussed.

1.2.1 Novel service paradigms and their impact

Widespread broadband residential access, together with the emergence of new digital convergence paradigms are reshaping traditional communication services usage and operation, allowing new levels of service diversity, quality and value to the subscriber. From the operators' standpoint, all these changes are twofold. On one hand, the transition from dial-up point-to-point, transient and low-bandwidth connections to high bandwidth, permanent links, the emergence of new services and devices in SOHO networks and the shift in the predominant network traffic models (with an ever growing significance of P2P-related traffic) increased the pressure on operators to constantly upgrade and improve their core and access networks in order to keep up. But, on the other hand, new opportunities were created in the process with even some of the most traditional operators becoming involved in the media content distribution, television and other content businesses.

Modern broadband-based access networks are redefining the architecture and operation of traditional communication media and services, bringing up new challenges and possibilities. Also, new business models emerged as it was the case of small and medium-scale ITSPs (Internet Telephony Service Providers), which propose lower rates for telephone telecommunications using VoIP technologies for service delivery over the internet.

We intend to study and understand to what extent new service paradigms have an impact on broadband network environments, and how these infrastructures can be improved and better leveraged in order to benefit the operation and reliability of such services. Also following this line of thought, new service concepts will be also explored and validated, in order to assess its feasibility.

1.2.2 The potential of frontier mechanisms and devices

The specific role and position that frontier devices and mechanisms fulfill in the context of broadband network architectures – as mediators responsible for the exchange of information between the boundaries of the provider and customer networks – give them a unique perspective of both sides. The fact that some of these devices (such as *home/residential gateways*) have a reasonable computational capability, available with no additional costs and constantly improved, thanks to the evolving nature of commodity embedded hardware, raises the possibility of leveraging their potential to exploit new applications.

In the scope of this thesis's work, we will study the possibility of using frontier devices and mechanisms in order to:

- develop a flexible management framework for broadband environments, designed around the basic premises of wide-reachability and integration with LAN devices, management protocols and services;
- enhance specific service integration and operation on broadband access networks;
- propose novel, more scalable approaches for security management in this type environments.

By taking advantage of the processing and remote management capabilities of such devices we intend, through functional decoupling, to develop a decentralized and distributed architecture capable of dealing with those challenges.

1.2.3 Management, service integration and security

Broadband access networks pose significant challenges for Internet Service Providers (ISPs), as a consequence of several factors, from the high bandwidth available for each connection to the frequent situation in which customers lack the necessary technical knowledge to be self-sufficient in the management of their equipment and infrastructure. In this thesis, these challenges will be divided into three main categories: management, security and service integration.

About Management:

Triple play networks brought new devices to the residential/SOHO LAN ecosystem, such as *set-top boxes* (STBs) and voice gateways, whose management is to be performed by the provider in order to configure and remotely administer them. This trend has had repercussions at the industry level, together with the tendencies towards standardization and service convergence.

Initiatives such as Home Gateway Initiative [HGI] and Broadband Forum [BBForum] make it possible to deploy an extensive and coherent set of remote management and security services in the customer network, capable of monitoring its internal network (if desired) and the traffic flowing between it and the provider network. Both Broadband Forum and HGI have developed technical standards and recommendations to define a set of interfaces for security and remote management [TR-124, HGI2008, HGI 2006] operations of devices located at the customer network (designated CPE - Customer Premises Equipment).

Nevertheless, conventional management architectures remain reminiscent of the dial-up era in many aspects, still treating all CPEs as target devices in a centrally-managed infrastructure with scalability and flexibility limitations that render them inadequate to deal with the new generation of high-speed access networks and related challenges.

This thesis will study the management architectures and standards currently in use by broadband providers, their weaknesses and handicaps, proposing alternative and/or complementary models adapted to the needs of current high-speed, multi-service access networks. These models will explore the use of frontier devices as agents capable not only of collecting information but also of performing an active role.

About Service Integration:

The subject of service integration can be analyzed from different perspectives, depending on the nature of the services in question, such as:

- **Content delivery services**: such as VoD, or streaming media from third-party providers, using unicast or peer-to-peer technologies.
- **Service-Oriented Architectures (SOA) and Cloud-based services**: such as some Web 2.0 applications, online cloud-based storage and backup services or streaming application providers.
- **Operator-provided complementary services over broadband**: such as triple-play, VoIP, VoD using broadband connections as the transport medium.
- **Small-Office, Home-Office (SOHO)/Subscriber infrastructure remote management services**: such as remote PC and LAN management.

These services and applications are going to be addressed in the scope of this dissertation, in order to understand how frontier mechanisms and devices can contribute to enhance their reliability and functionality.

About Security:

ISPs clearly define the security perimeter to end at their borderline equipment, with the customer being responsible for his own borderline equipment (such as *home gateways*) and everything beyond that. In these situations, the increasingly available bandwidth, combined with P2P applications and *everything-over-IP* convergence scenarios only contribute to substantially aggravate the potential risks involved for the customer or the ISP, more vulnerable to orchestrated DoS (Denial of Service) attacks and abusive uses of its network infrastructure.

Even if some of the risks already existed before the emergence of broadband access networks, the transient nature of classic dial-up connections and the reduced amount of available bandwidth helped ISPs, making it easier for them to detect and control potential security threats or incidents affecting their own network, customers or third-parties. Nowadays, the centralized security model in use by ISPs is, to a large extent, largely derived from the dial-up era and, as a direct consequence of the increasing diversity of services and applications supported over broadband, is suffering from scalability issues, hampering the detection and response to security threats and events. Since the customer expects voice and television services delivered over broadband to have equal, if not better, performance and reliability when compared to their conventional counterparts, the impact resulting of service interruptions will be stronger.

As a whole, those trends created the ideal conditions for rethinking the problem of security in SOHO and access networks. They constitute an emerging scenario which demands the identification and characterization of a set of new security threats associated with it, creating the opportunity to research and develop a new approach on how to deal with the customer network. This research work will address the need to rethink, and therefore revise the established security model in broadband networks by presenting a security solution devised in order to better integrate the existing security mechanisms at the ISP and customer levels, without compromising users' privacy or freedom.

1.3 Scientific contributions and expected impact

The scientific contribution of the PhD work developed in this thesis can be synthesized as:

Improvement of the management model used in broadband environments by incorporating support for a complete ecosystem of services and devices, while providing security and protecting users' privacy. This is mainly achieved by exploring the potential of residential gateways as frontier devices, making them able to play an active part in the scope of the management, security and service operation models in broadband networks. This will be demonstrated by proposing, developing and implementing a distributed management architecture capable of leveraging the potential of frontier devices (such as domestic/residential gateways) in order to build a flexible and scalable framework oriented towards broadband environments, and also by showing its application to security and service integration purposes.

In more specific terms, our contributions can be enumerated as follows:

- **Design and validation of a modular and dynamically extensible management platform for broadband access networks**, which complies with Broadband Forum's CWMP (CPE WAN Management Protocol) protocol framework. This platform is able to interact with services and devices not covered by the original specification, being able to support proxying and protocol translation mechanisms allowing it to manage non-CWMP compliant devices and interact with internal LAN management and media services which otherwise would be inaccessible to the operator. By using residential gateways as management proxies for the entire subscriber LAN, it promotes added security and confidentiality by reducing the need for operators to directly interfere in the customer premises environment, retaining all the advantages of the CWMP framework with the additional benefit of helping to overcome some of its limitations in NAT environments.

- **Design and validation of CWMP-integrated desktop management mechanisms**. This line of research followed two parallel paths: the first one was oriented towards the integration of APIs for classic Microsoft Windows-based devices (namely, the Windows Management Instrumentation API) into CWMP environments for operator-assisted desktop and appliance management, while the second one researched new paradigms of managed cloud-booting thin clients for domestic and SOHO usage, with a view into Total Cost of Ownership reduction and improved reliability and usability.
- **Integration of LAN management protocols and non-CWMP devices** into the proposed management framework. This concept was demonstrated using two different implementations: one focused on integrating with the Universal Plug and Play (UPnP) discovery, configuration and control mechanisms into the proposed management framework; the other was focused on providing a CWMP-compliant provisioning and management interface for Session Initiation Protocol (SIP) telephony endpoint devices.
- **Development of managed broadband service models**, specifically for multimedia content distribution and hosted/cloud storage integration. An operator-managed media distribution solution was developed, with the purpose of turning devices compatible with the Digital Living Network Alliance (DLNA) framework for media distribution across the LAN into receivers for content provided outside the customer premises LAN, by turning the residential gateway into a media hub.

  In terms of managed storage services, two alternatives were developed and tested: one oriented for data synchronization between storage repositories and the other for online storage service access. In both cases, domestic gateways are used as mediators that provide the service access point for the customer premises LAN.
- **Development of a managed, distributed security solution**, which turns home gateways into active security elements, both in terms of threat detection and countermeasure enforcement. In terms of security, home gateways have a privileged positioning, between the LAN and the access network, which provides them a unique view of both worlds. To leverage their potential and demonstrate this concept, a distributed IDS/IPS (Intrusion Detection System/Intrusion Prevention System) was conceived. This solution is able to detect, correlate and react to threats using a two-level coordination structure, designed as part of a shared security model concept based on the premise of providing user interaction mechanisms (both for configuration and feedback) as a part of its operation.

In terms of impact, the work developed in the scope of this thesis has been subject to publication and/or presentation in the following conferences, workshops and/or journals.

(*Published/accepted papers*)

- **T. Cruz**, T. Leite, P. Baptista, R. Vilão, P. Simões, E. Monteiro, F. Bastos, “Segurança em Redes de Acesso Triple-Play”, Proceedings of SINO 2008 (4ª Conferencia Nacional de Segurança Informática nas Organizações), Coimbra, Portugal, November 2008.
- **T. Cruz**, P. Simoes, T. Leite, P. Baptista, R. Vilão, E. Monteiro, F. Bastos, “How to Cooperatively Improve Broadband Security”, Proceedings of ECIW 2009 (8th European Conference on Information Warfare and Security), Lisbon, Portugal, July 2009.
- T. Leite**, T. Cruz**, P. Simoes, T. Leite, P. Baptista, R. Vilão, E. Monteiro, F. Bastos, “Uma Plataforma para Gestão de Configurações em Redes de Banda Larga”, Proceedings of CRC 2009 (9ª Conferência sobre Redes de Computadores), Oeiras, Portugal, October 2009.
- R. Vilão, **T. Cruz**, P. Simoes, T. Leite, P. Baptista, R. Vilão, E. Monteiro, F. Bastos, “Avaliação Empírica de um IDS Distribuído para Redes de Acesso de Banda Larga”,

Proceedings of CRC 2009 (9ª Conferência sobre Redes de Computadores), Oeiras, Portugal, October 2009.

- **T. Cruz**, P. Simoes, T. Leite, P. Baptista, R. Vilão, E. Monteiro, F. Bastos, "Um IDS Cooperativo para Redes de Acesso de Banda Larga, in Proceedings of CIBSI 2009 (5th Ibero-American Congress on Information Security)", Montevideo, Uruguay, November 2009.
- **T. Cruz**, P. Simões, J. Almeida, P. Batista, E. Monteiro, F. Bastos, A. Laranjeira, "CWMP Extensions for Enhanced Management of Domestic Network Services" (poster paper), Proceedings of LCN'2010 (35th IEEE Conference on Local Computer Networks), Denver, USA, September 2010.
- **T. Cruz**, P. Simões, E. Monteiro, F. Bastos, "Integration of PXE-based Desktop Solutions into Broadband Access Networks", Proceedings of CNSM 2010 (6th IEEE/IFIP International Conference on Network and Services Management), Niagara Falls, Canada, October 2010.
- **T. Cruz**, P. Simões, J. Almeida, P. Bastista, E. Monteiro, F. Bastos, A. Laranjeira, "Gestão de Redes Domésticas com Agentes CWMP Extensíveis", Proc. of CRC 2010 (10ª Conferência sobre Redes de Computadores), Braga, Portugal, November 2010.
- **T. Cruz**, P. Simões, J. Almeida, J. Rodrigues, E. Monteiro, F. Bastos, A. Laranjeira, "How to Provision and Manage Off-the-Shelf SIP Phones in Domestic and SOHO Environments", accepted for publication in the Proceedings of LCN'2011 (36th IEEE Conference on Local Computer Networks), Bonn, Germany, October 2011.
- **T. Cruz**, P. Simões, J. Rodrigues, E. Monteiro, F. Bastos and A. Laranjeira, "Outsourced Management of Home and SOHO Windows Desktops", accepted for publication in the Proceedings of CNSM 2011 (7th IEEE/IFIP International Conference on Network and Services Management), Paris, France, October 2011.
- **T. Cruz**, P. Simões, J. Rodrigues, E. Monteiro, F. Bastos, A. Laranjeira, "Using UPnP-CWMP Integration for Operator-assisted Management of Domestic LANs", accepted for publication in the Proceedings of CCNC 2012 (IEEE Consumer Communications and Networking Conference), Las Vegas, USA, January 2012.

(*Submitted, pending review*)

- **T. Cruz**, P. Simões, E. Monteiro, F. Bastos, A. Laranjeira, "A Framework for Internet Media Services Delivery to the Home Environment", submitted to the Journal of Network and Systems Management (JNSM).

(*Presentations and talks given outside the scope of conference-published papers*)

- P. Simões, **T. Cruz**, "Um IDS Distribuído baseado na Cooperação entre o ISP e os Utilizadores Domésticos", delivered at Jornadas sobre Segurança em Sistemas Informáticos, Covilhã, Portugal, March 2009.
- **T. Cruz**, P. Simoes, T. Leite, P. Baptista, R. Vilão, E. Monteiro, F. Bastos, "Scalable Approach to Data Collection in Broadband Access Networks", communication delivered at DSTIS 2009 (8th International Conference on Decision Support for Telecommunications and Information Society), Coimbra, Portugal, September 2009.
- F. Bastos, A. Laranjeira, P. Simões, **T. Cruz**, T. Leite, P. Baptista, R. Vilão, E. Monteiro, "Projecto S3P: Segurança em Ambientes Triple-Play", delivered at PT Inovação projects workshop for 2009-2011, Aveiro, Portugal, March 2010.

Thanks to the co-sponsorship of PT Inovação (the research arm of Portugal Telecom, a Portuguese telecommunications operator), the research work we developed in the context of this thesis has also spawned results in the form of products which are being used in production environments. Among those, there are two projects that stand out: a scalable and vendor-neutral CWMP management server (ACS – Auto-Configuration Server) and a CWMP-based VoIP telephone provisioning system capable of supporting non-CWMP endpoints.

1.4 Structure of this document

Figure 1.1 presents the structure of this dissertation. It is divided into four main parts.

Part I includes Chapter 1. Chapter 2, which presents a state-of-the-art about past and current trends on related topics, addressing access network technologies, LAN and operator management frameworks, residential gateway service/device management solutions and engaging trends.

Part II corresponds to Chapter 3, which presents the management framework for broadband access networks which was conceived in this dissertation and provides the groundwork for the research conducted thereafter.

Part III, which includes Chapters 4 to 7, will discuss application scenarios and service integration topics, their validation and evaluation of specific architectural concepts.

Finally, Part IV concludes this dissertation, presenting the conclusions of the research work that was undertaken.

As a consequence of the diversity and specificity of related subjects, it should be stressed that for each chapter on Part III, there are separated and specialized state-of-the art sections. For the same reasons, instead of a single section devoted to validation, the evaluation and validation discussions are spread across each sub-section and discussed in a subject-by-subject basis. This has to do with readability: for each application paradigm the validation stage evaluates specific functional and performance characteristics which are particular to each subject and, therefore, are better addressed separately.

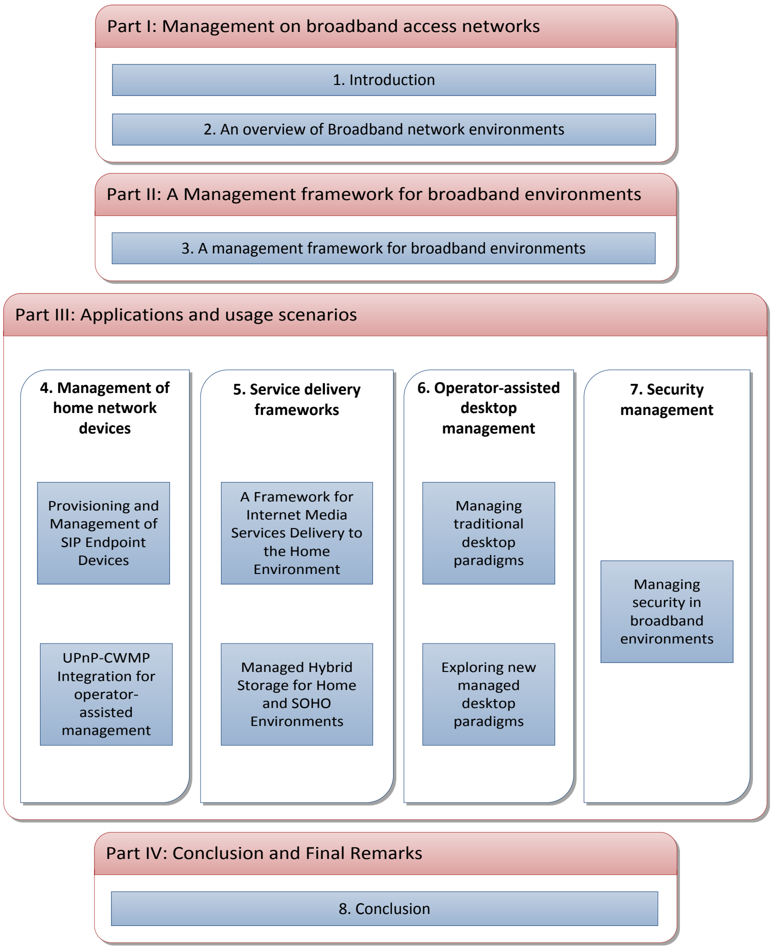
Part I: Management on broadband access networks
1. Introduction
2. An overview of Broadband network environments
Part II: A Management framework for broadband environments
3. A management framework for broadband environments
Part III: Applications and usage scenarios
4. Management of home network devices
Provisioning and Management of SIP Endpoint Devices
UPnP-CWMP Integration for operator-assisted management
5. Service delivery frameworks
A Framework for Internet Media Services Delivery to the Home Environment
Managed Hybrid Storage for Home and SOHO Environments
6. Operator-assisted desktop management
Managing traditional desktop paradigms
Exploring new managed desktop paradigms
7. Security management
Managing security in broadband environments
Part IV: Conclusion and Final Remarks
8. Conclusion

Figure 1.1: Document structure

2. An overview of broadband access networks and home networks

This chapter provides a look at the state of the art on the subject of broadband access networks and home networks, in terms of technologies, service paradigms and management. Its structure and contents are intimately linked to the context of the work presented in the next chapters. It is divided into five main sections.

Section 2.1 analyzes the evolution of the access network, giving special relevance to the aspects related to technologies, support infrastructures and service/usage paradigms. Whenever possible, concepts will be presented from an historical perspective, contextualizing significant milestones within their natural chronological evolution path. The purpose of this approach is to provide a clear insight into the reasons and circumstances which prompted the rise and/or demise of each technological milestone.

Section 2.2 deals with existing device and service interoperability frameworks for home environments. The purpose of this section is to present and detail the fundamental operation concepts behind these frameworks, most of which were designed for LAN operation but could also constitute a valuable management resource in a broader operator and service management scope.

Section 2.3 discusses the management frameworks for broadband access network environments from an operator-centric perspective, giving an insight about technologies either in use or with historical significance. The purpose of this discussion is to introduce Broadband Forum's CWMP (CPE Wan Management Protocol) framework, an emerging management standard which has become the *de facto* standard for device and service management in broadband environments..

Section 2.4 specifically deals with the role of residential gateways in broadband environments, analyzing their importance as mediator mechanisms in the intersection of the service, connectivity and LAN domains which form their operation environment. Special relevance will be given to the discussion of service and component management frameworks, since they play an important role in this context, enabling devices (and particularly, home gateways) to dynamically embed new functionality and service support features which are vital in IP-based consolidated service environments.

Section 2.5 discusses several yet-to-come technologies which may become important for residential gateways in the near/mid-term future.

Finally, section 2.6 concludes this chapter.

2.1 A perspective on access network technologies

With the possible exception of the telegraph, it is safe to assume that the first widely deployed access network infrastructure was the telephone system: whether it was manual or automatically switched, it was the earlier form of access network technology to be widely deployed and its expansion goals fueled the development of other technologies. Nevertheless, the first decades in the life of communication networks saw little change in the way they were used. Anybody used to the way telephones worked in the 1920s would find that things were not much different in the 1960s – in fact, one would only have needed to learn how to use a dial and, in certain cases, would inclusively have found that the devices were still connected to a manual exchange and, hence, still needed no dial. Actually what they would have noticed was a major change in the amount of people they could reach by telephone and the improvement on call and connection quality.

Later on, the introduction of transatlantic cables, microwave links and satellite communications, together with the introduction of digital processing technology contributed to improve the reach and quality of the telephone system, but its fundamental operational concept remained unchanged.

Historically, the second access network that came into existence was the Telex (used for remote teletype communications), whose first large-scale implementation was done in the 1930s, in Germany [Kimberlin1986]. It was created and maintained separately from the telephone network, and later on became a forerunner of electronic data exchange networks. Using the Telex, the output produced by a data processing system (e.g., an account transaction list) could be encoded on paper tape using a mutually agreed format – that tape would be transmitted by telex to a receiver that would subsequently produce the respective paper tape to be fed into the destination data processing system.

With the introduction of the first general-purpose computers in the 1950s rapidly came the need to remotely connect peripherals and other computers, thus creating the first data communication networks. With the invention of the Bell 101 dataset (110-baud) in 1958, it became possible to use the telephone network for TeletypeWriter eXchange (Bell TWX, a Telex competitor) – one of the first examples of technology convergence, allowing a teletype/Telex service to operate along with POTS (Plain Old Telephone System) on PSTN networks. It was the invention of the first full-duplex modem (the 300-baud Bell 103 dataset, a descendant of the 101 model), later on, in 1960, that launched the era of time-sharing computing. The modem, together with the arrival of the first fax service in 1962 (also based on modem technology), became the first use of the PSTN for anything else other than voice communications.

The next paradigm shift in terms of communications and computing only became possible thanks to the confluence of two technological landmarks – the Internet and the personal computer:

- The original ARPANET, which later grew into the Internet, was created in 1969. Still, its creation was significant only from the standpoint of a privileged minority, as it originally was a network designed for scientific and military purposes, its original usage context being heavily biased towards individuals with a technical way of life.
- The microcomputer, launched in 1974, started the era of personal computing. Every new computing paradigm has molded the way users perceive their status in relation to the computer. In opposition to the subservience and/or dependence users felt in the era of batch and time-sharing large-scale centralized systems, microcomputers gave them independence. For the first time, people could afford to own and use autonomous small-scale computers for their particular purposes.

Those two technologies had significant possibilities of their own but their full potential could only be achieved by converging them together. As it became evident with standalone PCs, whose data-processing potential is somewhat constrained by their isolation in terms of real-time data communication capabilities, the opposite is also true, as it was the case of Videotex [Krevitt-Eres1986] services. Launched with variable degree of success all around Europe in the 1980s and oriented towards data-presentation applications, with limited interactivity by using terminals no more intelligent than the dumb terminals of the 1970s, these services had a limited application scope and ultimately were superseded by the modern Internet.

Convergence ultimately started thanks to the opening of the Internet to public access in the early 1990s (albeit MILNET was already separated from ARPANET since 1983, the latter was restricted to academic, non-commercial usage being later integrated within NSFNET, which was handed over to private operators), together with the creation of the World Wide Web and the graphical browser [Berners-Lee1990]. Internet access and WWW browsers were the "killer apps" that allowed the shift towards the *connected* usage paradigm, by enabling new possibilities in terms of applications and interaction.

In those early days of public internet access (early-to-mid 1990s) residential users connected using dial-up modems through analog telephone lines. Apart from speed, dial-up had a number of other inconveniences, such as the fact of it taking up full use of the telephone line (people had to keep disconnecting and reconnecting if other members of the household wished to use the telephone, unless they had installed a separate phone line to connect the internet from) or the instability of connections. Because of the phone call rates, internet connections were of transient nature, taking only the time needed to download e-mail, browse some sites and little more.

However, even with a significant number of handicaps, dial-up was the dominant form of residential internet access through the 80s and 90s. Despite the emergence of alternatives such as ISDN (Integrated Service Digital Networks) technology, the price-performance ratio made the analog modem the dominant form of internet access for over more than 17 years. But the pace of change was about to accelerate: from the 1990s on, the Internet constantly evolved towards accessibility and content diversity, embracing a much broader audience in parallel with the growth in the number of domestic PCs, thanks to its widespread adoption in homes and small businesses.

This trend was the result of a symbiotic relationship between people and the technologies put at their disposal. In the early 2000's, the internet user base had been steadily growing and the use of online services became a way of life, sometimes to the extent of eliminating the need for daily commuting for many people. This scenario, together with the arrival of the dynamic-content web paradigms (based on Flash or Web 2.0 technologies), VoIP (Voice over IP), P2P (Peer-to-Peer), VoD (Video on Demand) or IPTV, among other services, led to the explosion of IP protocol traffic volume whose growth is constantly putting new demands on the access network in terms of changing traffic patterns (e.g. always-on connections as opposed to dial-up services), in terms of higher demand for bandwidth, and in terms of differentiated Quality of Service for different services over the same network. Those same requirements (and the introduction of flat-rate fees) were decisive to finally displace dial-up as the preferential internet access method, pushing the way for adoption of broadband technologies – available as early as 2000 [Anderson2002] but until then, suffering from a slow adoption curve because of cost and accessibility issues.

Initially, the residential broadband offer was primarily based on two technologies: ADSL and cable. ADSL was a technology that allowed for the delivery of broadband access over copper pairs. Cable networks, with their hybrid fiber and coax networks, were also competing for similar customers.

From the beginning to the mid part of the 2000s, and with the exception of some early-adopters, operators and service providers were not favorable to expand their residential data portfolio past Internet access. Meanwhile, several reasons drove the need for service diversification, as many operators with a wired access infrastructure found their revenue being eroded as a result of regulatory and market pressures. Soon became evident that deploying a broadband access network infrastructure to serve only as a means to deliver high-speed internet access without complementary value-added services was unattractive.

Triple-play and Multi-play[3] bundles delivered over broadband connections (Figure 2.1) were instrumental in helping operators to sustain their service margins in an increasingly competitive market. In a certain way cable operators had already helped establishing the idea of single-sourced service bundles in the customers' minds, which paved the way for other operators to also deliver their own services over a single copper pair or coax cable —the same cables used for telephone lines and cable television.

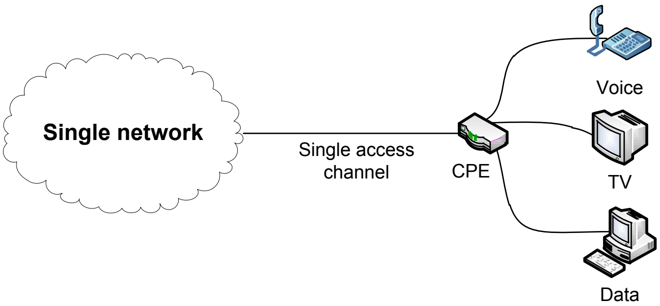

Figure 2.1: Converged, Triple-play services delivered over a single broadband connection

Anybody acquainted to the telephone system of the 1960s would be amazed by the services that are available now. In 1960, the telephone network was nearly exclusively used for voice, with the exception of insignificant amounts of data traffic from organizations that could afford it. Meanwhile, the telephone network started to carry voice, data, fax and internet traffic to residential users. By then, the nearest thing to mobility was the public payphone; nowadays over 60% of the population carry mobile phones, in some countries.

The widespread adoption of mobile phones was also a significant driver for change. Mobile phones started as basic voice devices, moved to providing short messaging facilities, and have morphed into digital assistants: thin computing platforms with text, video and voice communication capabilities, Internet access and organizer skills. Mobile phones own their popularity to a set of factors, namely attractive pricing packages, the feeling of security that constant contact brings, and even fashion. In some European countries there are now more mobile phones than fixed phone lines – in 2008 Finland had 61% of households exclusively served by mobile telephony versus 5% with fixed-only access (and 28% with both), the latter value corresponding to an 11% decrease from 2007 to 2008 [EU2008].

As for broadcast networks (TV and radio), they did not change significantly for a long time, when compared with their beginnings, and apart from the quality improvements brought by color TV and stereo sound reception (in FM radio and TV), they remained passive forms of communication, devoid of any interactivity. Digital TV, however, is helping change this scenario: nowadays broadcasters sell digital "pay-per-view" TV packages that provide interactivity, allowing viewers to choose the programming, consult electronic program guides or choose which camera angles to visualize.

In fact, one of the most important conclusions which comes to mind when we look back, is that communication networks evolved and were reshaped by both technological evolution and end-user demand. Nowadays, communication networks provide access to a diversified and versatile

[3] A Triple Play package is a bundle of Internet access, Television (IPTV), and Voice (VoIP) services delivered over IP. Multi-Play services are a natural extension of this concept.

array of services which would not even been thought as possible nor envisaged by those who placed the first copper pairs into the ground – this is especially true in the case of access networks, independently of the technology which supports them. These services are responsible for the increasing demands which are being imposed on the networks in terms of bandwidth, delay tolerance and mobility. The last mile, itself, is again being reshaped: provided over the same copper cables that supported the telephone (using analog or xDSL modems) or cable television systems, it has already reached its limits, and in some countries is already being gradually replaced by optical fiber [Morgan2011], as part of a strategy to provide higher bandwidth and service quality while securing a margin for accommodating both a growing subscriber base and technology advances.

Also, as ubiquitous services become the norm, the demand for mobile high-speed broadband services also increased. In response, the emergence of third generation wireless networks gave birth to the concept of mobile broadband, adding ubiquity and convenience to the mix, with basically the same benefits of conventional broadband. Femtocells [FemtoForum2011] are another example of how broadband access networks are helping to change the communications landscape, increasing coverage inside areas where mobile networks have limited reach and enabling the creation of fixed-mobile convergence scenarios.

In the face of this highly complex and dynamic scenario, the traditional symbiosis between network users and network operators calls for a balance:

- Users want a versatile, reliable and cheap access network – so they can immediately make use of the new services that emerge, or make a more versatile use of existing services, always at a reasonable cost. Any service, anytime, anywhere may be the motto that synthesizes this global trend.
- The service or network provider must be capable of reacting in a quick and economically-effective way to users' needs. As network operators strive to plan, create and constantly upgrade broadband services, there is a concern about ways of protecting their Return Of Investment, while expanding the subscriber base.

Until very recently operators had a separated vertical structure for each service, such as PSTN, Cable TV or Internet access, tailored to fit specific media requirements (Figure 2.2).

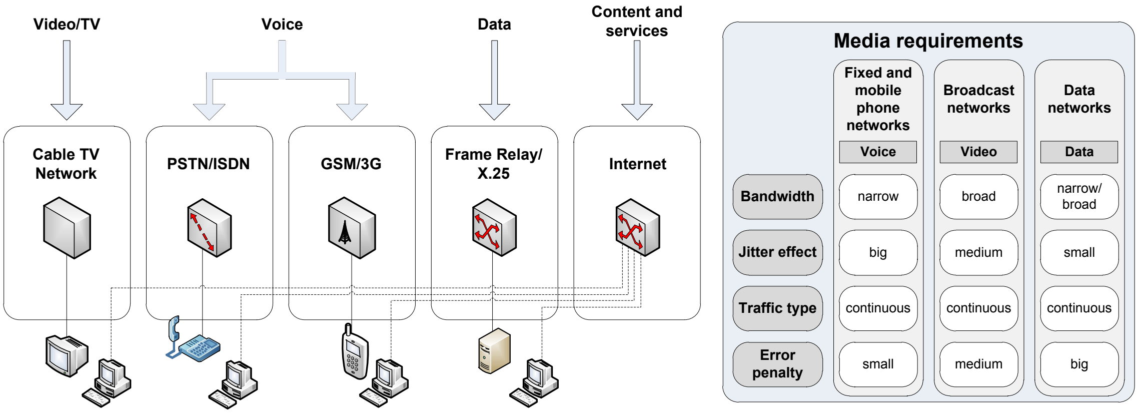

Media requirements

	Fixed and mobile phone networks	Broadcast networks	Data networks
	Voice	Video	Data
Bandwidth	narrow	broad	narrow/ broad
Jitter effect	big	medium	small
Traffic type	continuous	continuous	continuous
Error penalty	small	medium	big

Figure 2.2: Separated vertical network infrastructures for specific media requirements (adapted from [Fontes2005])

The introduction of Triple and Multi-play services, together with the increased pressure to compensate for decreasing revenue from traditional services (e.g., as is the case of the PSTN, where fixed voice is gradually giving place to mobile telephony) while trying to expand the subscriber base, called for a different approach and created the need for an integrated and rational development in the architecture of operator networks, geared towards optimization of diversified service delivery. This imposed a requirement on access networks for a great deal of flexibility that could only be achieved by shifting from the legacy vertical multiple-infrastructure model to an unified network suporting converged services using IP technology as the transport and, therefore, reducing logistic complexity (Figure 2.3).

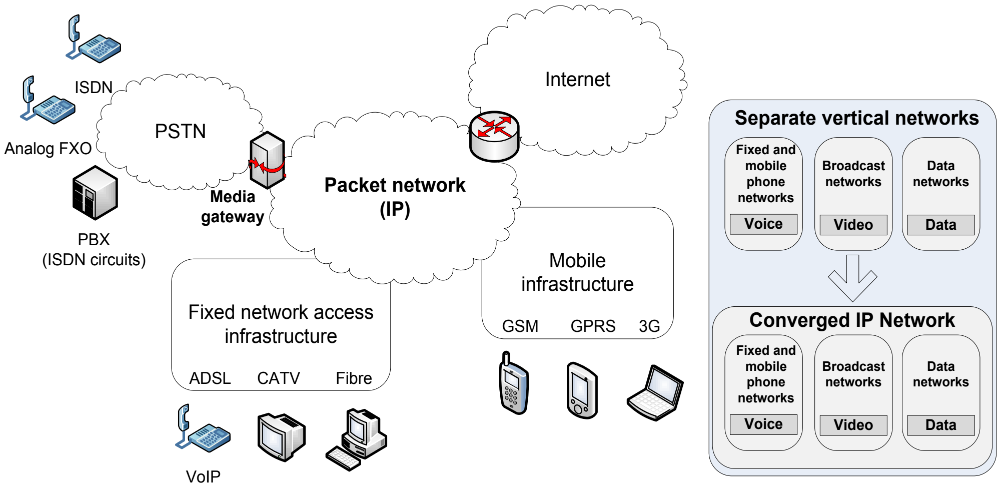

Figure 2.3: Converged IP network

This converged model, known as the NGN (Next Generation Network) [ITU2004], adopts a common access and transport layer based on IP packet switching/routing, with separate control and applications/services level. From a vertical perspective, this unified model shifts operator network architectures towards a different approach with separate horizontal planes (Figure 2.4).

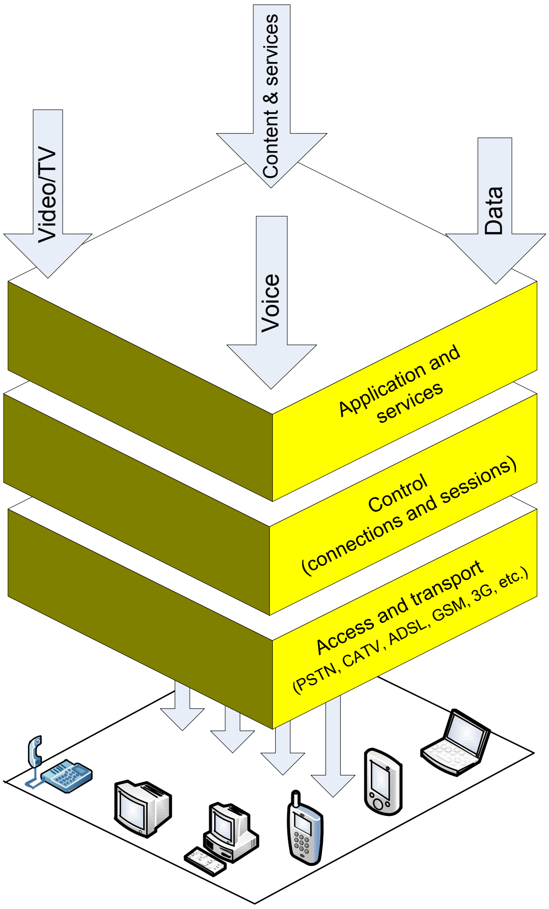
Video/TV
Content & services
Data
Voice
Application and services
Control (connections and sessions)
Access and transport (PSTN, CATV, ADSL, GSM, 3G, etc.)

Figure 2.4: Integrated horizontal planes

The service offer is decoupled from the network, supported on its own infrastructure, separated from the transport level. To enable a new service, providers only have to define it at the service layer level, without considering the transport layer – therefore, services become independent of the transport level. In this way, applications become progressively independent of the access network, with an increased degree of functionality being transferred to end-user devices (VoIP endpoints, PCs, Set-Top Boxes). In this architecture, the control plane for sessions and connections is also decoupled from the transport layer, assured by IP-based packet switching/routing.

In architectural terms, the move towards NGNs involves some changes, such as:

- Consolidation of legacy service-specific transport networks (either dedicated or overlay) into the core (IP-based) transport network. As a consequence, the PSTN migrates to VoIP and legacy services such as X.25 or Frame Relay to IP-based services, such as Virtual Private Networks.
- In the wired access network realm, NGN also brings convergence to an IP-based setup. For instance, in xDSL environments, where the legacy voice system is maintained in parallel to the xDSL infrastructure the move to converged services means that either DSLAMs provide voice ports or user might simply get a VoIP telephone (or adaptor) at his premises. The deployment of optical fiber in the last-mile is also a solution that is increasingly being adopted to ease migration towards NGNs, while ensuring a bandwidth capacity margin that can be used to provide new services in the future.

The NGN model is suited to support converged service offer over IP, such as triple-play or multi-play provisioned over a single broadband channel.

In face of the evolution towards the NGN model, some operators are starting to consider the possibility of extinguishing the traditional fixed line service and related infrastructure in order to eliminate the burden of maintaining two separate infrastructures. As of 2009, the FCC issued a notice of inquiry requesting comments from US operators regarding the possibility of moving the POTS to an all-IP infrastructure [FCC2009] that was generally well-received. Citing part of AT&T's answer [ATT2009]:

"Congress's goal of universal access to broadband will not be met in a timely or efficient manner if providers are forced to continue to invest in and to maintain two networks. Due to technological advances, changes in consumer preference, and market forces, the question is when, not if, POTS service and the PSTN over which it is provided become obsolete (...)

Indeed, perhaps the clearest sign of the transformation away from POTS and towards a broadband future is that there are probably now more broadband connections than telephone lines in the United States."

The demise of the fixed voice line is already accelerating in some countries (and not only in Europe, as it was already discussed). For instance, in the USA it is estimated that 22% of households had given up from their fixed lines in exchange for mobile phones, at a rate of 700,000 landlines being disconnected per month [ATT2009]. Also, as of 2010, VoIP services such as Skype had 663 million registered subscribers [TCPAPER2011] – and in the USA a single cable provider alone (Comcast) already provided broadband phone services to over 6.5 million consumers (14% total market penetration) since 2008 [Datamonitor2009].

Clearly, Sun Microsystems's late 1980s slogan "The network is the computer" [Weissman2011] was prophetic. From the early 2000s, a new trend towards service and content delivery is slowly displacing computing power as the main user concern and shifting it towards network accessibility and flexibility. Together with the evolution towards NGNs and decoupled connectivity, service, content and even application delivery over converged mediums, the emergence of a new *Cloud Computing* paradigm is contributing to reinforce this tendency.

The introduction, in the first years of the 2000s, of Web Service technology, rapidly followed by Service Oriented Architectures (SOA) and the Software as a Service (SaaS) deployment models enabled the possibility of on-demand software and service provision. Cloud Computing is an immediate offspring of these technologies, offering a consumption and delivery model for IT services based on the Internet, involving the provision of dynamically scalable and often virtualized resources as a service. A Cloud Computing environment must comply with five critical characteristics: on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service.

The distributed nature of the concept calls for computing and network core and access infrastructures capable of meeting the needs of communication, application, service and content providers, through clearly defined Service Level Agreement (SLA) terms ensuring that vital QoS, availability, flexibility and interoperability requirements are met.

Since 2009, and thanks to Cloud Computing, converged services and NGNs, a business model shift is underway for network operators, gradually evolving from the role of connection providers to service, application and content providers. More than ever, network access technologies play a vital role in this process, by fulfilling the role of a confluence point in which content, service, application and connection providers interact with the ones that ultimately demand innovation but also provide revenue: consumers.

2.2 Managing the customer premises LAN environment

The original home/SOHO LAN was a direct descendant of the workgroup network paradigm that became popular in the 90s. Sometimes not more than a mash-up of more or less loosely networked devices (mainly Desktop PCs and, occasionally, printers) the first home LANs were created to ease resource sharing. The idea of being able to share and access a storage device or a printer attached to another computer was the main reason behind their popularity.

Meanwhile, domestic LANs evolved into dynamic environments supporting services such as file and printer sharing, media distribution or domestic automation. Albeit small-sized, their configuration – a crucial issue when device interoperability is a must - is an increasingly complex task that the common user is frequently unable and, sometimes, even unwilling to deal with. In these unmanaged LAN environments, enterprise management frameworks and protocols are of little or no use.

The solution for this problem calls for a different type of approach, in which automation becomes the key to a friendlier user experience, therefore hiding the complexities of device configuration and interoperability from end-users. To this purpose, the industry has come with different proposals, among which the most representative are discussed in this section, namely: Jini/River, Zeroconf (and its incarnations), UPnP (Universal Plug and Play), DLNA (Digital Living Network Alliance), DPWS (Devices Profile for Web Services) and the Windows Rally framework. The section will end with a comparison of these proposals.

2.2.1 Jini/Apache River

Jini (Java Intelligent Network Infrastructure) [Sun2001] is a Java-based framework introduced by Sun Microsystems in 1998 and later transferred to the Apache foundation, where it was renamed to "River" [Apache2010]. On its essence, it is a service architecture framework for developing distributed systems that communicate using an RPC mechanism (the Jini Extensible Remote Invocation (JERI) protocol or classic Java Remote Method Invocation). It was originally developed to ease the creation of a network-centric ecosystem for domestic appliances, in order to supersede PCs as the primary computing devices on domestic environments.

Jini/River is based on Java, and although JERI is similar to Java Remote Method Invocation, it is more advanced, as it provides the means for locating services, through a process of discovery of published services (making Jini akin to the service-oriented architecture concept). However, it does not deal with network auto-configuration.

Jini/River makes use of a so-called "Plug and Work" model that enables a device to connect to the service network without user intervention, relying on two basic concepts: code downloading and remote interfaces. The first allows a device to automatically download data and functionality while the latter provides decoupling from the strict dependency of a service implementation, enabling a device to dynamically learn how to use a new service. In Jini/River, a service entity is the resource to be made available in the distributed environment, which may encompass physical devices (i.e., a printer) and software services (a query service).

For service discovery the Jini/River architecture relies on two mechanisms:

- **Discovery protocol:** allowing a device to poll the local network to locate a lookup service.
- **Lookup Service:** which provides service location mechanisms and provides the means for clients and servers to download and upload interface objects that provide the interfaces for services (clients use late binding).

The discovery protocol relies on a mix of announcement and request methods, which use multicast UDP whenever possible. A multicast request protocol is used when a service needs to discover all the lookup services present on the local network. A service sends multicast requests and receives unicast replies from lookup services around the network. There is also a unicast discovery protocol that can be used when a service or client already knows the location of a lookup service, allowing request to be performed directly – it may also be used to contact non-local lookup services which may be beyond the multicast reachability radius, or after a successful multicast discovery request.

There is also a multicast announcement protocol which is used by lookup services to advertise their presence to interested peers that may be within the multicast reachability radius. When a new lookup service is added to a local network, it will first announce itself using a multicast announcement message, which will also be sent periodically.

The Jini/River operation flow can be described as such (Figure 2.5):

- **Service provider discover & join:** the first step in creating a Jini service is to find the lookup service using the discovery methods (1). Once it is found, the lookup service returns a service registrar object, which is used to register the new service in the lookup service. When a provider is registered with the lookup service, it uploads a service proxy object to it which provides the interface for that service (2). To use a service, a client only needs to locate it using the lookup service and download the required proxy objects (stub objects) from it.
- **Client discover and lookup:** when a client wants to use a service it uses the discovery methods to locate a lookup server (3). Once found, a lookup service receives a request (4), to which it will answer with a lookup service registrar object, used to look up a particular service. The client will consult the catalog on the lookup service, using the type, name or service description as search criteria. Once the desired service is found, the lookup service will return a proxy object (a stub object), specifying how to connect directly to the service (5). The use of proxy objects allows clients to directly connect the implemented service without further interaction with the lookup service (6).

The use of service leases allows Jini/River to deal with the transient nature of services. Upon service registration, a lease is granted with a defined duration, which needs to be renewed before expire in order to remove a registration for a service that has failed or has become unreachable. Jini/River services can also be federated together, allowing a client to search for specific groups.

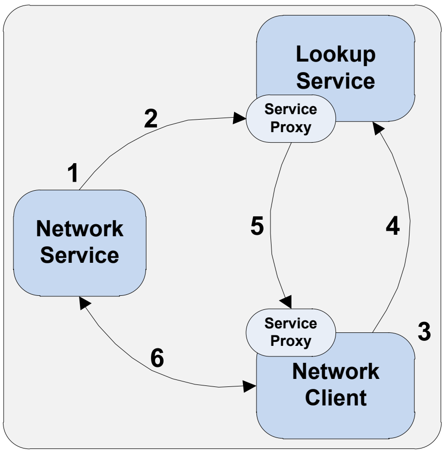

Figure 2.5: Jini/River operation flow diagram (adapted from [Sun2001])

2.2.2 Zeroconf

The IETF Zeroconf [Zeroconf, Williams2002] working group was formed in 1999. In 2003 it published its work on Dynamic Configuration of IPv4 Link-local Addresses (IPv4LL) [RFC3927]. IPv4LL was implemented and supported in several operating systems platforms such as Mac OS (9, X), the Microsoft Windows family (from 98 on) and Linux. Zeroconf relates to a set of techniques primarily targeted towards the automatic creation and configuration of IP networks, relieving users from the task of configuring all the related services (such as DNS or DHCP) which are needed to run a domestic network. It deals with three main tasks:

- **Automatic assignment of network addresses:**

  For **IPv4**, it uses the 169.254.0.0/16 address block [RFC3330] and is covered by RFC 3927 (IPv4LL). When a host wishes to configure an IPv4 Link-Local address, it selects an address using a pseudo-random number generator with a uniform distribution in the range from 169.254.1.0 to 169.254.254.255 inclusive. After it has selected an IPv4 Link-Local address, a host must test (using ARP probing) to see if the IPv4 Link-Local address is already in use before starting to use it.

  For **IPv6, it** uses the prefix fe80::/10 together with the device MAC address, as per [RFC4862].

- **Automatic resolution and attribution of device names:** covered by multicast DNS (mDNS) [Cheshire2006] or Microsoft's Link Local Multicast Name Resolution (LLMNR) [RFC 4795]. Both are very similar in nature (besides minor differences) and follow the basic multicast DNS operation principle: they use APIs which are similar to the unicast DNS system but with a different implementation. Each computer on the LAN stores its own list of DNS records (e.g. A, MX, PTR, SRV, etc) and when a mDNS client wants to resolve the IP address of a PC using its name, the PC with the corresponding A record replies with its IP address.

  mDNS allows a network device to choose a domain name in the ".local" namespace and announce it using a special multicast IP address. It is covered on an IETF draft [Cheshire2006]. The mDNS multicast address is 224.0.0.251 (FF02:0:0:0:0:0:0:FB for IPv6) and queries are performed on UDP port 5353.

  LLMNR allows a network device to choose any domain name, which is considered a security risk by some members of the IETF. The LLMNR multicast address is 224.0.0.252 (FF02:0:0:0:0:0:1:3 for IPv6) and queries are performed on UDP port 5355. It is not compatible with the DNS-SD Service discovery method next described.

- **Automatic location of network devices and related services:** covered by DNS-based Service Discovery (DNS-SD) [Cheshire2005], UPnP's Simple Service Discovery Protocol (SSDP, described in Section 2.2.3) or IETF's Service Location Protocol (SLP) [RFC 2608].

  In **DNS-SD**, service instance enumeration is performed by DNS queries (it is compatible with both unicast DNS and mDNS) for SRV records [RFC 2872] for a particular type of service (using the *<servicetype>.domain* format, where *servicetype* is from a list maintained by the DNS-SD task group – for instance, *_daap._tcp* refers to Apple's iTunes music streaming service). When a client needs to contact a particular service, identified by its service instance name, it queries for the SRV and TXT records of that name. The SRV record for a service gives the port number and target host name where the service may be found, while the TXT record gives additional information about the service.

  In **SLP**, client applications are called "User Agents", while services are advertised as "Service Agents" – "Directory Agents" also exist to provide scalability. When a service is to be located, a User Agent issues service request on behalf of a client application specifying the characteristics of the required service. The User Agent will get a reply specifying the location of all service instances on the network which satisfy the criteria. User Agent requests might be processed using multicast queries directly to Service Agents (to which Service Agents reply with unicast responses) or unicast queries to Directory Agents, which cache Service Agent information (from unicast requests or announcement information). Directory Agents announce to the network sending advertising messages using both multicast (when they start up) and unicast messages to User and Service Agents. User Agents must use Directory Agents when present, therefore avoiding the use of multicast requests, decreasing network utilization and increasing response speed.

There are several different incarnations, related to different manufacturers that implement each of these functionalities using a different protocol set, even if there are specific standards which address each one. For instance:

- Apple's Zeroconf implementation, called Bonjour [Apple2011] (born "Rendezvous") uses IPv4LL with mDNS and DNS-SD;
- Microsoft's implementation uses IPv4LL with LLMNR and SSDP;
- Linux's implementation, called Avahi, uses IPv4LL with mDNS and DNS-SD.

2.2.3 Universal Plug and Play (UPnP)

To simplify and/or automate device interoperability, the Universal Plug and Play (UPnP) Forum [UPnPForum] created a framework for the management, configuration and control of networked devices and appliances, which was introduced in 1999. This specification was mainly designed with the domestic LAN scope in mind, allowing for seamless device discovery, control and configuration, with minimum user intervention. It supports a wide range of devices, from printers to network equipment, having become an important mechanism for automatic configuration of devices within domestic LANs.

The UPnP ecosystem is a collection of networked devices using UPnP protocols to advertise, discover and access services in a seamless way, with minimum user intervention. Device and service profiles were developed for specific device categories, such as Audio/Video (AV), Internet Gateway Device (IGD), Printing and Lighting Control and even Heating, Ventilation and Air Conditioning (HVAC), among others. For each, a set of XML schemas was defined, describing the basic set of functions and services for that specific device type (the DCPs: Device Control Protocols [UPnPForum2011]). Figure 2.6 shows the UPnP protocol stack.

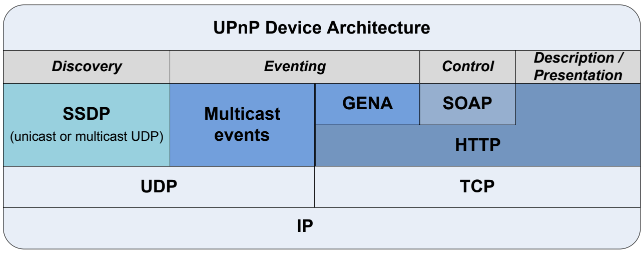

Figure 2.6: UPnP protocol stack

UPnP relies on the Simple Object Access Protocol (SOAP) [W3C2007] for device control, bringing together XML and HTTP to provide a Web-based messaging and remote procedure call mechanism. Specifically, UPnP device operation is a multi-stage process, comprising:

- **Automatic address configuration**, through which devices get a network address, typically using DHCP (Dynamic Host Configuration Protocol) [RFC2131], with fallback to AutoIP [RFC3927] when DHCP is not available.
- **Device discovery.** Devices advertise their presence to control points (managing entities) using Simple Service Discovery Protocol (SSDP) [Goland1999] multicast UDP messages with basic information about devices, services and a description URL which can be used to gather further information (Figure 2.7). Also, when a control point is started, it may send a multicast message to the network asking for other devices – when UPnP devices receive that message they respond to the requester, so that it can become aware of UPnP devices present on the LAN. Control points can also send unicast or multicast SSDP messages to search for specific instances (answered with unicast SSDP messages).
- **Description.** Using an URL provided in the discovery process, a control point can access an XML description of the UPnP device via HTTP (Figure 2.7). An UPnP device may embed other devices, each one with its own set of services. An UPnP service description includes a list of commands, or actions, to which the service responds, and parameters, or arguments for each action. It also includes a list of variables that model the state of the service at runtime, described in terms of their specific characteristics.
- **Command and control.** Once a control point has learned about the capabilities of a device, using its description, it is able to exert control by means of SOAP invocations of methods declared on its service descriptions (Figure 2.7).
- **Event generation.** By using the General Event Notification Architecture protocol (GENA – Figure 2.6) [Cohen1998], control points may subscribe to service events, for monitoring purposes. Such events are associated with the change of notification-enabled variables on service instances (Figure 2.7).
- **Presentation.** Devices might allow further user interaction via a HTTP presentation page, enabling direct control of specific device features, monitoring or other abilities.

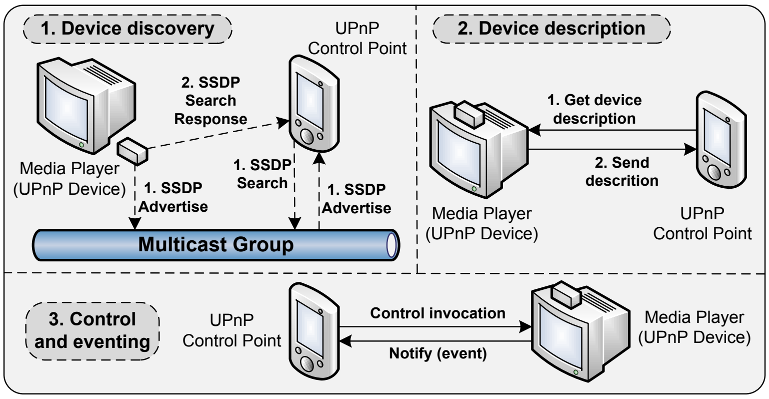

Figure 2.7: Basic UPnP operation and communication

2.2.4 Digital Living Network Alliance (DLNA)

The Digital Living Network Alliance (DLNA) [DLNA] framework is defined by a consortium formed in 2003 in order to develop interoperable digital media devices for the home, having among its members some of the biggest global brands in PC, Consumer Electronics (CE) and mobile electronics. Its goal is to allow seamless sharing of multimedia content within the home, among interoperable devices (like PCs, CE, and Mobile devices, among others) using already established industry standards.

The UPnP AV framework is the cornerstone technology for the DLNA architecture, which defines baseline guidelines, called Interoperability Requirements, organized into eight categories (see Fig. 2.8):

- **Network connectivity:** Ethernet, Wi-Fi and Bluetooth are specified as baseline network connectivity technologies.
- **Security:** mandatory support for WPA2 [IEEE2004] on WLAN.
- **QoS:** support for Wi-Fi Multimedia (WMM) [IEEE2005] QoS.
- **Device discovery:** by using UPnP protocols devices can find and identify each other. Media player devices will automatically search for Media Server devices that meet their specific criteria. DLNA specifies UPnP v1 as a baseline requirement for version 1.5.
- **Content discovery:** UPnP mechanisms are used to search and browse content in Media Server devices as requested by Media Players.
- **Media Transport:** HTTP (mandatory) or Real-time Transport Protocol (RTP) [RFC3550] (optional) might be used for stream transport.
- **Media format profiles:** DLNA version 1.5 mandates JPEG [ISO10918-1] for still images, Linear PCM for audio and MPEG-2 [ISO13818-1] for video. Optionally, Dolby Digital/AC-3 [Todd1994] audio or MPEG-4 [ISO14496-1] video are also supported.
- **Content protection:** DLNA 1.5 defines the use of DTCP-IP (Digital Transmission Content Protection over IP) [Hitachi2010] for link protection purposes.

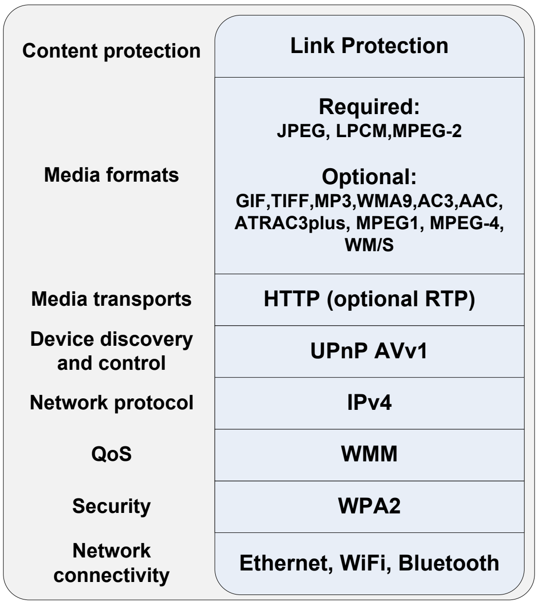

Content protection	Link Protection
Media formats	Required: JPEG, LPCM,MPEG-2 Optional: GIF,TIFF,MP3,WMA9,AC3,AAC, ATRAC3plus, MPEG1, MPEG-4, WM/S
Media transports	HTTP (optional RTP)
Device discovery and control	UPnP AVv1
Network protocol	IPv4
QoS	WMM
Security	WPA2
Network connectivity	Ethernet, WiFi, Bluetooth

Figure 2.8: DLNA version 1.5 interoperability guidelines (adapted from [Chou2006]).

The DLNA consortium also developed a certification process, complementary to these interoperability guidelines.

In the scope of DLNA, the UPnP AV (Audio/Video) specification provides mechanisms for content discovery, search and browsing of Media Server devices, as requested by Media Players and Control points. Its basic fundamental components are (see Fig. 2.9):

- **Multimedia control point**, which discovers media servers and media renderers, and connects them.
- **Media server**, which stores content on the network for access by media renderers.
- **Media player**, which renders content from a media server.

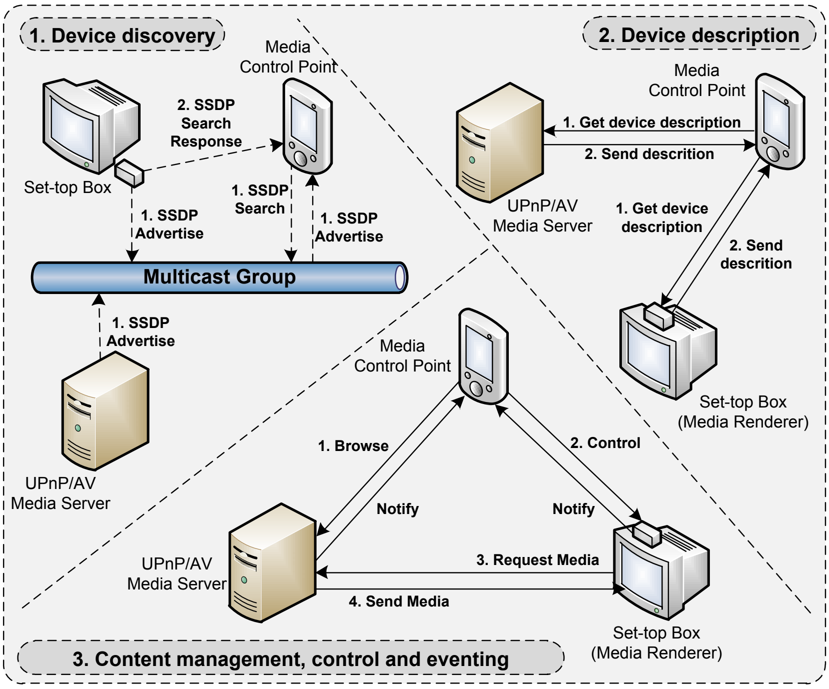

Figure 2.9: Basic DLNA (UPnP AV) operation and communication

When a multimedia control point is started (like an UPnP AV remote control, for instance), it may send a multicast message to the network asking for media servers and media renderers (otherwise, a control point may learn of a device of interest because that device sent discovery messages advertising itself). When the UPnP devices receive the message sent by the UPnP control point, each one responds to the control point requester. At this point, the UPnP control point knows about all of the connected UPnP devices (Stage 1 – Device discovery, in Fig. 2.9).

When a control point is interested in a device and wants to learn more about it and its capabilities, or to interact with it, it must retrieve a description of the device and its capabilities from the URL provided by the device in the discovery message (Stage 2 – Device description, in Fig. 2.9).

The user, using the interface provided by the control point, browses and selects any media content item listed by the UPnP media server. After the selection is finished, the user chooses a media renderer that will render (play) the multimedia item. Next, the user clicks on the Play button of the UPnP control point, and the media renderer starts to play the multimedia stream sent by the media server. Also, a control point might subscribe to service events, in order to receive notifications to track and monitor its status (Stage 3 – Content management, control and eventing, in Fig. 2.9).

2.2.5 Devices Profile for Web Services (DPWS)

DPWS [OASIS2009] (also called WSD – Web Services on Devices) is endorsed by the Organization for the Advancement of Structured Information Standards [OASIS] and defines a set of implementations guidelines enabling secure web service-based messaging, discovery, description and eventing. This standard was designed for resource-constrained devices such as embedded devices and other kinds of appliances, being originally positioned to succeed the UPnP 1.0 specification. While it may provide most of the functionality of UPnP, DPWS is fully conforming to the web service specification and has extensions for scenarios outside home LANs, such as enterprise environments.

In the DPWS scope devices run two types of services: hosting and hosted. Hosting services are directly bound to a device and relate with its discovery process, while hosted services provide functionality and depend on the hosting service of a device for discovery. DPWS also defines a set of built-in services, namely:

- **Discovery and description services:** providing mechanisms for device and service advertising, discovery and description.
- **Metadata exchange services:** providing access to device hosted services and their metadata.
- **Publish/subscribe services:** enabling other devices to subscribe to asynchronous events from a given service.

DPWS is based on existing Web Services standards such as XML, WSDL (Web Service Description Language) [W3C2001], XML Schema [W3C2004], SOAP or HTTP. It is implemented on recent versions of the Windows desktop (Vista, 7) and consumer platforms (Windows CE 6.0 R2) as part of the Windows Rally framework, which will be next discussed. It is also becoming popular in home and even industrial automation scenarios.

Its operation follows the same basic interaction paradigm also used by UPnP (discovery, description, command and eventing):

- **Discovery:** to find a device, the client sends an UDP multicast probe message. The probe message specifies the search criteria and whenever the client requires security. This probe message and the associated multicast discovery protocol are defined in WS-Discovery [OASIS2009a]. DPWS is designed to provide scalability for bigger, managed environments, enabling it to surpass the constraints of the multicast peer-to-peer model by using proxy mechanisms (a discovery proxy). When a discovery proxy is present on the network, multicast behavior is suppressed and clients communicate directly with it.

  The device replies with a unicast probe match message directly to the client, specifying: the SOAP-level address for endpoint reference; the available transports for reaching the device (SOAP over HTTP is mandatory, but other transports are supported); and security capabilities and requirements for the device. If security is needed, a secure channel will be negotiated, after device authentication.
- **Description:** a client sends a get metadata request directly to the discovered device, to which the device will respond either with the information embedded within the response message or by including a pointer to its location. This information specifies the hosted services, which can be queried by sending a get metadata message directly to them – each service will reply with the WSDL description for the control API of the service and other capability metadata.
- **Command and control:** control is exerted by SOAP invocations of the control methods described in the service WSDL description.
- **Eventing:** WS-Eventing [W3C2006] is based on a publish-subscribe model, in which clients can subscribe to receive service events (similar to UPnP's GENA).

2.2.6 Windows Rally Program

The Microsoft Windows Rally program [Microsoft2006] was an attempt to unify the device ecosystem, establishing a common framework for connectivity and interoperability. The result of this program found its way in more recent versions of the Microsoft Windows Operating System family. The Windows Rally Program addresses the following functionalities:

- **QoS diagnostics, network device and topology discovery:** Microsoft's proprietary Link-Layer Topology Discovery Protocol (LLTD) [Microsoft2011]. Its scope and aim is akin to IEEE's Link Layer Discovery Protocol (LLDP) [IEEE2005a]. LLTD operates at layer 2.
- **Network configuration:** *Windows Connect Now*, Microsoft's implementation of Wi-Fi Aliances' Wi-Fi Simple Configuration protocol, later renamed Wi-Fi Protected Setup (WPS) [WifiAlliance2007]. Windows connect now enables Ethernet out-of-band or wireless in-band configuration scenarios. It supports several usage models such as: Push-Button Configuration (PBC – where the user only has to bush a button on the Access Point (AP) and the client device for the initial setup to be performed), Personal Identification Number (PIN – which is placed on a sticker on the AP or client device and is used for establishing an initial configuration session); Media Transfer Protocol (MTP, using an out-of-band USB cable); USB (via an USB stick, albeit this method is being deprecated) or UPnP for wireless access point detection and configuration (for out-of-band wired configuration).
- **Device and service discovery:** DPWS and UPnP under the PnP-X [Microsoft2010a] platform (using DPWS/WS-Discovery and UPnP/SSDP discovery methods specific of each framework). PnP-X (Plug and Play eXtensions) adds device detection and configuration functionality for networked devices which is similar to the one that is provided for closely connected devices on USB or PCI buses.

The Windows Rally Program technologies are orthogonal to Microsoft's specific Zeroconf implementation for network auto-configuration (see Figure 2.10).

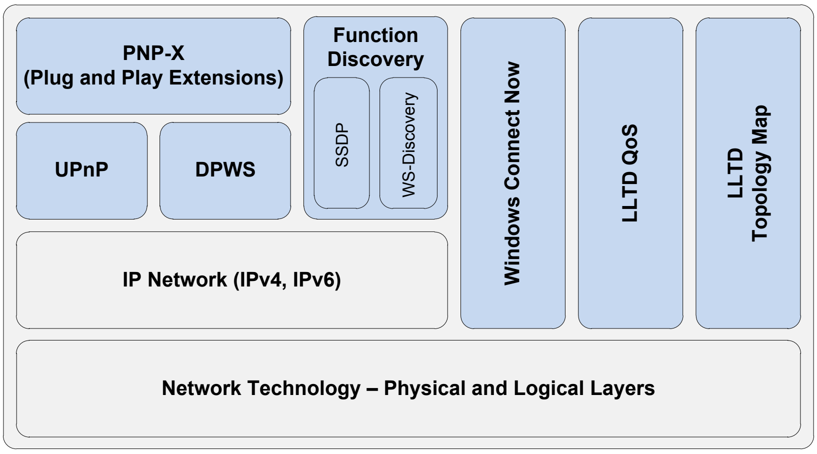

Figure 2.10: Windows Rally Project architecture (components highlighted in blue)

Unlike other technologies, such as UPnP or DPWS, LLTD is the significant novelty in the Windows Rally architecture, being used for two different purposes:

- **Network topology discovery:** for this purpose, LLTD uses a two-stage process. First, through quick discovery, it searches and enumerates all LLTD-capable stations on the network and their properties, called responders – to avoid network flooding, a special algorithm called RepeatBAND [Microsoft2011] is used. After quick discovery, the enumerator knows of available responders and the types of networks they are connected to (such as wired or wireless Ethernet). The behavior of the interconnection infrastructure of the network (switches, hubs) is inferred by the mapping session.

  Next, network topology tests are performed. For topology tests, stations might assume one of two roles: mapper and responder. In the quick discovery process, each responder associates with a mapper (if it doesn't already have such association) – this binding enables a responder to accept and answer topology test commands sent by the mapper. Since all responders report their association in the quick discovery stage, mappers are able to stop and release all their active associations if they detect a quick discovery packet from a responder associated to another mapper – the purpose of this behavior is to ensure that only a single mapper exists per Ethernet broadcast domain.

The LLTD API enables discovery of link details, but the construction of the network topology is up to the application layer, using specific algorithms. Such applications are able to request that a specific responder send LLTD frames with a predefined source and destination MAC Address (the source MAC may or may not be the one of the responder – in the latter case, a reserved range of MAC Addresses is used).

- **QoS diagnostics:** LLTD can be used for detection of network traffic characteristics or bottlenecks, such as available bandwidth, congestion problems or the existence of traffic prioritization mechanisms. A local application might assume the role of test controller, being able to use responders as traffic sinks for network path analysis. As such, an application is able to use LLTD to perform a cross-traffic analysis, identifying all responders from which it wants to obtain traffic counters, for capacity and performance auditing.

2.2.7 Framework Comparison

In terms of popularity, the UPnP and DLNA (based in UPnP AV) frameworks rank in the first place, having the biggest supported device base, followed by Apple’s Bonjour (Zeroconf), DPWS (widespread in home, health and industrial automation scenarios) and, finally, Jini/River.

Some frameworks, like UPnP or DPWS, are more device-oriented, while others, like Jini are service-oriented. Also the discussed frameworks differ significantly in terms of functional capabilities: while Jini, DPWS and UPnP focus primarily on device/service discovery, description and control/eventing, Zeroconf focuses on network auto-configuration and device/service discovery. The Windows Rally Project consolidates support for several frameworks (like UPnP and DPWS) to which it adds wireless network configuration (using Windows Connect Now technologies) and network topology discovery (using LLTD) functions.

Figure 2.11 compares all the frameworks previously discussed.

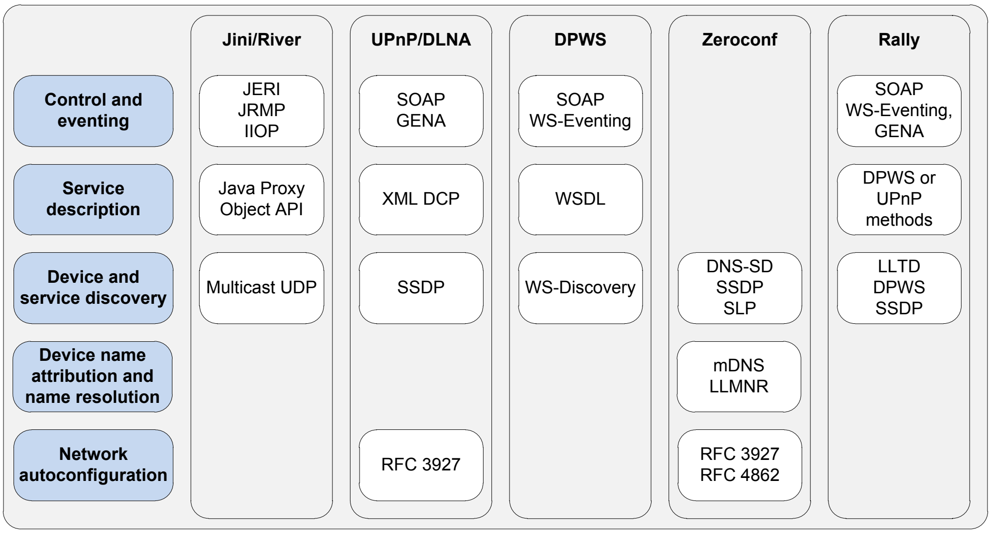

	Jini/River	UPnP/DLNA	DPWS	Zeroconf	Rally
Control and eventing	JERI JRMP IIOP	SOAP GENA	SOAP WS-Eventing		SOAP WS-Eventing, GENA
Service description	Java Proxy Object API	XML DCP	WSDL		DPWS or UPnP methods
Device and service discovery	Multicast UDP	SSDP	WS-Discovery	DNS-SD SSDP SLP	LLTD DPWS SSDP
Device name attribution and name resolution				mDNS LLMNR	
Network autoconfiguration		RFC 3927		RFC 3927 RFC 4862	

Figure 2.11: Interoperability framework comparison

Most of these solutions are designed with LAN environments in mind. Even if UPnP and DPWS have provision for internetworked operation, their reliance on multicast or link-local protocols for discovery methods (especially in simpler networks) means that functionality is lost in these scenarios – as it happens with DPWS directed discovery methods. In practice, for these protocols, internetworked operation is only supported through bridged protocol operation using unicast. Also, supported subscriber-based eventing mechanisms such as GENA (used for UPnP) or WS-Eventing (used for DWPS) are not able to operate when the subscriber is not directly addressable (as it happens in NAT environments).

Also, it should be noted that the Zeroconf approach only deals strictly with network device auto-configuration, name resolution and device discovery. Other, higher-level functions are deemed outside its scope: device control or eventing APIs are considered specific of each registered service or device and only concern the peers involved in the communication process. For instance, Apple's Airplay [Apple2011a, AT2Team2010] technology is an example of such services, which depend on the Bonjour framework for discovery.

The frameworks and standards for domestic LAN interoperability enumerated and described in this section are the most significant in terms of supported device base and popularity. However, other efforts should be mentioned such as TEAHA (The European Application Home Alliance – also covering device interoperability) [TEAHA], IGRS (Intelligent Grouping and Resource Sharing – whose scope is very similar to UPnP) [IGRS, ISO14543-5], ITU-T DHF (Digital Home-network) [ITU2002] and ISO/IEC HES (Home Electronic System – that uses IGRS as its main interoperability framework, as defined by the ISO/IEC 14543-5 series recommendations) [ISO14543-5]. However, despite their interest, all these efforts have failed to achieve significant expression in terms of popularity and adoption to present date.

Moreover, it should be stressed that the frameworks hereby discussed were limited to the ones that are designed to provide device interoperability and/or auto-configuration in an IP-based LAN scope. For this reason, other architectures such as HAVi (Home Audio/Video Interoperability - for AV device interoperability, which uses IEEE 1384/Firewire for transport but has been abandoned) [Teirikangas2001], LonWorks (for domestic and building control) [ISO14908], X10 (for domestic automation) [Rye2006], Z-Wave (for low-power wireless device communication) [Z-Wave], KNX (for home and building control – which is part of ISO/IEC 14543 HES standard) [KNX2009] or Zigbee (for low-power wireless device communication) [Zigbee] were deliberately omitted.

2.3 Management of home networks: an operator-centric perspective

As soon as Internet access became commonplace for home users, another role was added to the home LAN: sharing Internet access. Since it was economically unfeasible for most users to provide a dial-up line for each PC in the household, some users with a more advanced level of technical knowledge began configuring the PC which had the dial-up device attached (frequently an analog modem) to share its Internet connection with the rest of the LAN.

With the advent of broadband Internet access, dedicated dial-up devices designed for a conventional point-to-point connection paradigm were gradually replaced with affordable home routers providing routing and firewall capabilities (stateless at first and, later, capable of stateful packet inspection, as embedded hardware and software platforms evolved).

While home routers greatly simplified both the process of configuring a home/SOHO LAN and sharing an Internet connection, a wide range of issues remained. If the standalone LAN was frequently an unfriendly environment for the less technically-inclined users, the broadband connected LAN only contributed to aggravate the management and security concerns.

The expansion of residential broadband coverage, together with the introduction of IP-based services – such as VoIP and IPTV – raises concerns with the problem of managing the domestic LAN environment. The reliability and performance of IP-based services depend on proper configuration and operation of all the equipment in the service delivery path, including endpoints (e.g. set-top boxes) and residential routers. As most customers are not able nor willing to deal with the management of those devices, operators have no other choice than to remotely manage them (configuration management, monitoring, etc.) and the path between them and the access network (i.e. at least a segment of the customer LAN), in order to maintain adequate service levels.

In fact, and despite legitimate customer privacy concerns, the truth is this trend has already implicitly started – most triple play customers already have ISP-provided devices on their LAN with customized configurations and/or firmware which they cannot control. This creates the need for adequate management mechanisms capable of addressing the needs of service providers while safeguarding customer autonomy and privacy concerns

2.3.1 General-purpose management frameworks

Over time, the industry has proposed several management technologies in order to deal with the problem of remotely configuring, monitoring and troubleshooting devices across networks, with a variable degree of success. Among them, the ones next described stand out as being those that best represent the management ecosystem:

- **SNMP (Simple Network Management Protocol)** [RFC3410] is one of the most popular protocols, which is mainly used in LAN or MAN environments. It was defined by the IETF (Internet Engineering Task Force) and uses a data model based on the concept of a Managed Information Base (MIB) [RFC3410], using the encoding notation defined by ASN.1 [ITU2002a] and structured as a tree in which each node defines a unique naming subspace that can be extended. While several versions of the standard exist (version 3 is the most recent), SNMP suffered from chronic problems right from the start, related with data model (apart from standard MIBs, specific vendor extensions require the management application to be aware of them) and security (theoretically solved with SNMP v3, which was not widely adopted).

- **CMIP (Common Management Information Protocol**) is defined by ITU-T X.700 [ITU1997] and ISO [ISO9596] series recommendations. It is mainly used in telecommunication networks, and it's recognized by its complexity: CMIP over TCP/IP (CMOT) [RFC1189] is composed by 4 additional protocols whose description is spread among 7 different documents. CMIP used an object-oriented data model, in which each object encapsulates its own attributes, actions and notifications. While it had more features than SNMP, its complexity kept it from being successful.

- **JMX** [Sun2006] is a Java-based management technology that uses objects called MBeans (Managed Beans – which technically are probes) that represent a resource running on a Java Virtual Machine. Using connectors and adaptors (the first provide access to the MBeanServer API while the latter provide conversion for other protocols such as SNMP), applications can access the MBeans which are accessed through the MBeanServer API. The MBeanServer is the core of JMX, which mediates access between applications and MBeans. JMX found its way in some Java application servers where it is supported as a management mechanism.

- Also, **WSDM (Web Services Distributed Management)** [OASIS2006] is a web-service management standard issued by OASIS, originally in 2005. Its aim is to provide a web-service based management protocol for managing compliant devices, in some cases providing the same functionality as SNMP. WSDL is based on two specifications: MUWS (Management Using Web Services) [OASIS2005a] and MOWS (Management Of Web Services) [OASIS2005b] – the first provides the guidelines for representing and using management mechanisms as web-services plus a standardized event format, while the latter provides the guidelines for managing web-services as resources. Its latest version is 1.1, issued in September 2007 – it hasn't reached any significant expression and its usage is being deprecated in some products.

- **WBEM (Web-Based Enterprise Management)** [DMTF2011]: this designates a set of standards defined by the DMTF (Distributed Management Task Force) [DMTF] to address the management needs of distributed environments, such as CIM (Common Information Model – which defines the WBEM object-oriented data model) [DMTF2010], CIM-XML (a XML-based format for CIM data exchanges) [DMTF2009] and WS-Management (a SOAP-based web-services management standard) [DMTF2009a]. Several implementations exist for different operating systems such as Microsoft's WMI (Windows Management Instrumentation) [Microsoft2010b] that started to include WS-Management support in recent versions (the Distributed Component Object Model proprietary RPC mechanism was previously used). While WBEM is not specifically targeted towards any special device category, its implementation is too demanding for typical embedded devices – however, it became relatively popular for management of personal computers, applications, services and servers in enterprise environments.

- **UPnP:** some authors go to the extent of classifying UPnP as a management protocol for home environments [Royon2007] because its control and eventing mechanisms can be used for such purpose, even if the scope of UPnP is much broader. While UPnP's original design premises focused on automatic device interoperability, it recently got its own management extensions, documented in new DCPs for Device and Software Management [UPnPForum2010]. However, its adoption rate has been very slow and the majority of UPnP-compliant devices do not support them.

- **NETCONF** [RFC6241]**:** this is an IETF standard (originally published in 2006) that emerged as an alternative for both the aging SNMP, which was being progressively relegated to monitoring purposes, and for Command Line Interfaces (CLI), that have become popular among operators (which automated their use by using scripting languages such as Perl [Wall2000] or Expect [Libes1991]) for remote management of equipment. Some aspects of NETCONF are still in development – for instance, only recently (October 2010) the specification for its data model description language called YANG [RFC6020], based on XML, was published. NETCONF uses an RPC-based operation paradigm and supports several transports, such as Secure Shell (SSH, mandatory) [RFC4251], SOAP or TLS [RFC2246], using XML for device configuration data and protocol encoding. YANG allows data model extensibility (for vendor-specific parameters, for instance) and extends NETCONF with a new operation *(<get-schema>*) for retrieving data model schemas from a managed device, with versioning support. Also, the protocol itself can be extended with new capabilities (operations, in NETCONF terminology) which can be queried by a management client.

While this protocol diversity could suggest the opposite, the fact is that none of these frameworks seems adequate for device and service management in broadband access networks. Traditional management technologies such as SNMP or CMIP were deemed inappropriate, not only due to well know security and functional limitations [Neisse2004,CERT2008,Decius2010] but also because they do not cope well with specific requirements such as device-triggered management sessions and secure communication with managed devices behind NAT firewalls. Similarly, while other protocols are adequate for domain-specific usage, the same does not happen with these environments: even NETCONF, which is being developed by the IETF to target a broader management scope, has limitations in NAT environments, restricting their usage to manage CPE devices inside the customer LAN.

2.3.2 The CPE Wan Management (CWMP) protocol

The lack of adequate management and integration architectures for the connected home prompted the emergence of industry-led initiatives such as TEAHA (The European Application Home Alliance) [TEAHA], CENELEC Smart house [CENELEC2010], TAHI (The Application

Home Initiative) [TAHI], ITU-T DHF (specifically via the J.190 MediaHomeNet [ITU2002] recommendation), ISO/IEC HES [ISO14543-5], HGI [HGI] and the Broadband Forum [BBForum], with a variable degree of impact. In fact, HGI and the Broadband Forum seem like the most influential, having produced a number of technical standards and recommendations for remote management of customer premises equipment (CPE) that became widely accepted industry specifications.

Among those standards, Broadband Forum's CPE Wan Management Protocol suite (CWMP [TR-069]) emerged as a *de facto* standard, providing a framework for remote management of devices on the customers' premises. CWMP was originally conceived to remotely manage DSL modems, enabling secure auto-configuration, dynamic service provisioning, diagnostics, software/firmware management and status/performance monitoring of such devices. Later on the scope of CWMP was naturally extended to home gateways in general and, more recently, to all types of customer premises devices.

In fact, apart from being the most popular management protocol on broadband access networks, the CWMP protocol is currently endorsed by HGI, the Femto Forum [FemtoForum2008], the WiMAX Forum [WiMAXForum2009] (for WiMAX CPEs, albeit using a number of vendor-specific extensions, instead of a Broadband Forum ratified data model) and the DVB Forum [Deschanel2009][BBForum2007] (for DVB IPTV Set-top boxes) as its CPE management protocol.

The Broadband Forum (originally founded as the "ADSL Forum" and later renamed "DSL Forum", which finally merged with the IP/MPLS Forum to form the Broadband Forum) is a non-profit independent organization that aims to develop broadband network specifications. The forum now has approximately 200 members, including a considerable number of operators. The CWMP protocol is used for technical on-site assistance, remote diagnostics and configuration, while enabling the streamlining of automatic subscriber activation processes. It has grown up to the point of being used and supported by millions of devices. CWMP ensures full interoperability between vendors and providers, scalable, automatic management and provisioning.

The history of CWMP begins when the ADSL Forum faced the need to develop a protocol for CPE management. Initially, two approaches were considered for this purpose: the SNMP based approach of Cisco, or 2Wire's OGMP (Open Gateway Management Protocol) [Walko2002] protocol, which later became the basis for CWMP. The SNMP approach was discarded due to security concerns, the inexistence of standard MIBs for DSL CPEs management and scalability issues inherent to the proposed roles (the manager would connect to the agents instead of the agent connection to the master). As such, the decision was made to develop its own protocol, which in 2004 culminated with the publishing of the first release of the document which describes its operation: the Technical Report 069, also known as TR-069. TR-069 and related protocol documents are managed by Broadband Forum's BroadbandHome working group – there are other groups within the Broadband Forum that deal with operations and network management, core infrastructure (IP/MPLS) or physical technologies (metallic and fiber transmission).

The CWMP Protocol suite is formed by TR-069 and related technical documentation, such as TR-106 [TR-106], TR-181 [TR-181] and TR-157 [TR-157], which respectively describe the generic data model template, the TR-069 device data model and specific component objects for devices, they form the CWMP protocol suite. There are several other technical documents which cover subjects such as device model extensions, technology and interoperability requirements or protocol operation, among others (see Figure 2.12). In CWMP terminology, the management server is called the ACS (Auto-Configuration Server).

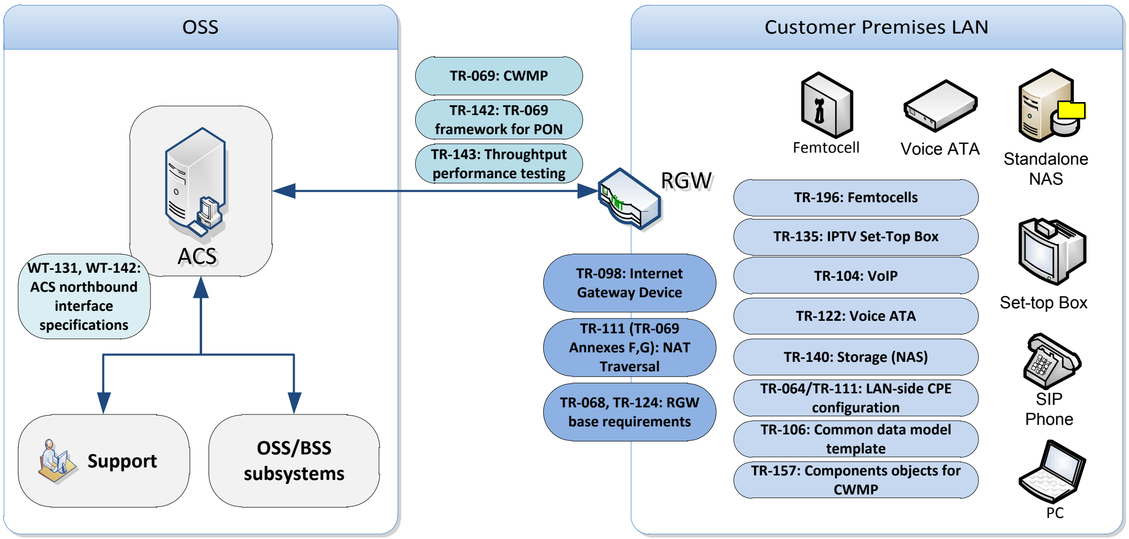
OSS
Customer Premises LAN
ACS
WT-131, WT-142: ACS northbound interface specifications
Support
OSS/BSS subsystems
TR-069: CWMP
TR-142: TR-069 framework for PON
TR-143: Throughtput performance testing
RGW
TR-098: Internet Gateway Device
TR-111 (TR-069 Annexes F,G): NAT Traversal
TR-068, TR-124: RGW base requirements
Femtocell
Voice ATA
Standalone NAS
TR-196: Femtocells
TR-135: IPTV Set-Top Box
TR-104: VoIP
TR-122: Voice ATA
TR-140: Storage (NAS)
TR-064/TR-111: LAN-side CPE configuration
TR-106: Common data model template
TR-157: Components objects for CWMP
Set-top Box
SIP Phone
PC

Figure 2.12: Broadband Forums' BroadbandHome remote management framework (adapted from [MR-239])

CWMP makes use of several known protocols and established standards, as indicated in Figure 2.13. CWMP enables service providers to manage CPEs in a complete and standardized way, whilst ensuring security by means of encrypted TCP/IP connections over SSL [Freier1996] or TLS [RFC2246] (TLS is preferred over SSL as of TR-069 Amendment 3 onwards). It is worth mentioning that the SOAP layer defined at TR-069 conforms to SOAP 1.1. However, it should be noted that TR-069 requires support for server and client role switching over persistent connections, which are unsupported SOAP 1.1 features and therefore require a custom implementation of SOAP to ensure compatibility.

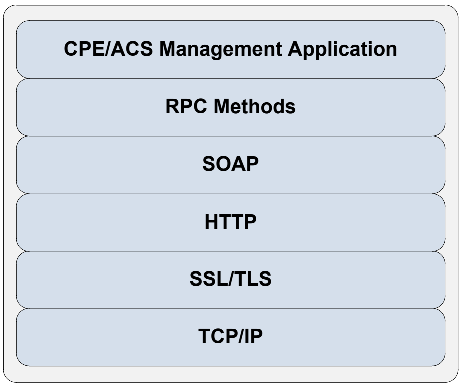

Figure 2.13: CWMP protocol stack (adapted from [TR-069])

CWMP operation makes uses of an RPC-based paradigm, in which both peers might reverse their roles invoking methods on each another, over the same persistent connection. Within this standard, some operations are performed in an unconventional way, such as SOAP bindings on HTTP, since the client/server roles can be inverted, with the server issuing "requests" (RPCs) to the client contained within a HTTP reply message. The fundamental reasoning behind this is that the CPE is in control of the communication process, since the entire session initiation depends on it – in CWMP, management sessions are always initiated by the CPE. This behavior is effective in dealing with the NAT traversal problem that hampers other management protocols, when used in an IPv4 network environment. The CWMP operation model still makes sense in IPv6 networks because even though IPv6 could in principle solve the the NAT problem, direct external access to devices in the customer LAN premises still cannot be taken for granted because of security measures (like firewalls).

There is an option for the ACS to request CPE to initiate connections, for asynchronous operations (such as real-time reconfigurations). This process requires the ACS to be able to reach the CPE to make the request, therefore being prone to interference from NAT mechanisms – for this purpose the TR-069 specification document includes Annexes (Annex F – Device-gateway Association and Annex G – Connection Request via NAT Gateway, also replicated on TR-111, parts I and II [TR-111]) that describe the methods used to deal with this situation.

In CWMP, the first thing the CPE does upon boot (and also subsequently, at periodic intervals) is to issue an Inform message to the ACS. As a response to this Inform message, which is also issued at the beginning of each CWMP session, the ACS can instruct the CPE to perform operations, change values or retrieve information.

Figure 2.14 presents an example of a CWMP management session. The CPE starts by establishing a connection to the ACS, optionally adding a security layer to the communication channel (it does that by issuing a HTTP Post containing the Inform request). Within this Inform message the CPE identifies itself and provides the ACS some additional information, such as serial number, uptime, etc. Figure 2.14 exemplifies a scenario of request without authentication: the received Inform is discarded by the ACS that demands the CPE to authenticate itself (in this case using basic authentication).

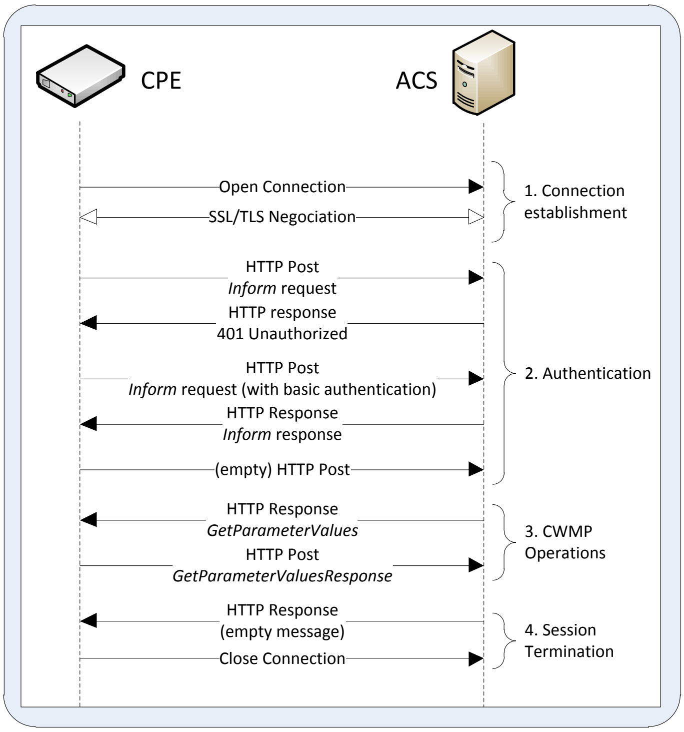

Figure 2.14: Example CWMP/TR-069 Session

Once authenticated, the CPE is cleared for access. The ACS responds to the HTTP POST with an *Inform* response message (in fact, an RPC invocation), letting the CPE know that the CWMP session was established successfully, to which it must respond with an empty message. After this the ACS is free to perform any RPC method invocation on the CPE, such as retrieval of stored parameters via the *GetParameterValues* RPC method (receiving a *GetParameterValues* Response message type, analogous to what was done with the Inform). The session is terminated when the ACS replies with an empty message, indicating that there are no pending operations to be performed.

The standard data model for a CWMP-capable device (see example on Figure 2.15) follows a common set of requirements for which the detailed structure, hierarchically organized like a directory tree, depends on the nature of the device. Data model information is structured using objects and parameters – the first are containers that group the latter, which are variables corresponding to specific configuration attributes. Some objects can be added and removed by the ACS (according to their nature).

In CWMP, a device always has a single root object. In the case of a home gateway, for instance, the root object must be *InternetGatewayDevice* (according to the data model specified in TR-098 [TR-098]). The root object usually contains three types of sub-elements: Common Objects as defined by TR-106 and TR-181; specialized Components from specifications such as TR-143 (throughput performance tests and statistical monitoring [TR-143]) and TR-157 (assorted management functions); and a "Services" object that contains all Service Objects associated with specific services (it should be noted that a single device might include more than one Service Object), either embedded on the managed device or located in other devices for which it provides proxied management services (therefore embedding their data model trees under the Services object of the proxy device). These features will be further discussed on Chapter 3.

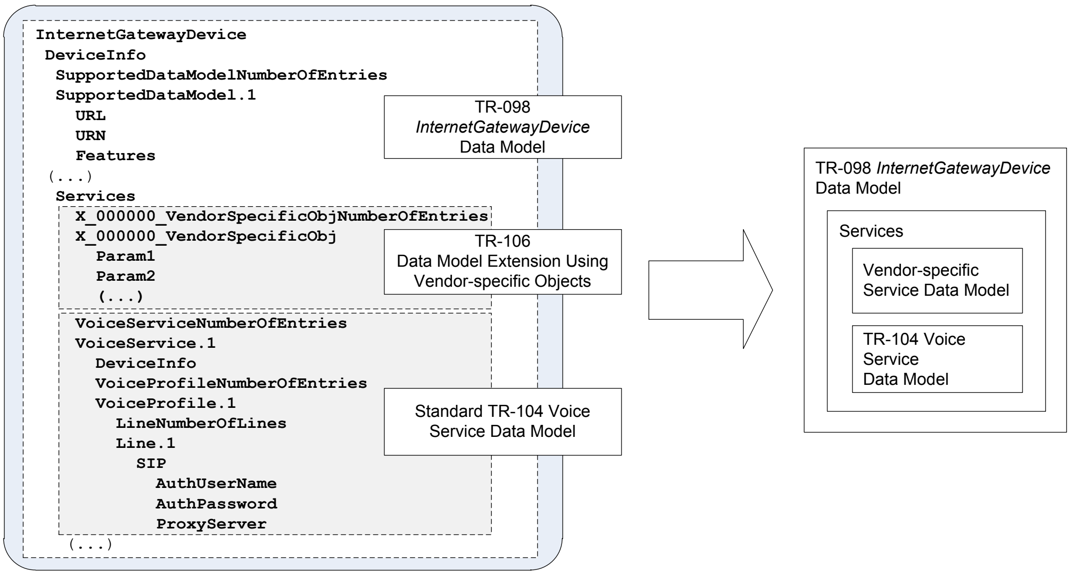
InternetGatewayDevice
DeviceInfo
SupportedDataModelNumberOfEntries
SupportedDataModel.1
URL
URN
Features
(...)
Services
X_000000_VendorSpecificObjNumberOfEntries
X_000000_VendorSpecificObj
Param1
Param2
(...)
VoiceServiceNumberOfEntries
VoiceService.1
DeviceInfo
VoiceProfileNumberOfEntries
VoiceProfile.1
LineNumberOfLines
Line.1
SIP
AuthUserName
AuthPassword
ProxyServer
(...)
TR-098 InternetGatewayDevice Data Model
TR-106 Data Model Extension Using Vendor-specific Objects
Standard TR-104 Voice Service Data Model
TR-098 InternetGatewayDevice Data Model
Services
Vendor-specific Service Data Model
TR-104 Voice Service Data Model

Figure 2.15: CWMP data model structure example

To avoid the same kind of problems that hampered SNMP (proprietary MIBs that do not follow a clear and understandable structure, making it very difficult to develop an overall solution for management of devices), the Broadband Forum has defined specific Technical Reports explicitly describing the data model structure for certain device classes, such as the TR-104 [TR-104] defining the structure used for VoIP services or TR-135 [TR-135], for set-top-boxes.

CWMP is also extensible, using two different mechanisms, in order to accommodate new functionality and services on managed devices:

- TR-106 specifies the ability of using vendor-specific parameters/objects (which extend standard data model components documented in other TR documents) for a given device, using the notation *X_<Vendor>_VendorSpecificName* for the identifier. The ACS is able to get knowledge about the supported data model of the managed device using the *SupportedDataModel* table with contains their XML description.
- Vendor-specific RPC methods can also be created, using the notation *X_<Vendor>_VendorSpecificName*. The ACS is able to get a list of supported RPC methods using the *GetRPCMethods* RPC call.

CWMP was designed to ease Operation, Administration, Maintenance and Provisioning (OAM&P) of devices on the customer premises in broadband access network environments, providing the means to enable *plug & play* device and service deployment – to a certain extent.

CWMP makes use of standard web technologies to provide dynamic service support, together with granular control, independent of the network or devices. It also encompasses extensibility mechanisms, allowing for vendor differentiation as well as spurring new types of devices and services, while providing adequate scalability and security. Also, the CWMP framework is constantly being updated, with a wide range of *technical work in progress* documents, such as PDs (Proposed Drafts) or WT (Working Texts) in line for publishing.

2.4 The role of home gateways in a multi-domain environment

Residential gateways play an important role in the scope of broadband environments. Apart from being responsible for mediating traffic and providing security, residential gateways are becoming bridgeheads for a multitude of purposes, from management to services. This is not happening by pure chance, as home gateways are placed in a strategic location, in the frontier between the access network (facing its northbound interface) and the residential LAN domains (facing its southbound interface).

In fact, as suggested by [Royon2007], it may be more than a simple question of physical location, since the home gateway is located at the intersection of three fundamental realms: the service realm, the connectivity/management realm and the LAN realm (see Figure 2.16).

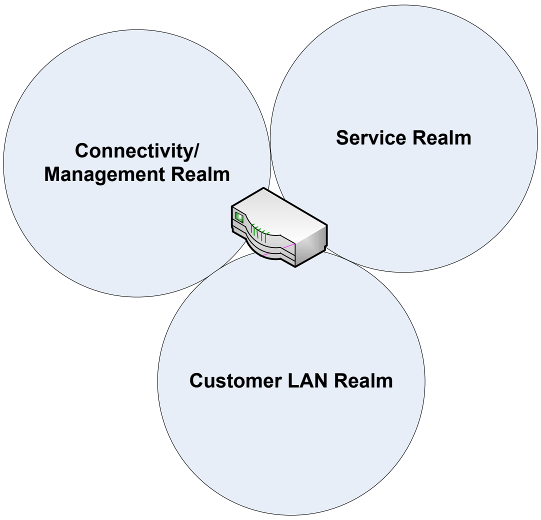

Figure 2.16: The residential gateway in a multi-domain environment (adapted from [Royon2007, Smedt 2006])

While the purpose and function of each realm is clearly defined, the same does not always happen to their constituting entities (see Figure 2.17).

As broadband service paradigms transitioned from the data-based, single-service model to the consolidated *n-play* service model, telecommunication operators were always wary of protecting their prevalence as the single service provider. This mindset is a legacy of the dial-up ISPs era, where telecommunication operators took advantage of the fact that the ownership of the access network infrastructure to gain a competitive lead over conventional ISPs, by providing their own ISP service. Also, due to years of regulatory limitations and lack of adequate policy management mechanisms, telecommunication operators had to endure the burden of constantly updating their infrastructure without getting a fair share of the revenue of the services that made money using it, a situation which also contributed to their unwillingness.

Even with the emergence in some countries of the multi-provider, multi-service paradigm, telecommunication operators still resist the idea of allowing third-parties to provide services using their infrastructure. We are presently in a transition stage, which is evolving at different paces: while in Europe the most common scenario is still based on single-sourced, operator-provided services (mainly data, voice and television), in some other countries (like USA), third-party services (like voice, TV, multimedia content distribution, among others) are beginning to emerge.

In this perspective, it should be mentioned that the broadband management framework that constitutes the subject of this document supports a management model which is able to deal with both single-sourced and multiple-provider scenarios. While its fundamental premise is based on the concept of letting the telecommunications operator in charge of infrastructure management, this does not restrain third-parties from shared access to resources and management mechanisms (limited accordingly with established agreements and policies).

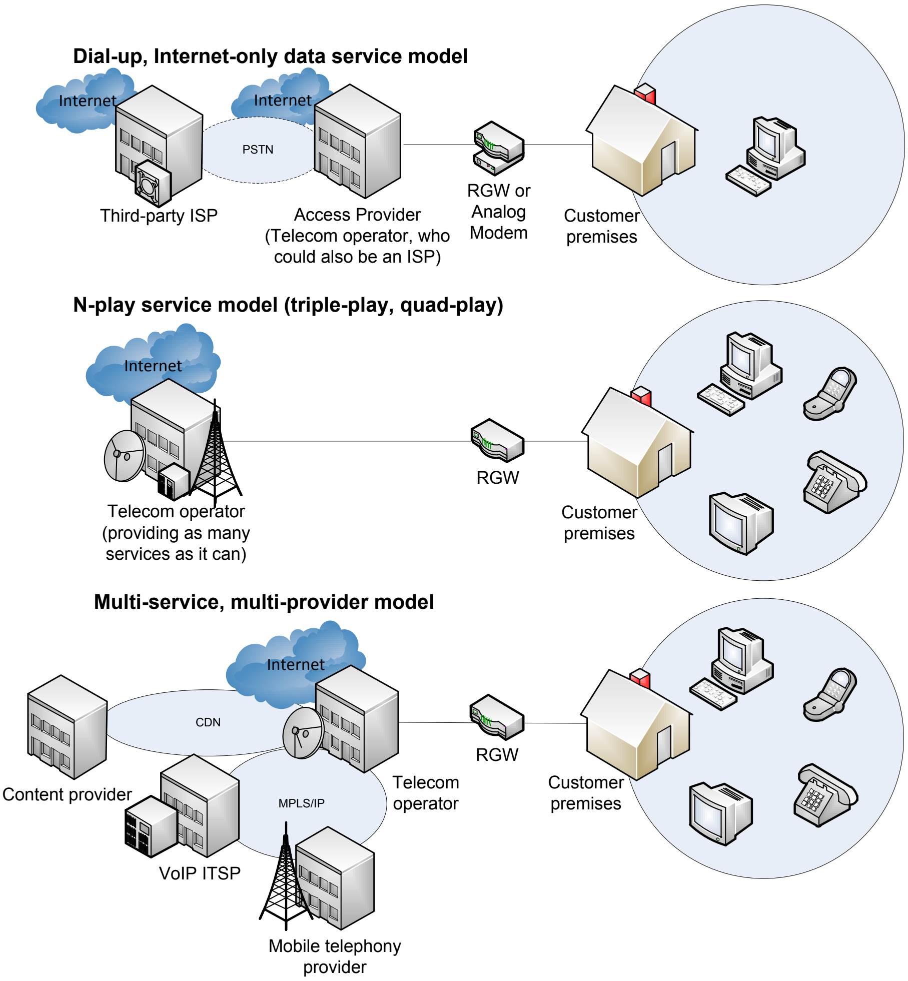
Dial-up, Internet-only data service model
Internet
Internet
PSTN
Third-party ISP
Access Provider (Telecom operator, who could also be an ISP)
RGW or Analog Modem
Customer premises
N-play service model (triple-play, quad-play)
Internet
Telecom operator (providing as many services as it can)
RGW
Customer premises
Multi-service, multi-provider model
Internet
CDN
Content provider
MPLS/IP
VoIP ITSP
Telecom operator
RGW
Customer premises
Mobile telephony provider

Figure 2.17: The evolution of service paradigms in broadband environments

The broadband service paradigm is expected to eventually evolve to a multiple service ecosystem where multiple providers coexist together with the telecommunications operator that provides connectivity and infrastructure management, together with its own value-added services. In this perspective, the residential gateway role is evolving from providing simple connectivity services (in the data-only service model) to a more sophisticated device fulfilling several different roles, actively involved in delivering converged *n-play* services and even multiple services from different providers, such as content distribution, or femtocell-based indoor mobile telephony (the multi-service, multi-provider model). This service model shift is proving to be disruptive enough to cause a profound change in services and its perception – and it’s not only about devices, because the end-user itself has started to produce its own content, instead of being a simple consumer.

As the market for these services is rapidly evolving and changing, the domestic/residential gateway, which was previously seen by manufacturers and operators as expendable equipment whose replacement was preferable to any upgrade, is becoming an increasingly capable device with considerable computing resources in terms of processing power and memory. This capability leap experienced by the embedded systems that constitute the basic residential gateway hardware has enabled the creation of managed execution environments allowing them to incorporate several different services and abilities in the form of software modules which can be added, removed or upgraded with minimum interference.

2.4.1 Managed execution environments

As residential gateways became able to host and execute specialized software components to add functionality and services to the gateway itself, arose the need for the development of execution environment frameworks for components, with inventory and lifecycle management, security and contention mechanisms.

Historically, the precursor of these frameworks was the **Java ME** (Micro Edition – previously J2ME) [Oracle2010] framework, which was specifically designed for embedded devices such as PDAs and intelligent appliances (were Jini played an important role into providing a plug-and-play distributed computing environment for devices). While, at certain point, it had become a dominant technology in the mobile device market, it has been meanwhile superseded by newer platforms such as Android (Dalvik VM) [Google2011] or iPhone's iOS [Apple2011b], among others.

Java ME suffered from many problems, such as fragmentation (until 2006 there was no freely available binary reference implementation), under exploitation of device capabilities, lack of dependency management and asymmetry in relation to the standard Java framework. These limitations made Java ME an effective platform for mobile devices managed by a single entity providing components and services (frequently the telecommunications operator, in the case of mobile phones).

To address these questions, the Open Services Gateway Initiative [OSGi] (nowadays called OSGi Alliance, since OSGi has moved beyond gateways) was formed to propose a service-oriented architecture, enabling a multi-service environment using software components from several providers.

OSGi runs on top of a Java Virtual Machine (JVM) with the CDC (Connected Device Configuration – for high-end consumer devices) [Sun2006a] or CLDC profiles (Connected Limited Device Configuration – for low-end consumer devices) [Sun2007], turning the device into an Operating System (OS) agnostic application server, which includes a servlet container. It enables bundles (pluggable components) to be delivered and installed on the fly, with lifecycle and permission management, without rebooting the device. It also provides a shared event bus and has monitoring and logging capabilities. Once a component is wrapped as an OSGi bundle, it can be installed in the device and registered as a service.

OSGi covers three areas (see Figure 2.18): bundles, lifecycle management, and services. The bundle is a JAR file (Java ARchive) embedding java classes, native code and metadata. Metadata (the manifest file) describe the contents of the JAR and provide information about the bundle, such as dependencies, which packages are exported to use by other applications or which is the activator class. In simple terms, a bundle is an encapsulation entity that can be installed and uninstalled.

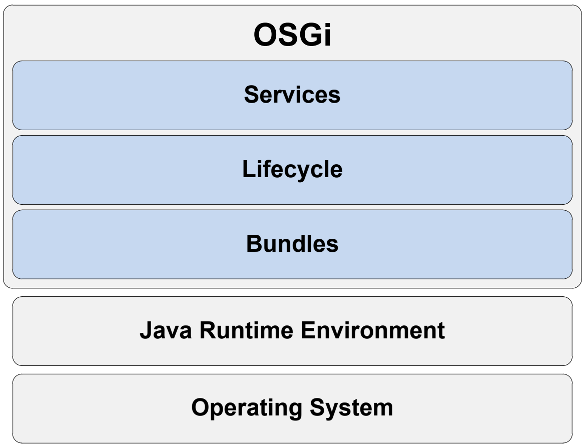

Figure 2.18: OSGi layers

Is up to the lifecycle layer to define the sequence of steps that bundles go through when installed, updated, uninstalled, starred and stopped (see Figure 2.19). Each bundle must implement *start()* and *stop()* endpoints methods for the lifecycle manager. Also, the lifecycle layer only allows a bundle to be used when all external dependencies are resolved. Runtime bundle updating is possible, albeit the operation implies that it must be previously stopped, together with its dependencies.

The services layer provides communication between bundles (see Figure 2.20) – it has a registry were all services are registered and where bundles can look for available services. In the OSGi service interaction model, service providers publish service descriptions in the service registry while service requesters discover services and bind to the respective providers.

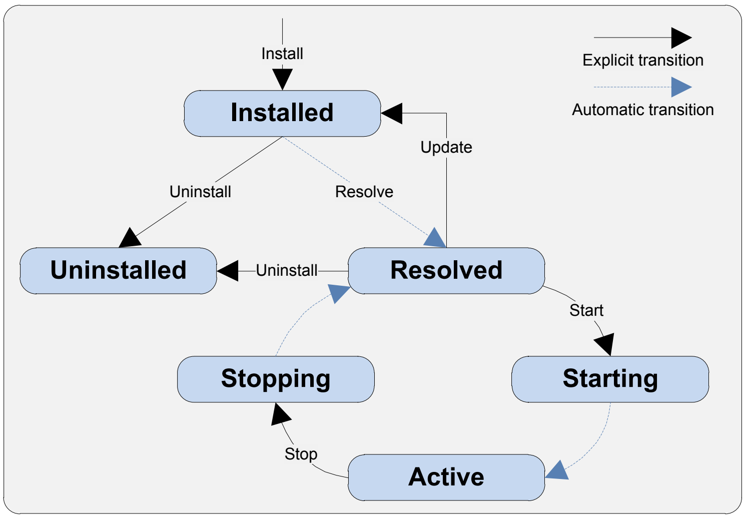

Figure 2.19: OSGi bundle lifecycle

In OSGi, a service is described as a Java class or interface (the service interface), together with their attributes (service properties, which allow to distinguish different service providers with the same exposed interface). The service registry allows service providers to be discovered using LDAP-syntax queries. The OSGi service layer also allows service implementations to change at runtime, using dynamic updating mechanisms.

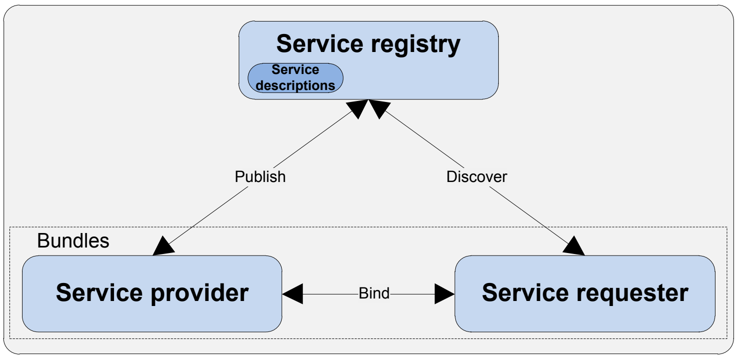

Figure 2.20: OSGi service interaction model

OSGi also provides bundle isolation, through the use of individual class loaders. This restricts bundle class visibility only to itself (even if the context of the same JVM), with independent namespaces.

While the Java Community tried to incorporate OSGi-like features for Java SE in the form of JSR 277 [Sun2006b], the specification was more limited, in comparison with OSGI. Eventually, OSGi R4 finally got in the Java Community process, being adopted as JSR 232 [Sun2006c] (for Java ME) and JSR 291 [Sun2007a] (for Java SE).

Microsoft's .NET CLR (Common Language Runtime) [Hamilton2003] it is a managed code environment similar to OSGi in terms of component management. While the full-fledged .NET framework is inadequate for embedded systems, Microsoft released two variants specifically targeted towards smaller-scale devices: the .NET Compact [Microsoft2011] and Micro Frameworks [Microsoft2011a]. The first was specifically designed for mobile applications (PDAs, smartphones) and embedded applications such as set-top boxes, requiring an OS (Windows CE, PocketPC and Mobile OS families), while the latter targets hardware platforms which are much more resource-constrained, being capable of running the CLR on the bare hardware. Both incarnations are not adequate for multi-service home gateways because of architectural restrictions, which make them better suited for windows-based device ecosystems with single service provider and management entities. And while some efforts have attempted to provide OSGi-like features for the.Net Framework and its subsets [Ntuba2004, Physalis, Escoffier2005, Wegner2008], they haven't been standardized nor succeeded into mainstream acceptation.

It is also interesting to note that Microsoft has developed OSGi-type features for the full .Net framework, in the form of Managed Extensibility Framework [Microsoft2010c]. However, they are not relevant in the scope of this discussing since the full .Net framework does not target embedded appliances and devices such as residential gateways or set-top boxes.

There are also **other execution environments**, apart from Java ME, OSGi or .NET CLR, like native Linux (which is used by several device manufacturers, but whose management mechanisms depend on specific implementations), CableLabs's OCAP (OpenCable Application Platform – specific for STBs in cable networks) [CableLabs2011], DVB-MHP (Multimedia Home Platform, Java-based middleware designed for interactive Set-Top boxes for DVB) [DVB2011], Ginga (a standard for the Japanese/Brazilian Digital Terrestrial TV system for interactive IPTV Set-top boxes) [ABNT2011] or Android's Dalvik Java VM. Even HGI has its own execution environment architecture [HGI2011], albeit very close to OSGi both in conceptual and functional aspects, (which is not surprising as it was defined in cooperation with the OSGi consortium). Nevertheless, OSGi remains one of the most representative and popular execution environment frameworks for residential gateways, also being considered universal middleware, with a usage context that spawns several different applications (including automotive and industrial usage, among others).

2.4.2 CWMP component management (TR-157)

While a managed execution environment (like OSGi) generally provides a middleware to deal with component and service management on top of an execution environment (like Java or Linux), its functions are orthogonal to remote management protocols. For instance several mobile platforms support OMA-DM [OMA2008] management for component download and update, which are, by their turn, locally managed by an OSGi middleware.

CWMP has also gone to the extent of standardizing multi-platform component management in its scope of operation, through the TR-157 component management specification. Appendix II of TR-157 (which was formerly specified by PD-194) [TR-157], deals with Software Module Management. It targets Linux, OSGi, .NET, Android, and Java ME execution environments, enabling component lifecycle management based on abstract models which are not tied to a specific platform and avoiding the need to develop tailored mechanisms for specific execution environments (EEs).

TR-157 recognizes the following entities (Figure 2.21):

- **Execution Environments (EEs),** the platforms that support dynamic loading, unloading and execution of software modules. TR-157 considers that EEs might be layered (one EE runs on top of other EE, for which it is the primary environment). As such, it is possible to support different concurrent execution environments on the same device (like OSGi and Linux-based EEs, and even several OSGi EEs on the same device).
- **Deployment Units (DUs):** are the entities that can be deployed in an EE. In encapsulates a set of resources such as functional EUs, configuration files and other resources. This is the fundamental entity that can be installed, updated or uninstalled. A DU might contain an EU, but it is not mandatory – also, an encapsulated EU cannot span across EEs (for the sake of isolation).
- **Execution Units (EUs)** are executable entities deployed by a DU, – scripts, services, software components fall into this category. An EU spawns processes to perform tasks or provide services, being an entity that can be started or stopped. EUs also expose configuration for implemented services, either via CWMP (TR-069 data model) or specific mechanisms.

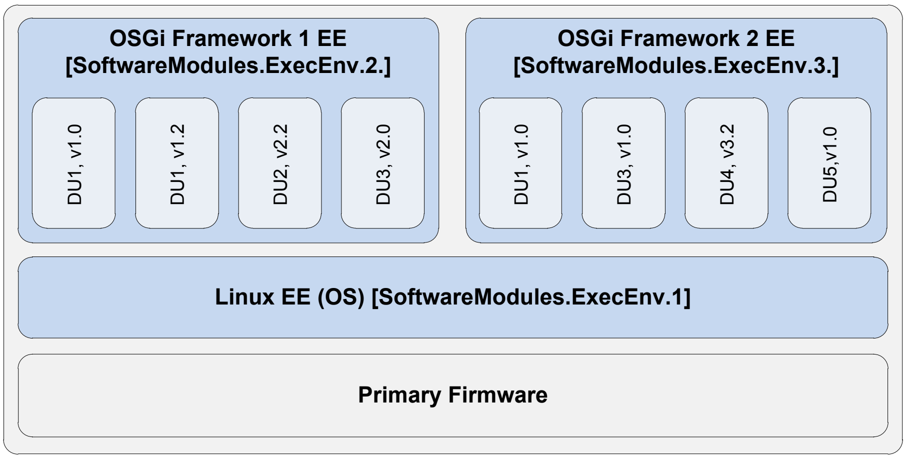

Figure 2.21: Managed entities in TR-157

Both the component and execution lifecycles follow a pattern along the lines of generic high-level processes which are common to all platforms (see Figure 2.22).

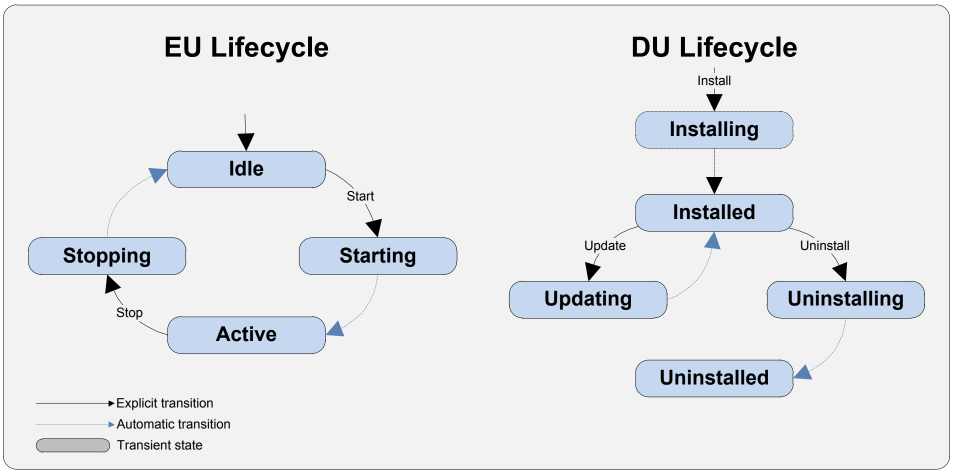

Figure 2.22: TR-157 Deployment and Execution Unit lifecycle

TR-157 does not deal with component dependencies, which must be resolved by the EE-specific management mechanisms prior to installing a DU – once again this has to do with ensuring independency from specific EE implementations. On the other hand, one of the benefits of TR-157 is that it can ease component management for EEs lacking a standardized module management facility, such as Linux – all the device manufacturer has to do is to add dependency and related fault management on top of the already provided mechanisms.

Also, TR-157 plays a critical part in future Broadband Forum plans: as of 2010, the co-chair of the BroadbandHome working group unveiled a vision where TR-157 component management plays an important part into enabling third-party software modules for CWMP-compliant CPEs (and not only residential gateways) using a distribution model akin to an application store, albeit for services and functionality [Kirsey2010].

Most of the solutions and prototypes devised and implemented in the scope of this research work were conceived for a Linux EE (albeit being based on portable Java and Python code, which can be run in almost any environment), relying on TR-157 mechanisms for component deployment and management in the prototype residential gateway devices used for testing.

2.5 Emerging trends: IPv6 and virtualized gateways

While IPv4 still dominates the majority of broadband access network environments, most operators are planning the move to IPv6. IPv6 brings new challenges both because of the need of infrastructure/protocol compliance and interoperability between IPv6 and IPv4. The IETF has already provided general guidelines [RFC 4779] and the Broadband Forum is also readying its guidelines for this process, some of which were already published as part of BroadbandSuite 4.0 (v4.1 is the most recent, at the time of this writing) [BBForum2011a] – however, further work is ongoing within the forum, such as WT-242 (*IPv6 Transition Mechanisms for Broadband Networks*) and IL-001 (*IPv6 Issue List*) [BBForum2011]. While the transition to IPv6 might have a mid-to-short term architectural impact (IPv4 NAT will probably disappear with time, albeit several customer LANs will remain IPv4-based for long), its impact does not significantly affect the management models and service paradigms addressed by our research work, which remain pertinent regardless of the IP protocol version.

Also, it is expected that virtualization will gain increased relevance in the context of broadband access networks, playing different roles at different levels of operation. Virtualization technologies, in their broader sense, were already being used in embedded devices for some time to provide hardware independence (JVM, .NET CLR) and/or component and namespace isolation (for security purposes). One application of this concept used virtualization to enable service co-location within a same gateway, with each service being associated to a specific provider [Royon2006] and using a JMX management framework on top of OSGi to enable remote management – in practice this provides support for multi provider remote gateway management.

While the use of virtualization techniques (full instruction-set, paravirtualized, user-space isolation, containers or application-level) was, until now, being confined to the device itself, a new trend is emerging towards its entire virtualization. In an effort to reduce deployment costs and also to take advantage of recent Fiber to the premises (Fiber-to-the-home/FTTH and Fiber-to-the-building/FTTB) deployments, operators are considering the possibility of reshaping the boundaries between the customer premises LAN and the access network, by moving the residential gateway (and most of the functionality of some devices, such as set-top boxes) to their data centers. In this scenario, the access network is bridged to the domestic LAN segment, with the virtualized gateway residing at the operator data center. This has a series of implications which are still being discussed by telecommunication operators and service providers, in order to define a reference architecture for its implementation (Figure 2.23).

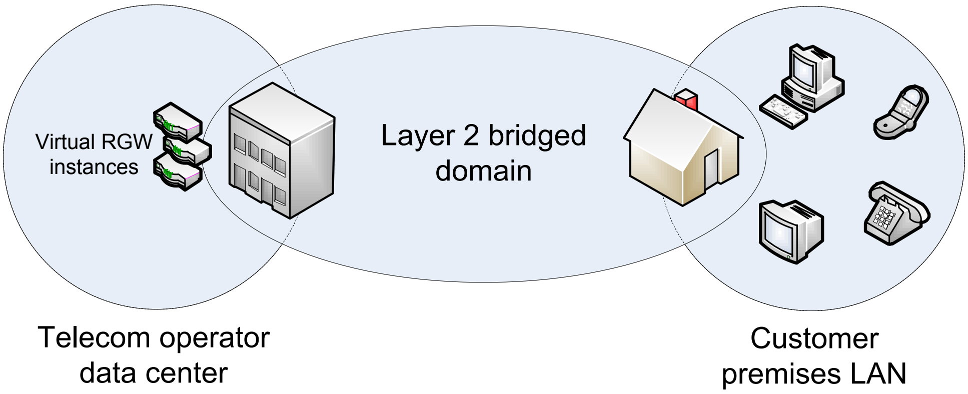

Figure 2.23: Virtualized RGW scenario

For instance, once the service provider infrastructure becomes bridged at layer 2 with the customer premises on Ethernet based PON (Passive Optical Networks) it is possible to use 802.11q VLANs to create per-subscriber home network domains which are able to spawn several households. In fact, with TR-101 [TR-101] the Broadband Forum had already been hinting at a similar possibility, albeit originally in the scope of a proposal to migrate ATM to Ethernet aggregation at the access nodes. Among several aspects, TR-101 already supported N:1 (1 VLAN per service, adequate for multicast service delivery scenarios such as IPTV) and 1:1 (1 VLAN or per subscriber) VLAN topologies at the aggregation level, in the scope of a multi-service offer (as originally specified by TR-058 [TR-058] and later updated by TR-144 [TR-144]). However, only with Ethernet PON technologies (like GPON) it became possible to bring full Ethernet connectivity right up to the RGW – TR-156 [TR-156] and TR-167 [TR-167] update TR-101 (which was a xDSL-centric specification) for these environments. Bridged or routed RGWs are supported in those scenarios as specified by TR-124 [TR-124], with dual-stack IPv6 support being added by TR-177 [TR-177].

Potentially, services such as CIFS/SMB file-sharing or UPnP services might be provided by the operator, as the subscriber LAN segment is extended to its infrastructure premises (Figure 2.24).

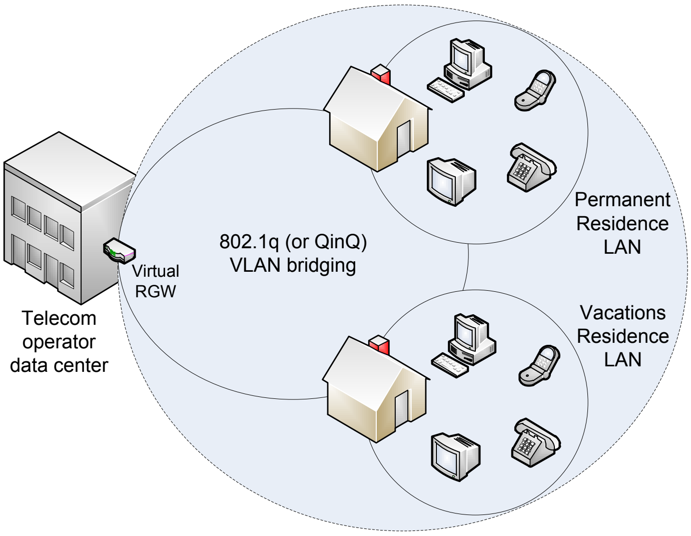

Figure 2.24: Single subscriber layer 2 domain in virtualized RGW scenario

At EURESCOM [EURESCOM], there is an undergoing study, launched in 2010 (EURESCOM P2055 – Virtual the CPE) [EURESCOM2011] with the participation of several European telecommunications operators and device manufacturers. While it is still in its early stages, available documentation shows that the main concerns are related to the process of defining a scalable architecture and decoupling RGW functionality. In this process, some of the components existing in current operator infrastructures might simply disappear, as it is the case with the BRAS (Broadband Remote Access Server), which has no place in a virtualized infrastructure.

Ultimately, functions such as L3 routing, NAPT (Network Address and Port Translation), ALG (Application Level Gateway) lightweight proxy services (i.e., for SIP VoIP signaling) and LAN DHCP might move to the operator infrastructure Virtual RGW pool, leaving behind a simple L2 bridge at the customer premises with simple shaping and Wi-Fi bridging capabilities, possibly embedded on the OLT (Optical Line Termination) equipment in fiber deployments. For advanced users, it is expected that mediation interfaces become available, providing an HTTP web user interface for virtual RGW configuration, akin to the ones that exist for standalone RGW.

In this scenario, several devices might also become virtualized, such as set-top boxes – moving all the interface and service logic to the operator cloud while leaving a simple appliance with basic display, networking and codec processing capabilities. Nevertheless, CWMP will maintain its relevance in virtualized infrastructure scenarios, since there will always be devices on the customer premises which will need to be managed. In fact, the virtualized broadband infrastructure might even favorably impact CWMP operation, solving the IPv4 NAT issues addressed by TR-111 and TR-069 Annexes F and G, therefore rendering them unnecessary since NAPT and ALG protocol helper functions are dealt with at the operator level.

As for services, a virtualized infrastructure does not remove the need for them. For legal (user and operator realm separation) and technical reasons execution environments will probably continue to exist, either embedded in virtual RGW instances or consolidated in separated service support infrastructures.

2.6 Conclusion

This chapter was written with two purposes in mind: to constitute a *state-of-the-art* section and to introduce the main topics which relate to the subjects addressed by this thesis. This approach has to do with the fact that the following chapters address a considerable wealth of disparate subjects, making them difficult for readers to grasp without an introductory discussion about the technologies and concepts involved. As such, the organization of this chapter reflects these concerns, expressed by its structure:

- First, it started with a discussion on the subject of broadband network paradigms, providing an an historical perspective on the evolution of the access network, related usage paradigms and recent trends.
- Section 2.2 was devoted to the customer premises LAN. Its purpose was to introduce a set of topics related to the interoperability frameworks and management standards for devices and services on residential LANs which will be revisited in Chapters 4 and 5.
- Section 2.3 provided an operator-centric perspective on the subject of home network management, starting with a brief discussion of general-purpose management frameworks, explaining how their limitations and handicaps make them inadequate for broadband access networks – Chapter 6, which has a section about protocol proxying mechanisms, will delve again on this subject. Later on, this section introduces the CPE WAN Management Protocol (CWMP), which constitutes the basis for the management framework presented in Chapter 3.
- Section 2.4 discussed the role of home gateways in the scope of a multi-domain environment involving the realms of connectivity, services and customer LAN. Its purpose is to explain how modern gateways can play a role in the context of the residential LAN that goes beyond providing connectivity, supporting innovative services and functionalities enabled by managed execution environments.
- Finally, section 2.5 presents an overview of emerging technology trends which can potentially shape the broadband access network scenario in the near future.

Part II:

A management framework for broadband access networks

3. A management framework for broadband access networks

This chapter presents the proposed management framework for broadband environments which constitutes the fundamental building block for the research work hereafter developed. This management framework is built on top of Broadband Forum's CPE Wan Management Protocol (CWMP/TR-069), the *de facto* industry standard for the management of devices that, despite located in the local network of domestic broadband users, still need to be directly managed by operators due to their critical relevance to added value services such as VoIP, IPTV, VoD or femtocell-based applications.

This management framework is designed to support an operator-centric device and service management model, addressing some of the shortcomings of the CWMP protocol while simultaneously providing an interoperability model for integration with LAN services, management protocols and non-CWMP compliant devices. This is possible by incorporating support for a generic agent extensibility mechanism which, whilst keeping full compatibility with the original CWMP specification, allows it to better adapt to the home network environment.

This chapter is structured as follows: Section 3.1 discusses the motivation and reasoning behind the management framework hereby proposed. Section 3.2 deals with the limitations and problems of using CWMP in current broadband environments. Section 3.3 presents the generic agent extensibility framework which is the basis of the proposed management architecture and section 3.4 discusses its potential application scenarios. Section 3.5 analyzes related work and section 3.6 concludes this chapter.

3.1 The case for a broad-scope management framework

With the current generation of service bundles and offer diversity, operators are becoming increasingly aware of the need for being able to manage a device and service scope as broader as possible. Faced with the impossibility of remaining confined to own infrastructure, operators have no other choice than to deal with the problem of management in broadband access networks to the maximum extent of its reach: i.e., the domestic LAN (Figure 3.1).

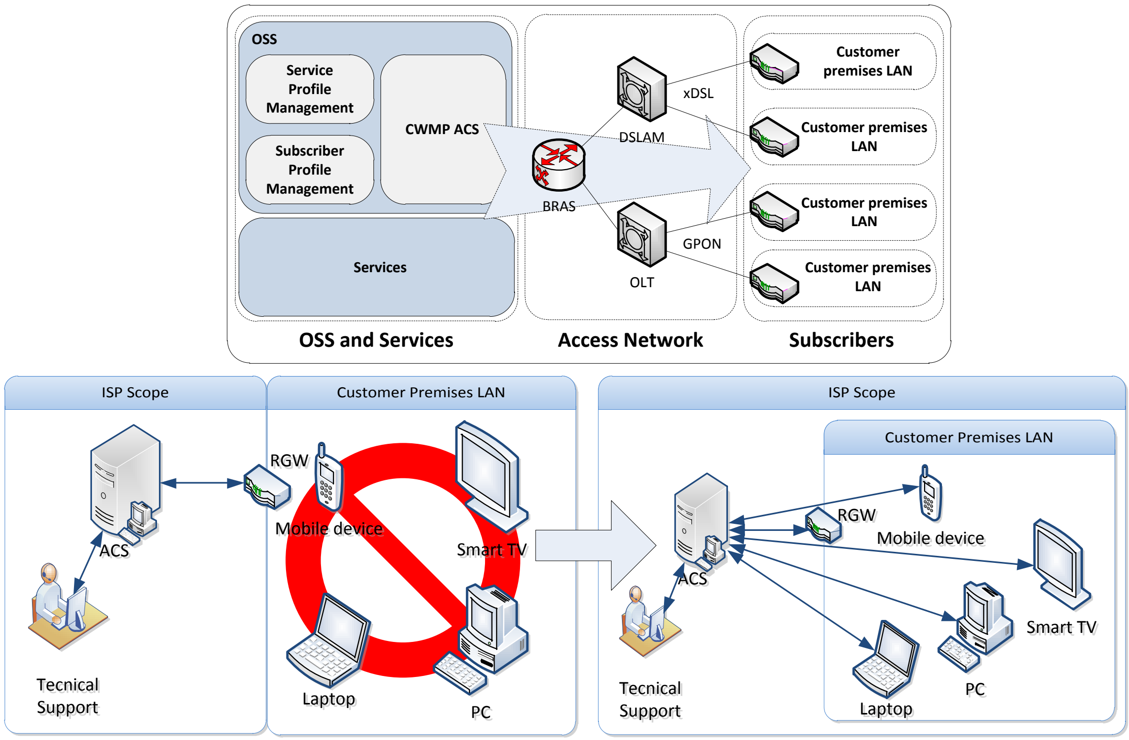

Figure 3.1: Management scope expansion towards the customer premises LAN

However, this poses several problems, among which the most important has to do with the inexistence of an integrated management framework capable of interoperating with the existing LAN protocol ecosystem (as discussed on Chapter 2), while being flexible enough to accommodate new services and applications.

Conceptually, CWMP is an obvious candidate to fit this purpose, since it was designed for CPE and service management on broadband access network environments, having become the *de facto* standard for these applications. However, it has some handicaps which severely limit its usability and prevent its use as the basis for a complete and wide-scope solution for management of devices and services on broadband environments, as its reach is not as wide as it would be desired – for instance, the effective adoption rate of CWMP is still slow for some device categories such as SIP telephones or UPnP/DLNA-compliant media player devices. In our opinion this results from multiple reasons:

- An understandable inertia, from vendors and operators, in the shift from proprietary platforms to CWMP solutions.
- The lack of generic integration mechanisms for both legacy devices and key LAN technologies like WMI and UPnP.
- The difficulties in the management of devices inside the domestic LAN and behind NAT Firewalls. While the protocol was designed with those scenarios in mind (for this reason the CPE always initiates a CWMP connection), it still poses some difficulties for situations where the ACS has to issue a connection request to perform an asynchronous operation. The CWMP suite does try to address this scenario, through TR-111 and TR-069 Annexes F and

G, but the proposed solution, which uses STUN [RFC5389] and registered device associations on the RGW, is cumbersome and fail-prone.

- The fact that, despite the increasing number of CWMP compliant devices available on the market, CWMP is still too heavy for implementation in the simpler CPEs.

Should these limitations be overcome, it would be possible to realize the full potential of CWMP to provide a complete management solution. Rather than starting from scratch, the framework proposed on this chapter chooses to improve the operation of CWMP in a completely transparent way, without breaking the operation of the protocol, and providing a solution which is compatible to the already existing CWMP management solution deployed on the operator infrastructure.

The management framework hereby presented leverages the potential of the CWMP protocol while addressing these shortcomings by providing a generic extensibility mechanism which improves the functionality and usability of CWMP-based solutions. The key advantage of this framework is the possibility of modular agent development, clearly separating the CWMP-specific modules from the set of management functionalities associated with each managed service within the CPE – thus simplifying the development of CWMP interfaces for new or already existing services. In addition, this framework also allows for dynamic addition or removal of management services during CPE runtime.

As such, the proposed solution allows operators to remotely manage services and features within both CWMP compliant (such as residential gateways) and non-compliant equipment, by building on the benefits of an established *de facto* standard to provide a wide-reach management environment able to uniformly address the needs of device and service management in broadband access network environments.

3.2 The challenges of bringing CWMP to the residential LAN

This section discusses the usage of the CWMP protocol in the perspective of an operator-centric management approach of the residential LAN, with special relevance on the crucial aspects which might prevent or limit their operation on such conditions.

Special attention is devoted to the CWMP-related standards which try to address the problems related with NAT and proxied management, discussing their limitations and their potential usage in real-world scenarios, with a view on providing alternative solutions, which will be proposed on Section 3.3.

3.2.1 Management of internal LAN devices on NAT environments

There is a growing number of CWMP-compatible CPE devices, especially broadband residential gateways (RGWs) or home/residential gateways (i.e. the devices which stand in the frontier between the access network and the domestic LAN). This is not a surprise, since CWMP was originally conceived for DSL modems. From a formal point of view CWMP can be used to manage any type of CPE, but it is more or less evident that it works better with "frontier devices" (directly accessible from the operator access network) than with devices located inside the domestic LAN. As already discussed, support for such topologies was added later – in response to the increasing need to manage devices such as set-top-boxes and IP phones – and still lacks fully satisfactory solutions capable of dealing with a wider range of applications.

Also, management of CWMP devices behind a NAT environment is still a complex issue, addressed by TR-111 Parts I and II (and also TR-069 Annexes F and G, with the same purpose):

- TR-111, Part I (TR-069, Annex F), defines a mechanism which uses DHCP to make the RGW aware of the CWMP-capable devices inside the private LAN. Those devices are listed in the RGW instantiated data model, on the *ManageableDevice* table (whose entries expire in sync with the DHCP lease mechanism). Similarly, the managed device also records the information about its LAN RGW, which is used on the connection establishment *Inform* calls to make the ACS aware of the fact that the managed device is behind a specific RGW device. However, since this mechanism is dependent on DHCP it does not cover CPEs with a manually defined IP address.

Figure 3.2 shows an example of how this mechanism might be used to activate and manage a CWMP device behind a CWMP NAT RGW.

First, specific device-related service profiles must be loaded on the ACS database (such as VoIP service profile), to match devices and service configurations (**1**). Once the device is powered, it will send a bootstrap inform to initiate a session with the ACS (**2**), in which it will be configured accordingly with the service profile (**3**). Also, the device will notify the RGW of its existence, becoming registered on its *ManageableDevice* table (**4**), as per TR-111, Part I.

If desired, the ACS can enable active notification properties on the *ManageableDevice* table of the RGW, in order to be automatically notified of any update, such as the presence of a new device (**5**). The ACS can therefore correlate the RGW with the service policy and device ID, configuring and optimizing the RGW for the service (**6**). Also, the ACS can send billing and eventing information for OSS/BSS systems (**7**).

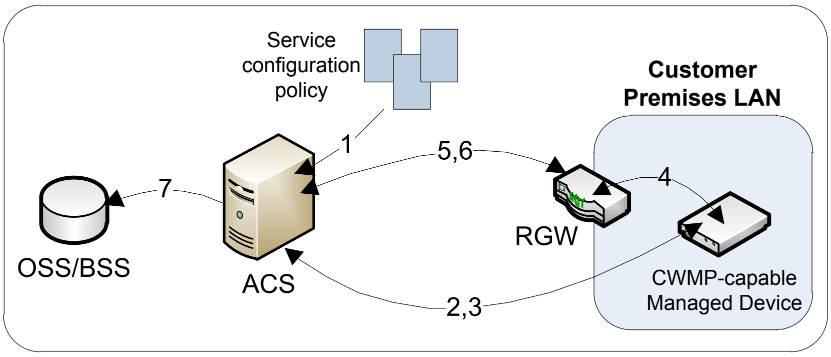

Figure 3.2: Usage of TR-111 device-gateway association mechanism

- TR-111, Part II (TR-069, Annex G) uses STUN to enable the ACS to make a connection request to a CPE inside a private LAN (for asynchronous configuration sessions). It requires the managed CPE to:
    - Discover that its connection to the ACS is via a NAT Gateway with an allocated private IP address.
    - Discover the NAT binding timeout and keep it alive.
    - Maintain an open NAT binding through which the ACS can send unsolicited UDP packets.
    - Determine the public IP address and port associated with the open NAT binding, and communicate this information to the ACS using the STUN mechanism [RFC5389].

STUN is not a reliable solution, something that becomes clear from its own specifications. Section 1 of the original STUN RFC [RFC3489] specifically stated "*This protocol is not a cure-all for the problems associated with NAT*" and the new RFC [RFC5389] refrains from considering STUN a complete NAT traversal solution ("*For these reasons, this specification obsoletes RFC 3489, and instead describes STUN as a tool that is utilized as part of a complete NAT traversal solution.*") along with a change of name reflecting the change in scope. Also, static port-forwarding rules are not viable to a multiple device environment. Therefore, this solution is prone to failure and the only alternative to configure a device inside the LAN is to wait for the next periodic Inform event to establish a management session. This strongly limits CWMP functionality.

3.2.2 Integration of LAN management technologies

CWMP is not directed towards LAN environments, like UPnP, DLNA or WMI. While such technologies focus on the management of devices and services within the LAN, CWMP provides remote device management services for the ACS run by the operator.

This is not a problem *per se*, since the two application fields are orthogonal to each other. Nevertheless, the absence of integration mechanisms means operators are often unable to take advantage of already existing management services. This means the dynamic environment of the domestic LAN – where the omnipresent UPnP and DLNA play an important role – is imperceptible to the operator. If CWMP is applied solely to the management of the home gateway this is not an issue, but in the context of management of internal devices this represents the loss of valuable management information (UPnP, for instance, could help solving the aforementioned NAT traversal problem). The same reasoning applies to WMI, if remote management of domestic PCs is at stake.

3.2.3 CWMP agents and proxies

Even recognizing that the number of CWMP-compatible devices keeps growing, it should be mentioned that integrating CWMP into some of the simpler CPE is not an easy task.

Some of the technologies used by CWMP – such as web-services and XML parsers – are not easy to port to the smallest embedded devices. Besides, the differences between CWMP web-services and the SOAP standards imply an additional implementation effort not easy in embedded systems. Despite the existence of commercial CWMP software for embedded devices [DIMARK, AXIROS] many CPE will not support CWMP, thus requiring some integration solution for legacy devices.

The CWMP data model provides some support for home gateways to act as management proxies (see Figure 3.3) for devices inside the subscribers LAN, treating each legacy device as a "Service" of the Home Gateway (by TR-106). The ACS communicates directly with the home gateway and is unable to distinguish between services provided by the home gateway itself and legacy devices. However, to the best of our knowledge, there are no implementations or reference guidelines taking advantage of those mechanisms.

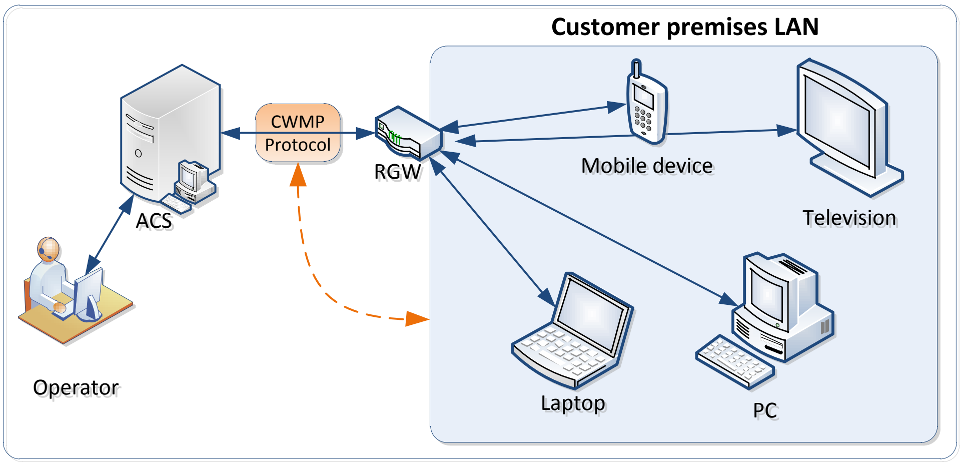

Figure 3.3: RGW-mediated CWMP proxied management model

3.3 A CWMP agent extensibility framework

CWMP documentation naturally focuses on the management protocol and expected ACS and CPE behavior, making no assumptions or recommendations on how to implement the ACS or the managed devices. This makes it possible to improve the agents without sacrificing CWMP compliance.

In order to better cope with some of the already mentioned problems, we propose a dynamic agent extensibility framework with clear decoupling between CWMP protocol services (concentrated in the so called “Master Agent”) and the device-specific management interfaces to be provided via CWMP (distributed across “Subagents”). Communication between the Master Agent and Subagents is based on a new protocol (extended CWMP: X-CWMP), much simpler than CWMP. Figure 3.4 illustrates this architecture.

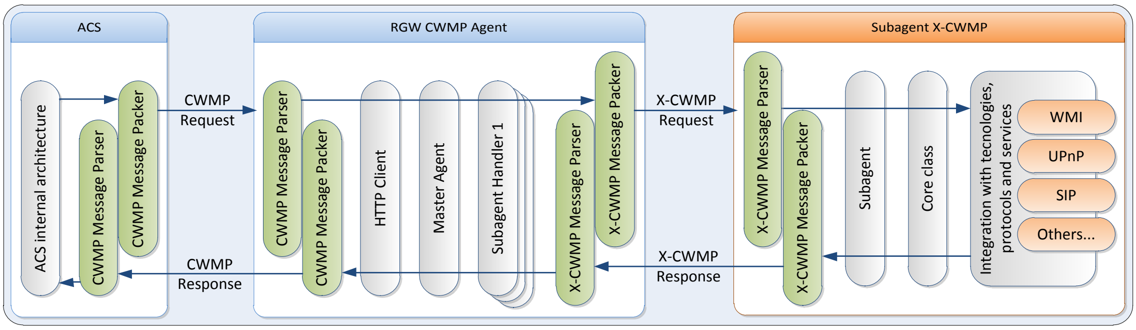

Figure 3.4: Proposed architecture for CWMP extensible agents

Even when implementing a classic CWMP agent – contained in a single CPE – the advantages are manifold, in comparison with monolithic approaches:

- Porting the master agent to new CPE devices is simpler, since it has no device-specific logic or code. The master agent just implements a “neutral” CWMP stack.
- Adding new services to the agent is only a matter of implementing new subagents, a task not requiring detailed knowledge of CWMP. When a new service is deployed in a CPE, for instance, the corresponding management interface can be quickly developed.
- Managed Services can be added or removed dynamically, without the need to restart the CWMP agent.

In addition, this approach also opens the way for less conventional configurations. For instance:

- Master agents and subagents can be located in different devices, allowing CWMP agents to proxy the access to those devices (which, for some reason, might not support CWMP). Figure 3.5, for instance, illustrates a CWMP agent in a home gateway also working as proxy for three internal devices. This allows for easier integration of legacy devices and also helps to solve the NAT traversal problem.
- Subagents can also be used as “Protocol Proxies” enhancing integration with existing LAN technologies.

According to the circumstances, X-CWMP subagents constitute a generic plug-in capability allowing the CWMP stack to communicate and abstract: (i) managed services of the host device; (ii) managed services of other CPE devices; and (iii) management services provided by other protocols and associated data models.

Incidentally proxied management also has the added benefit of further enhancing customers’ privacy, since the ISP doesn’t actually need to directly interfere in the LAN scope to perform management tasks, rather interacting with the RGW.

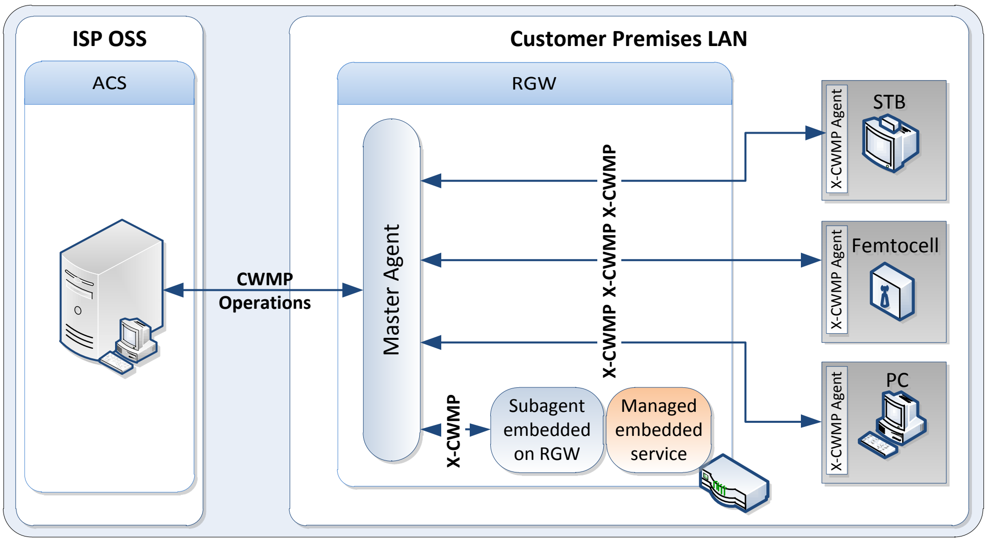

Figure 3.5: CWMP extensible agent in a device proxying scenario

3.3.1 Operation

The master agent is responsible for receiving, converting and forwarding requests. The effective processing of each request is performed by the subagent(s) associated with the related managed service(s). Master and subagent communication uses the specifically conceived X-CWMP protocol, which is discussed further on.

In the proposed architecture the master agent acts like an orchestrator and does not process any requests at all. On reception of a request it identifies the subagent(s) responsible for the specific objects(s) invoked in the request and forwards the request(s) to those agent(s).

Each subagent has a specific lifecycle, defined by an internal state machine (Figure 3.6). When a subagent is bootstrapped, it must register on the designated master agent. Once registered it becomes available for request processing, entering a loop where message exchanges and sync requests are processed as the master agent forwards them. Subagents may be unregistered by request or in failure situations.

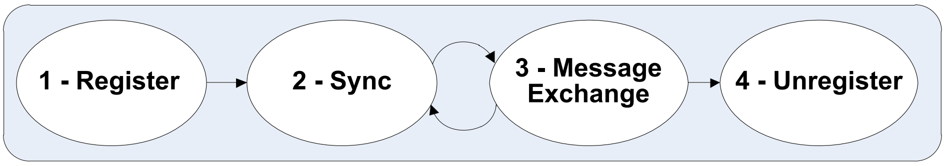

Figure 3.6: X-CWMP subagent lifecycle

Depending on the circumstances, subagents may have a pre-defined master agent (when the master agent and the subagent are located in the same device or when the network configuration is known in advance) or find the master agent using the discovery methods discussed later on.

Communication between master agent and subagents makes use of the X-CWMP protocol, based on XML messages (Figure 3.7) transported over TCP/IP, with optional use of SSL.

```
<Message>
<Seq>LAST+1</Seq>
<ID>5</ID>
<AgentID>Agent1</AgentID>
<ParameterName>InternetGatewayDevice.DeviceInfo.Manufacturer</ParameterName>
<ParameterValue>CorporateX</ParameterValue>
</Message>
```

Figure 3.7: Example of X-CWMP message (*SetParameters* request)

Every X-CWMP operation must be acknowledged. This makes it possible to identify faulty or unavailable instances whilst maintaining a convenient level of atomicity. Table 3.1 shows the list of defined X-CWMP operations.

Table 3.1: X-CWMP Protocol Messages

ID	Message Type	Message Description
1	*Register*	Register the subagent into the master agent
2	*Unregister*	Unregister the subagent from master agent
3	*Acknowledge*	Message acknowledge
4	*GetParameters*	Get parameter values
5	*SetParameters*	Set parameter values
6	*Synchronize*	Synchronize (agent synchronization for atomic operations)
7	*Fault*	A fault as occurred
8	*ReturnValue*	Return the request parameter value
9	*RunMethod*	Run a specific method

When a subagent registers on a master agent it becomes associated with the CWMP Data Model objects and properties for which it is responsible (see Figure 3.8). After a registration attempt passes validation the subagent is ready to receive requests.

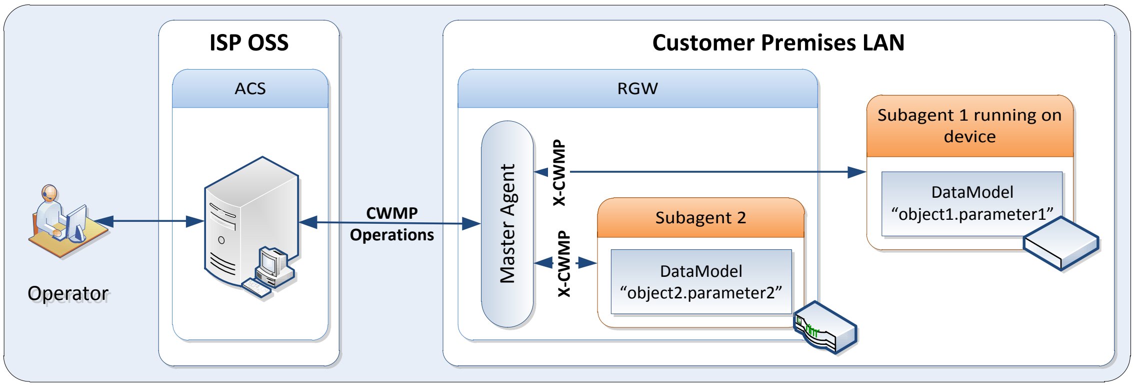

Figure 3.8: X-CWMP subagents data model association

When the master agent receives a CWMP message it checks for the registered subagent(s) responsible for handling that request. If they are found, the master agent will communicate with the respective subagents through X-CWMP requests for the processing of the original CWMP message. Next, the subagents reply to the master agent with the requested data. Finally, the master agent will create and send a combined CWMP message to the requester. Figure 3.9 presents a typical sequence diagram, including the initial registration of two subagents and, later on, the processing of a CWMP request that, taking into account its content, needs to be handled by those two distinct subagents.

Splitting CWMP requests across more than one subagent introduces the need for subagent synchronization whenever atomic processing of the CWMP request is required, implicitly or explicitly (either by the nature of the CWMP operation or by the nature of the involved objects). This means that in some situations X-CWMP requests need be processed using a two-step approach. Nevertheless, this does not impose added complexity in the development of managed services, since the problem is already present in monolithic CWMP agents.

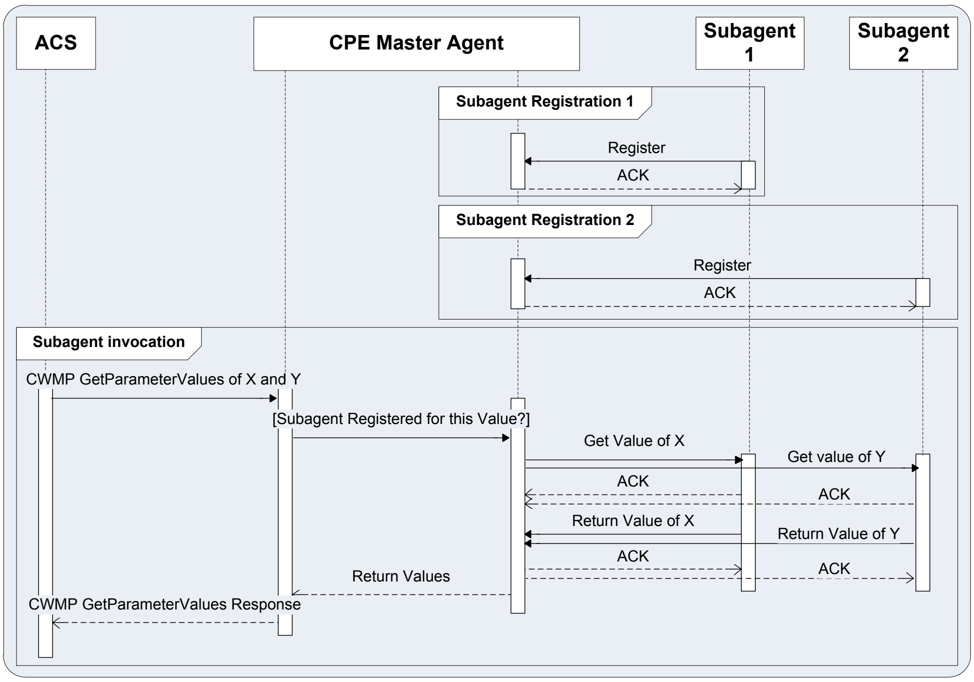

Figure 3.9: X-CWMP sequence diagram (Registration, *GetParameterValues* Operation)

Due to their nature, faults that occur in the context of X-CWMP (i.e. between the master agent and the subagents) may reflect on CWMP standard errors. Figure 3.10, for instance, illustrates a situation where a subagent fails to process a request (related with one of the managed objects it is responsible for), generating a CWMP fault.

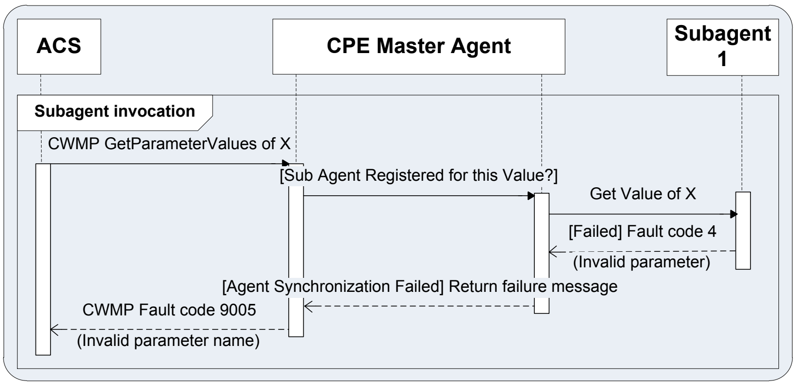

Figure 3.10: Handling of X-CWMP failure (example).

3.3.2 Implications on the CWMP data model

As already explained on Chapter 2, the standard data model for a CWMP-capable device is hierarchically organized like a directory tree with a single root node. Among the sub-elements of the root object, there is a "*Services*" object that contains all Service Objects associated with services embedded within the managed device or the data model for devices for which it provides proxied management services (it should be noted that a single device might include more than one Service Object).

X-CWMP subagents may integrate into the CWMP data model in two different ways, according to their nature (subagents representing an internal service of host CPE; subagents representing LAN devices proxied by the home gateway). In both cases CWMP standard specifications and principles are preserved in a coherent manner:

- In **conventional scenarios**, where the agent and the subagent are hosted in the same device (i.e. when the subagent implements the management interface for a specific module of the CPE) the subagent registers a node of the object tree and becomes responsible for every object and properties below that node (if the subagent is responsible for two independent tree nodes it registers twice in the master agent). This way the subagent does not need to specify every object is answers by.
- In **proxying scenarios** a similar method is applied, but with a twist in order to preserve compliance and integrity. A mechanism for registering Device-Gateway associations on a LAN is described in TR-111, using specific DHCP tags. According to TR-111, a CWMP-capable LAN device registers on the home gateway instantiated data model using a *InternetGatewayDevice.ManagementServer.ManageableDevice.{i}* object that contains basic information about the device (*ManufacturerOUI, ProductClass, SerialNumber*, etc.). This mechanism was conceived just to track device-gateway associations. Nevertheless, it also allows X-CWMP subagents representing proxied devices to discover and register themselves on the master agent without overlapping their tree of managed objects with the trees of the proxy CPE or other proxied devices.

Information and properties of LAN devices managed through X-CWMP agents are embedded on the home gateway CWMP data model through the use of dynamic TR-106 data model extensions. From a data model point of view, TR-106 allows an *InternetGatewayDevice* (home gateway) to act as management proxy for devices inside the subscriber LAN (Figure 3.11). Each proxied device can be modeled through a "Service" object instance which contains the correspondent data model. This mechanism was originally devised to enable a CWMP capable device to proxy the management functions for one or more other devices that are not CWMP capable. In this case, the ACS only communicates with the CWMP capable device, that incorporates the data models for the devices for which it is acting as a management proxy.

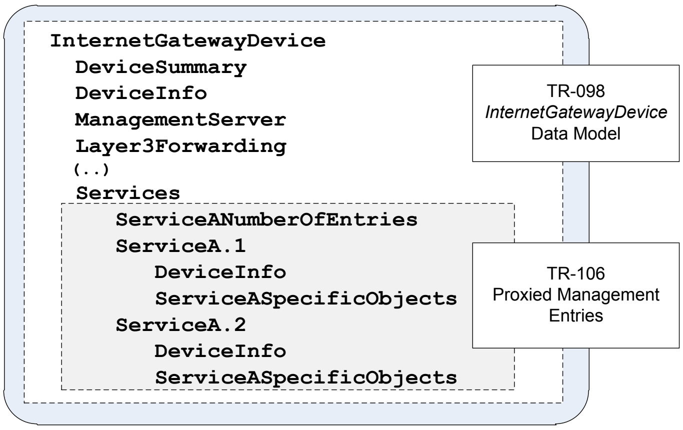

Figure 3.11: Example of TR-106 support for proxied management with two devices, both implementing a "ServiceA" object functionality.

TR-106 also specifies the possibility of using vendor-specific parameters/objects for a given device, using the notation *X_<Vendor>_VendorSpecificName* for the identifier. In line with this possibility, the parameters and objects related with each X-CWMP subagent are published as if they were a set of CWMP properties of a proxied device – using vendor specific extensions for particular needs not predicted by the attributes and objects defined in standard data model profiles. For each registered subagent, the master agent (i.e. the home gateway) creates an internal object handler instance to deal with the connection process and, if successful, another one to maintain information about the CWMP data model extensions provided by the subagent. For the ACS this process is completely transparent, since it still makes use of the standardized CWMP mechanisms.

3.3.3 Prototype implementation

In order to validate the proposed framework and to explore potential application scenarios, a proof-of-concept extensible agent was developed.

Implementing the X-CWMP master agent involved the development of a complete CWMP/TR-069 stack from the ground up. This option, however, was not motivated by technical reasons: it should not be difficult to adapt commercially available implementations of the TR-069 stack.

Figure 3.12 highlights the key modules of the master agent. Making use of the provided X-CWMP stack, creating new subagents is very simple, since the developer only needs to focus on implementing the managed services, providing request handlers for the corresponding objects.

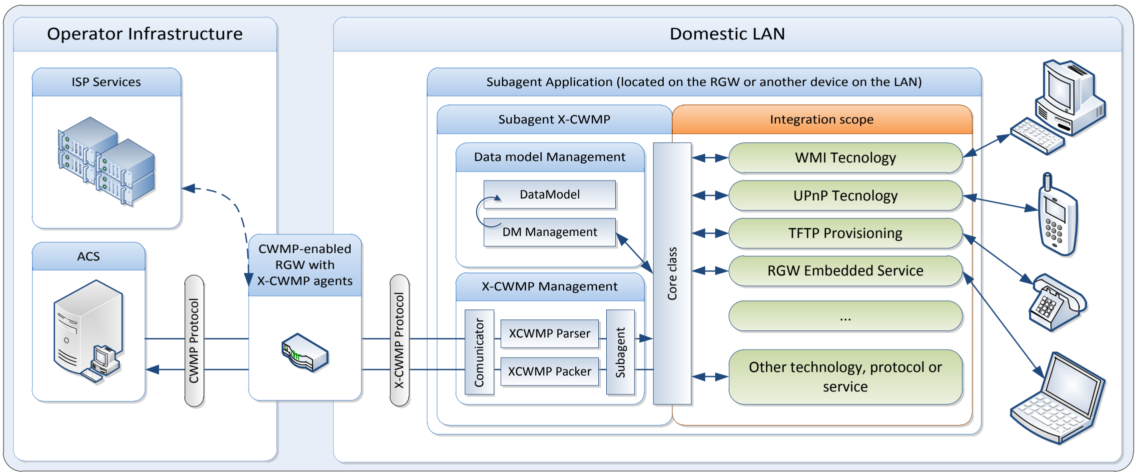

Figure 3.12: X-CWMP subagent module architecture

The Master Agent and a Reference Subagent were implemented both in Java and ANSI C, in order to foster portability. Support for the C language, in particular, allows for the development of very lightweight subagents capable of operating on very modest embedded systems hardware. Adding subagents to various OpenWRT-based devices [OpenWRT], for instance, was pretty straightforward.

CWMP compliance was successfully tested using several third-party ACS servers. Empirical tests also showed no relevant degradation in the response time during typical CWMP operations, even when using secure channels for communication between master and subagents. Although performance degradation may theoretically occur (due to the logical and sometimes even physical decoupling between master and subagents), in practice the relevant bottlenecks will remain in the ACS (that handles thousands of CPE) and the access network between the ACS and the master agent.

3.4 Application scenarios

In this Section we discuss possible application scenarios for the proposed X-CWMP framework, including the conventional scenario – where X-CWMP is used merely to increase the flexibility of the CWMP agent – and more advanced scenarios such as proxying for non-CWMP devices, integration with WMI and integration with UPnP.

3.4.1 Conventional scenario

The classic scenario corresponds to the usage of X-CWMP as a way to decouple the CWMP stack from the effective management interface of each managed service within a single CPE. The advantages of X-CWMP in this scenario were already identified in Section 3.3. Figure 3.12

is representative of this scenario, assuming the Master Agent and all Subagents reside on the same device.

3.4.2 Device proxying and NAT traversal

As already discussed, X-CWMP may be used to address two situations where CWMP has known shortcomings:

- One of these situations is device proxying. Non-CWMP devices placed inside the domestic LAN can become remotely manageable by adding X-CWMP subagents (which are much lighter and easier to deploy than full CWMP agents) binded with an external master agent – hosted for instance by the home gateway. Integration of the services provided by proxied devices in the data model can be implicit (those devices are transparently represented as services of the proxy) or explicit (using the already mentioned TR-106 vendor-specific parameters) but always coherent with CWMP specifications.
- Another situation relates with the problem of reaching managed devices behind NAT firewalls. As already discussed, the mechanisms devised by CWMP to such scenarios are unpractical and sometimes even unfeasible, since the STUN methods are often unable to deal with the variety of real world NAT implementations. X-CWMP subagents implicitly solve the NAT Traversal problem, transforming the home gateway in a management hub for all proxied devices, whose information is dynamically embedded in the data model of the master agent and therefore made accessible to the ACS. From a security standpoint, the X-CWMP approach has the benefit of completely eliminating the need for firewall rule exceptions to directly access devices inside the domestic LAN.

3.4.3 Integration of other management protocols and frameworks

As already mentioned, one of the side-benefits of the proposed extensible agent framework is the enhanced support for the development of generic layers to interface with popular LAN management technologies. This opens CWMP to a wide range of management information sources and mechanisms, allowing it to transcend its operation scope. The use of X-CWMP subagents, either located inside the RGW or the managed device allow them to provide protocol bridges for other management protocols or services (see Figure 3.11). For instance:

- Integration of desktop management APIs and protocols in CWMP management applications is interesting because it opens the way for a new class of added-value services related with remote management of desktop computers (see Figure 3.13). An ISP may, for instance, sell to domestic or SOHO users PCs bundled with software and remote management services (antivirus, software updates, remote recovery, etc.). Another example is the remote management of media center PC-based appliances, for increased integration with operator provided services such as IPTV and VoD.

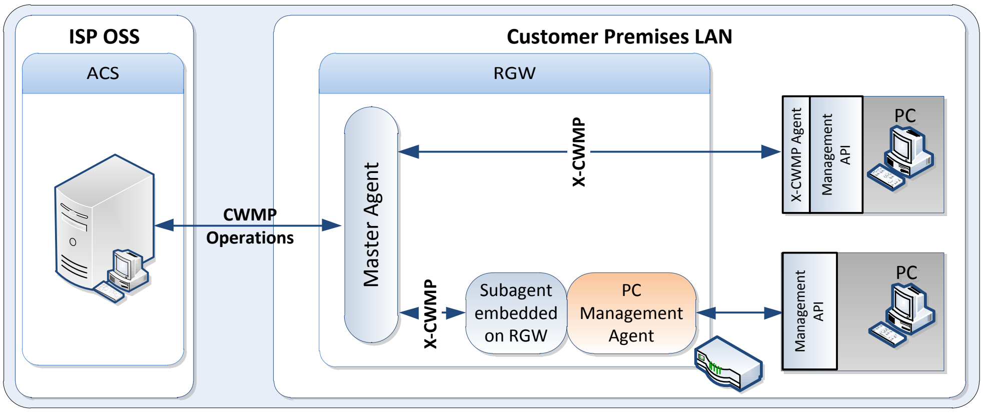

Figure 3.13: PC management using X-CWMP subagents

- X-CWMP subagent might be also used to enable CWMP to talk to other interoperability frameworks, such as UPnP or DPWS. From the operator point-of-view, being able to remotely access internal LAN services is extremely helpful to diagnose and configure the customer LAN (at least in the segments of relevance to the operator, for instance to provide connectivity and quality of service to IP phones and set-top-boxes) and relevant CPE devices, such as wireless access points, storage devices or media devices.
- The generic interoperability mechanism provided by X-CWMP subagent also enables CWMP to get information which can be used for other purposes besides configuration, device/service monitoring and/or diagnostics. By gathering information from different sources, it becomes possible to track device event status changes, for security monitoring purposes. For instance, UPnP is vulnerable to security issues, as shown by past history with serious threats such as Flash Vulnerability [Giobbi2008] and Conficker [Porras2009]. With UPnP-CWMP event traversal it is possible to track Internet Gateway Device (IGD) modifications at the ISP level, where security events from multiple domestic customers can be correlated in order to detect large-scale in-progress attacks (Figure 3.14). When such attacks are detected the ACS may use CWMP to activate defensive countermeasures at the level of each home gateway.

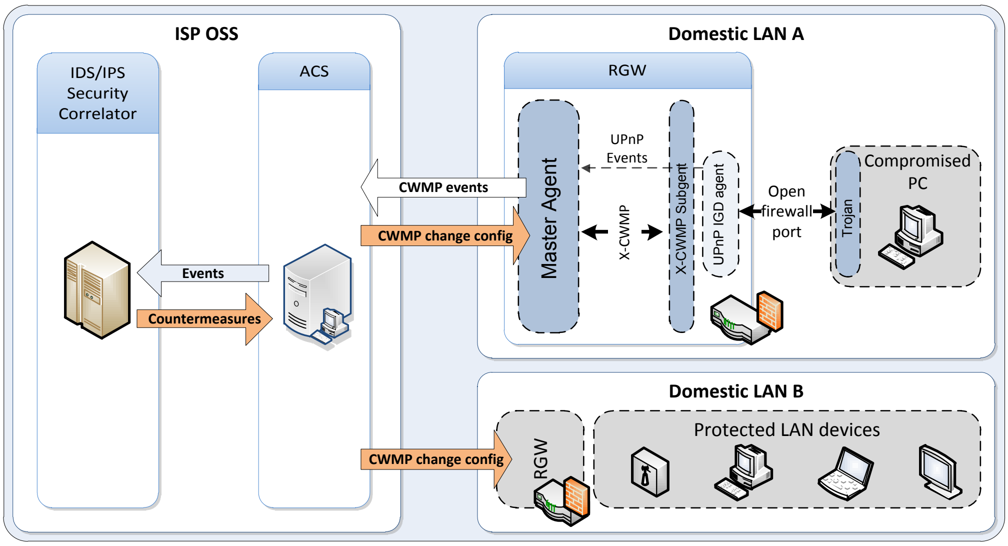

Figure 3.14: Foreign eventing integration for security purposes

On the next four chapters, specific examples of these application scenarios will be presented, explored and discussed, as this extensible management framework is orthogonal to all subsequent work.

3.5 Related work

To the best of our knowledge, X-CWMP is the first agent extension framework to be proposed for CWMP. Nevertheless, the concept of extensible management agents has been around for some time, especially in the area of SNMP, with at least three such proposals: SMUX [RFC1127], DPI [RFC1592] and AgentX [RFC2742]. Conceptually, X-CWMP is not far from AgentX, but the obvious differences between SNMP and CWMP (especially at the level of operational models and data models, but also underlying technologies such as XML) allow X-CWMP to be more elegant and easier to use.

The notion of device proxying based on CWMP-compliant gateways is also not novel. The idea is clearly latent in a number of specifications recently produced by the Broadband Forum, especially the already mentioned TR-106 data model. Nevertheless, those specifications appear to be incomplete and exploratory, in the sense that they provide valuable data modeling mechanisms but still lack an extensive discussion of device proxying approaches and exploitation scenarios. Our proposal, in this context, constitutes one of the first known contributions for such a discussion.

A similar observation can be made regarding integration with LAN-level management technologies. While the BroadbandHome working group of the Broadband Forum is working in on PD-174 (Remote Management of Non TR-069 Devices) [BBForum2011, MR-239], that eventually will be part of the upcoming TR-069 specification refresh (Amendment 4), this is a much-awaited work which was delayed for several years and whose details are unavailable to non-members. TR-157, published in late 2009, does provide support for representing UPnP and DLNA devices in CWMP data models, but the integration framework itself is not discussed nor explored by the Broadband Forum.

Projects B@Home [Freeband2004] and MUSE [MUSE] attempted to provide some mechanisms similar to the ones hereby described – work developed within these projects by Hillen [Hillen2009], Royon [Royon2007a], Delphinanto [Delphinanto2009] and Nikolaidis [Nikolaidis2007] already suggested bridging CWMP with other protocols (especially UPnP). Also independently, Minokoshi [Minokoshi2010] published a similar proposal. Axiros, a commercial provider of management solutions also has recently launched a CWMP management proxy product called AXPAND [AXIROS2010], which enables a CWMP ACS to manage devices which only support other protocols such as telnet, SSH, SNMP, among others.

Albeit sharing some similarities with these proposals, the solution hereby proposed goes one step further by integrating generic extensibility right into the CWMP stack, through means of runtime dynamically binded agents that communicate with the master CWMP agent (the CWMP agent residing on a home gateway) through the X-CWMP protocol. These agents can reside on the same CPE as the master agent, or embedded in managed devices, allowing the master agent to extend its data model to accommodate new devices and/or protocols. This allows for a more flexible operation mode and a wider range of applications. Also, most of the proposals developed within MUSE and B@Home target protocol-specific integration while the mechanism hereby proposed provides a generic extensibility mechanism.

Moreover, and specifically regarding AXIROS' product, the comparison must be further detailed. While apparently similar in conceptual terms to the proposal hereby described, both approaches are in fact rather different as AXPAND requires a proxy component to be deployed in a PC inside the user premises LAN – in practice, a dedicated CWMP client agent with protocol proxying capabilities. This is not a standardized nor practical solution for residential clients as it requires a dedicated PC to serve as a protocol proxy – while the concept itself could be adequate for occasional diagnostic situations (where a client might temporarily execute the proxy component on its PC to help the operator reach devices on his LAN), the fact that there is no version for Windows or MacOS operating systems limits its usability, even for these scenarios. Our proposal predates this product, as it is more flexible, since it does not require a dedicated proxy system, allowing for a hybrid distributed approach where proxying components might be located at the RGW or another device.

The recently started FIGARO project [FIGARO] intents to federate the residential network ecosystem in a way that is very similar to the one hereby proposed. However, as of August 2011, it was in its early stages, with no significant outcomes.

3.6 Final synthesis

CWMP is expected to become the predominant technology for remote management of devices that, despite located in the local network of domestic broadband users, still need to be directly managed by operators. Nevertheless, effective deployment of CWMP-based applications is still an ongoing process, with some success but slower than desirable.

In this chapter we proposed a CWMP-based extensible management framework. This framework promotes the clear decoupling between the CWMP protocol stack and the effective management interfaces associated with each service to be managed using CWMP. This solution is also a toolkit for developing CWMP agents, allowing for easier and faster deployment of new managed services, either for new or already existing agents.

In addition, X-CWMP can also become an invaluable tool in less conventional scenarios, such as device proxying, traversal of NAT firewalls and integration with LAN-level technologies such as WMI and UPnP. Those scenarios are already implicitly present in some of the more recent specifications of the Broadband Forum (especially in the form of data model extensions) but, up to this date, they have not been properly discussed or explored. In this chapter we discussed the potential of such scenarios and proposed specific solutions for their deployment, making use of X-CWMP.

Next chapters will deal with particular application scenarios of the proposed management framework, exploring new service and application paradigms built on top of it.

Part III:

Applications and usage scenarios

4. Management of home network devices

With the increasing need for Internet Service Providers to manage devices inside the customers' LAN (such as set-top-boxes, VoIP handsets, DVRs and media players, most of which not even supporting CWMP), there comes the need to deal and interoperate with a wide range of interfaces and management protocols that coexist on the customer premises network environment. Failing to do so means the loss of valuable resources which could complement existing device and service management frameworks, for remote equipment configuration, monitoring and diagnostics.

In this chapter we discuss how the framework proposed on Chapter 3 can be used to create a CWMP-based management ecosystem able to interoperate with non-compliant devices, bridging protocol and interface gaps in a transparent manner. Specifically, two use-cases of CWMP-integrated management of "foreign" devices will be presented:

- Integration of off-the-shelf SIP endpoint devices with CWMP, using a solution that supports the vast majority of commercially available SIP phones whilst maintaining full compatibility with the original CWMP specification.
- Integration of the Universal Plug and Play (UPnP) protocol framework (discussed on chapter 2) with CWMP-based architectures, to allow operators to remotely discover, manage and configure a wide range of UPnP devices on the customers' LAN, from printers to network equipment.

In both cases, the integration is totally transparent for the operator's ACS and for the managed devices, thus allowing ISPs to use the CWMP management infrastructure they already have to configure and provision "legacy" devices.

This chapter is structured as follows: Section 4.1 introduces the problem of foreign device management. Section 4.2 discusses CWMP-based provisioning and management mechanisms for non-compliant SIP endpoint devices and Section 4.3 presents a CWMP-UPnP integration solution for UPnP device management. Section 4.4 analyzes related work and Section 4.5 concludes the chapter.

It should also be mentioned that the content of this chapter is extensibly based on previously published materials [Cruz2011, Cruz2012].

4.1 Introduction

The home network environment (which has been introduced and discussed on Chapter 2) has been evolving into a dynamic ecosystem of protocols, devices and services (such as file and printer sharing, media distribution and domestic automation), whose configuration – a crucial issue when device interoperability is a must - is an increasingly complex task which the common user is frequently unable and, sometimes, even unwilling to deal with. This situation was further aggravated with the expansion of residential broadband coverage and the introduction of converged IP-based services (such as VoIP and IPTV, frequently bundled in triple-play service offers).

In converged-IP scenarios, the reliability and performance of IP-based services depend on proper configuration and operation of all the equipment in the service delivery path, including endpoints (e.g. set-top boxes, SIP phones) and domestic/residential gateways (home routers). This pushed for the proposal of adequate management mechanisms focused on the needs of service providers, among which CWMP emerged as the de facto standard. However, CWMP is not able (nor has any standardized provisions) to interoperate with non-compliant devices – this excludes a vast device ecosystem from the operator's management reach.

The dimension of the problem assumes bigger proportions as it becomes clear that, albeit successful, CWMP popularity varies accordingly to the type of device. For instance, devices that combine traditional media with new functionalities, such as "smart" TV sets, DVD or Blu-ray players are becoming more common in domestic LANs and most of them do not even support CWMP. Even most VoIP devices, which are increasingly present in the home network, are not CWMP-compliant, supporting instead other enterprise-oriented LAN management mechanisms – this is explained by historic reasons, since those devices were originally conceived for enterprise LANs.

This situation has several consequences. Apart from making operators unable to configure the customer LAN device ecosystem; it also means the loss of management information and mechanisms which are vital to other tasks, such as diagnostics or fine-tuning of device and service performance. This chapter presents two representative use cases that show how to overcome the problems of integrating non-CWMP devices in CWMP management frameworks:

- The first use case deals with integration of VoIP/SIP legacy endpoint devices within the CWMP management scope, allowing for remote, operator-driven provisioning, management and monitoring (using the residential gateway as protocol mediator).
- The second use case demonstrates the complete integration of the UPnP framework within CWMP, enabling operators to perform remote UPnP device control, configuration and monitoring.

4.2 Provisioning and management of SIP endpoint devices

Fixed telephony was among the first services to be offered in converged *n-play* bundles, supported using Session Initiation Protocol (SIP) [RFC3261] signaling over IP networks. As a result, the usage of SIP-based (Session Initiation Protocol) VoIP devices has known a significant growth in domestic environments, either in the form of standalone equipment (e.g. SIP telephones) or embedded devices (as it happens with some residential gateways, which embed analog-to-SIP adaptors).

In comparison with the ages-old analog telephone, which was a dumb piece of hardware, SIP devices require specific configuration in order to operate. However, while ATA devices embedded on residential gateways may potentially benefit from the management mechanisms already present at the gateway-level to enable operators to remotely configure them, standalone SIP devices inside the customer LAN are a different matter. For Internet Service Providers

(ISPs), the provisioning and management of those devices is a challenge – especially standalone SIP phones, since most of them were exclusively designed for corporate LAN usage, not supporting adequate mechanisms for remote management over broadband access networks.

Nevertheless, the need for proper management and provisioning of VoIP devices located in home networks is critical, especially considering that fixed voice services are regarded by users as crucial. For this reason, operators need to take over the responsibility of seamless managing all the equipment in the service delivery path, including home routers and the SIP endpoints located inside the home network.

To address this issue, the CWMP framework encompasses a standardized data model for SIP devices, specified by TR-104, which is commonly supported to manage Analog Telephony Adapter (ATA) devices embedded on CWMP-compliant home gateways and, potentially, VoIP SIP telephones. However, the effective number of CWMP-compliant standalone SIP telephones is still residual, to say the best. This situation happens because, prior to their introduction in domestic households, SIP telephones were already common in corporate environments, which were quick to grasp the benefits of converged IP communications.

As a result, the standard feature set in terms of management capabilities is frequently restricted to mechanisms exclusively oriented towards enterprise LAN environments. More specifically, the lowest common denominator supported by the vast majority of SIP telephones (provisioning based on the Trivial File Transfer Protocol, TFTP [RFC1350]) is not suitable outside managed LANs, due to security and reliability reasons. Even alternative mechanisms supported by some of the most recent equipment, such as HTTP(S) provisioning [RFC2616], require manual configuration or direct support from LAN-specific protocols, like DHCP [RFC2131].

In this section we propose to bridge this gap by means of an integration framework that allows operators to remotely provision and manage Commercial, Off-The-Shelf (COTS) SIP phones using CWMP as the frontend management interface. This extension builds on the generic extensibility mechanisms of CWMP to enable proxied management of SIP devices, using the home gateway as a mediator and preserving compliance with CWMP protocols and data models. As an added benefit, the proposed integration framework also provides a unified management interface to COTS SIP devices, abstracting the heterogeneous mechanisms and data models supported by each manufacturer.

4.2.1 An overview of legacy SIP device provisioning mechanisms

As already mentioned, SIP telephones are traditionally bundled with LAN-oriented management and provisioning features. Among the latter, TFTP provisioning became the most popular, due to its simplicity and ease of implementation, remaining the baseline mechanism supported by the majority of VoIP phones to the present day (with some, more recent models, also supporting HTTP or HTTPS).

By default, most telephones come with support for TFTP provisioning enabled. In general terms, the classic TFTP provisioning method used by different manufacturers consists of the same basic approach, with some slight differences from one to the other. It relies on TFTP and DHCP, sharing some similarities with remote boot protocols like PXE [Intel1999]. Its operation can be summarized as such (see Figure 4.1):

- Upon initialization of the network stack, DHCP options [RFC2132] are used to pass information about the location of the files to be used for the boot process, which are subsequently downloaded from a TFTP server. By default, SIP telephones request for the DHCP option 66 (TFTP Server Name) in the *DHCPREQUEST* stage. Additionally, options 67 (Boot File Name), 150 (TFTP Server Address) or 43 (Vendor specific Information) might also be used for similar purposes (1).
- The telephone gets the information about the TFTP server from the requested options embedded in the *DHCPACK* response (2).

- The device attempts to download its configuration files from the specified TFTP server (3,4) at predefined intervals (which can be changed), to check for configuration changes (5).

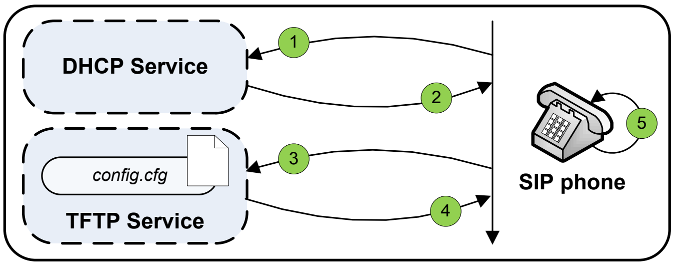

Figure 4.1: Generic TFTP provisioning mechanism for VoIP phones

Considering the characteristics of controlled LAN environments, TFTP provisioning is acceptable for large deployments, even if it lacks some security features (a situation which is normally dealt with by using VLANs and other protection methods). However, the same does not apply to access networks, due to the security and reliability issues imposed by the TFTP protocol. TFTP makes use of UDP for transport (which is an inherently unreliable protocol), suffering from several, implementation-dependent, limitations and inefficiencies (such as file size limitations). Also, TFTP is usually blocked on edge routers and does not offer any kind of protection for its payload, which is transferred in plaintext (even if some phones can optionally use more or less proprietary encryption schemes, difficult to set up). Another limitation relates with the need of explicit DHCP support, which still stands when replacing TFTP by HTTP or HTTPS.

4.2.2 Integrating provisioning mechanisms with CWMP

The proposed solution is based on the addition, at the home gateway, of a specific integration component (see Figure 4.2). This component acts as module of the gateway's CWMP agent, interacting with the ISP ACS (by means of CWMP operations) and with the home gateway TFTP and DHCP Servers (to indirectly interface with the managed SIP devices). In order to preserve the coherence of the CWMP data model, this specialized component is responsible for its own CWMP data model extensions (in accordance with TR-106 guidelines and the TR-104 standard), which are mapped and integrated on the global data model of the CWMP Agent.

In general terms, the proposed architecture integrates together a CWMP frontend interface with a TFTP-based provisioning backend. This extension is implemented as a module which maps the configuration properties of SIP endpoints into the CWMP data model of a domestic gateway, so that the operator ACS can remotely manage them using CWMP. The module will then configure the services needed for TFTP-based provisioning (TFTP and DHCP), which are embedded on the residential gateway and are exposed on the LAN-side interface, accordingly with the CWMP parameterization.

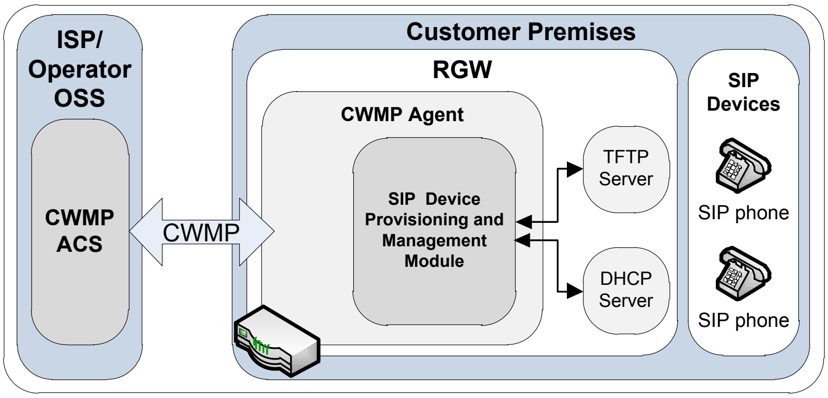

Figure 4.2: Generic architecture of the proposed CWMP extension

Integration with the CWMP data model

Considering the CWMP Data Model objects and parameters, the integration module will be responsible for two different branches of the supported CPE data model (the RGW, in this case): the standard branch for VoIP devices (as defined by TR-104) and another set of TR-106 compliant extensions for provisioning purposes (see Figure 4.3). Together, these two sets provide a common management interface, which is independent of the specific manufacturer of the managed SIP phone.

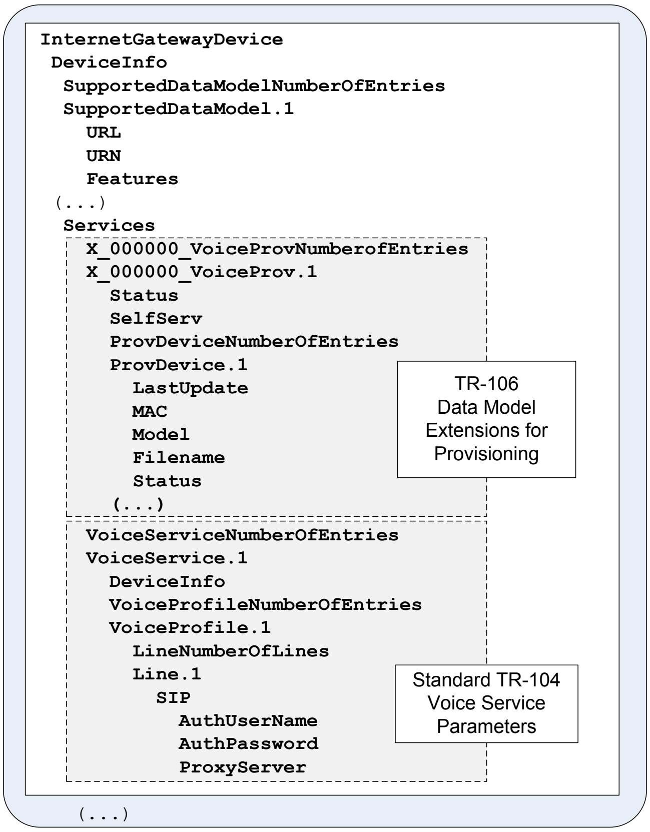

```
InternetGatewayDevice
 DeviceInfo
  SupportedDataModelNumberOfEntries
  SupportedDataModel.1
    URL
    URN
    Features
 (...)
  Services
    X_000000_VoiceProvNumberofEntries
    X_000000_VoiceProv.1
      Status
      SelfServ
      ProvDeviceNumberOfEntries
      ProvDevice.1
        LastUpdate
        MAC
        Model
        Filename
        Status
      (...)
    VoiceServiceNumberOfEntries
    VoiceService.1
      DeviceInfo
      VoiceProfileNumberOfEntries
      VoiceProfile.1
        LineNumberOfLines
        Line.1
          SIP
            AuthUserName
            AuthPassword
            ProxyServer
    (...)
```

Figure 4.3: Relevant data model structures for the CWMP extension

The data model extensions are organized using the *X_000000_VoiceProv* object on the CWMP data model (with an Organizational Unique Identifier (OUI) [IEEEOUI] defined as *000000* for test purposes). These data model extensions are also declared in a *SupportedDataModel* table entry (of the *DeviceInfo* root object), which points to an URL containing their XML description hosted on the CPE itself, so that the ACS can gain knowledge about the device data model supported by the managed device. The same happens for the TR-104 standard data model, whose CPE support has to be declared before being instantiated.

The information needed to control the provisioning mechanism for legacy VoIP devices is embedded on the CWMP data model of the residential gateway using a TR-106 compliant vendor specific entry named *X_000000_VoiceProv.{i}* for the management service, while specific SIP and device configuration parameters are managed using TR-104 data model structures.

This solution is possible thanks to TR-106, which allows for proxied management of noncompliant CWMP devices, further reinforced by TR-104, which goes one step further by explicitly allowing proxy management of non CWMP VoIP devices embedding their management information properties on the TR-098 data model of home routers (even though, to the best of our knowledge, our proposal is the first to make effective use of this possibility). This way, the operator ACS may seamlessly manage remote SIP devices using CWMP.

General operation

General operation can be split in three different phases: the initialization process (involving the ACS and the home gateway), the provisioning of the SIP device, and the runtime management of the SIP devices.

The **initialization process** goes as follows (see Figure 4.4).

If the operator has enabled voice services for the user in question and the CWMP provisioning extension is enabled (parameter *X_000000_VoiceProv.{i}.Status* is set to *Enable*), the ACS is able to add a new object to the *ProvDevice* table for each new provisioned SIP device on the LAN. Once the object is instantiated, the subagent will also enable the corresponding TR-104 data model extensions for the device (creating a new *VoiceService* object for it), immediately changing the value of the corresponding *ProvDevice.{i}.Status* parameter to *Ready*).

CWMP allows for definition of notification attributes on parameters, either by default or explicitly, using the *SetParameterAttributes* CWMP method. For instance, when the state of a given parameter with the Active Notification attribute enabled changes, the CWMP device will schedule the execution of an Inform method on the ACS to notify it. This will happen with the *ProvDevice.{i}.Status* parameter once its value changes to Ready as the result of the initialization process, since it has the CWMP Active Notification Attribute enabled by default.

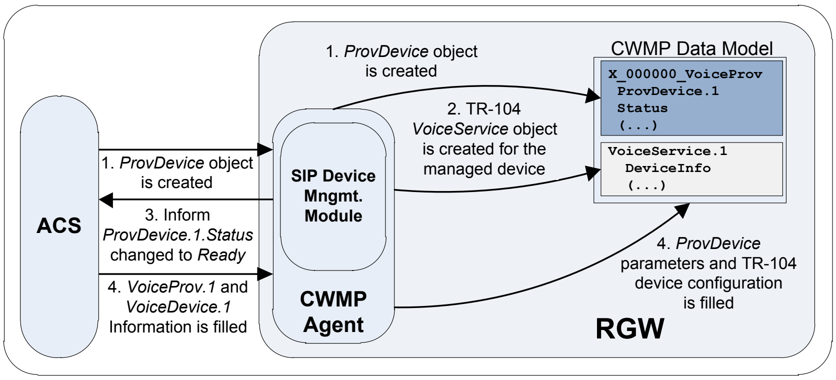
1. ProvDevice object is created
CWMP Data Model
X_000000_VoiceProv
ProvDevice.1
Status
(...)
2. TR-104 VoiceService object is created for the managed device
VoiceService.1
DeviceInfo
(...)
ACS
1. ProvDevice object is created
SIP Device Mngmt. Module
3. Inform ProvDevice.1.Status changed to Ready
4. VoiceProv.1 and VoiceDevice.1 Information is filled
CWMP Agent
4. ProvDevice parameters and TR-104 device configuration is filled
RGW

Figure 4.4: Initialization process

The **provisioning process** goes as follows (see Figure 4.5).

Once the ACS is notified of the successful initialization of the provisioning subsystem, it will be able to fill the TR-104 and remaining *VoiceProv.{i}.ProvDevice.{i}* parameters needed to configure the device (e.g., Model, MAC). Once it has finished it must change the corresponding *ProvDevice.{i}.Status* parameter to Provision, in order to enable the activation of the TFTP provisioning according to the corresponding *VoiceService* parameters. Once the process has been concluded, its state will change to *Enabled* (or *Fail*, if problems occur).

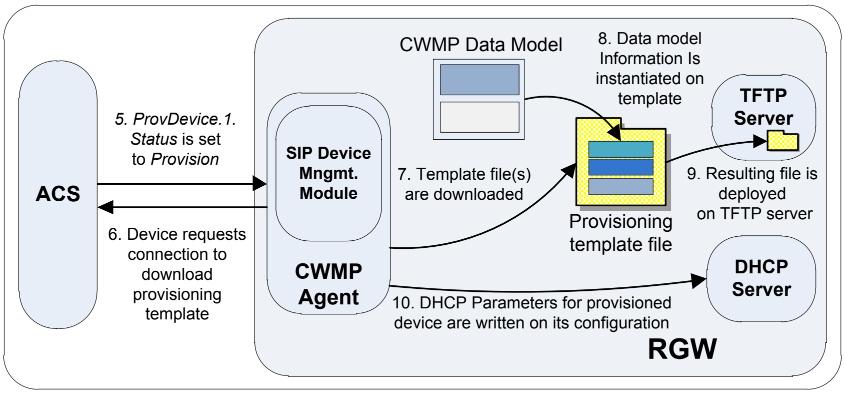
CWMP Data Model
8. Data model Information Is instantiated on template
TFTP Server
ACS
5. ProvDevice.1. Status is set to Provision
SIP Device Mngmt. Module
7. Template file(s) are downloaded
9. Resulting file is deployed on TFTP server
Provisioning template file
6. Device requests connection to download provisioning template
CWMP Agent
DHCP Server
10. DHCP Parameters for provisioned device are written on its configuration
RGW

Figure 4.5: Provisioning of the SIP device

Once the *ProvDevice.{i}*.Status parameter value is changed to *Provision*, the subagent will schedule an Inform to the ACS requesting the establishment of a session to download the device provisioning templates. The use of provisioning templates has to do with the fact that each hardware manufacturer has a different format and syntax for the provisioning files, therefore allowing the subagent to deal with them in a uniform way.

Each template comes as a simple text file whose name and content embeds predefined placeholders which are used by the subagent to instantiate TR-104 parameters by simple pattern-based text substitution, therefore creating the individual provisioning file for each SIP device. Once the template is downloaded and SIP parameters are instantiated on specific placeholders, the TFTP provisioning file is created and made available on the local TFTP service. Also, the corresponding DHCP static entries are created, biding the device MAC address to the TFTP server embedded on the residential gateway, and therefore to its predefined provisioning file.

This solution can be also used with HTTP(S) provisioning mechanisms, which in general also use DHCP option 66. By swapping the TFTP server with a HTTP server the same methods can be used to locally provision VoIP telephones. TFTP was used for the proof-of-concept since it is least common denominator in terms of provisioning methods among SIP telephones. Also, since most telephones check for their provisioning files at regular intervals to update their configuration, this mechanism is able to implement a complete remote device management proxied via the CWMP Agent of the residential gateway.

Runtime management goes as follows (see Figure 4.6).

For setting specific parameters on the TR-104 model, the ACS will interact with the CWMP agent on the domestic gateway using the *SetParameterValues* CWMP method. The SIP phone management module will repeat the same process used for initial provisioning, recreating a new file using the new configuration parameters, from the stored template (which was downloaded on the initial provisioning). Depending on the TFTP file resync interval (which most phones allow to customize), the phone will gather the new configuration file and instantiate the new configuration. After the process is concluded the corresponding *ProvDevice.{i}.Status* parameter will be set to *Enabled* (or *Fail*, if a failure occurs), and the *ProvDevice.{i}.LastUpdate* will be updated with the timestamp of the last TFTP download.

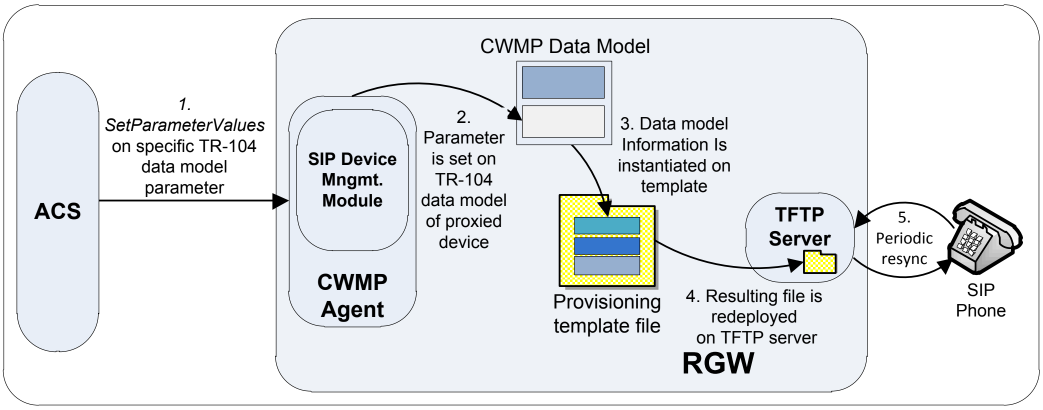
ACS
1. *SetParameterValues* on specific TR-104 data model parameter
SIP Device Mngmt. Module
CWMP Agent
2. Parameter is set on TR-104 data model of proxied device
CWMP Data Model
3. Data model Information Is instantiated on template
Provisioning template file
4. Resulting file is redeployed on TFTP server
TFTP Server
5. Periodic resync
SIP Phone
RGW

Figure 4.6: Setting a new configuration parameter for the SIP device

As for getting management information, there are two different situations. When management data is of static nature, the management information if directly gathered from the TR-104 internal data structures and sent to the ACS when it is requested (using a CWMP *GetParameterValues* method).

As for dynamic management information (such as statistical network traffic info or SIP session state), a different approach was adopted. When the SIP device is initially provisioned, and the model is on a list of supported models, an attribute mapping file is also downloaded which is used by the CWMP management module to map the web management interface on the SIP phone, converting it to TR-104 information. This process is taken care by a thread, which dynamically refreshes the information from the managed device (see Figure 4.7).

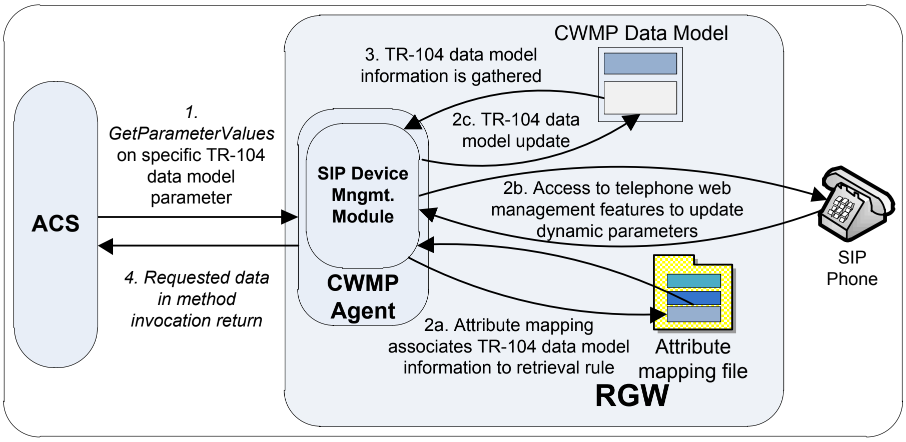
CWMP Data Model
3. TR-104 data model information is gathered
ACS
1. GetParameterValues on specific TR-104 data model parameter
2c. TR-104 data model update
SIP Device Mngmt. Module
2b. Access to telephone web management features to update dynamic parameters
SIP Phone
4. Requested data in method invocation return
CWMP Agent
2a. Attribute mapping associates TR-104 data model information to retrieval rule
Attribute mapping file
RGW

Figure 4.7: Dynamic parameter mapping on SIP devices

It should be mentioned that, unlike provisioning and setting operations, this feature is not supported for all SIP devices. An example of supported SIP devices with suitable features for implementing dynamic parameter mapping is the Linksys/Cisco SPA model range, which allows for downloading its entire active configuration and state, by accessing the URL *http://device/admin/spacfg.xml* – in this case, attribute mapping is only a question of XML [W3C2006a] parsing, easily performed using regular expressions or a XML parsing library. However, the module has experimental support for mapping rules in the form of *Curl* [Curl] expressions, for those situations where information is only accessible by interaction with a Web user interface.

Regarding run-time management, it should also be mentioned that some SIP phones also provide some support for SNMP-based management [RFC3410]. This feature is attractive for enterprise networks but, once again, not suitable for access networks and domestic environments. Nevertheless, in this section we do not address integration of SNMP-based management with CWMP, since SNMP support is restricted to a small number of SIP phones and implemented in very heterogeneous ways (e.g. proprietary MIBs). Our focus was instead to devise a simple way to provision and manage the largest possible range of COTS SIP phones for domestic users connected by broadband access networks.

Overall, the solution hereby proposed constitutes the first known implementation of CWMP proxied management for non-compliant SIP devices – TR-104 already hints for this possibility, but without discussing how to do it. Support for SIP device provisioning, management and monitoring makes this a complete solution, which offers a unified approach to remote management of heterogeneous devices (albeit, as already mentioned, with some of the runtime management mechanisms only being supported for a device subset, for the time being). Thus, this solution solves the problem of dealing with devices from different manufacturers, with distinct management interfaces and features (sometimes even between models from the same manufacturer), in an elegant way.

4.2.3 Application scenarios

This subsection discusses how the proposed solution can be used to implement device provisioning and management. Specifically, two scenarios are presented: operator-controlled management and provisioning; and self-service provisioning. Each scenario allows for a different balance between operator centralization and user-intervention on the process.

Operator-controlled provisioning and management

For provisioning, this scenario is based on the conventional, operator-driven process, which corresponds to the standard operation case described on the previous section.

Figure 4.8 illustrates this scenario. A customer might buy a SIP telephone directly from the operator, filling a service subscription form that will be internally submitted on the store front-desk system to be processed by the operator OSS, in order to pre-provision the device. Eventually, this will result in a series of CWMP operations that will create, at the customer's home gateway, the configuration parameters to be provided to the new SIP phone (identified by its mac address), once it gets connected to the home network. Once the customer arrives home he is able to plug the telephone and have it immediately configured and ready for use. In addition to conventional SIP and device configuration parameters, this method can also be used, for instance, to configure default entries on the telephone number memory list, speed dial keys or other relevant functionalities. This process requires no manual configuration of the SIP phone, neither at the store or at home.

Runtime device management is available using the methods already described in the previous section, which allow the operator to use the CWMP interface of the home gateway of the customer for proxied device management.

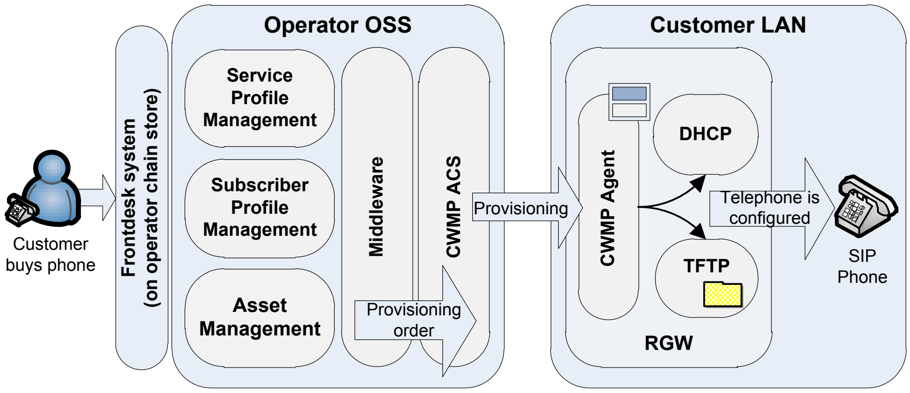

Figure 4.8: Operator-controlled provisioning

Self-service provisioning

The second method is designed to enable users to take care of the initial configuration process by themselves. A customer that subscribed a VoIP service from its ISP (likely bundled on a triple-play contract) can buy a COTS SIP telephone from a third-party and take care of its provisioning by answering a call on its newly bought telephone and using a simple and intuitive method based on an automated Interactive Voice Response (IVR) menu to provision the device. Its operation follows three different stages (Figure 4.9):

- Telephone is detected on the LAN and configured with a default self-service provisioning profile. This is possible because the subagent is able to deploy matching rules on the DHCP service configuration file which associate specific DHCP Vendor Class Identifiers (VCI) and MAC address prefixes to the supported device manufacturers' generic provisioning files.
- By means of CWMP component management operations defined by TR-157, the subagent is able to download and update these rules and files, accommodating new devices and manufacturers (1). When a new device is present on the LAN and issues a DHCPREQUEST, the service is able to apply a generic provisioning file that binds the device to a provisioning SIP extension used by the operator.
- Subsequently, the CWMP agent on the residential gateway creates a new *ProvDevice* entry and corresponding TR-104 parameters for the device. The *ProvDevice.{i}.Status* parameter (which has active notification attributes enabled) is changed to *SelfProv*, then scheduling an Inform to the ACS to signal the event. The ACS, which is coupled to the OSS infrastructure,

instructs the PBX to place a telephone call to the temporary extension assigned to the SIP phone that is to be provisioned (2). Once the extension is registered on the operator SIP PBX, the telephone will receive an incoming call from the service. By answering the questions posed by the automated IVR, the user provides the information required by the provisioning backend to configure the telephone (e.g. client ID and credentials). Optionally, the service can be configured to work in the reverse way (the user initiates the phone call).

- The IVR is used to collect user information for the OSS, which together with information from profile subsystems will be used to provision the device using the conventional method already described on the previous subsection (3).

This method gives customers the freedom to buy SIP phones from third parties (as long as they are included in the ISP hardware compatibility list) and to provision them in a very comfortable and cost-effective manner – the user only needs to interact, by phone, with an automated IVR.

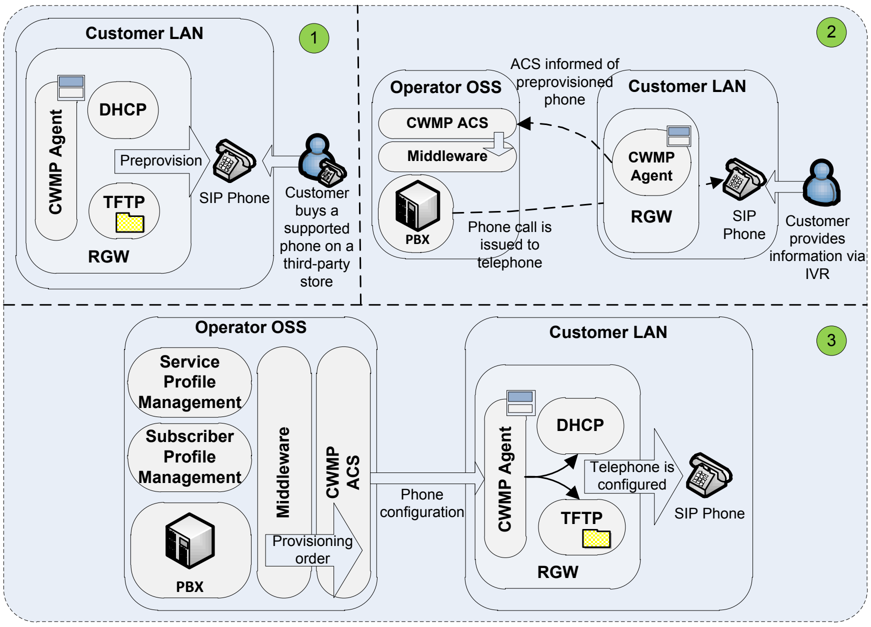

Figure 4.9: Self-service provisioning

4.2.4 Validation

In this subsection we discuss the proof-of-concept implementation of the home gateway CWMP agent (more specifically, of its SIP Device Management Module), the testbed used for experimental validation, and obtained results.

Proof-of-concept implementation

In order to create a proof-of-concept it was only necessary to extend the CWMP agent of a home gateway, in order to add the functionalities associated with management of SIP devices. As implicit in our proposal, no modifications were necessary at the levels of the ISP ACS or the SIP phones.

The proof-of-concept CWMP agent was developed in Java, using the dynamic X-CWMP agent extensibility framework presented in Chapter 3. Nevertheless, it could also be implemented using any other CWMP agent library, at the cost of less flexibility and dynamicity.

Figure 4.10 presents the architecture of the SIP Device Management Module. It maintains a basic set of module interface methods, as well as the state information about the data model elements for which it is responsible. It also interacts with the ISC DHCP [ISC] and ATFTP [ATFTP] services, already embedded on the domestic gateway (RGW) we used.

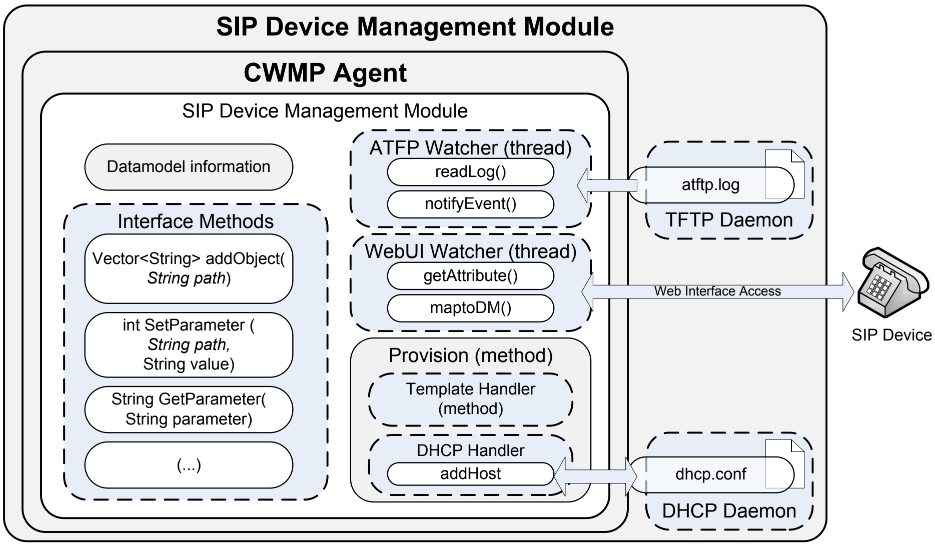

Figure 4.10: RGW implementation details

A TFTP watcher thread takes care of surveying the service log in order to track when devices download their provisioning files. This thread makes uses the *inotify* Linux kernel event mechanism [LDP] to get notifications when the log file is changed, therefore avoiding the use of polling mechanisms. When a device downloads its provisioning file, the *LastUpdate* parameter on its corresponding *ProvDevice* entry on the CWMP data model is updated with a timestamp. A provisioning method is also present, which takes care of template handling (for creating provisioning files) and DHCP configuration manipulation. The Web User Interface (Web UI) watcher thread provides the runtime management mechanisms for dealing with dynamic attributes, by accessing the web management interface of the SIP device and performing attribute mapping. For *Curl* expression support (used for gathering data from the device management Web UI by emulating user interaction), the *curl-java* package [Curl-java] is used.

As for the DHCP prefix matching system used in the self-service provisioning scenario, it was implemented using an ISC DHCP service feature which allows for the creation of device classes corresponding to specific pattern-matching conditions (see Figure 4.11). These rules allow for automatic provisional configuration of new SIP devices detected on the home network, enabling them to register on a PBX with a temporary extension for self-service provisioning purposes.

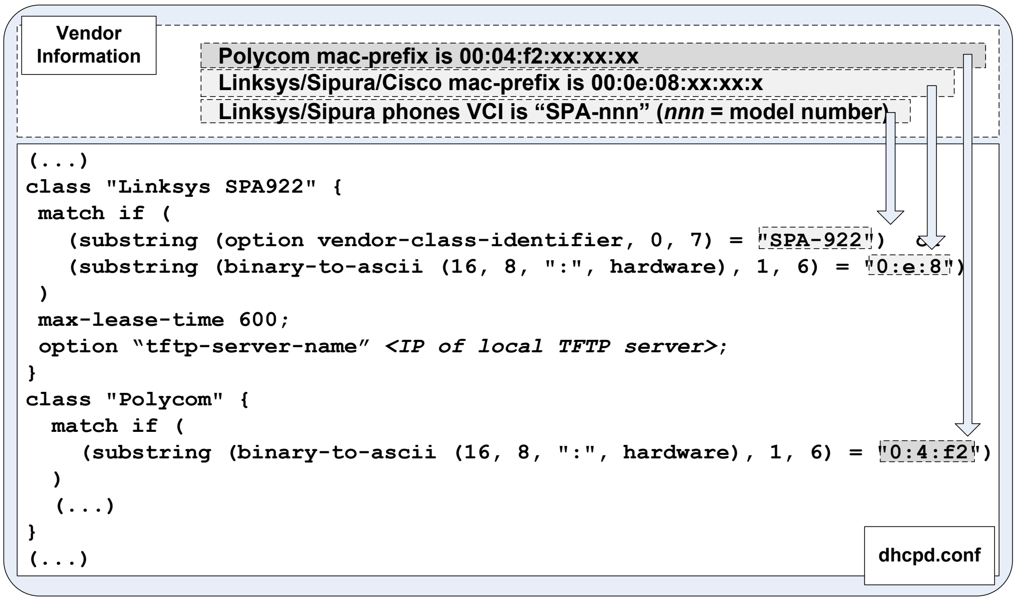
Vendor Information
Polycom mac-prefix is 00:04:f2:xx:xx:xx
Linksys/Sipura/Cisco mac-prefix is 00:0e:08:xx:xx:x
Linksys/Sipura phones VCI is "SPA-nnn" (*nnn* = model number)
dhcpd.conf

```
(...)
class "Linksys SPA922" {
 match if (
   (substring (option vendor-class-identifier, 0, 7) = "SPA-922") &&
   (substring (binary-to-ascii (16, 8, ":", hardware), 1, 6) = "0:e:8")
 )
 max-lease-time 600;
 option "tftp-server-name" <IP of local TFTP server>;
}
class "Polycom" {
  match if (
    (substring (binary-to-ascii (16, 8, ":", hardware), 1, 6) = "0:4:f2")
  )
  (...)
}
(...)
```

Figure 4.11: Example of DHCP vendor information matching for self-service pre-provisioning

Since DHCP host declarations take precedence over class declarations, once a device has been provisioned its specific DHCP entry will not conflict with the general pre-provisioning entry.

Testbed

To evaluate the solution, a testbed was assembled to reproduce the conditions of a managed broadband access network (see Figure 4.12).

The home gateway is a Linux system with an embedded CWMP agent and two network interfaces, supporting NAT, the ISC DHCP and ATFTP services. The VoIP PBX is an Asterisk 1.6 system [Digium] configured to validate endpoint registration. In order to mimic the conditions and restrictions imposed by the broadband access link, a transparent *Dummynet* bridge [Carbone2009] interconnects the customer with the ISP. A PC is connected to a monitor port on the Fast Ethernet switch, which mirrors all traffic coming to and from the port used by the ACS. This PC captures and measures network traffic using *Wireshark* [CACE2011].The ACS server is a Linux system configured with a CWMP ACS. Four of the most popular SIP telephone models were used for testing (all of them placed in factory reset state):

- Linksys (now Cisco) models SPA-922 and SPA-941;
- Yealink T-22P;
- and Polycom Soundpoint IP335.

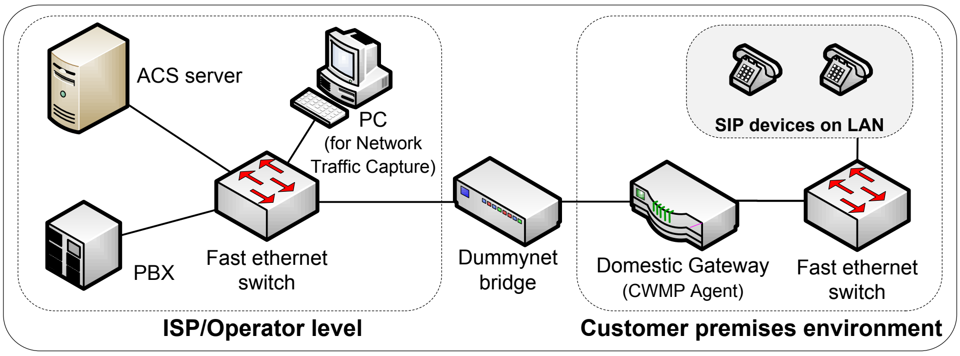

Figure 4.12: Testbed environment

Tests

Functional tests validated the integration schemes and applications scenarios previously proposed. Regarding performance evaluation, conducted tests focused on the operator-controlled provisioning model scenario. In this context, measurements were performed on the domestic gateway for two different aspects:

- Latency of the provisioning initialization process, from the initial addition of the *VoiceProv.{i}.ProvDevice.{i}* object to the point where the *VoiceProv.{i}.ProvDevice.{i}. Status* parameter changes to Enabled. This encompasses the download and processing of device templates, the automated creation of device-specific TFTP provisioning files and the configuration of the DHCP service for device provisioning.
- Latency of the complete device provisioning process, from the initial addition of the *VoiceProv.{i}.ProvDevice.{i}* object for the provisioned device to complete download of the TFTP provisioning file.

First, reference results were obtained on a scenario where all equipment were connected using 100Mb/s Ethernet (*Dummynet* bridge disabled). Next, tests were performed to measure latency on typical broadband access networks. The *Dummynet* bridge was configured to enforce bandwidth and traffic conditions representing typical commercial offers based on ADSL and GPON (see Table 4.1). The rationale for the specific configurations applied to *Dummynet* is discussed in Appendix A.

Table 4.1: Broadband test reference scenarios

	Nominal bandwidth (b/s) (Down/Up)		Effective bandwidth (b/s) (Down/Up)		RTT Latency	Pkt. Loss
ADSL	8M	512K	6.68M	427.5K	20ms	0.1%
	16M	1M	13.36M	835K	20ms	0.1%
GPON	20M	2M	18.6M	1.86M	5ms	0%
	100M	10M	93M	9.3M	5ms	0%
LAN	100M	100M	100M	100M	<1ms	0%

Measurements obtained for the Linksys SPA-922 are shown in Figure 4.13 (since initialization values obtained for the other models are similar, they will not be presented).

Results show the impact of access network conditions on the provisioning initialization process latency to be almost negligible. As for the apparent inconsistency for the full provisioning results, it has to do with the fact that the specific telephone model has a random default TFTP initial polling interval after a power up or reset of no more than 40 seconds.

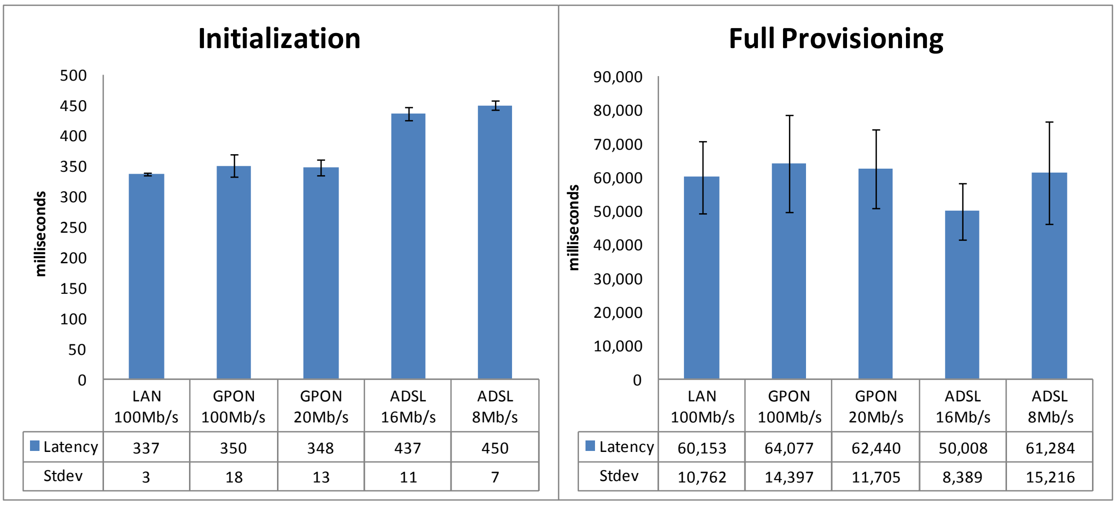

Figure 4.13: Latency for provisioning operations (Linksys SPA-922) (averaged values of 10 test runs; in milliseconds)

Nevertheless, the relevant conclusion from obtained results is that the latency remains well within an adequate timeframe for this type of operation, being more affected by the randomness of the initial TFTP polling (that also occurs in traditional enterprise LAN scenarios) than by the conditions of the access network or the devised integration framework – a situation that can be explained by the fact that measured network traffic between the ACS and the CPE, for the full provisioning process, was only around 30KB.

4.3 UPnP-CWMP integration for operator-assisted LAN management

The second use-case presented in this Chapter relates with UPnP-CWMP integration, for operator-assisted management of the home network.

The UPnP (Universal Plug and Play) framework was designed to simplify and/or automate the device interoperability and configuration in home networks. UPnP devices use UPnP protocols to advertise, discover and access services in a seamless way, with minimum user intervention. Device and service profiles were developed for specific device categories, such as Audio/Video (AV), Internet Gateway Device (IGD), Printing and Lighting Control. The scope and operation of the UPnP framework was already discussed with detail on Section 2.2.3.

UPnP support is popular across all classes of devices, especially media-related equipment. This is no coincidence, since UPnP is the basis for the DLNA media framework (also discussed on Chapter 2), being supported by a wealth of devices combining traditional media with new functionalities, such as “smart” TV sets, DVD or Blu-ray players, that are becoming more and more common in domestic LANs. As such, the UPnP mechanisms embedded on those devices would be of great value for remote diagnostic and configuration purposes. However, UPnP was not designed to operate in the environment of current broadband access networks (GPON, DSL, cable...), being of little or no use to remote management purposes.

On the other hand, CWMP, which is the standard for remote, operator-centric management, is only able to manage compliant devices. It lacks adequate integration mechanisms for key LAN technologies, among which UPnP stands out as one of the most important – its TR-157 [TR-157] data model extension does define profiles for embedding information about UPnP devices, but it does not allow device control.

In this section we present an extension to the CWMP protocol that allows operators to fully access and manage UPnP-compliant devices. This extension builds on generic extensibility mechanisms of CWMP to bridge with existing LAN technologies (as is the case with UPnP) thus providing more flexible and inclusive management topologies – whilst keeping compliance with the original CWMP framework.

4.3.1 An extension for bridging CWMP with UPnP

The UPnP standard and related protocols were designed for use on small, domestic LANs, even if their specifications allow some limited forms of remote bridging and access for roaming devices [UPnPForum2009]. In domestic LANs, limitations related to management protocol traffic overhead and scalability are not as critical as they are in WAN environments, due to close proximity, available bandwidth and the relatively reduced number of existing devices. As an example, the use of SSDP messages over Multicast UDP for UPnP device discovery or status update is inadequate for WAN environments, where the use of multicast UDP (inherently unreliable and frequently blocked on edge routers) and the considerable number of UDP messages sent to overcome unreliability create a traffic burden. Also, the existing DCP for security [UPnPForum2003] is disregarded by most device implementations, which openly expose their eventing and control mechanisms in an unsafe way.

Considering these limitations, we proposed a CWMP bridging subagent to map the description of UPnP devices into the CWMP data model of the residential gateway so that the operator ACS can remotely manage them using CWMP.

UPnP-CWMP bridging subagent architecture

The proposed bridging solution makes use of the CWMP dynamic agent extensibility framework that was presented on Chapter 3. This framework makes it possible to separate CWMP protocol services from the device and service-specific management interfaces to be provided via CWMP. The former role is fulfilled by a “Master Agent”, with the latter being handled by one or more “Subagents”, each one responsible for its own TR-106 CWMP data model extensions mapped and registered on the data model of the Master CWMP Agent.

The UPnP-CWMP bridging subagent maps the description of UPnP devices found on the domestic LAN on a gateway CWMP data model. For this purpose, the subagent embeds an UPnP control point which allows it to interact with all UPnP devices on the domestic LAN (Figure 4.14). The properties of the extensibility framework allow the subagent to be embedded either on the residential gateway or on another device on the domestic LAN (e.g. a PC temporarily used as a UPnP-CWMP bridge for diagnostics).

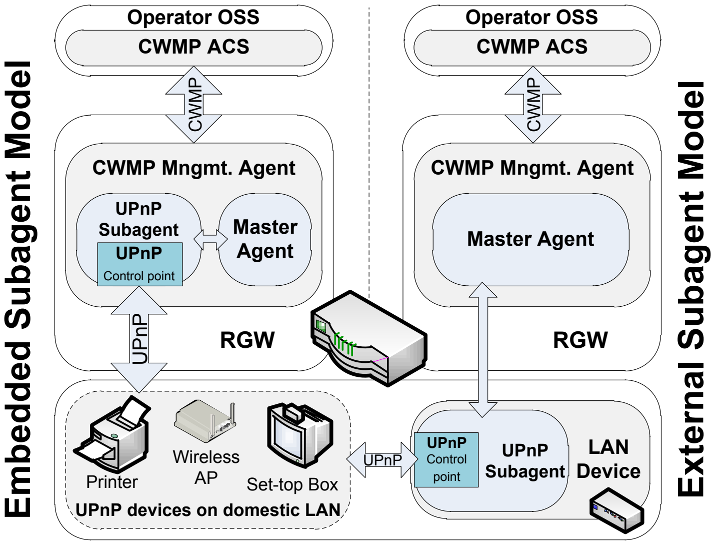

Figure 4.14: UPnP bridging subagent operation

Integration with the CWMP data model

Generically, when the CWMP bridging subagent registers on the master agent it becomes associated with the CWMP Data Model objects and parameters for which it is responsible.

In this specific case, the subagent will be responsible for two different branches of the supported CPE data model: the standard, TR-157-defined data model for UPnP device enumeration and another, TR-106 compliant, vendor data model extension for the objects and parameters used for UPnP device interaction.

TR-157 UPnP data model parameters and objects are only useful for discovery and enumeration purposes. They are embedded under the UPnP object and divided in two categories: UPnP properties for the managed device itself, positioned under the *Device* sub-object, and discovered UPnP devices on the domestic LAN, under the Discovery sub-object (Figure 4.15). Only the parameters under the *Device* object allow for some limited control of UPnP features of the host device.

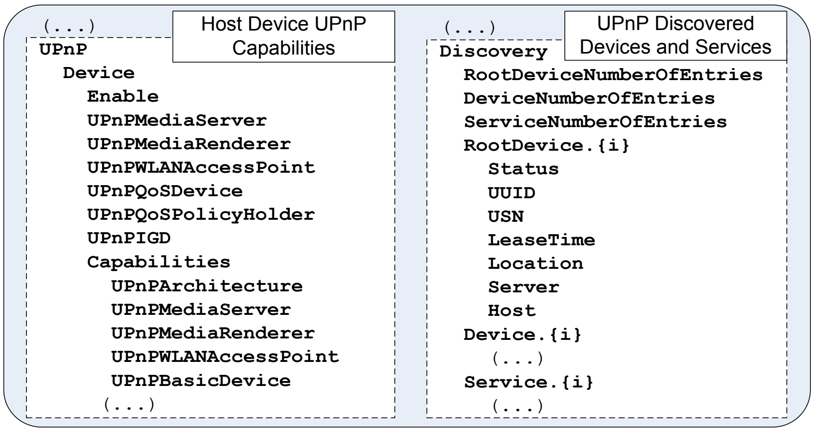
Host Device UPnP Capabilities
UPnP Discovered Devices and Services

```
(...)
UPnP
  Device
    Enable
    UPnPMediaServer
    UPnPMediaRenderer
    UPnPWLANAccessPoint
    UPnPQoSDevice
    UPnPQoSPolicyHolder
    UPnPIGD
    Capabilities
      UPnPArchitecture
      UPnPMediaServer
      UPnPMediaRenderer
      UPnPWLANAccessPoint
      UPnPBasicDevice
      (...)
```

```
(...)
Discovery
  RootDeviceNumberOfEntries
  DeviceNumberOfEntries
  ServiceNumberOfEntries
  RootDevice.{i}
    Status
    UUID
    USN
    LeaseTime
    Location
    Server
    Host
  Device.{i}
    (...)
  Service.{i}
    (...)
```

Figure 4.15: Standard TR-157 UPnP data model component objects

As already mentioned in Section 2.2.3, an UPnP service description includes a list of commands, or actions, to which the service responds, and parameters, or arguments for each action. It also includes a list of variables that model the state of the service at runtime, described in terms of their specific characteristics. To add UPnP device control functionality to CWMP, the device data model had to be extended to accommodate TR-106 compliant objects and parameters, placed under the *Services* standard object. The bridging subagent makes use of a strategy that explicitly maps UPnP device and service descriptions in CWMP terms, with UPnP structures (such as state variables, services and methods), being abstracted as CWMP objects and parameters (see Figure 4.16).

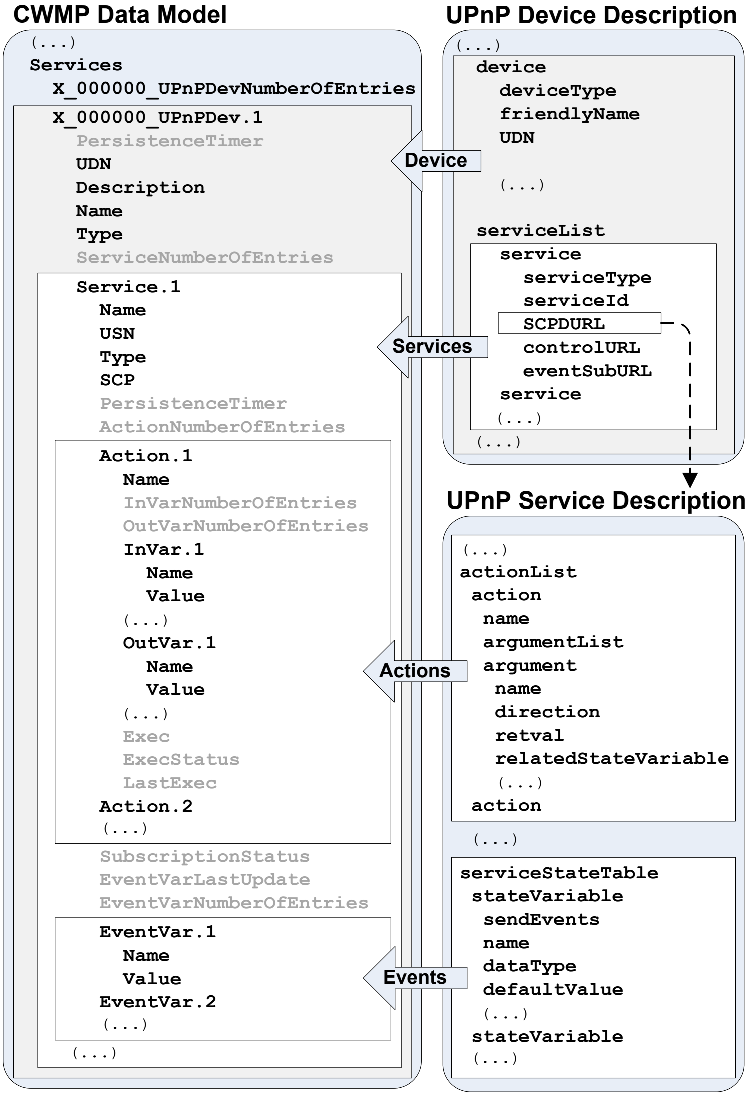

Figure 4.16: Extended CWMP data model mapping for UPnP (grayed entries do not directly relate to UPnP description, having specific purposes)

Each UPnP device, including embedded devices corresponds to a vendor-specific *X_000000_UPnPDev* CWMP object entry on the CWMP data model, with an Organizational Unique Identifier (OUI) [IEEEOUI] defined as *000000* for test purposes. Since standard TR-157 extensions already describe the relations between UPnP root devices, embedded devices and services, all modeled device instances are mapped as objects under the *Services* standard object.

The mapping technique that was implemented supports new or non-standard DCPs in a completely transparent way. This is important because, albeit standardized, DCPs for specific device categories can be, (and are frequently) expanded by manufacturers using non-standard actions and variables.

When a new UPnP device joins the domestic LAN, the subagent will process its description data, generating the corresponding instances on the CWMP data model. It uses a generic mapping template, which is simply applied to each mapped UPnP device or service. As such, this approach reduces processing time for initialization of new entries and simplifies the extension of the supported and instantiated data models of the CWMP device hosting the master agent.

Persistence of mapped information for a given UPnP device is controlled by using the CACHE-CONTROL information header on UPnP SSDP device announcements, corresponding to the TR-157 *LeaseTime* parameter on each discovered UPnP instance object (Figure 5.15). This information is used by UPnP control points to monitor the presence of a given UPnP instance – if an advertisement is not received prior to the expiration of the time interval specified in the CACHE-CONTROL header, it is assumed that it is no longer available. Nevertheless, the subagent will maintain the associated mapping scheme on a local cache, indexed by its individual UPnP Unique Device/Service Name (UDN/USN).

If an UPnP instance returns to the LAN, the processing time needed to generate the corresponding CWMP mapping will be reduced by reusing the cached data (updating dynamic components such as the *Location* parameter in the standard data model). Since TR-157 defines the ability to list inactive UPnP instances as optional, such entries will always be removed from the extended data model, but will remain listed on standard parameters until the cached data persistence timer expires (which defaults to 10 times the last lease duration and can be modified using the *PersistenceTimer* parameter).

UPnP-CWMP action invocation

For each service, action invocation is performed by loading the corresponding input parameters with required values, followed by an invocation action. Each input parameter is modeled as an *InVar* object, holding its instantiated value in the *Value* CWMP parameter. When all required input data is loaded, action invocation is performed by filling the *Exec* parameter with the "execute" value. Once executed, its SOAP return code and invocation timestamp will be returned on the *ExecStatus* and *LastExec* parameters, respectively.

Also, the *ExecStatus* parameter has a CWMP Forced Active Notification attribute enabled, which forces the managed device to inform the ACS (invoking the ACS *Inform* method) when its value changes, as a result of an UPnP action invocation. Therefore, it becomes possible to integrate event-driven UPnP action execution into CWMP (Figure 4.17).

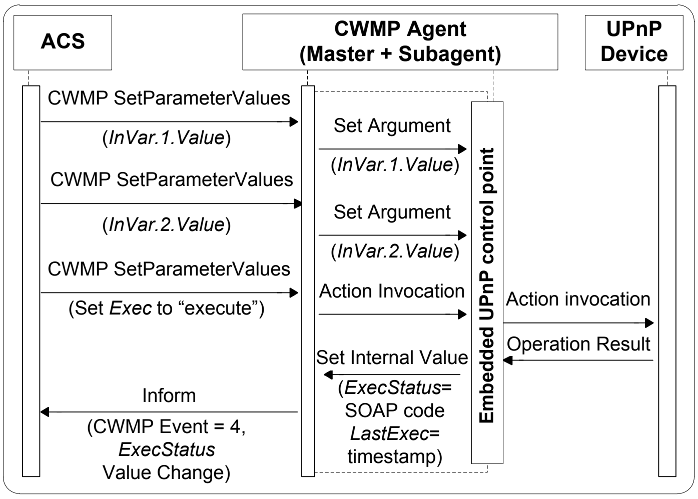

Figure 4.17: UPnP action invocation via CWMP

UPnP eventing (GENA) integration

Direct UPnP device state variable querying, as implemented via the UPnP *QueryStateVariable* action has been deprecated. Instead, the UPnP Forum recommends to query such variables using actions explicitly defined on the DCP. As such the CWMP data model for UPnP services only supports eventable state variables, which can be tracked and updated by means of the GENA subscriber-based eventing mechanism (see Section 2.2.3 for an in-depth discussion of UPnP).

When a service is subscribed, the embedded bridge will receive a special event updating the value of all evented variables – this is UPnP standard behavior. Subsequent events will only update the state of specific variables, as they change. The implemented data model maps eventable variables for an UPnP service directly on CWMP parameters, which are initialized with the defaults declared on the Service Control Protocol Description (SCPD) and are updated when the corresponding service is subscribed by the UPnP control point embedded by the CWMP subagent.

Service subscriptions are controlled by the *SubscriptionStatus* parameter (default: "*disabled*"). To enable service subscription, it can be set as "enable", enabling passive update semantics. When received, UPnP variable change notify events will be used to update the CWMP data model parameters mapped to the changed variable(s). To enable ACS notification of changed variables, the *SetParameterAttributes* CWMP method must be used to enable notification attributes on parameters corresponding to mapped UPnP variables which are to be tracked. In this situation, received UPnP variable change notify events will also generate a CWMP Inform invocation to notify the ACS.

4.3.2 Application scenarios

This UPnP-CWMP bridging extension was designed for remote diagnostics, configuration and management. Embedding UPnP control inside CWMP for other uses (such as DVR/media player session control across domestic LANs, bridged by broadband connections) is not viable since the latency requirements of normal CWMP operation are not adequate for user-device interaction. For instance, an ACS cannot directly initiate a connection. CWMP defines an ACS connection request mechanism to instruct the CPE to initiate a connection as soon as possible, incurring a latency penalty.

By using CWMP, the operator can remotely interact with UPnP devices on the customer LAN (Figure 4.14). Use scenarios include (but are not limited to):

- Remote enumeration of UPnP devices, for LAN topology discovery and detection of resource conflicts (such as IP address collisions).
- Troubleshooting of UPnP device interoperability problems.
- Remote configuration of UPnP devices (such as printers or networking devices).
- Integration of UPnP devices (e.g. lightning controls) in operator-assisted surveillance and automation services.

The UPnP Forum has also published a DCP for a *ManageableDevice* [UPnPForum2010], providing specific embedded management functions for UPnP devices such as log analysis, software and firmware management or network diagnostics, which can also become remotely accessible to the operator using the solution hereby described.

4.3.3 Validation

In this subsection we discuss the proof-of-concept implementation of the UPnP-CWMP bridging subagent, test methodologies and obtained results.

Proof-of-concept subagent, testbed and test plan

The proof-of-concept subagent was implemented in Java, following the interface requirements for the CWMP extensibility framework discussed in Chapter 3. The UPnP control point logic was implemented using the *Cling* [Cling] library.

The agent was evaluated in a testbed (Figure 4.18) reproducing the conditions of a managed broadband environment, including the customer premises LAN and the service provider infrastructure. The CWMP extensible framework used for supporting the UPnP-CWMP subagent allows it to be implemented either as a component embedded on the CWMP managed device or on another LAN device. Nevertheless, the former option was adopted for the sake of simplicity.

To emulate broadband access technology conditions, a transparent *Dummynet* bridge [Carbone2009] was used – its configuration details are discussed on Annex A. The domestic gateway is a Linux system with an embedded CWMP agent and two network interfaces, supporting NAT. A PC used for network traffic measurements (using the *Wireshark* analyzer [CACE2011]), is connected to a port on the Ethernet switch which mirrors traffic coming to/from the *Dummynet* bridge.

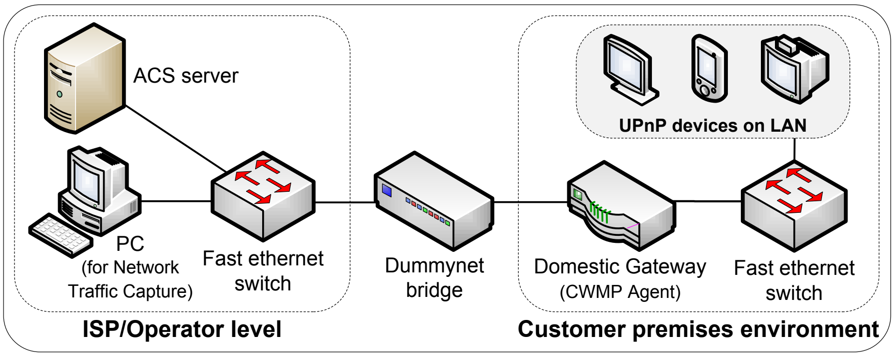

Figure 4.18: Testbed environment

The test plan is focused on evaluating performance indicators for subagent processing latency and protocol overhead for both network payload data and operation latency.

Subagent processing latency is measured as the elapsed time between the discovery of a new UPnP device or service on the customer LAN, the processing of its description and the initialization of its associated CWMP data model entries. For this purpose, 15 test interactions were performed, consisting of the same UPnP device and related services (the *Cling* UPnP media renderer [Cling]) being activated on the customer LAN. Measurements showed an average data model initialization delay for a new UPnP device of 1096ms (611ms standard deviation), which seems adequate.

For protocol overhead and latency tests, measurements were performed for 15 invocations of a service action on an UPnP device (the *BinaryLight* demo device on the *Cling workbench* application [Cling]) within both single (one UPnP invocation per CWMP session) and pipelined (several UPnP invocations per CWMP session) CWMP sessions.

For network traffic, obtained results (Figure 4.19) show a clear difference between native UPnP operations and both single and pipelined CWMP sessions. The difference between the amount of generated traffic on CWMP-embedded and native UPnP operations has to do with the way they are performed: on CWMP a set of tasks, such as loading CWMP data model parameters corresponding to UPnP input action parameters, have to be performed prior to action invocation, generating additional traffic in comparison to native UPnP SOAP operations. Also, the difference between single and pipelined session results has to do with CWMP session setup overhead, which is only performed once for the latter case. Still, the CWMP data overhead penalty is not significant, from a practical point of view.

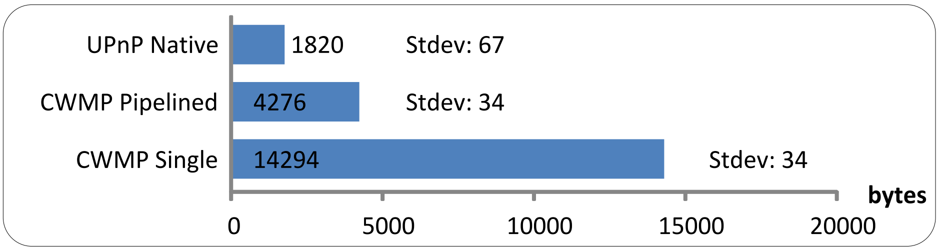

Figure 4.19: Average generated network traffic

For latency tests, the *Dummynet* bridge was configured to emulate typical broadband access network conditions (Table 4.2). Native 100Mb/s Ethernet LAN is used for reference.

Table 4.2: Broadband test reference scenarios

	Nominal bandwidth (b/s) (Down/Up)		Effective bandwidth (b/s) (Down/Up)		RTT Latency (ms)	Pkt. Loss
ADSL	8M	512K	6.68M	427.5K	20	0.1%
	24M	1M	20.04M	835K	20	0.1%
GPON	100M	10M	93M	9.3M	5	0%
LAN	100M	100M	100M	100M	<1	0%

In terms of total latency overhead (Figure 4.20), it becomes clear that CWMP processing is consistently responsible for the biggest share of total elapsed time, with pipelined operations having a smaller penalty than single-session.

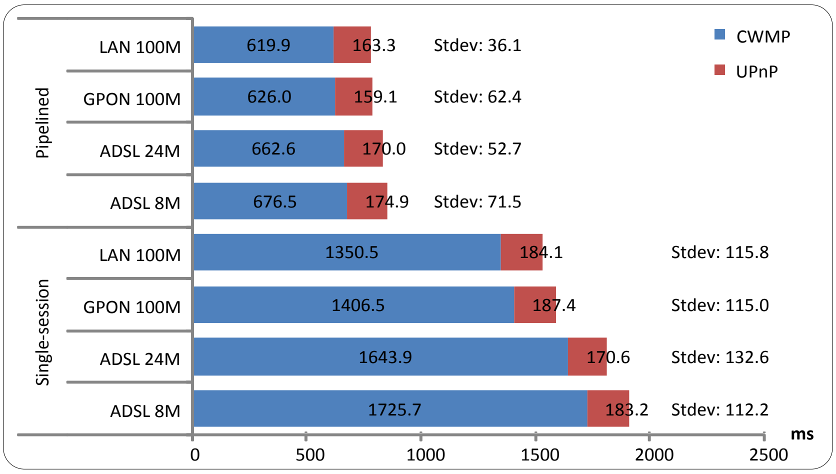

Figure 4.20: CWMP-UPnP protocol latency overhead

The number of operations involved, together with generated network traffic influence overall latency, which improves with better access network conditions. Once again, and similarly to what happened in terms of protocol data overhead, overall latency values reflect the different nature of the two protocols.

4.4 Related work

In this section we discuss previous work related with the two presented use cases (Section 4.2 and Section 4.3).

To the best of our knowledge, the mechanism we propose in Section 4.2 is the first solution that bridges the commonplace TFTP provisioning mechanisms found on VoIP telephones with a CWMP management framework, thus implementing a proxied management model. In the specific context of CWMP-based management of VoIP services, there are already a few commercial implementations of the TR-104 data model [Thomson2008] – for the sake of completeness it should be also mentioned that TR-122 [TR-122] mandates TR-104 support for CWMP-compliant standalone ATA devices. However, in general most TR-104 implementations are used for ATA interfaces embedded on the home gateway (to connect analog legacy phones). This is due to both technical and commercial reasons – as already discussed, currently available VoIP COTS devices lack CWMP support. Additionally, in a transition phase it is easier and cheaper to allow the customer to keep its old analog phone. However, in the long turn, excluding standalone SIP phones precludes many attractive voice service functionalities that can only be supported by dedicated SIP devices.

Also, to the extent of our knowledge, the solution presented in Section 4.3 is the first proposal for full TR-157-compliant UPnP integration, taking advantage of standardized data model components which enumerate discovered UPnP instances on the domestic LAN and their relations to reduce complexity of control extensions. This has the benefit of reducing data model parameter and structure redundancy, while simplifying agent development and improving standards compliance.

Prior to CWMP UPnP TR-157 component standardization, some work had been done to integrate generic device discovery and control features of UPnP on CWMP environments, using CWMP-UPnP dedicated bridges on domestic gateways [Delphinanto2009] [Hillen2009][Nikolaidis2007]. This work has been partially superseded by TR-157, which enabled the inclusion of information about UPnP devices discovered on the domestic LAN on the CWMP data model. However, UPnP control remains absent on TR-157.

For the sake of completeness, it should be also mentioned that TR-064 [TR-064] (later extended by TR-133 [TR-111]) describes a UPnP-based LAN-side CPE configuration procedure. However, it only defines an UPnP DCP interface for device configuration performed from inside the customer premises (which may be used for initial device setup performed from a PC, for instance), without any kind of protocol interoperability features.

Outside the realm of CWMP, the OSGi [OSGi] management framework also incorporated support for UPnP device control points and generic devices [OSGi2011] – an example implementation can be found on the Apache Felix project [Apache2011].

Work has also been done to extend the reach of UPnP devices beyond the domestic LAN [Martinez2009], with the proposal of a specific DCP for such purposes [UPnPForum2009]. While this is implicitly one of the basic premises behind this chapter, these solutions are focused on bridging disparate devices or domestic networks together over broadband access networks for device interoperability in an unmanaged way, not focusing on operator-driven management operations.

Once again, it is important to realize that the notion of protocol proxying based on CWMP-compliant gateways (a feature which is explored by both use cases presented on this chapter) is clearly latent in a number of specifications and updates recently produced by the Broadband Forum, especially in TR-106 and TR-104, which explicitly foresee the existence of proxying management support on the CPE data model. It is known that proxying extensions for CWMP are planned and will be made available in the 4th amendment of TR-069 (MR-239) specification (mainly targeting UPnP integration). However, no documentation has been disclosed to now.

4.5 Conclusion

In this chapter we have dealt with the problem of operator-centric management of customer LAN devices. While CWMP does exist for such purpose, in practice its scope it does not cover all classes of devices commonly found on home LANs, therefore creating the need for operators to deal with management of foreign devices. This chapter discussed two use cases which made use of the framework presented on Chapter 3:

- **Management and provisioning of legacy VoIP SIP devices**, by using integration between CWMP based management and classic TFTP-based provisioning mechanisms, this way allowing for initial device provisioning and runtime configuration updates. This integration is achieved by turning the domestic gateway into a management proxy for legacy VoIP devices on the LAN, using the TR-104 data model together with a system for generating provisioning files through the combination of CWMP data model properties with generic templates. The proposed integration framework can also support remote monitoring and general management of SIP devices. A number of application scenarios for this integration framework were also discussed, including two distinct approaches for provisioning (operator-based and self-service). Overall, feasibility of the proposed approach was demonstrated, paving the way for future extensions of the concept, such as firmware upgrades through TR-157 component management methods.

- **Management of UPnP devices**, by using an extension to the CWMP protocol to allow operators to access and manage UPnP devices on the customer's LAN.

  The advantages of this solution are manifold. First, it implements TR-157 data model profiles for discovered UPnP devices and services, also making use of this information to reduce redundancy and simplify the data model extension for control purposes. Second, it was designed with scalability and performance in mind, using caching mechanisms to reduce the time needed to generate the data model entries for new devices and also allowing for fine grained UPnP event integration. Third, it uses neutral and generic UPnP-CWMP mapping strategies allowing integration of new UPnP devices, regardless of the existence of non-standard extensions. Finally, it is flexible enough to be decoupled from the CWMP managed host and deployed on another device on the domestic LAN, enabling innovative management topologies.

  In terms of security, this solution does not intend nor attempts to solve the security issues of the UPnP protocol on LAN environments, as it does not affect its operation inside this scope. Nevertheless, the bridging mechanism hereby presented does ensure that an adequate level of security is provided for all management operations which are performed outside the context of the customer premises LAN, encapsulated and secured by the CWMP protocol.

Those use cases have therefore demonstrated the feasibility of operator-based multi-protocol, device and service management in the customer premises LAN. The proposed approach has the benefit of complete integration within existing CWMP management infrastructures, while providing mediation mechanisms which enable interaction with foreign devices and device management APIs, embedded within CWMP operations.

5. Managed service delivery frameworks

Modern broadband access networks act as enablers for a whole new array of services and opportunities both for operators and third-parties. From the operators' perspective, added-value services are complementary to its broadband access offers, being instrumental to sustain service margins and compensate for the decreasing revenue from traditional services (like PSTN voice). In this context, a managed service paradigm in which the operator has control over service delivery mechanisms makes sense. The same applies to scenarios where third-party providers use the ISP infrastructure as a service carrier (with shared revenue models).

The main purpose of this chapter is to explore the concept of managed services while proposing and exploring new and innovative service paradigms to enrich the already existing offer. As such, two specific operator-managed service delivery frameworks are proposed:

- A managed framework that enables internet service providers (or associated third-parties) to provide multimedia content to DLNA-compliant devices located inside the domestic networks of their customers (such as PCs and generic media players) in a seamless manner
- An operator-managed storage service which explores the existing complementarity between traditional/local network storage appliances and emerging cloud-storage services.

This chapter is structured as follows. Section 5.1 discusses the problems surrounding service delivery for home networks. Section 5.2 presents a managed service for content delivery to the residential LAN using DLNA. Section 5.3 presents and explores a managed hybrid storage service for domestic users. Section 5.4 analyzes related work and section 5.5 concludes this chapter.

5.1 Introduction

Chapter 2 already discussed with detail some of the problems of service delivery on broadband access networks. While avoiding, to the possible extent, to take a stance in the Net Neutrality debate, the fact is that telecommunication operators are frequently treated as data pipe suppliers by service providers. While, in some cases, operators might host specific service provider gateways on their infrastructure to enhance service quality, usually they do not provide quality of service mechanisms with a scope reaching as far as the customer premises. Third-party service traffic is a burden which most operators have to deal with without any compensation – the same burden that is responsible for a considerable share of the pressure to perform constant infrastructure upgrades in order to keep up with customers' demands.

Instead of facing this situation as a threat, it might be considered an opportunity for operators to get involved, by providing end-to-end management services for third-party service providers, or even offering new services of its own. This a consequence of the fact that operators, as pipe suppliers, have complete end-to-end control of the infrastructure that is used to pass service traffic, going from the border gateway router to the access network and, consequently, the residential gateway (and even further, considering the proposals already discussed in Chapters 3 and 4).

By making use of their already existing management infrastructure and through service-level agreements, operators could be able to go further than the current practice of hosting specific gateways from service providers (as in content delivery networks) to the point of offering end-to-end service-oriented management capabilities to ensure better performance (for instance offering QoS mechanisms customized across all its infrastructure up to the residential gateway configuration)

This chapter proposes two service delivery frameworks that can demonstrate both concepts of operator single sourced services and of operator-mediated third-party services:

- A content delivery service for DLNA devices and ISP or third-party provided content. This service turns the residential gateway in a managed UPnP AV/DLNA (Universal Plug and Play/Digital Living Network Alliance) media server, which can be dynamically updated by the broadband operator using CWMP extensions specifically designed for this purpose. This framework enables the domestic/residential gateway to become a mediator for both operator-provided and Internet media content, provided through UPnP services visible inside the domestic LAN. This architecture uses plugins to abstract each service, making them independent of the domestic gateway platform and allowing ISPs to easily add new media services while better coping with protocol updates.
- A hybrid managed storage service which makes use of domestic/residential gateways as storage service hubs. This solution offers a balance between cost, reliability and accessibility by eliminating the need for a dedicated appliance inside the customer premises and transforming the residential gateway in a storage hub, which has its own local storage capabilities and is kept synchronized with a virtualized storage container on the operator or third-party storage service infrastructure, providing redundancy and reliability by replicating data to a virtual container located outside the customer premises. This solution relieves users from the (sometimes daunting) task of configuring and managing their own storage devices, while taking advantage of the fact that the domestic gateway is a device which is permanently powered on (especially in triple-play environments) in order to provide connectivity services for the LAN, thus eliminating the need for a separate appliance.

For both cases, service management interfaces and operational management mechanisms are provided the management framework that was discussed on chapter 3.

5.2 Delivery of internet media services to the home environment

The proliferation of interconnected media devices inside domestic households (such as media players or similar appliances), whose function is to provide access to contents such as video or music has contributed to transform the domestic Local Area Network (LAN), making it evolve from its former role supporting file, printer and Internet sharing for PCs to a private media distribution infrastructure. Specifically, devices that combine traditional media with new functionalities, such as "smart" TV sets, Game Consoles or Blu-ray players are becoming more common in domestic LANs.

Nevertheless, some of those appliances are frequently designed to operate exclusively inside LAN boundaries, accessing local media repositories located on PCs or other devices. Even with some devices becoming increasingly capable, with feature sets supporting other kinds of media services available from external sources, like Internet media (e.g., *Youtube* [Youtube], *Flickr* [Flickr]) or content providers (e.g., *Netflix* [Netflix], *Hulu* [Hulu]) a simple protocol or Application Program Interface (API) change can render them useless until a firmware update is available – which may never happen, depending on the support status of the device.

Also, as broadband operators strive to enhance their service offerings, with some even becoming content providers, the separation between the LAN and the operator media distribution infrastructures becomes an issue which some providers are solving by adding new devices (supporting specific protocols and APIs for media delivery) to an already very heterogeneous environment. As a result, the domestic LAN becomes cluttered with several media rendering and control devices that support different protocols and services which, by their turn, might become useless as their end-of-life support status is reached and updates cease.

As such, it would be desirable to provide broadband media services in a seamless way, integrated with existing protocols and media delivery frameworks, such as UPnP AV [UPnPForum2008] (Universal Plug and Play Audio Video) and DLNA [DLNA] (Digital Living Network Alliance), designed from the ground up for such purposes. However, the DLNA specifications and their core UPnP AV functionalities were designed for use on LAN environments, being unable to properly operate across the LAN boundaries, over broadband access networks.

In this section we propose a managed architecture that makes use of CWMP and UPnP/DLNA framework components to enable a residential gateway to become a Media Server for contents remotely provided and delivered over broadband networks. This architecture enables the domestic gateway to become a mediator for both operator-provided and Internet media contents provided as UPnP resources visible inside the domestic LAN and managed using a CWMP protocol extension designed for this purpose.

Also, the adoption of a neutral and portable UPnP architecture that uses plugins to abstract each service makes the UPnP layer independent of the domestic gateway platform (hardware and software), allowing operators to easily add new media services while better coping with protocol updates.

5.2.1 Delivering operator-provided media services using UPnP AV and CWMP

The idea of delivering media contents over broadband networks using the UPnP AV/DLNA frameworks has the main advantage of allowing for seamless and secure distribution to a wide range of devices which already support those protocols, without the need for pushing another protocol stack or specific device into the domestic LAN for such purpose. In this context, it would be desirable that the UPnP AV/DLNA framework could natively operate across access network boundaries, which is not possible. As previously mentioned, the UPnP framework and related protocols were designed for use on small, domestic LANs. In such environments, the burden of management protocol traffic and scalability are secondary issues due to available bandwidth resources and the relatively reduced number of devices involved. For example, the use of the Simple Service Discovery Protocol (SSDP) over Multicast UDP for UPnP device discovery has two serious shortcomings that are incompatible with operation on WAN environments: the use of multicast UDP (which is frequently blocked on edge routers) and the sheer number of UDP messages involved in the device discovery process (which are sent to overcome the unreliability of UDP), creating a traffic burden.

Also, UPnP does not contemplate any kind of security measures for protecting its eventing and control mechanisms from malicious abuse, a vital issue in WAN environments. As such, and despite the fact that the UPnP Forum has published a Device Control Protocol specification for bridging UPnP devices with an UPnP network across the Internet [UPnPForum2009] (using an unmanaged method with a limited application scope), the use of the UPnP framework on WAN environments is neither secure nor desirable. Consequently, this also applies to DLNA devices.

As such, we propose a solution that overcomes these limitations to provide managed media delivery using CWMP and UPnP AV/DLNA, based on a managed architecture in which the domestic gateway abstracts media content provided by sources outside the LAN using a DLNA-compliant embedded UPnP AV media server (see Figure 5.1).

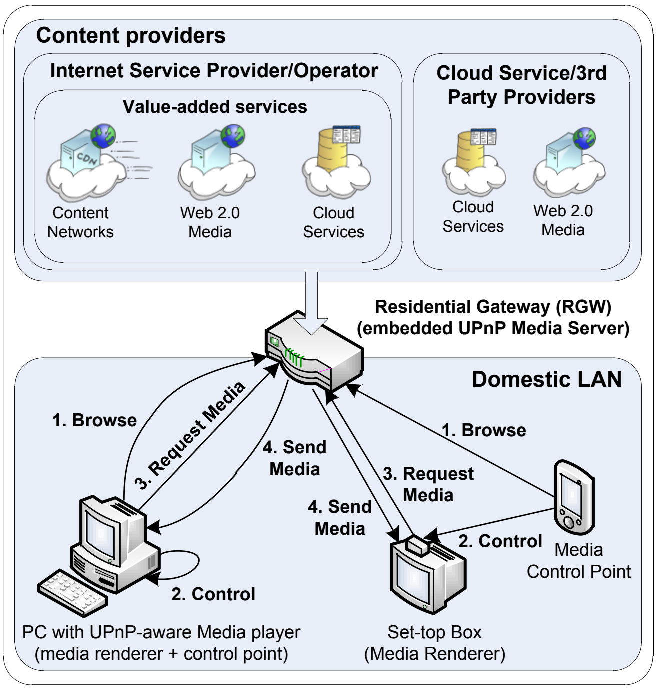

Figure 5.1: Using UPnP AV services to deliver media content over access networks

This approach solves the aforementioned problems, extending the reach of the UPnP AV/DLNA device ecosystem beyond the domestic LAN premises, allowing providers (operators and/or third-parties) to deliver media content seamlessly. For clients, it means that existing DLNA/UPnP AV devices can be used to access a new range of contents and services which are external to the domestic network.

The media server embedded on the domestic gateway is based on a modular architecture that makes use of plugins to deliver a wide range of contents and services available outside the LAN environment – abstracted as media items residing on a regular UPnP AV media server which is advertised inside the domestic LAN (see Figure 5.2).

More than one instance of a specific backend type may be available, each one configured to provide content from a different source, identified and advertised on the UPnP Network as different UPnP Media Servers.

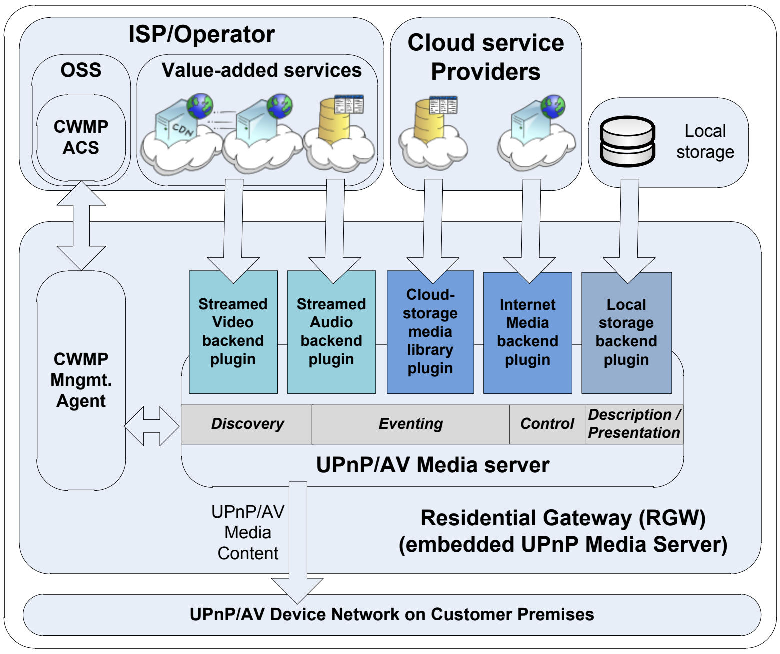
ISP/Operator
OSS
CWMP ACS
Value-added services
CDN
Cloud service Providers
Local storage
CWMP Mngmt. Agent
Streamed Video backend plugin
Streamed Audio backend plugin
Cloud-storage media library plugin
Internet Media backend plugin
Local storage backend plugin
Discovery
Eventing
Control
Description / Presentation
UPnP/AV Media server
UPnP/AV Media Content
Residential Gateway (RGW) (embedded UPnP Media Server)
UPnP/AV Device Network on Customer Premises

Figure 5.2: UPnP AV modular media server embedded on the home gateway

Once again, the CWMP protocol was chosen to provide the service management interface, bringing together the benefits of a mature, tried and tested technology while leveraging existing CWMP-compliant Operation Support Systems (OSS) from operators.

In this scenario, CWMP is used to manage the UPnP AV media server embedded on the gateway, allowing the operator to customize which specific backends are available and the configurations of each active instance. For this purpose, a CWMP bridging subagent (build using the chapter 3 management framework toolset) was developed with the purpose of mapping the configuration of UPnP AV media server and its associated plugins into the CWMP data model of the domestic gateway, so that the operator ACS can remotely access and configure all media server properties using CWMP. CWMP service objects and attributes are mapped into the RGW data model using TR-106 compliant vendor-specific extensions.

Once again, the master-subagent topology was adopted. In our specific case, the subagent resides on the domestic gateway itself (Figure 5.3) and acts as a bridge which maps the configuration of the UPnP AV media server, associated plugins and existing instances on its CWMP data model.

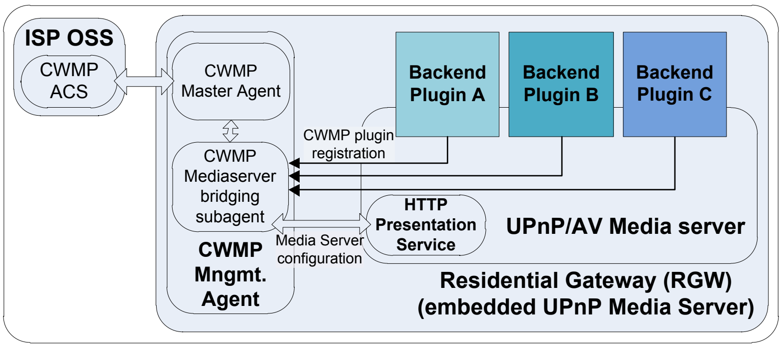

Figure 5.3: CWMP bridging subagent general operation model

Upon initialization of the UPnP media server, all available plugins are enumerated and initialized, one by one. Each plugin registers its configuration parameters on the CWMP bridging agent, which updates the CWMP data model for the UPnP media service with the new parameters.

This solution allows the ACS to remotely access and configure all existing media server parameters and related backend plugins by using the CWMP API. When the ACS creates a new instance of a backend service or modifies any existing media server backend associated object or parameter, the CWMP subagent on the home gateway will update the media server configuration.

CWMP data model extension mechanism

Generically, when the CWMP bridging subagent registers on the master agent it must become associated with the CWMP Data Model objects and parameters for which it is responsible. After a registration attempt passes validation the subagent is ready to receive requests.

Standard UPnP AV/DLNA capabilities embedded on the residential gateway are declared using the proper TR-157 data model parameters under the UPnP object, albeit only for strict compliance. Information and properties of the UPnP AV media server and its plugins are embedded on the home gateway CWMP data model through the use of TR-106 data model service objects and parameters. This way, the CWMP subagent maps the UPnP AV media server configuration namespace data into the CWMP data model. Specifically, TR-106 enables the possibility of embedding such information by using specific service object instances which contain the correspondent data model information.

The *Services* object on the CWMP data model (see Figure 5.4) of the domestic gateway will have a vendor-specific entry for configuration of the UPnP AV media server framework (embedded on *the X_<OUI>_UPnPMediaServer* object, where the Organizational Unique Identifier (OUI) is defined as 000000 for test purposes). Inside, each backend instance will be represented by a specific object, named *<Backend Name>.{instance}*.

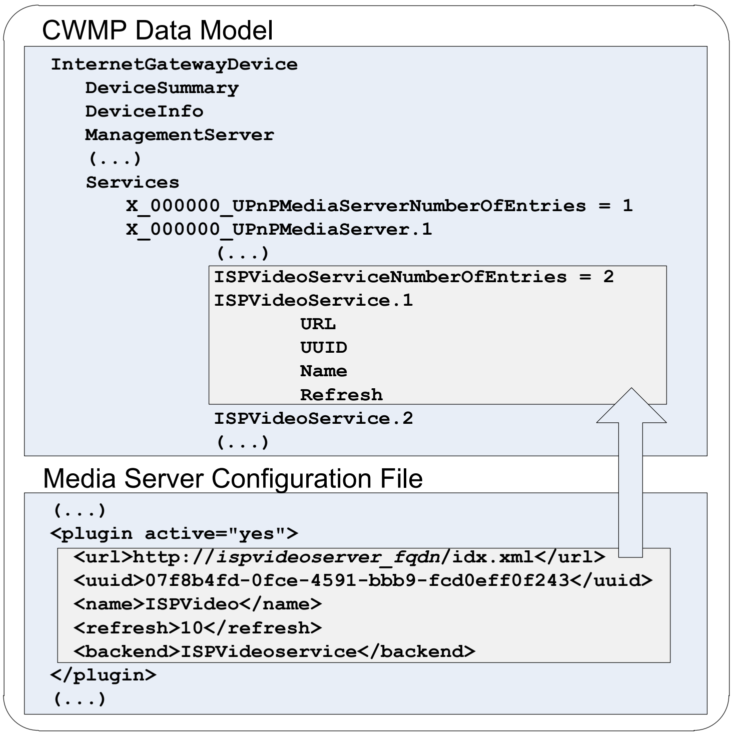

CWMP Data Model

```
InternetGatewayDevice
   DeviceSummary
   DeviceInfo
   ManagementServer
   (...)
   Services
      X_000000_UPnPMediaServerNumberOfEntries = 1
      X_000000_UPnPMediaServer.1
             (...)
             ISPVideoServiceNumberOfEntries = 2
             ISPVideoService.1
                    URL
                    UUID
                    Name
                    Refresh
             ISPVideoService.2
             (...)
```

Media Server Configuration File

```
(...)
<plugin active="yes">
  <url>http://ispvideoserver_fqdn/idx.xml</url>
  <uuid>07f8b4fd-0fce-4591-bbb9-fcd0eff0f243</uuid>
  <name>ISPVideo</name>
  <refresh>10</refresh>
  <backend>ISPVideoservice</backend>
</plugin>
(...)
```

Figure 5.4: CWMP TR-106 parameters for the embedded media server

Integration of the UPnP media server data model into CWMP uses static mapping. Specific CWMP parameters and objects inside each *<Backend Name>* and respective plugin instances are mapped into the media server configuration attributes. In this case the ACS user only manipulates conventional CWMP parameters.

Specific mapping rules are stored on a registry accessed by the CWMP Subagent, which contains all mapping rules and extensions which are used to update the CWMP supported data model of the domestic gateway (accessible to the ACS).

Backend plugin and instance management

When the media server first starts, all backend plugins must be enumerated and initialized. Upon initialization, each one has to register itself on the CWMP subagent, using an existing mapping profile or one provided by the plugin itself. Following, the subagent will initialize all backend plugin instances, creating the associated CWMP objects and parameters. New plugin instances can be created using the CWMP *AddObject* method to create a new object instance of the supported data model object associated with a specific backend, for which the subagent will translate the information, mapped to the native media server configuration.

The plugin management solution is designed to support different mapping profile versions for different versions of the same backend plugin, which can coexist peacefully (using *the <plugin name><version>* naming convention to identify each plugin version and mapping profile entry).

5.2.2 Application scenarios

In this subsection we will present several application scenarios for the managed media delivery framework previously presented, based on specific use cases which reflect actual implementation possibilities for managed services.

Operator-controlled media distribution

Using the proposed architecture, the residential gateway can become a distribution hub for media content provided by the operator or specific content providers with whom it has established Service-Level Agreements (SLA). This content becomes available to all UPnP AV-DLNA media renderer devices inside the customer premises (Figure 5.5).

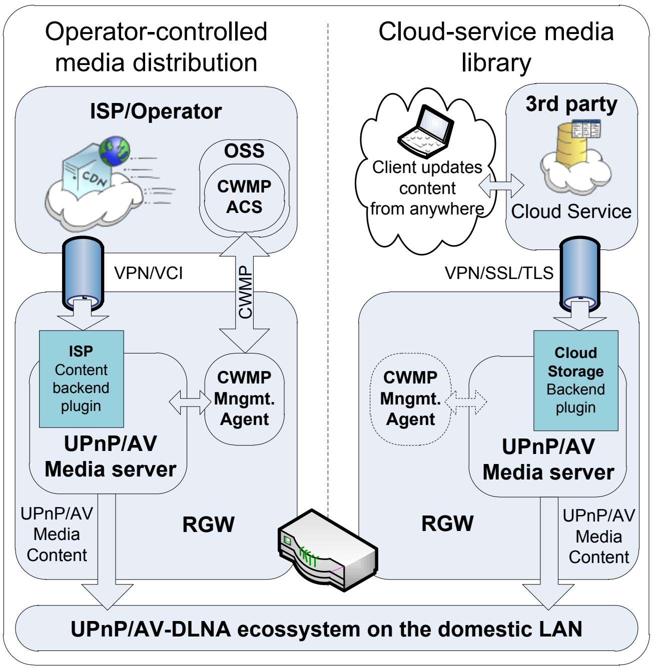
Operator-controlled media distribution
Cloud-service media library
ISP/Operator
CDN
OSS
CWMP ACS
Client updates content from anywhere
3rd party
Cloud Service
VPN/VCI
CWMP
VPN/SSL/TLS
ISP Content backend plugin
CWMP Mngmt. Agent
UPnP/AV Media server
CWMP Mngmt. Agent
Cloud Storage Backend plugin
UPnP/AV Media server
UPnP/AV Media Content
RGW
RGW
UPnP/AV Media Content
UPnP/AV-DLNA ecossystem on the domestic LAN

Figure 5.5: Application scenarios

Once the customer subscribes the media service, two operations might take place.

First, if the media server plugin that allows access to the service is outdate or not deployed on the domestic gateway, the operator will upload the most recent version. Once the update is detected, the embedded media server will reload its configuration, enabling the new plugin which, by its turn, will update its own configuration parameters on the CWMP management subsystem. The recently-introduced TR-157 CWMP software module management methods (discussed on Section 2) are used for this purpose.

Second, the operator will remotely configure and enable the new plugin, from the ACS, by using CWMP. Additional security and QoS transport features might be ensured by using a VPN (Virtual Private Network) or VCI (Virtual Circuit Interface) tunnel, which can also be configured using CWMP standardized data model extensions. Those mechanisms not only provide traffic separation and increased security, but also a simple way to differentiate traffic at the gateway level (e.g. for QoS policy enforcing purposes).

This approach allows the operator to deliver media content to all kinds of UPnP AV or DLNA-enabled media playing devices inside the customer premises in a controlled way, without directly exposing its backend infrastructure.

Cloud service integrated DLNA media libraries

Similarly to the previous application scenario, a plugin might be developed to access a storage container on a cloud service, exposing its media content to the customer premises though the media server (Fig. 8). Content can be updated from anywhere, becoming available to the UPnP AV-DLNA ecosystem on the domestic network of the service subscriber. For instance, a user travelling abroad might sync its photo collection on the cloud storage or media library service (using his notebook, from any place with an available Internet connection) to share with his family back at home, which will be able to watch the pictures on a Smart TV or other capable UPnP AV-DLNA media renderer. This feature can be available directly as a customizable feature to the end-user or as a value-added service from the service provider (with whom the third-party provider might have an established SLA).

When compared with a typical cloud-based storage service, such as Dropbox [Dropbox], the proposed scenario has several advantages:

- It does not require a dedicated client application – any DLNA device is able to access and reproduce media content without any modifications.
- Dropbox client implementations are inadequate for integration within media devices (which are embedded devices with modest resources, in comparison to PCs). Generally, Dropbox clients for smaller embedded devices (such as smartphones) have several limitations in comparison with the full-fledged PC version (e.g. the absence of local caching).
- Last but not least, Dropbox is simply not adequate for real-time media streaming, as it is basically a file-oriented synchronization service, with best-effort performance and no support for QoS. Instead, the proposed solution builds on the operator OSS infrastructure to provide adequate performance.

Seamless Internet media services access for DLNA devices

The same solution might be used to provide Internet media service content (from a service like Flickr, for instance) for devices unable to directly access it. This solution offers a simple way for end-users to retrieve and play on their DLNA-compliant devices media content from Internet services. Similarly to the previous application scenario, this feature might be available as an end-user customizable feature or as a value-added service from the service provider (with whom the third-party provider might have an established SLA).

The prototype we have developed (see Section 5) already bundles proof-of-concept plugins for some of these services (like Flickr and Apple Movie Trailers, among others), fully integrated with the managed solution.

When compared with unmanaged solutions for Internet-to-DLNA streaming, such as Skifta [Qualcomm2011] and TVersity [TVersity], the proposed solution offers several advantages:

- Skifta and TVersity require specific applications to be installed in the devices that mediate content transfer. While Skifta was designed for smartphones (requiring a PC client application for specific use cases, such as remote DLNA library access), TVersity requires installation in a PC. Our proposal is based on a platform-neutral solution for domestic gateways, not requiring further devices or software installations.
- Skifta and TVersity devices (PCs, smartphones) with the installed applications are required to be permanently available (i.e. powered on and accessible). In contrast, the proposed solution relies on domestic gateways, a device which is already permanently powered on.
- Skifta and TVersity do not offer any QoS support. Our solution builds on the operator OSS infrastructure to provide adequate performance.

5.2.3 Validation

In this subsection we will delve into validation work, focusing first on specific details of the implemented prototype, its specific components and their integration, together with a proof-of-concept managed media delivery service used for validation purposes. Next, the validation of the proposed framework will be addressed, including test methodologies and discussion of obtained results.

Prototype implementation details

The UPnP AV media server

The media server component of the domestic gateway is one of the basic building blocks of the proposed architecture, which has to fulfill three fundamental requisites: to be open-source, portable and modular. Along the selection process to find an adequate UPnP framework for this purpose, the *BriSA* [Brisa], *Coherence* [Coherence] and *Cling* [Cling] projects stood out as potential candidates, since they complied with the three basic premises. However, among the three, the Coherence UPnP framework was found to be better suited: first, Coherence is entirely developed in Python [Python] , a language which was intended from its conception to be platform-neutral and portable. Also, Python is available for a wide range of operating systems and environments, including embedded systems; second, Coherence is based on a modular architecture that uses plugins to extend its functionality. Also, plugins can be easily added or updated by simply importing their source file into the plugin library directory; third, Coherence supports the most complete set of specifications, including UPnP v1, v2 and DLNA v1.5.

The Coherence framework architecture is divided into three parts: (see Figure 5.6).

- **The core** provides the building blocks for UPnP operation: an SSDP server; an MSEARCH [Goland1999] client; server and client for HTTP/SOAP requests; server and client for GENA eventing. It also supports interfaces for implemented UPnP services and devices.
- **Device implementations** link to the core using their respective UPnP service interfaces. For example, a media server must, at least, bind to the *ContentDirectory* and *ConnectionManager* service interfaces when registering to the core.
- **The backends** are the interfacing points of the UPnP framework with the exterior. Backends are plugins that implement the methods to deal with external sources such as filesystems, hardware or user interfaces. Backend plugins can be used to integrate new media sources on the media server, such as local storage devices, Internet media services (such as *Flickr* or *Youtube*), cloud services, among others.

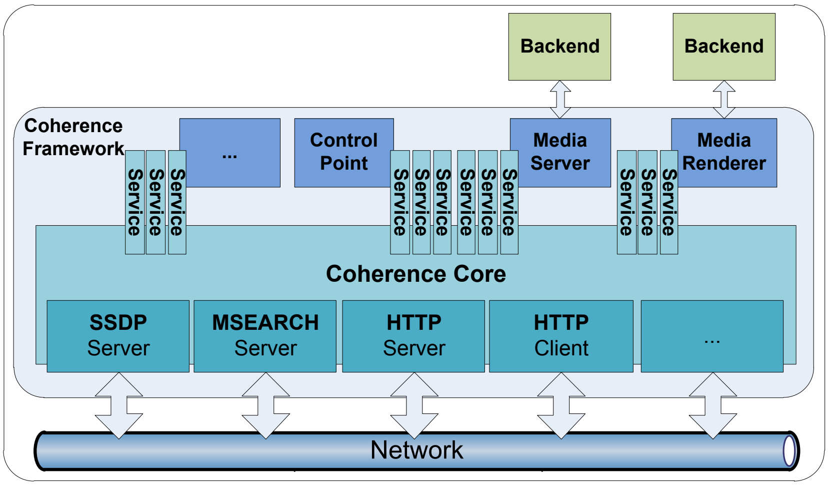

Figure 5.6: Coherence framework architecture (adapted from [Coherence2011])

When an UPnP device implementation registers with the core, it declares which service interfaces it uses, attaching hooks that map service actions to its backend. After that, the core generates a Universally Unique Identifier (UUID) [ISO11578] for the device and creates the data structures that support it and its UPnP device and service description. From them on, the core will take care of the "frontend" UPnP work, sending periodic SSDP announcements, managing GENA subscriptions for state variables and interfacing SOAP service action calls for the respective backends. From a user point of view, each backend plugin instance is seen and announced on the LAN as a separate UPnP AV device with its own UUID.

Plugin and instance management via CWMP

One of the limitations of the Coherence framework has to do with the lack of manageability mechanisms, a weakness shared by all evaluated UPnP frameworks. Coherence does not provide interfaces for integration on operator-scale management infrastructures, relying instead on an XML-formatted text file to store its configuration. To solve this problem, we developed a bridging mechanism to enable configuration of the embedded media server using CWMP.

Also, to integrate the configuration of the media server with the CWMP management framework, the code of the Coherence plugin subsystem had to be adapted so that each plugin has the ability to register on the CWMP management subsystem. Since each plugin is responsible for registering itself on the CWMP bridging agent, which updates the CWMP data model for the UPnP media service, the Coherence plugin subsystem was also modified to integrate the CWMP registration mechanism, so that each backend plugin is responsible for declaring its own configuration options to the CWMP bridging logic.

When the media server first starts (see Figure 5.7), the directory containing all the backend plugins will be scanned (1) and initialized (2). Upon initialization, each plugin will attempt to register its availability on the CWMP subagent (3), which will check the mapping registry (4) for an existing mapping profile (which will be requested from the plugin, if not found).

The subagent will scan the media server configuration file for all backend plugin instances (5), creating the associated CWMP objects and parameters under the *X_000000_UPnPMediaServer* service object (Figure 5.7). Also, when a new instance of a specific backend plugin is created (using the CWMP *AddObject* method to create a new object instance of the supported data model object associated with a specific backend), the subagent will translate the information, which will be written to the media server configuration file (5a), refreshing the media server configuration (5b), and creating a device UUID for the new instance.

For each operating backend instance, the CWMP subagent will use the UPnP presentation HTTP interface provided by the media server to query and manage runtime operation (6), accessing an URL of the form *http://localhost:<configuredport>/<media server instance device*

UUID>/config (see Figure 5.7). All media server configuration changes are written to persistent storage, maintained on a local file, allowing the service to use the last known configuration, if the management infrastructure temporarily fails.

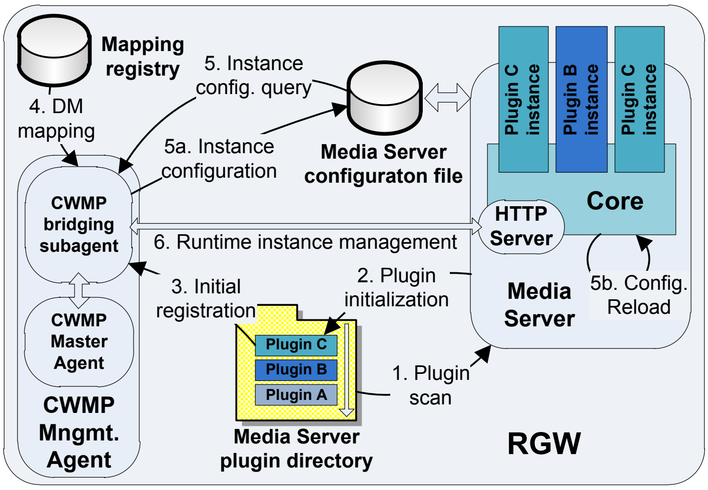

Figure 5.7: Initialization of the embedded media server

Plugin updates follow similar rules (see Figure 5.8). When a new plugin is uploaded to the system (using software module management methods provided by TR-157), it will be dropped on a specific directory, which is monitored using the *inotify* API [LDP] (1). This will generate an event (2) forcing the media server to reload its configuration (3) and initialize the new plugin (4), which will proceed with its initial registration (5).

Upon the initial registration process, the CWMP subagent will require the plugin to additionally declare its mapping profile to the mapping registry (6), which also will be used to update the CWMP supported data model (7). To create an instance of the newly available backend plugin, the ACS invokes the CWMP method *AddObject* to create an object of the type declared for the specific plugin on the supported data model namespace of the domestic gateway (available to the ACS) and fill its configuration properties.

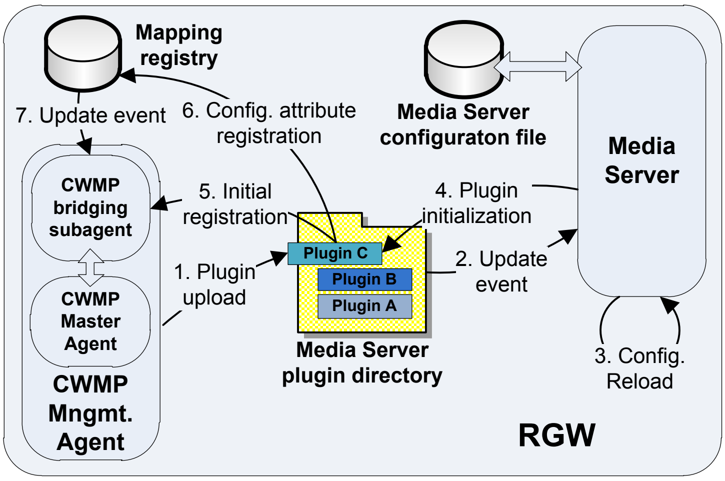

Figure 5.8: Backend plugin update

Proof-of-concept plugins - ISPVideo

One of the implemented proof-of-concept plugins (called *ISPVideo*) was designed to provide managed access to MPEG video streams available on an operator video content delivery infrastructure to DLNA-UPnP AV devices (see Figure 5.9).

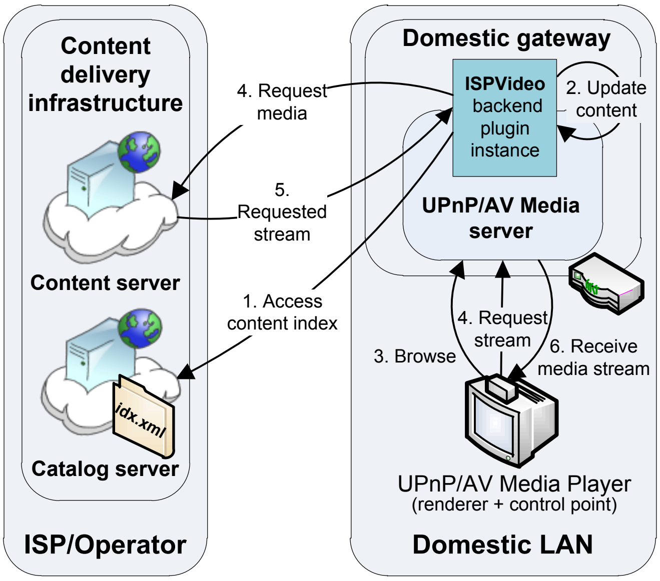

Figure 5.9: The ISPVideo backend plugin (CWMP management infrastructure omitted for clarity)

The *ISPVideo* backend plugin uses a simple approach: the content/service provider publishes and periodically updates an XML document (*idx.xml*) available via HTTP on the service catalog server (1), describing available content, its location and metadata (such as title, duration, URLs for locating thumbnails of movie posters, and such). This information is used by the plugin to build the UPnP AV media server catalog, instantiated on the DIDL-Lite [UPnPForum2011a] scheme structures for its *ContentDirectory* service (2).

From the domestic user standpoint, the DLNA-compliant device on the LAN is able to browse (3) the entries available on the selected media server instance, in the same manner as it would in a conventional UPnP AV media server. Once a specific entry is selected (4), the plugin will follow the provided URL source for the selected media stream which will be downloaded from the content server (4, 5) and proxied for the DLNA device on the LAN via HTTP transport (6).

It should be mentioned that the adopted media server framework has support for transcoding, implemented using the *gstreamer* [Gstreamer] software. However this feature was judged to be overwhelming to the modest processing capabilities of domestic gateways. As such, content providers might provide several streams encoded using different codecs (i.e. MPEG-2 and MPEG-4) to match the capabilities of the UPnP AV/DLNA CE devices.

Testbed and test plan

A prototype home gateway platform was implemented, integrating the CWMP management support, media server and proof-of-concept plugins. For validation purposes, this prototype was deployed on a testbed (Figure 5.10), specifically conceived to reproduce all relevant aspects of the proposed managed media service architecture.

On the operator level, there is an CWMP ACS and a webserver, which will hold both roles of content and catalog server for the *ISPVideo* service.

On the customer premises side, the domestic LAN is equipped with the prototype home gateway hereby described together with some commercial off-the-shelf DLNA/UPnP AV renderers and control points (used for compliance testing) and a PC equipped with a UPnP-enabled media player (for performance testing). The prototype home gateway is a Linux system with two network interfaces, a CWMP agent (with the bridging subagent) and the media server, were the *ISPVideo* plugin is deployed. It is based on a platform with modest computing capabilities (using cheap, off-the-shelf components and a single-core Atom CPU clocked at 1.6GHz paired with 512MB of RAM) to mimic as close as possible the constraints of typical embedded systems used on commercial hardware to build home gateways. In order to emulate the conditions and restrictions imposed by the broadband access link, a transparent *Dummynet* emulator interconnects both environments (ISP/Operator and Customer premises).

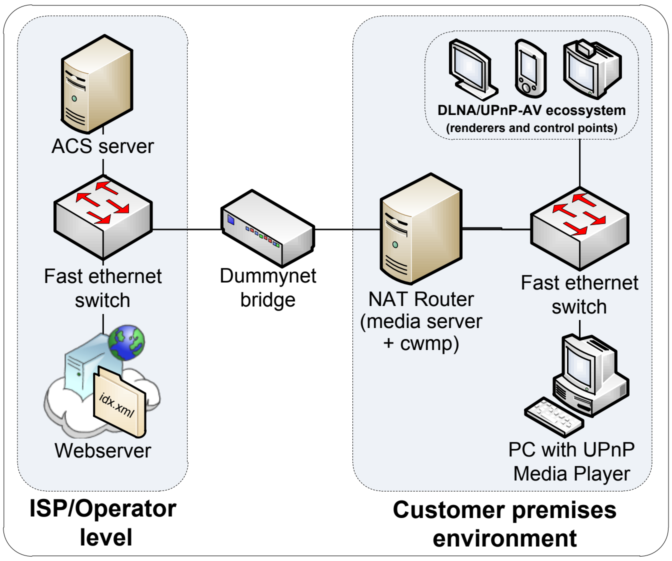

Figure 5.10: Testbed architecture

This testbed hosted a series of measurements, following a test methodology devised to validate three different aspects of the proposed architecture:

- **Functional validation.** Functional interoperability tests were performed using the prototype implementation in conjunction with commercial off-the-shelf DLNA devices, with the purpose of verifying its compatibility requirements.
- **Performance.** Operation latency was measured on different scenarios, in order to assess the performance of the solution on different broadband access network technologies. The *Dummynet* system (Figure 5.10) was used to emulate typical access network scenarios (see Table 5.1), taking into account the technology and protocol overhead for specific access network technologies (see Annex A for a more detailed discussion).
- **Resource usage.** Overall resource usage (CPU and memory) of the architecture components embedded on the prototype router was measured in order to verify its ability to operate on more constrained embedded platforms typically used on commercial routers.

Table 5.1: Broadband test reference scenarios

	Nominal bandwidth (Mb/s) (Down/Up)		Effective bandwidth (Mb/s) (Down/Up)		RTT Latency (ms)	Pkt. Loss
ADSL	12	1	10.02	0.835	20	0.1%
	20	1	16.7	0.835	20	0.1%
GPON	30	3	27.9	2.79	5	0%
	100	10	93	9.3	5	0%
LAN	100	100	100	100	<1	0%

Because the test plan intended to evaluate the performance of the prototype using a stream within the capabilities of all emulated scenarios, a 480p MPEG-4 encoded video stream with a bitrate of 2957Kb/s and 9:56 of duration was used for all performance and resource test runs. This stream is publicly available [Blender2008].

Results and discussion

Functional validation was performed using a set of UPnP AV-compliant and DLNA-certified devices and software components, deployed on the LAN side of the testbed network. For each one, the same basic interoperability test was performed consisting of browsing, selecting and playing a set of audio and video files. The availability and usability of basic play controls were also tested. The results are shown on Table 5.2.

Table 5.2: List of devices and software solutions used for functional validation

Product	Type	Platform	Status
Microsoft Xbox 360	Game Console	Dedicated hw	Works
Popcorn Hour A-110	Media Player	Dedicated hw	Works
Storex Slimbox	Media Player	Dedicated hw	Works
EGREAT R1B	Media Player	Dedicated hw	Works
Sony PlayStation 3	Game Console	Dedicated hw	Works
XBMC	Media Center Software	Windows	Works (issues with content list refresh)
Foobar 2000	Audio Player	Windows	Works
Windows Media Player 12	Media Player	Windows	Works
Windows Media Center	Media Center Software	Windows	Works
Samsung LCD TV LE32C550	LCD TV with DLNA support	Dedicated hw	Works (basic controls – no seek)
LG BD570	Blu-ray Player	Dedicated hw	Not working
Western Digital WD TV Live	Media Player	Dedicated hw	Works
ASUS O!Play HD2	Media Player	Dedicated hw	Works
AndroMote	Remote Control Software	Android	Works

Apart from one device (LG BD 570), all tests performed satisfactorily, with all devices and software components providing the expected minimum functionality, albeit with minor problems in a few cases – which can be tracked down to incomplete or incompatible DLNA/UPnP AV implementations. Measured results are shown in Figure 5.11.

For performance assessment, runs of 20 tests were executed for each emulated network access scenario, to estimate the average latency of media player operations. Each test consisted of a mix of basic control commands (Play, Pause, Stop, Seek to 50%) starting with a content browsing operation, on a container with 1000 entries. For this purpose, the *Cling Workbench* [Cling2010] tool was used on the test PC to control the XBMC media player and measure the latency of media control operations.

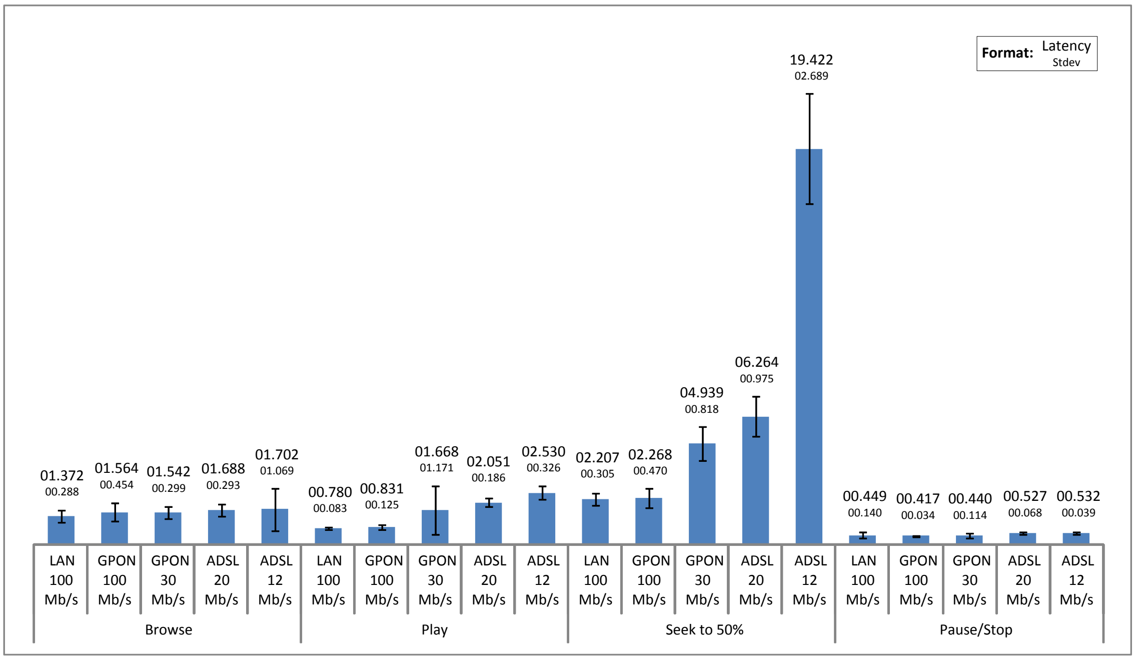

Figure 5.11: Media player operation performance (values in seconds)

Special care was taken to obtain a reliable measure of user-interactivity, instead of simple operation latency measurements. For instance, instead of simply measuring the elapsed time between an UPnP (standard) *AVTransport.Play* method invocation and the reception of the corresponding event confirming the change of state on the media player, it was ensured that no subsequent buffering-related immediate pauses were detected (which would mean that from the user standpoint, the video stream was not playing yet).

Results show the average operation latency to be adequate for normal usage, with some specific operations, like Play and Seek, to be more prone to influence from access network technology conditions, as is the case for the 19.4 second pause averaged on seek operations on the emulated 12Mb/s ADSL scenario, due to buffering. This can be related to the proof-of-concept plugin implementation, which acts as a simple connectivity proxy – most UPnP AV/DLNA media players are designed to operate on LAN environments and do not implement adequate buffering techniques needed for streaming content across the access network. This problem could be handled with an improved plugin implementation, adopting pre-buffering techniques together with support of HTTP range requests and partial responses (both in the operator media server and plugin) to improve seek performance, albeit with a possible increase on resource footprint usage. In addition, this matter becomes less of an issue as the bandwidth on broadband connections increases and latency decreases. Replacing HTTP by RTP for content streaming between the operator infrastructure and the domestic gateway is another alternative to consider. RDP is more efficient for media delivery purposes and its adoption would be especially adequate for operator-specific media content service plugins not requiring interaction with third-party content providers.

The resource footprint of the prototype components was measured, for every access network scenario and also for conventional 100Mb/s Fast Ethernet, for reference purposes. Five series of tests were run, each one composed by five individual runs of the same video streamed to the XBMC (XBox Media Center) media player [XBMC] on the test PC. Data collection was performed using *top* [LeFevre2011] (an utility which allows for system resource monitoring in real-time) with a sampling rate of 2 seconds, averaged over a period of 10 minutes, to cover the entire duration of the video stream. For each series of 5 runs, the corresponding average and standard deviation were calculated (see Figure 5.12).

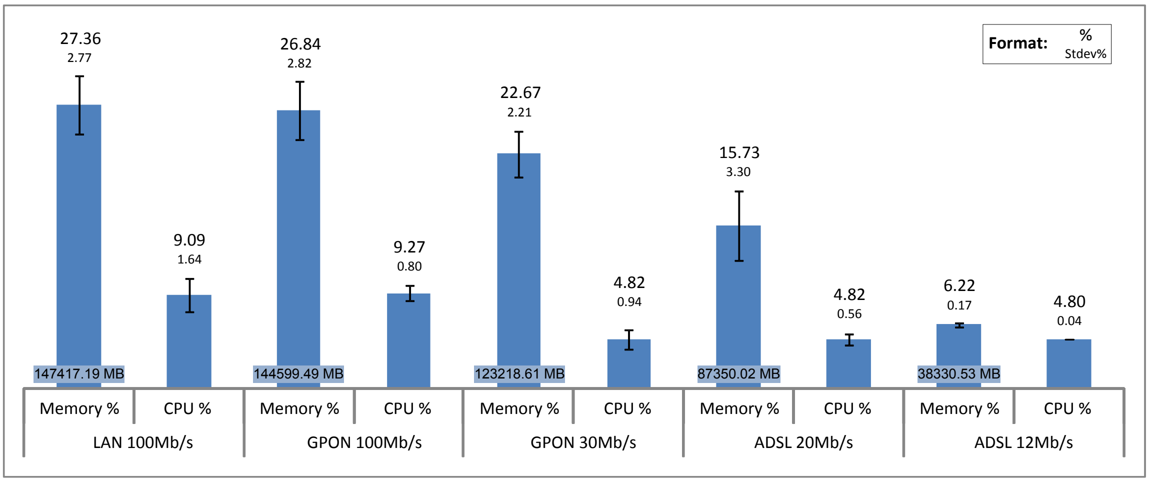

Figure 5.12: Average resource consumption on domestic gateway

Results show the resource consumption of the media gateway to be compatible with the capabilities of the current generation of domestic gateway devices. Further investigation showed the DLNA media server framework to be responsible for using the largest share of memory and processing power, consistently above 70% of the total component resource usage, both for CPU and memory, reaching peak values of 85% in some situations. This is partly due to the nature of the Python environment, but it can also be traced to the implementation of the proof-of-concept plugin which does not include any kind of throttling or adaptive streaming mechanisms (such as dynamic buffering) which could help mitigate this situation.

5.3 Managed hybrid storage for home and SOHO environments

The second service delivery framework proposed in this chapter relates with storage services.

Storage, in the form of simple PC-to-PC file sharing, was one of the most important services behind the emergence of home and SOHO LANs. As users' needs evolved, basic workgroup file sharing began to show its inherent weaknesses in terms of reliability, convenience and availability: standard PCs did not offer redundancy or failover mechanisms. Soon, more sophisticated networked storage devices began to appear on the LAN, such as Windows Home Servers or dedicated storage appliances. By using dedicated client applications or network protocols such as CIFS/SMB [Microsoft2011b], FTP [RFC959] or even HTTP(S) [RFC2616], these devices provide consolidated storage, frequently combined with features like redundancy, file versioning, media streaming or printer sharing. However, the majority of those devices were designed to operate inside the customer premises network, being inaccessible from the outside – this is partly because most supported access protocols (as is the case with CIFS/SMB) are not adequate for usage outside the LAN scope, for security and design reasons. Some users have solved this problem with the help of Virtual Private Network (VPN) technologies but most lack the technical expertise to do so. Also, some firewalls block VPN traffic, making remote storage access difficult or impossible.

More recently, cloud storage services emerged to fill the accessibility gap left open by traditional dedicated storage devices, providing a virtualized storage container where users can store their files and which is accessible from almost anywhere with an Internet connection (and not just inside the user premises, as it happens with the majority of dedicated storage appliances). This is the case for services such as *Dropbox* [Dropbox], Amazon's *Cloud Drive* [Amazon2010], Apple's *iCloud* [Apple2011c] or Ubuntu *One* [Canonical2010], which use dedicated client applications or browser interfaces, to enable users to access their cloud file shares in a transparent and convenient way. However, when it comes to comparing dedicated storage appliances and cloud services, from the home and SOHO users' perspective, each one has their specific advantages and drawbacks:

- **Accessibility**. Cloud storage services are accessible from almost any place with an Internet connection, while dedicated appliances are frequently inaccessible outside the customer premises LAN scope. However, most cloud storage services tend to use proprietary interfaces and APIs which require the installation of software components, barring unsupported devices from accessing them.
- **The Capacity/Speed ratio** is another problem. While both dedicated appliances and cloud services might provide high storage capacities, the latter are seriously hampered by available broadband access network speeds and specific usage requirements, which may render them unpractical, above certain capacity thresholds. Inside the customer premises LAN, cloud storage is no match for dedicated devices when it comes to capacity and speed.
- **Cost** is yet another question. Dedicated appliances come with an ownership cost which encompasses electric power consumption, maintenance and other software/hardware issues. Cloud services, on the other hand, charge a regular fee to relieve the final user of such burden. On the long term, the cost advantage might favour cloud storage services, but individual user mileage may vary.
- **Reliability**. While cloud storage providers have their own location-independent internal distributed redundancy and backup mechanisms, which are transparent to the end-user, dedicated appliances are different. By being self-contained, dedicated appliances are more prone to data loss in the case of a catastrophic event (such as a fire or theft), in which case device-specific redundancy techniques (such as RAID volumes) may become completely useless.

- **Manageability** is another aspect which is orthogonal to both storage paradigms – albeit for different reasons. For dedicated network storage, it has to do with device configuration complexity, making end-users more or less prone to neglecting critical parameters that compromise basic accessibility and reliability by simple ignorance and contributing, up to some point, to nullify the benefits of such devices. For cloud storage services, management issues exist at a different level: while operators might host specific service provider gateways on their infrastructure to enhance service quality, usually they do not provide quality of service mechanisms with a scope reaching as far as the customer premises. Cloud storage traffic is a burden which most operators have to deal with without any compensation.

While both approaches have their share of advantages and drawbacks in terms of accessibility, capacity, cost and reliability, it can be observed that, in several aspects, the two paradigms do not overlap, rather complementing one another. This situation hints at the possibility of developing a hybrid solution capable of integrating the benefits of both in order to provide a balanced storage framework, also prefiguring an opportunity for operators to get involved, by providing their own managed storage service, using their own storage backend infrastructure or resourcing to a third-party storage service providers. Akin to triple-play services delivered over broadband access networks, this service can be provisioned and managed by operators using standard protocols, therefore leveraging their already existing management infrastructure.

In line with this reasoning, this section presents an operator-managed hybrid storage solution that offers a balance between cost, reliability and accessibility by eliminating the need for a dedicated appliance inside the customer premises and transforming the residential gateway in a storage hub, which has its own local storage capabilities and is kept synchronized with a virtualized storage container on the operator storage service infrastructure. The advantages are manifold: first, it relieves users from the (sometimes daunting) task of configuring and managing their own storage devices; second, it takes advantage of the fact that the domestic gateway is a device which is permanently powered on (especially in triple-play environments) in order to provide connectivity services for the LAN, therefore eliminating the need for a separate appliance; and third, it provides redundancy and reliability by replicating data to a virtual container located outside the customer premises, on the storage provider infrastructure.

Service management interfaces and eventing mechanisms are ensured by CWMP. While this section focuses on the service-specific aspects of the storage service, it should be mentioned that the adoption of CWMP enables the operator to make use of its existing infrastructure to provide a complete and secure end-to-end management solution.

5.3.1 An architecture for using residential gateways as storage hubs

The proposed operation model aims to combine the benefits of cloud storage with dedicated appliances, namely in what refers to accessibility, speed, capacity and reliability, while eliminating the cost associated to ownership of dedicated appliances. This is achieved by transforming the domestic gateway in a storage hub. In fact, it turns out to be the ideal device for this purpose because of the following aspects:

- **Hardware capabilities.** It is frequent for domestic gateways to have one or more Universal Serial Bus (USB) ports, which can be used to attach external hard disks. In terms of processor and memory capabilities, most modern embedded systems (a category in which domestic routers fit) are computationally able to deal with the requirements posed by this solution.
- **Location.** Domestic gateways are placed on the boundaries between the access network and the customer premises LAN, holding a privileged position mediating all traffic to and from both domains. By having a network interface natively on the access network, the domestic router is a privileged management entry point for the LAN behind it. Also, since the role of the customer LAN firewall is normally performed by the home router, service deployment and troubleshooting are simplified.

- **Full-time operation.** Domestic gateways are supposed to be operating on a full-time basis, without interruption. This is even more true on triple-play environments where they are critical to telephony and television services.
- **Efficiency.** Turning domestic gateways into storage hubs eliminates the need for dedicated storage service devices, therefore reducing ownership costs (nevertheless, the proposed solution could be based on a dedicated local device, if desired).

These properties turn residential gateways into the perfect candidates to provide a hybrid storage service for customer premises LANs, whose operation will be next described.

Operation model

Each RGW has its own local storage capabilities in the form of flash memory (embedded flash, USB pen or removable memory card) or hard disk (embedded or external). The storage device contents is made accessible to the interior of the customer premises LAN using protocols such as SMB/CIFS, FTP or HTTP(s) (albeit other protocols might be used) and synchronized with the cloud storage container on the operator infrastructure (see Figure 5.13).

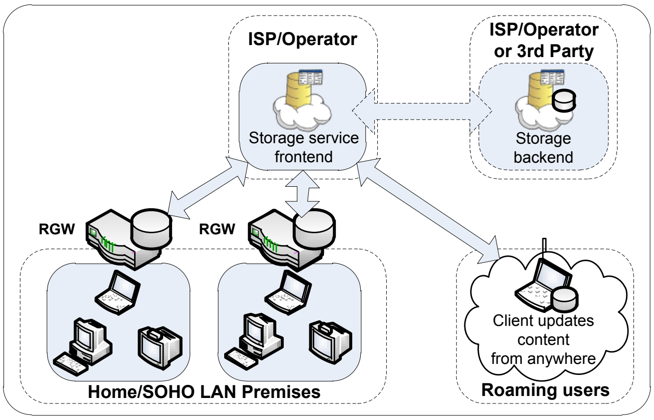

Figure 5.13: Hybrid storage service operation model

Every RGW associated with the same storage service subscription is allowed to synchronize with the central repository. Roaming users might also use native client applications to access its virtualized storage container while on the road (still, this section focuses on the CPE-level implementation and does not detail this specific use case). Each CPE has the ability to detect changes on its local storage space and start a synchronism operation with the main service repository which will be synced with the most recent version. Subsequently, the service will trigger a global sync operation with all CPEs which are registered to the same user, to update their local storage from the main storage service container.

While the operator is in charge of the storage service frontend and its management, the backend might be provided by a 3rd party storage provider, like Amazon's S3 service [Varia2010] (with whom a Service Level Agreement might be established) or single-sourced by the operator itself.

This solution has the ability to stream content to the customer premises network at the same speed than a dedicated NAS device would do (provided the content is already synchronized and replicated on the local storage), while supporting disconnected operation, remaining accessible even in the event of an access network failure. Also, it does not require that devices have specific support for using the service, since data is accessed using standard protocols (such as SMB/CIFS). Redundancy is provided in a cost-effective way, since data is replicated to the service main storage repository (whose backend might be assured by a third-party cloud storage service), instead of only relying on self-contained device data replication methods, such as RAID (which, nevertheless, can be used together with this solution to achieve extra redundancy), making this solution adequate for both multi-branch SOHO infrastructures and home users with a service subscription for a single household. Versioning is also supported, therefore allowing the user to recover previous versions of a specific file.

5.3.2 Management and eventing

The hybrid storage service proposed on this section uses CWMP to handle service management and storage synchronization events, thus allowing operators to use their already existing CWMP-based operations support systems (OSS) to manage the service (Figure 5.14).

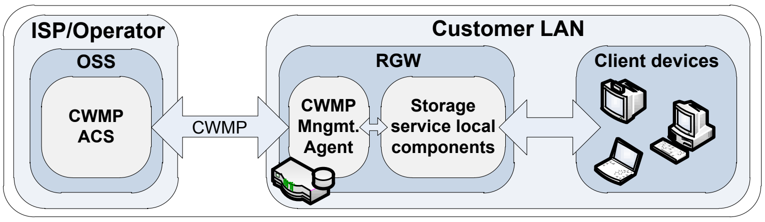

Figure 5.14: Usage scope of CWMP in the proposed architecture

The CWMP framework encompasses a standardized data model for NAS devices, specified by TR-140 [TR-140], which allows for management of networked storage services on CWMP-compliant devices, being used as the referential data model for the proposed storage service. It should be however mentioned that, contrary to the service hereby proposed, practically none of the currently available COTS NAS devices support CWMP.

Eventing interfaces for synchronism between local storage spaces and the operator are also provided by the CWMP protocol, in a secure fashion. CWMP has a native eventing mechanism which allows for the definition of notification attributes on parameters, either by default or explicitly, using the *SetParameterAttributes* CWMP method. For instance, when the state of a given parameter with the Active Notification attribute enabled changes, the CWMP device will schedule the execution of an Inform method on the ACS to notify it. This mechanism is used to pass synchronization events between the subscriber devices providing the storage service.

Integration architecture

The proposed solution is based on the addition, at the home gateway, of a specific integration component. This component acts as module of the gateway CWMP agent, interacting with the ISP ACS (by means of CWMP operations) and with the local storage service components. In order to preserve the coherence of the CWMP data model, this specialized component is responsible for its own CWMP data model extensions (in accordance with TR-106 guidelines and the TR-140 standard), which are mapped and integrated on the global data model of the CWMP Agent.

In general terms, the proposed architecture integrates together a CWMP frontend interface with the storage service frontend. This extension is implemented as a module which maps the configuration properties of the storage service into the CWMP data model of a domestic gateway, so that the operator ACS can remotely manage them using CWMP. The module will then configure the components needed for enabling the storage service, which are embedded on the domestic gateway and are exposed on the LAN-side interface, accordingly with the CWMP parameterization (see Figure 5.15).

On the operator side, the ACS implements the CWMP management server, which interfaces with the storage service frontend by means of a middleware layer based on a message-oriented queuing system for passing events back and forth between both components.

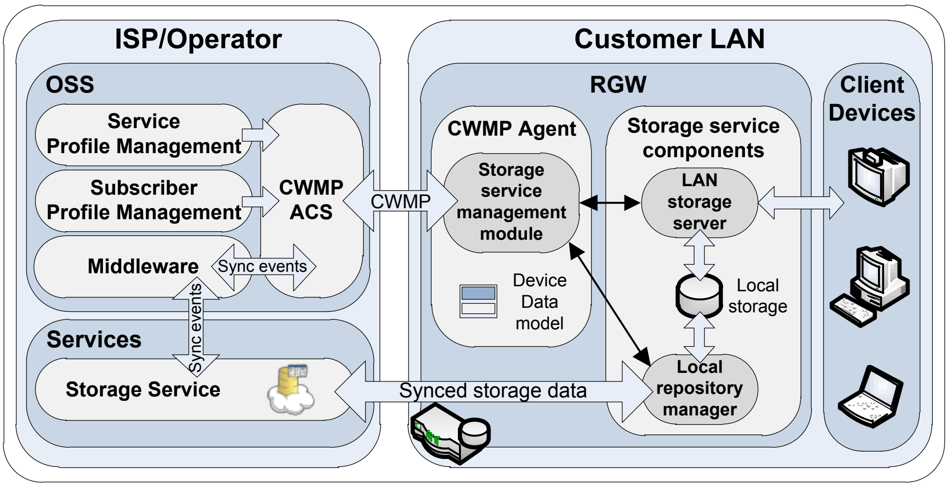

Figure 5.15: Integration architecture for the hybrid storage service

Integration with the CWMP data model

Considering the CWMP Data Model objects and parameters, the integration module will be responsible for two different branches of the supported CPE data model: the standard branch for NAS devices (as defined by TR-140) and another set of TR-106 compliant extensions for service provisioning and management purposes (see Figure 5.16).

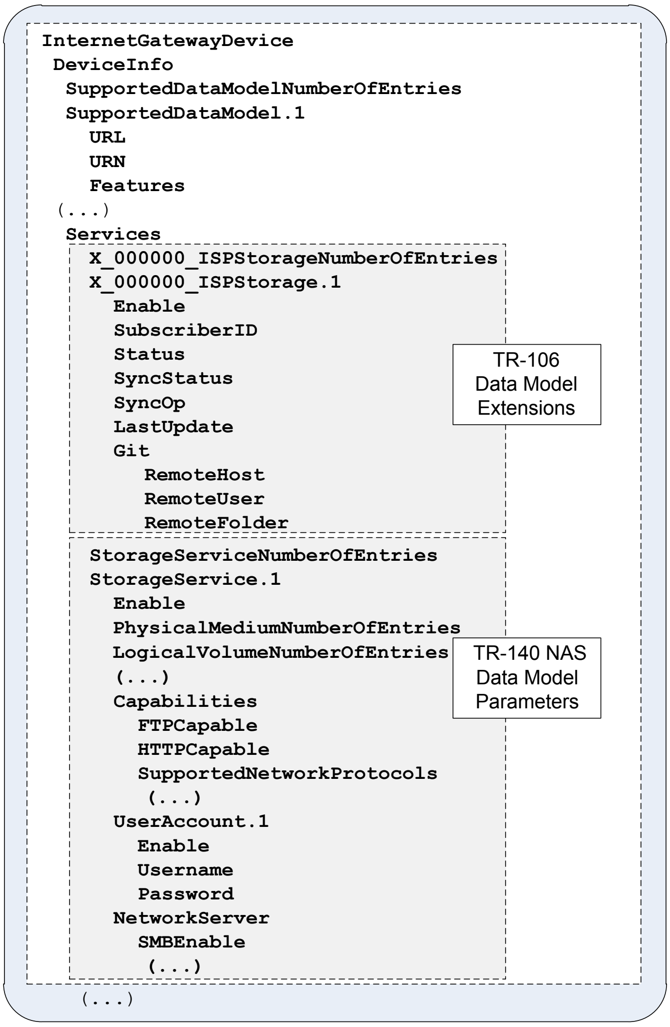

```
InternetGatewayDevice
 DeviceInfo
  SupportedDataModelNumberOfEntries
  SupportedDataModel.1
    URL
    URN
    Features
 (...)
  Services
    X_000000_ISPStorageNumberOfEntries
    X_000000_ISPStorage.1
      Enable
      SubscriberID
      Status
      SyncStatus
      SyncOp
      LastUpdate
      Git
         RemoteHost
         RemoteUser
         RemoteFolder
    StorageServiceNumberOfEntries
    StorageService.1
      Enable
      PhysicalMediumNumberOfEntries
      LogicalVolumeNumberOfEntries
      (...)
      Capabilities
        FTPCapable
        HTTPCapable
        SupportedNetworkProtocols
         (...)
      UserAccount.1
        Enable
        Username
        Password
      NetworkServer
        SMBEnable
         (...)
   (...)
```

TR-106 Data Model Extensions
TR-140 NAS Data Model Parameters

Figure 5.16: CWMP data model integration

Both TR-140 and TR-106 explicitly allow for embedding the information of managed services within the TR-098 data model of home routers (even though, to the best of our knowledge, our proposal is the first to make effective use of this possibility to embed storage features on a domestic gateway), under the *Services* root object.

The data model extensions are organized using the *X_000000_ISPStorage* object on the CWMP data model (with an Organizational Unique Identifier (OUI) defined as *000000* for test purposes). These data model extensions are also declared in a *SupportedDataModel* table entry (of the *DeviceInfo* root object), which points to an URL containing their XML description hosted on the CPE itself, so that the ACS can gain knowledge about the device data model supported by the managed device. The same happens for the TR-140 standard data model, whose CPE support has to be declared before being instantiated.

The information needed to control the storage service components is embedded on the CWMP data model of the home gateway using an *X_000000_ISPStorage* entry, while specific local scope storage service and device configuration parameters are managed using TR-140 data model structures. The TR-140 *StorageService* entry is used to manage the LAN storage service which will provide access to the repository data from the customer premises LAN. For this purpose the *Baseline:1, NetServer:1, FTPServer:1, HTTPServer:1, UserAccess:1, VolumeConfig:1* and *VolumeThresh:1* profiles are supported, as per TR-140, in order to provide support in the data model for the features of the LAN network server.

Service management using CWMP

Service management can be split in three different phases: the provisioning process, service initialization, and the runtime management of the service.

The **provisioning process** goes as follows (see Figure 5.17): once the CWMP storage service management extension is active, the *X_000000_ISPStorage* and the *StorageService* objects will be instantiated and enabled, for management of the storage service and the associated network storage server component. Then, the ACS will configure the service, which will be followed by its activation by setting of the *X_000000_ISPStorage.1.Enabled* parameter to *True*.

The initialization process goes as follows (see Figure 5.17): on each initialization, the CWMP storage service management module will configure the repository manager and the network storage server using the information which is present on the CWMP data model instances. Once the internal modules are successfully initialized, the CWMP storage management module will change the *X_000000_ISPStorage.1.Status* parameter to *Ready*, (with the Active Notification attribute enabled). When the state of a given parameter with the Active Notification attribute enabled changes, the CWMP device will schedule the execution of an *Inform* method on the ACS to notify it. This will happen with the *Status* parameter once its value changes to *Ready* as the result of the initialization process. For its turn, a middleware layer will pass the information to the Storage Service Frontend, assembling a subscription request on behalf of the CPE. If the CPE is allowed to access the service, a confirmation message will be passed to the ACS, which will change the *Status* parameter to the *Subscribed* state. After each initialization, the repository manager attempts to synchronize its data with the storage service, in order to update its content to the latest version.

Runtime management goes as follows (see Figure 6.17): to set specific parameters on the CWMP data model objects for the storage service (*X_000000_ISPStorage.1* and *StorageService.1*), the ACS will interact with the CWMP agent on the domestic gateway using standard CWMP methods. For each ACS-instructed configuration change, the CWMP Storage Service management module will deal with the related internal modifications.

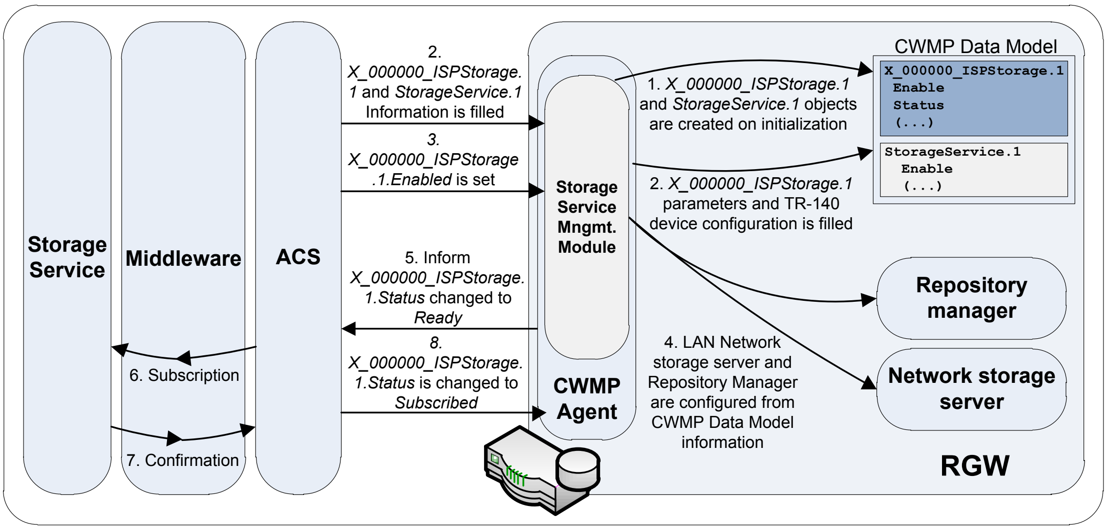
Storage Service
Middleware
ACS
CWMP Data Model
X_000000_ISPStorage.1
Enable
Status
(...)
StorageService.1
Enable
(...)
2. X_000000_ISPStorage.1 and StorageService.1 Information is filled
3. X_000000_ISPStorage.1.Enabled is set
5. Inform X_000000_ISPStorage.1.Status changed to Ready
8. X_000000_ISPStorage.1.Status is changed to Subscribed
6. Subscription
7. Confirmation
Storage Service Mngmt. Module
CWMP Agent
1. X_000000_ISPStorage.1 and StorageService.1 objects are created on initialization
2. X_000000_ISPStorage.1 parameters and TR-140 device configuration is filled
4. LAN Network storage server and Repository Manager are configured from CWMP Data Model information
Repository manager
Network storage server
RGW

Figure 5.17: Service management via CWMP

Eventing using CWMP (synchronization operations)

While the eventing mechanism for signaling synchronism operations between local storage repositories and the operator service could be implemented using a custom signaling method, the fact is that CWMP proved to be adequate for the purpose, since it already offers a secure and appropriate notification mechanism[4]. The middleware provides operation translation and queueing between the Storage service frontend and the ACS, similarly to a decoupling layer which can be implemented in a distributed fashion using a publish-subcribe message queueing system such as Apache ActiveMQ [Apache2010a].

The synchronism mechanism for CPEs is based in a push-pull model. However, for subscribers with a single storage service access point (like regular home users without roaming access, for instance) only push operations occur, since there is only a single CPE involved.

Push operations (see Figure 5.18) occur when the content of the local repository is changed. The user accesses the local storage device by means of the network storage server (for instance, a SMB/CIFS server), being able to modify its contents. A filesystem watcher on the repository manager is able to detect changes and, as soon a file handle is closed, an update operation is pushed into the internal stack (a FIFO – First In, First Out).

Every time a content change is detected, it is pushed into the storage service frontend to update its version – as soon the operation is finished, the CWMP storage service management module is notified, in order to update the contents of the *SyncStatus* parameter (see Figure 5.16), which has Active Notification properties enabled, therefore forcing the CWMP agent to notify the ACS of its change. A middleware layer will then pass this information to the storage service frontend in order to inform of the sync operation completion.

[4] In our case CWMP was used also for eventing/synchronization purposes (in addition to service management) because it was already available and provided all the required functionality (therefore providing a satisfactory solution for rapid prototyping). In our opinion, this option does not compromise the overall interest of the proposed hybrid storage service or the relevance of the mechanisms proposed in Chapter 3 in the provisioning of this service.

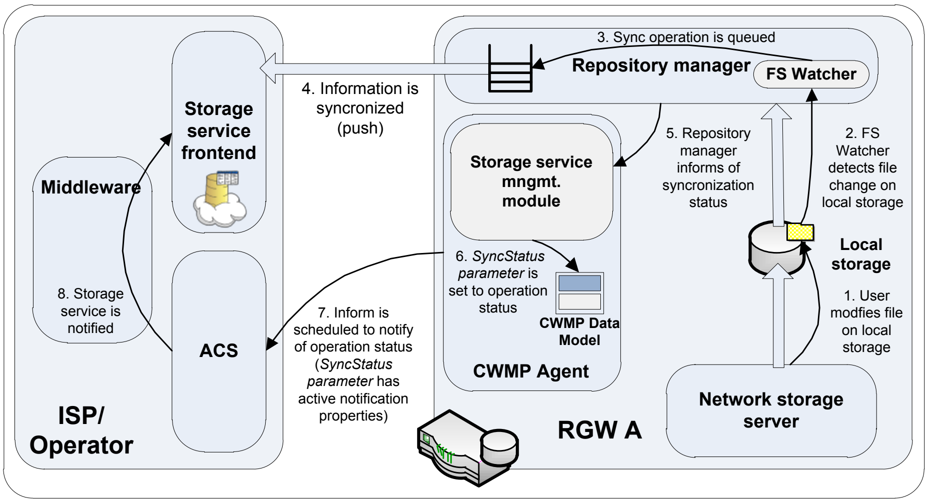
3. Sync operation is queued
Repository manager
FS Watcher
4. Information is syncronized (push)
Storage service frontend
Middleware
8. Storage service is notified
ACS
ISP/ Operator
Storage service mngmt. module
5. Repository manager informs of syncronization status
2. FS Watcher detects file change on local storage
6. SyncStatus parameter is set to operation status
CWMP Data Model
CWMP Agent
7. Inform is scheduled to notify of operation status (SyncStatus parameter has active notification properties)
Local storage
1. User modfies file on local storage
Network storage server
RGW A

Figure 5.18: Push event notification via CWMP

The storage service frontend keeps track of all subscribers. When there is more than one service subscriber entity, a push operation must be issued to all involved subscribers in order to update their local information. **Pull operations** (see Figure 5.19) are issued by the storage service frontend and enqueued by the middleware layer, which will generate an internal operation sequence number (*opsec*) to be appended to each operation. The ACS will be instructed to execute a *SetParameterValues* operation on the *SyncOp* parameter, which will be filled with the string pull <*opseq*>. As a result, the CWMP storage service management module of the CPE will enqueue a pull operation on the local repository manager.

Once the pull operation is completed the repository manager makes the updated information available on the local storage device and will communicate the operation status to the CWMP storage service management module, which will update the *SyncStatus* parameter with the operation result, appended by its *opseq* sequence tag. Since this parameter has *Active Notification* properties, the CWMP agent will notify the ACS of its change (which will pass the information into the middleware layer, therefore dequeuing the operation).

Conflicts are resolved locally by each subscriber agent entity, which has the ability to detect whenever the file on the main service container is more recent than the original version of a local file which was also meanwhile updated. In this case, a collision is registered and the local file is renamed to its original name concatenated with a timestamp and queued for future update on the main storage container.

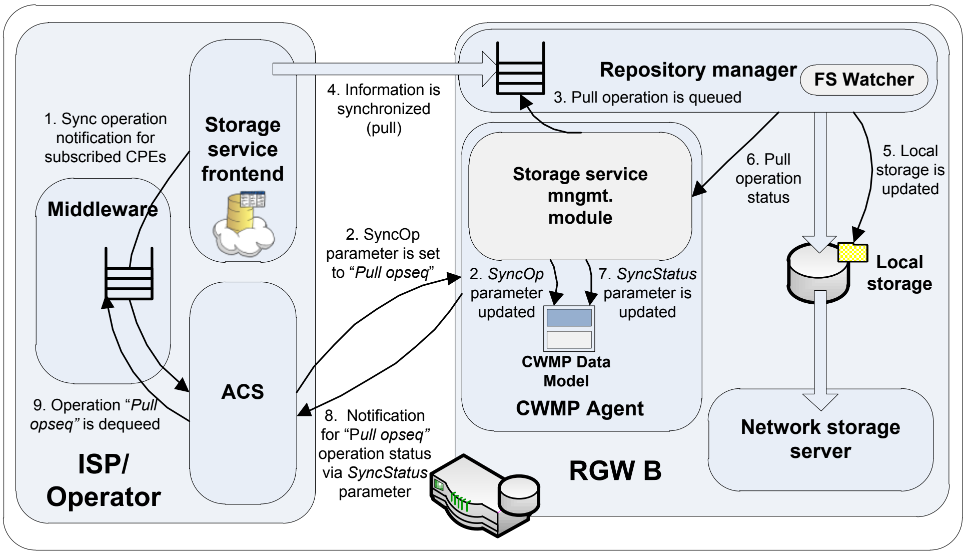
1. Sync operation notification for subscribed CPEs
Storage service frontend
4. Information is synchronized (pull)
Repository manager
FS Watcher
3. Pull operation is queued
Middleware
2. SyncOp parameter is set to "Pull opseq"
Storage service mngmt. module
6. Pull operation status
5. Local storage is updated
Local storage
2. SyncOp parameter updated
7. SyncStatus parameter is updated
CWMP Data Model
CWMP Agent
ACS
9. Operation "Pull opseq" is dequeued
8. Notification for "Pull opseq" operation status via SyncStatus parameter
Network storage server
ISP/ Operator
RGW B

Figure 5.19: Pull event notification via CWMP

5.3.3 Application scenarios

This subsection discusses how the proposed solution can be used by operators to provide a portfolio of storage services. Specifically, three scenarios are presented, addressing different use cases for both SOHO and home users.

Hybrid NAS for home users

Most low-cost dedicated NAS appliances are not redundant by nature, being little more than network-attached disks with a special firmware component. The hybrid storage model hereby presented could be used to add redundancy, replicating data to an external storage container (see Figure 5.20). Since home routers could provide this functionality, an operator could market this service in a converged service bundle.

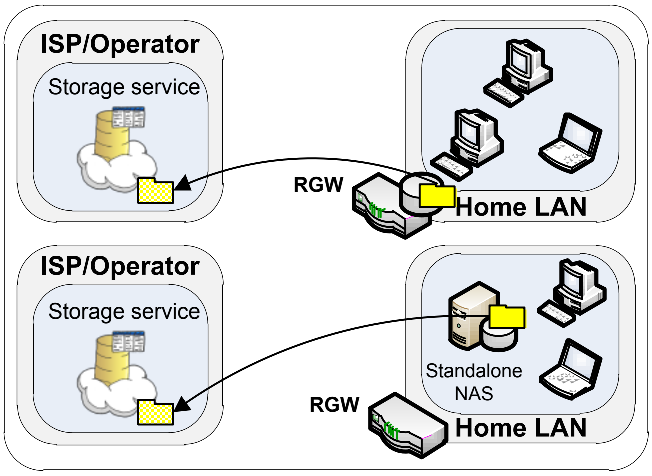

Figure 5.20: Hybrid NAS storage service for home users

However, even dedicated appliances with RAID support could benefit of the storage model hereby proposed, adding an extra redundancy layer (akin to cold site storage) to its native mechanisms, therefore protecting data from catastrophic events. While the proposed framework is targeted towards home routers, it is equally possible to implement in dedicated NAS devices (as depicted in the second half of Figure 5.20).

Content distribution mechanism for home users

This storage service could be used by an operator as a content delivery mechanism. Users could, for instance, buy music from an online store and have it placed on its storage container, being propagated to all devices associated with the subscriber account. Also, operators could use this service to improve support, automatically providing software and firmware updates to devices which the user might have acquired right to the customer premises, in a transparent and convenient way.

Enterprise storage synchronization for multiple branches

An organization with several dispersed branches could benefit from the hybrid storage service hereby described as a means to keep a shared document repository synchronized (and safe) between all workgroup networks. This service could be also used to propagate software updates, making them available at local network speeds, once the related files are propagated (see Figure 5.21).

While not a new idea in conceptual terms, its implementation would otherwise require a considerate amount of hardware and expertise. An operator could market this service for Small-to-Medium enterprises in a service bundle. Being the access provider the same entity as the storage service provider, it becomes possible to provide a complete managed solution with end-to-end QoS enforcement.

This use case also applies to home subscribers with multiple households, which could benefit from a transparent storage service with content synchronization. For instance, a media library could be synchronized between a permanent and a holiday residence of a service subscriber, enabling users to access their media content in a transparent and efficient manner.

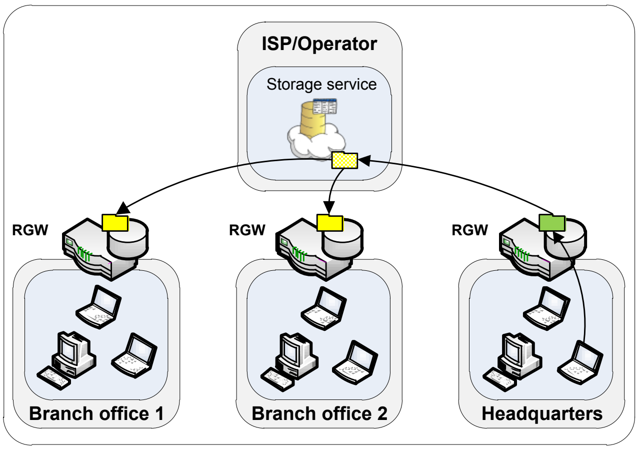

Figure 5.21: Storage repository synchronization across branches

5.3.4 Proof-of-concept and validation

In this subsection we will discuss a proof-of-concept implementation of the hybrid storage service presented on this section, the testbed we used for experimental validation, and obtained results.

Proof-of-concept implementation

The proof-of-concept implementation was designed with the main purpose of providing an adequate test ground for the concept. Its building blocks are depicted in Figure 5.22.

The basic data synchronism system is based on the *git* [Git] version control system, which is able to use bandwidth and CPU resources efficiently, by only transferring file changes (binary deltas) on each update, and by using compression. Also, because it was designed for revision control, it can be used to implement file versioning mechanisms. Session encryption is ensured by an SSH tunnel.

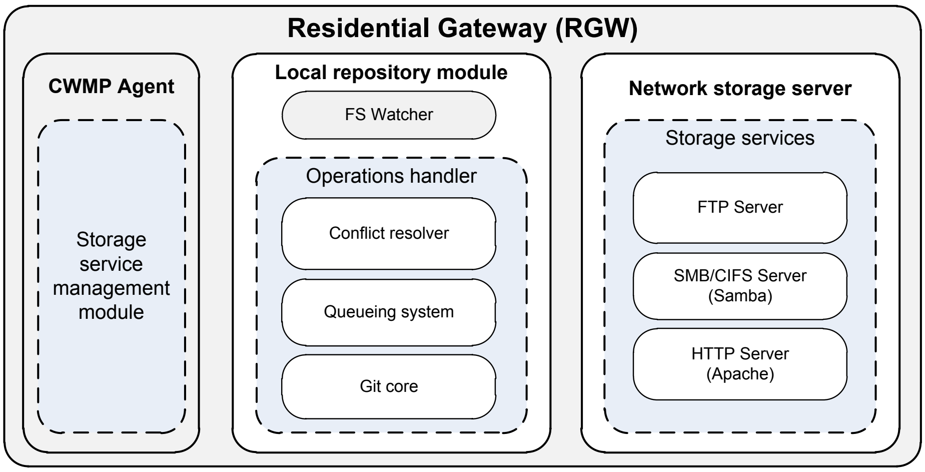

Figure 5.22: Proof-of-concept implementation

The local repository module includes the *git* core, together with the operation queue manager and an external conflict resolver, which complements *git's* native conflict detection and resolution mechanisms – while *git* is able to deal with most update conflicts in an automatic fashion, it provides enough information to enable resolution in the remaining situations. The

File System (FS) watcher is implemented using the *inotify* Linux kernel call to detect changes on the synchronized volume.

The network storage server component provides access to the synchronized storage volume to the LAN clients, supporting HTTP, FTP and SMB/CIFS.

The proof-of-concept CWMP agent was developed in Java, using the dynamic CWMP agent extensibility framework described in Chapter 3. Nevertheless, it could also be implemented using any other CWMP agent library.

The architecture and proof-of-concept hereby described are focused on RGW/CPE-level implementation and, therefore, does not cover the specific use case of roaming clients. However, such situations could be covered by using a dedicated client in the lines of the recently presented *SparkleShare* solution [SparkleShare], which also supports *git*.

Testbed

To evaluate the solution, a reference testbed scenario was established, in order to reproduce the conditions of a managed broadband access network (see Figure 5.23).

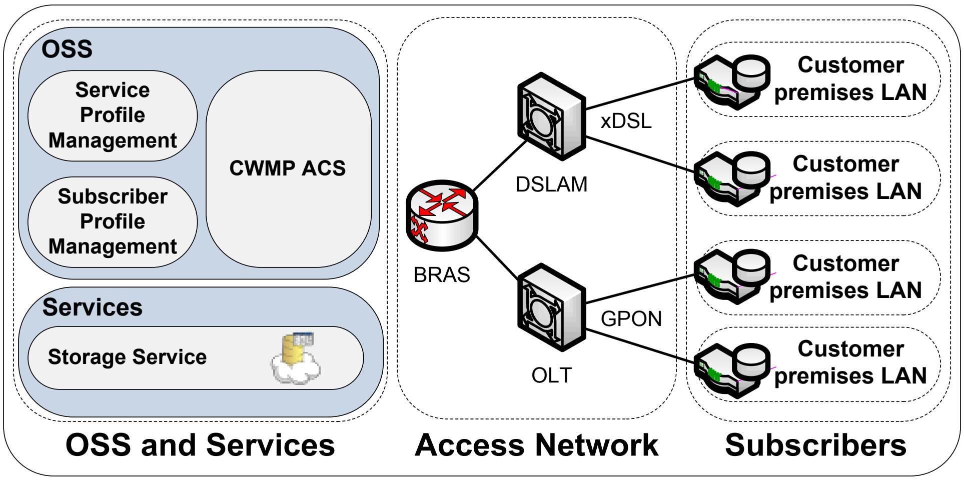

Figure 5.23: Reference scenario

This corresponds to typical broadband access networks using GPON or xDSL technologies, where the operator (OSS and services), access network and customer premises domains are clearly identified (the operator core network is absent for simplicity). Since it was unfeasible to reproduce this exact scenario, we opted to emulate it by implementing the testbed which is presented on Figure 5.24.

To mimic the conditions and restrictions imposed by the broadband access link, a transparent *Dummynet* bridge interconnects the customer with the ISP – this system is hosted by a PC with two PCI Express Gigabit Ethernet interfaces, capable of a (measured and tested) sustained forwarding throughput of 800Mb/s (enough to avoid interference with the performed experiments). A PC is connected to a monitor port on the Ethernet switch, which mirrors all traffic coming to and from the port used by the ACS. This PC captures and measures network traffic using *Wireshark*.The ACS server is a Linux system configured with a CWMP ACS – related profile management services are also incorporated, in a simplified form.

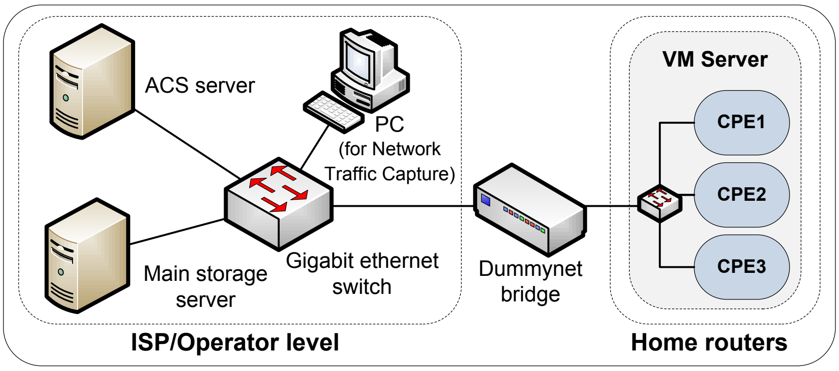

Figure 5.24: Implemented testbed

As for the CPE/RGWs (Home routers), the ideal scenario imposed the use of identical hardware platforms, software component versions and configurations, in order to achieve significant results. That requisite proved difficult to comply with because of the number of variables involved – a problem which was solved by using virtualization techniques. Each home gateway is emulated by a virtual machine hosted by a quad-core Intel Core 2 Quad Q8400 (2.66GHz) with 12GB of RAM and a PCIe Gigabit Ethernet interface. By using the VMware ESXi [VMware2011] hypervisor, it became possible to emulate each home gateway as an exact replica of the other: thus, processor (with affinity of 1 core for each CPE/RGW, limited to 1.5GHz), memory (512MB), storage (5GB for the system image, 10GB to emulate the attached storage device) and hardware specifications were identical among all instances. Moreover, the virtualized internal switch of the hypervisor provides equal network QoS access and throughput shares for each VM instance [VMware2007, VMware2011a].

General considerations about the tests

As for the test plan, it was divided into two stages: functional and performance tests.

Functional tests were successfully performed to validate the operation model and integration schemes proposed in previous subsections.

For performance evaluation, conducted tests focused on overall performance and data overhead, following the following methodology:

- First, reference results were obtained on a scenario with all the equipment connected using 100Mb/s Ethernet (*Dummynet* bridge configured to limit traffic to 100Mb/s in both directions, for each CPE/RGW). In this test phase, results were obtained for local storage overhead, network traffic overhead and data synchronization performance.
- Next, tests were performed to measure data synchronization performance on typical broadband scenarios. The *Dummynet* bridge was configured to enforce bandwidth and traffic conditions representing typical commercial offers (see Table 5.3). As already mentioned, the rationale for the specific configurations applied to Dummynet is discussed in the Annex A.

Table 5.3: Broadband test reference scenarios

	Nominal bandwidth (b/s) (Down/Up)		**Effective bandwidth (b/s) (Down/Up)**		**RTT Latency**	**Pkt. Loss**
ADSL	8M	512K	6.68M	427.5K	20ms	0.1%
	16M	1M	13.36M	835K	20ms	0.1%
GPON	30M	3M	27.9M	2.79M	5ms	0%
	100M	10M	93M	9.3M	5ms	0%
LAN	100M	100M	100M	100M	<1ms	0%

All tests involving data synchronization were performed using a scenario with 3 CPE/RGWs, were one of them triggers an update which is propagated to the main repository and to the remainder CPE/RGWs, referred hereafter as the “1-to-2” scenario (see Figure 5.25).

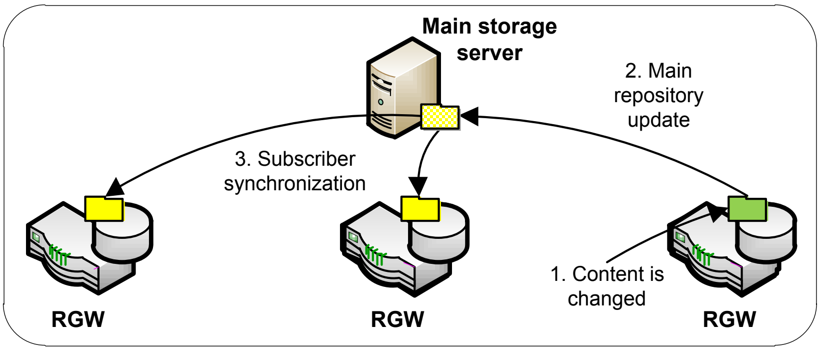

Figure 5.25: "1-to-2" synchronization reference scenario for tests

Overhead on local storage

Overhead on local storage was measured to assess how much space is actually used on the attached storage device on each CPE/RGW, for a certain amount of nominal storage. The basic unit for storage comparison was based on decimal multiples of 100KB, both for total aggregate storage and individual file sizes. Tests were performed for two situations: **individual files**, to assess how compression influences the local storage overhead (since *git* uses compression on the local repository) and **multiple files**, to understand how the overhead differs for several files, for the same aggregate storage size. Both tests were performed on the reference 100Mb/s LAN scenario, since bandwidth does not affect measured results.

For the **individual file test**, different sizes were used for three different content types: random content (a file difficult to compress and generated using the */dev/urandom* Linux pseudo-device as source), random text (using a *lorem ipsum* text generator [Loremipsum], resulting on an intermediate compression ratio), and zeroed files (best-case compression).

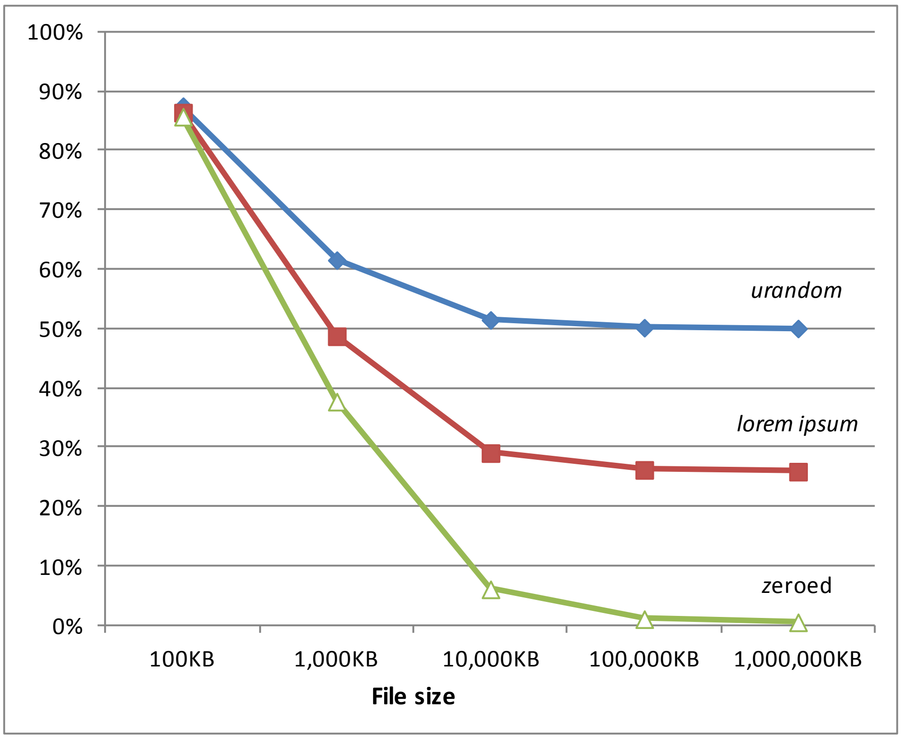

Figure 5.26: Overhead for individual files

Figure 5.26 presents measured results. For small file sizes, the higher overhead has to do with the extra information used by *git* to control versioning and integrity. For bigger file sizes, local overhead will depend on their content: since random files have a low compression ratio, local overhead is expected to be higher than for text files. While it may seem a high penalty, one must remember that this tradeoff is typical of versioning mechanisms in general.

As for the **multiple file test**, the approach was to use the random-sourced (*urandom*) files for each aggregate storage multiple of 100KB (see Table 5.4), based on a worst compression rate case scenario.

Table 5.4: Multiple file test scenarios

	Number of files (*urandom*)				
Aggregated size	**1**	**10**	**100**	**1,000**	**10,000**
1,000,000KB	1,000,000KB	100,000KB	10,000KB	1,000KB	100KB
100,000KB	100,000KB	10,000KB	1,000KB	100KB	
10,000KB	10,000KB	1,000KB	100KB		
1,000KB	1,000KB	100KB			
100KB	100KB				

The results are shown in Figure 5.27. The difference of local overhead between single and multiple files, for the same amount of used space, is negligible, showing the efficiency of *git* in handling such payloads. Similarly to the individual file tests, the higher overhead for smaller file sizes is caused by the extra information used by *git* to control versioning and integrity, which is negligible in absolute terms.

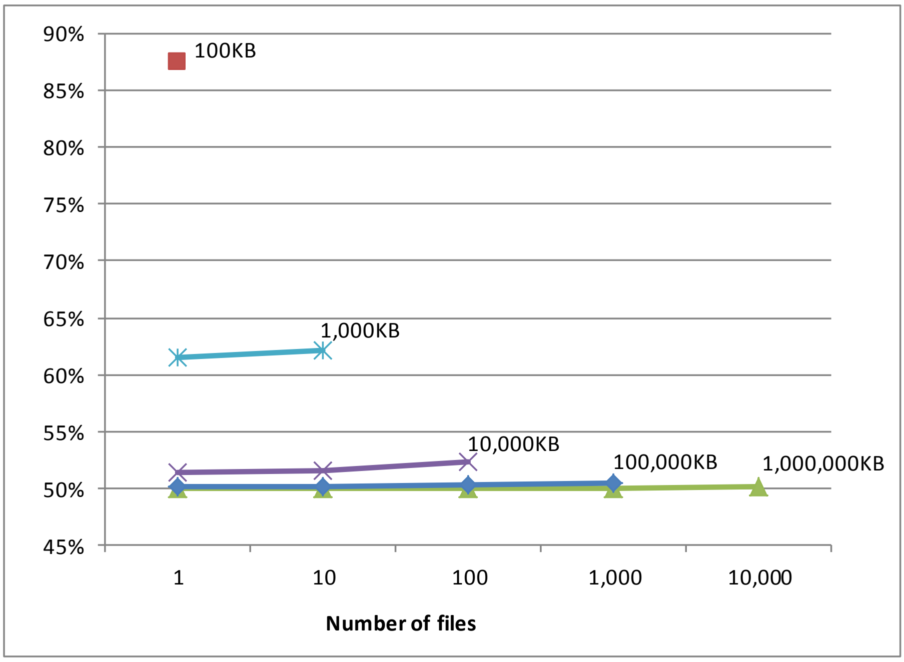

Figure 5.27: Multiple file overhead for the same aggregated size

Overhead of network traffic

Network traffic overhead was measured on the reference 100Mb/s LAN scenario with the *Wireshark* tool, making a comparison between the amount of synchronized data (aggregated upstream and downstream synchronization flows for both single and multiple files with random content). This test was performed using the "1-to-2" reference scenario (see Figure 5.25).

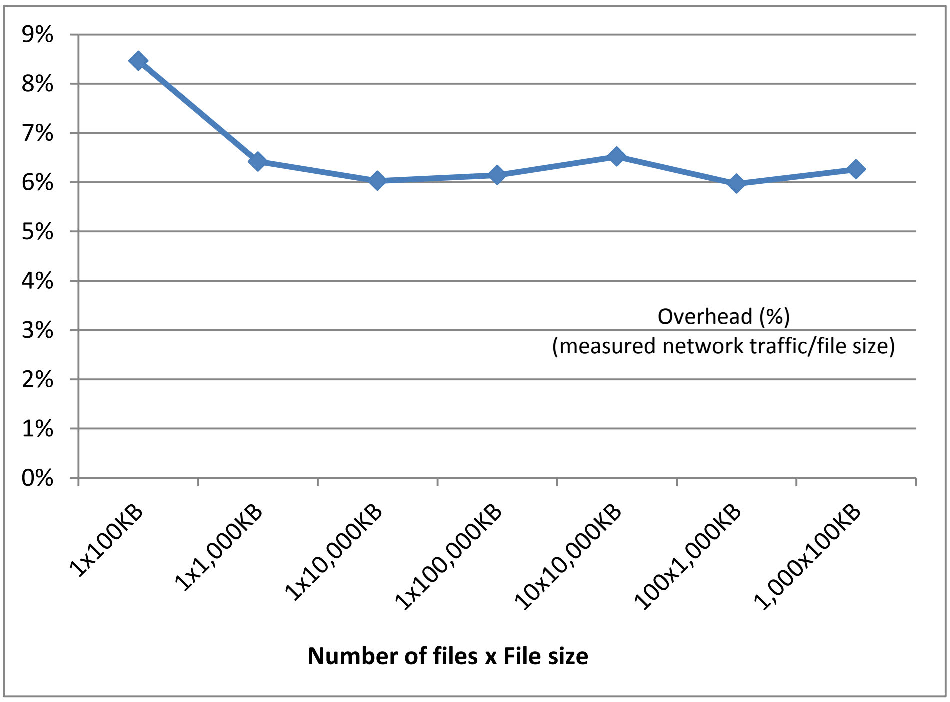

Figure 5.28: Network traffic overhead

Figure 5.28 presents measured results. Except for the case were a single 100KB file was transferred (where the real amount of overhead is amounts to a mere 50KB, which is negligible in absolute terms) the overhead remains consistently between 6% and 7%, regardless of the file size or the number of synchronized files. CWMP protocol overhead was also accounted for and measured, both for provisioning and synchronization operations (see Figure 5.29).

In both cases, 10 tests were performed: in the first case for a complete provisioning operation and, in the latter, for the 1-to-2 scenario. The amount of CWMP-related traffic is very small in either situation, especially in comparison with data synchronization traffic.

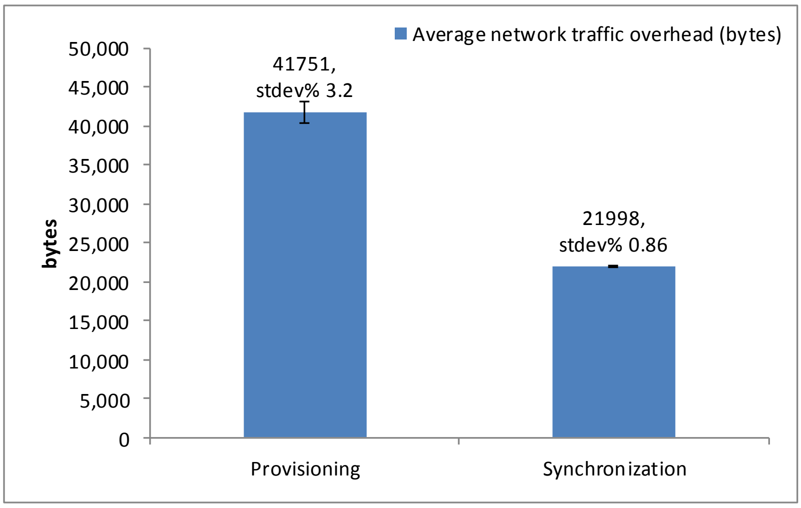

Figure 5.29: CWMP network traffic overhead

Latency of synchronization operations

For data synchronization, the evaluated scenario consisted on measuring the complete delay for the “1-to-2” reference scenario (see Figure 5.25) from the ACS standpoint (using CWMP events for timing). Again, a two-stage test was performed, for multiple and single files with random content (sourced from the already mentioned *urandom* pseudo-device), with 5 interactions per test.

Results for the **multiple file synchronization** tests (see Figure 5.30) demonstrated that, for the same aggregated amount of data (100,000KB), the distribution mechanism is not significantly affected by file sparsity.

However, a detailed analysis of synchronization performance, separated for download (*pull*) and upload (*push*) operations can reveal more details about specific issues not revealed by the simple overall delay measurements. For each case, the operation delay is compared with the theoretical time needed to transfer the payload (network transfer time).

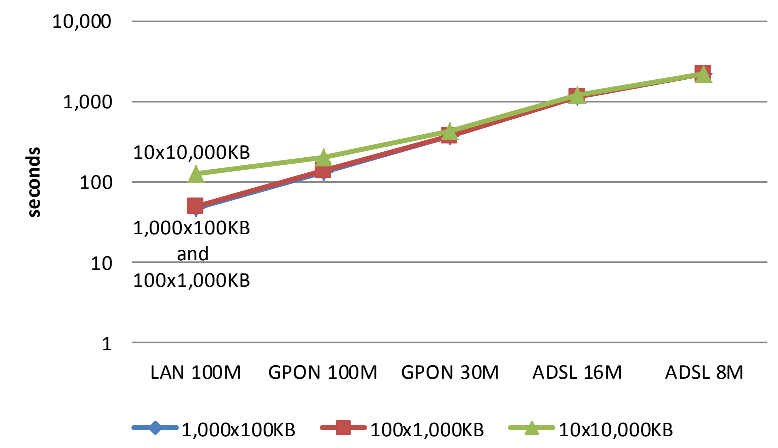

# Files	LAN 100M	GPON 100M	GPON 30M	ADSL 16M	ADSL 8M
1,000x100KB (stdev%)	46.98 (0.45)	131.33 (0.34)	365.35 (0.09)	1129.77 (0.13)	2165.13 (0.02)
100x1,000KB (stdev%)	49.24 (0.23)	135.60 (0.45)	367.53 (0.02)	1131.06 (0.21)	2167.23 (0.01)
10x10,000KB (stdev%)	122.71 (0.68)	197.08 (0.52)	424.59 (0.12)	1186.42 (0.13)	2221.10 (0.01)

Figure 5.30: Synchronization results for multiple files (values in seconds)

For *pull* operations (see Figure 5.31) obtained results are, in general, close to the network transfer time needed to download the payload, with a small overhead for each network technology. The only exceptions are the LAN and GPON scenarios (especially GPON 100Mb/s) for the 10x10,000KB tests, due to processing limitations of the CPE/RGWs – *git* can not process the payload fast enough, on par with network throughput.

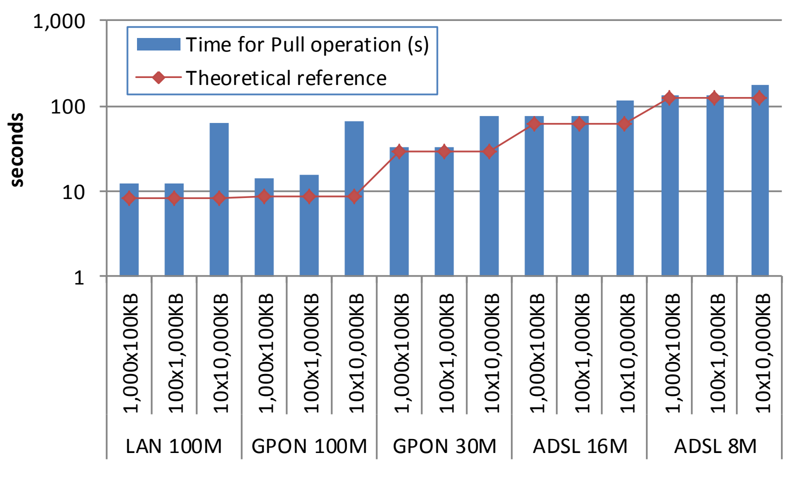

Figure 5.31: Performance of pull operations for multiple file synchronization operations (average for 2 CPE/RGWs)

For *push* operations (see Figure 5.32), the overhead remains generally consistent, independently of the number of files on the payload. However, it can be observed that for the reference LAN scenario there is an overhead increase – again due to *git* not being able to process the data fast enough to maintain throughput. The same behavior was not observed on GPON scenarios due to asymmetric upload/download speeds.

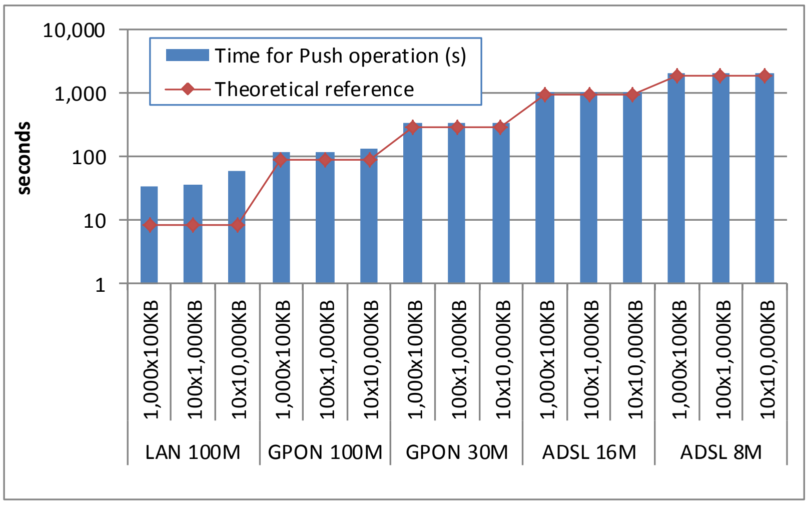

Figure 5.32: Performance of push operations for multiple file synchronization

Results for the **single file synchronization tests** (see Figure 5.33) have shown that, for a tenfold increase in file sizes above 100KB, the synchronization time also increases on the same order of magnitude. The 100KB exception is caused by the git transfer and processing overhead with is more noticeable on small file sizes.

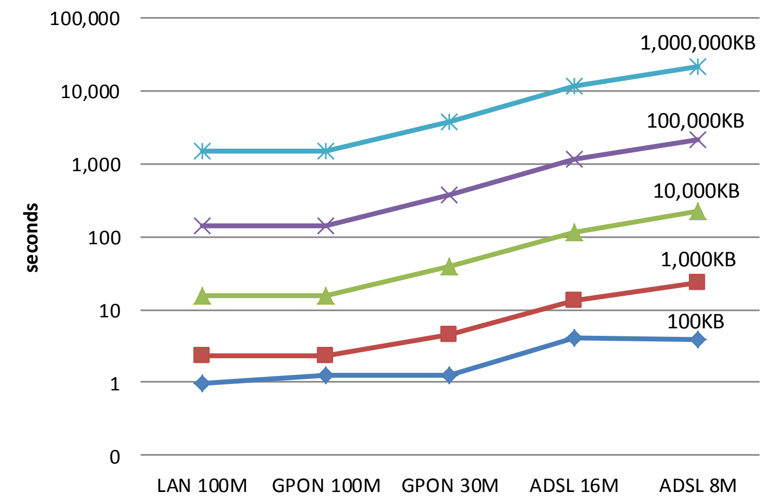

File size	LAN 100M	GPON 100M	GPON 30M	ADSL 16M	ADSL 8M
100KB (stdev%)	0.96 (1.24)	1.24 (6.20)	1.28 (9.51)	3.99 (8.40)	3.94 (0.34)
1,000KB (stdev%)	2.32 (0.78)	2.32 (1.02)	4.58 (0.55)	13.25 (0.50)	23.76 (1.52)
10,000KB (stdev%)	15.35 (0.38)	15.40 (0.36)	38.56 (0.08)	116.37 (0.52)	219.34 (0.09)
100,000KB (stdev%)	142.81 (0.12)	142.34 (0.11)	380.57 (0.18)	1142.08 (0.11)	2180.03 (0.07)
1,000,000KB (stdev%)	1478.95 (0.60)	1491.25 (1.26)	3842.45 (0.28)	11472.88 (0.32)	21844.30 (0.07)

Figure 5.33: Single-file synchronization results (values in seconds)

Once again, a detailed analysis of synchronization performance, separated for download (*pull*) and upload (*push*) reveals more details. For each case, the operation delay is compared with the theoretical time needed to transfer the payload (network transfer time).

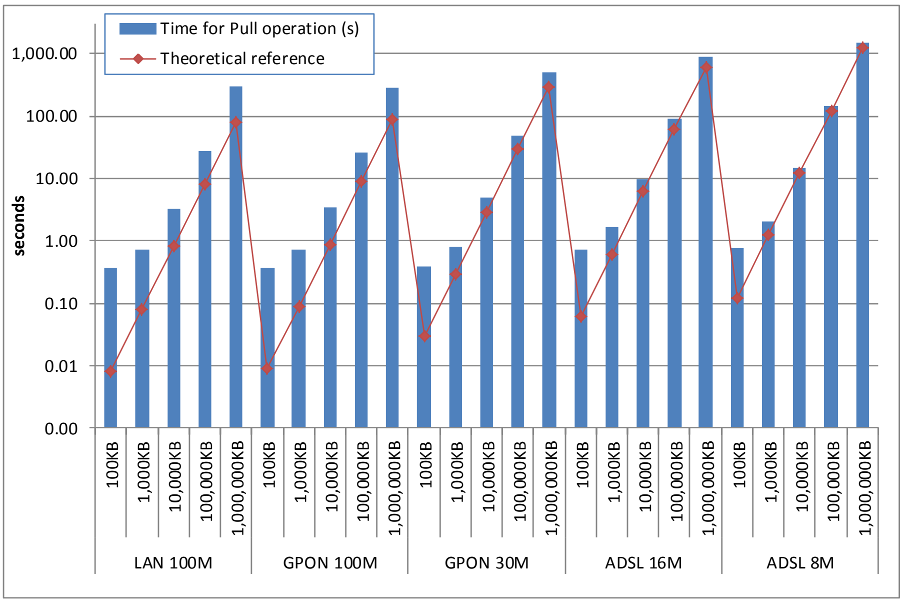

Figure 5.34: Performance of pull operations for single file synchronization (average for 2 CPE/RGWs)

Pull operations (see Figure 5.34) suffer from the same problem observed on the multiple file synchronization test. While obtained results are generally close to network transfer time for the payload (with a small overhead), the LAN and GPON scenarios (especially GPON 100Mb/s) reflect the processing limitations of the CPE/RGWs.

For *push* operations (Figure 5.35), only the LAN tests deviate significantly from the reference network transfer time. This is due to both CPE/RGW processing limitations (which reflect on *git* performance) and the asymmetrical nature of the other access network technologies, that have lower upload speeds.

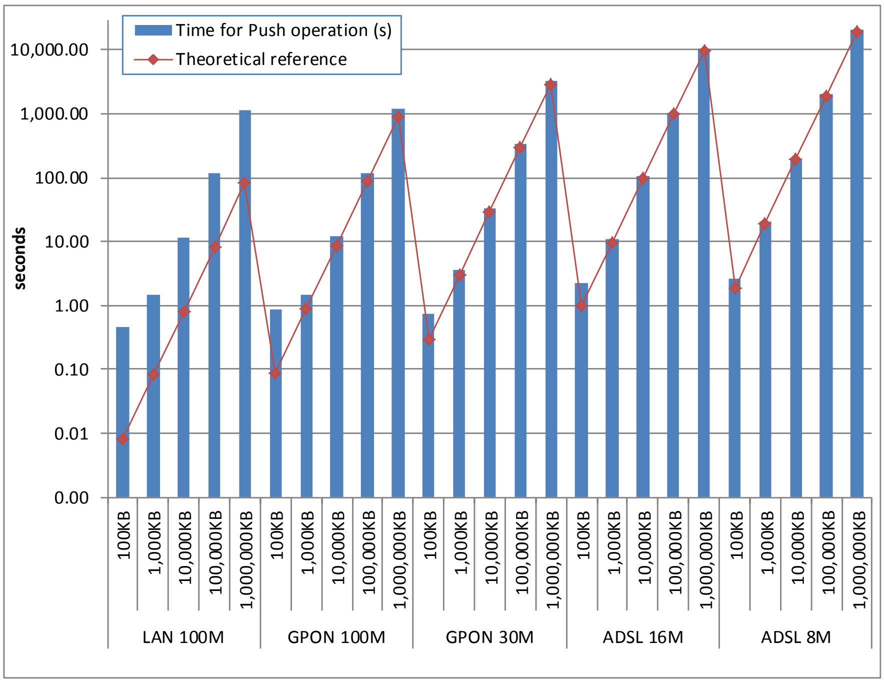

Figure 5.35: Performance of push operations for single file synchronization

Overall conclusions

The obtained results demonstrate that the proposed approach has potential for real-world usage.

The worst-case overhead for the proposed synchronization mechanism is well within adequate levels and might be further mitigated in several ways: first, embedded system performance is increasing each day, eventually reaching the point where it will cease to be an limitation factor; second, it is not expected that the synchronization mechanism will use all available bandwidth at all times (this is especially critical in triple-play scenarios), rather using the remaining capacity; third, since this a fully managed solution, the operator is able to enforce QoS and other adequate measures to maximize performance.

5.4 Related work

This section addresses previous work related with the service delivery frameworks discussed in Section 6.2 and Section 6.3.

The DLNA media-distribution service presented in Section 5.2 proposed a solution to overcome the limitations of the DLNA/UPnP frameworks outside the LAN. The idea of extending the reach of UPnP devices beyond the domestic LAN scope, which is one of the main concepts presented on Section 5.2, has been the subject of some work. The use of VPNs and DLNA proxies to enable cross-LAN interoperability is proposed by [Kamil2009], optionally using social networking mechanisms for access control. In this solution, management can be provided by SNMP [RFC3410] and/or NETCONF [RFC6241] plugins. A proposal based on deployable OSGI components to enable cross-LAN UPnP interoperability, using SIP for inter-gateway signaling, is presented in [Martinez2009]. Another solution for device interoperability, using P2P protocols to create an extended home space for devices and services, is provided by [Park2008]. These three approaches are based on gateway-oriented peer-to-peer solutions for device interoperability, which are not suitable to establish structured managed media delivery architectures as proposed by our approach.

The UPnP forum has also published a Device Control Protocol specification for bridging UPnP devices with an UPnP network across the Internet [UPnPForum2009], having announced an evolution of this specification, scheduled for late 2010 but not yet published. Similarly to [Kamil2009] and [Martinez2009], this solution is focused on bridging disparate devices on domestic networks over broadband access networks, in order to allow UPnP devices on each side to interoperate in a transparent way. This makes it possible to constitute ad-hoc topologies for media sharing, but lacks integration with operators, content-distributors and other kinds of service providers.

In terms of CWMP-UPnP integration, TR-157 standardized a set of data model elements allowing it to embed some information about UPnP devices existing on the customer premises. It has also been proposed to further integrate the generic device discovery and configuration features of UPnP on CWMP environments, using control points embedded on domestic gateways [Delphinanto2009, Nikolaidis2007]. However, those proposals did not encompass other kinds of service-specific integration approaches. Instead, they embed UPnP operation semantics within CWMP for remote configuration, troubleshooting and diagnostic purposes. For instance, in [Nikolaidis2007] it is discussed how the QoS properties of a wireless router could be remotely configured to solve streaming media performance problems.

Outside the CWMP scope, the OSGi management framework also incorporates provisions for accommodating UPnP device control points and generic device implementations as part of its specifications [OSGi2011]. Projects like Apache Felix [Apache2011] have developed generic components to bridge UPnP discovery and control with OSGi platforms (compliant with the OSGi UPnP specification). However, to our knowledge, no one has attempted to use such mechanisms to implement the kind of managed service hereby proposed.

As for content-oriented device and service interoperability, which is closer to the spirit of our proposal, there are several proposals, mostly based on unmanaged ad-hoc sharing topologies. Qualcomm, for instance, recently announced the Skifta service [Qualcomm2011], which aims to tunnel content from local or external sources to DLNA devices on domestic networks ("media shifting", according to Qualcomm's terminology). Skifta uses Android smartphones as intermediaries. It is a locally installed application which bypasses operator and service provider management infrastructures. While the Skifta content-adaptation mechanisms are exclusively hosted at the service provider, our solution allows them to be totally or partially moved to the domestic gateway. TVersity [TVersity] and Playon [Playon] are two software-based server solutions that are able to transcode internet content for DLNA/UPnP AV devices inside the LAN. However, they are based on an unmanaged approach requiring a dedicated storage and streaming server (a Windows PC) inside the customer LAN to operate.

An architecture that uses UPnP and OSGi to enable users to access content from multimedia servers outside their home network is presented by [Kang2005]. This approach transforms the home gateway in a proxy media server for multimedia providers reachable outside the LAN. Another proposal, designed for large-scale media broadcast, is presented by [Hwang2011]. This proposal enables multicast video delivery for DLNA devices inside the customer LAN, using domestic gateways as media proxies. However, these two proposals lack an integrated operator-oriented management approach, leaving management mechanisms out of the discussion.

Solutions to allow domestic LAN content (stored on UPnP media servers) to become available to remote clients are proposed by [Belimpasakis2008] and [Song2009]. Both proposals integrate in a social networking context for access control of shared content. While an interesting proposition, these solutions do not encompass any management mechanisms, being targeted towards *ad-hoc* usage.

The Openiptv forum [OpenIPTV2009] has been developing a specification that supports optional gateway capabilities for allowing set-top boxes to stream IPTV content to DLNA consumer electronics devices [OpenIPTV2008]. Apart from being IPTV-centric, the specification has not yet been finalized and no compliant products exist as of 2011 [OpenIPTV2011].

To our knowledge, the architecture hereby presented is the first proposal to integrate the UPnP AV/DLNA and CWMP frameworks together to provide media distribution over broadband networks as a managed service.

As for the storage service proposed on Section 5.3, the research field of distributed systems, namely distributed storage, has known significant activity in recent years. Support for disconnected access by using replication, which was one of the main novelties of *CODA* [Kistler1992] (a descendent of AFS-2 [Howard1988], which already used disk caching to speed up access) is supported by some other distributed filesystems, like *Intermezzo* [Braam1999]. Also, some storage systems operate using total or partial peer-to-peer data propagation and replication in weakly connected scenarios, like *Bayou* [Terry1998] or *EnsemBlue* [Peek2006]. In most cases, these are either outdated or unfit for production environments.

Clustered file systems have appeared in recent years, such as RedHat's *GFS* [OKeefe2005] (shared-disk), *GlusterFS* [Gluster2011] or *Lustre* [Yu2006] (a distributed evolution of *Intermezzo*). The field of supercomputing research has also contributed with distributed mechanisms such as caching and distributed hash tables that became the basis for the *MapReduce* [Dean2004] framework, which has spawned a new generation of distributed filesystems such as Google's GFS [Ghemawat2003] or *Hadoop DFS* [Apache2009]. However, most of those approaches to distributed storage privilege specific features over others, such as access performance, fault-tolerance or parallelism. Even within restricted the class of distributed storage mechanisms with local replicas, most of them were not deemed adequate for our purposes, in some cases because of the reliance on close coupling of nodes (being designed with high-speed interconnection scenarios in mind), in other cases because supported replication schemes are incompatible with our purposes.

For our specific needs, there was no need for real time data synchronization, metadata servers or complex distributed lock control mechanisms. The problem was therefore reduced to a question of file-level synchronization, for which *Unison* [Pierce2004] and *Rsync* [Tridgell1996] were obvious choices. However, they lacked support for file versioning, something that *git* supported, together with compression and an optimized transfer algorithm based on the use of file deltas. By using *git*, replicas are reconciled using client-provided merge procedures for conflict resolution.

The storage service hereby proposed somehow relates with one of the use cases specified on TR-140: the remote backup storage service. However, the TR-140 specification is targeted towards devices providing storage services for the LAN, not providing any explicit mechanisms for management of such services nor efficient synchronization mechanisms – a conclusion which is reinforced by the *File Retrieval Theory of Operations* section of the same document, which states that a operator might upload or download files using native CWMP download and upload methods, with FTP and HTTP as an option for third-parties. Both methods are inefficient for large quantities of data. File transfer operations over CWMP are affected by both protocol and encoding overhead, also suffering from another problem which also affects FTP and HTTP: they do not take into account if a file was changed or not, transferring the whole contents in each transaction, without any kind of encryption or compression.

To our knowledge, the use of home routers/RGWs as participant nodes in a distributed storage topology is a novelty. While router-assisted distributed storage mechanisms were proposed in the past [Li2010], they do not consider their use as storage nodes. Work has also been done in terms of distributed storage for home environments, but it has been focused on mechanisms designed for the LAN scope or mobile device accessibility [Karypidis2006], [Peek2006].

The commercial offer for cloud storage solutions is mainly based on proprietary protocols targeted toward general-purpose client devices (such as PCs, tablets or smartphones), which require specific libraries and/or client applications for access. This is the case with *Dropbox* [Dropbox], Ubuntu *One* [Canonical2010] or Apple's *iCloud* [Apple2011c] – curiously the first two use Amazon's S3 [Varia2010] service for backend storage, only providing a mediator frontend. While the paradigm hereby proposed supports this operation model, it also provides access to the storage service for devices on the home/SOHO LAN without requiring any add-on component.

Very recently, one manufacturer [LaCie2011] started to offer a hybrid cloud drive in its product portfolio, in the form of a dedicated appliance with an internal hard drive and network connectivity, being able to synchronize its contents to an external provider – however it is designed with only one way synchronization for a single device in mind and, like most distributed storage solutions, it is unmanaged and transparent to the operator. Our proposal goes further by integrating its management and signaling mechanisms within the CWMP framework of the operator, making use of its native management, security and eventing mechanisms to provide a complete end-to-end solution.

5.5 Conclusion

This chapter had the purpose of exploring new managed service paradigms, taking advantage of the management framework proposed on Chapter 3. To this purpose, two proposals were presented, each addressing innovative managed service paradigm implementations:

- **A media distribution service** enabling UPnP AV/DLNA consumer electronics devices to extend their reach beyond the domestic LAN in a managed way, providing seamless access to media distribution services provided by operators or associated third-party content providers.

Overall, the advantages of the presented approach are manifold: it allows already existing UPnP AV/DLNA devices to access and use new services without any modifications, while allowing content providers and operators to reach a wider device base, instead of only targeting purpose-build devices, equipped with specific firmware support for each service.

- **An operator-managed hybrid storage model**, which makes use of domestic routers as storage service hubs to explore the existing complementarity between the traditional/appliance-based and the cloud-storage service models. It provides an integrated solution eliminating the need for special-purpose storage appliances, by turning the home gateway in a storage gateway that can be accessed by all kinds of devices inside the LAN without need for explicit support (like the installation of a client application), while offering redundancy, availability, speed and support for disconnected access.

In both cases, the solutions are fully managed by the operator, that is able to remotely provision and manage the services. These functionalities were implemented by integrating the proposed solution with the CWMP management framework we previously proposed. The use of CWMP also enables the operator to provide complete end-to-end monitoring and QoS management.

Overall, feasibility of the proposed approaches was also demonstrated, through experimental studies focused on validating the solutions – addressing functional, performance and compliance validation.

6. Operator-assisted desktop management

This chapter will deal with the thematic of desktop management over broadband access networks, by pursuing an approach which makes use of the management framework presented in Chapter 3 to propose two novel and complementing solutions to the problem: the first one targeted towards the existing (and still predominant) windows-based PC desktop paradigm and another one designed around the concept of remote *cloud boot*, enabling the creation of stateless, *cloud-based* desktop paradigms (albeit it can also be used to help managing conventional PCs).

The first proposal, discussed on Section 6.2, allows traditional desktop management technologies designed for the corporate LAN environment – namely the widely available Windows Management Instrumentation (WMI) [Micrososft2010b] – to operate in domestic and SOHO environments. This integration extends the applicability of desktop management technologies to a wide range of Windows-based devices located in home networks, small offices and small organizations. Consequently, it opens the way for a number of novel application scenarios, which will also be addressed. This solution is based on the integration of WMI with the Broadband Forum's CPE Wan Management Protocol (CWMP), being fully compliant with CWMP standards and therefore, able to integrate with existing CWMP platforms (a benefit of the management platform discussed in Section 3).

The second proposal, discussed on Section 6.3, is based on the concept of cloud-boot, using an established standard for remote boot on LAN environments – the PXE (Preboot eXecution Environment) protocol [Intel1999]. By overcoming PXE limitations on access networks, the proposed solution integrates PXE support in the Internet Service Provider (ISP) management infrastructure – also making use of the CPE WAN Management Protocol (CWMP) to control PXE-related service parameters on the user's residential gateway. Application models for this solution will also be presented, through specific use cases where it can be used to remotely manage desktops, providing boot support for install, recovery or update procedures. We also discuss how to implement a completely stateless thin-client – thus enabling a complete end-to-end *Desktop-as-a-Service* model (DaaS) [Fisher2008] based on remote desktop technologies and boot-time downloaded OS.

The rest of this chapter is organized as follows: Section 6.1 discusses the problem of desktop management in broadband environments in more detail, with Sections 6.2 and 6.3 dealing with the two proposed solutions. Section 6.4 discusses related work, while Section 6.4 will conclude this chapter.

It should also be mentioned that the content of this chapter is extensibly based on previously published material [Cruz2010, Cruz2011a].

6.1 The case for operator-assisted desktop management

Large organizations always had to deal with the problem of managing hundreds or thousands of desktop PCs, each one with a Total Cost of Ownership (TCO) which tends to largely exceed its initial acquisition cost, when accounting for maintenance and indirect costs. As a result, considerable amounts of money and effort have been invested in enterprise desktop management solutions, motivating industry standards such as the well-known Web-Based Enterprise Management initiative (WBEM [DMTF2011]), promoted by the DMTF (Desktop Management Task Force) [DMTF]. The WBEM framework was developed for distributed enterprise management scenarios and has spawned several implementations, like Microsoft's Windows Management Instrumentation (WMI).

A different reality faces domestic and SOHO (Small Office, Home Office) users, which are required to directly manage their own PCs, despite frequently lacking the required technical expertise to do so. As opposed to corporate users, those users do not have proper tools or technologies at their disposal – in fact, the vast majority of desktop management standards focuses on enterprise LAN management paradigms, excluding domestic and SOHO environments or even small organizations served by commodity broadband Internet services, due to design and practical limitations. For these situations, existing alternatives are limited:

- Unmanaged standalone PCs, with a significant TCO overhead and unable to be remotely diagnosed or recovered from bare metal, in case of critical failures.
- Standalone Intel *vPro*–certified PCs with embedded out-of-band desktop management capabilities [Intel2010]. However, vPro is only available on limited hardware and firmware combinations.

As a result, there is a lack of remote desktop management solutions for domestic and SOHO users connected to broadband access networks. This contrasts with the enterprise LAN environment, where there are several standards, resources and frameworks for PC or thin-client management.

Industry-led initiatives such as HGI [HGI] and the Broadband Forum [BBForum] did produce a number of recommendations and technical standards for remote management of devices on broadband environments (such as CWMP), covering a considerably large array of managed devices located inside the customer premises, such as home gateways, network equipment, set-top-boxes, VoIP devices, web terminals and all sorts of storage and media devices. However, the desktop is still not covered by CWMP – in fact, one could almost say that the PC is the last device standing out.

As an alternative to a device-centric approach (i.e. PC management), other paradigms have emerged which try to address the problem of desktop management by decoupling the physical equipment from the desktop work environment – thin clients, remote desktop protocols, web terminals/appliances and related concepts. These solutions provide access to desktop sessions using remote access protocols, in a device-independent fashion, namely:

- In the case of telecommuters or remote branches of larger organizations, the use of remote desktop protocols – such as Microsoft's Remote Desktop Protocol (RDP) [Microsoft2010d], Citrix Independent Computing Architecture (ICA) [Harder2009] or Virtual Network Computing (VNC) [Richardson2009] – over SSH [RFC4251], SSL [Freier1996], TLS [RFC2246] or VPN tunnels on PCs or thin-clients. However, depending on the VPN technology, thin clients can be tricky to configure – alternatively the VPN client could be configured on the broadband router, an operation that might also require technical skills lacked by most users.
- Web appliances. However, despite recent developments, these appliances are still unsuitable for many applications and may represent, by themselves, a management problem.

To deal with these problems, we propose two different but somehow solutions: one that addresses the needs of the conventional desktop paradigm and another one based on a remote-booting, thin-device concept that targets the post-PC paradigm (albeit it can also be used to help managing conventional PCs), both integrated within the CWMP management framework.

6.2 Managing the traditional desktop paradigm through CWMP-WMI integration

To deal with the lack of alternatives for the management of traditional Windows-based systems on broadband environments, this section proposes a solution which bridges the gap between desktop management technologies – namely WMI, due to its wide installed base of Windows PCs – and the CWMP framework. More specifically, we present an extension to the CWMP protocol that allows broadband operators to remotely access and manage Windows-based PCs, servers and appliances by using the WMI management API.

This extension is fully compliant with the CWMP standard and therefore easy to integrate with already existing CWMP management infrastructures. Depending on the way it is implemented, it may also be transparent to managed devices, resulting in instant support by millions of Windows machines – desktops, servers, Windows Home Servers [MicrosoftHS], Windows XP embedded devices, etc.

The potential applications leveraged by this integration are also discussed in this section. Integration of WMI with the CWMP management framework for commodity Internet access obviously requires a number of adjustments in the traditional desktop management paradigm – which was focused on the corporate LAN. Several novel application scenarios are presented, introducing the concepts of outsourced and cooperative desktop management.

6.2.1 Introducing WMI

As already mentioned, WMI is an implementation of DTMF's WBEM standard for Microsoft Windows environments. It is supported in all Windows versions since Windows 98. WMI can be used to automate administrative tasks on remote computers – where WMI is also used by the operating system for internal management purposes. WMI uses the DMTF Common Information Model (CIM) to represent systems, applications, networks, devices, and other managed components [10], federating information from several sources.

By defining a model of status, configuration and operational aspects of a Windows-based environment, WMI provides a remote management interface which can be used to gather information about installed software and updates, drivers and devices, network interfaces configuration and user profiles, among others. It is also extensible in order to accommodate new sources of management information.

WMI has two types of information providers: event providers (which generate notifications) and data providers (that deal with management data). In WMI terminology, data is structured in classes that contain properties. WMI Query Language (WQL) [Microsoft2008] – with a SQL-style syntax – can be used to generate and receive event notifications or to retrieve instances of class data and class definitions.

WMI is clearly oriented towards LAN-based environments, an option that reflects on its design. Until Windows OS versions 7 and Server 2003 R2 the WMI API was exposed (for direct access) through Distributed COM (DCOM) [Microsoft2008a], a Windows-specific RPC mechanism not suitable for use outside the LAN environment. Later versions also support Windows Remote Management (WinRM) [Microsoft2009], a firewall-friendlier implementation of the WS-Management protocol based on SOAP. Still, both DCOM and WinRM have trouble operating on NAT environments, requiring the configuration of port mappings for inbound connections on the NAT gateway (such as, for instance, a home gateway), a solution which only works for one managed device at a time.

Still, while its design limitations make it unsuitable for use in broadband environments, the resources provided by WMI are valuable in the context of operator-assisted remote management scenarios. To overcome these limitations, we propose to integrate WMI with the Broadband Forum's CWMP protocol.

6.2.2 CWMP-WMI integration architecture

As discussed before, CWMP is not a protocol targeted towards managing devices and services within LAN environments, like WMI. Instead, CWMP focuses on providing remote device management services for the ACS run by the operator. Since the two application fields are orthogonal to each other, this is not a problem by itself. Nevertheless, this isolation between WAN and LAN management mechanisms means that operators are often unable to take advantage of management services that are present at the LAN level (as is the case of WMI). As a result, the dynamic environment of the domestic LAN becomes invisible to them, representing the loss of valuable management information.

CWMP specifications have somewhat anticipated this problem by including data model extensions that allow the implementation of proxy management mechanisms. TR-106 provides support for home gateways to behave as management proxies for devices and services inside the residential LAN, instantiating each one as a "Service" of the Home Gateway [TR-106]. The ACS communicates directly with the home gateway, being unable to distinguish between services provided by the home gateway itself and proxied services.

The solution hereby presented makes use of the generic extensibility capabilities of the management framework presented on Chapter 3 to bridge CWMP and WMI, abstracting WMI operations within the CWMP API.

As such, the proposed CWMP-WMI extension framework allows two alternative integration scenarios (Figure 6.1):

- The integration is performed at the home gateway. In this situation the specialized component responsible for the integration is a full part of the CWMP agent of the home gateway (or, using a different terminology, a "local CWMP/WMI subagent"). The WMI service is remotely accessed using DCOM or WinRM.

  This option is natively supported by all Windows devices and can be implemented using a generic CWMP stack or the framework discussed in Chapter 3.
- The CWMP-WMI integration component is located in the managed Windows device, acting as an X-CWMP subagent. In this case the access to the WMI services is performed locally, also using DCOM or WinRM.

  This option simplifies a number of restrictions typically found in WMI services (e.g. access control mechanisms). Furthermore, as discussed in Section 6.2.4, it provides a performance advantage by avoiding the inefficient use of DCOM over the network. However, it does require X-CWMP and the installation of software in the Windows PC.

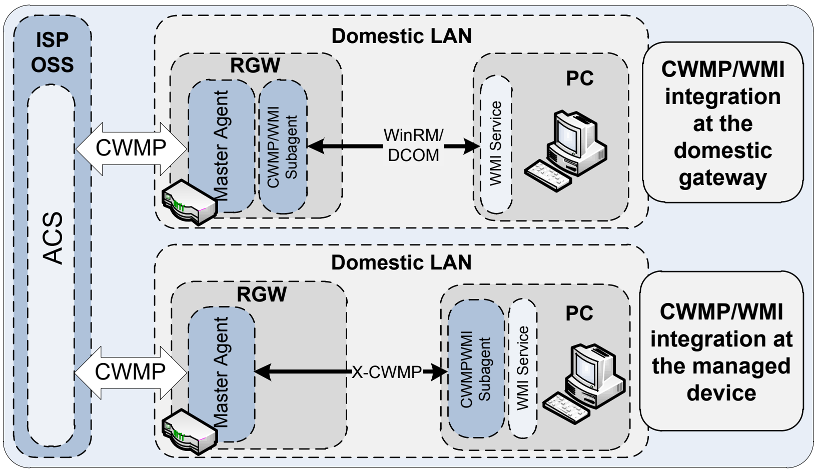

Figure 6.1: CWMP/WMI integration alternatives.

In either case the CWMP/WMI integration component is responsible for declaring which WMI attributes are translated in to CWMP data model parameters and objects for each Windows device instance. These two topologies are both implementable using the framework proposed in Chapter 3, due to the fact that it was designed to allow a hybrid distributed management topology were CWMP agents can proxy the access to devices that, for some reason, do not support CWMP – Master Agents and Subagents may be located in different devices, and X-CWMP subagents may also be used as "Protocol Proxies", enhancing integration with existing technologies like WMI.

Integration of WMI on the CWMP Data model

When the CWMP/WMI subagent registers on the master agent it becomes associated with the CWMP Data Model objects and parameters for which it is responsible. After a valid registration the subagent is ready to receive requests.

Information and properties of LAN devices managed through the CWMP/WMI agents are embedded on the home gateway CWMP data model through the use of dynamic TR-106 data model extensions. This way, the CWMP-WMI subagent maps WMI CIM namespace data into the CWMP data model.

Specifically, TR-106 allows an *InternetGatewayDevice* (home gateway) to act as management proxy for devices inside the subscriber LAN. Each proxied WMI device is modeled through a "*X_<OUI>_WMIService*" vendor-specific object instance – for which a default OUI [IEEEOUI] (Organizational Unique Identifier) of 00000 was defined, for test purposes – which contains the correspondent data model. In this case, the ACS only communicates with the CWMP capable device, which incorporates the data models for the devices for which it is acting as a management proxy (Figure 6.2).

```
InternetGatewayDevice
     (..)
     Services
     X_000000_WMIServiceNumberOfEntries = 1
     X_000000_WMIService.1
         DeviceInfo
             Manufacturer
             ManufacturerOUI
             (..)
         WMIServiceWQL
            Query
            Result
```

Figure 6.2: CWMP TR-106 data model for proxied WMI agent

Integration of the WMI data model into CWMP is performed using one of two methods:

- **Static CWMP-WMI mapping.** Specific CWMP parameters inside each *X_<OUI>_WMIService* instance are statically mapped into WMI class attributes. In this case the ACS user only manipulates conventional CWMP parameters (see Figure 6.3). Each mapping can be configured as write-protected or read-only, at the subagent level. Also, WMI classes (attribute containers) can be mapped as CWMP objects, for better structuring of mapped data.
- **Embedded queries.** Embedded WQL queries are also supported by asynchronous calls. By using a *X_<OUI>_WMIService.{i}.WMIServiceWQL* object, it becomes possible to directly embed WQL queries (placed on the *WMIServiceWQL.Query* parameter) whose result is stored on the *WMIServiceWQL.Result* parameter. The ACS still manipulates CWMP parameters, but the ACS user needs to explicitly know WQL syntax and semantics to perform operations. The *WMIServiceWQL.Result* parameter is defined on the CWMP data model with the Forced Active Notification attribute enabled, allowing asynchronous operations using change notification events (see Figure 6.4).

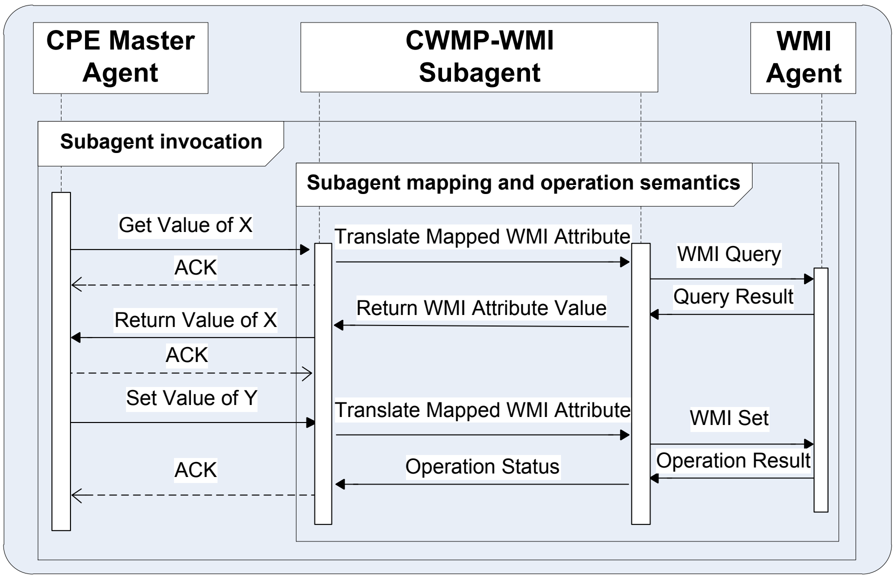

Figure 6.3: Read and write operations on statically mapped attributes

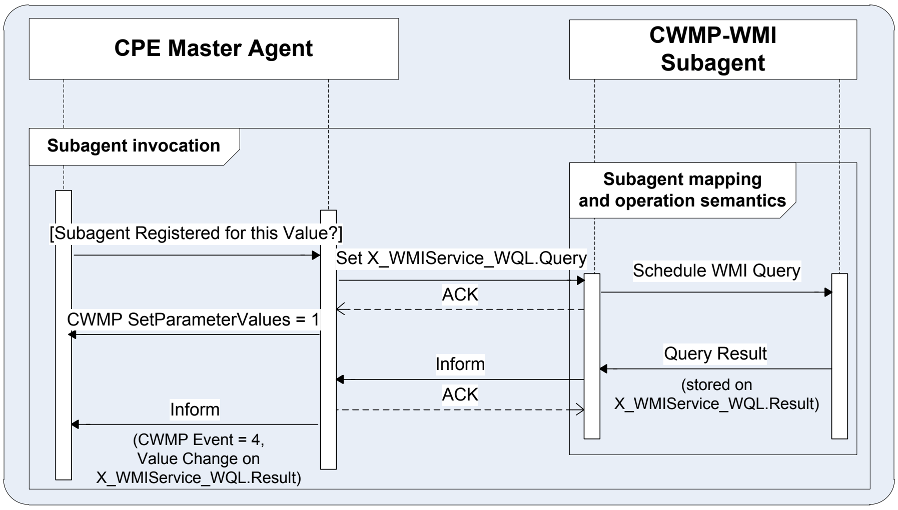
CPE Master Agent
CWMP-WMI Subagent
Subagent invocation
Subagent mapping and operation semantics
[Subagent Registered for this Value?]
Set X_WMIService_WQL.Query
Schedule WMI Query
ACK
CWMP SetParameterValues = 1
Query Result
Inform
(stored on X_WMIService_WQL.Result)
ACK
Inform
(CWMP Event = 4, Value Change on X_WMIService_WQL.Result)

Figure 6.4: Embedded WQL query operations

Static CWMP-WMI mappings are configured at the subagent level, with 3 different types of relationship semantics (depending on the nature of the mapped WMI attributes).

If the mapped WMI attribute is static (e.g. *OS version*), it is queried and stored on an internal data structure, which is only updated at startup. If the value of the WMI attribute changes over time, its mapping can be configured as such:

- as a **direct-mapping attribute** which is queried/set by the subagent each time the corresponding CWMP parameter is requested or set. This is the default behavior.
- or as a **deferred-update attribute** stored on an internal data structure updated by a refresher thread. This behavior is independently configurable for *Get* and *Set* operations on each mapped WMI attribute. A parameter corresponding to a mapped attribute with deferred write properties will always be defined as requiring *Forced Active Notification*, in order to allow asynchronous write operations using parameter change notification events, as per the CWMP standard (Figure 6.5).

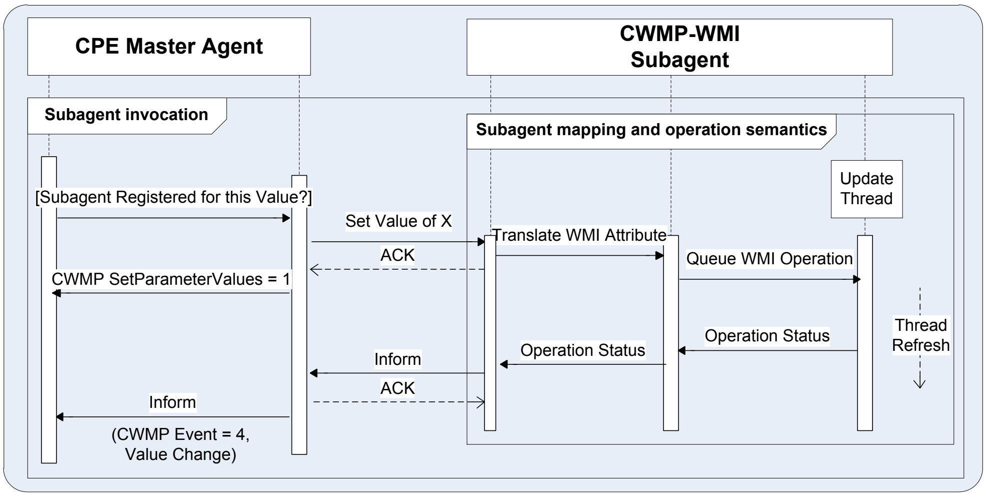

Figure 6.5: Write operations on deferred-update attributes

The refresher thread used for deferred updates must be dynamically scheduled, in order to make use of spare processing power on the host (Home Gateway or Windows Device), minimizing the overhead of the Subagent. Since operations on deferred-update attributes do not need to wait for the result of a WMI operation, they will take less time to execute, at the expense of reduced resolution on attribute data updates. This thread may also be integrated with WMI Events (which are a trigger-like mechanism) in order to enable native event-driven update solutions.

CWMP notification change parameter attributes are supported on the CWMP-WMI subagent data model instances, either being preconfigured (using *Forced Active Notification*) or defined at runtime, using a CWMP *SetParameterAttributes* operation. Thus, it is possible to adjust static attribute mapping behavior in order to balance the tradeoff between fine-grained temporal resolution on specific attributes, speed and subagent load penalty on the managed PC, while complying with CWMP operation semantics.

6.2.3 Application scenarios

As already mentioned, the proposed CWMP/WMI integration framework obviously requires a change in the traditional desktop management paradigm.

While in the corporate LAN the desktops and the management platform usually belong to the same entity (the corporation), here the internet service provider – that controls the ACS – plays an important role, either as the direct manager of the customers desktops or as a mediator for third party management services. These outsourced or cooperative models call for a trust relationship between involved parties. They also bring legal and ethical concerns. However, such concerns are not much different from privacy and security issues brought by popular cloud-based services such as Google Docs [Google2011] and Dropbox [Dropbox]. In fact, typical users take more risks installing third-party software than they would by trusting the management of their desktops to providers they already know. In addition, those providers are already managing other devices *inside* the customer LAN, such as home gateways and set-top-boxes. As long as users are given the choice of controlling which Windows devices are remotely managed and which management information is retrieved, this seems like an acceptable compromise, when compared with other situations.

Next, we discuss two possible application scenarios that demonstrate how the CWMP-WMI extension can be used to provide a new class of added-value services related to remote management of desktop computers.

Operator-assisted management of windows devices

Internet Service Providers may sell to domestic or SOHO customers PCs bundled with software and remote management services (antivirus, software updates, remote recovery, etc.). They may also remotely manage Windows-based media centers, for increased integration with operator provided services and content such as IPTV or VoD (Figure 6.6).

Using the CWMP-WMI Subagent, operators can remotely diagnose and solve common issues on Windows networks, such as domain/workgroup or network stack misconfiguration. Also, the Subagent can also be used to enable remote registry scans to detect a wide range of issues, including Trojans and other security vulnerabilities.

Information about the update level of the OS, its services, running processes and even access event logs can be fed to operator-level Intrusion Detection/Prevention Systems (IDS/IPS) to support active security policies and mechanisms (using correlation platforms similar to the one proposed on Chapter 7).

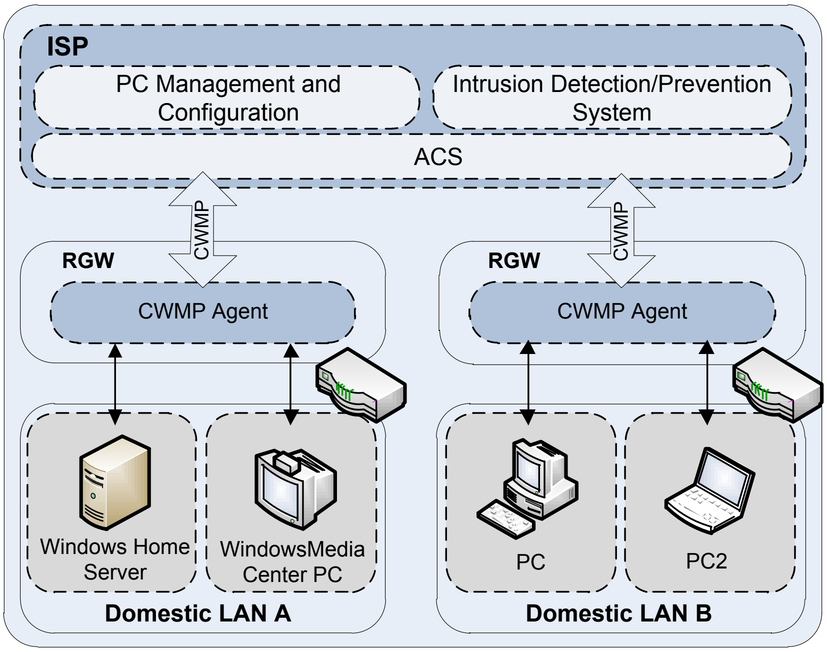

Figure 6.6: Operator-assisted management of Windows devices

Desktop management as a service

While the CWMP management framework was designed with the needs of Internet service providers in mind, it is possible to extend its scope for other uses. Typically, broadband access provider Operations Support Systems (OSS) already integrate the CWMP ACS within their infrastructure, as a component that interfaces with asset management, billing or provisioning systems. In this line of though, it is possible to envision a scenario were an Internet provider may provide an interface for third-party service providers specialized in desktop management (Figure 6.7). Such a service would be especially attractive for small corporations, which would be able to outsource – or cloud-source – the remote management of their desktops and Windows servers.

The proposed usage scenario makes use of a middleware layer, which interfaces with the ACS (and possibly with other OSS services) and exposes an API designed to provide managed access to WMI agents on devices. Third-parties can use this API to offer a remote desktop management service with asset, inventory policy or software lifecycle management capabilities. User privacy is protected by granular non-repudiation and access control mechanisms.

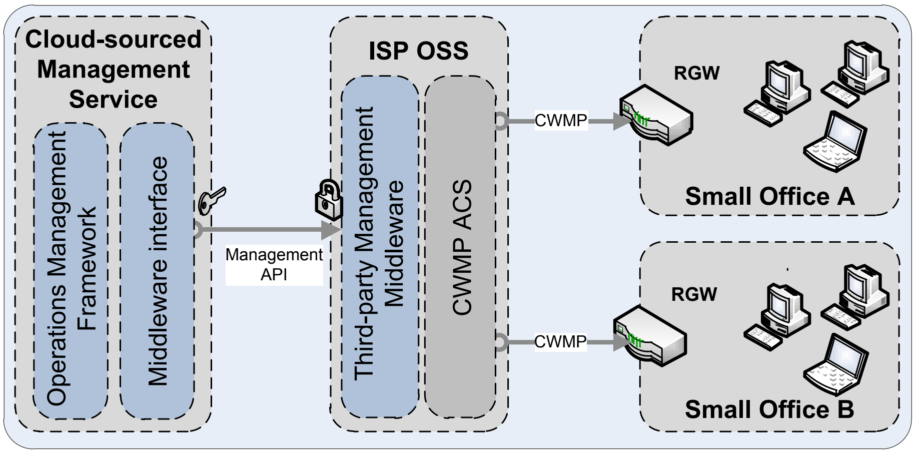
Cloud-sourced Management Service
Operations Management Framework
Middleware interface
Management API
ISP OSS
Third-party Management Middleware
CWMP ACS
CWMP
RGW
Small Office A
CWMP
RGW
Small Office B

Figure 6.7: Cloud-sourced desktop management

6.2.4 Validation

For validation purposes, a proof-of-concept CWMP-WMI subagent was implemented in Java, using *J-Interop* [Dimentrix2008] to access the WMI API through DCOM. In principle WinRM would be more suited to our purposes, since it is an interoperability-focused protocol based on SOAP. However, WinRM support is only available on more recent versions of the Windows OS family, whereas DCOM is universally supported since Windows 2000. As such, DCOM was used to maximize the span of supported Windows OS versions (including, for instance, devices with Windows-embedded versions).

CWMP-WMI validation testbed

The implemented agent was evaluated in a testbed emulating the ACS, the customer premises LAN and the service provider infrastructure (Figure 6.8). In order to mimic the conditions and restrictions imposed by the broadband access link, a transparent *Dummynet* bridge [Carbone2009] interconnects the customer with the ISP. The residential gateway is a Linux system with two network interfaces, supporting NAT. It has an embedded CWMP agent (the “Master Agent” for the CPE environment). A Windows PC acts as a WMI managed PC, with the CWMP-WMI Subagent installed. Another PC is connected to a monitor port on the Fast Ethernet switch, which mirrors all traffic coming to and from the port used by the *Dummynet* bridge. This PC captures and measures network traffic using *Wireshark* [CACE2011].

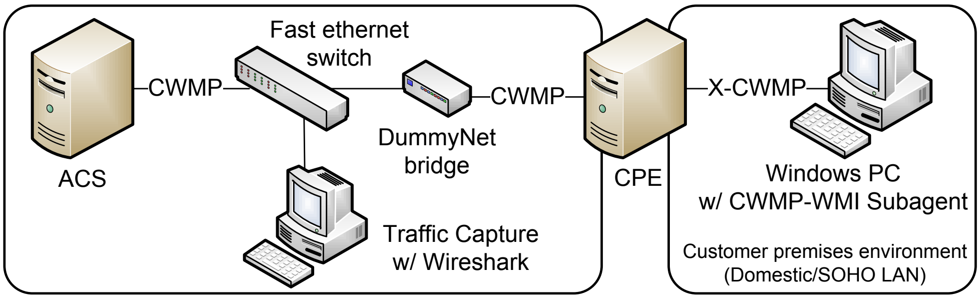
ACS
CWMP
Fast ethernet switch
DummyNet bridge
CWMP
Traffic Capture w/ Wireshark
CPE
X-CWMP
Windows PC w/ CWMP-WMI Subagent
Customer premises environment (Domestic/SOHO LAN)

Figure 6.8: Testbed architecture

CWMP-WMI validation tests

Testing methodology has focused on two specific aspects: latency and traffic protocol overhead. For this purpose, the test plan entails both static-mapping (with direct-mapped attributes) and embedded queries.

The following static-mapping attributes were defined:

- **X_<OUI>_WMIService.{i}.WMIServiceDM.Uptime:** maps to the *SystemUpTime* WMI attribute of the *Win32_PerfFormattedData_PerfOS_System* class.
- **X_<OUI>_WMIService.{i}.WMIServiceDM.FreeMem:** maps to the *AvailableKBytes* WMI *attribute* of the *Win32_PerfFormattedData_PerfOS_Memory* class.
- **X_<OUI>_WMIService.{i}.WMIServiceDM.TotalMem:** maps to the *TotalPhysicalMemory* WMI attribute of the *Win32_ComputerSystem* class.

For embedded queries, a specific "bulk" WQL query (*SELECT * FROM meta_class*) was selected in order to assess CWMP-WMI performance for a larger payload (roughly 1MB).

All repetitive CWMP queries (averaged for 10 tests) were performed using both single/isolated operations (one operation performed on a single CWMP session) and pipelined operations (several operations performed on a single CWMP session).

A two-stage test plan was followed:

- First, reference results were obtained on a scenario where all equipment were connected using 100Mb/s Ethernet (*Dummynet* bridge disabled), with the CPE/RGW configured for packet forwarding only, instead of NAT. Latency and traffic data were obtained both for end-to-end WMI-native queries and for proxied CWMP-WMI access.
- Next, tests were performed to gather latency data for CWMP-WMI Subagent usage on typical access network scenarios. The *Dummynet* bridge was configured to enforce bandwidth and traffic conditions representing typical commercial offers (see Table 6.1). The specific configurations applied to *Dummynet* are discussed on Annex A.

Native WMI/DCOM tests were performed using a simple application written in Java that used the *J-Interop* package to directly query the WMI API using DCOM. The Subagent code was modified to record local WMI query execution times in order to calculate CWMP latency overhead (which is obtained by subtracting the local WMI query latency from the total time needed to perform the full CWMP operation).

Table 6.1: Broadband test reference scenarios

	Nominal bandwidth (b/s) (Down/Up)		Effective bandwidth (b/s) (Down/Up)		RTT Latency	Pkt. Loss
ADSL	4M	512K	3.34M	427.5K	20ms	0.1%
	16M	1M	13.36M	835K	20ms	0.1%
	24M	1M	20.04M	835K	20ms	0.1%
GPON	20M	2M	18.6M	1.86M	5ms	0%
	100M	10M	93M	9.3M	5ms	0%
LAN	100M	100M	100M	100M	<1ms	0%

CWMP-WMI validation test results (100Mb/s LAN)

Figure 6.9 presents results for queries performed on the reference scenario (100Mb/s LAN).

Static, single-session operations perform significantly worse than native or pipelined session queries, because single-session queries are more affected by CWMP session establishment latency. This is partly due to the fact that an ACS cannot directly initiate a connection – CWMP defines an ACS connection request mechanism that instructs the CPE to initiate a connection as soon as possible.

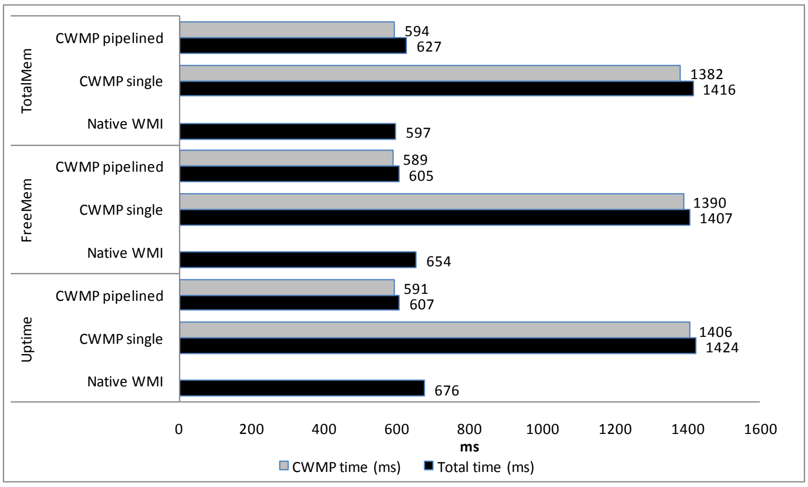

Figure 6.9: WMI/DCOM vs. CWMP latency for static attribute queries (100Mb/s LAN; values averaged for 10 experiments)

With pipelined operations the results are greatly improved, to the point where they surpass native WMI queries over the network. The reason for this has to do with two factors:

- Local DCOM/WMI queries, used by the X-CWMP Subagent, are faster than queries performed over the LAN.
- The CWMP RPC protocol is more efficient than DCOM in terms of operation latency (more data exchanges over the network amplify the network latency penalty).

For single-attribute operations (Figure 6.10), native WMI has a higher traffic overhead, in comparison with CWMP. Also, CWMP pipelined operations generate less network traffic, in comparison with native WMI and single-session operations.

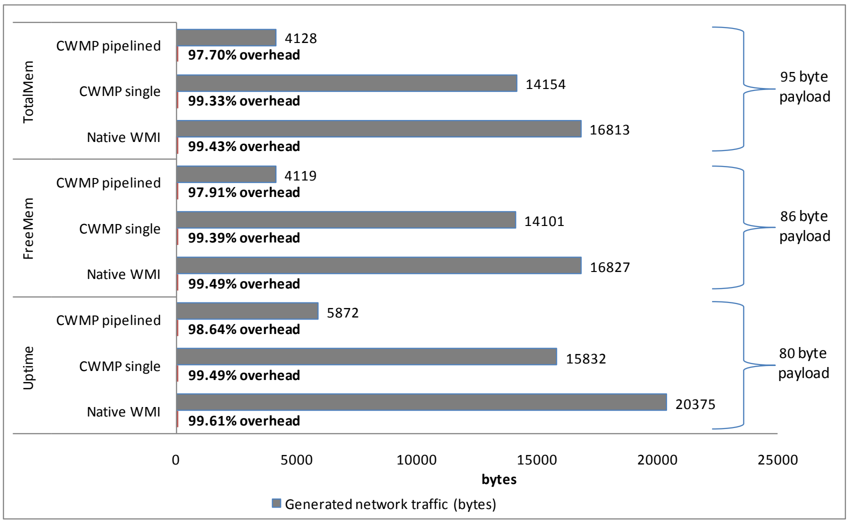

Figure 6.10: WMI/DCOM vs. CWMP traffic overhead for static attribute queries (100Mb/s LAN; values averaged for 10 experiments)

For embedded “bulk” queries, the results for the LAN reference scenario show a situation similar to static queries. WMI queries are significantly faster when performed locally on the managed PC (Figure 6.11) by the X-CWMP Subagent, rather than over the LAN. This makes the CWMP latency overhead comparable with native WMI.

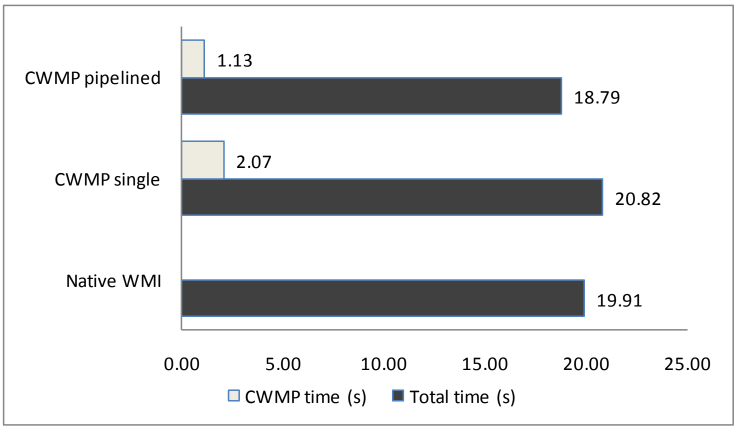

Figure 6.11: WMI/DCOM vs. CWMP latency for embedded bulk queries (100Mb/s LAN; values averaged for 10 experiments)

For a bigger payload, WMI/DCOM inefficiencies are aggravated in terms of data overhead, amounting to 74%, versus 7.7% and 2.4% for CWMP operations (Figure 6.12). Again, CWMP showed higher efficiency than DCOM for RPC operations.

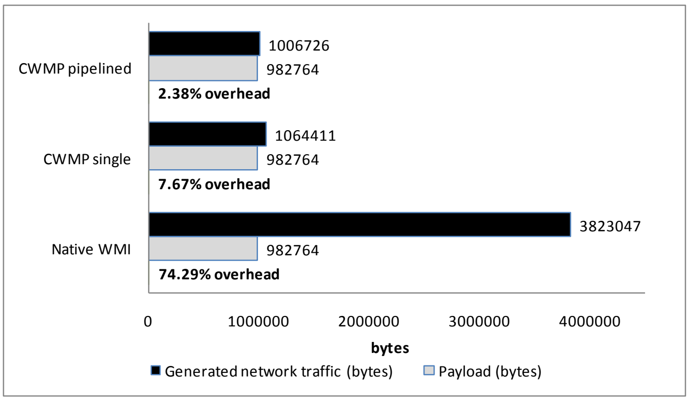

Figure 6.12: WMI/DCOM vs. CWMP traffic overhead for embedded bulk queries (100Mb/s LAN; values averaged for 10 experiments)

CWMP-WMI validation test results (broadband access networks)

For each of the defined CWMP-mapped WMI attributes (*Uptime*, *FreeMem*, *TotalMem*) a set of 10 queries was performed both for single and pipelined CWMP sessions, over different emulated broadband access network conditions.

Average results for total operation time and CWMP latency (Figure 6.13) show significant difference between single and pipelined operations. The average individual operation penalty is 1637ms for single-session and 641ms for pipelined sessions, a difference explained by session setup overhead, which accounts for most of the difference between pipelined and single-session results. The influence of the access network technology is clearly higher on single-session CWMP operations – specifically in terms of latency, which penalizes session setup because of the higher number of protocol exchanges that are performed. Yet, for such small payloads, available upstream bandwidth has little or no influence.

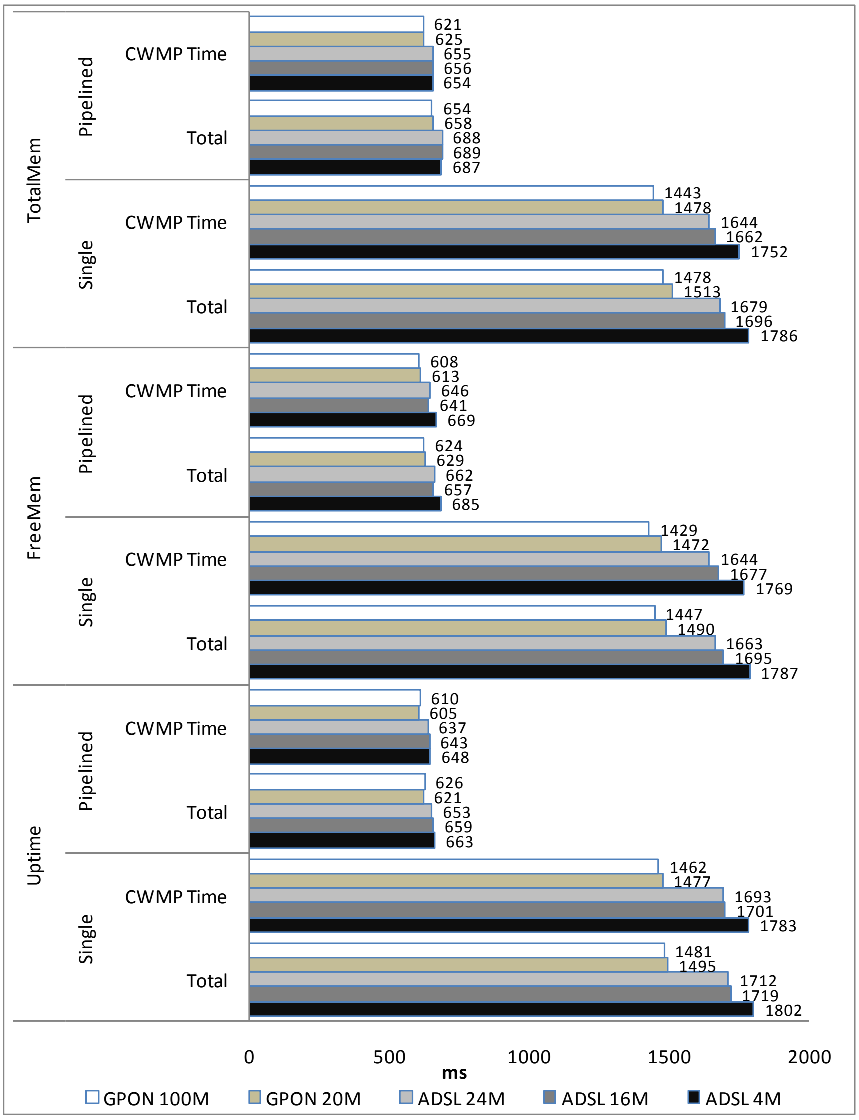

Figure 6.13: CWMP performance for static attribute queries (values averaged for 10 experiments)

For embedded "bulk" queries, the results show significant differences (Figure 6.14), due to the payload encoding and marshalling. In this situation, available upstream bandwidth negatively impacts data transfer times[5].

In comparison with the 100Mb/s LAN reference results for native WMI/DCOM (19.9s – Figure 6.11), CWMP on 100Mb/s GPON (with asymmetric data transfer rates, see Table 6.1) has an average operation latency of 21.6s, with a CWMP penalty of 2.9s for single-session operations and 19.4s with a 1.9s penalty for pipelined sessions.

[5] Extrapolating these results, it becomes clear that native WMI/DCOM operation performance over broadband access networks would be seriously impacted because of the protocol overhead.

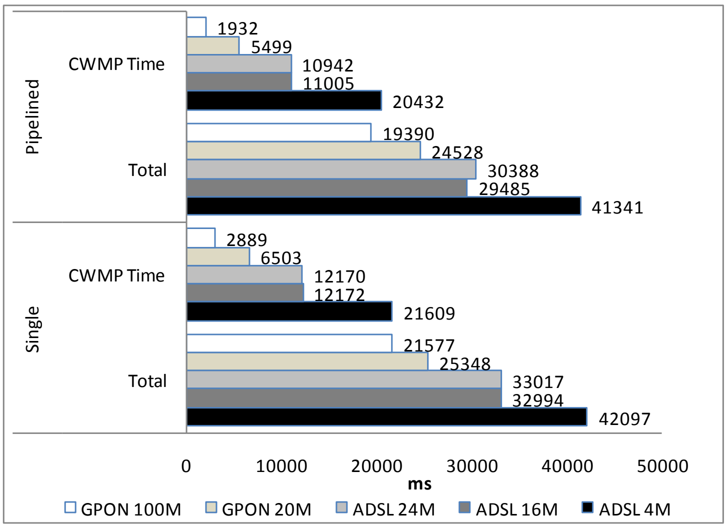

Figure 6.14: CWMP performance for embedded bulk queries (values averaged for 10 experiments)

Overall, these results demonstrate the viability of integrating WMI operations over CWMP. Additionally, they could be further improved by adopting some of the encoding and traffic management techniques proposed by [26].

6.3 Support for novel managed desktop paradigms

Section 4.2 addressed the runtime management of Windows-based devices. While proposed mechanisms can cover a large range of Windows devices (such as Microsoft-based set-top boxes and Windows Home Servers), in practice most of its potential applications relate with the management of classic PCs. However, a number of alternative desktop paradigms are emerging and will play an important role in the home environment.

In this section we complement the (Windows) runtime management solutions from Section 6.2 with the generalized concept of cloud-boot, which opens the way for a number of novel desktop models.

Regarded as the logical evolution beyond the centralized (mainframe-based) and client-server paradigms, cloud computing is a somewhat vague term (more of a metaphor) that encompasses a wide array of technologies and concepts that work together to allow the delivery and consumption of services hosted and supported by remote data centers (providing dynamically scalable and often virtualized computing resources) to another service or an end-user, generally using a web browser as a universal client.

However, the fundamental cloud computing concept of delivering *everything-as-a-service* is heavily dependent on the existence of reliable and capable data pipes connecting providers to service consumers. As such, cloud computing owes much of its success to the increasingly available high-speed commodity broadband access networks (fixed and mobile) without which it would be an unfeasible proposition.

Also, the widespread availability of broadband network access, together with cloud services, spawned a new breed of thin-computing devices that heavily rely on such services. Instead of using the traditional model – where data and applications reside on the device itself – those devices store user data and access applications in the cloud. Netbooks, tablets and some smartphone platforms are now specifically designed as thin computing devices.

At the present evolution stage, even if many users could already permanently live and work on a cloud environment, traditional applications are still the norm on many usage categories. As an example, even if some office and productivity suites are already offered as a service (e.g.

Google Docs [Google2011], Zoho [Zoho]), they remain somewhat limited in comparison to their traditional, locally deployed counterparts, in terms of features, functionality and usability. Also, some traditional applications are starting to embed support for cloud service components (e.g. Microsoft Office suite [Microsoft2010e]) but without replacing the traditional desktop computing model by cloud computing. Instead, these two approaches will likely coexist, cooperate and merge with each other.

The mainstream desktop device is still the standalone PC, with a TCO that largely exceeds its initial acquisition cost. Organizations with dozens, hundreds or thousands of PCs feel this problem in a much bigger scale. For them, as already discussed, the industry created specific standards and tools for enhanced desktop management. An important component of desktop management frameworks relates with the support of remote boot operations, using technologies such as the well-known Preboot eXecution Environment (PXE) protocol [Intel1999], which provides support not only for traditional managed PCs but also for devices like thin-clients. However, remote boot is also an exclusive of corporate LANs.

Until recently, network boot over access networks was unfeasible due to bandwidth limitations. However, with broadband access networks bandwidth steadily increasing, this restriction is disappearing, shifting the focus to the remaining obstacles – like the fact that PXE, the standard remote network boot protocol, uses mechanisms such as the Dynamic Host Configuration Protocol (DHCP) [RFC2131] and the Trivial File Transfer Protocol (TFTP) [RFC1350] in ways that make it unsuitable for naked use over access networks or WAN links.

In this context, integrating PXE-based solutions into broadband access networks would allow novel management paradigms, targeting not just domestic end-users but also telecommuters working from their homes and small businesses which are too small for local deployment of full-fledged enterprise desktop management platforms.

In this section we propose a solution that brings the benefits of managed desktop computing to home users, telecommuters and small businesses by integrating PXE technologies into broadband access network environments to allow for better management of existing desktops and to enable the creation of completely stateless thin-client devices capable of securely booting a remote OS over broadband links.

In addition, we also propose a desktop services delivery model capable of efficiently providing a secure and quality managed desktop experience to domestic and SOHO end-users, using a PXE-based thin-client platform for broadband environments that can replace a full-fledged PC whilst maintaining most of its benefits.

6.3.1 Integration of PXE on access environments

PXE is a Network Boot firmware extension for PC BIOS created in the context of the Intel Boot Initiative. Supported by most Network Interface Cards, it is a *de facto* standard for network boot. It was originally conceived as a special piece of firmware (the PXE boot ROM) that allowed to use the network adapter to download and execute an agent – the Network Bootstrap Program (NBP) – over a LAN at boot time, for deployment, diagnostic or bare metal recovery. However, PXE can also be used to support completely stateless thin-clients [Cruz2003] whose operating environment is downloaded from the network when powered up, instead of using local firmware.

PXE operation model

PXE boot ROMs offer a set of APIs that allow NBPs to use network resources independently of the network adapter hardware, through a Universal Network Device Interface (UNDI) API, complemented by UDP and TFTP APIs (Figure 6.15).

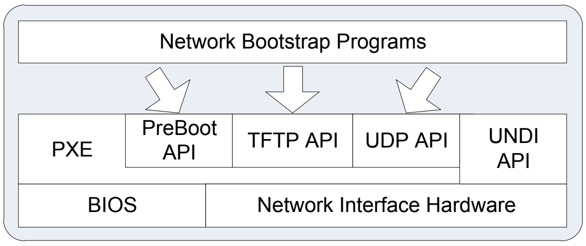

Figure 6.15: PXE boot ROM API.

PXE-compliant Boot ROMs provide the means to control the boot process in order to download and execute either a full-blown OS or just a small pre-boot management agent for diagnostic or pre-staging purposes. Through the use of PXE it is possible to configure a desktop PC boot sequence to be preceded by a PXE boot attempt before using local mass storage devices, making it possible to download and boot a remote OS or a remote agent (in order to initiate maintenance tasks) or to proceed with the normal boot sequence from local storage. Alternatively, it is possible to configure PXE boot to be attempted only in case of local storage device failure (as a recovery mechanism). The operation of a PXE boot ROM follows a simple three-stage process (Figure 6.16):

- IP subsystem initialization. The PXE Boot ROM gets a valid IP via the DHCP protocol (1,2), together with DHCP option tags that identify the presence of PXE support at the DHCP server level, together with the location of the TFTP server and the file name of the NBP to be downloaded.
- TFTP download of the NBP (3,4), using TFTP.
- Execution of the downloaded NBP agent (5,6,7).

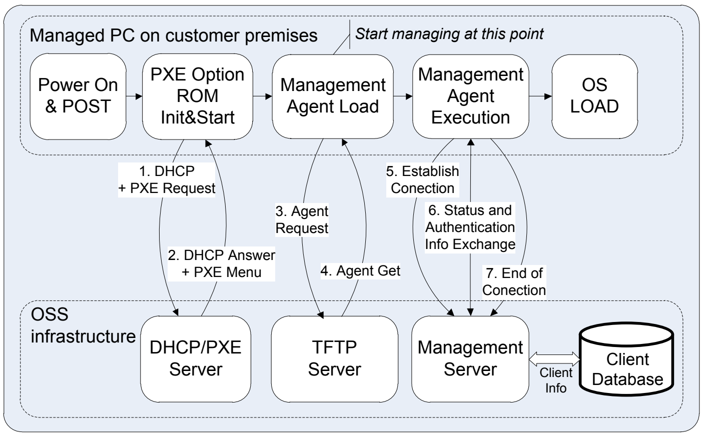

Figure 6.16: PXE agent download process.

The problem with PXE on access networks

When originally conceived, using PXE outside LAN environments was not envisaged, since the download latency would be too high, even for small boot agents with limited functionality. Meanwhile this assumption has been challenged by broadband Internet access technologies like ADSL and GPON. Yet, other problems remain:

- PXE integrates with DHCP in a way that makes it a LAN-specific protocol. Besides depending on DHCP to get a valid IP address, it also receives PXE-specific information from the DHCP server, in the form of DHCP option tags passed to the boot ROM upon initialization. In access networks the ISP DHCP server only manages IP addresses up to the residential router – the customer LAN is managed by its internal DHCP server.
- TFTP (used to download the NBP) is unsuitable outside LAN environments for a number of reasons. First, TFTP is frequently blocked on edge routers, requiring some sort of VPN or tunnel to operate. Second, it is based on UDP and implements a very simplistic transport and

session support, without any windowing mechanism, operating in lock-step mode with only one packet (acknowledgement or data) on the network at any time (the exception being the Microsoft pipelined TFTP service [Gorman2009]), resulting in low throughput over high latency links. Third, redundancy or load balancing is difficult to achieve using TFTP – there are no recovery mechanisms if a boot server is down. Finally, TFTP has no authentication mechanisms and, depending on the implementation and protocol version, can be limited to a file size of 32MB or 4GB.

- PXE is not secure. PXE does encompass the Boot Integrity Services (BIS) [Intel1998], which supposedly provide server verification and validation. However, BIS are not supported by most PXE implementations, making it possible to impersonate the server and provide tampered boot images.

Enabling PXE over broadband access networks

Enabling PXE over broadband access networks implies addressing each of the problems already identified:

- DHCP integration can be solved by dynamically configuring the customer LAN DHCP server on the broadband router to deliver the correct BOOTP option tags [RFC2132], even if pointing to a boot file located on a remote server outside the customer LAN (Figure 6.17).

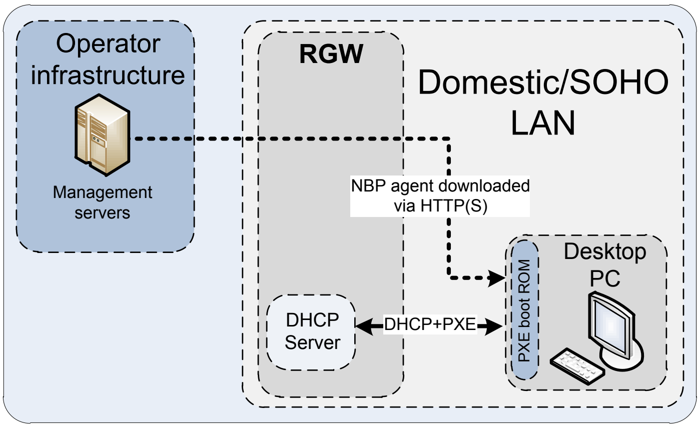

Figure 6.17: PXE operation on broadband environments.

- TFTP can be replaced by the Secure Hypertext Transfer Protocol (HTTPS). The idea of replacing TFTP by HTTPS is not novel, even if so far with LAN environments in mind. One specific network boot loader (gPXE [Etherboot2008], from the Etherboot Project) already provides this feature as an option, supporting PXE with TFTP or HTTPS. gPXE can be used as a drop-in replacement for an existing boot ROM (by flashing it over the existing PXE firmware). Alternatively, gPXE can also be chain loaded by legacy PXE boot ROMs (thus implying no firmware or hardware modifications) or executed from an USB stick.

  HTTPS solves several problems at the same time. Since it is based on TCP sockets it is more reliable than TFTP over broadband links, which are prone to higher bit-error rates than LAN connections. Second, it becomes possible to achieve load balancing on HTTPS/PXE boot servers using mechanisms as simple as DNS round-robin. Finally, because gPXE allows for a chain of boot URLs to be passed on, it is possible to implement a redundancy mechanism that allows for the PXE ROM to sequentially attempt booting from a series of boot servers.

- PXE Security might be enhanced with HTTPS, to guarantee the security of the downloaded stream (avoiding the use of VPN technologies and using SSL instead). The PXE process is still somehow vulnerable during the initial stages (carried inside the domestic/SOHO LAN and based on DHCP), but at least there are no additional security risks introduced by accessing a server across a public network.

In our proof-of-concept prototype, the original PXE boot ROMs were replaced with gPXE. However, it is also possible to deliver gPXE as a NBP agent from the local router itself via TFTP chain loading, keeping the legacy PXE ROM untouched. Since the gPXE binaries are very small (around 30KByte), they do not impose a significant overhead penalty in terms of performance or local flash storage on the local residential gateway (Figure 6.18).

Either way (native gPXE or chainloaded gPXE), all related services (such as DHCP) and parameters must be correctly embedded and configured on the broadband router. This implies some kind of configuration framework enabling the ISP to remotely enable and manage PXE support on the broadband router. The next section will deal with this topic.

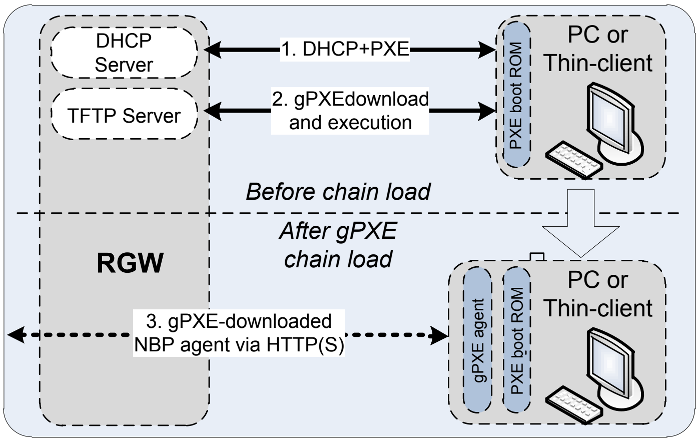

Figure 6.18: gPXE chain loading using a standard PXE boot ROM.

6.3.2 Embedding PXE support on CWMP-based management frameworks

It makes no sense to have PXE-support over access networks if there is no supporting infrastructure on the other side (ISP and/or third-party provider). According to the application scenario, this infrastructure encompasses connectivity, boot services, desktop management services and remote desktop services. This makes it necessary to coordinate PXE with the management framework used by the ISP.

Modern broadband residential gateways are multi-service gateways capable of delivering services such as DHCP, NAT, firewall or DNS caching to the internal LAN. They also perform the role of DHCP servers for internal LANs. Since PXE depends on DHCP, it is necessary to take a step further and incorporate support for PXE broadband boot into those devices, preferably using the standard management protocols already in place.

CWMP can be used by the ISP to configure all PXE-related parameters on the CPE/RGW, enabling the provider to configure which agents and/or which boot images each managed desktop may use and download upon boot. Specific PXE service entries can be added to the CWMP data model, allowing the ACS to configure PXE parameters on each managed CPE/RGW. It is possible to include relevant PXE and managed desktop attributes in the home gateway CWMP data model through the use of the dynamic TR-106 extensions (as discussed in Chapter 3). Following this approach, for each managed device (*thin client* or PC) a *PXEdevice* service entry might be added to the TR-106 data model, containing:

- The identification of the device.
- Its MAC address (optional – for DHCP static leases only).
- Its IP address (optional – for DHCP static leases only).
- Its specific BOOTP DHCP option tags (option tag 66 - *boot server* and option tag 67 - *filename*).

The CWMP agent of the broadband router uses this data to reconfigure the embedded DHCP server, so that it can provide the information to the PXE boot ROMs (through option tags).

To enhance PXE operation, an ISP may also use CWMP to configure a private virtual circuit pipe in order to offer QoS assurance to PXE, related management traffic and remote desktop services. This makes it possible to establish SLA agreements between ISPs and third-party providers of desktop services (commercial providers or private companies serving their own telecommuters and remote offices) to allow end-to-end differentiation of desktop service traffic, for security and/or QoS purposes. Depending on the circumstances the ISP might provide just simple management support (using its CWMP platform to configure PXE parameters at the residential gateway), QoS-enabled virtual channels or security services (e.g. ISP-assisted VPNs). Figure 6.19 illustrates this scenario, both for managed PCs and for thin clients.

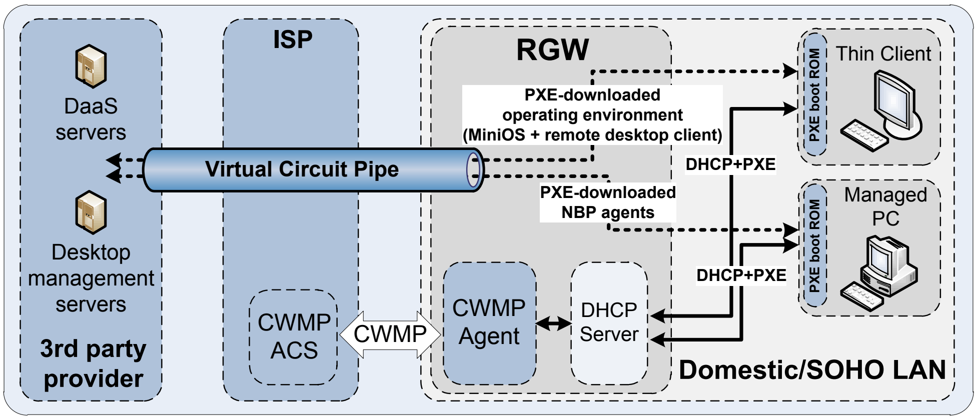

Figure 6.19: PXE-based broadband desktop management.

6.3.3 Application paradigms

The increasingly available bandwidth on broadband access networks – together with PXE support and CWMP-PXE integration – enables ISPs and service providers to deliver a whole new class of services over broadband, designed to minimize desktop TCO while fulfilling users' needs. In this subsection we discuss three possible application paradigms.

Managed PCs

The first scenario corresponds to the usage of classic PCs with locally stored data and applications. In this context, enabling PXE over broadband access networks makes it possible to provide new services:

- Small Businesses may subcontract PC management to third party providers, which, using servers located in remote data centers, may provide a number of PC management services, including bare metal recovery of OS images, automated OS upgrades and remote diagnostics. In this context PXE drastically cuts operations costs, since the need for on-site interventions is strongly reduced and there is no need to place one server on the premises of each customer.
- Large corporations may use this model to manage PCs of their telecommuters and small remote offices, as an alternative to the current VPN-based solutions.
- ISPs might start bundling managed PCs to their commercial offers, addressing both domestic and small business users. Many ISPs already include bundled PCs in their offers (desktops, netbooks), and in that context adding PXE-enabled management mechanisms to those PCs would reduce after sales costs and increase customer satisfaction.

Desktop-as-a-Service and thin-client computing

Thin Clients are well known in enterprise LANs. Being little more than appliances based around low-cost commodity hardware, bundled with a remote desktop protocol client embedded in firmware, they are in many ways the modern counterparts to the old dumb terminals.

When properly managed, enterprise thin-client computing has considerable TCO savings in comparison with typical PCs [Davis2008]. Thin-clients are less prone to critical hardware failures, consume less energy and produce less noise. Since data and applications reside in remote servers, there is no locally stored state information to save and backup in case of replacement. The entire management burden (backups, software and OS updates, etc.) moves to server level.

The recently coined DaaS concept [Fisher2008] – an offspring of thin-client computing and virtualization – applies to the delivery of a desktop environment as a subscribed service [Fisher2008]. DaaS delivery models can be classified in three categories:

- **Hosted desktop session.** This is roughly the same concept that has been used for years in traditional thin-clients. A remote server (e.g. a Microsoft Windows Server with licensed Terminal Services [Microsoft2008b] or a Nomachine NX [Regis2009] infrastructure with a X/Windows [RFC1198] or Windows Server backend) provides simultaneous remote sessions over a specific remote desktop protocol, on a shared server.
- **Hosted virtualized PC/desktop instance.** Instead of sharing a server instance, complete desktop PC instances are self-contained and virtualized on specific platforms (e.g. VMware ESX hypervisor platform [VMware2009]) or physically hosted in datacenters, as blade PCs.
- **End-device local virtualization.** Virtualized desktop instances run locally on the end-device (a PC), on a locally deployed hypervisor. The desktop instance image may be streamed or kept on local storage, using disconnected computing capabilities in the latter case to resynchronize user data and applications when the client device is back on a suitable corporate network – Citrix XenDesktop [Citrix2010], for instance, supports this operation model. Although developed for corporate LAN environments, this approach may also become suitable for commodity broadband access networks, with the increasing available bandwidth.

Broadband users might also benefit from DaaS, replacing standalone PCs with thin clients. This way desktop computing becomes a secure and managed commodity service (see Figure 6.20).

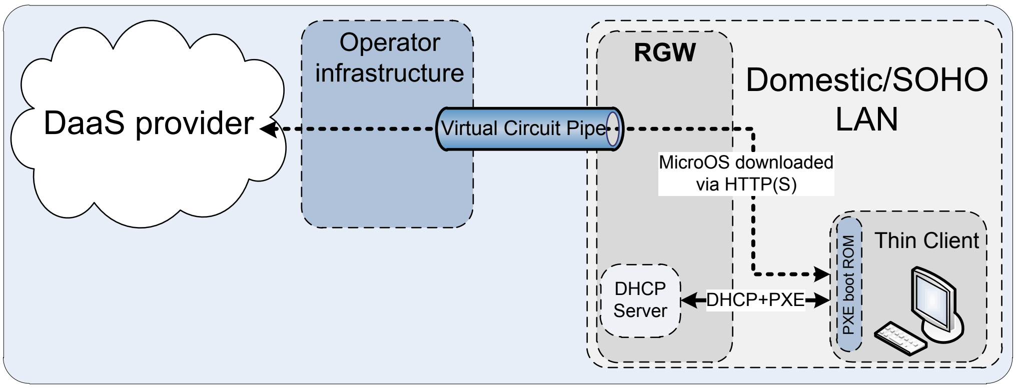

Figure 6.20: PXE-based broadband desktop management.

In this context, instead of using conventional thin-clients with specific connectivity protocols (e.g. RDP, NX, ICA or PCoIP [Black2010]) pre-loaded in the firmware, PXE might be used to build a generic "empty shell thin-client" without local firmware (downloading its entire operating environment from a provider at boot, instead of relying on internal firmware). This type of thin-client does not require firmware updates or replacement to support protocol updates or migration.

This results in increased flexibility and manageability, since even the lightweight operating environment of the thin client is downloaded at boot time. This operating environment is assembled by the DaaS provider – which, as already mentioned, might correspond to the ISP itself or to third-party providers – according to the needs of each customer profile. This means that using a common hardware base it is easy to dynamically customize the experience of each user.

Autonomous “empty-shell” appliances

Until recently, conventional remote desktop protocols used on hosted environments were not media-friendly, especially with video streams [Yang2002]. A new generation of protocols designed to overcome those limitations – such as ICA/HDX [Citrix2010], PCoIP or RDP 7 [Microsoft2010f] – has finally made possible to deliver a PC-like desktop experience on thin-clients (however, those new protocols are not equally suitable for hosted DaaS services over broadband).

As an alternative, the already mentioned “empty shell thin-client” paradigm may enable a particular variation of the concept where some applications are locally downloaded together with the OS at boot time (packed inside the PXE-downloaded boot image) while the user personal data are stored on cloud storage services. This scenario is inline with the advent of so-called “Cloud OS” proposals, like the upcoming Google ChromeOS [Pichai2009] – where applications are supposed to run inside a browser. This approach is also media-friendly from the ground up, provided that media support is included on the OS image (and associated browser).

The concept may be further extended, by using a mix of local applications (delivered on the OS image), remote applications (accessed using remote desktop protocols in application delivery mode) and browser-based applications that use local processing power for media handling without stressing the network.

The “end-device local virtualization” DaaS model fits perfectly in this paradigm. An important benefit from the proposed “empty shell” approach relates with stability and security: since the terminal device stores no operating firmware, if the image of the operating environment becomes corrupted (either by failure or by security attacks) the next reboot will simply load a new image, supplied by the provider. This way, enabling PXE over broadband access networks extends the range of such devices beyond enterprise LANs.

6.3.4 Validation

A proof-of-concept prototype was implemented using the extensible CWMP management platform discussed on Chapter 3 and gPXE [Etherboot2008]. This prototype was then integrated in a testbed emulating the ISP CWMP management server, the customer premises LAN and the service provider infrastructure. In order to mimic the conditions and restrictions imposed by the broadband access link, a transparent *Dummynet* bridge [Carbone2009] interconnects the customer with the ISP and the DaaS provider, emulating the access network link (see Annex A) by enforcing the bandwidth and traffic conditions of typical commercial offers (Table 6.2).

Table 6.2: Broadband test reference scenarios

	Nominal bandwidth (b/s) (Down/Up)		Effective bandwidth (b/s) (Down/Up)		RTT Latency	Pkt. Loss
ADSL	4M	512K	3.34M	427.5K	20ms	0.1%
	8M	512K	6.68M	427.5K	20ms	0.1%
	16M	1M	13.36M	835K	20ms	0.1%
	24M	1M	20.04M	835K	20ms	0.1%
GPON	20M	2M	18.6M	1.86M	5ms	0%
	100M	10M	93M	9.3M	5ms	0%
LAN	100M	100M	100M	100M	<1ms	0%

The broadband router is a linux system with two network interfaces, supporting NAT, a CWMP agent and the ISC DHCP service [ISC] configured to assist PXE boot.

The thin client is a proof-of-concept "empty shell" built around cheap, of-the-shelf components similar to those normally found on commercial thin-clients (x86 1.2MHz CPU, 1GB RAM, embedded LAN, audio and graphics).

On the DaaS provider side there is an HTTP boot server for "remote" PXE. There is also a group of servers providing desktop environments based on RDP, Nomachine and PCoIP. These servers had no direct role in the experimental measurements we present in this section (specifically focused on PXE and the desktop boot process) but showcase the application paradigms which will be next discussed.

Experimental measurements

To test PXE over broadband, a proof-of concept mini OS was created, with a compressed payload around 30MByte (kernel and file system). It is based on the Slitaz Linux Distribution [Slitaz], including a browser, a media player, basic tools and desktop clients for RDP and Nomachine. This mini OS is downloaded by the thin-client at boot, via PXE.

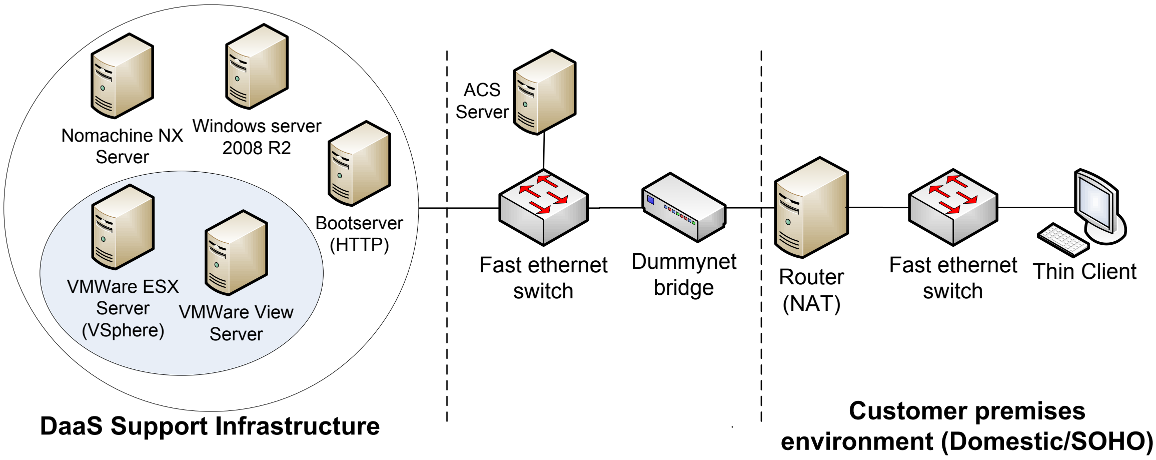

Figure 6.21: Experimental testbed.

The Mini OS image is deployed on the HTTP boot server (Figure 6.21). Its location is advertised to the PXE boot ROM on the thin-client using DHCP tags provided by the DHCP service embedded on the emulated broadband router located in the user premises. The ISP previously defined these tags, using the ACS Server and CMWP agent of the router. Performance measurements encompass four phases:

- Hardware power-on-self-tests (which do not depend on the network and take around 15 seconds for all cases).
- PXE initialization elapsed time.
- Download of the MiniOS image.
- And boot of the MiniOS image (which does not depend on the network and takes around 43 seconds for all cases).

The first round of measurements was disappointing, since the MiniOS download was ostensibly slower in access networks, when compared with the LAN reference (Table 6.3). Even with 100 Mb/s GPON it took 9 times longer to download the MiniOS image, when compared with Fast-Ethernet.

Table 6.3: PXE performance over access networks (mm:ss)

(average of 10 experiments)	**PXE initialization**	**MiniOS download**	**Total time since power up**	**Stdev (%)**
LAN 100Mb/s	00:07	00:07	01:12	0.18
ADSL 4Mb/s	00:08	04:37	05:43	0.24
ADSL 8Mb/s	00:08	04:15	05:21	0.25
ADSL 16Mb/s	00:07	03:34	04:40	0.12
ADSL 24Mb/s	00:08	03:33	04:39	0.03
GPON 20Mb/s	00:07	01:03	02:08	0.43
GPON 100Mb/s	00:07	01:02	02:07	0.77

A subsequent analysis revealed efficiency issues on the gPXE stack, which was unable to properly scale with increased bandwidth, especially in ADSL scenarios (measurements with the same gPXE stack, but with 10ms RTT latency, showed significant improvements). Further analysis revealed the cause: due to memory space limitations and other constraints, the gPXE TCP/IP stack lacked TCP features such as out-of-order packet recovery, selective ACK, window scaling or congestion control. As such, it has a default TCP window size of 4KB, penalizing latency. A TCP window size of 4KB makes the download performance very dependent of link latency – with an RTT delay of 20ms, for instance, the maximum throughput is around 1.6Mb/s [Mahdavi1997].

The gPXE TCP/IP stack was therefore tuned in order to better scale with increased bandwidth in broadband scenarios. The TCP window size was increased to 32KB, enough to cope with access network scenarios but still short from the optimal size of approximately 61KB (due to intrinsic limitations of the current implementation of gPXE).

Performance measurements (Table 6.4) clearly show that PXE over broadband access networks is viable, both from a functional point of view and from a usability perspective: for GPON users the elapsed time from power-up to a fully operational desktop is similar to the 100Mb/s LAN reference.

Even users using slow ADSL 4Mb/s connections have tolerable performance: 2 minutes and 28 seconds from power-up to a fully usable desktop, compared to 1 minute and 12 seconds of the reference LAN scenario. This happens since differences in MiniOS download times are attenuated by long network-independent phases, such as hardware power-on-self- tests and MiniOS boot (it should be noted, however, that bigger OS images might degrade this ratio).

Table 6.4: PXE performance over access networks (mm:ss)

(average of 10 experiments)	**PXE initialization**	**MiniOS download**	**Total time since power up**	**Stdev (%)**
LAN 100Mb/s	00:07	00:07	01:12	0.18
ADSL 4Mb/s	00:08	01:22	02:28	0.32
ADSL 8Mb/s	00:08	00:44	01:49	0.91
ADSL 16Mb/s	00:07	00:31	01:37	0.45
ADSL 24Mb/s	00:08	00:30	01:36	0.75
GPON 20Mb/s	00:07	00:15	01:20	0.69
GPON 100Mb/s	00:07	00:07	01:12	0.54

Considerations about hosted DaaS models and protocols

The performance study previously discussed is directly focused on the boot process, where PXE plays a direct role. Nevertheless, the success of the application paradigms proposed in Section 6.3.3 also depends on runtime performance and the quality of the user experience, when using remote desktop services over broadband access networks.

The feasibility of delivering hosted DaaS services over broadband depends on variables such as the adopted model, application usage and, particularly, network performance (with bandwidth and latency coming on top). Apart from bandwidth, latency is crucial when evaluating remote desktop performance because it accounts for a considerable part of the overall desktop session response time, which ideally should be below the human perception threshold of 50-150 ms [Schneiderman1992]. Preliminary results are summarized in Figure 6.22, which illustrates the suitability of discussed remote desktop protocols accordingly to network latency conditions and usage requirements.

Classic remote desktop protocols (like ICA, RDP until version 6.1 or Nomachine), originally designed for use in session-based hosted environments, had very limited multimedia handling capabilities (low-rated, unidirectional audio with no support for video streams), being only suited for DaaS delivery on business environments. Meanwhile a new generation of protocols (e.g. ICA/HDX, RDP version 7) delivers a more complete desktop experience for both session-based and virtual instance hosted models, with features such as bidirectional audio and media redirection for local rendering.

In some cases, as with RDP 7, newer protocol versions perform significantly better on high-latency links, while consuming less bandwidth and offering increased functionality. By using techniques such as request-reply round-trip elimination, adaptive bandwidth usage, caching or progressive build (providing lossy-compressed images which are progressively built to a full lossless state), those protocols are narrowing the gap between traditional standalone PCs and remote desktop computing. In some cases (e.g. ICA/HDX) explicit support for remote office DaaS delivery is provided, using a branch repeater appliance that uses caching techniques and de-duplication of data for hosted applications and local staging for streamed applications, thus being able to improve scalability and delivery speeds, whilst decreasing bandwidth usage on normal desktop sessions and increasing the number of simultaneous users supported on a single WAN connection.

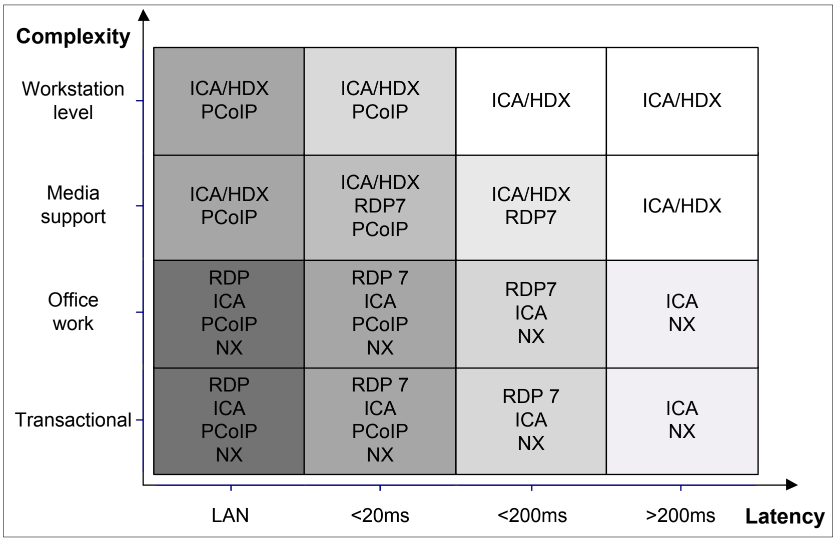

Complexity	LAN	<20ms	<200ms	>200ms
Workstation level	ICA/HDX PCoIP	ICA/HDX PCoIP	ICA/HDX	ICA/HDX
Media support	ICA/HDX PCoIP	ICA/HDX RDP7 PCoIP	ICA/HDX RDP7	ICA/HDX
Office work	RDP ICA PCoIP NX	RDP 7 ICA PCoIP NX	RDP7 ICA NX	ICA NX
Transactional	RDP ICA PCoIP NX	RDP 7 ICA PCoIP NX	RDP 7 ICA NX	ICA NX

Figure 6.22: Suitability of remote desktop protocols accordingly to network latency and usage.

As of the time of the writing of this document, a comparison of remote desktop protocol performance on broadband environments was being conducted. Although this study falls outside the scope of this section, it should be mentioned that its preliminary conclusions show that not all protocols might be equally suitable to DaaS delivery over broadband access networks. While RDP 7, ICA/HDX and PCoIP seem capable of delivering a near-complete PC experience to the end user – with perfect network conditions – they behave differently on broadband environments, with PCoIP performing better on low-latency, high-bandwidth situations and ICA/HDX on the opposite. Classic RDP (up to 6.1), ICA and NX are adequate to business usage (where media support is not a critical issue) with the latter two performing well on high latency links. The ongoing study also suggests that, for accessing hosted DaaS services, QoS traffic prioritization and bandwidth management can make a significant difference in terms of user experience.

6.4 Related work

In this section we discuss previous work related with the two topics we addressed.

To the best of our knowledge, our proposal is the first attempt to integrate CWMP with desktop management technologies such as WMI. Recently, Minokoshi et al., [Minokoshi2010] proposed a supposedly generic approach for CWMP-based LAN protocol proxying, limited to UPnP, SNMP and LLDP. All these approaches are targeted towards management of dedicated appliances and network topology mapping, being conceptually different from the solution hereby presented, which targets desktop computing devices.

For PXE, the situation is akin to WMI, since to the best of our knowledge this is the first proposal to extend PXE – or equivalent remote boot protocols – over broadband access networks. In this context, the closest related work is probably the Etherboot Project, with gPXE [Etherboot2008], which has been actively trying to extend PXE functionalities in multiple directions. Nevertheless, despite implicitly including some building blocks (such as the replacement of TFTP by HTTPS), gPXE does not address PXE over access networks, lacking for instance the integration with CWMP management for remote configuration of PXE parameters.

Intel Active Management Technology (iAMT), which constitutes the core of the vPro Management Engine [Intel2010] may be an alternative for centralized remote management of conventional standalone desktop PCs over broadband, since it supports roaming communications and integrates with several management technologies, including PXE. However, it is only supported on a very limited subset of business-oriented PCs with specific firmware and intel-only hardware, excluding the majority of existing PCs and, more important, practically all thin-clients. iAMT also requires specific infrastructure support to handle roaming users, though a *management presence agent* on the corporate firewall. Moreover, it raises concerns because most users are not able to detect remote access to their PCs via iAMT, turning a management tool into a potential backdoor.

While thin-client computing has been employed primarily on enterprise LANs, a few providers attempted to provide similar services over WAN by using conventional remote desktop protocols [Desktone], [SCC], either for selected applications or complete desktop environments,. In the same line, some organizations rolled-out their own remote desktop support infrastructures, for telecommuters and remote branches [Microsoft2008c].

However, there are few independent studies comparing the performance of thin-client solutions on such environments. Howard [Howard2000] analyzed the performance of several hardware thin-clients using the i-Bench benchmark suite, albeit centered on server-side performance and foregoing client-side performance. Nieh [Nieh2003] and Yang [Yang2002] compared the performance and efficiency of thin-client protocols, analyzing their performance relation with network bandwidth by using slow-motion benchmark techniques to assess loading latency for visual elements. Lai [Lai2002, Lai2006] studied thin-client WAN performance, including efficiency with packet loss and high latency situations, finding no significant impact until a 4% packet loss level was reached, well above typical commodity broadband ratios. These studies also confirmed shortcomings on media usage (especially video) with classic RDP and ICA protocols. Nevertheless, a deeper analysis on the effects of congestion and loss on thin-client performance was left to further study.

6.5 Conclusion

In this chapter we proposed two different but somehow complementary approaches for desktop management on home networks.

First, we introduced the concept of **CWMP-WMI integration**, proposing two alternative solutions for the integration architecture (integration at gateway level or at desktop level). Integration between the WMI and CWMP data models has also been addressed, including support for conventional static attribute mappings and embedded dynamic queries without disrupting CWMP operation semantics. Overall, the proposed mechanisms allow for integrating WMI within CWMP, overcoming the limitations of WMI operation over firewalls, NAT environments and broadband access networks.This chapter also demonstrated the feasibility of integrating WMI operations on the CWMP protocol, paving the way for the creation of operator-assisted management services for Windows desktops, Servers and Appliances.

Next, we proposed **a framework to use PXE on commodity broadband access networks**. This framework makes it possible to provide a number of desktop management solutions (previously limited to enterprise LANs) to small businesses, telecommuters and domestic users.

In that context, we identified the limitations of using PXE outside enterprise LANs and discussed how to overcome those limitations. Next, we discussed the integration of PXE support in a CWMP management framework, allowing for centralized management of PXE boot support at the ISP level, with remote configuration of protocol-related attributes and other properties – such as virtual circuits for QoS provisioning of PXE and remote desktop session network traffic. This also eliminates the need to involve users in the configuration process, therefore reducing the potential of service disruption by misconfiguration or undesirable tweaking attempts from inexperienced users. We also discussed how this approach integrates with the ecosystem constituted by end-users, ISPs and (possibly third-party) providers of DaaS and/or desktop management services.

Then we discussed potential application paradigms for this ecosystem, including managed PCs, DaaS, Thin-Client computing, and autonomous "empty-shell" appliances.

Finally we addressed implementation issues and presented the results of an experimental study focused on validating the use of PXE on broadband access networks – addressing performance and functional validation. Complementing that experimental study, we also discussed the performance of remote desktop protocols over broadband access networks, a fundamental requisite for the previously identified application paradigms.

Overall, the proposed PXE extension to broadband networks – when combined with the proposed DaaS service models and thin-client solutions – provides increased security, total protocol independence and dynamic provider-driven customization. This CWMP-managed, PXE-enabled DaaS model brings together the benefits of thin computing desktop solutions and broadband environments, while promoting an integrated view of the service, involving the multiple providers to deliver the services to the end-users.

7. Security management

In this chapter we address the management of security in a scope that intersects home networks, broadband access networks and the management framework we have previously proposed.

From an ISP perspective, modern broadband access networks pose significant and ever increasing challenges in terms of security management. The growing number of permanently connected home networks, with a myriad of poorly managed devices, imposes significant security risks – not only to the domestic customers, unable to defend themselves from security attacks, but also to the ISP and third-parties potentially targeted by large-scale distributed botnet attacks fed by swarms of zombie domestic PCs.

In this context, the traditional delimitation of customer and ISP perimeters is no longer effective. Home networks became too complex and vulnerable to be autonomously managed by the average customer, and the scale and sophistication of distributed security attacks make it more and more difficult for the ISP to properly manage security without intervening outside the boundaries of its own network.

Considering this state of affairs, we propose an alternative architecture for security management. This architecture increases the level of integration and cooperation between the domains of the ISP infrastructure and the home network. At the same time, it potentially improves the scalability and granularity of traditional intrusion detection and prevention mechanisms.

This chapter is structured as follows. Section 7.1 discusses security in the context of home networks and broadband access networks. Section 7.2 presents a generic operation model for a distributed IDS (DIDS) strongly supported by Residential Gateways (RGW). Next, Section 7.3 discusses evaluation work, based on a proof-of-concept implementation of the managed RGW-based DIDS. Section 7.4 analyzes related work and Section 7.5 concludes the chapter.

7.1 Motivation

As already mentioned, widespread broadband residential access and the advent of new digital convergence paradigms are the driving forces reshaping traditional communication services usage and operation – allowing new levels of service diversity, quality and value. This paradigm shift is a natural outcome of the remarkable evolution that started with the introduction of broadband access technologies to the consumer, in the early 90s.

However, not all aspects of the ISP infrastructure evolved to scale in the same way. Specifically, the mildly distributed security frameworks in use by ISPs are, to a large extent, largely derived from the dial-up era, with limited reach and based on traffic barriers, sinkholes and probes placed at strategic locations to detect and contain potential attacks (see Figure 7.1). These security frameworks are typically based on two assumptions:

- Customers are fond of their privacy and do not easily tolerate external interference with their own equipment and networks – therefore, they are autonomously responsible for their own infrastructure. As such, ISPs clearly define their intervention perimeter to end at their borderline equipment, with the customer being responsible for his own borderline equipment (such as home gateways) and everything beyond that.
- In line with its intervention perimeter, the focus of the ISP is the protection of its own infrastructure, not the protection of the customer or third parties. Cooperation mechanisms are inexistent or, at most, have limited functionality.

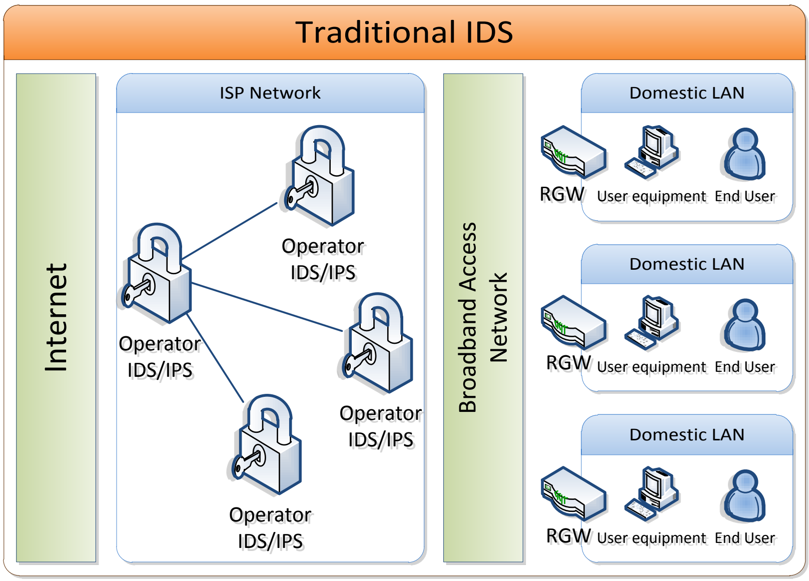

Figure 7.1: Traditional operator approach to IDS

The problem is that those two principles are no longer adequate. Most threats arise from compromised customer equipment inside the home network. Typical DDoS (Distributed Denial of Service) attacks come from botnets formed by swarms of compromised hosts attacking in a coordinated way [Douglieris2003, Thomas2007, Higgins2011, Arora2011]. Spammers use similar techniques to flood mail servers worldwide. In the past, such attacks could be mitigated by limiting traffic in the barriers deployed inside the ISP network, but nowadays these measures are ineffective because the nature of the attacks has changed: instead of using a small number of compromised nodes to flood a target at high rates, now there is a much higher number of compromised nodes (spread among several ISPs and geographic locations) individually generating directed network traffic at almost insignificant rates. Besides, the total volume of traffic handled by the ISP (number of connected homes, traffic generated by each home network) has increased to the point where is it difficult or even impossible to perform detailed traffic analysis at the aggregation points.

Even if some of these risks already existed before the advent of modern broadband access networks, the transient nature of classic dial-up connections and the reduced amount of available bandwidth made it easier to detect and control potential security threats or incidents affecting their own network, customers or third-parties. Nowadays, the increasingly available bandwidth, combined with P2P (Peer-to-Peer) applications and everything-over-IP convergence scenarios only contribute to substantially aggravate the potential risks involved either for the customer or the ISP – which became more vulnerable to orchestrated DoS (Denial of Service) attacks and abusive uses of its network infrastructure.

As a result, the traditional security model begins to suffer from scalability issues, hampering the detection and response to security threats and events. Even when designed to protect the ISP and/or the customer, conventional security measures frequently have some kind of drawback affecting one or both sides – for instance, bandwidth limiting and traffic shaping have been subject of discussion due to ethical and legal concerns. Such collateral effects make it difficult to achieve the right balance between security and privacy. Converged service bundles worsen this situation, since the customer expects VoIP and IPTV services to have equal, if not better, performance and reliability when compared to their conventional counterparts. The impact resulting of service interruptions is more noticeable and inconvenient – even when they are caused by security problems within the customer own network.

Also, as customers voluntarily accept some degree of ISP interference in their networks, the cornerstone premises of the traditional security model begin to erode: with the emerging broader-scope management paradigm shift that was discussed on Chapter 2 and Chapter 3, operators are becoming increasingly pervasive inside the customer premises LAN, managing services and devices and therefore relieving the customer from the burden of being the sole responsible for the management of its own network. However, traditional security models are not designed to take advantage from this increased tolerance.

Considering this situation, we propose an alternative security management framework that takes advantage of the specific role and position of residential gateways – as devices responsible for the exchange of information between the ISP infrastructure and the customer network – to develop a vastly distributed IDS/IPS (Intrusion Detection System/Intrusion Protection System), allowing the ISP to deploy more granular and scalable security mechanisms at the gateway level, thus enabling it to become the first defense layer of its own network.

According to this framework, the RGWs work in a coordinated way, actively monitoring network traffic (based on configurations predefined by the provider), notifying the ISP of suspicious events and enforcing preventive or corrective countermeasures, according to the instructions issued by the ISP. In order to increase scalability, each RGW may also have some degree of autonomy, locally performing event correlation and/or directly deciding on local countermeasures.

From the customer point of view, partially outsourcing the management of security services to the ISP (which has both the needed expertise and a broader view of the access network) potentially results in increase security and reduced maintenance efforts. Besides, it does not restrict or compromises its freedom or privacy, as discussed in the next section.

From the ISP point of view, this is a natural extension. Security mechanisms are remotely managed like other services and devices already controlled by the ISP, using the same technologies (such as CWMP). This architecture makes possible to deploy an extensive and coherent set of remote management and security services in the customer LAN, capable of monitoring and acting on the internal network (if desired) and the traffic flowing between it and the provider network.

7.2 Generic operation principles

In response to the increasing difficulty to scale conventional ISP-level security solutions, the proposed architecture presents a distributed security model taking advantage of the processing and remote management capabilities of home gateways, in a way that allows some of the security functions to be transferred to the client's equipment. Modern home gateways are sophisticated embedded systems with reasonable computational capabilities, already available with no (additional) costs and located at a privileged point of the network – between the home network and the access network.

The fundamental operation concept of the proposed security model is based on the idea of turning the RGW into an active security appliance, able to collect information and statistics related to network traffic that can be used as input for semi-centralized (supervised on unsupervised) training and decision inference and correlation systems, which can take actions to secure the network by distributing countermeasures to the RGWs and other network devices (like selectively filtering of network traffic in response to possible attacks). Optionally, RGWs might have some level of autonomy, with local aggregation, correlation and decision-making mechanisms. Figure 7.2 illustrates the general layout of the proposed framework.

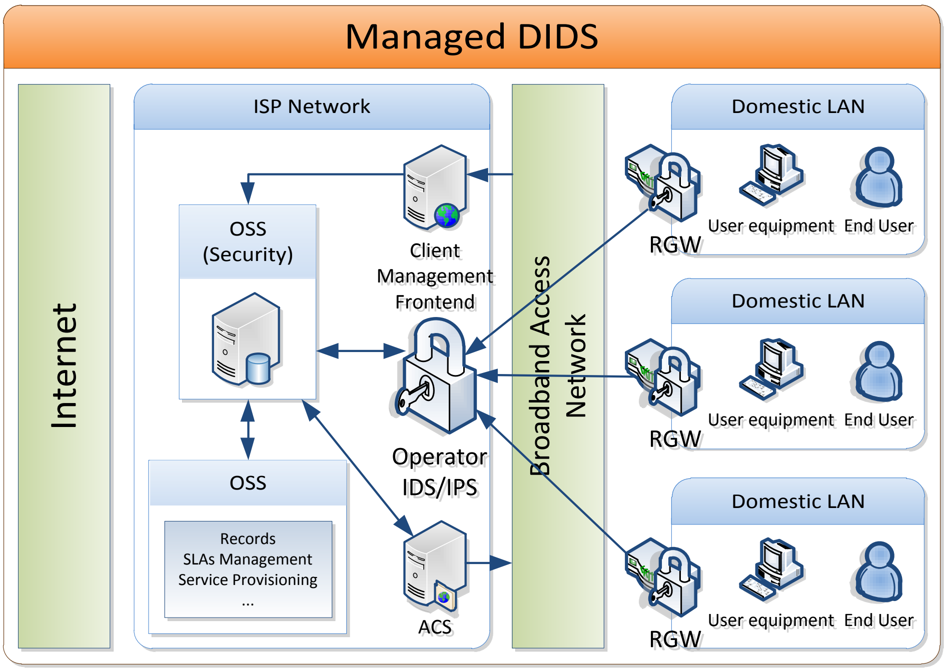

Figure 7.2: Proposed RGW-based distributed IDS solution

This framework supports the concept of operator-assisted home LAN security, provided as a service based in a shared management model, where the ISP extends its reach to the residential gateway and (optionally) even further up to the customers LAN:

- Allowing the ISP to reach the RGW alone creates the conditions to deploy "intelligent" security and monitoring mechanisms at the RGW level (therefore, in the strategic border between the ISP and the customer networks instead of being centralized in the ISP infrastructure, with added scalability benefits) without interfering with the customer LAN. Because the RGW mechanisms are being exclusively used to monitor and act on the traffic between the ISP and the client, the customer's privacy is not compromised: the ISP could always do a similar monitoring inside its network, albeit with significantly higher costs and scalability limitations.
- If the user explicitly authorizes the ISP to do so (for instance by means of End-User License Agreements similar to those already used by antivirus and other security applications), the

RGW can extend its reach to the internal LAN, probing for a wide range of vulnerabilities via optional customer monitoring mechanisms, interfacing with operating system integrated resources to gather information about installed updates and active security parameters.

This approach introduces a paradigm shift by making use of a distributed structure in which domestic gateways act as the frontline. Table 7.1 compares this new approach with traditional solutions.

Table 7.1: Security model paradigm comparison

Traditional security model	Shared management security model
ISP is restricted to its own network	Customers define the ISP's scope of influence. Minimum ISP reach extended to RGW boundaries. Customer LAN can be monitored by the RGW if explicitly accepted by the customer
Static barriers inside the ISP infrastructure traffic/bandwidth limiting and/or shaping	RGW-level mechanisms allow more granular traffic control
RGW has limited capabilities, security effectiveness heavily depends on customer knowledge	RGW integrated into the security infrastructure, shared management between the ISP and customer
Traffic monitoring is possible at the ISP infrastructure, with scalability limitations	Traffic monitoring at the RGW level (possibly complemented by monitoring at ISP level)
Available detection and contention mechanisms in case of distributed attack are limited, slow and mostly ineffective	Capable of detecting attack in progress before it spreads further, via local RGW probes. RGW-embedded capabilities allow for deployment of countermeasures in a fast, effective way.

The architecture of the proposed DDIS has three fundamental component entities (Figure 7.3):

- **Distributed monitoring:** performed by specialized components at the RGW level that act as probes, monitoring service activity and network traffic.
- **Distributed inference and correlation:** inference and correlation mechanisms exist at the ISP level, to process the information from the RGWs. For enhanced scalability, local inference and correlation mechanisms might also exist at the RGW level, filtering and aggregating relevant events before they are forwarded to the upper level.
- **Security policies and countermeasure enforcing:** through distributed "push-back" mechanisms based on autonomous or semi-autonomous decision processes. Also, the RGW may have some limited decision-making abilities using an optional event correlation and inference mechanism.

It should be stressed that there will always be customers whose gateways will not cooperate with the IDS of the ISP (e.g. customers that replace the ISP-provided RGWs with other models) and there will always be a risk of having compromised gateways inside the structure. Consequently the ISP may never absolutely trust the information provided by each RGW/probe. The ISP platform must have enough flexibility to simultaneously deal with customers with cooperating gateways (presumably the vast majority), customers with compromised gateways and customers with no unmanaged gateways (for instance selectively monitoring network traffic within its own access network). In spite of that, and from a global standpoint, the potential benefits in terms of granularity and scale are considerable.

The proposed architecture is not only based on the idea of using the RGW to extend the reach of the ISP IDS. It also tries to effectively improve the articulation between the home network and the ISP network. To achieve this, the security policies adopted by the distributed IDS take into account the customer profile (e.g. user records, subscribed services, typical usage patterns) and they also allow the customer some level of customization – by means of a portal where one can, for example, configure explicit support for some applications or specify more detailed usage profiles and policies (e.g. parental control of Web contents).

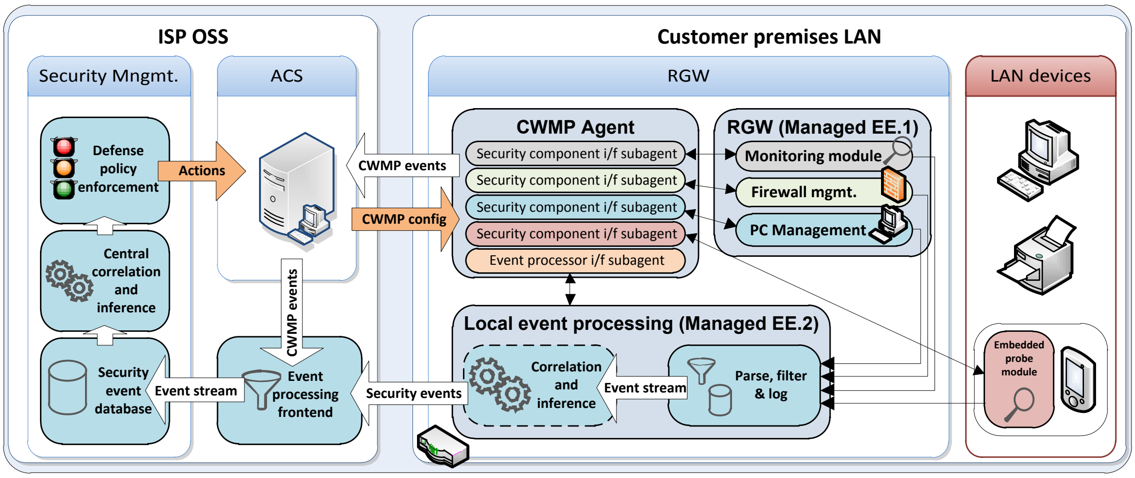
ISP OSS
Security Mngmt.
ACS
Customer premises LAN
RGW
LAN devices
Defense policy enforcement
Actions
Central correlation and inference
Security event database
Event stream
Event processing frontend
CWMP events
CWMP config
CWMP events
Security events
CWMP Agent
Security component i/f subagent
Security component i/f subagent
Security component i/f subagent
Security component i/f subagent
Event processor i/f subagent
RGW (Managed EE.1)
Monitoring module
Firewall mgmt.
PC Management
Local event processing (Managed EE.2)
Correlation and inference
Event stream
Parse, filter & log
Embedded probe module

Figure 7.3: Generic (and simplified) architecture for the proposed DIDS solution

The proposed architecture includes dynamically deployable components that provide security-related services at the RGW level (Figure 7.3). These components are remotely managed using CWMP (very much like other services discussed in previous chapters), building on TR-157 component management mechanisms to provide module management on the Execution Environment (EE) of the RGW. Thanks to the extensible agent framework introduced in Chapter 3, security components might even be hosted on other devices on the customer LAN, being proxied through the RGW. These entities may have a passive or active role, ranging from monitoring/detection (probes) to actuation/decision capabilities based on analysis of the data flows or other information that can be obtained from the databases of the ISP.

The security and event processing components are maintained on different execution environments (EE.1 and EE.2), for isolation and ease of management. This also has the benefit of enabling separate resource throttling policies for each of the EE, avoiding RGW resource starvation. Component package management is performed using TR-157 mechanisms, making it possible to install, update or replace security or event processing modules, with specific component configuration being performed through CWMP data model extensions mapped on the RGW instantiated data model.

The proposed model is not fully distributed by nature – even if it concedes RGWs the ability to embed autonomous detection and action mechanisms, all the operation is centrally orchestrated on the ISP infrastructure. A management infrastructure is kept on the provider's side to coordinate the various participants on this process, orchestrating its operation based on the correlation of the pieces of information collected from the ISP network and from the different RWGs. The existence of a profile management backend makes it relatively simple for the ISP to deploy new components, rules or configurations for extended user groups, taking on account the specific profiles of the installed equipment (RGW models, set-top-box models, etc.) and contracted services. Also, while some control remains centralized at the ISP level, this architecture is compatible with both supervised and unsupervised anomaly detection systems, as it was not designed with a single IDS paradigm in mind.

Additionally, this architecture may coexist and cooperate with existing traditional IDS mechanisms – in fact, some of the limitations induced by the already mentioned need to support non-cooperating RGWs can only be solved by integrating additional IDS/IPS mechanisms inside the ISP network.

7.2.1 Security mechanisms and event management

The detection and correct handling of security incidents or anomalies is a key element in the proposed architecture. The detection and treatment of security events may occur at two different levels (see Figure 7.4):

- **At local level (RGW).** For efficiency and scalability reasons, and in order to allow a high granularity in the monitoring process, the RGW may host a local event correlation engine (fed by events generated by traffic analysis tools, such as intrusion detection systems, port scan detectors or log records). Optionally, all events might also be pre-processed by an internal RGW correlation engine, making it possible to activate counter-measures at the local level, based on the application of rules and procedures belonging to the security mechanisms of the RGW itself and/or with notification to the ISP level.
- **At the ISP level.** The event correlation engine that exists at the ISP level processes the events received from the multiple involved RGWs. This allows, for example, the detection of combined attacks either affecting or coming from several customers of the ISP. The provider can react to these events taking preventive measures on its network (coordinating traditional IDS mechanisms) and/or adjusting the RGW configurations.

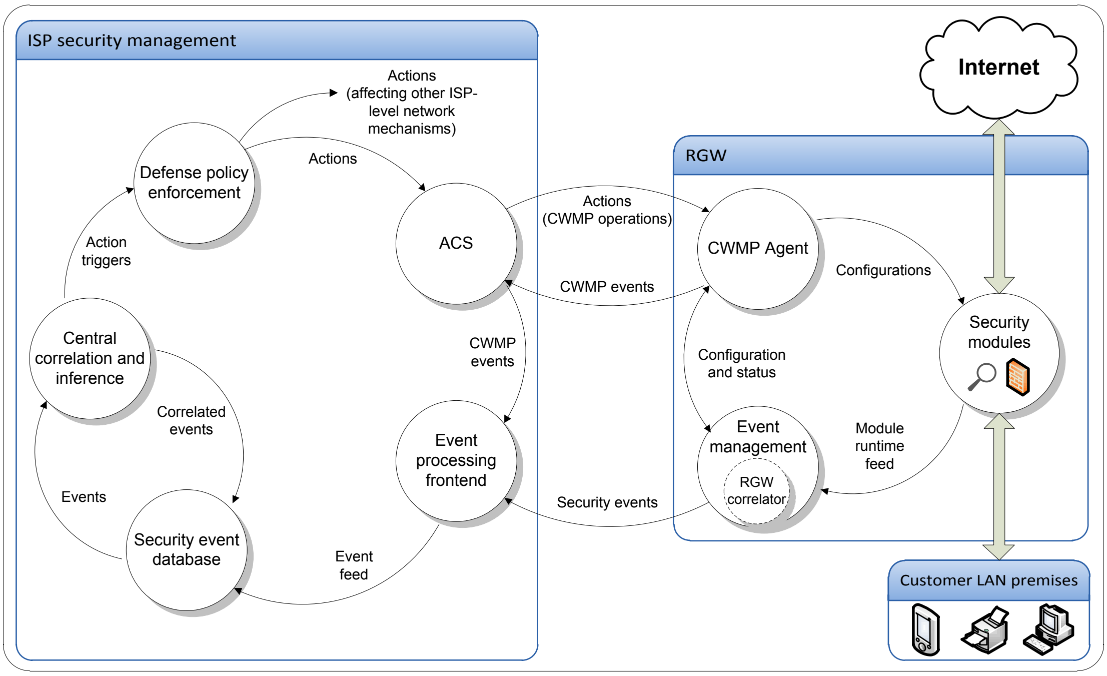

Figure 7.4: Generic model of the event-based decision making process

Locally generated events can be treated in three different ways:

- Events are filtered (using simple pattern-matching mechanisms) and ignored either because of their naturally harmless nature or because their threat level is ineffective against countermeasures already in place.
- Events are locally correlated (if the system has enough processing capabilities to do such), originating a set of reaction measures which can be locally decided and executed by the RGW. This correlation process might be based on supervised or unsupervised methods, with a variable degree of autonomy. This has the added benefit of enhancing scalability, since the ISP will only receive processed (inferred, correlated and composed) events that are considered of interest, resulting in a reduced amount of event traffic in comparison with an event stream filtered using simpler and less sophisticated (i.e., patter-matching) mechanisms.
- Events are aggregated and reported to the ISP, in order to feed the global correlator and inference mechanisms. Depending of the outcome of the global correlation and inference, the

system can engage the proper reaction measures, either at the access network (backbone traffic filtering, sinkholes [Greene2002]) and/or at local RGWs (e.g. reactive or preventive blockage of selected traffic), through mass distribution of configuration parameters.

The proposed platform is able to deal with two levels of operation (RGW/microscopic, ISP/macroscopic), while covering a broad scope, from the provider's network to the entrance/exit points of the clients' network. As an example, let's consider a group of customer networks that have just been attacked and infected by a Trojan and are taking part in a synchronized DDoS attack (see Figure 7.5). The RGWs might directly report the abnormal activity pattern or notify the provider through a locally correlated/inferred event. On the provider's side, when the correlation mechanism detects a global pattern, it performs the distribution of new security rules (e.g. to restrict the usage of specific TCP/IP ports at the RGW level) to prevent the propagation of the attack to other clients and contain the outgoing traffic from affected customers. This example shows why the distributed event management system is relevant in the context of the presented architecture.

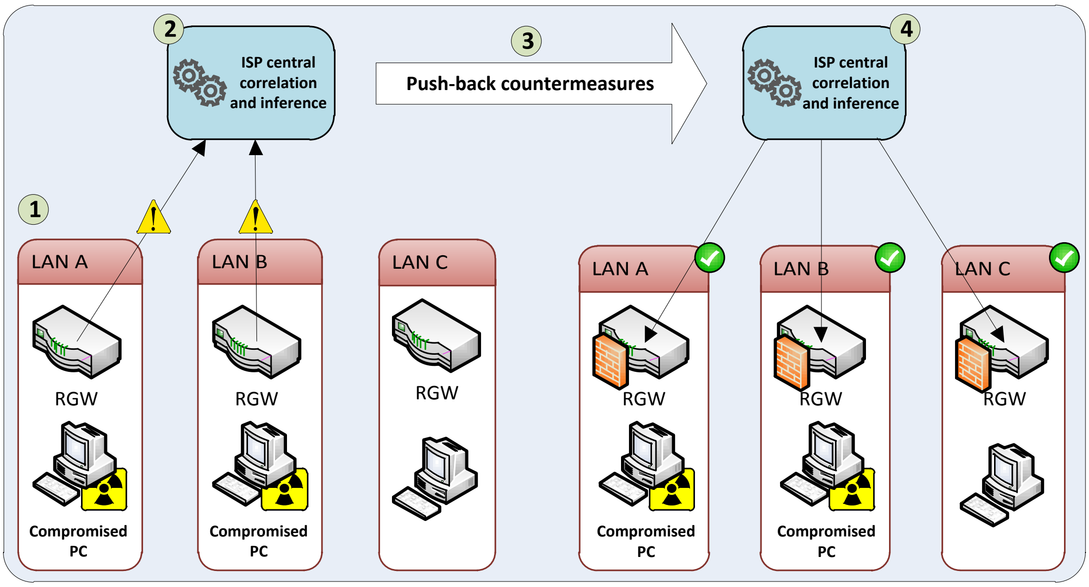

Figure 7.5: Operation example with distributed push-back countermeasure

7.2.2 Compatibility with existent IDS paradigms

DARPA's Common Intrusion Detection Framework Architecture (CIDF) [DARPA1998, Staniford-chen1998, Tung2001, Kahn1998] was an effort to develop standard protocols and APIs, allowing intrusion detection systems to share information and resources. Many of the ideas developed within the CIDF effort were also the basis for IETFs Intrusion Detection Working Group (IDWG) work, such as the Intrusion Detection Message Exchange Format (IDMEF [RFC4765]), used for interchange of security events in the proof-of-concept platform discussed on Section 7.3.

One of the most noteworthy results of the CIDF effort consisted on the definition of a generic IDS architecture, which builds on discrete functional blocks with clearly defined functions:

- **E-blocks (event-boxes):** generic sensor elements that acquire information to be processed by other functional blocks. Network traffic probes, for instance, are an example of such elements.
- **D-blocks (database-blocks):** generic data persistence elements which store and persist information from *E-blocks*, for subsequent processing. Without these elements, IDS architectures would be limited to a simple real-time reactive operation.

- **A-blocks (analysis-boxes):** generic processing modules which analyze, correlate and infer information from D, E and even other A-blocks to detect anomalies or suspicious behavior, being able to generate alarms.
- **R-blocks (reactive-blocks):** generic action enforcement blocks, which implement specific actions and countermeasures to deter or avoid a threat. An R-box might be fed by D and A-boxes.

The CIDF model is of particular interest because it offers a generic decomposition tool to analyze the modules constituting the proposed architecture (Figure 7.6).

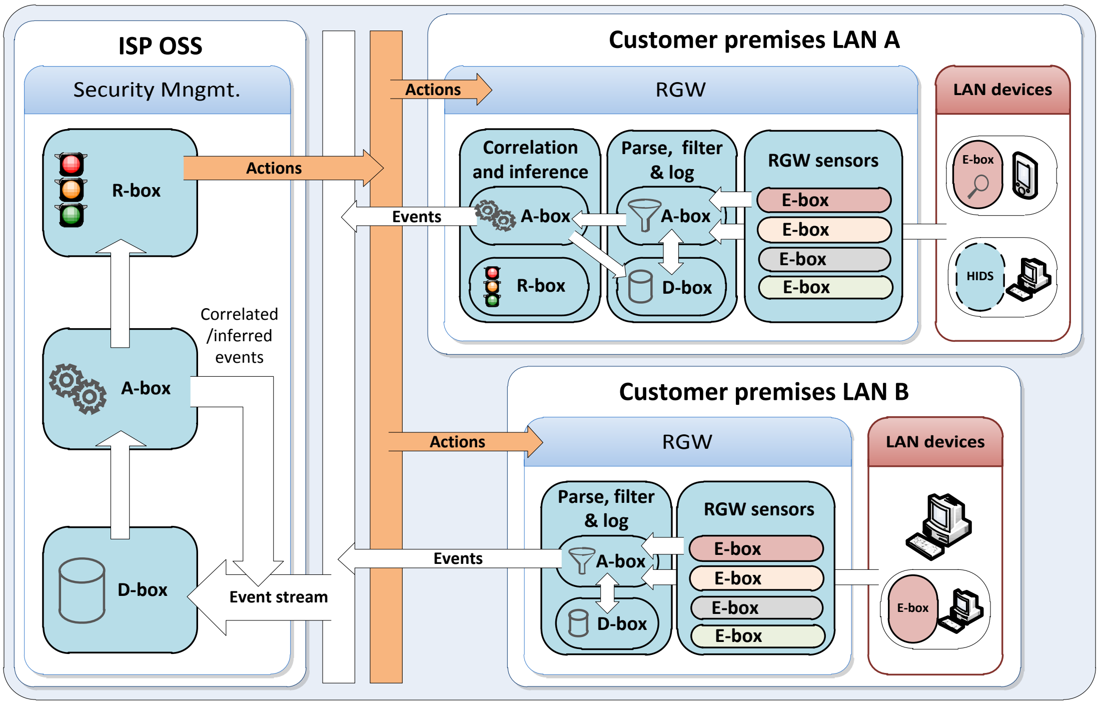

Figure 7.6: CIDF-inspired functional decomposition of the proposed architecture

The main innovation of the proposed model resides in the placement of functional blocks (which in traditional IDS architectures reside exclusively inside the ISP infrastructure) within the RGW, with distributed coordination. While the basic architecture supports basic A and D-blocks on the RGW, it may optionally go to the extent of hosting a complete mini-IDS functional model on each RGW (as is the case of the *Customer A* RGW, in Figure 7.6). The proposed architecture offers a slight twist in the CIDF model, since it considers that D-boxes might also be fed by the output of A-blocks, enabling complex correlation patterns over extended time periods.

The CIDF-inspired decomposition of the proposed architecture also shows that its operation model was deliberately conceived to enable integration with the majority of the IDS/IPS paradigms in existence, to the point of allowing for combined use of both Host (HIDS) and Network IDS (NIDS) techniques to enhance event detection capabilities.

A-boxes frequently contain sophisticated analysis, correlation or inference mechanisms which commonly distinguish IDS paradigms from each other, whose implementation is the subject of intensive study in the last years. Existing methodologies are usually classified in two main groups [Garcia-Teodoro2009, Douglieris2003]:

- **Signature/fingerprint-based detection** is based on characteristics extracted from traffic flows, such as statistical variations of specific parameters (frequently related to traffic volume) or patterns such as the distribution of involved IP addresses or ports. These methods are unsuccessful in identifying unknown anomalies, requiring supervised analysis and/or training to incorporate new signatures in the IDS – this has the side effect of letting the network unprotected from rogue threats for a variable amount of time. Tools such as Snort [Snort] fall into this category when used in its simplest configuration (without plugins as SPADE [Staniford2002] or OSSEC [OSSEC]).

- **Anomaly-based detection** consists on finding deviant behavior from established "normal" usage patterns. Several techniques have been researched on this field, based on statistical, knowledge-based or machine-learning techniques, using IP-flows, single-link or network-wide data with signal-processing techniques (such as wavelets) [Barford2002, Brutlag2000], Kalman filters [Soule2005], PCA (Principal Component Analysis) [Lakhina2004] or Sketches [Khrishnamurthy20073, Dewaele2007]. However, there are two fundamentally different approaches to anomaly detection which distinguish one from another in what respects to their autonomy.

  Anomaly detection based on supervised learning requires training based on labeled traffic, which is normally inconvenient to produce. This helps establishing a baseline model which corresponds to "normal" traffic – any deviating pattern is considered anomalous (in practice this corresponds to behavioral profiling). This method is able to detect unknown anomalies and rogue threats – however the training process is time-consuming and requires a regular feed of anomaly-free data sets (a complex and error-prone task) which must be kept up to date to be effective. The URCA tool [Silveira2010], for instance, uses both signature-based and supervised learning techniques. Another example is presented in [Perdisci2010].

  Autonomous/unsupervised anomaly detection is a somewhat recent trend, based on the assumption that an IDS should not rely on previous knowledge to operate, rather being able to autonomously detect and characterize threats. While some authors [Mazel2011, Mazel2011a] propose that modern networks should rely on completely unsupervised detection and reaction methods, common sense dictates otherwise as a failure could rend inoperable significant sections of the network infrastructure (due to automatic misjudgment and consequent decision). Opinions aside, the proposed platform is nevertheless able to incorporate such mechanisms, providing access to both fine-grained (single link, on a single RGW) and wide-scope mechanisms to monitor/correlate traffic flow information and enforce security policies. Botminer [Gu2008] is an example of a tool that uses these methods, performing cross-cluster correlation to identify hosts with similar suspicious activity patterns.

While virtually all A-box techniques fit into the proposed architecture, their choice and positioning criteria must obey some restrictions, especially when used at the RGW-level. The home gateway is an embedded systems platform with reasonable but limited computing resources. For instance, among anomaly-detection methods for RGWs, those based on real-time IP flow analysis using time-slots are found to be particularly adaptable and flexible enough for integration on RGW A-boxes.

Section 7.3 shows an example of how signature-based systems can be incorporated on a proof-of-concept platform built for evaluation purposes. However, anomaly-based detection systems equally fit in the proposed operation model, whether supervised or autonomous. Their integration within this platform is also a straightforward process, whether using supervised and autonomous detection methods. Still, some observations should be pointed out:

- Integration of supervised anomaly detection A-boxes is similar to signature-based methods, without any significant difference from the implementation to be described in Section 7.3. As for training, the proposed architecture does offer the needed flexibility, since existing mechanisms could be easily repurposed or adapted for training purposes.
- As for unsupervised detection schemes, the majority of published work on the subject is based on sub-space and inter-space clustering anomaly detection methods. Since the proposed platform enables the deployment of mechanisms for local pre-processing of single-link traffic at the RGW, it is possible to perform clustering analysis at each RGW, for instance using different flow levels for time series analysis (as proposed by [Mazel2011]). Subsequent

correlation of anomalies from multiple RGWs might be performed at the ISP level, enabling the possibility of network-wide meta-correlation. As an example, [Mazel2011] proposes performing anomaly correlation from single-link multiple flow aggregations to estimate their impact by finding if it they are visible at different flow levels – this idea could be further extended to network-wide scope if performed at the ISP level (as suggested by [Lakhina2004a]). This concept might also be applied for anomaly characterization and autonomous reaction techniques, in which case R-boxes must also be able to generate the adequate action for an autonomously generated threat response.

Other techniques, which are of particular use in HIDS systems, such as target monitoring (used by tools such as tripwire [Kim1994], which control and report changes on internal system files and parameters) can also be supported using OSSEC (although they are not covered in this discussion). Several authors classify some hybrid approaches as new IDS categories, such as the case of *stealth probes* [Marinova-Boncheva2007], which consist of global correlation and inference procedures carried along prolonged periods of time (months) to detect attacks prepared and executed over an extended time frame – these techniques also fit naturally within the scope of the proposed framework.

Individually, each IDS category has its particular set of benefits and drawbacks, which can be overcome with a combination of different techniques for correlation of data obtained from signature-based and anomaly-based detection mechanisms. By being neutral in terms of IDS-paradigms, the proposed architecture is able to cope with autonomous and semi-autonomous, supervised or unsupervised anomaly detection and reaction mechanisms. The decision of which combination to use is up to the third-party security management provider (eventually, the ISP) which may choose the combination that best fits its purpose and the RGW capabilities (if optional correlation and inference mechanisms are deployed on the RGW).

7.3 Evaluation

In order to validate the proposed framework, a proof-of-concept implementation was built. This implementation was then evaluated, from a perspective of performance and functional behavior, to assess the feasibility of the proposed solution. In this Section we present the key characteristics of the developed prototype and discuss the evaluation results.

7.3.1 Proof-of-concept implementation

To describe the proof-of-concept implementation we start by introducing the Prelude IDS Platform (which became a key building block for the event management and correlation engines, both as RGW level and ISP level). Next, we describe how the security-related components were introduced in the RGW, and finally we discuss the infrastructure developed in the side of the ISP.

Prelude as a key building block

The Prelude IDS platform [Vandoorselaere2008] was used for event management and correlation. Prelude is a hybrid IDS platform that combines host-based (HIDS) and network-based (NIDS) capabilities, while also including sophisticated programmable event processing and correlation mechanisms [Chifflier2008]. It is extensible and includes a library with APIs for various programming languages (*libprelude*). In the Prelude platform all security events are encoded using IDMEF, a standard format that can, for instance, be easily forwarded to SQL databases.

Figure 7.7 illustrates the Prelude IDS platform architecture. The central entity for event processing is the Prelude Manager, which accepts data from sensors managed by the *libprelude* library. Sensors are small agents capable of collecting information from several sources, like the Prelude LML (*Log Management Lackey*) that generates events from processing system logs (generating IDMEF events in the form of IDMEF messages that are sent to the event management module through secure SSL connections). Integration with the OSSEC HIDS is also possible, thanks to a plugin agent that directly generates IDMEF events, enabling cross-platform event gathering (which could be used to gather information from Windows-based desktops in the home network, for instance).

The event correlation engine (Prelude Correlator) is able to process IDMEF messages and generate alerts (also in the format of IDMEF messages), using a set of rules described using the LUA [LUA] or Python [Python] programming languages.

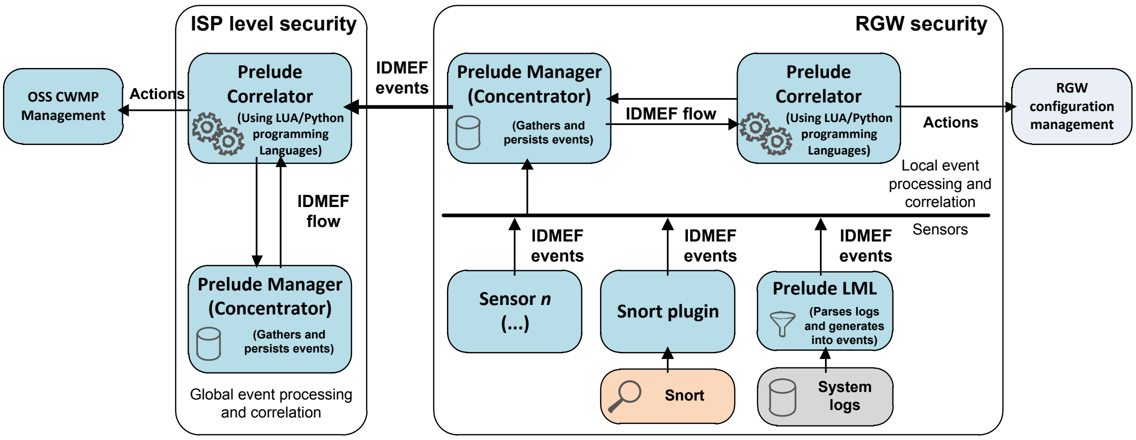
ISP level security
RGW security
OSS CWMP Management
Actions
Prelude Correlator (Using LUA/Python programming Languages)
IDMEF events
Prelude Manager (Concentrator) (Gathers and persists events)
IDMEF flow
Prelude Correlator (Using LUA/Python programming Languages)
Actions
RGW configuration management
Local event processing and correlation
Sensors
IDMEF flow
Prelude Manager (Concentrator) (Gathers and persists events)
IDMEF events
IDMEF events
IDMEF events
Sensor *n* (...)
Snort plugin
Prelude LML (Parses logs and generates into events)
Snort
System logs
Global event processing and correlation

Figure 7.7: The Prelude IDS platform

RGW Components

Figure 7.8 presents the architecture of the platform, from the RGW perspective. The three main modules are the **Security System** (modules on the EE.1), the **Event Management** (modules on the EE.2 – including the optional local correlation engine) and the embedded RGW **Configuration Management** system, which integrates with the CWMP agent and complies with TR-157.

Customer network monitoring and LAN device HIDS modules were omitted, albeit they could be easily implemented, for instance using the WMI subagent from Chapter 4 (for Windows device management), the OSSEC HIDS or using the UPnP module from Chapter 5 (with full GENA integration, as showcased in the application scenarios from Chapter 3).

The **Security System** (EE.1) aggregates the environmental active defense and network/activity monitoring mechanisms. It includes components for firewall management [Netfilter], web filtering (proxy [Squid] and parental controls [Squidguard]), intrusion defense/protection system [Snort], portscan detector [Scanlogd] and UPnP IGD control service [LinuxIGD]. Some of these components will play a passive role (portscan detectors, IDS), only generating events to feed the local event management engine. Others will play an active role, with its configuration being dynamically modifiable by local or remote (ISP) decisions, for instance in reaction to security incidents. Configurations (rules, ACLs, etc.) may be remotely controlled by the ISP, using CWMP.

Event management is hosted on EE.2, running the Prelude components for local event processing and correlation.

The **Configuration Management** system is responsible for managing the RGW configurations (installed services, active configuration profiles). Configuration updates can be distributed remotely by the ISP or by the RGW itself (as a reaction to a security incident dealt locally by the local RGW correlation engine). Package configuration wraps the Debian *apt* [Debian] package management system within a TR-157-compliant component management framework.

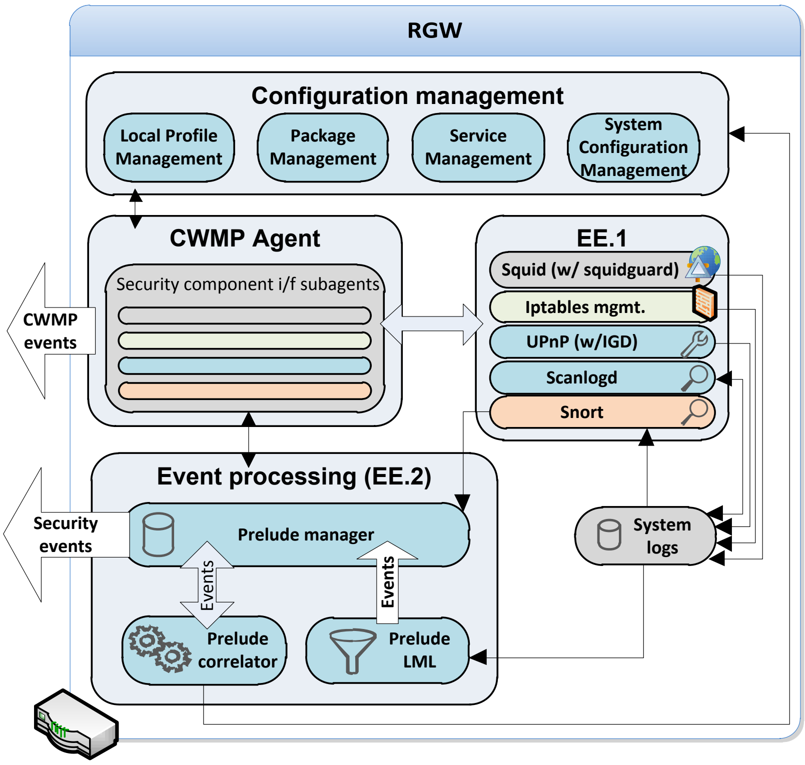

Figure 7.8: Prototype security platform (RGW-hosted components)

Security-management infrastructure on the ISP side

Figure 7.9 presents the prototype architecture, from the ISP standpoint. The main components are the Security Event Management, the CWMP Management and the Profile Management.

The **Security Event Management** module is a security event correlation system (also based on the Prelude IDS platform), fed by the local event management modules of each RGW and also by events detected by sensors positioned in the ISPs own network. It provides event correlation and activation of the adequate orchestrated actions (traffic filtering on the ISP own network level and/or firewall configuration updates distributed to selected RGWs). Relevant CWMP events are also fed in the event queuing and processing frontend.

The **CWMP Management** system (ACS + Profile Management) deals with the remote configuration of the RGW units (distribution of the applications and configurations to be used by each RGW). It also monitors of the operation of each RGW, detecting and reacting to operation failures. These functions fall within the scope of the CWMP protocol framework, which natively provides the means to manage configurations in a broader, generic perspective (e.g. updates, inventory) and is also used as a mechanism of remote intervention by the ISP, in order to dynamically update/change the configuration of the RGW, in response to or prevention of security incidents.

The **Profile Management** system ensures the maintenance of a database with equipment (CPE manufacturers, models and versions installed in each location) and user profiles (subscribed services, preferences defined in the customer portal, etc.). These profiles are essential to define which configurations must be sent to each RGW.

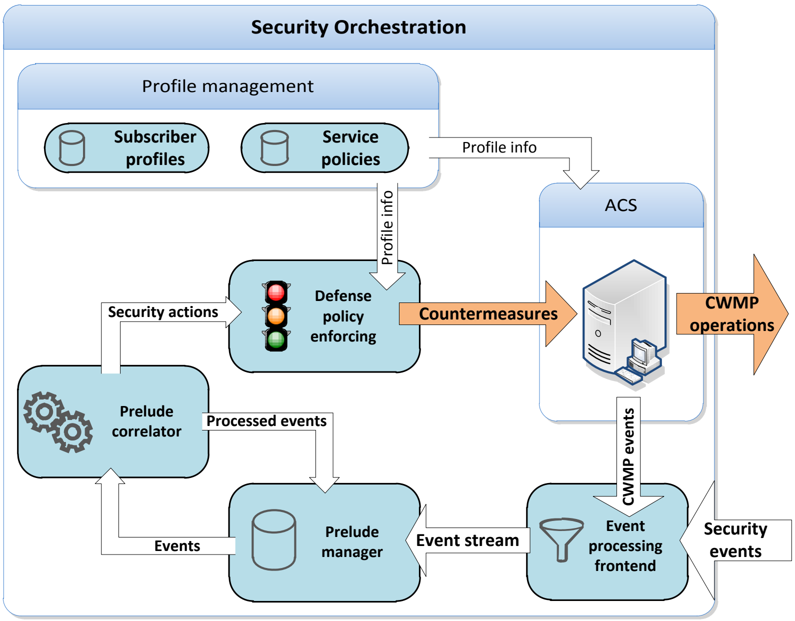

Figure 7.9: Prototype security platform (ISP infrastructure)

7.3.2 Validation

The validation process was constrained due to the practical impossibility of deploying the prototype in a real world scenario with a significant number of nodes (at least a few thousands). The lack of simulation models and data capable of faithfully characterizing typical user and attacker behavior patterns was also an obstacle – there are no established mathematical network models for triple-play traffic, in order to simulate the heterogeneous and elastic nature of such environments [Gudkova2011].

The evaluation started with a number of simple functional tests in small-scale scenarios (5 to 10 gateways), checking how the platform reacted to injected network attacks. Those tests successfully demonstrated the capability to react to typical security incidents, even when those incidents were only possible to detect combining several local sources (Snort, Scanlogd, IGD) or correlating events from multiple RGWs. However, those tests did not assess the behavior of the platform in larger, real world scenarios.

Considering the general lack of traffic and security models (or even faithful network traces) for broadband access networks, an empirical method was devised: two distinct testbeds reproducing a simple home network were created and connected to the internet, one of them adequately protected (updated versions of operating system and applications) and the other one left with intentionally vulnerable desktops (Windows machines without security updates).

Using both testbeds with typical applications (e-Mail, messaging, Web, P2P, FTP, VoIP, etc.) it was possible to collect and analyze a set of network traces that were later used as representations of broadband internet usage. Considering the specific scope of our work, this approach was considered as fairly acceptable, despite its shortcomings.

In parallel to that process, the global event manager component was subject to performance load testing (using synthetic conditions), in order to assess its capacity in terms of RGW event correlation and reporting.

Finally, combining the data collected on those two processes made it possible to analytically estimate the overall capacity of the built prototype.

Next, we discuss each of these three steps.

Empiric traffic characterization

The original intention was to make use of third-party network traces (such as the ones publicly available from [MIT] and [UMass]) as a basis to characterize typical network usage of domestic broadband users. Soon became evident that such traces were inadequate for the intended purposes, mainly because of two factors: age (most captures are significantly old, predating the days of widespread P2P, VoIP and similar services) and the fact that most of them contained data captured at the backbone level, being very difficult to extract traces originated from specific domestic LANs. Furthermore, since most captured traces are anonymized, it becomes almost impossible to establish adequate traffic correlation patterns.

Considering this situation, a laboratory testbed was built in order to emulate conditions similar to those of typical domestic LANs. This network was directly connected to the Internet (without added routers or firewalls in-between) and was used to support two different data collection scenarios (Figure 7.10):

- **Domestic LAN with 3 PCs** using Windows XP as the Operating System, with current and updated antivirus software, active firewall and up-to-date OS patches and service packs.

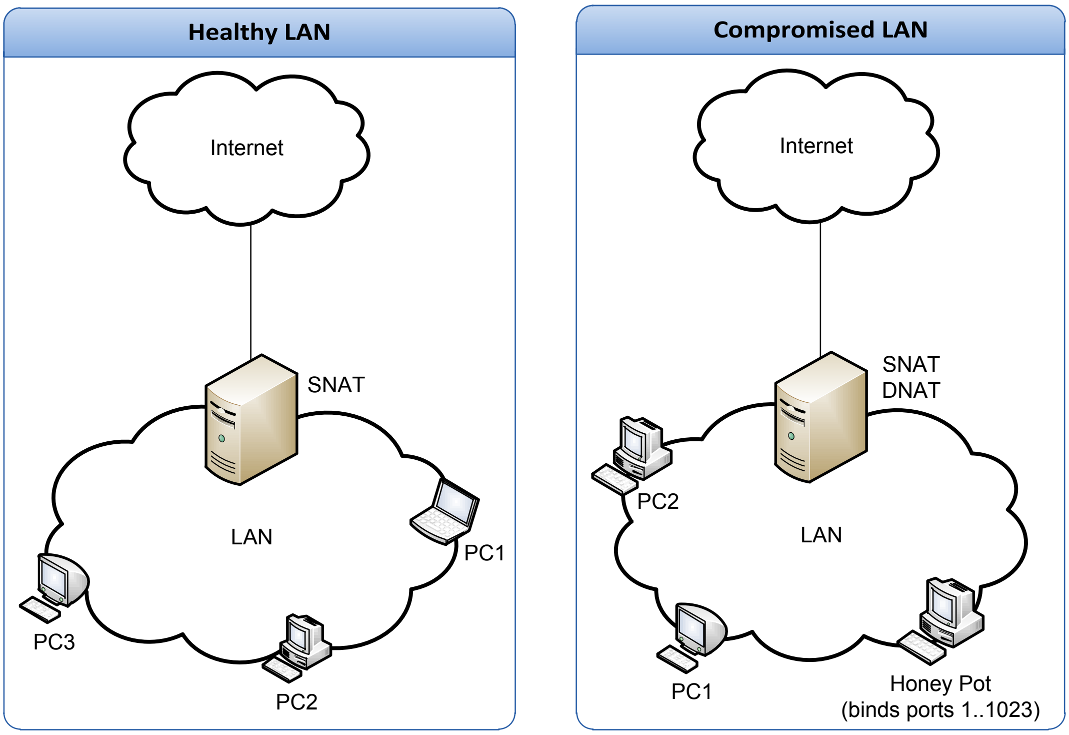

Figure 7.10: Testbed scenarios

- **The same domestic LAN with 2 PCs and 1 honeypot:** all the PCs are using Windows XP in the same condition as in the previous scenario. The honeypot PC has the same OS as installed out-of-the box, without any updates or protection measures configured and/or installed, in order to behave like a *honeypot* [Pickett2003].

Three voluntary users were asked to make normal use of the PCs, in consecutive periods of up to 7 hours. They were asked to use the machines as they normally do, by surfing the web, checking their e-mail, using FTP and P2P applications. No special precautions were taken in order to guarantee identical usage patterns in each sampling period (in a specific day a user could generate more P2P traffic than in previous days) – yet, further observation showed some degree of consistency among the sampling periods, in this context of empiric validation.

The capture of network traffic capture between the domestic LAN and the Internet was done with the *TCPDump* [TCPDump] utility. The security events/alarms detected by the router's

Defense System (see Figure 7.8, components of EE.1) during the sampling periods were also registered.

To evaluate whether a given event constitutes or not a relevant security threat is something that depends on the context in which it was generated (vulnerability level of the systems that were involved, services available on the network, etc.). For evaluation purposes, a worst case scenario was adopted: all events generated by the Defense System are considered as relevant, independently of the results of their specific correlation analysis.

The results of the network traffic captures, summarized in Table 7.2, show the differences between the "healthy" and the "compromised" scenarios. The "compromised" scenario roughly generates 60 times more security events – overloading the platform and also "attracting" more than twice the network traffic, as a side-effect of exploits and similar threats. The percentage of IP addresses involved in security events is very low on both cases. This shows that most of the attack potential comes from a reduced set of external sources, in a repetitive pattern.

Table 7.2: Traffic characteristics (for a period of 7 hours)

	Healthy scenario	**Compromised (honeypot) scenario**
Security-related events per hour	29	1.686
# of external IP addresses involved in the traffic capture	50,960	278,433
% of external IP addresses involved in security-related events	0,19%	0,32%
# of events per 100MB of network traffic	7	102
Average traffic rate, not considering IPTV	178 Kbps	470 Kbps

Scalability

The tests that were performed showed that traffic analysis and event processing at the level of each RGW did not suffer from capacity limitations. For the considered levels of traffic, there appears to be no significant capacity bottleneck at local level (and even if there was, hardware upgrades would likely solve this).

For this reason, in order to assess the scalability of the proposed platform, special attention was given to the event processing component at the ISP level, since it is probably the most sensible to load (especially the correlation engine).

Tests were conducted using two different machines: one with the role of simulating the domestic RGW and the other the ISP-level service components. Table 7.3 presents the characteristics of such machines, which are considerably dated even when compared with the current generation of desktop equipment – production, carrier-level ISP servers are substantially more capable that the kind of equipment used in this validation procedure.

Table 7.3: Hardware characteristics

	Domestic RGW	**ISP Server**
Processor	Pentium 4 CPU 3.00GHz	Intel Pentium 4 CPU 3.00GHz
Memory	512 MB	2 GB
L1/L2 Cache	16KB / 2MB	16KB / 2MB
Operating System	Ubuntu Linux	Ubuntu Linux
Network Interfaces	82541GI Gigabit Ethernet	3Com 3c905C-TX Fast Ethernet

To assess the event manager at the ISP level, a group of synthetic tests was executed as per the following methodology:

- Previously, several groups of events were generated at the RGW level, configured to directly forward all events to the ISP event processing system. Those events were artificially generated (without using the defense system) and were not subject to any local processing, therefore avoiding any limitations of the RGWs that feed the ISP event manager. These events, encoded in the IDMEF format, are sent in bursts through secure channels (as per the standard specification).

- The event manager at the ISP level was configured to apply correlation procedures to all the events received, in order to also stress and evaluate the correlation module and not only the event reception and storage capabilities. This configuration is a worst-case scenario since it demands more processing power at the ISP level than it would be needed in a normal usage context, where it is normal for several events not to be correlated because of their specific nature.
- The processing power for each group of events was measured in the time interval between the correlation of the first and last event. Figure 7.11 presents the results obtained for sets of 100, 1,000 and 10,000 events.

Measurements shown in Figure 7.11 demonstrate that the average capacity of the ISP event manager, using the specified hardware and subjecting all events to a set of demanding correlation rules, is relatively stable and close to 92 IDMEF messages (or events) per second. It also becomes clear that the system maintains the same processing capacity in transient overload situations, thus being possible to use buffers to accommodate activity peaks without losing events. The estimated message processing rate will be used to extrapolate the number of clients supported by the platform in the conditions here described.

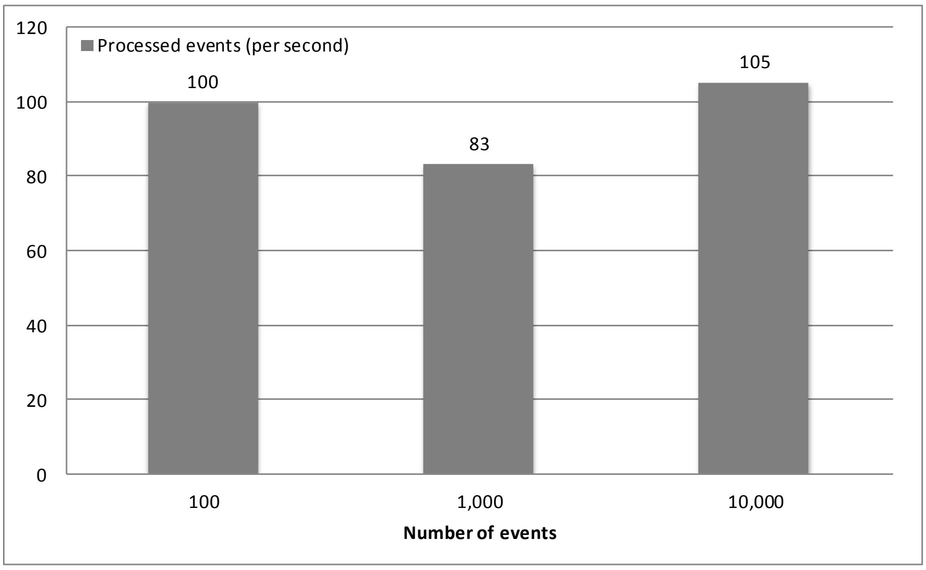

Figure 7.11: IDMEF event processing capacity

Estimated Capacity of the Prototype (without event filtering)

In order to estimate the number of customers supported at the ISP level, the empirically estimated capacity of the ISP event manager (92.6 events per second – Figure 7.11) was crossed with the empirical characterization of traffic (Table 7.2).

Figure 7.12 shows the projected capacity of the presented platform, considering three scenarios:

- 100% "healthy" customers (the traffic handled by each RGW generates around 29 security events per hour – see Table 7.2);
- a mix of "healthy" customers (80%) and "compromised" customers (20%);
- 100% "compromised" clients.

It should be mentioned that these results are very conservative, since they assume there is no filtering, aggregation and correlation at local level (as discussed next). Furthermore, carrier-level servers are expected to have much larger capacity that the old hardware used as reference.

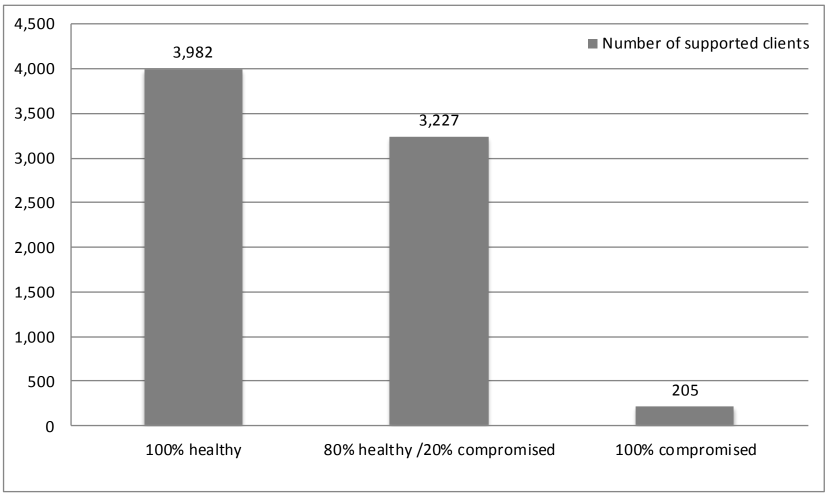

Figure 7.12: Estimated capacity of the ISP-side Infraestructure

The communication between domestic RGWs and the ISP event manager can also be responsible for introducing some network load. Considering the number of events generated per hour and the average size of an IDMEF message (estimated as 4KB, from traffic capture traces), an average estimation of the IDMEF traffic flowing through the domestic LAN is presented in Figure 7.13, showing that that they do not constitute overhead: with 1,000 RGWs and the 80/20 scenario it would result in an aggregate debit of 3.2Kbps.

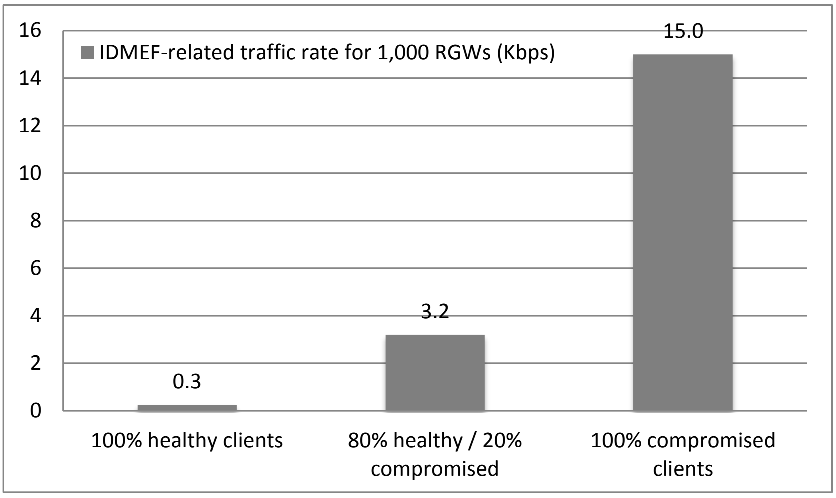

Figure 7.13: IDMEF network traffic debit rate (1,000 RGWs)

Event traffic rate per RGW (with event filtering)

Estimates presented in Figure 7.12 and Figure 7.13 assume that all events are reported back to the ISP level. This worst case scenario is very unlikely, since several events will be discarded locally by the RGWs (false alarms, harmless warnings, etc.), others will be locally processed and other ones will be correlated, resulting in the transmission of a single event resulting from a correlated sequence of local events. Thus, only a fraction of locally generated events is sent back to the ISP, positively impacting the performance of the proposed architecture. Empirically, we believe that the proportion between ISP-reported events and local events is less than 10%, even in scenarios without intensive local processing usage.

Figure 7.14 presents the estimates of capacity, considering a ratio of events reported to the ISP of 50%, 40%, 30%, 20% and 10% of the total number of locally generated events. Even assuming a conservative estimation (80% healthy clients / 20% compromised clients; reporting of 10% of local events to the ISP level), the capacity rises to 30,000 clients using modest hardware for the ISP server. This value could be extensively improved by upgrading the server hardware, by adopting proper distribution and load-balancing mechanisms on the ISP-side, and by giving more autonomy to each RGW (in order to report less events to the ISP).

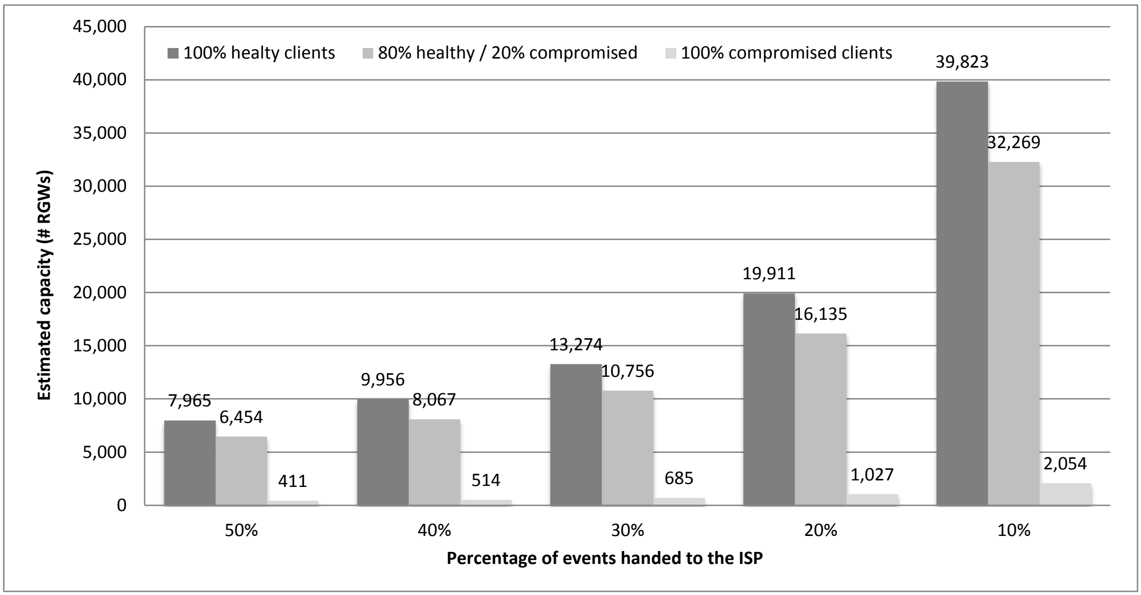

Figure 7.14: Capacity estimations for different scenarios

Synthesis

Complete validation of the proposed security management framework is not a trivial task and probably falls outside the scope of this thesis. The lack of proper traffic models and representative networks traces makes it difficult to perform an extensive analytical evaluation. The developed prototype was important to validate functional aspects and to provide a reference implementation able to prove the feasibility of the concept. Nevertheless, without a large-scale deployment in real-world scenarios, it is difficult to assess its scalability.

Considering this situation, a set of empirical experiences and estimates was performed, trying to evaluate the scalability of the concept. Those experiences and estimates are, by nature, based on arguable heuristics. Nevertheless, we believe they provide a reasonable idea of the system capacity, showing that it would be feasible to deploy the platform in real world scenarios.

7.4 Related work

The idea of assigning a more active role to domestic gateways is not new. Nowadays firewalls and QoS management capabilities are standard features in most gateways and can be configured by users through Web interfaces. However, this approach is limited by a set of factors:

- The user has the responsibility to configure these tools, and he often lacks the adequate technical preparation.
- The gateway works standalone. There is no correlation of attacks with other users of the same ISP or with specific provider or LAN services.
- The capabilities of these tools are relatively limited, and they may not be robust enough for dealing with attack patterns which are getting more and more sophisticated each day.

Also, the concept of distributed IDS with cooperative functions spread among distinct locations is not new: [Cuppens2001] already proposed a DIDS using IDMEF as a means to allow event reduction and interchange. [Antoniadis2002] also presented a IDS based on a distributed probe architecture, with centralized event processing and correlation capabilities. [Ioannidis2002] proposed a router-based DDoS defense system built with cooperating routers which communicate using a special *pushback* protocol – attacks are traced step-by-step closer to their sources and their bandwidth allocation controlled. [Koutepas2004] described an IDS/IPS framework for DDoS attacks built around a distributed management architecture and based on communities of peers, using multicast to exchange IDMEF messages – this approach assumes that each peer (designated as a Cooperative DDoS Entity) corresponds to an ISP. Attack alerts are communicated within the Distributed IDS using a flooding mechanism – once attack detection has been established they install rate-limiting filters to fight. [Wan2002] also proposed a DDoS defense solution using strategically placed probes that communicate events through IDMEF messages and are able to activate traffic-limiting mechanisms in order to deter an incoming threat. [Alfaro2006] proposed a DIDS based on a publisher/subscriber model that uses IDMEF for event information exchange, decoupling the publisher, consumer/subscriber and broker/router functions in order to enhance its scalability.

Concerning the optimization of probe distribution and location in distributed IDS, an interesting study was presented by [Suh2005].

Even if several of the approaches that were previously listed explore the idea of spreading traffic probes along the network infrastructure, most of them do not solve the problem of processing and correlating all the collected data in real-time, which constitutes a significant bottleneck in such scenarios. This issue has been specifically addressed in the scope of DIDS distributed inference mechanism research, which was originally born from the need to both detect coordinated attacks at a global scale [Katti2005] and enhance the scalability of those systems, being classified in two different categories [Feamster2010]:

- **Data-sharing techniques** use collaborative information sharing [Allman2006, Allman2008] related to several aspects, from network-level indicators [Bailey2005, Cooke2005] to message contents (for mail systems) [Damiani2004, Kong2006, Razor, Pyzor, DCC]. In the scope of these solutions, aggregation and data reduction techniques were also researched, in order to enhance scalability.
- **Anomaly-based detection techniques** basically attempt to improve centralized traffic anomaly detection systems by using a central coordinator point that performs large-scale inference [Lakhina2004].

The use of correlation in autonomous detection and decision scenarios was also proposed by [Cuppens2002], who also incorporates historic mechanisms and the capability of analyzing incomplete event chains, through virtual alerts.

The provided list of methods and techniques is by no means exhaustive. Nonetheless, apart from some similarities, none of the reviewed previous work proposed a hybrid DIDS using a hierarchic architecture with two-level correlation in the same terms as the proposal hereby described does.

Already afterwards our own proposal, [Feamster2010] introduced the conceptually similar idea of outsourced home network security, using programmable hardware in a scheme coordinated by a central controller (more specifically switches with OpenFlow API capabilities [OpenFlow] that are able to collect statistical data and enforce actions). However, this work does not specifically address the problem of distributed inference/correlation, only hinting at existing approaches. Besides, it requires programmable networking hardware with specific capabilities, instead of using the commodity, already-installed RGW platform. Finally, while the use of aggregated data is suggested, in order to reduce overall overhead, it is not clarified how programmable network hardware might do that.

7.5 Conclusion

The security management architecture hereby presented distinguishes itself by taking advantage of the domestic gateway – as a borderline device between the access network and the home network – to create a distributed security platform, with potential benefits for the provider and for the client. Even if this approach seems to go against the traditional vision of the Internet service – with the borderline boundaries on the ISP access network – it is adequate in face of recent developments like the introduction of Triple Play networks and the evolution of residential CPE management paradigms.

When compared with more traditional approaches, this framework has the following distinguishing features:

- Widespread adoption of more sophisticated local security mechanisms, including packet filters, proxies, intrusion detection and prevention (HIDS and NIDS), port scan detection and many other mechanisms that usually are found in bigger networks.
- Support for sophisticated event management, allowing the system to correlate incidents (both at local and global level) and enable adequate countermeasures, both on the ISP core network and on the RGWs.
- Support of sophisticated security configurations, taking into account the expertise of the ISP, the profile of the customer (contracted services, installed equipment, typical usage profiles, etc.) and the customer preferences (e.g. by means of dedicated web portal).

In terms of capacity and scalability, it was empirically shown that the proposed architecture is capable of delivering adequate performance, even with modest hardware requirements on the RGW-side.

Part IV:

Conclusion and final remarks

8. Conclusion

This final chapter provides an overview of the work that was perform, the problems that were addressed and the contributed developments to the field of study.

It is organized as follows. Section 8.1 provides a synthesis of the work hereby presented, followed by an overview of the research goals that were accomplished (Section 8.2). Finally, Section 8.3 discusses ongoing and future research directions.

8.1 Synthesis

The field of device and service management in broadband access networks is surrounded by a diversity of issues, related to technologies, standards, devices and services. Ultimately, this work is a reflex of all this complexity and entropy, being deemed from the start to approach a diversity of disparate subjects that, nonetheless, are somehow congregated in a common set of goals.

For those reasons, the structure of this document (and the research work itself) was somewhat unconventional. It was necessary to maintain an adequate level of isolation between subjects, each one with its own intrinsic set of research challenges and related ecosystem of standards and technologies. As a result, the adopted structure devotes each chapter of Part III (Applications and Usage Scenarios) to a separate research topic, maintaining the evaluation of proposed solutions and concepts as close as possible to their presentation and discussion.

Following this logic, the three initial chapters (corresponding to Parts I and II of this document) were structured and organized as a contextual foundation for this dissertation, with the objective of providing the necessary background to support the concepts and ideas which are next introduced.

Chapter 1 provided an introduction explaining the context of this dissertation and establishing its research topics and main objectives. It also provided an overview of the document structure.

Chapter 2 was devoted to the *state-of-the-art* in the field of broadband access networks. It provided a discussion of several subjects, structured along the topics of access network technologies, management protocols, service integration and future trends. This chapter also intended to provide a technical introduction for some of the standards and technologies more deeply addressed in the following chapters.

Chapter 3 presented the proposed management framework, which constituted the basis for the managed device and service integration solutions that were discussed on Chapters 4-7, (application scenarios). This management framework is based on Broadband Forum's CWMP protocol, but goes one step further by introducing agent extensibility features for proxied management, in order to support novel operation models and abstract interaction with services and protocols not originally envisioned by the specifications.

Next, the work diverged into four different areas (corresponding to Part III of this document), namely:

Chapter 4 was dedicated to management of non-CWMP compliant home network devices. This chapter showed how the extensible management framework can be used to incorporate management capabilities for non-CWMP capable devices on a common management infrastructure. This was demonstrated by integrating UPnP and VoIP endpoint device management on a CWMP management solution, using RGWs as management proxies.

Chapter 5 addressed the use of the proposed framework for service management, also exploring two innovative service paradigms: one for integrating broadband media content delivery for DLNA devices and the other for RGW-assisted distributed storage. Both services constitute proposals for value-added managed services which could be provided by communications operators, third-parties, or a combination of both.

Chapter 6 dealt with desktop management, presenting two different solutions: the first one integrated together the CWMP framework and the WMI management API to enable operator-based management of Windows-based devices, while the second proposed a new stateless desktop paradigm using the concept of cloud booting over broadband networks.

Chapter 7 proposed an innovative security architecture for broadband environments. It relies on a distributed detection and defense solution involving RGWs as both data collection and security enforcement mechanisms, prefiguring a distributed IDS topology with ISP-level coordination and event correlation. This approach was designed to deal with modern, distributed threats such as DDoS or Botnets, whose effectiveness and damage potential was substantially increased by the evolution and penetration of broadband networks in common households, to the point of affecting customer-sensitive services such as voice communications or television (VoIP and IPTV, respectively) in triple-play environments.

In global terms, we consider that this dissertation fulfilled its initial goals, demonstrating the potential of a management approach that allows the operator to reach the customer premises LAN to take care of configuration, diagnostics and troubleshooting operations. Also, this work emphasized the importance of the domestic gateway as an instrumental device in a multi-service scenario, positioned as the privileged security, management and service interface between the domestic consumer LAN and the communication and service providers.

The subject of management in broadband access networks is the pervasive topic of this dissertation, influencing all chapters (and therefore keeping the research work faithful to its initial proposition). Nevertheless, it should be mentioned that the application scenarios that spawned from its development (addressed in Chapters 4-7) revealed a degree of value and consistency that made them stand out as autonomous propositions. In fact, most of these scenarios could be implemented outside the scope of the management framework hereby discussed (albeit with loss of flexibility and functionality), increasing their potential interest. Moreover, the development of the application scenarios was one of the most engaging and gratifying aspects of this work, mainly because of their potential to be transposed, with relative ease, to innovative market products.

8.2 Accomplished research objectives

This section details how the specific research subjects initially identified (see Section 1.2) were addressed, discussing obtained results and other relevant aspects.

8.2.1 New service paradigms and their impact

In line with the research plan, innovative managed service paradigms for broadband environments were explored and evaluated. The idea of this topic was to study and understand to what extent new service paradigms have an impact on broadband network environments, and how much the operation and reliability of such services could be improved, while also researching and evaluating new service concepts. In the scope of this goal, the topics of media content delivery (Chapter 5), distributed storage (Chapter 5), desktop computing (Chapter 6) and security (Chapter 7) were addressed using a managed service approach.

8.2.2 The potential of frontier mechanisms and devices

RGWs are the frontier devices in the context of broadband network architectures, fulfilling the role of mediators responsible for the exchange of information between the boundaries of the provider and customer networks. By exploring the computational capability of RGWs, which is potentially available with no additional costs, we intended to leverage their potential to exploit new applications and services.

This research topic was approached in two different ways. First, the management framework proposed on Chapter 3 was designed to take advantage of RGWs as management mediators for the customer premises LAN devices and services. Second, all services and application paradigms discussed on Chapters 4-7 use the RGW as a service hub for the residential LAN, for management, service support or security. In fact, this topic is orthogonal to all the work that was developed and published, whether related to device, service or security management.

8.2.3 Management, service integration and security

As explained in Chapters 2 and 3, management of devices and services in broadband access network environments poses significant challenges for Internet Service Providers (ISPs), as a consequence of several factors. In this dissertation, these challenges were divided into three main categories: management, security and service integration.

This work explored the potential of a management model designed for residential and SOHO broadband environments, which builds on the strength of an already existing management standard – namely Broadband Forums' CWMP – to provide a complete and wide-reach solution for operators and service providers to manage devices and services.

As part of this effort, a complete CWMP management solution encompassing a modular and scalable ACS, together with an extensible agent for client devices, was developed. Later, PT Inovação (the research arm of Portugal Telecom and our industrial partner) required further enhancements to be added to the ACS and the SIP endpoint management solution (Chapter 4) which are now being used in a production environment. This management framework, which was presented and discussed with detail on Chapter 3, addressed the limitations of CWMP by bridging the operator management infrastructure with the device and service ecosystem of the residential LAN environment, using an extensible agent architecture and corresponding data model extensions – and therefore bringing together CWMP and RGWs with the customer LAN ecosystem of devices. This approach also supports the creation of new value-added managed services that can enhance service and communication operators' portfolios (further explored in Chapters 4-7).

Security was another topic that was addressed in the scope of this proposition, through the definition of the distributed security architecture presented in Chapter 7. This architecture, which perfectly integrates within the proposed management framework, enables RGWs to become part of a distributed threat detection and reaction system with a variable degree of autonomy. Also, thanks to its modular CIDF-inspired architecture, this solution is able to accommodate in a cooperative fashion a wide variety of mechanisms and security paradigms.

8.3 Future work

While this thesis documents the research effort that was done so far, within the subject of service and device management over broadband access networks, it does not constitute an ending point. To the author this continues to be an engaging and gratifying track, with prospects of future and promising developments.

With the impending release of PD-174 (Remote Management of Non TR-069 Devices) in the near future (probably within TR-069 Amendment 4), proxied management of non-CWMP devices will finally be addressed by the Broadband Forum specifications. While the management framework hereby described will remain compatible with Amendment 3 (and probably Amendment 4), it is necessary to wait for the specification to become available to fully assess which modifications might be required in each of the proposed applications, in order to keep protocol compliance.

M2M (Machine-to-Machine) interaction is another line of development that will be followed in the near future, with the aim of providing seamless integration with smart metering, intelligent home applications and sensors. Similarly, integration with mobile devices (smartphones, tablets, mobile internet devices) will be pursued for applications such as IP telephony or media content delivery, with a view towards integration in a multi-service environment.

In the service domain, two lines of development are already being followed: the development of RGW-assisted online (i.e. *cacheless*) storage services, and an evaluation of remote desktop protocols for Desktop as a Service (DaaS) services over broadband access networks. The latter effort is complementary to the thin-client proposal presented in Chapter 6 and intends to infer to which extent those protocols are adequate to provide a complete desktop experience using broadband networks for connectivity.

Finally, accompanying the industry trend towards virtualization, we are currently studying possible architectures for virtualized RGWs, encompassing service and functional aggregation at the operator-level. The objective is to explore the potential of such development, in terms of services, flexibility and usage models.

A. *Dummynet* Configuration

Dummynet was extensively used in several of the experimental evaluation studies presented in this thesis, in order to emulate the typical conditions of traditional access network environments.

A reference a set of configurations was therefore defined, corresponding to typical commercial offers based on ADSL and GPON. In order to adequately emulate the access network conditions of those offers, technology and protocol overhead were taken into account.

In the case of ADSL it was assumed ATM-based encapsulation. Considering the use of PPP over Ethernet (PPPoE) [Mamakos1999] this results in Ethernet frames on the ATM over ADSL data link layer. After PPPoE, next encapsulation stages add overhead as follows:

- First, there is RFC 2684 encapsulation (PPP over ATM AAL5 [Grossman1999]), with up to 10 octets overhead.
- Next, the packet is placed on an AAL5 Common Part Convergence Sub-layer (CPCS) Protocol Data Unit (PDU), which accounts to an extra 8 octets header. A CPCS PDU can have a payload of up to 65,535 bytes, padded to fit on 48-byte ATM cells.
- When converted to an ATM fixed-size cell (using AAL5), the stream is broken into multiple cells of 53 bytes (48 bytes of data and 5 bytes of header information). This accounts for the "ATM cell tax" around 10.4%.

Since its value is not fixed and varies with packet size, average overhead must be estimated with base on a packet size distribution. For a uniform packet size distribution on PPPoE/LLC without FCS MAC, Aken [Aken2003] estimates an efficiency value of 83.5%. This value was adopted as the reference for our ADSL test cases.

In the case of GPON it was assumed the use of GPON Encapsulation Method (GEM) [ITU2009]. GEM supports a native transport without an added encapsulation layer, with an estimated efficiency of 93% (calculated assuming collected data with a distribution of 53.3% for 64-byte, 28.1% for 512-byte and 15.6% for 1518-byte packets). For testing purposes, native IP traffic was assumed (straight IP over Ethernet, without PPPoE, as used by some GPON providers).

Table A.1 lists the effective bandwidth, round-trip latency and packet loss parameters defined for each test scenario. For ADSL emulation the *Dummynet* uplink and downlink queues were set with a depth of 10 and 30 packets, respectively. Also, an MTU of 1492 was configured to reproduce PPPoE-induced fragmentation. For GPON emulation the uplink and downlink queues were configured with a depth of 40 packets.

Table A.1: Broadband test reference scenarios

	Nominal bandwidth (b/s) (Down/Up)		**Effective bandwidth (b/s) (Down/Up)**		**RTT Latency**	**Pkt. Loss**
ADSL	4M	512K	3.34M	427.5K	20ms	0.1%
	8M	512K	6.68M	427.5K	20ms	0.1%
	16M	1M	13.36M	835K	20ms	0.1%
	24M	1M	20.04M	835K	20ms	0.1%
GPON	20M	2M	18.6M	1.86M	5ms	0%
	100M	10M	93M	9.3M	5ms	0%

References

[ABNT2011] Brazilian Association for Technical Norms (ABNT), “15606:2 - Digital terrestrial television - Data coding and transmission specification for digital broadcasting, Part 2: Ginga-NCL for fixed and mobile receivers - XML application language for application coding”, ABNT std. 15606-2, July 2011

[Aken2003] D. Aken, S. Peckelbeen, “Encapsulation Overhead(s) in ADSL Access Networks”, Thomson SA, June 2003.

[Alfaro2006] J. Alfaro, I. Barrera-Caparròs, “Intercambio distribuido de alertas para la gestión de ataques coordinados“, IX Reunión Española sobre Criptología y Seguridad de la Información , Barcelona, Spain, September 2006

[Allman2006] M. Allman, E. Blanton, V. Paxson, S. Shenker, “Fighting Coordinated Attackers with Cross-Organizational Information Sharing”, in Proc.of Hotnets V (5th ACM Workshop on Hot Topics in Networks), Irvine, CA, November 2006.

[Allman2008] M. Allman, C. Kreibich, , V. Paxson et al., “Principles for Developing Comprehensive Network Visibility”, in Proc. of HotSec 2008 (3rd USENIX Workshop on Hot Topics in Security), San Jose, USA, July 2008.

[Amazon2010] Amazon.com Inc., Amazon Cloud Drive, available at: https://www.amazon.com/ clouddrive/learnmore (last retrieved on 10/08/2011)

[Anderson2002] B. Anderson et al, “Domesticating broadband - what consumers really do with flat-rate, always-on and fast Internet access”, BT Technology Journal, Vol. 20, nº1, January 2002

[Antoniadis2002] D. Antoniadis, “LOBSTER: A European Platform for Passive Network Traffic Monitoring”, in Proc. of TRIDENTCOM 2008 (4th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities), Innsbruck, Austria, March 2008

[Apache2009] Apache Hadoop Distributed File System project, available at: http://hadoop.apache.org/hdfs/

[Apache2010] Apache Foundation, “About Apache River”, available at: http://river.apache.org (last retrieved on 07/06/2011)

[Apache2010a] Apache ActiveMQ project, available at: http://activemq.apache.org/

[Apache2011] Apache Felix Project, available at: http://felix.apache.org

[Apple2011] Apple Corp., “Bonjour Overview”, Apple MAC OSX Developer Library, available at: http://developer.apple.com/library/mac/#documentation/Cocoa/Conceptual/NetServices/Introduction.html#//apple_ref/doc/uid/10000119i, March 2011 (last retrieved on 10/08/2011)

[Apple2011a] Apple Corp., “iTunes- Airplay”, available at: http://www.apple.com/itunes/airplay (last retrieved on 10/08/2011)

[Apple2011b] Apple Corp., “Apple iOS”, available at: http://www.apple.com/ios/

[Apple2011c] Apple Corp., “Apple iCloud”, available at: http://www.apple.com/icloud/ (last retrieved on 10/08/2011)

[Arora2011] K. Arora, “Impact Analysis of Recent DDoS Attacks”, Published on the International Journal on Computer Science and Engineering (IJCSE), Vol. 3, February 2011.

[AT2Team2010] Airtunes 2 Team, “Specification of the RAOP protocol”, available at: http://git.zx2c4.com/Airtunes2/about (last retrieved on 10/08/2011)

[aTFTP] aTFTP Project, available at: http://freshmeat.net/projects/atftp.

[ATT2009] AT&T, “COMMENTS OF AT&T INC. ON THE TRANSITION FROM THE LEGACY CIRCUIT-SWITCHED NETWORK TO BROADBAND”, December 2009

[AXIROS] AXIROS, “TR-069 Appliance”, available at: http://www.axiros.com

[AXIROS2010] Axiros GmbH, “Axiros AXPAND”, available at: http://axiros.com/offerings/ competenceintr-069andbeyond/axpand-non-tr-069-to-tr-069-conversion-server-cwmp-proxy.html (last retrieved on 08/08/2011)

[Bailey2005] M. Bailey, E. Cooke et al., “ Data Reduction for the Scalable Automated Analysis of Distributed Darknet Traffic” in Proc. of IMC’05 (ACM SIGCOMM Internet Measurement Conference 2005), New Orleans, USA, October 2005.

[Barford2002] P. Barford, J. Kline, D. Plonka, and A. Ron, “A signal analysis of network traffic anomalies” in Proc. of IMW ’02 (2nd Internet Measurement Workshop 2002), Marseille, France 2002

[BBForum] Broadband Forum, available at: http://www.broadband-forum.org

[BBForum2011] Broadband Forum, “Technical Work in Progress”, available at: http://www.broadband-forum.org/technical/technicalwip.php (last retrieved on 08/08/2011)

[BBForum2011a] Broadband Forum, “BroadbandSuite 4.1”, available at: http://www.broadband-forum.org/technical/releaseprogram.php#Release41 . (last retrieved on 17/08/2011)

[Belimpasakis2008] P. Belimpasakis,S. Moloney, V. Stirbu, J. Costa-Requena,, “Home media atomizer: remote sharing of home content - without semi-trusted proxies.”, published in IEEE Transactions on Consumer Electronics,vol.54, no.3, pp.1114-1122, August (2008)

[Berners-Lee1990] T. Berners-Lee, R. Caillau, “WorldWideWeb: Proposal for a HyperText Project”, available at: http://www.w3.org/Proposal (last retrieved on 10/08/2011)

[Black2010] A. Black, “PCoIP Display Protocol: Information and Scenario-Based Network Sizing Guide”, 2010.

[Blender2008] Blender Peach Project, “Big Buck Bunny”, available at: http://www.bigbuckbunny.org (last retrieved on 10/08/2011)

[Braam1999] P. Braam et al, “The InterMezzo File System”, In Proc. of The Perl Conference 3, O'Reilly Open Source Convention, Monterrey, USA, August 1999.

[BriSA] BriSA UPnP Framework Project, available at: https://garage.maemo.org/projects/brisa (last retrieved on 10/08/2011)

[Brutlag2000] J. Brutlag, “Aberrant behavior detection in time series for network monitoring” in Proc. of LISA 2000 (the 14th USENIX conference on System administration), Louisiana, USA, December 2000

[Cablelabs2011] Cablelabs Inc., “OpenCable™Application Platform Specifications: OpenCable Application Platform (OCAP)”, Document OC-SP-OCAP1.2-110512, May 2011

[CACE2011] CACE Technologies, Wireshark Network Protocol Analyzer, available at: http://www.wireshark.org

[Canonical2010] Canonical Inc., “Ubuntu One”, available at: https://one.ubuntu.com/

[Carbone2009] M. Carbone, L. Rizzo, “Dummynet revisited”, SIGCOMM CCR, Vol. 40, No. 2, November 2009.

[CENELEC2010] CENELEC, An Interoperability Framework Requirements Specification for the Smart Home”, Document CWA50560:2010, May 2010.

[CERT2008] Carnegie-Mellon Computer Emergency Response Team (CERT), “Advisory CA-2002-03: Multiple Vulnerabilities in Many Implementations of the Simple Network Management Protocol (SNMP)”, available at: http://www.cert.org/advisories/CA-2002-03.html (last retrieved on 08/08/2011)

[Cheshire2005] S. Cheshire et al, “DNS-Based Service Discovery”, IETF draft-cheshire-dnsext-dns-sd-03, June 2005

[Cheshire2006] S. Cheshire et al, “Multicast DNS”, IETF draft-cheshire-dnsext-multicastdns-06, August 2006

[Chifflier2008] P. Chifflier, and S. Tricaud, “Intrusion Detection Systems Correlation: a Weapon of Mass Investigation”, in Proc. of CanSecWest 2008 (CanSecWest Applied Security Conference 2008), Vancouver, Canada, March 2008, available at: http://www.prelude-ids.com /fileadmin/templates/pdf/correlation-womi-cansec2008.pdf

[Chou2011] J. Chou, T. Simerly, “DLNA, UPnP pave way for home video nets”, EE Times Asia, available at http://www.eetasia.com/ARTICLES/ 2006MAR/PDF/EEOL_2006MAR16_DSP_NETD_OPT_TA.pdf, March 2006 (last retrieved on 02/04/2011)

[Citrix2010] Citrix Corporation, “XenDesktop Modular Reference Architecture”, January 2010.

[Citrix2010a] Citrix Corporation, “Optimizing HDX Technologies for XenDesktop 4 Whitepaper, Revision 1.0”, March 2010.

[Cling] Cling Project, available at: http://teleal.org/projects/cling/support.

[Cling2010] Cling, “Cling workbench”, available at: http://teleal.org/projects/cling/workbench (last retrieved on 10/08/2011)

[Cohen1998] J. Cohen, S. Aggarwal, “General Event Notification Architecture Base”, IETF draft, July 1998.

[Coherence] Coherence UPnP Framework Project, available at: http://coherence.beebits.net

[Coherence2011] Coherence Project,“Coherence Architectural Overview”, available at http://coherence.beebits.net/wiki/ArchitecturalOverview (last retrieved on 10/08/2011)

[Cooke2005] E. Cooke, F. Jahanian, D. McPherson, “The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets”, in Proc. of SRUTI 2005 (1st USENIX Workshop on Steps to Reducing Unwanted Traffic on the Internet), Cambridge, USA July 2005.

[Cruz2003] T. Cruz, P. Simões, "Enabling PreOS Desktop Management", Proc. of the IM’2003 (IFIP/IEEE Int. Symposium on Integrated Network Management), Colorado Springs, May 2003.

[Cruz2010] T. Cruz, P. Simões, E. Monteiro, F. Bastos, “Integration of PXE-based Desktop Solutions into Broadband Access Networks”, Proceedings of CNSM 2010 (6th IEEE/IFIP International Conference on Network and Services Management), Niagara Falls, Canada, October 2010.

[Cruz2011] T. Cruz, P. Simões, J. Almeida, J. Rodrigues, E. Monteiro, F. Bastos, A. Laranjeira, “How to Provision and Manage Off-the-Shelf SIP Phones in Domestic and SOHO Environments”, accepted for publication in the Proceedings of LCN’2011 (36th IEEE Conference on Local Computer Networks), Bonn, Germany, October 2011.

[Cruz2011a] T. Cruz, P. Simões, J. Rodrigues, E. Monteiro, F. Bastos and A. Laranjeira, “Outsourced Management of Home and SOHO Windows Desktops”, accepted for publication in the Proceedings of CNSM 2011 (7th IEEE/IFIP International Conference on Network and Services Management), Paris, France, October 2011.

[Cruz2012] T. Cruz, P. Simões, J. Rodrigues, E. Monteiro, F. Bastos, A. Laranjeira, “Using UPnP-CWMP Integration for Operator-assisted Management of Domestic LANs”, accepted for publication in the Proceedings of CCNC 2012 (IEEE Consumer Communications and Networking Conference), Las

Vegas, USA, January 2012.

[Cuppens2001] F. Cuppens, “Managing Alerts in a Multi-Intrusion Detection Environment”, in Proc. of
ACSAC 2001 (17th Annual Computer Security Applications Conference) New-Orleans, , USA, December 2001.

[Cuppens2002] F. Cuppens, “Correlation in an intrusion detection process”, in Proc. of SECI’02 (1st INRIA’s Securité des Communications sur Internet Workshop), Tunis, Tunisia, September 2002

[Curl] cURL Project, available at: http://curl.haxx.se.

[Curl-java] curl-java Package, http://curl.haxx.se/libcurl/java.

[Damiani2004] E. Damiani, S. Vimercati, P. Samarati, “P2P-Based Collaborative Spam Detection and Filtering”, in Proc. of P2P 2004 (4th IEEE Conference on Peer-to-Peer Computing), Zurich, Switzerland, August 2004.

[DARPA1998] Common Intrusion Detection Framework, available at: http://gost.isi.edu/cidf/

[Datamonitor2009] Datamonitor Inc., “Comcast Corporation: Company Profile”, September 2009

[Davis2008] E. Davis, “Green Benefits Put Thin-Client Computing Back On The Desktop Hardware Agenda”, Forrester Research, March 2008.

[DCC] P. Vixie, “Distributed Checksum Clearinghouse”, available at: http://www.rhyolite.com/anti-spam/dcc/

[Dean2004] J. Dean, S. Ghemawat, “MapReduce: Simplified Data Processing on Large Clusters”, in Proc. of OSDI'04 (Sixth Symposium on Operating System Design and Implementation), San Francisco, USA, December, 2004.

[Debian] Debian Project , “Apt Package Management”, available at: http://wiki.debian.org/Apt (last retrieved on 10/08/2011)

[Decius2010] Decius, “Exploiting Internet Surveillance Systems”, Presented at DEFCON 18, July 2010

[Delphinanto2009] A. Delphinanto et al., “Remote discovery and management of end-user devices in heterogeneous private networks”, Proc. of the 6th Annual IEEE Consumer Communications and Networking Conference (CCNC 2009), Las Vegas, USA, January 2009.

[Deschanel2009] M. Deschanel, “DVB IPTV Standartization”, Presented at DVB World 2009, Berlin, 2009, available at: www.dvb.org/documents/modules/DVB-World-2009-DVB-IPTV-r2.ppt (last retrieved on 08/08/2011)

[Desktone] Desktone Inc., available at: http://www.desktone.com/

[Dewaele2007] G. Dewaele, K. Fukuda, P. Borgnat et al., “Extracting hidden anomalies using sketch and non gaussian multiresolution statistical detection procedures,” in Proc. of LSAD ’07 (ACM SIGCOMM 2007 Workshop on Large-Scale Attack Defense), Kyoto, Japan, August 2007

[Digium] Digium Inc, “Asterisk IP PBX”, available at: http://www.asterisk.org.

[DIMARK] DIMARK, “TR-069 Embedded Client”, available at: http://www.dimark.com/tr-069_client.html

[Dimentrix2008] Dimentrix Technologies, “j-Interop: Pure Java – DCOM Bridge”, http://www.j-interop.org

[DLNA] DLNA Consortium, “DLNA Networked Device Interoperability Guidelines”, 2009

[DMTF] Distributed Management Task Force, available at: http://www.dmtf.org

[DMTF2009] DMTF, “Representation of CIM in XML version 2.3.1”, DMTF Specification DSP0201, July 2009

[DMTF2010] DMTF, “CIM Infrastructure Specification version 2.6.0”, DMTF Specification DSP0004, March 2010

[DMTF2010a] DMTF,” Web Services for Management (WS Management), version 1.1.0”, SMTF Specification DSP0226, March 2010

[DMTF2011] Distributed Management Task Force (DMTF), “Web-Based Enterprise Management Specifications”, available at: http://www.dmtf.org/standards/wbem (last retrieved on 08/08/2011)

[Douglieris2003] C. Douglieris, A. Mitrokotsa, “DDoS attacks and defense mechanisms: classification and state-of-the-art”, Published on Computer Networks, Vol. 44, Elsevier, 2004

[Dropbox] Dropbox Inc., available at: http://www.dropbox.com

[DVB2011] DVD Project Office, “Digital Video Broadcasting (DVB): Globally Executable MHP (GEM) Specification 1.3 (including OTT and hybrid broadcast/broadband)”, Document A153, March 2011

[Escoffier2005] C. Escoffier, D. Donsez, and R. S. Hall, “Developing an OSGi-like service platform for .NET”, in Proc. of CCNC 2006 (3rd IEEE In Consumer Communications and Networking Conference), Las Vegas, USA, 2006

[Etherboot2008] H. P. Anvin, M. Connor, “x86 Network Booting: Integrating gPXE and PXELINUX”, 2008 Ottawa Linux Symposium.

[EU2008] European Comission, “E-Communications Household Survey”, Special Eurobarometer 293, June 2008

[EURESCOM] Eurescom GmbH, available at: http://www.eurescom.eu/about-us.html

[EURESCOM2011] Eurescom GmbH, “P2055 – Virtual CPE”, 2001, available at: http://www.eurescom.eu /services/eurescom-study-programme/list-of-eurescom-studies/studies-launched-in-2010 /p2055-virtual-cpe.html (last retrieved on 08/08/2011)

[FCC2009] Federal Communications Commission, “COMMENT SOUGHT ON TRANSITION FROM CIRCUIT-SWITCHED NETWORK TO ALL-IP NETWORK NBP Public Notice # 25”, DA 09-2517, December 2009

[Feamster2010] N. Feamster, “Outsourcing Home Network Security”, in Proc. of HomeNets’10 (2010 ACM SIGCOMM workshop on Home networks), New Delhi, India, 2010

[FemtoForum2008] Femtocell Forum, “Femto Forum adopts field-proven management protocol to facilitate large-scale femtocell deployments”, available at: http://femtoforum.org/fem2/ pressreleases.php?id=234, July 2008 (last retrieved on 08/08/2011)

[FemtoForum2011] Femto Forum, “What is a Femtocell?”, available at: http://femtoforum.org/fem2/about-femtocells.php?id=207 (last retrieved on 11/08/2011)

[FIGARO] Figaro FP7 Project, “Future Internet Gateway-based Architecture of Residential Networks: State-of-the-art of energy management, e-Health and community-service requirements on common service delivery frameworks”, March 2011

[Fisher2008] J. Fisher, ”Desktone and Desktops as a Service (DaaS) – Transforming the Corporate PC”, Desktone Inc., April 2008.

[Flickr] Flickr, Flickr photo sharing service, available at: http://www.flickr.com

[Fontes2005] F. Fontes, “Operator’s Network Evolution and NGN”, in Proc. of AICT 2005 (Advanced Industrial Conference on Telecommunications), Lisbon, Portugal, July 2005

[Freeband2004] Freeband Communications, “B@Home: Broadband, multimedia applications in the home; business models and architectures”, 2004, available at: http://www.freeband.nl/uploadedFiles/Projectflyer%20B@Home%20EN.pdf (last retrieved on 08/08/2011)

[Freier1996] A. Freier, P. Karlton, P. Kocher, “The SSL Protocol Version 3.0”, IETF Internet Draft, November 1996

[Garcia-Teodoro2009] P. García-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernandez, E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges", published in Computers & Security, Vol. 28, No. 1-2, pp. 18-28, February-March 2009

[Ghemawat2003] S. Ghemawat, H. Gobioff, S. Leung, "The Google File System", The 19th ACM Symposium on Operating Systems Principles, Lake George, USA, USA October, 2003.

[Giobbi2008] Giobbi, R., "Vulnerability Note VU#347812: UPnP enabled by default in multiple devices", US-CERT, January 2008

[Git] Git version control system project, available at: http://git-scm.com

[Gluster2011] Gluster Inc., "An Introduction to Scale-Out NAS for Cloud and Virtual Environments", Technical Whitepaper, March 2011

[Goland1999] Y. Goland, T. Cai, P. Leach, Y. Gu, S. Albright, "Simple Service Discovery Protocol/1.0 Operating without an Arbiter", IETF draft, October 1999.

[Google2011] Google Inc., "Android Developers Guide: Android Runtime", available at: http://developer.android.com/guide/basics/what-is-android.html#runtime (last retrieved on 08/08/2011)

[Gorman2009] T. Gorman, "Windows Deployment Services Technical Reference", Micrososft Corporation, March 2009.

[Greene2002] B. Greene, "Remote Triggering Black Hole Filtering", The ISP Essentials, Cisco Press, 2002

[Grossman1999] D. Grossman, J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", IETF RFC 2684, September 1999.

[Gstreamer] Gstreamer open-source multimedia framework, available at: http://gstreamer.freedesktop.org (last retrieved on 10/08/2011)

[Gu2008] G. Gu, R. Perdisci, J. Zhang, and W. Lee, "BotMiner: Clustering analysis of network traffic for protocol-and structure-independent botnet detection", in Proc. of 17th USENIX Security Symposium, Vancouver, Canada, August 2008.

[Gudkova2011] I. Gudkova, K. Samouylov, "Approximating Performance Measures of a Triple Play Loss Network Model", Published in Proc. of NEW2AN 2011(11th Next Generation Wired /Wireless Networking Conference), St. Petersburg, Russia, August 2011

[Hamilton2003] J. Hamilton, "Language integration in the common language runtime", ACM SIGPLAN Volume 38 Issue 2, February 2003

[Harder2009] J. Harder et al., "Technical Deep Dive: ICA Protocol and Acceleration", July 2009.

[HGI] Home Gateway Initiative (HGI), available at: http:///www.homegatewayinitiative.org

[HGI2006] HGI, "Home Gateway Technical Requirements: Release 1", July 2006.

[HGI2008] HGI, "Home Gateway Technical Requirements Residential Profile Version 1.0", April 2008.

[HGI2011] HGI, "Requirements for Software Modularity on the Home Gateway, Version 1.0", Document RD008-R3, June 2011

[Higgins2011] K. Higgins, "DDoS Attacks Evolve And Spread", Security Dark Reading, May 2011, available at: http://www.darkreading.com/smb-security/167901073/security/attacks-breaches/229403058/ddos-attacks-evolve-and-spread.html (last retrieved on 10/08/2011)

[Hillen2009] B. Hillen et al, "Remote Management of non.TR-069 UPnP end-user devices in a private network", Proc. of CCNC'2009 (The 6TH IEEE Conf. on Consumer Communications and Networking Conference), Las Vegas, USA, January 2009

[Hitachi2010] Hitachi, Ltd., Intel Corp., Panasonic Corp., Sony Corp., Toshiba Corp., "DTCP Volume 1 Supplement E: DTCP Volume 1 Supplement E - Mapping DTCP to IP (Informational Version) Revision 1.31", 2010

[Howard1988] J. Howard et al, "Scale and performance in a distributed file system." ACM Transactions on Computer Systems Vol. 6, February 1988, pp. 51-81.

[Howard2000] B. Howard, "Thin is back", PC Magazine, Ziff Davis Media, April 2000.

[Hulu] Hulu Inc., Hulu video-on-demand service, available at: http://www.hulu.com

[Hwang2011] T. Hwang, H. Park, E. Paik, J. Chung, "EAFR-based DLNA proxy for high-quality video distribution in extended home space.", published in IEEE Transactions on Consumer Electronics, vol.57, no.1, pp. 120- 125, March (2011)

[IEEE2004] IEEE Computer Society, "IEEE 802.11i-2004: Amendment 6: Medium Access Control (MAC) Security Enhancements", 2004

[IEEE2005] IEEE Computer Society, "IEEE 802.11e-2005: Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements", 2005

[IEEE2005a] IEEE, "802.1AB: Station and Media Access Control Connectivity Discovery", IEEE Std 802.1AB, May 2005

[IEEEOUI] IEEE, OUI Registration Authority, available at: http://standards.ieee.org/ develop/regauth/oui/public.html (last retrieved on 08/11/2010)

[IGRS] IGRS Alliance, available at: http://www.igrs.org

[Intel1998] Intel Corporation, "Boot Integrity Services Application Programming Interface Version 1.0", December 1998.

[Intel1999] Intel Corporation, "Preboot Execution Environment (PXE) specification version 2.1", September 1999.

[Intel2010] Intel Corporation, "Intel vPro Technology Reference Guide, Rev. 2.1", February 16, 2010.

[Ioannidis2002] J. Ioannidis, S. Bellovin, "Implementing pushback: Router-based defense against DDoS attacks", in Proc. of NDSS 2002 (Network and Distributed System Security Symposium 2002), San Diego, California, February 2002.

[ISC] Internet Systems Consortium, "What is ISC DHCP?", available at: http://www.isc.org/software/ dhcp/about (last retrieved on 12/06/2010)

[ISO10918-1] ISO/IEC Joint Pictures Expert Group, "Information technology -- Digital compression and coding of continuous-tone still images: Requirements and guidelines: Recommendation T.81", ISO/IEC International Standard 10918-1, 1994

[ISO11578] ISO/IEC: Information technology --Open Systems Interconnection -- Remote Procedure Call, ISO/IEC standard 11578:1996 (1996)

[ISO13818-1] ISO/IEC, "Information technology -- Generic coding of moving pictures and associated audio information: Systems", ISO/IEC International Standard 13818-1:2007, 2007

[ISO14496-1] ISO/IEC Motion Pictures Expert Group, "Information technology -- Coding of audio-visual objects – Part 1: Systems", ISO/IEC International Standard 14496-1:2010, 2010

[ISO14543-5] ISO/IEC,"14543-5-1: Information technology – Home Electronic System (HES) architecture – Part 5-1: Intelligent grouping and resource sharing for HES Class 2 and Class 3 – Core protocol", FCD Draft, 2007

[ISO14908] ISO/IEC, "14908-1: Interconnection of information technology equipment -- Control network protocol -- Part 1: Protocol stack", FCD Draft, August 2011

[ISO9596] ISO/IEC, "ISO/IEC 9596-1: Information technology -- Open Systems Interconnection -- Common management information protocol", March 1998

[ITU1997] ITU-T,"Recommendation X.711: Information technology -- Open Systems Interconnection -- Common management information protocol:

Specification", October 1997

[ITU2002] ITU-T,"Recommendation J.190: Architecture of Media HomeNet that supports cable-based services", 2002.

[ITU2002a] ITU-T, "Recommendation X.680: Information technology – Abstract Syntax Notation One (ASN.1): Specification of basic Notation", July 2002

[ITU2004] International Telecommunications Union, "NGN Working definition", available at: http://www.itu.int/ITU-T/studygroups/com13/ngn2004/working_definition.html (last retrieved on 11/08/2011)

[ITU2009] International Telecommunication Union, "Gigabit-capable passive optical networks (GPON): General characteristics, Amendment 1", ITU-T G.984, October 2009.

[Kahn1998] C. Kahn, P. Porras, S. Staniford-Chen, and B. Tung, "A Common Intrusion Detection Framework", July 1998.

[Kamil2009] A. Kamil, A. Hamid et al, " Network Management Architecture Toward Universal Communication", In Proceedings of the IUCS'2009 (3rd International Universal Communication Symposium), Tokyo, Japan, December, 2009

[Kang2005] D. Kang, K. Kang,S. Choi , J. Lee, " UPnP AV architectural multimedia system with a home gateway powered by the OSGi platform.", published in IEEE Transactions on Consumer Electronics, vol.51, no.1, pp. 87- 93, February (2005)

[Karypidis2006] A. Karypidis, S. Lalis, "Omnistore: A system for ubiquitous personal storage management", In Proceedings of the 4th IEEE International Conference on Pervasive Computing And Communications, Pisa, Italy, March 2006, pp. 136–147.

[Katti2005] S. Katti, B. Krishnamurthy, D. Katabi, "Collaborating Against Common Enemies", in Proc. of IMC'05 (ACM SIGCOMM Internet Measurement Conference), New Orleans, USA, October 2005.

[Kim1994] G. Kim, E. Spafford, "The Design and Implementation of Tripwire: A File System Integrity Checker", in Proc. of SIGCOMM '94 (2nd ACM Conference on Computer and Communications Security), London, UK, November 1994

[Kimberlin1986] D. Kimberlin, "Telex and TWX History - Document Notes", 1986, available at: http://www.baudot.net/docs/kimberlin--telex-twx-history.pdf (last retrieved in 10/08/2011)

[Kirsey2010] H. Kirsey, "Interview with Heather Kirksey, co-chair of BroadbandHome Working Group", available at: http://vimeo.com/10264548, Italy, March, 2010. (last retrieved on 07/06/2011)

[Kistler1992] J. Kistler, M. Satyanarayanan, "Disconnected operation in the Coda file system", ACM Transactions on Computer Systems Vol. 10, February 1992.

[KNX2009] KNX Association, "KNX System Specifications: Architecture, Version 3.0", June 2009

[Kong2006] J. Kong et al., "Scalable and Reliabile Collaborative Spam Filters: Harnessing the Global Social Email Networks", in Proc. of WWW 2006 (15th International World Wide Web Conference) Workshop on the Weblogging Ecosystem, Edinburgh, UK, May 2006.

[Koutepas2004] G. Koutepas, F. Stamatelopoulos, B. ,Maglaris , "Distributed Management Architecture for Cooperative Detection and Reaction to DDoS Attacks", published in JNSM (Journal of Network and Systems Management), Vol. 12, No. 1, March 2004

[Krevitt-Eres1986] B. Krevitt-Eres et al, "A decision-makers' guide to videotex and teletext.", UNESCO, 1986

[Krishnamurthy2003] B. Krishnamurthy, S. Sen, Y. Zhang, and Y. Chen, "Sketch-based change detection: methods, evaluation, and applications" in Proc. of IMC '03 (Internet Measurement Conference 2003) , Miami, USA, October 2003

[LaCie2011] LaCie S.A., “LaCie CloudBox”, available at: http://www.lacie.com/us/products/product.htm?id=10563.

[Lai2002] A. Lai, J. Nieh, “Limits of Wide-Are Thin-Client Computing” in Proc. of ACM SIGMETRICS 2002, Marina del Rey, USA, June 2002.

[Lai2006] A. Lai, “On the Performance of Wide-Area Thin-Client Computing”, ACM Transactions on Computer Systems, Vol. 24, No. 2, May 2006.

[Lakhina2004] A. Lakhina, M. Crovella, and C. Diot, “Diagnosing network-wide traffic anomalies” in Proc. of ACM SIGCOMM ’04, Portland, USA, August 2004

[Lakhina2004a] A. Lakhina, K. Papagiannaki, M. Crovella, et al., “Structural analysis of network traffic flows”, in Proc. of SIGMETRICS’2004 (ACM Conference on Measurement and Modeling of Computer Systems 2004), New York, USA, June 2004.

[LDP] Linux Kernel Documentation, “Why Not dnotify and Why inotify" http://www.kernel.org/pub/linux/kernel/people/rml/inotify/README.

[LeFevre2011] W. LeFevre, Unix Top, available at: http://www.unixtop.org

[Li2010] J. Li et al, "Router-supported Data Regeneration in Distributed Storage Systems", in Proc. of FAST’10 (8th USENIX Conference on File and Storage Technologies), San Jose, CA, February 2010.

[Libes1991] D. Libes, "Expect: Scripts for Controlling Interactive Processes," Computing Systems Vol. 4, No. 2, University of California Press, 1991.

[LinuxIGD] Linux-IGD, available at: http://linux-igd.sourceforge.net.

[Loremipsum] Lorem ipsum generator, available at: http://code.google.com/p/lorem-ipsum-generator

[LUA] The LUA Programming Language, available at: http://www.lua.org.

[Mahdavi1997] J. Mahdavi, “Enabling High Performance Data Transfers on Hosts: (Notes for Users and System Administrators)”, Pittsburgh Supercomputing Center Technical note, December 1997.

[Mamakos1999] L. Mamakos, “A Method for Transmitting PPP Over Ethernet (PPPoE)“, IETF RFC 1516, February 1999.

[Marinova-Boncheva2007] V. Marinova-Boncheva, “A Short Survey of Intrusion Detection Systems”, Published in Problems of Engineering, Cybernetics and Robotics, Vol. 58, Institute of Information Technologies - Bulgarian Academy of Sciences, 2007

[Martinez2009] J. Martínez, N. Madrid, R. Seepold, "End to End UPnP AudioVisual Service Provisioning and Management ", Intelligent Technical Systems - Springer LNEE, Vol. 38, Part I, pp 45-58, 2009

[Mazel2011] J. Mazel, P. Casas, and P. Owezarski, “Sub-space clustering & evidence accumulation for unsupervised network anomaly detection” in Proc. of TMA ’11 (3rd COST TMA International Workshop on Traffic Monitoring and Analysis 2011), Vienna, Austria, April 2011

[Mazel2011a] J. Mazel, P. Casas et al., “Sub-Space Clustering, Inter-Clustering Results Association & Anomaly Correlation for Unsupervised Network Anomaly Detection”, in Proceedings of CNSM 2011 (7th International Conference on Network and Service Management), Paris, October 2011

[Microsoft2006] Microsoft Corp., “Windows Rally Technologies: an Overview”, March 2006

[Microsoft2008] Microsoft Corporation, "Querying with WQL”, MSDN Library, 2008.

[Microsoft2008a] Microsoft Corporation, “ [MS-DCOM]: Distributed Component Object Model (DCOM) Remote Protocol Specification”, MSDN Library, 2008

[Microsoft2008b] Microsoft Corporation, “Technical Overview of Windows Server 2008 Terminal Services”, January 2008.

[Microsoft2008c] Microsoft Corporation, “How MSIT Uses Terminal Services as a Scalable Remote Access Solution“, Technical White Paper, Feb. 2008.

[Microsoft2009] Microsoft Corporation, “Windows Remote Management”, 2009.

[Microsoft2010] Microsoft Corp., “Link Layer Topology Discovery Protocol Specification”, September 2010

[Microsoft2010a] Microsoft Corp., “PNP-X: Plug and Play Extensions for Windows Specification”, August 2010

[Microsoft2010b] Microsoft Corporation, “Windows Management Instrumentation Remote Protocol Specification v10.1”, March 2010.

[Microsoft2010c] Microsoft Corp., “Microsoft .NET Framework Managed Extensibility Framework”, available at: http://mef.codeplex.com (last retrieved on 08/08/2011)

[Microsoft2010d] Micrososft Corporation, “Remote Desktop Protocol: Basic Connectivity and Graphics Remoting Specification”, Revision 19, June 2010.

[Microsoft2010e] Microsoft Corporation, “Office Web Apps”, available at: http://www.microsoft.com /office/2010/en/office-web-apps/default.aspx (last retrieved on 08/11/2010)

[Microsoft2010f] Microsoft Corporation, “Remote Desktop Connection Protocol Performance and Improvements in Windows Server 2008 R2 and Windows 7“, January 2010.

[Microsoft2011] Microsoft Corp., “.NET Compact Framework”, Microsoft Developers Network (MSDN), available at: http://msdn.microsoft.com/en-us/library/f44bbwa1.aspx (last retrieved on 08/08/2011)

[Microsoft2011a] Microsoft Corp., “.NET Micro Framework”, available at: http://www.microsoft.com/en-us/netmf/about/default.aspx (last retrieved on 08/08/2011)

[Microsoft2011b] Microsoft Corp., “Microsoft SMB Protocol and CIFS Protocol Overview”, March 2011

[MicrosoftHS] Microsoft Corp., Windows Home Server, available at: http://www.microsoft.com/ windows/products/winfamily/windowshomeserver (last retrieved on 08/11/2010)

[Minokoshi2010] R. Minokoshi et al, “A Study on CWMP (TR-069) Proxy with Home Network Protocols”, Proc. of IEICE Technical Committee on Information and Communication Management (IM) 2010, Hokaido, Japan, July 2010.

[MIT] MIT Lincoln Laboratory traces, available at: http://www.ll.mit.edu/mission/ communications/ist/corpora/ideval/data/index.html (last retrieved on 10/08/2011)

[Moore1965] G. Moore, “Cramming more components onto integrated circuits”, Electronics Magazine, Vol. 38, nº8, April 1965

[Morgan2011] S. Morgan, “Fibre to the Cabinet: The Solution to Superfast Broadband or a Convenient Stopgap?", White Paper, InterConnect Communications/ Telcordia, 2011

[MR-239] Broadband Forum, “MR-239: Broadband Forum Value Proposition for Connected Home”, April 2011

[MUSE] Multi-Service Access Everywhere (MUSE) Project, EU FP6 Project, available at: http://www.ist-muse.org/

[Neisse2004] R. Neisse et al., “Implementation and bandwidth consumption evaluation of SNMP to web services gateways”, Proc. of the IEEE/IFIP Network Operations and Management Symposium (NOMS 2004), Seoul, Korea, April 2004.

[Netfilter] Netfilter, available at: http://www.netfilter.org/projects/iptables/index.html.

[Netflix] Netflix Inc, available at: http://www.netflix.com

[Nieh2003] J. Nieh et al., “Measuring Thin-Client Performance Using Slow-Motion Benchmarking”, ACM Transactions on Computer Systems, Vol. 21, No. 1, February 2003.

[Nikolaidis2007] A.Nikolaidis et al. “Local and Remote Management Integration for Flexible Service Provisioning to the Home”, IEEE Communications Magazine, pp. 130-138, October 2007.

[Ntuba2004] J. Ntuba. “Design and Implementation of an OSGi Service Architecture for the .NET Platform,” MsC. Thesis, Free University Berlin, July 2004.

[OASIS] Organization for the Advancement of Structrured Information Standards, available at: http://www.oasis-open.org

[OASIS2005a] OASIS,”Web Services Distributed Management: Management Using Web Services (MUWS) 1.0 Part 1”, OASIS Standard, March 2005

[OASIS2005b] OASIS,”Web Services Distributed Management: Management of Web Services (WSDM-MOWS) 1.0”, OASIS Standard, March 2005

[OASIS2006] OASIS, “An Introduction to WSDM”, Committee Draft wsdm-1.0-primer-cd-1, February 2006

[OASIS2009] OASIS, “ Devices Profile for Web Services Version 1.1”, OASIS Standard, July 2009

[OASIS2009a] OASIS, “Web Services Dynamic Discovery (WS-Discovery) Version 1.1”, OASIS Standard, July 2009

[OKeefe2005] M. O’Keefe, P. Kennedy, “Enterprise data sharing with Red Hat Global File System”, Red Hat Magazine issue 9, July 2005

[OMA2008] Open Mobile Alliance, “OMA Device Management Protocol Approved Version 1.2.1”, OMA Document OMA-TS-DM_Protocol-V1_2_1-20080617-A. June 2008

[OpenFlow] OpenFlow Specification v0.8.2., available at: http://yuba.stanford.edu/openflow/ documents/openflow-spec-v0.8.2.pdf. (last retrieved on 10/08/2011)

[OpenIPTV2008] Open IPTV Forum, ”OIPF Service and Platform Requirements, v2.0”, 2008

[OpenIPTV2009] Open IPTV Forum, ”Open IPTV Forum Whitepaper”, available at http://www.openiptvforum.org, 2009 (last retrieved on 10/08/2011)

[OpenIPTV2011] Open IPTV Forum, ”Questions & Answers”, available at: http://www. openiptvforum.org/questionsanswers.html (Deliverables section - last retrieved on 10/08/2011)

[OpenWRT] OpenWRT Project, available at: http://www.openwrt.org

[Oracle2010] Oracle Corp., “Java ME Technology Overview”, available at: http://www.oracle.com/ technetwork/java/javame/java-me-overview-402920.html (last retrieved on 08/08/2011)

[OSGi] OSGi Alliance, available at: www.osgi.org

[OSGi2011] OSGi, “OSGI Service Compendium, Release 4, version 4.3”, available at: http://www.osgi.org /Specifications/HomePage

[OSSEC] OSSEC HIDS, available at: http://www.ossec.net

[Park2008] H. Park, I. Lee, T. Hwang, N. Kim, “Architecture of home gateway for device collaboration in extended home space”, published in IEEE Transactions on Consumer Electronics, vol.54, no.4, pp.1692-1697, November (2008)

[Peek2006] D. Peek, J. Flinn,”EnsemBlue: Integrating Distributed Storage and Consumer Electronics”, in Proc of OSDI’06 (7th USENIX Symposium on Operating Systems Design and Implementation) , Seattle, USA, November 2006.

[Perdisci2010] R. Perdisci, W. Lee, and N. Feamster, “Behavioral Clustering of HTTP-Based Malware”, in Proc. of NSDI’10 (7th USENIX Symposium on Networked Systems Design and Implementation), San Jose, USA, April 2010.

[Physalis] Physalis Project, available at: https://developer.berlios.de/projects/physalis/

[Pichai2009] S. Pichai, “Introducing the Google Chrome OS”, http://googleblog. blogspot.com/ 2009/07/introducing-google-chrome-os.html, July 2009.

[Pickett2003] M. Pickett, “A guide to the honeypot concept”, SANS Institute, 2003.

[Pierce2004] B. Pierce, J. Vouillon,”What’s in Unison? A formal specification and reference implementation of a file synchronizer”, Tech. Rep. Technical Report MS-CIS-03-36, Dept. of Computer and Information Science, University of Pennsylvania, 2004.

[Playon] Playon media server, available at: http://www.playon.tv

[Porras2009] Porras, P. et al., “An Analysis of Conficker’s Logic and Rendezvous Points”, Technical Report, SRI International, February 2009

[Python] The Python Programming Language, available at: http://www.python.org

[Pyzor] Pyzor. http://pyzor.sourceforge.net

[Qualcomm2011] Qualcomm Inc, “Skifta media-sharing service: What is media shifting?”, available at http://www.skifta.com/ support/media-shifting (last retrieved on 10/08/2011)

[Razor] V. Prakash ,”Vipul’s Razor”, available at: http://razor.sourceforge.net

[Regis2009] S. Regis, “Introduction to NX technology”, July 2009.

[RFC1189] U. Warrier et al., “The Common Management Information Services and Protocols for the Internet (CMOT and CMIP)”, IETF RFC 1189, October 1990

[RFC1198] B. Schleifer, “FYI on the X Window System”, IETF RFC 1198, 1991.

[RFC1227] M. Rose, “SNMP MUX Protocol and MIB”, IETF RFC 1227, 1991.

[RFC1350] K. Sollins, “The TFTP Protocol (Revision 2)”, IETF RFC 1350, July 1992

[RFC1592] B. Wijnen et al., “Simple Network Management Protocol Distributed Protocol Interface Version 2.0”, IETF RFC 1592, 1994.

[RFC2131] R. Droms, “Dynamic Host Configuration Protocol”, IETF RFC 2131, March 1997.

[RFC2132] S. Alexander, R. Droms, “DHCP Options and BOOTP Vendor Extensions”, IETF RFC 2132, March 1997.

[RFC2246] T. Dierks, C. Allen, “The TLS Protocol, Version 1.0”, IETF RFC 2246, January 1999

[RFC2608] E. Guttman, C. Perkins et al., “Service Location Protocol, Version 2”, IETF RFC 2608, June 1999

[RFC2616] R. Fielding et al., “Hypertext Transfer Protocol – HTTP/1.1”, IETF RFC 2616, June 1999.

[RFC2742] L. Heintz, S. Gudur, M. Ellison, “Definitions of Managed Objects for Extensible SNMP Agents”, IETF RFC 2742, January, 2000.

[RFC2872] Y. Bernet, R. Pabbati, “Application and Sub Application Identity Policy Element for Use with RSVP”, IETF RFC 2872, June 2000

[RFC3261] J. Rosemberg et al., “SIP: Session Initiation Protocol“, IETF RFC 3261, June 2002.

[RFC3330] IANA, “Special-user IPv4 Addresses”, IETF RFC 3330, September 2002

[RFC3410] J. Case et al., “Introduction and Applicability Statements for Internet Standard Management Framework”, IETF RFC 3410, December 2002

[RFC3489] J. Rosenberg et al., “Simple Traversal of User Datagram Protocol Through Network Address Translators (STUN)”, IETF RFC 3489, March 2003

[RFC3550] H. Schulzrinne et al,”RTP: A Transport Protocol for Real-Time Applications”, IETF RFC3550, 2003

[RFC3927] S. Cheshire et al, “Dynamic Configuration of IPv4 Link-Local Addresses”, IETF RFC 3927, May 2005

[RFC4251] T. Ylonen, C. Lonvick, “The Secure Shell (SSH) Protocol Architecture”, IETF RFC 4251, January 2006

[RFC4765] Debar H. et al. “The Intrusion Detection Message Exchange Format (IDMEF)”, IETF RFC 4765, March 2007.

[RFC4779] S. Asadullah et al, “ISP IPv6 Deployment Scenarios in Broadband Access Networks”, IETF RFC 4779, January 2007

[RFC4795] B. Aboba et al, “Link-Local Multicast Name Resolution (LLMNR)”, IETF RFC 4795, January 2007

[RFC4862] S. Thomson et al,”IPv6 Stateless Address Autoconfiguration”, IETF RFC 4862, September 2007

[RFC5389] J. Rosenberg et al., “Session Traversal Utilities for NAT (STUN)”, IETF RFC 5389, October 2008

[RFC6020] M. Bjorklund, Ed., “YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)”, IETF RFC 6020, October 2010

[RFC6241] R. Enns, M. Bjorklund et al., “Network Configuration Protocol (NETCONF)”, IETF RFC 6241, June 2011

[RFC959] Postel, J. et al, File Transfer Protocol, IETF RFC 959, October 1985

[Richardson2009] T. Richardson, “The RFB Protocol Version 3.8”, RealVNC Limited, November 2009.

[Rickli2010] T. Rickli, “A component model for the .NET CLR”, Masters Thesis, ETH Zurich, September 2010

[Royon2006] Yvan Royon et al.“Virtualization of Service Gateways in Multi-provider Environments”, in Proc. of CBSE 2006 (The 9th International SIGSOFT Symposium on Component-Based Software Engineering), Västerås, Sweden, 2006

[Royon2007] Y. Royon, “Environments d’exécution pour paserelles domestiques”, Institut National des Sciences Apliquées (INSA/INRIA), PhD. Thesis, December 2007

[Royon2007a] Y. Royon et al., “Multi-service, Multi-protocol Management for Residential Gateways”, in Proc. of Broadband Europe 2007, Antwerp, Belgium 2007

[Rye2006] D. Rye, “Standard and Extended X10 Code Protocol”, Technical Paper, available at: http://software.x10.com/pub/manuals/xtdcode.pdf (last retrieved on 08/08/2011)

[Scanlogd] Scanlogd, available at: http://www.openwall.com/scanlogd/

[SCC] SCC Inc., available at: http://www.scc.com

[Schneiderman1992] B. Shneiderman, “Designing the User Interface: Strategies for Effective Human-Computer Interaction, 2nd Ed.” Addison-Wesley, 1992.

[Silveira2010] F. Silveira and C. Diot, “Urca: Pulling out anomalies by their root causes,” in Proc. of INFOCOM ’10 (the 29th IEEE Conference on Computer Communications), San Diego, USA, March 2010

[Slitaz] Slitaz GNU/Linux, available at: http://www.slitaz.org.

[Smedt2006] Smedt et al., “The multi-play service enabled Residential Gateway”, in Proc. of Broadband Europe 2006, Geneva, 2006

[Snort] Snort IDS, available at: http://www.snort.org.

[Song2009] T. Song, Y. Kawahara, T. Asami, “Using SNS as Access Control Mechanism for DLNA Content Sharing System”, In Proceedings of CCNC’2009 (The 6th IEEE Consumer Communications and Networking Conference), Las Vegas, USA, January, 2009

[Soule2005] A. Soule, K. Salamatian, and N. Taft, “Combining filtering and statistical methods for anomaly detection”, in Proc. of IMC ’05 (Internet Measurement Conference 2005), Berkeley, USA, October 2005

[Sparkleshare] Sparkleshare project, available at: http://www.sparkleshare.org

[Squid] Squid web proxy, available at: http://www.squid.org

[Squidguard] Squidguard, available at: http://www.squidguard.org.

[Staniford2002] S. Staniford, J. Hoagland, J. McAlerney, “Practical automated detection of stealthy portscans”, Published on the Journal of Computer Security (JCS), IOS Press, Vol. 10, 2002

[Staniford-chen1998] S. Staniford-Chen, B. Tung, and D. Schnackenberg, “The Common Intrusion Detection Framework (CIDF)”, in Proc. of ISW’98 (1998 Information Survivability Workshop), Orland, Florida, October, 1998.

[Suh2005] K. Suh, Y. Guo, et al. “Locating network monitors: complexity, heuristics, and coverage”, in Proc. of INFOCOM’2005 (the 24th IEEE Conference on Computer Communications),Miami, USA,March 2005

[Sun2001] Sun Microsystems, “Jini Network Technology”, Datasheet, March 2001

[Sun2006] Sun Microsystems, “Java Management Extensions (JMX) Specification, version 1.4”, November 2006

[Sun2006a] Sun Microsystems, ”JSR-000218: Connected Device Configuration 1.1.2”, March 2006

[Sun2006b] Sun Microsystems, “JSR-000277: Java Module System”, November 2006

[Sun2006c] Sun Microsystems, “JSR-000232 Mobile Operational Management: Final Release”, October 2006

[Sun2007] Sun Microsystems, ”JSR-000139 Connected Limited Device Configuration 1.1”, August 2007

[Sun2007a] Sun Microsystems, “JSR-000291 Dynamic Component Support for Java SE: Final Release”, August 2007

[TAHI] The European Application Home Alliance (TAHI), available at: http://www.theapplicationhome.com

[TCPAPER2011] Telecom.Paper BV, “Skype grows FY revenues 20%, reaches 663 mln users”, available at: http://www.telecompaper.com/news/skype-grows-fy-revenues-20-reaches-663-mln-users (last retrieved on 09/08/2011)

[TCPDump] TCPDump/Libpcap, available at : http://www.tcpdump.org/

[TEAHA] The European Application Home Alliance (TEAHA) EU FP6 Project, available at: ftp://ftp.cordis.europa.eu/pub/ist/docs/ka4/au_fp6_teaha_en.pdf (last retrieved on 10/08/2010)

[Teirikangas2001] J. Teirikangas, “HAVi: Home Audio Video Interoperability”, Technical report, Helsinki University of Technology, 2001.

[Terry1998] D. Terry et al, “The Case for Non-transparent Replication: Examples from Bayou", IEEE Data Engineering, December 1998, pp. 12-20.

[Thomas2007] V. Thomas, N. Jyoti, “Bot countermeasures”, published on the Journal of Computer Virology, Vol. 3, 2007

[Thomson2008] Technicolor Thomson, “Thomson Gateway TR-069 Configuration Guide R7.4 and higher”, Document E-DOC-CTC-20071119-0003, May 2008

[Todd1994] C. Todd, G. Davidson, M. Davis, L. Fielder, B. Link, S. Vernon, ” AC-3: Flexible Perceptual Coding for Audio Transmission and Storage”, in Proc of the 96th Convention of the Audio Engineering Society, February-March 1994

[TR-058] Broadband Forum, “TR-058: Multi-Service Architecture & Framework Requirements”, September 2003.

[TR-064] Broadband Forum, “TR-064: LAN-Side DSL CPE Configuration”, May 2004

[TR-069] Broadband Forum, “TR-069 - CPE WAN Management Protocol specification v1.2, Amendment 3”, November 2010.

[TR-098] Broadband Forum, “TR-098: Internet Gateway Device Data Model for TR-069, Issue 1, Amendment 2”, September 2008

[TR-101] Broadband Forum, “TR-101: Migration to Ethernet-Based Broadband Aggregation, issue 2”, July 2011.

[TR-104] Broadband Forum, “TR-104: DSLHome Provisioning Parameters for VoIP CPE”, September 2005

[TR-106] Broadband Forum, “Data Model Template for TR-069 Enabled Device, TR-106 Amendment 5, November 2010)

[TR-111] Broadband Forum, “TR-111: Applying TR-069 to Remote Management of Home Networking Devices”, December 2005

[TR-122] Broadband Forum, “TR-122: Base Requirements for Consumer-Oriented Analog Terminal Adapter Functionality”, November 2006

[TR-124] Broadband Forum, “Functional Requirements for Broadband Residential Gateway Devices (TR-124) issue 1.0”, 2006

[TR-124] Broadband Forum, “TR-124: Functional Requirements for Broadband Residential
Gateway Devices, issue 2”, May 2010.

[TR-133] Broadband Forum, “TR-133 DSLHomeTR-064 Extensions for Service Differentiation”, September 2005

[TR-135] Broadband Forum, “TR-135: Data Model for a TR-069 Enabled STB, Version: Issue 1”, December 2007

[TR-140] Broadband Forum, “TR-140: TR-069 Data Model for Storage Service Enabled Devices, Issue 1.1”, December 2007.

[TR-143] Broadband Forum, “TR-143 Enabling Network Throughput Performance Tests and Statistical Monitoring, Issue: 1, Corrigendum 1”, December 2008

[TR-144] Broadband Forum, “TR-144: Broadband Multi-Service Architecture & Framework
Requirements, issue 1”, August 2007.

[TR-156] Broadband Forum, “TR-156: Using GPON Access in the context of TR-101, issue 2”, September 2010.

[TR-157] Broadband Forum,” Component Objects for CWMP, TR-157 Amendment 3”, November 2010.

[TR-167] Broadband Forum, “TR-167: GPON-fed TR-101 Ethernet Access Node, issue 2”, September 2010.

[TR-177] Broadband Forum, “TR-177: IPv6 in the context of TR-101, issue 1”, November 2010.

[TR-181] Broadband Forum,” TR-181: Device data model for TR-069, issue 2”, May 2010.

[Tridgell1996] A. Tridgell, P. Mackerras, “The rsync algorithm.”,Tech. Rep. TR-CS-96-05, Department of Computer Science, The Australian National University, Canberra, Australia, 1996.

[Tun2001] B. Tung, et al. , “The Common Intrusion Detection Framework Specification”, November 2001.

[TVersity] TVersity media server, available at: http://tversity.com

[UMass] UMass Trace Repository, available at: http://traces.cs.umass.edu

[UPnPForum] Universal Plug and Play (UPnP) Forum, available at: http://www.upnp.org

[UPnPForum2003] UPnP Forum: Device Security:1 Service Template, November 2003

[UPnPForum2008] UPnP Forum: UPnP AV Architecture:1 for UPnP Version 1.0 (2008)

[UPnPForum2008a] UPnP Forum: UPnP Device Architecture 1.1 (2008)

[UPnPForum2009] UPnP Forum, “Remote Access Architecture:1”, September 2009

[UPnPForum2010] UPnP Forum, “ManageableDevice:1 DCP Version 1.01”, July 2010.

[UPnPForum2011] UPnP Forum, “Standards: Device Control Protocols”, available at: http://upnp.org/sdcps-and-certification/standards/sdcps (last retrieved on 01/04/2011)

[UPnPForum2011a] UPnP Forum, ”DIDL-Lite schema for UPnP A/V ContentDirectory services, version 2.0.”, available at: http://www.upnp.org/schemas/av/didl-lite-v2.xsd (last retrieved on 10/08/2011)

[Vandoorselaere2008] Y. Vandoorselaere, “Prelude Universal SIM: State of the Art”, presented at the Libre Software Meeting 2008, Mont-de-Marsan, France, July 2008, available at: www.prelude-technologies.com/fileadmin/templates/pdf/RMLL_2008.pdf (last retrieved on 10/08/2011)

[Varia2010] Varia, J. “Amazon Web Services - Architecting for the cloud: best practices”, White paper, January 2010

[VMware2007] VMware Inc., “VMWare Virtual Networking Concepts Information Guide, revision 20070718”, July 2007

[VMware2009] VMware Inc., “Getting started with ESX”, 2009.

[VMware2011] VMWare Inc., “Migrating to VMWare ESXi”, White paper, 2011

[VMware2011a] VMware Inc., “VMWare VSphere 4.1 Networking Performance, Performance Study”, April 2011

[W3C2001] W3C Consortium, “Web Services Description Language (WSDL) 1.1”, W3C Note, March 2001

[W3C2004] W3C Consortium, “XML Schema Part 0: Primer Second Edition”, W3C Recommendation, October 2004

[W3C2006] W3C Consortium, “Web Services Eventing WS-Eventing”, “W3C Member Submission, March 2006

[W3C2006a] W3C Consortium, “Extensible Markup Language (XML) 1.1 (Second Edition)”, http://www.w3.org, August 2006.

[W3C2007] W3C Consortium, “SOAP Version 1.2 Part 1: Messaging Framework (Second Edition)“, www.w3.org, April 2007.

[Walko2002] W. Walko, “Spec proposed for managing CPE routers and gateways”, EETimes, October 2002, available at: http://www.eetimes.com/electronics-news/4143648/Spec-proposed-for-managing-CPE-routers-and-gateways (last retrieved on 08/08/2011)

[Wall2000] L. Wall et al, “Programming Perl, Third Edition”, O'Reilly Media, 2000

[Wan2002] K. Wan, R. Chang, “Engineering of a global defense infrastructure for DDoS attacks”, in Proc. of ICON 2002 (10th IEEE International Conference on Networks), Singapore, August 2002.

[Wegner2008] T. Wegner, “Opening OSGi to the world Simple integration of services not written in Java”, OSGi Alliance Community Event, Berlin, June 2008

[Weissman2011] S. Weissman, “Sun Microsystems Was Right: The Network IS the Computer!”, available at: http://www.aiim.org/community/blogs/community/Sun-Microsystems-Was-Right-The-Network-IS-the-Computer! (last retrieved on 07/08/2011)

[WifiAlliance2007] Wi-fi Alliance, “Wi-fi Protected Setup Specification 1.0”, January, 2007

[Williams2002] A. Williams, “Requirements for Automatic Configuration of IP Hosts”, IETF draft-ietf-zeroconf-reqts-12, September 2002

[WiMAXForum2011] WiMAX Forum, “WiMAX Forum Network Architecture: Architecture, detailed Protocols and Procedures - Over-The-Air Provisioning & Activation Protocol based on TR-069 Specification”, Document WMF-T33-105-R015v01, November 2009

[XBMC] XBMC Media Center Project website, available at: http://xbmc.org

[Yang2002] S.Yang et al., “The Performance of Remote Display Mechanisms for Thin-Client Computing”, Proceedings of USENIX 2002 Annual Technical Conference, Monterey, USA, 2002.

[Youtube] Youtube, “About Youtube”, available at: http://www.youtube.com/t/about_youtube (last retrieved on 10/08/2011)

[Yu2006] W. Yu et al, “Benefits of High Speed Interconnects to Cluster File Systems: A Case Study with Lustre”, In Proc. of IPDPS 2006 (IEEE Parallel and

	Distributed Processing Symposium), Rhodes Island, Greece, 2006.
[Zeroconf]	Internet Engineering Task Force, Zeroconf Working Group, available at: http://www.zeroconf.org (last retrieved on 07/06/2010)
[Zigbee]	Zigbee Alliance, available at: http://www.zigbee.org
[Zoho]	Zoho Corporation, available at: http://zoho.com
[Z-Wave]	Z-Wave Alliance, available at http://www.z-wavealliance.org