ETSI TR 101 569 V1.1.1 (2012-12)

Technical Report

Access, Terminals, Transmission and Multiplexing (ATTM); Integrated Broadband Cable and Television Networks; Cable Network Transition to IPv6

Reference

DTR/ATTM-003018

Keywords

cable, DOCSIS, HFC, IPv6

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Copyright Notification

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://ipr.etsi.org).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Report (TR) has been produced by ETSI Technical Committee Access, Terminals, Transmission and Multiplexing (ATTM).

Introduction

The present document accommodates an urgent need in the industry to assess the current status of Cable Equipment standards that are implemented and deployed in the components that comprise integrated broadband cable and television networks with regard to their readiness for IPv6. Considering the depletion of IPv4 addresses, transition to IPv6 is required in order to enable continued growth of the customer base connected to Cable Networks and ensure service continuity for existing and new customers. High-quality connectivity to all kinds of IP-based services and networks is essential in today's business and private life.

A plethora of transition technologies have been proposed in IETF, other standardization organizations and by manufacturers of IP technology to allow coexistence of IPv4 and IPv6 hosts, access and core networks as well as services. Each of these technology options is specified, implemented and deployed in various forms and stages. The present document analyses the transition technologies, provides technical summaries and derives recommendations for one or more technologies depending on demographic and Cable Network architecture requirements. The results of the technical analysis can be used for further standardization of Cable Network transition to IPv6.

1 Scope

The present document assesses the current status of Cable Equipment standards that are implemented and deployed in the components that comprise integrated broadband cable and television networks and the approaches for their transition to IPv6. Since the time-to-market is a factor considering the depletion of IPv4 addresses, the present document accommodates an urgent need in the industry and provides the fundamental analysis for further standardization work.

The present document assesses the IPv6 transition technologies to support basic customer services, voice and data.

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are necessary for the application of the present document.

Not applicable.

2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ETSI TR 102 881 (V1.1.1): "Access, Terminals, Transmission and Multiplexing (ATTM); Cable Network Handbook", June 2010.

[i.2] ETSI TS 103 161 (all parts) (V1.1.1): "Access, Terminals, Transmission and Multiplexing (ATTM); Integrated Broadband Cable and Television Networks; IPCablecom 1.5".

[i.3] ETSI ES 201 488 (all parts) (V1.2.2): "Access and Terminals (AT); Data Over Cable Systems", October 2003.

[i.4] ETSI ES 202 488 (all parts) (V1.1.1): "Access and Terminals (AT); Second Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems", September 2003.

[i.5] ETSI EN 302 878 (all parts) (V1.1.1): "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems", November 2011.

[i.6] IEEE 802.11a: "IEEE Standard for Telecommunications and Information Exchange between Systems - LAN/MAN Specific Requirements - Part 11: Wireless Medium Access Control (MAC) and physical layer (PHY) specifications: High Speed Physical Layer in the 5 GHz band", 1999.

[i.7] IEEE 802.11b: "IEEE Standard for Information Technology - Telecommunications and information exchange between systems - Local and Metropolitan networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Higher Speed Physical Layer (PHY) Extension in the 2.4 GHz band", 1999.

[i.8] IEEE 802.11g: "IEEE Standard for Information Technology - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Further Higher Data Rate Extension in the 2.4 GHz Band", 2003.

[i.9] IEEE 802.11n: "IEEE Standard for Information technology - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput", 2009.

[i.10] IETF I-D draft-bcx-behave-address-fmt-extension-02: "Extended IPv6 Addressing for Encoding Port Range".

NOTE: Available at http://tools.ietf.org/id/draft-bcx-behave-address-fmt-extension-02.txt.

[i.11] IETF I-D draft-despres-intarea-4rd-01: "IPv4 Residual Deployment across IPv6-Service Networks (4rd) ISP-NAT's Made Optional".

NOTE: Available at http://tools.ietf.org/id/draft-despres-intarea-4rd-01.txt.

[i.12] IETF I-D draft-ietf-behave-nat64-discovery-heuristic-11: "Discovery of IPv6 Prefix Used for IPv6 Address Synthesis".

NOTE: Available at http://tools.ietf.org/html/draft-ietf-behave-nat64-discovery-heuristic-11.

[i.13] IETF I-D draft-ietf-v6ops-6204bis-10: "Basic Requirements for IPv6 Customer Edge Routers".

NOTE: Available at http://tools.ietf.org/html/draft-ietf-v6ops-6204bis-10.

[i.14] IETF I-D draft-massar-v6ops-ayiya-02: "AYIYA - Anything In Anything".

NOTE: Available at http://tools.ietf.org/id/draft-massar-v6ops-ayiya-02.txt.

[i.15] IETF I-D draft-mdt-softwire-map-dhcp-option-02: "DHCPv6 Options for Mapping of Address and Port".

NOTE: Available at http://tools.ietf.org/id/draft-mdt-softwire-map-dhcp-option-02.txt.

[i.16] IETF I-D draft-mdt-softwire-mapping-address-and-port-03: "Mapping of Address and Port (MAP)".

NOTE: Available at http://tools.ietf.org/id/draft-mdt-softwire-mapping-address-and-port-03.txt.

[i.17] IETF I-D draft-xli-behave-divi-04: "dIVI - Dual-Stateless IPv4/IPv6 Translation".

NOTE: Available at http://tools.ietf.org/id/draft-xli-behave-divi-04.txt.

[i.18] IETF I-D draft-xli-behave-divi-pd-01: "dIVI-pd - Dual-Stateless IPv4/IPv6 Translation with Prefix Delegation".

NOTE: Available at http://tools.ietf.org/id/draft-xli-behave-divi-pd-01.txt.

[i.19] IETF RFC 868: "Time Protocol", May 1983.

NOTE: Available at http://www.ietf.org/rfc/rfc868.

[i.20] IETF RFC 1142: "IS-IS Protocol Specification", February 1990.

NOTE: Available at http://www.ietf.org/rfc/rfc1142.txt.

[i.21] IETF RFC 1350: "Trivial File Transfer Protocol", July 1992.

NOTE: Available at http://www.ietf.org/rfc/rfc1350.txt.

[i.22] IETF RFC 1912: "Common DNS Operational and Configuration Errors", February 1996.

NOTE: Available at http://www.ietf.org/rfc/rfc1912.txt.

[i.23] IETF RFC 1918: "Address Allocation for Private Internets", February 1996.

NOTE: Available at http://tools.ietf.org/html/rfc1918.txt.

[i.24] IETF RFC 2131: "Dynamic Host Configuration Protocol", March 1997.

NOTE: Available at http://www.ietf.org/rfc/rfc2131.txt.

[i.25] IETF RFC 2178: "OSPF Version 2", April 1998.

NOTE: Available at http://www.ietf.org/rfc/rfc2178.txt.

[i.26] IETF RFC 2453: "RIP Version 2", November 1998.

NOTE: Available at http://www.ietf.org/rfc/rfc2453.txt.

[i.27] IETF RFC 2473: "Generic Packet Tunneling in IPv6 Specification".

NOTE: Available at http://www.ietf.org/rfc/rfc2473.txt.

[i.28] IETF RFC 3123: "A DNS RR Type for Lists of Address Prefixes (APL RR)", June 2001.

NOTE: Available at http://tools.ietf.org/rfc/rfc3123.txt.

[i.29] IETF RFC 3142: "An IPv6-to-IPv4 Transport Relay Translator".

NOTE: Available at http://www.ietf.org/rfc/rfc3142.txt.

[i.30] IETF RFC 3633: "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) v6", December 2003.

NOTE: Available at http://www.ietf.org/rfc/rfc3633.txt.

[i.31] IETF RFC 4213: "Basic Transition Mechanisms for IPv6 Hosts and Routers".

NOTE: Available at http://tools.ietf.org/rfc/rfc4213.txt.

[i.32] IETF RFC 4361: "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)", February 2006.

NOTE: Available at http://tools.ietf.org/rfc/rfc4361.txt.

[i.33] IETF RFC 4380: "Teredo - Tunneling IPv6 over UDP through Network Address Translations (NATs)", February 2006.

NOTE: Available at http://www.ietf.org/rfc/rfc4380.txt.

[i.34] IETF RFC 4861: "Neighbor Discovery for IP version 6 (IPv6)", September 2007.

NOTE: Available at http://tools.ietf.org/rfc/rfc4861.txt.

[i.35] IETF RFC 4862: "IPv6 Stateless Address Autoconfiguration", September 2007.

NOTE: Available at http://tools.ietf.org/rfc/rfc4862.txt.

[i.36] IETF RFC 5572: "IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP)".

NOTE: Available at http://tools.ietf.org/rfc/rfc5572.txt.

[i.37] IETF RFC 5625: "DNS Proxy Implementation Guidelines", August 2009.

NOTE: Available at http://tools.ietf.org/rfc/rfc5625.txt.

[i.38] IETF RFC 5969: "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification".

NOTE: Available at http://tools.ietf.org/rfc/rfc5969.txt.

[i.39] IETF RFC 6052: "IPv6 Addressing of IPv6/IPv4 Translators".

NOTE: Available at http://tools.ietf.org/rfc/rfc6052.txt.

[i.40] IETF RFC 6106: "IPv6 Router Advertisement Options for DNS Configuration", November 2012.

NOTE: Available at http://tools.ietf.org/rfc/rfc6106.txt.

[i.41] IETF RFC 6145: "IP/ICMP Translation Algorithm", April 2011.

NOTE: Available at http://tools.ietf.org/rfc/rfc6145.txt.

[i.42] IETF RFC 6146: "Stateful NAT64 - Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", April 2011.

NOTE: Available at http://tools.ietf.org/rfc/rfc6146.txt.

[i.43] IETF RFC 6147: "DNS64 - DNS Extension for Network Address Translation from IPv6 Clients to IPv4 Servers", April 2011.

NOTE: Available at http://tools.ietf.org/rfc/rfc6147.txt.

[i.44] IETF RFC 6164: "Using 127-Bit IPv6 Prefixes on Inter-Router Links", April 2011.

NOTE: Available at http://tools.ietf.org/rfc/rfc6164.txt.

[i.45] IETF RFC 6204: "Basic Requirements for IPv6 Customer Edge Routers", April 2011.

NOTE: Available at http://tools.ietf.org/rfc/rfc6204.txt.

[i.46] IETF RFC 6333: "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", August 2011.

NOTE: Available at http://tools.ietf.org/rfc/rfc6333.txt.

[i.47] IETF RFC 6563: "Moving A6 to Historic Status", March 2012.

NOTE: Available at http://tools.ietf.org/rfc/rfc6563.txt.

[i.48] Broadband Forum TR-069 Issue 1 Amendment 4: "CPE WAN Management Protocol".

NOTE: Available at http://www.broadband-forum.org/technical/download/TR-069_Amendment-4.pdf.

[i.49] CableLabs CM-SP-DOCSIS2.0-IPv6-I06-120809: "Data-Over-Cable Service Interface Specifications; DOCSIS 2.0 + IPv6 Cable Modem Specification", August 2012.

[i.50] CableLabs CM-SP-OSSIv3.0-I19-120809: "Data-Over-Cable Service Interface Specifications; Operational Support System Interface Specification", August 2012.

[i.51] CableLabs CM-SP-eRouter-I08-120329: "Data-Over-Cable Service Interface Specifications; IPv4 and IPv6 eRouter Specification", March 2012.

[i.52] IEEE 802.1Q-2011: "IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks".

[i.53] IETF RFC 6334: "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite".

[i.54] ISO/IEC 8802-3:2000: "Information technology -- Telecommunications and information exchange between systems -- Local and metropolitan area networks -- Specific requirements -- Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications".

[i.55] IETF RFC 2185: "Routing Aspects of IPv6 Transition".

[i.56] IETF RFC 3053: "IPv6 Tunnel Broker".

[i.57] IETF RFC 2460: "Internet Protocol, Version 6 (IPv6) Specification".

[i.58] IETF RFC 4193: "Unique Local IPv6 Unicast Addresses".

[i.59] IETF RFC 5942: "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes".

[i.60] IETF RFC 3971: "SEcure Neighbor Discovery (SEND)".

[i.61] IETF RFC 4443: "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification".

[i.62] IETF RFC 4605: "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding (IGMP/MLD Proxying)".

[i.63] IETF RFC 4632: "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan".

[i.64] IETF RFC 5308: "Routing IPv6 with IS-IS".

[i.65] IETF RFC 2740: "OSPF for IPv6".

[i.66] IETF RFC 4798: "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)".

[i.67] IETF RFC 4659: "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN".

[i.68] IETF RFC 3315: "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)".

[i.69] IETF RFC 3646: "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)".

[i.70] IETF RFC 3736: "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6".

[i.71] IETF RFC 4191: "Default Router Preferences and More-Specific Routes".

[i.72] IETF RFC 4075: "Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6".

[i.73] IETF RFC 3596: "DNS Extensions to Support IP Version 6".

[i.74] IETF RFC 6434: "IPv6 Node Requirements".

[i.75] IETF RFC 2464: "Transmission of IPv6 Packets over Ethernet Networks".

[i.76] IETF RFC 2827: "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing".

3 Symbols and abbreviations

3.1 Symbols

For the purposes of the present document, the following symbols apply:

µs	Microsecond
Gbit/s	Gigabit per second
Mbit/s	Megabit per second
MHz	Megahertz
ms	Millisecond

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

3GPP	Third Generation Partnership Project
4rd	IPv4 Residual Deployment
6over4	IPv6 over IPv4
6PE	IPv6 Provider Edge
6rd	IPv6 Rapid Deployment
6VPE	IPv6 VPN Provider Edge
A	(DNS) A Resource Record
A+P	Address Plus Port
A6	(DNS) A6 Resource Record (historic)
AAA	Authorization, Authentication, Accounting
AAAA	(DNS) AAAA Resource Record
ACK	Acknowledgement
AFT	Address Family Translation
AFTR	AFT Router
AICCU	Automatic IPv6 Connectivity Client Utility
ALG	Application Layer Gateway
ALG-P	Application Layer Gateway - Proxy
ALG-TS	Application Layer Gateway - Transaction Services
ALP	Application Layer Proxy
ALTS	Application Layer Translation Service
API	Application Programming Interface
APL	Access Prefix List
ARP	Address Resolution Protocol
ARPA	Address & Routing Parameter Area
AS	Autonomous System
ASIC	Application Specific Integrated Circuit
ATM	Asynchronous Transmission Mode
AV	Audio Video
AYIYA	Anything-In-Anything
B2B	Business to Business
B4	(DS-Lite) Basic Bridging BroadBand element
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
BIS	Bump-In-the-Stack
BNG	Broadband Network Gateway
BOOTP	Protocol for assigning addresses
BRAS	Broadband Remote Access Server
BSS	Billing Support System
CAPEX	Capital Expenditure
CBSA	Customer Block Set Allocation
CDN	Content Distribution Network
CDP	Cisco Discovery Protocol
CE	Customer Edge
CEF	Cisco Express Forwarding
CGA	Cryptographically Generated Addresses
CGN	Carrier-Grade NAT
CIDR	Classless Inter-Domain Routing
CLAT	Customer-side XLAT
CLNS	ConnectionLess-mode Network Service
CM	Cable Modem
CMTS	Cable Modem Termination System
CoPP	Control Plane Policing
CPA	Customer Port Allocation
CPE	Customer Premises Equipment
DAD	Duplicate Address Detection
dCEF	Distributed CEF
DCU	Destination Class Usage (accounting method in routers)

DF bit	Don't Fragment flag (in IPv4 header)
DF	Don't Fragment
DHCP	Dynamic Host Configuration Protocol
dIVI	Dual Stateless IPv4/IPv6 Translation
DMZ	DeMilitarized Zone
DNS	Domain Name System
DNSSL	DNS Search List
DOCSIS	Data Over Cable Service Interface Specification
DPI	Deep Packet Inspection
DR	Data Retention
DR/LI	Data Retention & Lawful Intercept
DSCP	Differentiated Services Code Point
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexer
DS-Lite	Dual Stack Lite
DSTM	Dual Stack Transition Mechanism
DTV	Digital Television
DUID	DHCP Unique Identifier
eBGP	External BGP
ECMP	Equal Cost MultiPath
ECN	Explicit Congestion Notification
EGP	Exterior Gateway Protocol
E-MTA	Embedded Multimedia Terminal Adapter
EUI	Extended Unique Identifier
FIN	End of TCP state
FTP	File Transfer Protocol
GGSN	Gateway GPRS Support Node
GIADDR	Gateway IP Address
GPON	Gigabit Passive Optical Network
GRE	Generic Routing Encapsulation
GUA	Global Unicast Address
GW	GateWay
HE	HeadEnd
HFC	Hybrid Fiber-Coax
HSRP	Hot Standby Router Protocol
HTML	HyperText Markup Language
IA_NA	Identity Association for Non-Temporary Addresses (DHCPv6 option)
IANA	Internet Assigned Numbers Authority
iBGP	Internal BGP
ICMP	Internet Control Message Protocol
ICXF	(IPv6-IPv4) InterConnection Function
ID	Identity
IEEE	Institute of Electrical and Electronics Engineers
IESG	Internet Engineering Steering Group
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IKE	Internet Key Exchange
IMAP	Internet Message Access Protocol
IP	Internet Protocol
IPE	Internal Provider Edge
IPFIX	IP Flow Information Export
IPoE	IP over Ethernet
IPsec	IP Security
IPTV	Internet Protocol Television
IPv4	IP version 4
IPv6	IP version 6
IRC	Internet Relay Chat
ISA	Internet Security and Acceleration Server
ISAKMP	Internet Security Association and Key Management Protocol
ISATAP	Intra-Site Automatic Tunneling Addressing Protocol
IS-IS	Intermediate System To Intermediate System

ISP	Internet Service Provider
ISSU	In-Service Software Upgrade
IVI	Stateless IPv4/IPv6 Translation
L2	Layer 2
L2TP	Layer 2 Tunneling Protocol
L3	Layer 3
LAN	Local Area Network
LDP	(MPLS) Label Distribution Protocol
LER	(MPLS) Label Edge Router
LI	Lawful Interception
LLA	Link Local Address
LLDP	Link Layer Discovery Protocol
LNS	L2TP Network Server
LSA	Link State Advertisement
LSN	Large Scale NAT
LSP	(MPLS) Label Switched Path
M2M	Machine to Machine
MAC	Media Access Control
MAP	Mapping of Address & Port (new transition protocol for IPv6)
MAP-E	Mapping of Address and Port - Encapsulation Mode
MAP-T	Mapping of Address and Port - Translation Mode
MFIB	Multicast Forwarding Information Base
MIB	Management Information Base
MLD	Multicast Listener Discovery
MLD/L	Multicast Listener Discovery
MP-BGP	MultiProtocol BGB
MPEG	Motion Picture Experts Group
MPLS	MultiProtocol Label Switching
MSB	Most Significant Bit
MSO	Multiple Service Operator
MSS	(TCP) Maximum Segment Size
MTA	Mail Transfer Agent
MTU	Maximum Transmission Unit
NAPT	Network Address and Port Translation
NAPT-PT	NAPT Protocol Translator
NAT	Network Address Translation / Network Address Translator
NAT-PT	NAT Protocol Translator
NBMA	Non-Broadcast Multiple Access
NCC	Network Coordination Center
ND	Neighbor Discovery
NDP	Neighbor Discovery Protocol
NIC	Network Interface Card
NMS	Network Management System
NPU	Network Processing Unit
NS	Name Service
NSF/GR	Non-Stop Forwarding / Graceful Restart
NTP	Network Time Protocol
NUD	Neighbor Unreachability Detection
OAM	Operation, Administration and Maintenance
OLT	Optical Line Terminal
OPEX	Operational Expenditure
OS	Operating System
OSI	Open System Interconnection
OSPF	Open Shortest Path First
OSS	Operational Support System
OUI	Organizationally Unique Identifier
PC	Personal Computer
PCP	Port Control Protocol
PD	Prefix Delegation
PE	Provider Edge
PE-PE	Provider Edge to Provider Edge
PGW	PDN GateWay

PIM	Protocol Independent Multicasting
PIO	Prefix Information Option
PLAT	Provider-side XLAT
PMP	Port Mapping Protocol
PMTU	Path MTU
PMTUD	PMTU Discovery
POP	Point of Presence
POP3	Post Office Protocol version 3
PPP	Point-to-Point Protocol
PPPoX	PPP over Media X (e.g. PPPoE - PPP over Ethernet)
PPTP	Point-to-Point Tunneling Protocol
PRL	(ISATAP) Potential Routers List
PTR	(DNS) Pointer Record
PVST	Per-VLAN Spanning Tree
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
QPPB	QoS Policy Propagation via Border Gateway
QPSK	Quadrature Phase Shift Keying
RA	Router Advertisement
RADIUS	Remote Authentication Dial-In User Service
RARP	Reverse ARP
RD	Route Distinguisher
RDNSS	Recursive DNS Server
RDP	Remote Desktop Protocol
RF	Radio Frequency
RF/HFC	Radio Frequency/Hybrid Fibre Coax
RFC	Request For Comments
RG	Residential Gateway
RIP	Routing Information Protocol
RIPE NCC	RIPE Network Coordination Centre
RIPE	Réseaux IP Européens
ROAD	Routing and Addressing
RR	(DNS) Resource Record
RSTP	Rapid Spanning Tree Protocol
RSVP	Resource ReSerVation Protocol
RTSP	Real-Time Streaming Protocol
SASL	Simple Authentication and Security Layer
SCU	Source Class Usage (accounting method in routers)
SEND	Secure Neighbor Discovery
SHA	Secure Hash Algorithm
SI	(DS-Lite) Softwire Initiator
SI-ID	Softwire Initiated Identifier
SIIT	Simple IP/ICMP Translation
SIP	Session Initiation Protocol
SLAAC	StateLess Address Auto Configuration
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SNTP	Simple Network Time Protocol
SP	Service Provider
SPF	Shortest Path First
SRV	Service
SSH	Secure Shell
SSL	Secure Sockets Layer
SSO	Single Sign-On
STB	Set-Top Box
STDPA	Static/Dynamic Port Allocation
STUN	Session Traversal Utilities for NAT
SYSLOG	Syslog Protocol
TACACS	Terminal Access Controller Access Control System
TB	Tunnel Broker
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol

TIC	Tunnel Information and Control protocol
ToD	Time of Day protocol
TRT	Transport Relay Translation
TSP	Tunnel Setup Protocol
TV	Television
UDP	User Datagram Protocol
UI	User Interface
ULA	Unique Local Address
UPnP	Universal Plug and Play
URL	Uniform Resource Locator
VLAN	Virtual LAN
VoIP	Voice over IP
VPN	Virtual Private Network
VRRP	Virtual Router Redundancy Protocol
WAN	Wide Area Network
Web-UI	Web User Interface
WEP	Wired Equivalent Privacy
WPA	Wi-Fi Protected Access
xDSL	Digital Subscriber Line of variant x (e.g. VDSL - Very High Speed DSL)
XLAT	(Address) Translator
XML	eXtensible Markup Language

4 Transition to Next Generation IP

4.1 General Considerations

Internet Protocol version 6 (IPv6) is a technology which is beginning to be adopted by network operators and ISPs across global markets as IPv4 addresses become depleted in all registries. However, despite its maturity as a standard and the gradual momentum for its adoption, the integration of IPv6 within an end-to-end network is being found to be in its early years given the significant investments in legacy network architectures. Technologies that were developed in order to work around the limitations of the available IPv4 address pool such as NAT and CIDR provided stopgaps towards the depletion of IPv4 addresses and, thus, slowed the process of transitioning to a more advanced protocol. All organizations involved in the development and application of Internet protocols taken as a community have failed to come together in a single voice of reasoning. The potential of an IPv6 all-in day where IPv6 should have been adopted as a worldwide agreement allowing all industries to move together never materialized. Instead, market reality necessitates consideration of transitional functionality and network migration.

The present document - focusing on issues encountered by Cable Network operators when considering IPv6 - studies the potential transition technologies that may be utilized within the MSO end-to-end network to effectively and efficiently enable a transition to IPv6 from IPv4. The evolution in the development of IPv6 has failed to consider interoperability with IPv4 as a key requirement. With hindsight, had this been a consideration during the development of IPv6 then the industry would have been equipped with the means to integrate the IPv6 solution alongside the existing IPv4 deployed network architecture. For example, it may have been possible to use a /96 address range to embed the full IPv4 range in the IPv6 addressing structure. However, since no consideration had been given to interoperability it meant the industry was offered two protocols, IPv4 and IPv6, that are not interoperable with each other.

The industry today is left to manage network topologies using IPv4 and IPv6 since IPv6 had not been designed to be backward compatible with IPv4. Consequently, a network topology may be designed to have end-to-end architectures of both IPv4 and IPv6, unless the industry were to develop a suitable transition technology. The former approach of both protocol versions existing in parallel is not always economic. Therefore, network operators are assessing with some urgency appropriate transition technologies to maintain continuity of their customers' services once the IPv4 addresses are depleted. This presents the industry with challenges to find ways to support network operators with their currently deployed equipment that may not be IPv6 capable by developing suitable transition technologies to allow for a smooth entry into IPv6.

This will result in not just a major change in the basic network infrastructure of the Internet but all of its associative requirements and functions and is poised to have far-reaching effects that one cannot possibly calculate in manpower and administration and that are mainly due to the ubiquity of the Internet today. However, one can place a guess in the direction of some 16 billion US dollars of cost worldwide that would not have to be met without the need to change. In the present document, basic issues of the transition from the current IPv4 networks towards IPv6 are addressed giving a brief overview of how the transition can happen and an introduction to the relevant technical issues in this area along with appropriate solutions.

The primary problem that is being addressed by moving to IPv6 is the lack of IPv4 addresses with the current addressing scheme although this is not the only issue. IPv6 offers 128 bit addresses, which are foreseen to be large enough for future purposes providing some 5 x 10^{28} addresses. While changing the address structure where each address is represented by eight sections containing two bytes each typically separated by colons, various other features are being built into the new Internet Protocol to easily enable features like security, QoS, mobility, etc. (e.g. CoPP, Netflow, ECN, ECMP).

NOTE: These features can be used with IPv4 also, but IPv6 is optimized for their application.

The change in the IP address structure impacts the whole networking stack in transit and end nodes (Layers 2 and 3 but not Layer 1 in the OSI layer model). IPv6 does not just disturb the Layer 3 functionality, new protocols such as NDP also change Layer 2 functions to a certain degree with multicast being used instead of broadcast while link-local neighbor tables and destination cache replace the ARP/RARP function. All intelligent protocols that are located within Layers 2 and 3 as well as some protocols associated with Layer 4 should change to accept the new structure and role IPv6 plays within the stack.

Some systems and applications may not have an upgrade path to IPv6 and, thus, will remain IPv4 clients until they are taken out of service, while others may be expensive or impractical to upgrade in a timely manner. Hence, there is a requirement for co-existence of the two networking technologies and internetworking requirements have emerged.

Since parts of the existing Internet backbone may continue to remain an IPv4 network, the scenario of how IPv6 will be deployed needs to be understood.

Next generation wireless networks are a hot topic in today's world. MSOs are being required to consider integrating mobile services into their own portfolio leveraging technologies and protocols developed in the wireless world. This does actually help to mitigate the cost of potential transition solutions in CPE and other network equipment. Irrespective of whether European or international standards are followed, it is clear that the IPv6 enabled end-user terminal will have to support both voice and data features. While the characteristics of data traffic do not require the device to be a uniquely addressable entity (it is possible to have a scenario of private addresses and usage of an address translation gateway), a voice channel requires the setup of a peer-to-peer connection where the terminal needs to be directly addressable in order to be contacted by external applications. With the arrival of broadband technologies and the always-on Internet, even data applications (for example multi-party gaming) bring in requirements for unique addresses. With the current, rapidly diminishing pool of IPv4 addresses, it would be infeasible to support such requirements. Hence, wireless networks become the most important business case to serve as a driver to migrate towards IPv6.

Even within wired networks, IP has become an essential part of the network particularly for MSOs as well as increasingly for incumbent network operators even for voice services. There is a move away from circuit switched networks towards packet switched IP networks while attempting to retain some of the features of circuit switching (for example end-to-end connectivity) through additional features and applications built on top of IP.

The pressure to migrate to IPv6 is being felt in all parts of the world to various degrees. In some countries such as Japan the scarcity of IPv4 addresses is felt more urgently and, hence, it is taking the lead to move towards IPv6. Countries like the United States, which received huge blocks of IPv4 addresses during the early days of the Internet, are feeling less immediate pressure to move towards IPv6, although that might not be equally applicable to all types of ISPs. As a result, the global Internet will not migrate along a uniform transition timeline, but instead will first transition to IPv6 in regional pockets where the need is greatest.

To briefly cover the deployment status of IPv6, it is widely deployed in test networks especially in Japan and Europe which are being followed by China and Korea. Japanese ISPs have started commercial sales of IPv6 addresses. In the United States, the largest ISPs (including MSOs) have begun deployment of IPv6 to their customers. However, overall deployment is still limited. The U.S. Department of Defence had established a goal of transitioning to IPv6 by 2008, and as a first step it had announced that any new equipment purchased after September 2003 should support IPv6. This illustrates the preparations being made by some organisations that recognise the move to IPv6 as being a necessity and, therefore, taking steps to ensure investment protection in purchased equipment. The availability of early test networks such as the 6bone network (consisting of independent sets of IPv6 networks across the globe, linked together by special tunneling mechanisms over the IPv4 Internet) provided the industry not only direction to the importance of moving to IPv6 but also provided a test bed for native IPv6 implementations and transition testing.

Equipment vendors all over the world are today including IPv6 support as a checkbox item in any network equipment that is being manufactured allowing the ISP the freedom to decide whether to use it. One of the issues facing MSOs today is the unavailability of applications supporting IPv6 - the porting of existing applications and creation of new applications is yet to happen on a large scale. The unavailability of applications supporting IPv6 has been one of the key missing components from the IPv6 toolbox.

4.2 Development of Transition Technologies

Originally, eighteen different transition or co-existence technologies were considered, fully or partially developed, and with the possibilities of interactions between different networks. Implementing different mechanisms of the mixed IPv4/IPv6 Internet presented a very complex and unadvisable picture driving the industry to look for potentially simplistic solutions of a single methodology.

The most widely deployed mechanism is to make changes to the network gateway since this scheme is a tried and tested method and can be directly applied to the scenario of having a small IPv6 island network connecting to the IPv4 Internet. Some of the techniques available are:

Configured tunnels:
Manually configured tunnels are used to set up permanent pathways for IPv6 across the IPv4 Internet. Much of the 6bone operates in this fashion. Routing decisions require information on all tunnel end points (which IPv6 addresses are available behind which IPv4 address), and the tunneling software needs to be installed on the gateway.

Automatic tunnels:
Automatic tunnels are set up when required and torn down at the end of the traffic relay. This mechanism requires "IPv4 compatible" addresses that are available to each IPv6 host which are used to derive the tunnel endpoints, thereby forcing a major limitation on the usage of this technique.

6to4:
This mechanism uses unique 6to4 prefixes in the IPv6 address to determine the IPv4 address of the egress gateway. It is widely used as a tunneling mechanism.

NAT-PT:
NAT-PT is a well-established translation mechanism. It is merely an extension of the NAT mechanism, which allows a private IPv4 network to exist behind a few defined public IPv4 addresses. In the case of NAT-PT the private network can be an IPv6 network. SIIT (Simple IP/ICMP Translation) is the technique defined to translate between IPv4 and IPv6 payloads. As in the case with NAT, ALGs (Application Layer Gateways) have to be defined for each application, which requires passing through the NAT-PT gateway.

ISATAP:
The Intra-Site Automatic Tunneling Addressing Protocol connects IPv6 hosts and routers within IPv4 networks. ISATAP treats the site's IPv4 infrastructure as a non-broadcasting multiple access link layer and tunnels the IPv6 payload in an IPv4 packet.

Tunnel Brokers:
When an isolated IPv6 host within an IPv4 only network would like to contact an external IPv6 network, it uses the facilities of a tunnel broker which ensures the management of the tunnel - IPv6 address allocation to the host, DNS reachability information propagation to remote IPv6 networks, etc. Teredo is an example for a tunnel broker.

Shipworm/Teredo:
Shipworm or Teredo is a technique for the transport of UDP packets across NATs, which works well in the scenario when there is a private IPv6 network behind the NAT machine. Teredo servers and relays are designated devices, which allow the overlay of the Teredo network over the existing IPv4 infrastructure. IPv6 packets are encapsulated as UDP payload and are relayed by the Teredo relay which is available within the local network to the connected Teredo server. From there the packet is routed to the appropriate Teredo server nearest to the ultimate destination where the decapsulation of IPv6 is handled by the Teredo relay. Teredo is a less preferred mechanism to be used where 6to4 or other tunneling mechanisms are unavailable since there is an overhead due to the encapsulation into UDP.

DSTM:
The best mechanism for individual hosts to communicate in the absence of a gateway is to implement both IPv4 and IPv6 stacks on the device. All the devices in the network are configured for IPv6 but not for IPv4, a single machine within the network operates as a DSTM server allotting temporary IPv4 addresses to devices when required to enable communications over the IPv4 network.

BIS:
Bump-in the-stack is a technique for non-IPv6 compliant applications on an IPv4 only host to communicate with IPv6 hosts. The IPv4 stack on the host has a "bump" added which converts the specific packets designated to the IPv6 destination into IPv6 packets and performs the reverse mapping for the returned packets.

6over4:
This technique encapsulates IPv6 in IPv4 without explicit tunnels. Other than the requirements for dual stack support on hosts, a gateway port also requires specific configuration to work as an interface for 6over4.

The present document discusses not only the historical development of IPv6 but also assesses the transitional landscape that appears technically different. The above technologies were the original choices within the industry until end of 2008 with significant focus given on development of solutions resulting in an array of proposals considering the requirements of ISPs. Since then, the industry looked at developing practical, economic and more realistic solutions for the market deciding on three main technologies that were considered to be in the forefront for potential deployment. The development focused on DS-Lite, dIVI and NAT64. These three technologies are at present the most widely used transition technologies particularly among MSOs, with DS-Lite presenting an estimated 90 % of deployments worldwide.

Although this is the current situation, the present document provides more detailed studies, researching and technically analyzing a wide range of options for transition technologies to give a balanced technical summary for the evaluation of actual transition strategies.

4.3 End-to-end IPv6 and Network Support Systems and Applications

Support systems and transport networks are a main issue and concern for IPv6 deployment as is transitioning to IPv6 as a whole a major concern for MSOs. Most technologies require the use of IPv6 in some manner.

Routing is the most important function in the Internet which relies completely on IP addresses for the propagation of reachability information. With the transition towards IPv6 and in a mixed network scenario, existing routing protocols for both interior and exterior routing require upgrades.

The IGPs like OSPF, RIP and IS-IS have been defined in newer versions which operate on IPv6 addresses as the basic identifiers.

There is only one EGP used in the Internet - BGP version 4. BGP-4 is merely used to propagate reachability information between autonomous systems and leaves the specifics of routing within the autonomous system to IGPs. Extensions have been defined on BGP-4, which allow the exchange of information on non-IPv4 protocols including IPv6. Also, BGP-4 operates over TCP sockets and the protocol can function irrespective of whether the network layer operates IPv4 or IPv6. BGP-4 has been extensively used in the 6bone for propagating the routing information of IPv6 networks.

DNS support is an indispensable requirement when transitioning to IPv6. The existing DNS mechanism for IPv4 networks provides the name to address lookup and the reverse mapping from address to name. The DNS name servers now require to store the associated IPv6 address of a name (in addition to the IPv4 address). They should be able to perform the reverse lookup functionality as well.

The DNS records for IPv6 currently use the AAAA format (in early IPv6, the A6 format has been defined alternatively, but its status is historic [i.47] and not many servers support this format). Hence a DNS client requiring the IPv6 address of a host would request the server for the AAAA entry of that host. If the entry is found, it is used directly, otherwise the A entry corresponding to the IPv4 address is retrieved and mapped into an IPv6 address by using standard procedures.

4.4 Network Management

SNMP is the de-facto standard for a management protocol used in the current Internet. As new standards are being defined for the IPv6 protocol and all other related technologies, the corresponding SNMP MIB definitions for these are also being made. However, vendor adoption of these MIBs has been slow; only recently vendors have started implementing some of the relevant MIBs. The standard transport mechanism for SNMP is over UDP which could run either over IPv4 or IPv6 with appropriate changes to the socket layer. Standard SNMP management platforms like HP OpenView have basic support for IPv6 available and are also indicating roadmaps with full feature support for IPv6. Unless sufficient management tools are available, the commercial deployment of IPv6 would be difficult since ISPs and enterprise network managers require the tools to configure and monitor IPv6 networks. The tools become very important especially in a mixed network scenario where the network manager will have to keep track of tunnels, routing issues, DNS configurations, etc. across both IPv4 and IPv6 networks.

4.5 MSO Considerations

The industry has dealt with some of the potential pitfalls of IPv6 and its transition; but in considering MSO deployments, hundreds of potentially affected applications and services can be listed/ Portals, DHCP and lease times, OSS and BSS in general. These challenges also extend into monitoring technologies and security. Transitional compliance within a standard should consider all of these aspects with a final gap analysis on IPv6 and transitionally required technologies.

There is no straightforward answer to the question what direction MSOs should take; and difficulty compounds when the vast list of possible techniques suggested in previous clauses as a basis is considered and becomes clear. Further, since multiple transition techniques are defined, it is likely that multiple techniques can be used within a local network and, hence, the network architect has to consider issues arising out of combinations of techniques if so required.

The mechanisms and the issues involved in the transition from pure IPv4 networks into a mixed IPv4-IPv6 networks and the further progress into pure IPv6 networks and their transition need to be studied comprehensively. Since many aspects are still undergoing standardization, commercial implementation of the transition got delayed. Several experimental networks working individually and the 6bone on a global level have established the viability of creating IPv6 island networks within the IPv4 network and proven the usability of the technology. However, applications and management/network configuration tools are required to be updated to work in mixed network scenarios before the transition becomes a reality.

An additional challenge for the industry is the fact that the United States and a few other larger countries have a few IPv4 address blocks still available that will serve their needs until 2015 to 2018. This may be considered by IPv6 proponents as a significant delay or block to IPv6 deployment. However, that seems to be shortsighted as it is overlooking an important fact. Today's industry operates in a world economy where customers access network provider's services from many different geographic locations worldwide. Consequently, the fact that IPv4 address blocks are available in some countries and not available in other countries means that there is a demand to move to IPv6. Therefore, the "chicken and egg" situation that might be assumed in a situation with IPv4 addresses still being available in limited amounts is not applicable from a customer perspective. As the world moves to IPv6 with significant drive, the situation arises where an increasing number of customers will not be able to reach the network providers' services if they are not provisioned on IPv6. Investors can expect to react to market pressure as customer churn results from the lack of IPv6 connectivity, thus, driving a more rapid adoption of IPv6 in the near future.

5 Background and Concept of Transition

IPv6 transition is critical to the long-term sustainability of European and global networks in order to ensure business continuity. With more and more services and industries come to rely on the global Internet as a fundamental platform the need for ubiquitous connectivity of devices and services becomes very urgent. MSOs will not have an option not to move to IPv6. Such near-term strategic areas like Mobile Internet and Smart Grids as well as the continued growth in residential and business broadband access services are poised to introduce massive numbers of devices that require network connectivity, which may not easily be provided by the current Internet (IPv4) networks with their depleting address space. More details on RFCs that represent important steps in the history of the development of IPv6 are provided in Annex B.

Widespread adoption of IPv6 has been identified as the best way forward to address the exhaustion of the IPv4 address space. Prompt and efficient adoption offers potential for innovation and leadership in advancing the Internet, while delayed adoption of IPv6 would lead to disadvantages for all users and a weaker competitive position of the industry. In the meantime, the IANA Unallocated IPv4 Address Pool was exhausted on 3 February 2011, and the RIPE NCC IPv4 Address Pool is dwindling (http://www.ripe.net/internet-coordination/ipv4-exhaustion). The urgency to transition broadband Internet networks to IPv6 is becoming critical.

Device manufacturers, software developers and network operators are adopting IPv6. Both MSO investments and the impact from depletion of IPv4 addresses drive the need for its implementation into the customer network. However, the vision of an Internet running IPv6 only will not become a full reality any time soon. For a considerable period of time, significant numbers of devices and services will exist that customers want to use and that were designed to require IPv4 connectivity. Among the more prominent examples for such devices and services are IPTV sets with static firmware, IP-connected refrigerators and other appliances, security devices and devices that cannot be upgraded through software to enable an IPv6 stack. These devices have a longevity of up to eight or ten years and, thus, transitional requirements may be needed at least until 2020.

An immediate replacement of these IPv4 hosts and networks may not be feasible or not desirable for various technical and economic reasons. It is particularly the task of access network operators and broadband service providers to ensure customer choice in terms of IPv6 technology and services, as well as the continuation of IPv4 connectivity. Appropriate transition technologies enable and allow the coexistence of IPv6 and IPv4 in various parts of the end-to-end network. In this way, services may be consumed and customer premises equipment may be used transparently while fostering a smooth transition to the required extended address space provided by IPv6. Simultaneous connectivity among IPv4 and IPv6 hosts and services by employing appropriate transition technologies ensures e.g. the ability to offer IPv4 services even though the address pool has been depleted and the consumer is connected via IPv6.

5.1 Broadband Cable Network Providers

As of today, integrated broadband cable and television networks go into the home of more than 73 million customers in the European Union providing Digital TV, Broadband Internet and Telephony services. Broadband Internet services provided by Cable Networks utilizing DOCSIS cable modem technology enable about 23,2 million subscribers in Europe (2010) to connect to the Internet with download speeds of currently up to 200 Mbit/s. This figure has grown by at least 12 % annually and is expected to continue to grow. Multiple vendors are pushing the equipment supply chain whilst Cable Network operators provide the platform to satisfy fundamental entertainment, communication and information needs to consumers. Furthermore, the industry is anticipating a transition to delivery of digital television using broadband cable modem technology, which will dramatically increase the number of broadband connected households.

Cable Networks are recognised as a key enabler in supporting Europe's Digital Agenda. To continue to meet the demand of accelerating connectivity of digital devices, a standardised approach for the cable eco-system to rapidly transition to IPv6 is required. A failure for an effective standards driven transition would impair the ability to achieve cost effective solutions on a large scale.

5.2 IP Connectivity in Cable Networks

Integrated broadband cable and television networks are built against various international and ETSI standards. Figure 1 depicts the fundamental architecture of a Cable Network as it is currently deployed with a hybrid fibre-coax approach. [i.1] provides a complete overview on Cable Network architectures and services. The IP communication system in Cable Networks is based on the series of ETSI DOCSIS [i.3], [i.4], [i.5] and PacketCable standards [i.2]. Since 2006, the current version of the DOCSIS technology (DOCSIS 3.0) has natively supported IPv6. For various technical reasons, service providers in Europe have not implemented IPv6 support for their customers.

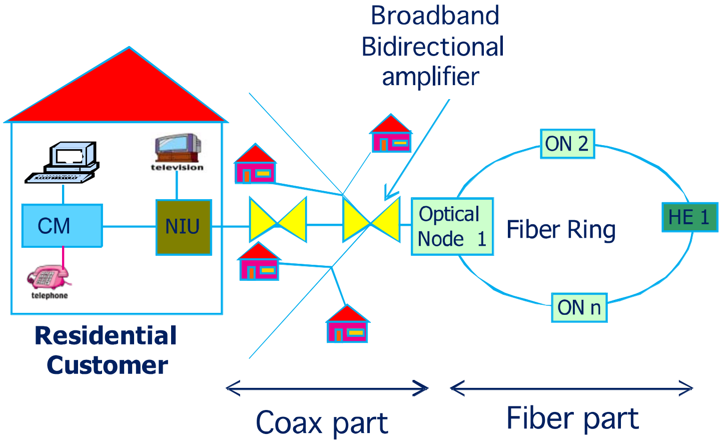

Figure 1: Principle architecture of a hybrid fibre-coax (HFC) Cable Network

In order to achieve end-to-end connectivity, the Cable Network is interconnected at the Headend (HE) with a backbone network and the Internet. The backbone network may or may not be operated and managed by the same entity as the access network. The home network within the customers' premises is typically installed, configured and operated by the customer.

For the purposes of the analysis in the present document, the end-to-end network is subdivided into various parts, each of which currently supports IPv6 and/or IPv4 with various degrees of probability. The network parts also distinguish themselves by the number of times they occur globally and by the uniformity of their operation and management. Both aspects have an impact on how difficult it is to homogeneously transition to IPv6.

5.2.1 Customer Host

This is the device that the customer uses to consume an IP-based service. It is the final destination and the originating source of IP packets. As Customer Premises Equipment (CPE) it is typically owned and operated by the customer. Cable Network operators are unlikely to be able to support and manage the wide variety of host devices that customers are deploying in their homes.

5.2.2 Home Network

The Home Network extends from the Customer Host to the Access Gateway. It may be constituted by a simple wired or wireless link between the two devices (in which case it becomes irrelevant for this analysis since it is not addressable and does not process protocol messages) or may consist of a complete infrastructure of routers, wireless access points, etc. It may connect multiple Customer Hosts to the Access Gateway and it is typically owned, installed and operated by the customer. Cable Network operators may impact the capabilities of the Home Network up to a certain extent by customer information and/or delivery of devices.

5.2.3 Access Gateway

This device is installed at the customer location and constitutes the separation between the Home Network and the Access Network. While the device may be owned by the customer or the Cable Network operator it is typically authorized for usage by the latter. It terminates the Cable Network on the customer side.

5.2.4 Access Network

The Access Network extends from the Access Gateway to the Headend. It is operated and managed by the Cable Network operator. It typically uses DOCSIS technology to establish IP connectivity and packet transport.

5.2.5 Headend

This device separates the Access Network from the Backbone Network. It terminates the Cable Network on the operator side. For the purposes of the present document it also terminates the portion of the end-to-end IP connection that traverses the operator-managed network. As such it will be the location of devices that are providing transition technology services. However, that does not mean that all Headend functionality is concentrated at a single geographical location.

5.2.6 Internet

This is the Internet cloud that is considered to be the host of IP-based services. It is the final destination and the originating source of IP packets which may be delivered as IPv4 or IPv6. The actual delivery of the service or the path the packet has taken through the Internet cloud is irrelevant for the present document, only the way of addressing is taken into account.

Due to the inherent incompatibility of IPv4 and IPv6 addressing schemes, special measures have to be taken (i.e. transition technologies have to be deployed) in order to ensure connectivity between equivalent network parts supporting different Internet protocols (e.g. IPv4 host talking to IPv6 service across IPv4 home network and IPv6 access network and backbone). The items listed below summarize the types of connectivity that are considered in the present document.

- IPv4 Customer Host via IPv4 Access Network connecting to IPv4 Internet
  This is out-of-scope of the present document as it is provided by native IPv4

- IPv6 Customer Host via IPv4 Access Network connecting to IPv4 Internet
  This is out-of-scope of the present document as it is within customer responsibility

- IPv4 Customer Host via IPv6 Access Network connecting to IPv4 Internet

- IPv6 Customer Host via IPv6 Access Network connecting to IPv4 Internet

- IPv4 Customer Host via IPv6 Access Network connecting to IPv6 Internet

- IPv6 Customer Host via IPv6 Access Network connecting to IPv6 Internet
  This is out-of-scope of the present document as it is provided by native IPv6

- IPv4 Customer Host via IPv4 Access Network connecting to IPv6 Internet

- IPv6 Customer Host via IPv4 Access Network connecting to IPv6 Internet

6 Analysis of Cable Network Deployments

This clause defines the main parts of the MSO network as they are currently deployed. Developing technologies such as CDN are not taken into account.

6.1 Overview

For the purposes of analysing currently deployed Cable Network architectures with regard to their support of IPv6 and with the goal to identify connectivity gaps that are to be addressed by transition technologies, the following network components are considered.

- End-user CPE

- Access Network

- Core Network
- Datacentre
- Support Services

Whilst this list is not exhaustive, and does not cover irregular or bespoke components used by individual Cable Network operators, it provides an overview of the main areas that can be mapped to individual operator deployments.

Figure 2 shows the main component areas of a Cable Operator network providing end-to-end Internet access. There are a number of additional services that may be offered by the Cable Operator, including business services, customer site-to-site Virtual Private Networks (VPNs), etc. The details of such configurations are beyond the scope of this analysis.

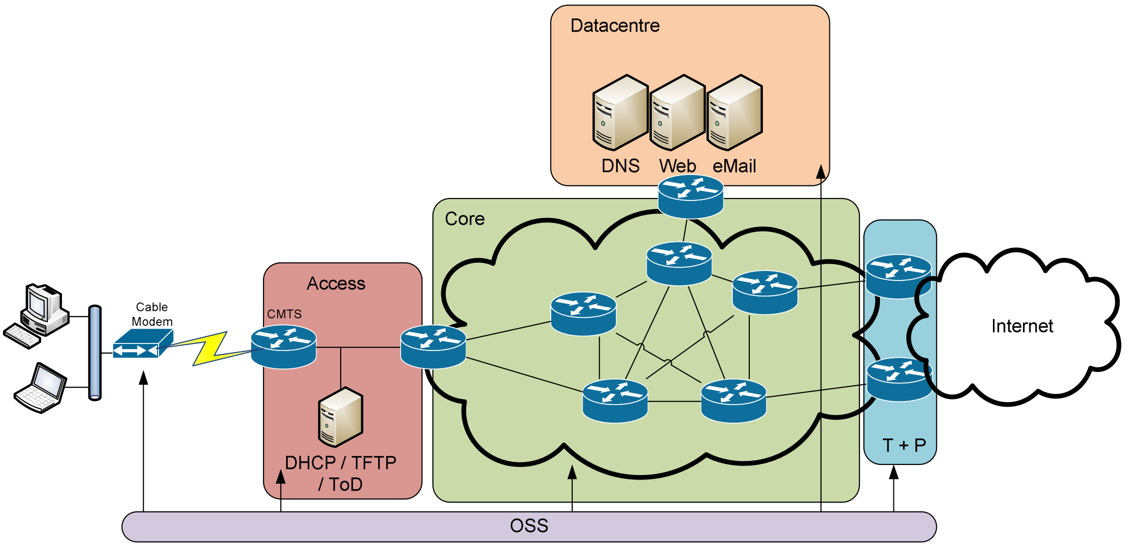

Figure 2: Main component areas of a Cable Operator network for end-to-end Internet access

6.2 End-user CPE

Figure 3 gives a view of the most common CPE installed at customer sites.

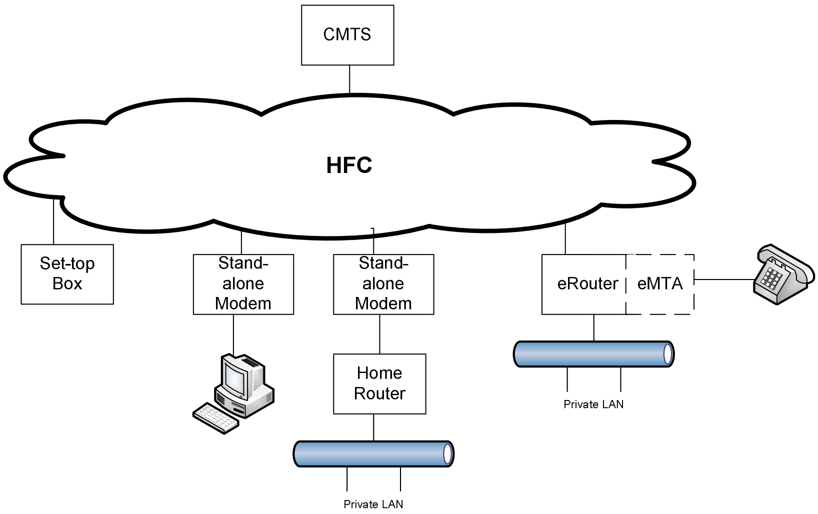

Figure 3: CPE installed at a customer site

6.2.1 Stand-alone Cable Modem

A stand-alone Cable Modem is a CPE device that provides connectivity to the Cable Operator's HFC network. It has a single Ethernet interface on the customer-facing side. Stand-alone Cable Modems do not provide any routing facility and are typically either used for connecting a single host (e.g. PC) or for connecting to the WAN interface of a stand-alone home router. The customer equipment connected to the Cable Modem is assigned a public IP address from the operator's DHCP server.

6.2.2 Home Router

A stand-alone home router provides LAN functionality within the customer home. It has one Ethernet WAN interface, which is connected to the Cable Modem, and one or more wired Ethernet LAN interfaces. It generally includes a Wireless LAN interface using IEEE 802.11a/b/g/n standards [i.6], [i.7], [i.8] and [i.9].

The IP addressing of the LAN generally uses private IP addresses [i.23], most commonly in the class C range 192.168.x.x. There is an extensive range of home routers available, all with differing functionality, however the most common features include:

- LAN DHCP server
- DNS (Domain Name System) forwarding
- Stateful Packet Inspection firewall
- Network Address Translation (NAT) with features to allow single or multiple port forwarding (manually or via UPnP)
- Wireless LAN configuration to support one or more BSSIDs (Basic Service Set Identifiers) and wireless security (WEP, WPA, WPA2)

6.2.3 eRouter

An eRouter is a combined Cable Modem and home router. These devices generally provide all of the features as detailed in clauses 6.2.1 and 6.2.2. However all functionality is enclosed in a single housing which requires a single power supply and takes up less space. Features and functionality vary widely between manufacturers. The CableLabs eRouter specification [i.51] defines the minimum set of requirements for these devices as they are deployed by Cable Operators.

6.2.4 Set-top Box

A set-top box is used to receive and decode digital television signals delivered by the Cable Operator. In addition, the set-top box may also provide interactive services and be able to retrieve content from the Internet. The broadcast signal does not use IP, however interactive functions and access to the Internet are provided using IP. Where the set-top box is provisioned with a private IP address, access to the interactive functions either remain internal to the operator network, using web proxies to connect to the Internet, or operate behind a Network Address Translator with NAT-44. A web proxy or NAT allows multiple privately addressed clients to access content through a single public IPv4 address. Interactive and Internet connectivity on set-top boxes is generally limited to HTML based content.

6.2.5 Consumer Equipment

There is a wide and ever-increasing range of equipment that is becoming "internet-ready" by having either a wired Ethernet interface or a built-in wireless IEEE 802.11a/b/g/n interface. This device category includes Personal Computers (PCs) and laptops with a wide range of operating systems and versions as well as tablet devices, eReaders, games consoles (static and handheld), network storage devices, printers, PDAs, televisions, Smartphones, etc.

6.2.6 E-MTA

An E-MTA (Embedded Multimedia Terminal Adapter) is deployed by Cable Operators to provide telephony services using VoIP across the HFC network. In the strict sense, the E-MTA is a combination of two separately provisioned devices, an embedded Cable Modem (eCM) and an embedded MTA (eMTA). The eMTA is provisioned with an IP address from the Cable Operator's DHCP server. Additional options such as a SIP (Session Initiation Protocol) gateway address are provided with the DHCP response. Some Cable Operators offer VoIP services using PacketCable [i.2] to provide dedicated bandwidth for VoIP services.

6.3 Access Network

The Access Network provides the physical connectivity from the Cable Modem to the Cable Modem Termination System and is interconnected at the latter's location with the Core Network. This can be achieved using Layer 2 (e.g. Ethernet), Layer 3 routing or directly into the MPLS backbone. IP addressing is commonly provided from within the Access Network.

6.3.1 DHCP

DHCP (Dynamic Host Configuration Protocol) is used extensively in Cable Networks to provide IP addressing for both the Cable Modem and the public side of the customer LAN. The DHCP protocol is defined in [i.24].

The CMTS is used as a relay agent to receive broadcast DHCP messages from the clients connected to the HFC network and to forward them unicast to a DHCP server. The CMTS also inserts the following additional options into the DHCP message (Table 1).

Table 1: Additional DHCP options inserted by the CMTS

GIADDR	The IPv4 address of the interface of the CMTS on which the DHCP DISCOVER was received.
Option 82, suboption 1	The MAC address of the device via which the DHCP DISCOVER was broadcast.
Option 82, suboption 2	The Circuit ID which is generally a CMTS interface ID.

The DHCP server is often used to authenticate the request using a database or directory and assign a class of service. The service class is used to specify the reply options and their values. Service classes are used, for example, to specify a DOCSIS configuration file to download or to provide a DNS server IP addresses.

The DHCP server will make a DHCP OFFER to the client which answers with a DHCP REQUEST. After the server acknowledges the request with a DHCP ACK the process is complete.

Additional functionality on the CMTS can be used to improve security by intercepting the OFFER and encrypt the configuration file name.

A further security mechanism can be enabled on the CMTS to perform Source Address Verification. The client IP address is extracted from the DHCP OFFER and / or is learnt by sending DHCP LEASEQUERY messages to the DHCP server to confirm the MAC identity for the IP address.

DHCP resilience is generally provided by two or more DHCP servers being specified in the configuration of the CMTS. The CMTS will forward DHCP DISCOVER messages to all defined servers. The client chooses which DHCP OFFER it responds to with a DHCP REQUEST. This is generally the first DHCP OFFER received.

6.3.2 TFTP

The Trivial File Transfer Protocol [i.21] is used in Cable Operator networks to provide DOCSIS configuration files to the Cable Modems and set-top boxes. As part of its initialisation process, a Cable Modem downloads the configuration file (specified in a DHCP option) from a TFTP server (also specified in a DHCP option). TFTP can also be used for downloading a new firmware revision to a Cable Modem.

6.3.3 ToD

The Time of Day protocol [i.19] is used by Cable Devices during initialisation to acquire the current time. The time server IP address is provided as a DHCP option.

Under European regulatory requirements and guidelines, Cable Operators are required to provide real-time and historical data on which IP address has been assigned to a customer. This data is generally stored in a database and is retained for the time specified by the local authority.

6.4 CMTS

The Cable Modem Termination System (CMTS) is a router that hosts the physical interface to the HFC network and provides the routed IP connectivity to the Access Network.

On the HFC interface, the following standards are supported:

- DOCSIS 1.0 / DOCSIS 1.1
  The first generation DOCSIS cable transmission standard which was originally developed in 1997 supports upstream and downstream channels. The main difference between the North American technology option (also known as DOCSIS) and the European technology option (also known as EuroDOCSIS) is the channel bandwidth in downstream. It is based on the bandwidth required to broadcast a single analogue television channel as defined by the regional television standards. The DOCSIS standard supports 6 MHz channels, whereas the EuroDOCSIS standard supports 8 MHz channels.
  Using 256-QAM modulation, an 8 MHz channel allows for a maximum data capacity of 55 Mbit/s. DOCSIS 1.0 supports QPSK or 16-QAM modulation in the upstream using a 3,2 MHz channel. This provides a maximum upstream data rate of 10,2 Mbit/s.
  The DOCSIS1.1 [i.3] enhancement was introduced in 1999 and added QoS functionality, secure firmware upgrade and X.509 certificate-based authentication to the standard.

- DOCSIS 2.0 [i.4]
  This standard was introduced in 2001 to provide improvements in the upstream bandwidth available over the HFC network. It supports QPSK, 8-QAM, 16-QAM, 32-QAM or 64-QAM modulation over upstream channels that can be configured to be up to 6,4 MHz wide. It provides a maximum upstream data rate of 30,72 Mbit/s.

- DOCSIS 3.0 [i.5]
  The DOCSIS 3.0 standard, introduced in 2006, added several new features and enhancements to functionality. The primary improvements included support for multiple bonded channels in both upstream and downstream and support for IPv6.
  Whilst the maximum data rate for each channel has not changed compared to DOCSIS 2.0, by combining multiple channels and sharing the data across them it is possible to multiply the total data capacity available to the end user. There is no defined maximum number of bonded channels, although current technology provides up to 16 bonded downstream channels and 8 upstream channels.

6.5 Core Network

The Core Network provides the physical connectivity between multiple Access Networks, Datacentres and the transit (T)/ peering (P) partners.

The topology of the Core Network is different for each service provider, but many use MPLS (Multiprotocol Label Switching) in a full-mesh configuration to ensure resilience and reliability across the network.

6.6 Datacentre

The Datacentre infrastructure is generally used to host the Cable Service Provider's own or third party network services. These include authoritative and recursive DNS servers, email servers, web portal services, provisioning systems, OSS platforms, data storage and backup.

6.6.1 DNS

DNS (Domain Name Service) is used in IP networks to resolve domain names into host IP addresses. There are two primary functions of DNS, authoritative and recursive.

6.6.1.1 Authoritative DNS

An authoritative DNS server holds the zone files that store the mapping between IP address and host name. It is the IETF best practice [i.22] to define a reverse mapping (PTR record) for every IP address. This is normally achieved by pre-populating the forward and reverse zones with unique names mapped to the IP addresses. The structure of the name is decided by the Cable Operator, but generally takes the form <customer_number>.<region>.<serviceprovider>.com.

Zone files will also exist for all other domains that the Cable Operator is responsible for, such as for hosting and for its own email, web content, etc.

6.6.1.2 Recursive DNS

Recursive DNS servers are deployed in the Cable Operator network and are used by the clients to resolve IP addresses and names from the Internet. These servers generally also cache the results of lookups to improve performance and reduce overall DNS query rate.

Service providers generally provide additional DNS services to customers, either from internal systems or via third-party agreements.

6.6.2 Email

An email service may be supplied to a customer using standard protocols POP, SMTP and/or IMAP, as well as via a webmail portal over HTTP/S.

6.6.3 Web Portals

The service provider may allow for customer self-care functions including the ability to subscribe to, change service tier or profile or subscribe to additional services offered by the service provider or an approved third party. These functions are generally provided via a web-based portal using HTTP/S.

6.7 Support Services

Within the Cable Operator network, management of subscriptions and services as well as provisioning of devices is provided by a wide range of Operational Support Systems (OSS) tools. These tools often use industry standard network management protocols such as SNMPv1 or SNMPv2 and more recently the CPE WAN Management Protocol TR-069 [i.48].

Other OSS tools have been developed internally by the service provider.

6.8 Lawful Intercept / Data Retention

Lawful intercept is the lawfully authorized interception and monitoring of communications. Cable Service Providers should comply with the legal and regulatory requirements for the interception of voice as well as data communications in IP networks.

The IP data should be retained by the Cable Service Provider for a defined duration and made available to the legal and regulatory authorities when requested. The defined data retention period varies from region to region.

6.9 Other Functions

Other functions exist within many Cable Operator networks that are irregular or of a bespoke nature. It is impossible to provide a comprehensive list. Some examples, though, may include:

- Web content filtering: This is the blocking of illegal / immoral web content delivered by a specified list of Internet hosts. The specified list may be provided by the local legal or regulatory authority.
- Deep Packet Inspection: This can be used to enforce traffic restrictions on end users based on the type of traffic. One example is the rate-limiting of peer-to-peer file transfer. Restrictions can be imposed on overall network traffic or to individual end users based on policy rulesets.

In summary, there is extensive use of IPv4 across every component within the Cable Operator network. This includes a complex combination of various addressing schemes:

- private, non-globally routable IP addressing within the end-user LAN, which uses network address and port translation to a single globally routable IPv4 address;
- private, non-globally routable IP addressing for control and connectivity of the Cable Devices;
- globally routable IPv4 addresses used in all of the fundamental network components across the Access and Core Networks and within the Datacentre.

It is common to find a mixture of several routing protocols within the Cable Operator network. Routing across the Core Network and with transit and peering partners is generally achieved using BGP-4. However, the Access Network often uses interior routing protocols such as IS-IS [i.20], OSPFv2 [i.25] or RIPv2 [i.26] to provide routing across the CMTS to the Core Network, Datacentres or other networks.

7 IPv6 Readiness Market Analysis

This market analysis summarizes the present state of the measures MSOs are considering to date with regard to managing the transition to IPv6 and outlines the status of deployment of IPv6 within live networks and test environments. Although due care has been applied in collecting and presenting this information, it has to be taken into account that it is solely based on an assessment gathered during the study as given by the present document and reflects impressions and experiences expressed. It is mainly focused around European MSOs although some information regarding the worldwide deployments may be indicated as well.

Thorough testing by manufacturers and network operators in the lab and in the field is considered a crucial prerequisite for production-level deployment of a new technology such as IPv6. Still, it seems that many vendors, especially those focusing solely on the interests of single national or regional markets, have defined individual strategies and product propositions and have carried out limited testing themselves. This, however, only accounts for 40 % of the total market. Other vendors have been rather diligent in their aims to validate IPv6 on live networks and to establish interoperability to current configuration. Although it may be considered a deficiency that IPv6 is not inoperable with IPv4, the fact still has a single benefit. Due to the necessity to implement IPv6 in a completely separate architecture, it could be placed on a network with only a small amount of threat of affecting the present IPv4 setup. Notable exceptions are link performance and capacity values as those resources are still shared by the two protocols if configured on the same node.

It can be concluded that approximately 60 % of the market are in the process or have long completed testing and verification of a base IPv6 configuration on current products and setups and potentially even on new projects. This places European network operators in the favourable position of knowing that any kit deployed recently, currently or in the near future is capable of delivering IPv6 at this moment in time. Still the amount of resources to be shared has to be verified to be sufficient.

IPv6 Day of 2012 had less than the significant impact that was expected at least from the point of view of an access network provider. The ISPs - although well ached in testing - did not push IPv6 into a major percentage of the customer CPE footprint. The reason was mainly the lack of confidence in the capabilities of transition technologies that are required to ensure that all services continue to be available also to customers connected via IPv6. The CPE representing the highest value in terms of capital investment and customer impact on the one side and the technologies for transition still being tested with no fixed expectation on which would win the development race and would allow full customer service on switch over on the other side, left European MSOs in a state of preparation instead of full readiness for deployment.

For the Core Network, products have been ready to a certain degree for deployment since 2003 and to a certain degree before that. The only features not completely implemented to date are mainly security features. On the Access Network side, there is still a lack of development completion, with many features not available for IPv6. In general, however, the market is ready for IPv6 native deployment with products being available. In terms of deployment, though, the status is significantly less advanced. The majority of ISPs have implemented IPv6 in the Core and Access Network or at least have deployed network components that are IPv6 capable and are preparing to implement. CPEs are to follow, but there are less than 25 % of the homes connected on IPv6 within Europe at present. Full deployment to the customer within Europe is expected to be around 40 % by mid-2013. To the date of the present document there is no reason why most ISPs cannot deploy IPv6 in a native fashion with increasing function and complexity later on.

8 Support of IPv6 on Cable Network Components

8.1 Overview

This clause provides a high-level view of the present status of IPv6 development for various aspects of the technologies that are used to establish end-to-end IP connectivity in a Cable Network. Particular attention is given to:

- Peering
- Routing and Core Network
- Access Aggregation
- CMTS/Headend
- HFC and CPE
- OSS and BSS
- DHCP and DNS
- Firewalls and Load Balancers
- Backend Support Services
- Web services
- DTV services
- Network management and monitoring services
- Data Retention and Lawful Interception
- Home services

Some aspects of technologies and services deployed in Cable Networks are considered beyond the scope of the present document with regard to their readiness for IPv6 transition. Those include SIP services and Voice services. They are not part of the IPv6 development analysis due to the lack of an immediate requirement for MSOs to provide them over IPv6.

In the course of the analysis of the components of end-to-end IP connectivity justification for the implementation of transition technologies is provided by comparing their effect to the level of IPv6 deployments within the Cable Industry. The analysis reveals the restrictions in terms of technical or logistical aspects that are forcing MSOs to consider deploying technologies that enable a transition to IPv6 instead of deploying full IPv6/IPv4 dual stack from the beginning.

The market expectation for ISPs is the delivery of end-to-end connectivity to enable consumption of services by their subscribers. That expectation leaves them in a position of little to no control over the final transition path to IPv6 due to the dependency on the IPv6 readiness and IPv6 support of software and devices that are not managed by the ISP or only in a very limited way. That applies for example to the operating systems that are running on customer end devices as well as on applications and on other software stacks in the home. No matter what Cable Operators do to their network they will, finally, be responsible to perform a gap analysis and, at some point, face the decision of service deprecation. It is important to understand what devices and services have to be IPv6 ready in the end-to-end delivery chain before IPv4 connectivity of the home can be dropped without the risk of services becoming unavailable to the customer. It requires an in-depth analysis to ascertain the current level of capabilities with currently deployed equipment. Generally, the market is doing well on the development of IPv6 and transitioning technologies across products and services; but the issue is the timely identification of the gap that will be left and the development and testing of technologies to close it. Current efforts to identify the gap and to develop standardized solutions to address it are delayed mostly due to challenges in establishing coordination and agreements among industries towards a target date for IPv6 readiness. There is even a lack of consensus whether transition technologies are required at all or if the approach to move to IPv6 natively or to consider dual-stack solution should be preferred.

Consequently, the challenge for MSOs is to accept that the introduction of IPv6 is prone to a gradual transition. By planning an IPv6 deployment early, an MSO could minimize risks, costs and complexities. However, the need for a transition technology is unlikely to be eliminated completely. Thus, the most important IP network standard today - and this is not only applicable to the Cable Industry but to any ISP - is the transition standard.

The requirements that would influence the decision for an MSO to transition to IPv6 are essentially twofold:

- Depletion of available globally-routable IPv4 addresses, with no access to a new pool.
- Desire to enable customers to reach IPv6-only services that are being deployed on the Internet, which requires support of IPv6 connectivity.

With a gradual transition starting early, MSOs acquire time to settle into the new technology and find the gaps in the deployment. The present document considers the technologies that will establish the basis for this transition, providing details on their state of development, analyse their features, applicability and restrictions in the context of a Cable Network and provide details on deployment considerations.

8.2 Peering

When discussing peering from a Cable ISP's viewpoint the focus would be on the ability for the customer to access the general Internet from within the MSO's network. Peering establishes this connectivity, so it is a logical place to start an end-to-end consideration although it - strictly speaking - does not represent the far end of a communication session. Peering might not be considered a "hardware" based technology or a "feature" based technology as such but it is a fundamental technology in an MSO's transition to IPv6. IPv6 peering is the logical extension of any transitional connectivity within an MSO's network.

Current IPv6 routing tables in eBGP are approaching sizes of 9 000 routes. This average figure increases per day. IPv4 is, however, increasing by leaps and bounds to an even greater degree. IPv4 routing tables now have hit sizes of 400 000 routes, which is potentially accompanied by resource issues, and will have an impact on transitional technologies and hardware requirements. As the IPv4 address space gets further depleted the few remaining addresses will get assigned in an increasingly less structured way. "supernets" will be broken up, longer prefixes will be accepted by peers and subnets will generally be displaced. As a result, the lack of contiguous prefixes in IPv4 will increase exponentially. As the IPv6 table erupts into the 100 000 mark so the IPv4 table may push into 500 000 to 600 000 routes. Even if the resource issue is solved via increased memory, the convergence times and route-dampening implications are remaining and are already expected to become extreme.

It is useful to give some consideration to the question why IPv6 will increase and the gradient to the transition will be steeper than might have previously been expected. Transition technologies still provide a service degradation and add latency at least in some flows. Especially peer-to-peer traffic will be affected, although the technologies are designed well and are expected to have little or no impact in general. As a result, it is likely that the native IPv6 transport will be preferred over the private IPv4 solution provided by the Cable Service to enable IPv6 connectivity which is delivering an almost but not exactly equivalent performance. So the mass use of the transition technologies may be measurably shorter than would otherwise have been expected. This is accelerated by the fact that more services that are delivered via the Internet are being placed on IPv6.

It should be noted that eBGP ingress and egress filtering changes dramatically with problems concerning ICMPv6 (due to the nature of NDP). There are also major issues with PMTU that may force the MTU down to a fixed value of 1 280 Bytes to prevent black-holing. This is owed to the fact that there is no fragmentation on IPv6 in transit.

In summary, peering is IPv6 capable. It has issues, but in general it works well. For the few remaining issues, there are workarounds that accommodate the required functionality.

8.3 Core Network

NOTE: In the general case, the Core Network implements centralized iBGP. However, some MSOs operate a Core Network with MPLS where special considerations in terms of IPv6 support apply.

Routing is the set of Internet functionality that provides forwarding capabilities between hosts that are located on separate L3 networks. For a Cable Network, two types of routing can be distinguished. Firstly, the Core Network (which is discussed in this clause) and secondly the Access Aggregation. MSOs that operate an MPLS core that currently has no native IPv6 implementation on RSVP or LDP use instead a method of IPv6 "tunneling" implementing 6PE or 6VPE using the IPv4 FECs.

Routing is enabled by IPv6 as a L3 protocol. However, IPv6 enables functionality on L2 as well and leverages resulting information. Within NDP, IPv6 translates into a new EtherType, the protocol uses ICMPv6 for messaging.

Above Layer 3, transport services on the source host pass data in the form of TCP segments or UDP datagrams down to the layer where IPv6 operates. This actually does not change. Transitioning considerations are mainly applicable to the applications that focus on the IP packet, L3 and local link functionality. This functionality is typically implemented in hardware establishing a forwarding plane that has equivalent or better performance compared to IPv4.

At present, 6PE, 6VPE and IPv4 IGPs for MPLS are used to deliver IPv6 through the Core Network. These technologies establish a suitable platform that allows a provider to implement a dual stack for BGP with both L3 protocols. The only requirement is to align the feature sets. For example, IPv4 CoPP and QoS are well defined and encompass significant functionality, whereas for IPv6 considerable gaps have been identified for these features especially in tunnelled topologies.

8.4 Access Aggregation

PE iBGP to the Access IGP is well understood within the data-networking domain. 6PE and IPv6 IS-IS/OSPFv3 provide ample functionality to interconnect to the CE Access Aggregation. Route summarization and aggregation are fully configurable/implemented within most IGPs. Intermediate L2/L3 switches are also generally well defined for IPv6 with profiling and forwarding carried out in hardware in most cases at both L3 and L2.

Access Aggregation, on the other hand, is still under development and cannot be considered stably defined and implemented on the CMTS. However, similar to the PE functionality as described above implementations of IPv6 functions are available for most features although with limitations in some areas. After completion of the development, a fully stable Access Aggregation topology from PE to CMTS running native IPv6 IGPs with an extremely close proximity to the IPv4 feature set is expected to become available.

As a result, the capability to transport IPv6 packets through the CMTS on to the PE routers is determined by the state of CMTS development. This is the bottleneck in IPv6 feature progression. However, with a limited feature set IPv6 transport can be deployed for Access Aggregation at present. It can be expected that Access Aggregation will become capable of supporting a full, non-static IPv6 topology environment on the PEs and intermediate switches in the short timeframe.

8.5 CMTS/Headend

After some delay compared to other network components, CMTSs can now be considered fully developed for IPv6. IPv6 support is constantly improving in its stability, feature sets and interoperability on all currently deployed platforms.

Native IGPs, at least IS-IS, is ready for deployment in most topologies, with filtering, NDP and many other interfaces fully complying to the applicable standards. Native and dual stack interfaces with PD and SLAAC capabilities are commonly available.

8.6 HFC and CPE

DOCSIS 3.0 [i.5] is the IPv6-ready technology for Cable Modems. The original specification was published in August 2006 and, subsequently, standardised by ETSI. The widely used DOCSIS 2.0 [i.4] even today does not fully support IPv6. The standard has been augmented by a DOCSIS 2.0 + IPv6 Cable Modem specification [i.49]. However, most ISPs that support DOCSIS 3.0 have not deployed IPv6 across their networks, yet, and those that have are only looking at testing or field-trials for friendly customer validations. There are very few full IPv6 deployments among Cable Operators today.

It is considered that for most networks transitional requirements are imposed in order to manage the development into DOCSIS 3.0 with IPv6 or DOCSIS 2.0 + IPv6. Even where only the IPv6 EtherType is required for transit traffic while management traffic stays on IPv4, transitional requirements have to be taken into account. Management interfaces on CPE are likely to be the highest priority for IPv6 implementation at most MSOs as not having to hand out IPv4 addresses for management interfaces anymore is expected to yield the most benefit without the risk of affecting customer services.

In conclusion, DOCSIS does not create a barrier to IPv6 deployments.

The HFC distribution network as a Layer 1 physical transport platform using RF technology has little concern with IPv6. It interprets the stream of IP packets to be physical signals and has no visibility to its logical structure. Multiplexers, ILAs and cross connects on fiber or coax will not be affected by the IPv6 transition or have to be changed. Therefore, the only component in the HFC infrastructure other than the CMTS that is impacted by IPv6 is the CPE. In the past, L2 and L3 Cable Modems, E-MTAs and eRouters have followed quite unpredictable development paths depending on individual MSOs' requirements and vendor strategies. Only recently, development has been pushed towards IPv6 capable devices. Even bridging modems (i.e. Cable Modems that are limited to L2 functionality) face requirements for IPv6 forwarding which is enabled by knowledge of the frame structure and IPv6 EtherType.

The CPE represents the largest part of equipment cost for any Cable Operator which exceeds the investment for CMTSs, Core Networks and backend devices multiple times, up to 10 times if not more in some cases. CPE also represent the largest number of pieces of hardware deployed in the network. The industry has seen since 2002 CPE vendors taking some responsibility towards IPv6 but without actual native compatibility. Tunneling was the optimal choice but Teredo pass through was the most valued choice among home devices connected to Cable Networks. Taking into account the inherent incompatibility of the two protocols, IPv6 proved itself to be a rather complex implementation impacting almost every functional component within the CPE.

Currently, complete tunnel implementations, IPv6 friendly firewalls, DHCPv6 PD and SLAAC, IPv6 accessible Web GUIs, RAs for RDNSS and DNSSL, Prefix Suppression within the RA and finally UPnP support are available from the CPE market. So, CPEs can be considered to be almost completely developed for a basic dual stack or native IPv6 implementation. It can be stated that IPv6 deployments will not be slowed down due to development on the CPE or the L3 integration of IPv6 anymore. Aspects of logistics, mass swap outs, cost and confidence on the end-to-end delivery will be final reasons why IPv6 transitional technologies may be required for the next 3 to 8 years in the cable industry.

8.7 OSS and BSS

All Cable ISPs provide advanced IP services - whether it is IP telephony, tiered high-speed data services, home networking or subscription media. These services cannot scale without tools to efficiently install, provision and manage network devices, as well as activate and support subscriber services. In short, Operational Support Systems (OSSs) are the oil that keeps the broadband IP infrastructure and services engine running smoothly. OSSs are by nature extremely valuable to the service the ISP offers.

High-speed Internet access is the dominant IP service offering in Cable Networks. During the past five years, MSOs have made significant progress in improving Cable Modem installations and subscriber service activation. When done manually, the service activation process requires a telephone call from the prospective customer to a Cable Operator service center. However, MSOs are increasingly automating the ordering and activation process through the use of Web-based tools that allow user self-registration. Additionally, many Cable Operators offer consumers self-installation options, so they can purchase and connect Cable Modems themselves.

During the service provisioning process over IPv4, services are linked to the subscriber's account and then instantiated in the network through device provisioning processes. Service provisioning can either be done manually by a MSO customer service representative or directly by the consumer through a Website. The latter is, obviously, preferable as it provides a customer self-service option, reducing costs and time to service activation.

For high-speed Internet access, service options that should be provisioned include the customer's Internet service provider (ISP), service class (defined by access speed or other attributes), E-Mail addresses and personal Web space. A home networking service adds enhanced data features, such as a managed firewall or virus protection and end-device authentication for a wireless network. Content service provisioning may include enabling access to targeted video, audio or gaming services. Telephone service provisioning is significantly more complex. In addition to provisioning the underlying data service, telephone lines and numbers should be assigned along with enhanced services such as voice mail, caller ID, call forwarding and more.

Once services have been selected and assigned to the customer, devices on the network should be provisioned to deliver these services, typically via Simple Network Management Protocol (SNMP). For all Cable IP services this includes configuring DNS, Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), SYSLOG and Time of Day (ToD) servers.

Considering the degree of integration and complexity of these services, IPv6 is a requirement to future-proof such a valuable component as OSS. The definition and implementation of application and network requirements for OSS is, actually, well advanced on IPv6 support and can be deployed today.

BSS is just as valuable as OSS. Defining the requirements for BSS as a whole within the MSO's infrastructure is extremely difficult as the degree of topologies, applications and type of deployments makes it impossible to generalize. BSS have to be considered individually as per MSO. But as a coarse estimate, it can be assumed that the development for IPv6 support is in progress or already complete to the point of delivery for certain topologies.

8.8 DHCP and DNS

DHCP is currently the standard host configuration protocol for the TCP/IP protocol suite for IPv4 with development work being done since the early 1990s on DHCPv6. There are still some gaps to be addressed in future revisions of DHCPv6 but, generally, the protocol is functional and deployable.

While most of the changes that IPv6 brings impact technologies at the lower layers of the TCP/IP architectural model, the significance of the modifications means that many other TCP/IP protocols are also affected. This fact often gets neglected. It applies particularly to protocols that work with addresses or configuration information including DHCP. For this reason, a new version of DHCP became necessary when IPv6 was introduced.

One of several enhancements introduced in IPv6 is an overall strategy for easier administration of IP devices, including host configuration. There are two basic methods defined for auto configuration of IPv6 hosts:

- Stateless Autoconfiguration
  A method defined to allow a host to configure itself without help from any other device.

- "Stateful" Autoconfiguration
  A technique where configuration information is provided to a host by a server.

The decision which of these methods to use depends on the characteristics of the network. DHCPv6 provides "Stateful" Autoconfiguration for IPv6. As with regular DHCP, DHCPv6 may be used to obtain an IP address and other configuration parameters or just to get configuration parameters when the host already has an IP address.

DNS has to deal with a similar resource limitation as described in clause 8.3 in relation to the Core Network. IPv6 capable DNS servers are expected to handle requests for both AAAA and A records simultaneously and then define their priority. The DNS can be considered as a fully functional and proven IPv6 service worldwide whilst lacking some security gaps and "white listing" capabilities. The complete validity of the service on IPv6 has been verified for several years with the experimental network 6bone.

8.9 Firewalls and Load Balancers

Firewall products are generally well advanced on IPv6 implementation. Full rule base, packet decryption, (specific packet field integer rules,) acceptance of all lengths of prefix, NDP ICMPv6 separation, VRRP, all interface configurations and general policies all based on IPv6 are among the important features that are available. Basically, Firewalls and Load Balancers are completely IPv6 compliant.

8.10 Backend Support Services

OSSs as considered in clause 8.7 cover the DHCP, CPE, CMTS backend requirements while BSSs deal with the management of customer services. In addition, MSOs are deploying an extensive backend infrastructure that is required as a key function to manage the communication functionality, provide services such as backup, data center services and general corporate applications to allow reporting, administration and management. These systems tend to be neglected when setting up an IPv6 strategy due to their fundamentally invisible operation and lack of direct effect on subscriber connectivity.

For this device category, it can be determined that OSs, backup software, corporate finance application clusters as examples are actually fully compliant with IPv6. There are only a few holes in application awareness on IPv6, but this is identified as a minor gap.

8.11 Web Services

When migrating Web services from IPv4 to IPv6 within an MSO's infrastructure, simply making a Web server ready for IPv6 has been identified to raise serious problems whereby for example users could experience errors when they request services that were not yet supported by IPv6. In addition, a suitable solution requires that the organization's network be protected from malicious IPv6 accesses, although there are no comprehensive security solutions for IPv6 unless a separate firewall is deployed. Reverse proxy solutions have been implemented as a workaround. However, the indication is that Web services are still not fully stable with IPv6.

It was identified that one of the simplest solution to make Web services IPv6-ready is a Web server that supports both IPv4 and IPv6, e.g. IPv6/IPv4 dual stack. Since the essence of a Web service is to allow to connect various services to each other, making a single server ready for IPv6 may cause various errors by the services that are not yet migrated to IPv6. In addition, the DNS may have issues with potential mismatches for double entries.

Another solution for Web services is to duplicate the present IPv4 system for IPv6. This is a solution that allows for the prevention of errors by modifying links to the services that are not ready for IPv6. However, this solution increases the content management cost because double sets of Web contents has to be managed consistently.

The most consistent topology is not the IPv6 infrastructure that has been isolated not only logically but also physically from IPv4 services. Actually, IPv6 has been designed with much regard to unhindered coexistence with IPv4 in the sense that there is no integration. This approach would have created no extra need for transitional technologies if the IPv4 address pool was not being exhausted. The ultimate issue is that private or public IPv4 addresses cannot simply be "translated" into an IPv6 address within the protocol and, thus, protocol translators/transitional devices/services are required.

To minimize the influence of duplicated IPv6 Web services on still existing services based on IPv4, a reverse proxy server was identified as a best current practice. There can be many reasons for this, such as:

- A reverse proxy server does not require any modification to the running IPv4 services.
- Cache technologies can be applied on the proxy server to keep the number of requests forwarded to the Web servers extremely low. Since the total volume of contents the reverse proxy server handles is relatively small (usually below several Gigabyte), the reverse proxy server can achieve a high hit ratio.
- It is possible to preserve the security of an IPv4 ecosystem by setting up the proxy server outside the firewall of the IPv4 network. This prevents the IPv6 proxy server from being a backdoor host that breaks the security of the network.

Features and topologies become available for Web services and are generally simpler to add on top of the current IPv4 setup, irrespective of the deployment option chosen.

8.12 DTV Services

Digital Television Networks are shifting from the mere transport of native MPEG-2 Transport Streams to the provision of IP-based services, either unidirectional or fully interactive. The inclusion of an IPv4/IPv6 stack in a DTV terminal requires it to receive configuration information from the network such as host, gateway and DNS IP addresses. The DNS IP addresses are configured either statically or dynamically. DTV generally has experienced significant variations concerning the architecture and integration of the IP stack within a set-top box. Typically, these devices were deployed with overlay IP stacks that are completely ignoring the OS stack and bypassing it through the kernel. Consequently, it becomes tedious to apply standard transition technologies to each of the proprietary IP implementations. Upgrading the firmware is in general not possible to the degree required for the introduction of IPv6 support, in addition to the inability of the connected network components to perform such a software upgrade. Despite this, the alignment of TV and data networking in terms of transport platforms and protocols is improving with more IP deployment options and integration of services. Considerations are given to the use of an all-IP approach with an evolved version of the DOCSIS protocol as the fundamental transport platform. In any evolution in network access, MSOs have to take into account legacy DTV equipment that is and will continue to be operating in Cable Networks for a considerable amount of time.

For simplicity reasons in terms of this analysis, the term "Digital Television Network" is assumed to refer to a single DTV broadcast service conveying a single MPEG-2 Transport Stream (or other services) to a certain coverage area. Even if delivered across IP, the Digital Television service still refers to a single broadcasting service. It seems that IPv6 is not given significant consideration by MSOs for these services unless it is made available by default in new and upcoming kit. Indications are that the technology used for IP-based delivery of Digital Television is well delayed in its consideration of IPv6 with some notable exceptions for new DTV platforms and devices.

8.13 Network Management and Monitoring Services

Network Management Systems (NMS) are fast becoming IPv6 compliant and are reaching a similar level of feature richness and functionality as in an IPv4 environment. There are, generally, certain missing items such as MIBs within SNMP, SYSLOG events on NDP DAD, IPFIX/Netflow v9 (IPv6 and IPv4 integrated templates) and IPv6 DPI. Although there are significant gaps, MSOs are capable of monitoring their IPv6 network effectively.

It has been identified that there remain continued issues with element management and IP Service Level Agreements for 6PEv4 FECs since it is not possible to monitor an LSP for IPv6 forwarding capability. This applies to a single vendor scenario and even more in an environment with multiple vendors. With the present development of most NMS solutions used in Cable Networks, operators are unable to monitor the IPv6 network infrastructure beyond the bare minimum but we hope to have IPv6 LDP native RFC draft completed by the end of the year for acceptance and deployment.

8.14 Data Retention and Lawful Interception

With Data Retention regulations applicable at a national level in several European and other countries, requirements are placed on the ISP to track potential user activities and log a record of the traffic and subscriber details for each communication. This enables the ISP to make available a reconciled history of communication traffic on their network.

IPv6 puts new challenges for network administrators in the context of user identification and traffic interception especially with the lack of security capabilities within the protocol or features employed in the network. Unlike IPv4, an IPv6 address no longer uniquely identifies a user or host. The IPv6 address can be randomly generated and can keep changing all the time dependent on the network configuration. Hosts with IPv6 stack can also communicate via predefined tunnels over IPv4 infrastructure without the ISP being able to perform DPI on every single flow due to the capacity limitations of most of the network equipment. Consequently, tunnelled traffic mostly bypasses network security implemented via firewalls, monitoring and LI. The solution for IPv6 requires an extended set of monitoring data to be collected from network devices with extended features without effecting neutrality of the network. These new solutions are either a duplicate of the present IPv4 structure or add Netflow with its "smart" logging similar to the situation within certain implementations of DS-Lite.

IPv6 DR and LI are well covered in most cases by both CMTS deployments and separately on other network devices. The new deployments are being executed and the regulatory authorities are cooperative on the requirements during the IPv4 to IPv6 transition. There are deployments for both DR and LI on IPv6 but the full extent of the capabilities and the differences in the topologies makes those deployments still in its early days for traffic inception and retention on IPv6.

8.15 Home Services

Home services are an area where most of the technology issues for IPv6 readiness have to be considered significant. It is the most numerous and most diverse section of the end-to-end communication path, although not owned nor deployed nor managed by the ISP itself. Home content, applications and devices controlled by the customer provides the strongest challenge to IPv6 deployments.

Consequently, IPv6 deployment strategies have to take the state of development of devices in the home and of the software and application as a given. Even those parameters, however, are rather undefined and changing over time. Still, it means little even if there are 95 % of the applications and home devices IPv6 ready. It remains a marketing decision for a Cable ISP when it is feasible to not enable connectivity for legacy devices that are only capable of an IPv4 stack. Similarly, issues remain with old Smartphones that are intended to use the connectivity provided by the home network but cannot be upgraded to support IPv6. These are examples of some of the major issue that are being faced by a Cable ISP in the context of home services. As an implication, MSOs are more significantly restricted in their ability to deploy IPv6 not necessarily by the degree of IPv6 support on the Cable Network but by legacy home devices owned and operated by its customers.

Despite the legacy market which encompasses not only old devices but also newer devices lacking an IPv6 stack, handheld devices are good examples of straight-forward IPv6 native and dual stack capabilities. Taking into account the proliferation of this device class particularly in the home the legacy issue may be mitigated in the near future. Another challenge is accommodating the somewhat 96 000 approved applications running on PCs and workstations and their current OS releases whereby these applications are not designed to support IPv6. The issue becomes even more significant when considering the 1,3 million approved handheld applications for the mobile market.

It can be summarised that the home network, customer controlled devices and a plethora of applications restrict Cable Operators to transition to a full native deployment of IPv6 end-to-end. With depletion of IPv4 addresses to provide connectivity to these IPv4 only devices and applications, an IPv4 gateway is required to be provided that enables transport of IPv4 traffic between the home and the Internet across whatever IPv6 infrastructure the Cable Operator chooses to deploy.

9 IPv6 Standardization

While the IETF started working on IPv6 in 1990 limited IPv6 adoption to date in residential environments has been seen. Although governments/institutions have been pushing service providers to adopt IPv6, there has never been a real business driver for service providers to introduce IPv6. There were no new applications that required IPv6 and the cost of introducing IPv6 was perceived too high. All of this has resulted in the situation that service providers had very little focus on introducing IPv6 in residential broadband networks to date. However, with the exhaustion of public IPv4 addresses and the very rapid adoption of Smartphones and M2M devices resulting in increasing demand of transport capacity and addressing resources, the introduction of IPv6 should be pursued as quickly as possible in order to sustain the growth and to provide customers a proper service.

Technically it can be seen that IPv6 is incompatible with IPv4 and has introduced several new concepts that change the present mode of operation of broadband networks:

- IPv6 addressing schemes: unicast addressing including LLA (Link Local Addresses), GUA (Global Unicast Address) and ULA (Unique Local Address), multicast addressing, depreciation of broadcast addressing;
- IPv6 header format: Next Header, etc.;
- SLAAC: Stateless Autoconfiguration for IPv6 address assignment without involvement of a DHCP server;
- Default router support using Router Advertisements (RAs);
- DHCP PD: DHCP Prefix Delegation to assign prefixes to home networks;
- Neighbor Discovery (ND), MLD (Multicast Listener Discovery), etc. supported through ICMP.

Although there have been many good reasons for these changes, the concepts have implications on how IPv6 can be offered to residential subscribers.

In addition, IPv6 poses multi-dimensional issues with requirements over which service providers have no or very little control. IPv6 support has implications on multiple elements as given below and as illustrated by Figure 4:

- End devices hardware/Operating System
  - PC: MAC OS, Linux, Windows Vista/7 have good IPv6 support while Windows XP operates only in dual stack mode and Windows 98 has no IPv6 support
  - IPv6 support on mobile devices is emerging (Symbian, iOS, Android, etc.)
  - VoIP protocols have little IPv6 support to date
  - IPTV OS/STB have little IPv6 support to date
- CPE/RG
  - Support for IPv6 is emerging on xDSL/GPON/Ethernet/Cable
- Access nodes
  - DSL/GPON/Ethernet: most vendors start supporting certain architectures for IPv6
  - CMTS: most vendors support IPv6
- Aggregation/Edge/Core Network elements
  - Most of the devices support IPv6 since many years and have been deployed for a significant amount of time
- Fixed (BNG/BRAS), mobile (GGSN/PGW) edge nodes
  - BNG/BRAS: Most vendors start supporting IPv6 for PPPoX, IPoE (DHCPv6/DHCPv6 PD), LNS
  - GGSN/PGW: Most vendors have support for IPv6 architectures defined by 3GPP Release 8 and Release 7
- Applications
  - IPv6 support of end-user applications is provided:
    - If they are supported on the proper OS
    - If they use the IPv6 API(s) to support the IPv6 network connectivity
  - Websites can be considered IPv6-ready:
    - If they support IPv6 addressing/connectivity
  - CDN support IPv6:
    - If they support IPv6 addressing/connectivity

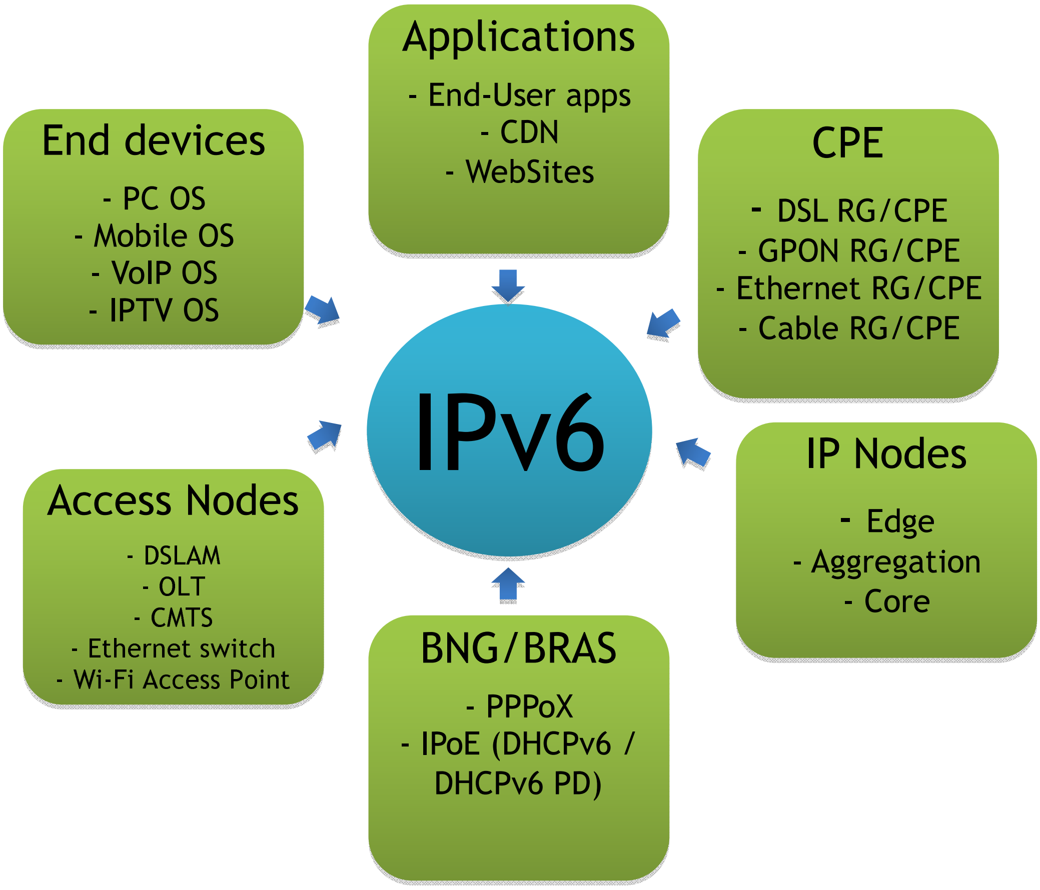

Figure 4: IPv6 introduction, the multi-dimensional problem

Some of the implications on these elements are dependent on the network design chosen for introducing IPv6. The following clause highlights the implications of certain scenarios. The analysis focuses on unicast IPv6 connectivity as IPv6 for multicast is still less relevant since most multicast IPTV platforms are not ready yet to move to IPv6 and most IPTV solutions do not require public IPv4 address space.

A list of the applicable standards and a technical summary of the RFCs relevant for IPv4, IPv6 and transition technologies is provided in Annex A and Annex B, respectively.

In summary, most technologies are well advanced in their development of IPv6 support or have completed their IPv6 implementation. There are technologies that still significantly lack consideration of IPv6 but this does not prevent MSOs from pursuing a dual stack topology and executing their IPv6 strategies. A notable exception to this positive resume is the home network with legacy equipment and applications not willing or not capable to transition to IPv6 and the potential for many IPv4 Internet services being requested for a long period of time. This requires the Cable ISP to maintain an IPv4 address within the home and IPv4 connectivity to the home which forces the infrastructure to support transition technologies once the MSO has depleted its IPv4 address pool. This situation will be faced by most of the MSOs within the next 5 to 8 years.

10 IPv4 to IPv6 Transition Technical Analysis

Global IPv4 address space is currently projected to be depleted around the middle of 2012 to 2015 for most MSOs. As part of the resulting rollout of IPv6 in the worldwide networking footprint specific measures should be taken to allow a smooth transition and coexistence between IPv4 and IPv6 capable network devices and services. This clause presents a technical summary of the key transition technologies available today as listed in the table below.

In order to provide a comprehensive overview of transition technologies that have been proposed, all of that are developed to a certain level of maturity were considered without pre-judging their suitability for MSO requirements. The benefits and drawbacks are analyzed and prioritized against the specific characteristics of Cable Networks. Subsequently, a judgement was made to determine which transition technologies should be analyzed in greater detail.

With this process even deprecated technologies could be considered for re-development if they provide increased benefit for the MSOs' transitional needs.

Technology	Technical Description	Summary Analysis
For further consideration within the present document		
NAT64	*Using NAT64, only native IPv6 connectivity is provided to the end customer. End-to-end IPv6 connectivity is provided natively between end-hosts and Internet services. To provide communication from IPv6 to an IPv4 host or Internet service a NAT64/DNS64 service needs to be deployed in the network.*	*NAT64 is a transition technology using IPv6 in the home network. The main caveats being* *1 - Applications who do not use DNS will not work across the network.* *2 - It requires the complete home network to be IPv6 enabled and it does not support IPv4 only devices in the home network.* *3 - It requires an IPv6 native transport from CPE to the NAT64 device.* *4 - Extensive ALGs are required due to the fact that the NAT64 device translates IPv6 to IPv4 and vice versa* *5 - It requires PCP to prevent service deprecation on subscribers who previously had IPv4 public addresses and provided Internet access to their home.* *6 - Deployment requires NAT64 to be carrier-grade in terms of its performance and features.* *NAT64 is not widely adopted as a transition technology so far given the above complications mainly mentioned in items 1 and 2.*
dIVI	*The CPE router is provisioned with an IPv6 prefix, a public IPv4 address and a port range. Where IPv4 address sharing is not required, the port range specifies all ports. For IPv4 hosts, IPv4 to IPv6 address translation is performed at the CPE router using the stateless mapping scheme as defined in [i.39], whereby the IPv4 address is embedded within the IPv6 address. The suffix extension is used to include the port set ID and the port range as defined in [i.10] within the IPv6 address. The packet is forwarded to the XLAT. The XLAT restores the IPv4 header by extracting the address and port from the IPv6 address. The restored packet can be forwarded to its destination.* *In reply, the XLAT replaces the IPv4 header with an IPv6 header based on the same port mapping algorithm. In this way, no state information is required to be maintained in the XLAT. The CPE router replaces the IPv6 header with an IPv4 header with address and port extracted from the IPv6 address.*	*dIVI is a dual stateless IPv4/IPv6 translation mechanism. It is an extension of the 1:1 stateless IPv4/IPv6 translation (IVI) mechanism. The major benefits of the extensions are using public IPv4 addresses more effectively and removing the requirements of DNS64 and ALG, without losing the stateless, end-to-end address transparency and bidirectional-initiated communications. The use of two stateless IVI translations allows IPv4 traffic to be transported across an IPv6 network.*
Teredo	*Teredo is an address assignment and automatic tunneling technology that provides unicast IPv6 connectivity across the IPv4 Internet. As such, it is an alternative to 6to4. However, 6to4 works well when a 6to4 router exists at the edge of the provider domain. The 6to4 router uses a public IPv4 address to construct the 6to4 prefix and acts as an IPv6 advertising and forwarding router. The 6to4 router encapsulates and decapsulates IPv6 traffic sent to and from site nodes.*	*Teredo is a client technology created by Microsoft to deploy IPv6 on servers and relays using brokers. This technology can be used within the own network, but it does not facilitate a Cable Operator to deploy IPv6. It only facilitates the ability to do IPv6. This technology in its current form is not worth for further consideration. But if it is extended to the reverse methodology it should be analysed further.*

DS-Lite	*DS-Lite is based on a tunnel between the CPE and the AFTR. This tunnel is created through a SI of an IPv6 encapsulation of the IPv4 packet from CPE to AFTR. The IPv4 customer side is based on private addressing and is translated to public addresses at the AFTR. The CPE, therefore, is only required to perform a single encapsulation and the AFTR is performing two, one to encapsulate and decapsulate an IPv4 within an IPv6 packet and the second one to translate private to public IPv4 addresses via NAT.*	*DS-Lite is a feature-rich technology with a few caveats.* *1 - It requires a new CPE to be delivered to the customer location supporting the features required for SI.* *2 - It requires an IPv6 native transport from CPE to AFTR.* *3 - Deployment requires DS-Lite to be carrier-grade in terms of performance and features.* *4 - It requires PCP to prevent service deprecation on subscribers who previously had IPv4 public addresses and provided Internet access to their home.* *DS-Lite is a commonly supported in devices connected to Cable Networks due to the functionality potentially built in. DS-Lite does not require to further assign IPv4 private addresses in the network and uses a common form of NAT for IPv4 which is a proven technology.*
Not for further consideration within the present document		
4in6	*4in6 refers to tunneling of IPv4 in IPv6. It is an Internet interoperation mechanism allowing IPv4 to be used in an IPv6 only network. 4in6 uses tunneling to encapsulate IPv4 traffic over configured IPv6 tunnels as defined in [i.27]. 4in6 tunnels are usually manually configured but they can be automated using protocols such as TSP to allow easy connection to a tunnel broker.*	*4in6 is a dual stack transition technology. The main caveats being* *1 - The end device (CPE router) needs to support 4in6 and a 4in6 router needs to be added to the network* *2 - The 4in6 tunnel endpoints need to be manually configured which provides an operational overhead* *3 - The network needs to be upgraded to support full IPv6* *4in6 is mainly used in enterprise environments to provide IPv4 connectivity over IPv6. The principles are adopted in DS-Lite, which is one of the more generalized transition mechanisms in the residential markets. This technology is not for further consideration as its main principles are covered in more developed transition technologies.*
6rd	*When deploying 6rd, hosts are supplied with dual stack addressing. Both IPv4 and IPv6 addresses are provided to the customer. IPv4 connectivity in this model is provided the same way as it is done today, using private or public IPv4 addressing. IPv6 connectivity is provided using 6to4 tunneling where the standard 6to4 prefix 2002::/16 is changed to an IPv6 prefix that belongs to the ISP-assigned address space. The IPv6 prefix allocated to the end customer is derived from the IPv4 address assigned to the CPE. The v4suffix-length, v6prefix-length, 6rd Border Relay IPv4 (anycast) address and 6rd service provider prefix are provided to the CPE using DHCP. To deploy 6rd, a 6rd CPE needs to be deployed in conjunction with a 6rd border relay which provides 6to4 tunneling [i.38].*	*6rd is a dual stack transition technology. The main caveats being* *1 - The CPE needs to support 6rd and a 6rd border relay needs to be added to the network* *2 - Since the IPv6 address is provided from an IPv4 address, while IPv4 exhaustion is happening, the IPv6 address is subject to NAT with its implications and a second transition to native IPv6 support will be required, which adds costs* *3 - It requires PCP to prevent service deprecation on subscribers who run services from their home* *4 - Deployment requires the 6rd border relay to be carrier-grade in terms of performance and features.* *6rd is mainly used when the access network cannot transition to IPv6. Since a second transition is required to native IPv6, the technology is not widely adopted.*

6to4	*6to4 is a system that allows IPv6 packets to be transmitted over an IPv4 network without the need to configure explicit tunnels. Special relay servers are in place that allow 6to4 networks to communicate with native IPv6 networks. 6to4 is especially relevant during the initial phases of deployment of full, native IPv6 connectivity, since IPv6 is not required on nodes between the host and the destination. However, it is intended only as transition mechanism and is not meant to be used permanently. 6to4 may be used by an individual host, or by a local IPv6 network. When used by a host, it should have a global IPv4 address connected. The host is responsible for encapsulation of outgoing IPv6 packets and decapsulation of incoming 6to4 packets. If the host is configured to forward packets for other hosts often connected to a local network, it is a router. 6to4 does not facilitate interoperation between IPv4-only hosts and IPv6-only hosts. 6to4 is simply a transparent mechanism used as a transport layer between IPv6 nodes.*	*6to4 is a dual stack transition technology which was generalized into 6rd. The main caveats being* *1 - The end device (PC or CPE) needs to support 6to4 and a 6to4 relay needs to be added to the network* *2 - It is a temporary solution due to the fact that the IPv6 addresses are provided using an IPv4 address* *3 - Deployment requires 6to4 border relays to be carrier-grade in terms of performance and features including fragmentation/MTU handling.* *6to4 is mainly used in initial phases to provide IPv6 connectivity but is not meant to be used permanently.*
6in4	*6in4 uses tunneling to encapsulate IPv6 traffic over explicitly configured IPv4 links as defined [i.31]. The 6in4 traffic is sent over the IPv4 Internet inside IPv4 packets whose IP headers have the IP protocol number set to 41. This protocol number is specifically designated for IPv6 encapsulation. In 6in4, the IPv4 packet header is immediately followed by the IPv6 packet being carried. This means that the encapsulation overhead is simply the size of the IPv4 header of 20 bytes. With an Ethernet MTU of 1 500 Bytes, an IPv6 packet of 1 480 Bytes can be transported without fragmentation. 6in4 tunneling is also referred to as proto-41 static because the endpoints are configured statically. Although 6in4 tunnels are generally manually configured, tools like the utility AICCU can configure tunnel parameters automatically after retrieving information from a TIC server.*	*6in4 is a dual stack transition technology. The main caveats being* *1 - The end device (CPE router) needs to support 6in4 and a 6in4 router needs to be added to the network* *2 - The IPv4 tunnel endpoints need to be manually configured which provides an operational overhead* *6in4 is mainly used in enterprise environments to provide IPv6 connectivity but is not meant to be used permanently given the above caveats. 6in4 does not have any valid deployment with its present definition.*
4rd	*The CPE is provisioned with an IPv4 identifier and port set as defined in [i.11]. These are selected based on the requirements for the service and can be an IPv4 prefix (subnet) or a single IPv4 address or an IPv4 address and port set for shared access. The IPv4 address and port set is mapped into a CE index, which is appended to the IPv6 domain prefix to complete a 64-bit prefix. The IPv6 suffix is all zero. The CE router uses this as its 4rd source address and computes a destination address based on the border relay prefix. It forwards the 4rd packets to the border relay. The border relay decapsulates the 4rd addresses and replaces the header with an IPv4 header containing source and destination IPv4 addresses.* *In reply, the border relay generates the IPv6 addresses from the received IPv4 header and forwards the packet back to the CE router. The CE router in turn derives the IPv4 header and forwards the reply to the IPv4 host.*	*4rd is a transition mechanism that tunnels IPv4 traffic across an IPv6 network. 4rd provides a good alternative to DS-Lite as it uses stateless tunneling and there is no requirement for a CGN in the network to perform NAT44. 4rd allows for an IPv4 prefix (i.e. an IPv4 subnet) or a single IPv4 address or a shared IPv4 address with a static port allocation. This technology will not be further considered due to its technical nature and lack of deployment and coverage in the market.*

TRT	*For an IPv6 client to communicate with an IPv4 server, a DNS-ALG is used to perform translation of the A record response into an AAAA record. The DNS-ALG builds an IPv6 address using the IPv4 address contained in the returned A record. [i.29] specifies the use of the prefix C6::/64 followed by 32 zeroes plus the 32-bit IPv4 address. However, IANA has not allocated the C6::/64 prefix. Thus a locally configured prefix would be required. The TRT proxies the request to the IPv4 server using the embedded address for the destination and its own IPv4 address as the source. A session table maintains state. ALGs would be required for some applications to use this method.* *In reply, the TRT gateway replaces the IPv4 header with an IPv6 header based on information held in its session table. The reply is sent back to the originating host.*	*TRT is the translation gateway used to allow IPv6 clients to communicate with IPv4 servers. TRT is a stateful gateway device that relays packets between an IPv6 client and an IPv4 server. The TCP/UDP connection from the client is terminated on the TRT, and the TRT creates a separate connection to the destination server and relays between the two connections. This technology is already clarified / extended within NAT64 and therefore is redundant for further consideration in its own right.*
TSP	*TSP uses XML messages which can be transported using TCP or UDP across either an IPv4 or IPv6 network. It is defined in [i.36]. The XML messages provide SASL authentication for the username and password of the tunnel source as well as tunnel creation and handshaking control. Once the tunnel parameters have been negotiated and the tunnel is established, TSP is used to send keepalive messages between the tunnel endpoints to ensure the tunnel connectivity is maintained.*	*TSP is an experimental protocol used to negotiate the setup of a static IP tunnel between a client and a tunnel server. TSP does not itself provide any IPv4 or IPv6 connectivity, but provides the control mechanisms for tunnel creation, either to transport IPv6 encapsulated in IPv4 or IPv4 encapsulated in IPv6. The technology is not completely defined and not mature enough for further consideration.*
NAPT	*Deprecated*	
NAPT-PT	*Deprecated*	
SIIT	*SIIT is not a standalone transition technology but rather a mechanism that is leveraged in many other technologies.*	
ISATAP	*ISATAP is used for automatic deployment of IPv6 in IPv4 sites. ISATAP typically builds its PRL by consulting the DNS. Hence, it is a lower-layer protocol that relies on a higher layer. A circularity is avoided by relying on an IPv4 DNS server, which does not rely on IPv6 routing being established. However, this approach is a violation of network design principles. ISATAP carries the same security risks as 6over4: the IPv4 virtual link should be delimited carefully at the network edge, so that external IPv4 hosts cannot pretend to be part of the ISATAP link. That is normally done by ensuring that proto-41 cannot pass through the firewall.*	*ISATAP is not a potential technology for IPv6 transition due to its limited deployment and limited vendor support. ISATAP specifies an IPv6-IPv4 compatibility address format as well as a means for site border router discovery. ISATAP also specifies the operation of IPv6 over a specific link layer. The defined link layer is based on IPv4 which bears significant limitations for deployment in today's networks.*
6over4	*6over4 is not a standalone transition technology but rather a method of link-local addressing. In 6over4 any host wishing to participate in 6over4 over a given IPv4 network can set up a virtual IPv6 network interface. The link-local address is determined as follows:* *- It starts with fe80:0000:0000:0000:0000:0000 or fe80:: for short. The lower-order 32 bits to the binary value should be set to the IPv4 address of the host. For example, host 192.0.2.142 would use fe80:0000:0000:0000:0000:0000:c000:028e as its link-local IPv6 address. A shortened notation would be fe80::c000:28e.*	*6over4 is not an end-to-end technology but rather an addressing method that has the potential to be further developed into a transition technology. There are numerous issues with the technology mainly based around the fact that link-local addresses are not routable across the network. This technology cannot be used without further development and deployment.*

AYIYA	*An IPv4 client establishes a tunnel session with the tunnel server using the TIC (or TSP) protocol to pass authentication parameters. The tunnel server uses the authentication to determine the IPv6 address to assign and prefix to route across the tunnel as defined in [i.14]. Once the tunnel is established, all IPv6 packets pass directly through the tunnel being encapsulated / decapsulated by the tunnel client and tunnel server. All communication between the tunnel client and server are via UDP/IPv4 only. The AYIYA header includes a SHA-1 signature which is used to ensure the integrity of the message. Where the AYIYA tunnel client is configured on a router, the LAN behind it becomes dual stack and any connected LAN device has full access to either IPv4 or IPv6. No ALG or translation is required. NAT traversal can also be achieved with no additional configuration as the tunnel uses standard UDP datagrams. The AYIYA protocol has an overhead of 52 Bytes in addition to the IPv4 20-Byte header. This allows an IPv6 MTU of 1 428 Bytes across an Ethernet network.*	*AYIYA is a tunneling protocol that is intended to be used to transport any protocol within a tunnel across a network that uses a different protocol. The most common implementation is AICCU which is used by a tunnel broker to provide IPv6 access to IPv4-only clients on the public Internet using a static tunnel. This is not a standalone technology and needs further development or needs to be used as a basis for another transition technology.*

10.1 DS-Lite

One of the most promising candidate technologies for deployment in the transition from IPv4 to IPv6 in Cable Networks is DS-Lite. As demonstrated by the results of the analysis provided in the present document, the main advantage of DS-Lite over its alternative technologies is the lack of a requirement for inter-communication between IPv4 and IPv6. This technology provides almost 0 % service deprecation comparative to IPv4 native on some vendors.

10.1.1 Technical Summary

DS-Lite enables customers to access services natively over IPv6 and through translation over IPv4 and is, thus, a key technology to ensure complete service continuity when introducing IPv6 into the Access Network. As IPv6 is not "backward compatible" to IPv4, i.e. the two protocols exist independently of each other without any interaction in most topologies, DS-Lite is introduced to allow for a smooth transition towards IPv6 once no more IPv4 addresses are available and a co-habitation of IPv4 and IPv6 on the same network infrastructure is required. DS-Lite also provides the ability for IPv4-only home devices, applications and OSs to continue to access the Internet with minimal use of IPv4 public addresses from the operator's remaining pool. It also avoids dual or multiple layers of address translations which otherwise introduces its own set of problems.

A potential standardization of DS-Lite technology would need to address the following objectives:

- To define the logical and physical parameters allowing customers to access the public Internet across an operator managed IPv6 network using automatic tunneling
- To define the specific features to use

DS-Lite can be summarised as providing the following:

- IPv4 connectivity to IPv4 hosts over home routers (CPE) and Access Networks that are provisioned with IPv6 addresses only
- Dual stack connectivity for hosts connected to IPv6-only Access Networks
- Less need to maintain an IPv4 or dual stack Access Network
- A lightweight solution
- Single NAT - i.e. no need to have multiple layers of NAT
- Avoidance of protocol translation and need for ALG-TSs - sole use of few ALG-Ps
- Sharing of a limited number of public IPv4 addresses among a large number of customers using port translation

- Automatic tunnel establishment (between CPE and AFTR)
- Port forwarding capability on an AFTR using technologies such as Web-UI, NAT-PMP, UPnP, A+P

10.1.2 Technical Requirements Summary

Two network components are involved in the end-to-end DS-Lite communication, the AFTR and the CPE. Considerations for the definition of requirements for the AFTR and the network itself are listed below. A detailed analysis of CPE requirements is following from clause 10.1.9.

10.1.2.1 AFTR

The AFTR is the LSN device that is placed at the edge of the network (IPE connected to the internal MPLS PE/LER) as the IPv6/IPv4 gateway. For egress packets, the AFTR performs de-capsulation from a 4in6 packet to a IPv4 packet while translating ports and address.

Requirement considerations for AFTR deployment are:

- Hardware topology
- Logical topology
- Software/hardware features
- Scalability
- Resilience and redundancy
- IP address allocation and DHCP specific features (v4 and v6)
- Forwarding/convergence performance
- Monitoring, management, reporting and access
- DR specific technical requirements

10.1.2.2 Core Network, Access Aggregation and Internet

For the purposes of the present document, the network is assumed to be composed of a Core Network, Access Aggregation and Internet.

The Core Network is the central BGP or/and MPLS network edge-to-edge to the peering points.

Access Aggregation is the portion of the network between the internal PE and the CMTS IP side. When analyzing requirements for Access Aggregation, consideration of DOCSIS or RF/HFC aspects is not included. Access Aggregation begins at the external facing IP interface of the operator's network towards the Internet.

10.1.3 Technical Definition

DS-Lite uses IPv6-only links between the provider and the customer to carry IPv4 privately addressed packets.
The DS-Lite home gateway (CPE) is provisioned with only an IPv6 address on its WAN interface. At the LAN-side interface, the CPE operates its own DHCP server handing out RFC 1918 private IPv4 addresses [i.23] to home devices. The CPE does not perform NAT; the NAT function is located on a carrier-grade NAT device in the provider's network, which is also a tunnel terminator for the IPv4-in-IPv6 tunnel. This device is called AFTR.

The IPv4 packet from the home device to an external destination is encapsulated in an IPv6 packet by the CPE and handed to the provider network. The packet is decapsulated at the AFTR, and NAT44 is performed to map the host's private IPv4 address to a public IPv4 address. The IPv6 tunnel source address is added to the NAT table, along with an IPv4 source address and port in order to both disambiguate the customer private address and provide the reference for the tunnel endpoint. If a home device needs to access an IPv6 service, packets are transported "as-is" and routed to an Internet server. With DS-Lite technology, the communications between endpoints remain within their address family without requiring protocol family translation.

NAT services allow service providers to conserve IPv4 addresses and maintain IPv4 Internet access while migrating to IPv6. If implemented in components of the Core Network (such as Microsoft's ISA), an optimized software implementation provides scale, improved transaction rates and complete logging and accounting.

NAT services operate in two modes: Network Address and Port Translation (NAPT) optimized to provide scale and a subscriber-aware, Layer 2-aware NAT. In NAPT mode, an operator can deploy centralized devices to provide IPv4 service continuity with minimal changes to the Access Network and host devices. Layer 2-aware NAT enhances NAPT to create a virtual NAT table per subscriber. This allows for customized NAT policies and integrated RADIUS accounting; and it uniquely permits all subscribers to share a common internal IP address to simplify IPv4 address assignment and administration.

Example features of a carrier-grade centralized NAT include:

- NAT44 with high scalability, high transaction rate and N+1 redundancy
- Two modes of operation:
  - NAPT with traditional inside and outside addresses
  - Layer 2-aware NAT allowing for overlap and reuse of subscriber (inside) IP addresses
- Assigned port range blocks for each host to improve scalability and reduce logging requirements
- Per-subscriber port limits with application priority based on traffic forwarding class
- Configurable high/low prioritization of applications traversing the NAT based on traffic forwarding class
- Concurrent support for L2TP network server on the same platform

If deployed as described above, the following benefits for the service provider apply:

- Mitigates service provider IPv4 address exhaustion
- Allows IPv4 services to continue and evolve during the migration to IPv6 services
- Layer 2-aware NAT removes requirement for unique private IP addresses per subscriber
- Allows for unique NAT subscriber policies such as PPP, PPPoE, L2TP LNS and IPoE (DHCP)

B4 / DS-Lite CPE / Home Gateway – 'Basic Bridging BroadBand'-device @ the client site.
- Dual-stack IPv4/IPv6 (IPv4 private and IPv6 public/global addresses).
- 'Normal' IPv6 traffic forwarding and IPv4 traffic encapsulation over an IPv4-in-IPv6 Tunnel.
- Receives LSN/AFTR tunnel-endpoint through DHCPv6 (option not standardized yet).
- Initiates the softwire which terminates at the AFTR.

AFTR / CGN / LSN – Address Family Translation Router as a network element in the ISPs network.
- Terminates IPv4-in-IPv6 Tunnels.
- Does the (private to public and v.v.) NAT44 translations.
- Maps IPv6 src-address, IPv4 src-address & IPv4 src-port to uniquely identify customers

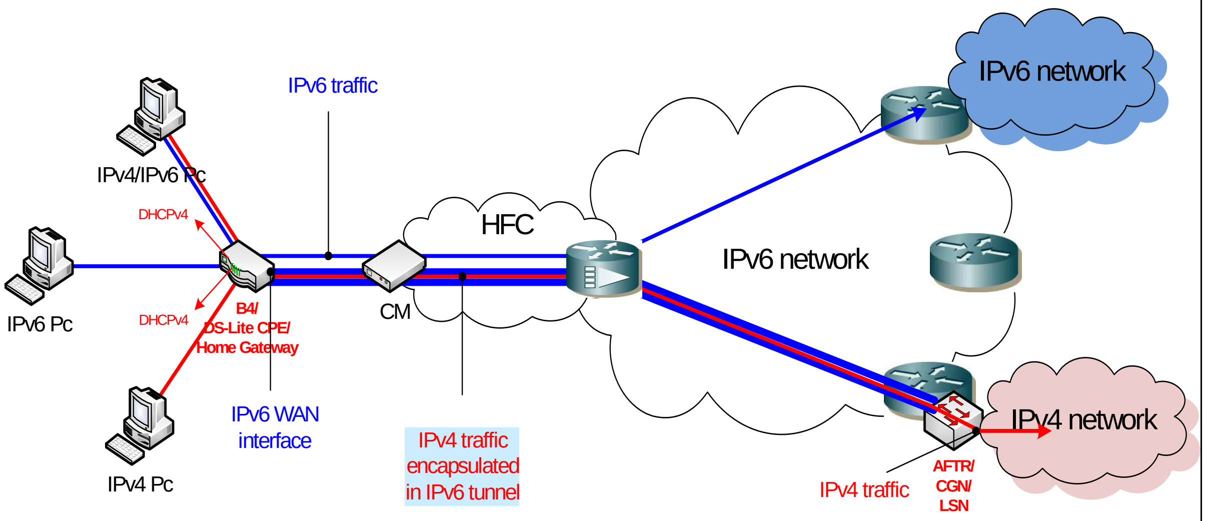

- CPE initiates softwire to AFTR once an IPv4 packet comes from the LAN and is destined to a public IPv4 destination.
- CPE encapsulates IPv4 packet into an IPv6 packet and sends across the tunnel.
- LSN decapsulates IPv6 packet, then does NAT44 from a private IPv4 to a public IPv4 address (each new session logged, 9 fields, for DR) and creates table-entry for uniquely mapped source IPv6 address (tunnel endpoint) & source IPv4 address & tcp/udp-port.

- Public IPv4 src-address/port triggers lookup in LSN to identify customer's IPv6 address (tunnel endpoint) and original private IPv4 address/port.
- IPv4 destination gets rewritten into original IPv4 private source, packet gets encapsulated into IPv6 and is sent across the tunnel.
- CPE decapsulates packet and forwards IPv4 packet to original source/port.

Figure 5: Technical definition of DS-Lite

The following aspects should be taken into account when considering deployment of AFTRs. They have been derived from an analysis of several AFTR implementations available on the market which has been performed in the course of the preparation of the present document:

1) The feature sets for the AFTR are not complete for deployment in certain aspects. Currently, there is active development work going on that will evolve into a fully functional, non-service deprecating AFTR design. Some dependencies are related to CPE development timelines (e.g. implementation of PCP), however this is expected to be resolved by late 2012.

2) Until the completion of the development of a complete, deployable technology, AFTRs are available for testing and field-trials. PCP and RADIUS are the two main features that are required to avoid major deployment risks whereas other features are not crucially required to deploy DS-Lite.

3) Technically, the DS-Lite solution is completely new for use by Cable Operators and, thus, care has to be taken to create awareness of the lack of previous experience. Cooperation of operators, manufacturers, training organizations and field engineers ensure constant exposure during testing and field-trials to local and support staff to allow for enhancement of technical troubleshooting skills.

4) DS-Lite comes with significant advantages one of them being that it is architected to not to cause any service deprecation or tangibly different experience by the customer. An assessment of DS-Lite indicates that the throughput of the AFTR is higher than with some standard routers only transiting IP packets. Typical capacities are well above max requirements even early product phases. DS-Lite dynamically generates tunnels which are torn down if not used. So there is no a priori resource assignment to a customer. When considering the utilization of IP addresses, it is much more efficient to share public addresses with and AFTR than to try to continue the present IPv4 allocation available for IPv4 native customers. This characteristic should be considered in utilization algorithms.

According to industry sources, implementations of DS-Lite are extremely stable. At present, no service effecting systematic bugs have been reported. Its development and maintenance seems to be well established by manufacturers which should affect any risk management assessment in a positive way.

In order to fully assess the availability and stability of required features such as PCP, extensive application testing is required. Table 2 depicts an example for the extent of such testing. A partial pass means that the application could be verified with some functionality but some components were missing that would not work without PCP or needed static configuration which can generally not expected to be performed by a normal subscriber.

Table 2: DS-Lite application validation

Seq#	Application/protocol	Pass/fail	Max Port binding	Avg Port binding
7	UDP	pass		
2	DNS	pass		
4	ICMP	pass	1	0
16	NTP	pass	2	1
26	POP3	pass	4	0
48	SMTP	pass	2	0
	FTP ^{TM}	pass		
	FTP ^{TM} passive	pass	2	2
66	FTP ^{TM} active	pass	3	2 (+ 1 alg)
15	Flash Video	pass		
81	IMAP	pass	2	1
56	STUN	pass		
93	Myspace ^{TM}	pass		
23	Windows Live ^{TM}	pass		
13	Facebook ^{TM}	pass	31	20
18	YouTube ^{TM} Web	pass		
39	Twitter ^{TM}	pass		
53	Mega Upload	partial pass		
70	Flickr ^{TM}	pass		
72	Dailymotion ^{TM}	pass	31	15
87	Rapidshare	partial pass		
	Hotfile	partial pass		
122	Rutube	pass		
125	Justin.tv	pass		
131	Metacafe ^{TM}	pass		
132	YouTube ^{TM} HD	pass	31	18
149	Yahoo ^{TM} Video	pass		
156	Baidu Hi ^{TM}	pass		
159	Zshare	pass		
165	Bliptv	pass		
170	Break Videos	pass		
176	Pandora Tv	pass		
205	Deezer ^{TM}	pass		
318	Baidu ^{TM}	pass		
344	Godtube	pass		
14	Shockwave ^{TM} Flash	pass		
24	Google ^{TM} Video	pass		
98	IPsec	pass		1
95	ISAKMP	pass		

Seq#	Application/protocol	Pass/fail	Max Port binding	Avg Port binding
111	SSH Daemon	partial pass		
111	SSH	pass		
6	SSL	pass		
28	Gmail TM	pass	12	
30	Hotmail TM	pass		
88	Yahoo TM Mail	pass		
10	µtorrent	partial pass		40
89	Vuze (bitttorent)	partial pass	> 100	40
34	eDonkey	partial pass	40	10
12	MSN TM Messenger	pass	68	22
12	MSN TM Messenger File Transfer	pass		
	MSN TM PC-toPC Voice	pass	10	6
124	MSN TM Messenger TM Webcam	pass	22	13-16
350	MSN TM Messenger TM Remote Assistance	pass		
68	Google TM Talk	pass	2	1
71	Jabber TM	pass		
202	Google TM Talk Data	pass	5	3
180	Google TM Talk Voice	pass	4	3
92	IRC File Transfer (client)	pass		2
92	IRC	pass		1
92	IRC File Transfer (server)	pass		2
440	Skype TM Phone-to-PC	pass	16	
441	Skype TM PC-to-PC AV Chat	pass	10	10
248	Skype TM Generic	pass	6	2
21	Skype TM PC-to-PC	pass	17	9
44	PC: Valve's Steam Service	pass	17	5
90	PC: World Of Warcraft TM	pass		
109	PC: Counter-Strike TM	pass	23	3
25	Windows TM Update	pass	4	2
	WindowsTM Activation	pass	3	3
49	Avg Update	pass	2	1
123	RDP	pass		
297	Nintendo Wii TM Web Browsing	pass	4	2
324	Nintendo Wii TM Control	pass	5	1
78	Nintendo Wii TM Data	pass	5	1
	PlayStationTM	pass	39	16

10.1.4 AFTR Hardware Features/Topology

- Role / Location
  In some cases the AFTR performs the role of a 6PE. The AFTR should be placed at the network edge as close to the external peering points as possible. In case the network realm has multiple exit points, a balanced path should be implemented to all exit points (AFTR nodes will only be added based on capacity requirements).

  In case the AFTR does not perform the role of a 6PE, the AFTR is placed at the network edge as close to the internal PEs as possible.

- Node Type
  The type of the node should be chassis-based to guarantee later scalability and hot-swap capability. Features and functionality should be distributed on blade or port level.

- Physical Ports / Traffic Balance
  Physical ports should have sufficient capacity to sustain the expected traffic requirements. Ports capable of 10 Gbit/s are recommended. The actual design may depend on aggregation of node capacity.

- Memory
  Memory capacity should be sufficient to hold the full IPv6 and IPv4 BGP routing table; minimum requirements are based on the operator's PE requirements.

- Integrated Forwarding
  The ICXF and AFTR mechanisms should be integrated on blade or ingress port.

- Forwarding Architecture
  Forwarding should be implemented in hardware with a minimum number of 4 million sessions per blade. Node latency is expected to be below 1 ms.

10.1.5 AFTR DS-Lite Specific Software Requirements

- Tunnel Identifiers / Client-Customer ID
  For data retention purposes, tunnel identifiers should be uniquely associated with a single CPE.

- DS-Lite Timers
  Timers have to be configurable. Default values are defined in [i.46].

- Softwire Initialization (DS-Lite Timers Extended)
  Quick drop quick pickup approach is preferred.

- Port Block Allocation per IP Address
  Port block allocation is configurable to allow for any ratio assignment.

- Static/Dynamic Port Allocation (STDPA) Deterministic NAT
  DS-Lite requires efficiency and scalability not only with regard to throughput and node latency but also in terms of IP address ratios and customer port allocation (CPA - the amount of ports within the block assigned to a specific B4 address (CBSA)). This also requires the ability to configure a "static" block allocation that is assigned to a customer for a time period ranging from 12 to 172 hours. Dynamic allocations are required as well and work in conjunction with the static allocation of a port block (possibly 200 ports per block as an example). This initial block will be static staying the same for the customer no matter the circumstance as long as his Ipv6 public CPE address does not change for a fixed amount of time. Then any further blocks he might request over that 200 initial port block would be dynamic and end on TCP FIN, etc.

- NDP / ECMP Integration
  Due to the nature of an AFTR load balancing on links is extremely important. ECMP should be seen as a requirement and be integrated into NDP to allow for effecting events and ECMP use of best path matches.

- TCP MSS
  TCP MSS support is mandatory for the AFTR due to the removal of an end-to-end MTU sizing functionality. This will avoid the need for excessive fragmentation.

- Fragmentation
  Fragmentation should be placed on the ASIC running in the linecard and should be pre-tunnel SI. Fragmentation of IPv6 packets under certain requirements according to [i.46] should be supported.

- Load Balancing
  Load balancing across aggregated interfaces should be supported. This includes IP address and port allocation, ECMP, PVST load balancing, LDP hashing, etc.

- IPv4 Private Subnet Segmentation
  The AFTR should be able to segment IP address blocks into smaller blocks for the local interfaces.

- Non-ALG Deployment
  The AFTR should have a non-ALG approach for decapsulation where possible. However, ALGs are a requirement with FTP, RTSP and SIP.

- Traffic Prioritization
  Traffic prioritization should be possible to e.g. support B2B customers running video or voice to the SI.

- DS-Lite IPFIX/Netflow
  Netflow v9 or IPFIX should be a common template allowing for both IPv4 and IPv6 to be demonstrated in a single entry. Note that 1:1 ratios should be applicable and both a DS-Lite and standard traffic flow template should be configurable.

- Clustering
  AFTR clustering allows redundancy between two AFTRs and should optimize capacity for NAT cache and configuration matching. Clustering should use the maximum efficiency for IP address allocation.

- Shared Resource (Single AFTR GW Address)
  The AFTR should be a shared resource as much as possible with a single AFTR address. Backup AFTR addresses can be configured but one node should only use one address at a time.

- Physical Redundancy
  Physical redundancy in the service NPUs is required allowing for stand-by or active allocation during failure of a linecard or NPU.

- Thresholds / Watermarks
  Thresholds should be aligned to resources and be present in watermarks leading to events and changes in resource allocation.

- Relative Buffering
  Buffering should allow for multiple entry points into single or multiple NPU for NAT entries, packet buffers and logging.

- AFTR Address Withdrawal
  The AFTR should have at least five points of AFTR GW address withdrawal occurrence. The list includes loss of route-out, loss of all BGP/IGP sessions, loss of forwarding, loss of NPU capacity and certain errors in the NAT caching. Any of the failures should be detectable based on configurable timers with 15 seconds being the default setting.

- Transitional Co-existence
  DS-Lite should be able to co-exist and share resources with other transitional methodologies such as NAT44.

- IPv4 Private SI ID for Block Resource Allocation
  At times, there will be more than five Internet users behind a CPE. This means that standard port block allocation may not be enough. To mitigate this risk it is appropriate to establish a threshold of IPv4 private addresses associated with the B4 as a feature requirement. This will allow extra blocks to be dynamically allocated to heavy workstation B4 setups enabling higher levels of port allocation.

- SI Monitoring
  In order to supervise SI activity either for a return-path connection or for PCP, the tunnel establishment should be monitored and have an event.

- Anycast
  Anycast AFTR gateway addresses are a requirement to allow simplicity of deployment for a single address across multiple AFTRs.

- Data Retention (DR) RADIUS
  For DR purposes two potential methodologies exist, NAT caching with RADIUS or Netflow. The requirement is to hold the following parameters for each flow:
  - Source IPv4 address
  - Source IPv6 address
  - Source port (internal IPv4)

  - Remote destination port (external IPv4)
  - Time stamp
  - Remote IPv4 address
  - Remote IPv6 address (this will be the AFTR address) (optional)
  - Incoming interface (optional)
  - Outgoing interface (optional)
  - Client ID (if different from the IPv6 address).

- UPnP/PMP/Port forwarding/PCP
  The AFTR has to allow UPnP/PMP forwarding and port mapping through the tunnel.

The AFTR software requirements depend on the deployment topology. In most cases, two topologies are required. In the "integrated topology", the AFTR functions as a full MPLS 6PE router. Alternatively, the AFTR functions as a L3 router, hairpinning connections through an external 6PE router. This topology is called "hairpin topology".

10.1.6 Integrated Topology/Transport Requirements

The following list of features has to be supported by the DS-Lite architecture for both IPv4 and IPv6 / L2 and L3 in an integrated topology solution. Features marked with an asterisk (*) are topology dependent. Generally, all forwarding functionality should be implemented in hardware to prevent performance issues. However, there are some features where software based forwarding can compete with the equivalent hardware architecture.

- MP-BGP (including 6PE and 6VPE)
- BGP community / 32 bit AS
- MPLS LDP (currently only IPv4 is supported natively but the requirement for native IPv6 in MPLS is becoming indispensable)
- ECMP
- QoS (IPv4/IPv6) - classification, priority queuing, etc.
- QPPB/SCU/DCU
- SNMP v1/v2/v3 (transport over IPv4 and IPv6)
- ACLs/prefix lists/filtering (both IPv4/IPv6)
- TACACS/RADIUS (IPv4/IPv6)
- SYSLOG (event reporting for IPv4 and IPv6 as well as transport over both protocols)
- CoPP (IPv4/IPv6)
- Netflow v9 (potentially also previous versions will be required depending on the state of the implementation of Netflow)
- XML* (IPv4 and IPv6 reporting and transport)
- MAC accounting
- IEEE 802.1Q [i.52]
- EtherChannel
- Ethernet OAM
- NSF/GR (IPv4/IPv6)

- Policy Based Routing (IPv4/IPv6)
- IS-IS (potentially MT* for IS-IS as well if MPLS IPv6 LDP allows for dual stack) (IPv4/IPv6)
- Static Routing (IPv4/IPv6)
- OSPFv2/v3
- CDP/LLDP (IPv4/IPv6)
- VRRP/HSRP (IPv4/IPv6)
- VLAN mapping/Double Tagging
- L3 multicasting/MFIB (IPv4/IPv6)
- IPv6 forwarding (hardware)
- IPv4 forwarding (hardware)
- Ethernet technologies
- Virtual interfaces (IPv4/IPv6)
- AAA* (IPv4/IPv6) (extended beyond RADIUS)
- BFD (IPv4/IPv6)
- MLD/L2 multicast
- Full NDP (ICMPv6, DAD, NUD...)
- PIM/IGMPv2/v3
- CEF/dCEF
- Anycast
- Route reflection (IPv4/IPv6)
- Standard IPv4 VPN
- ISSU/SSO technologies
- NTP (IPv4/IPv6)
- SEND*
- IPsec
- DNS (IPv4/IPv6 server and client)
- DHCP relay (IPv4/IPv6)
- Graphical traffic and threshold monitor
- NMS
- Jump-off services / Terminal console services

10.1.7 Feature Development Requirements

Table 3 presents features that have been identified to be in development aiming at enhancement of the AFTR capabilities. They are generalized and more analysis needs to be done to validate their scope and their necessity for deployment.

Table 3: AFTR features under development

FEATURE	DESCRIPTION
IPv6 Filtering	IPv6 filtering has been released for testing.
PCP	PCP is in development and a first release is expected in June. Further development will be required.
AFTR Clustering	AFTR clustering allows for the synchronization of the NAT cache states/entries across the network with limited additional network traffic. First release is expected by end 2012.
IPv4 Private SI-ID	The SI-ID is based on the IPv6 B4 address. It will eventually change to using both the IPv6 source address and the IPv4 private source address. There are issues with this feature due to performance changes and resource utilization. If available, port allocation can be carried out per host device and not per CPE.
DS-Lite Fragmentation	With this feature IPv6 packets can be fragmented even if the encapsulated IPv4 packet has the DF bit set. This also requires re-assembly in the upstream.
Static and Dynamic Port Block Allocation	With this feature it will be possible to configure the mix of static and dynamic port allocations. Standard DS-Lite allocates ports based on dynamic tunnel creation and dissipates allocations of the AFTR NAT cache once the last session drops on timeout or finalization. This creates issues with logging and processing if the hardware cannot forward and register a DS-Lite flow. In order to mitigate this limitation, the first port block allocation (or more) for a defined amount of customers (possibly all) is static for a period of time. It can cause IPv4 public address efficiency issues but allows to balance dynamic and static allocations.
RADIUS	RADIUS is tested for DS-Lite DR requirements. However, it will also serve its purpose for monitoring. RADIUS only effects performance by about 5 % for transit node latency. It is, therefore, the preferred method where legally available.

10.1.8 Performance Requirements and Comparison

Table 4: Downstream throughput example measurement for an MSO

Country	A	B	C	D	E	F	G	H	I	J
Hub										
Site1	27 036	296	89	42	1 841	473	136	88	18	34
Site2	37 645	364	97	42	1 982	576	140	143	23	33
Site3	17 235	239	50	34	1 176	303	113	65	11	17
Site4	10 918	100	31	13	757	210	49	33	4	17
Site5	2 066	231	65	31	1 394	452	119	98	5	22
Site6	32 435	337	115	66	2 150	526	144	117	24	32
Site7	13 740	117	41	18	860	264	49	50	8	8
Total (Mbit/s)	141 075	1 683	488	247	10 160	2 805	749	594	93	162

Table 5: Upstream throughput example measurement for an MSO

Country	A	B	C	D	E	F	G	H	I	J
Hub										
Site1	2 006	29	5	7	53	86	6	6	10	6
Site2	1 408	16	8	6	42	53	5	4	8	3
Site3	1 411	15	7	8	55	75	5	6	9	3
Site4	194	2	1	0	6	16	1	1	1	1
Site5	1 780	20	9	6	52	81	9	5	6	6
Site6	1 894	21	19	5	55	80	9	13	10	6
Site7	1 431	11	12	6	41	69	11	6	8	12
Total (Mbit/s)	10 124	114	61	38	303	460	46	41	51	37

Tables 4 and 5 show an example of a national deployment for an MSO. The performance figures are based on expected traffic within any given Cable ISP.

The AFTR performance is evaluated against lowest acceptable benchmarks and the actual delivery of throughput, convergence, failover and latency for all aspects and features within the chassis. All performance requirements are based on peak capacity and average throughput which are considered for the capacity the platform is designed for. Capacity estimates are determined by expected growth of subscriber numbers and increase of throughput per subscriber.

- Throughput interfaces
  All interfaces are required to operate at line rate with 10 Gbit/s. The interface should be compatible with the operator's PE router connectivity.

- Node latency
  The required maximum node latency is 100 µs.

- Flow throughput
  The flow throughput is defined by three main performance figures:

  - CPE initialization
    This is the initialization of the port allocation per subscriber (i.e. the ports that are allocated when a CPE comes up for the first time) with a single external IP address per CPE.

  - Primary flow initialization
    Primary flow initialization is performed after the CPE has already been granted a port allocation. The flow is established as a "new" flow in the NAT cache. New flows are characterized by the lack of any entry except a source IPv6 address already in the cache. The whole flow needs to be allocated and registered into the NAT cache. The requirement is to be able to handle 800 000 flows per 40 Gbit/s chassis throughput capacity.

  - Secondary flow initialization
    The requirement is 1 million flows per 40 Gbit/s chassis throughput capacity.

- Convergence
  Convergence of routing and link failure should be well within 10 ms.

10.1.9 DS-Lite CPE Requirements

DS-Lite is used to provide end-to-end IPv4 connectivity across an IPv6 network. The CPE Router is enabled to establish IPv4 connectivity to remote servers even if the access network is IPv6 only; and the CPE Router is not provisioned with a public IPv4 address. Since the encapsulation functions are performed in the CPE Router, DS-Lite is agnostic to the type of connectivity in the Home Network. All LAN functionality within the customer premise can be IPv4-only, IPv6-only or dual-stack (IPv4 and IPv6).

DS-Lite is not involved in the end-to-end communication between two IPv6 nodes. These packets pass directly as native IPv6. All communication within the Home Network remains as native connectivity, IPv4 to IPv4 or IPv6 to IPv6.

In order to enable IPv4 connectivity for a host, the CPE Router encapsulates IPv4 packets it receives from its LAN interface into IPv6 packets that are forwarded to the AFTR. Thus, IPv4-in-IPv6 tunnels are established between the B4 function in the CPE Router and the AFTR. The private IPv4 addresses are translated in the AFTR to public IPv4 addresses in order to route the packet through the IPv4 Internet to its final destination.

When configured to use DS-Lite, all native IPv4 functionality on the CPE Router should be disabled. The use of private addressing [i.23] should be enforced on the Home Network as all IPv4 packets with a public IPv4 address will be dropped by the AFTR as a security measure.

No changes are required on the Home Network. The IPv4 network can be addressed using private addresses, and the CPE Router is the default gateway.

The Access Network side of the CPE Router, however, does not require an IPv4 address either private or public. All communication via the Access Network is performed over IPv6.

The CPE Router also does not require any Network Address Translation (NAT) functionality as this is moved onto the AFTR.

Due to the increased size of the IPv6 header caused by the encapsulated IPv4 packet, it is expected that the overall packet will become oversized and, therefore, requires fragmentation. The B4 function of the CPE Router should perform fragmentation and reassembly on the encapsulated IPv6 packet. It should not fragment the inner IPv4 packet.

The encapsulation into IPv6 exists only between the CPE Router (B4) and the AFTR. The communicating hosts have no awareness of the tunnel; so any fragmentation and reassembly should be managed by the CPE Router and AFTR.

When the CPE Router (B4) receives an IPv4 packet with the DF bit set which would require fragmentation due to its size alone, it should drop the packet and send an ICMP "packet too big" message back to the host with the recommended MTU size set to 1 460 Bytes. The reduction of 90 Bytes compared to the commonly used MTU size of 1 550 Bytes accommodates for the required encapsulation signaling in the surrounding IPv6 packet.

If the DF bit is not set, the entire packet will be encapsulated with the IPv6 header; and the IPv6 packet is fragmented to fit the MTU of the path between the B4 and AFTR.

MSS clamping is also required at the B4 such that it notifies the hosts of the maximum IPv4 TCP segment size that can be sent.

As the customer-facing IP addressing and routing is independent of the device management, either IPv4 or IPv6 can be used for the management of the CPE Router particularly if it is embedded into a Cable Modem. Changes to the OSS systems are therefore not mandatory.

Figure 6 depicts the addressing scheme that is used by DS-Lite in the Access Network.

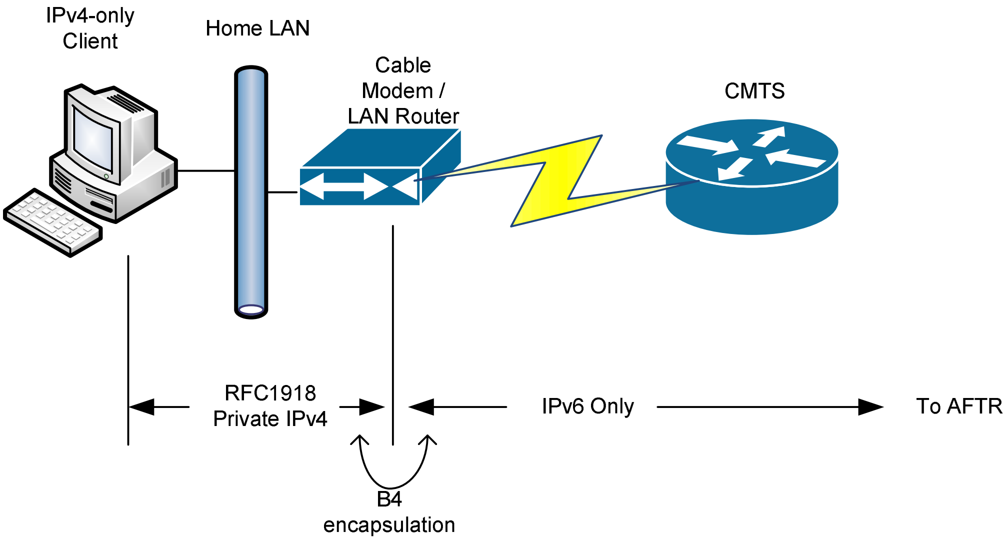

Figure 6: DS-Lite addressing schemes in Access Network

10.1.9.1 Feature Synopsis for DS-Lite in the Access Network

Requirements on the host are determined by the available connectivity in the Home Network. The Home Network may be dual stack or IPv4 only or IPv6 only. In the case that IPv4 is used, the Home Network has to support private addressing [i.23].

The following requirements apply to the CPE Router:

- WAN interface facing the Access Network
  - Request IPv6 global address via DHCPv6
  - Request B4 IPv6 prefix via DHCPv6-PD
- LAN interface facing the Home Network
  - IPv4 private addressing
  - DHCPv4 server for LAN addressing (stateful DHCP service may be included)
  - IPv4 DNS Proxy
  - IPv4 MTU set to 1 460 Bytes

  - TCP MSS clamping to 1 420 Bytes

- B4 router
  - Receive AFTR IP address via DHCPv6
  - Use DNS IPv6 server received via DHCPv6
  - Encapsulate LAN IPv4 packets in IPv6 header
  - Decapsulate IPv6 packets for LAN IPv4
  - Fragment encapsulated IPv6 packets
  - Reassemble received IPv6 fragments
  - PCP
  - Implement the well-known B4 IPv4 address

Assuming that the CPE Router is embedded in a Cable Modem, additional requirements apply:

- Should support bridging of IPv6 packets
- May support IPv4 or IPv6 management
- If IPv6 is implemented, NDP [i.34] and SLAAC [i.35] have to be supported as well
- If the Cable Modem is implementing DOCSIS 2.0 and IPv6 it has to comply to:
  - DOCSIS 2.0+IPv6 Cable Modem Specification [i.49]
  - IPv6 management as defined in the OSSIv3.0 Specification [i.50]
  - Provisioning Mode Override as defined in [i.50] and [i.49]
- If the Cable Modem is implementing DOCSIS 3.0 it has to comply to:
  - MAC and Upper Layer Protocols Interface Specification [i.5]
  - Upstream Drop Classifiers
  - OSSIv3.0 Specification [i.50]
  - Provisioning Mode Override as defined in [i.50]
- If the embedded device implements eRouter it has to support:
  - IPv6 provisioning of CPE devices as specified in [i.51]
  - IPv6 address assignment as specified in [i.51]
  - Identity of DHCPv4 client identifier and DHCPv6 DUID as defined in [i.32]
  - TR-069 [i.48] via native IPv6 transport to allow remote management of WiFi and router parameters

In addition to the requirements listed above, the CPE Router should support the following features on its interfaces facing the Home Network (wired, wireless):

- SLAAC
- DNS resolver information as defined in [i.40]
- Stateless DHCPv6 (INFORM) server or, potentially, stateful DHCPv6 server (IA_NA) with the option to switch from SLAAC
- IPv4 NAT/NAPT

- IPv4 static NAT
- IPv4 inbound port forwarding
- IPv4 stateful packet firewall (5-tuple filters), enabled by default
- IPv6 stateful packet firewall (5-tuple filters), enabled by default
- Recursive DNS server option as defined in [i.40]
- Requirements for IPv6 customer edge routers as described in [i.45] and [i.13]

For purposes of a unified user experience some requirements apply to the user interface:

- Web-UI should be accessible on the LAN IP interface
- Login is initially presented in the format of requesting username and password
- Localization of UI for different languages should be supported
- A Status tab should contain subpages on Software, Connection, Security, Diagnostics
- A Router Basic tab should contain subpages on WAN Setup, LAN&DHCP Server, Backup
- A Router Advanced tab should contain subpages on Option, IP Filtering, MAC Filtering, Port Filtering, Forwarding, Port Triggers, DMZ Host
- A Firewall tab should contain subpages on Web Filter, Local Log, Remote Log
- A Parental Control tab should contain subpages on User Setup, Basic Setup, Content Filter, ToD Filter
- A Wireless tab should contain subpages on Radio, Security, Advanced, Access Control
- An MTA tab should contain subpages on Status, DHCP, QoS, Provisioning, Event Log

10.1.9.2 Feature Development Requirements

Additional functionality to allow UPnP NAT Traversal and static port allocations to function correctly require the use of the Port Control Protocol (PCP). Additional development on the CPE Router is required to allow the customer to configure these features or to pass UPnP/PMP messages from the hosts to the AFTR.

10.1.9.3 Software and Hardware Requirements

All DS-Lite packets that traverse the CMTS are IPv6. The CMTS should therefore support native IPv6 connectivity on the Access Network. There are no additional specific requirements on the CMTS to support DS-Lite.

DS-Lite requires additional processing on the CPE Router to encapsulate / decapsulate the IPv6 packets. To ensure that maximum throughput is achievable on the device, the encapsulation / decapsulation should be done in hardware on the router. It should not be application-based.

The Home Network can continue to operate as IPv4-only, however it is recommended that dual stack is enabled. IPv6 hosts should be provisioned with global addresses in order to avoid issues with double-encapsulation for simultaneously existing Teredo hosts.

Where the Home Network is IPv4-only, any host with Teredo enabled will attempt to establish a connection over UDP/IPv4 to a Teredo server. This will cause the CPE Router to encapsulate the IPv4 packets into IPv6 to forward to the AFTR.

10.1.9.4 Performance Requirements and Comparison

The DS-Lite B4 functionality should not introduce any degradation in performance from a native IPv4 connection. There is a significant amount of additional processing required on the CPE to encapsulate and decapsulate the IPv4 packets. This will introduce additional latency in the end-to-end communication.

The B4 functionality should be implemented in hardware on the CPE Router to ensure the maximum performance can be achieved. This will minimise the latency, and improve throughput.

Additional consideration should be given to the provisioning of DNS. If a public IPv4 DNS resolver is used, then the B4 will be required to encapsulate every DNS request from an IPv4 client. This will have significant performance overhead on the CPE. It is essential that the CPE Router is configured as a DNS proxy for all LAN IPv4 clients, and forward all queries to an external resolver using native IPv6.

DS-Lite requires the specific B4 functionality to be available on the CPE Router.

Where an eRouter is used, this may require the replacement of existing devices with hardware that is capable of operating the B4 function in hardware.

Where a separate CPE Router is used, the additional B4 functionality does not affect the Cable Modem. However, it may also require the replacement of the router hardware.

The B4 functionality will also require configuration to enter the AFTR address. This should be set via the option defined in RFC 6334 [i.53].

Where IPv6 has been deployed, there are no additional changes required, and therefore no cost for the CMTS.

10.1.10 Development and Deployment Status

DS-Lite has been issued as standards-track by the IETF [i.46]. There are several CPE Router vendors that are able to provide DS-Lite, both in hardware or application-based implementations. The B4 router implementation is also available as open source.

DS-Lite is well advanced in development compared to other technologies. Several AFTR solutions are available on the market. Therefore, the standardisation of the AFTR becomes increasingly valuable to ensure interoperability.

10.1.11 Failings/Issues of DS-Lite

As with all NAT solutions, DS-Lite has restrictions on inbound connections. DS-Lite also has a restricted payload due to the increased IP header used for IPv6. Fragmentation is more likely to occur.

As the NAT functionality is removed from the CPE, existing features such as UPnP NAT traversal and port forwarding can no longer be used to direct inbound packets on predetermined IPv4 TCP or UDP ports to a specific host.

Application Level Gateways on the CPE cannot be used when DS-Lite is employed. The AFTR should provide the additional port mapping requirements for each required ALG.

Other issues are mostly well-known and present topics for further development.

- The cost, particularly of the CPE router, is expected to be increased due to the extension of the required functionality. MSOs that are not deploying WiFi solutions and stick to bridging CPE devices will face the issue of having to exchange CPE just for the reason of introducing DS-Lite.

- DS-Lite requires ALGs (ALPs to be exact in most cases) which require the AFTR to do some form of intelligent vicissitudes to the transit packets bearing the risk of drop in performance and limits in functionality. The risk of service-deprecation is minimized if all functionality is included natively in the AFTR.

- With DS-Lite, all functional intelligence is located in the AFTR. Thus, functionality that requires a public address in the local network has to be performed in the AFTR itself. An example of this is PCP. In DS-Lite, it is placed on the AFTR due to UPnP 1 and 2 requirements.

- To predict scaling requirements comparative to IPv6 utilization can be misleading. This is not specific to DS-Lite but occurs similarly in other transition technologies. The issue is that it is largely unknown what portion of the traffic originated by a certain client will be IPv6. Thus, the system has to be scaled for a constant maximum of all clients and designs cannot adapt based on exact traffic predictions.

- DS-Lite is a tunnel technology and, thus, suffers from MTU requirements beyond the norm. This means that MSS clamping, IPv4 fragmentation and fragmentation according to [i.46] are required in order to avoid ICMP blocking. Fragmentation resends and general PMTU control can have performance effects on customer services. Therefore, there is a major requirement to optimize implementations such that they do not cause a large impact on current RTTs and node latency.

- The requirement for Data Retention capacity is rather large and, thus, making the capacity available may become a major issue for the receiving logging server. Normally, on public addresses there is a single entry or two per week depending on lease times and utilization. Within DS-Lite, each time an individual subscriber receives a port and IP address assignment the event should be logged. This can amount to 200 assignments a day with a start and stop time resulting in 400 logs per user per day. Assuming that an MSO could have 10 % of its customer base online at one time, that would mean a huge amount of logs.

10.1.12 Summary Assessment

DS-Lite is extremely usable as a transition technology particularly if the cost for exchanging the CPE can be absorbed by churn or other projects that require upgrading the CPE. The AFTR itself, if deployed centrally, does not represent a major investment. Example figures indicate the expected range to be around 200 000 EUR per 1 million customers, which translates into a cost of 20 cents per subscriber.

Technically, the solution has been continuously improved in recent years with many of the known issues resolved and general performance improved. Current implementations are expected to provide 40 µs node latency on average and still less than 1 ms with logging, fragmentation and ALGs.

With the cost issue resolved, DS-Lite represents a fully functional solution. It now has less than 0,2 % service deprecation with the present PCP implementation and will be almost 0,1 % at the end of the year in some products.

10.2 NAT64

10.2.1 Technical Summary

Global IPv4 address space is currently projected to be depleted around the middle of 2012 to 2015 for most MSOs. As part of the resulting rollout of IPv6 in worldwide networks, specific measures should be taken to allow a smooth transition and coexistence between IPv4 and IPv6.

NAT64 is one of the technologies the current report recommends to consider for standardisation mainly due to the lack of a need for inter-communication between IPv4 and IPv6. This technology will allow customers to access services natively over IPv6 and through translation over IPv4.

In order to enable connectivity between IPv6 hosts and the Internet, NAT64/DNS64 presents always an IPv6 address to the host independently if communication is to be established with an IPv6 or IPv4 addressable device. Communication to an IPv4 device is enabled by synthesizing the DNS A record into a AAAA record (DNS64) and by IPv6 to IPv4 address translation via a NAT64 device. As such, the technology is dependent on DNS and requires devices in the home to be natively IPv6 capable. IPv4-only devices and non-DNS based applications will not work in this environment.

A potential standardization of NAT64 technology would need to address the following objectives:

- To define the logical and physical parameters allowing customers to access the public Internet across an IPv6 network using automatic tunneling.

- To define the specific features to use.

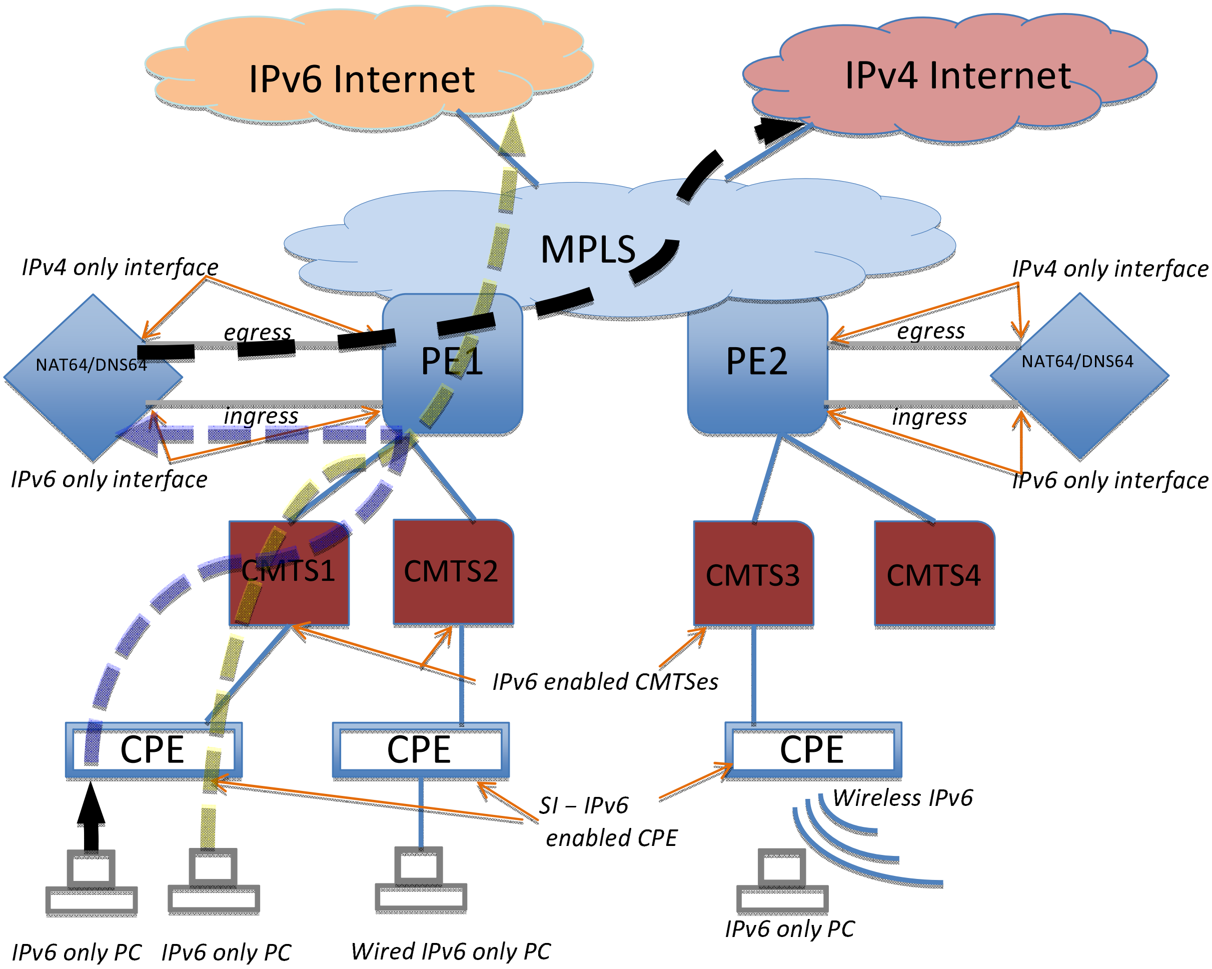

Figure 7: Packet flow with NAT64

NAT64 can be summarized to provide the following features:

- Connectivity of IPv6 hosts to IPv4 or IPv6 Internet via home routers (CPE) and an Access Network that is provisioned with only IPv6 addresses
- No need to maintain IPv4 or dual stack Access Networks
- A lightweight solution
- Single NAT - i.e. no need to have multiple layers of NAT
- Avoidance of protocol translation and need for ALG-TSs - sole use of a few ALG-Ps
- Sharing of a limited number of public IPv4 addresses among a large number of customers using port translation
- Automatic communication establishment to NAT64 endpoints via DNS64
- Port forwarding capability using technologies such as: Web-UI, NAT-PMP, UPnP

NAT64 allows a client in the IPv6 domain to initiate communication with a server in the IPv4 domain by translating source IPv6 address and port to IPv4 address and port. It works in conjunction with a modified DNS known as DNS64. NAT64 relies on DNS64 to provide an AAAA record (corresponding to the server in the IPv4 domain) to IPv6-only hosts initiating communication with the IPv4 server. The AAAA record is created from the A record for the IPv4 address. The IPv4 address is mapped to an IPv6 address prepending a well-known IPv6 prefix assigned to the NAT64 gateway. NAT64 manages a pool of public IPv4 addresses and performs a NAPT function by translating IPv6 source address and port to IPv4 source address and port. This is shown in Figure 8.

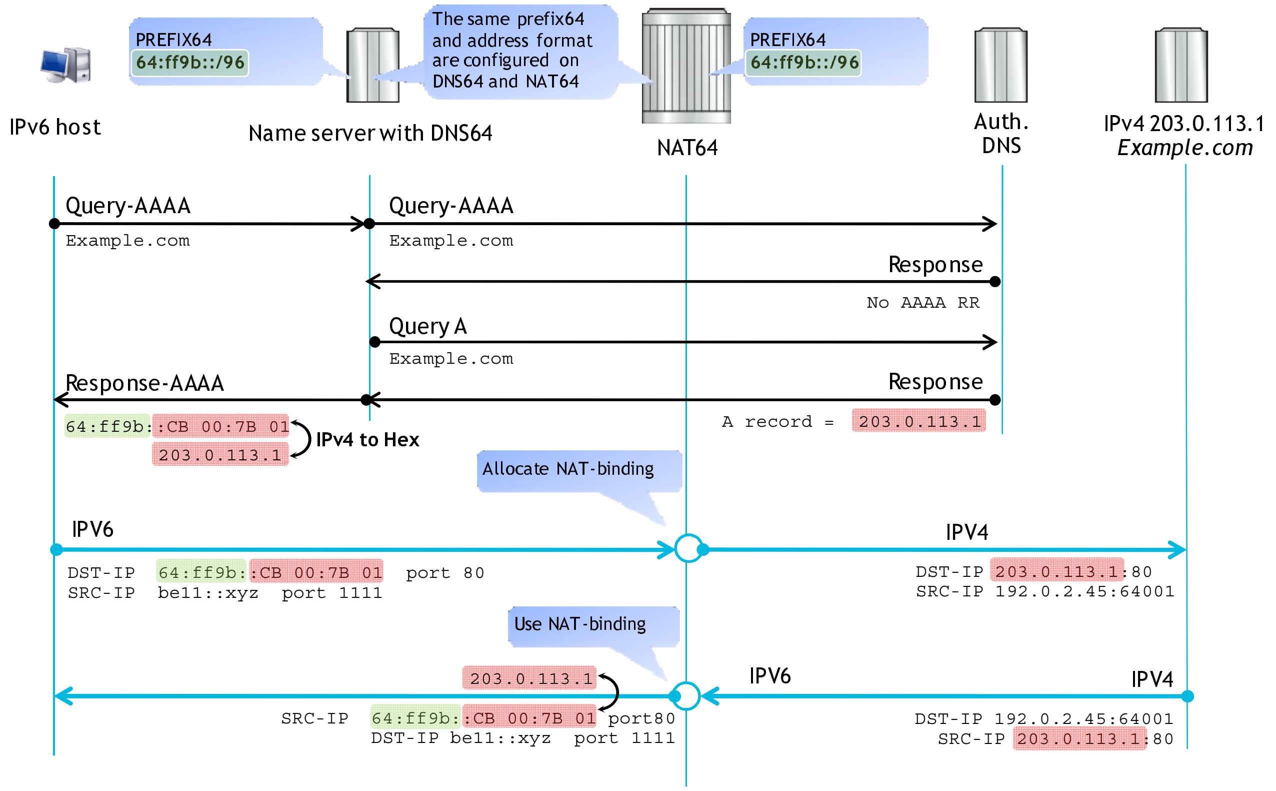

Figure 8: Addressing scheme in NAT64

For TCP and UDP flows, NAT64 maintains mapping between the IPv6 transport address and port and the IPv4 transport address and port and performs header translations. For ICMP, stateful NAT64 needs to maintain mapping between the IPv6 transport address and ICMPv6-identifier and the IPv4 transport address and ICMPv4-identifier.

The NAT64 prefix can be:

- By default the well-known prefix 64:ff9b::/96 (with fixed prefix length of 96 bits).
  This is best practice.

- A network specific prefix.
  The addressing scheme defined for NAT64 [i.39] allows subnet lengths for the NAT64 prefix to be 32, 40, 48, 56, 64 or 96 bits.

Depending on the prefix length, the IPv6 address with embedded IPv4 address is formatted according to Table 6.

Table 6: Embedding IPv4 addresses in IPv6 addresses with different prefix lengths

0-15		16-31		32-47		48-63		64-79		80-96		96-111		112-128	
Prefix Prefix	:	Prefix Prefix	:	Prefix Prefix	:	Prefix Prefix	:	0 Prefix	:	Prefix Prefix	:	IPv4 IPv4	:	IPv4 IPv4	/96
Prefix Prefix	:	Prefix Prefix	:	Prefix Prefix	:	Prefix Prefix	:	u IPv4	:	IPv4 IPv4	:	IPv4 Suffix	:	Suffix Suffix	/64
Prefix Prefix	:	Prefix Prefix	:	Prefix Prefix	:	Prefix IPv4	:	u IPv4	:	IPv4 IPv4	:	Suffix Suffix	:	Suffix Suffix	/56
Prefix Prefix	:	Prefix Prefix	:	Prefix Prefix	:	IPv4 IPv4	:	u IPv4	:	IPv4 Suffix	:	Suffix Suffix	:	Suffix Suffix	/48
Prefix Prefix	:	Prefix Prefix	:	Prefix IPv4	:	IPv4 IPv4	:	u IPv4	:	Suffix Suffix	:	Suffix Suffix	:	Suffix Suffix	/40
Prefix Prefix	:	Prefix Prefix	:	IPv4 IPv4	:	IPv4 IPv4	:	u Suffix	:	Suffix Suffix	:	Suffix Suffix	:	Suffix Suffix	/32

Bits 64 to 71 (u) should always be set to zero even when using a /96 prefix.

The NAT64 translation causes a change in MTU. In addition to the minimum length of 40 Bytes for the IPv6 header, 20 Byte length of the IPv4 header have to be taken into account. If after IPv4 to IPv6 translation the IPv6 link MTU is exceeded, it is recommended to fragment the IPv4 packets before they enter the NAT and to set the Max Outside MTU of the NAT accordingly.

Max Outside MTU = IPv6 MTU - 40 Bytes IPv6 header - 8 Bytes IPv6 fragmentation header + 20 Bytes IPv4 header

For QoS consistency, the Traffic Class of the inbound IPv6 packet should be copied into the DSCP field of the outbound IPv4 packet. The DSCP of the inbound IPv4 packet should be copied into the Traffic Class field of the outbound IPv6 packet.

Some considerations should be taken regarding ALGs. In general, any application that communicates an IPv4 address in its downstream payload or a non-local IPv6 address in the upstream payload will be affected by NAT64. More specifically, there are some limitations to the functionality of the Application Layer Gateways (ALGs) in combination with NAT64 due to the way the ALG carries out translations. When translating inside information into outside information, IPv6 addresses are translated into IPv4 addresses without any issues. On the other hand, when an IPv4 address is received in the payload of an incoming message, this address will not be translated because it is a random outside address and not a NAT address. This has an impact on FTP, SIP and RTSP ALGs.

- SIP
  The connection information in a SIP message describes the IP addresses and ports to be used to connect to the other party of the call. From the perspective of a client behind a NAT64 gateway, his own IP address will be translated correctly. But the IP address received from the other side may be an IPv4 address and will not be translated into an IPv6 address. Thus, the NAT64 client will not be able to initiate a connection to the other client. If only one client is behind a NAT64 gateway, SIP calls are still possible. When client A (IPv4) can connect to client B (NAT64), client B can use this connection to connect back to client A. If both clients are behind the NAT64 gateway (the same or different), both clients will receive each other's IPv4 outside addresses and no client will be able to start the connection.

- RTSP
  Connection information in an RTSP message describes the IP address and ports to be used by the client to receive the media traffic. If the client is behind the NAT64 gateway, the server will receive correctly translated connection information and the client will be able to receive the data sent out by the server. If the server is behind the NAT64 gateway, the server will not receive translated connection information and the server will not be able to send out the data to the client.

- FTP
  Some servers may abort the connection when they receive the wrong type of address according to their current connection.

10.2.2 Technical Definition

NAT64 uses IPv6-only links between the provider and the customer to carry IPv6 or IPv4 packets. The home gateway (CPE) is provisioned with only an IPv6 address on its WAN-side interface. At the LAN-side interface, the CPE operates its own DHCP server, handing out IPv6 addresses to home devices. The CPE performs native IPv6 packet forwarding; the IPv6 to IPv4 translator is located on a NAT64 device in the provider's network, which performs NAT to enable connectivity with the IPv4 Internet.

The IPv6 packet from host in the home to an external destination is either natively forwarded for IPv6 destinations or is forwarded to a NAT64 device for IPv4 destinations. Packets to an IPv4 destination are assigned a public IPv4 address at the NAT64 device. The IPv6 source address is added to the NAT table along with the port such that the pair is mapped to the customer address. If a home device needs to access an IPv6 service, related traffic is transported "as-is" and routed to its Internet destination.

The following features are specified for NAT64:

- Carrier-grade address translation with high scale, high transaction rate and N+1 redundancy
- Two modes of operation:
  - NAPT where a single customer is sharing a public IP address with other users by unique assignment of port ranges
  - NAT 1:1 where a single customer is mapped to a single public IPv4 address
- Port range blocks assigned to each host for greater scalability and reduced logging requirements
- Per-subscriber port limits with application priority based on traffic forwarding class
- Configurable prioritization of applications traversing the NAT based on traffic forwarding class

- PCP for dynamic port requests which is required to enable initiation of communication from inside the Home Network such as in P2P applications and when servers (Web, E-mail, etc.) are operated

If deployed as described above, NAT64 will have the following benefits for the provider:

- Mitigate IPv4 address exhaustion
- Allow IPv4 services to continue to be accessed during the migration to IPv6 services
- Enable scalable logging

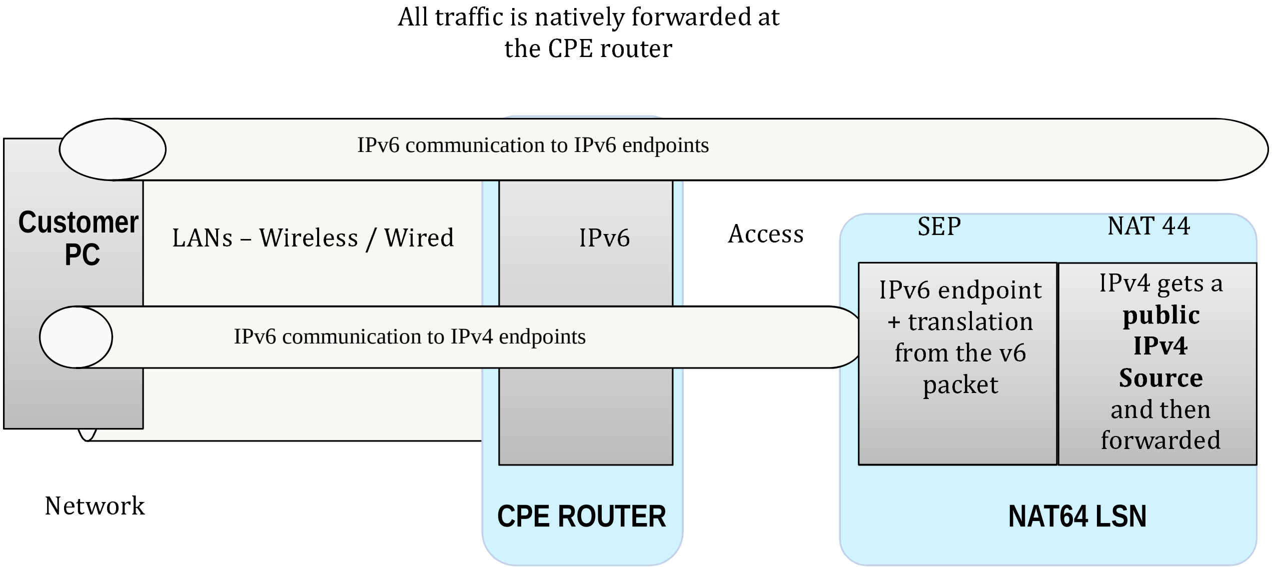
All traffic is natively forwarded at the CPE router
IPv6 communication to IPv6 endpoints
Customer PC
LANs – Wireless / Wired
IPv6
Access
SEP
NAT 44
IPv6 endpoint + translation from the v6 packet
IPv4 gets a **public IPv4 Source** and then forwarded
IPv6 communication to IPv4 endpoints
Network
CPE ROUTER
NAT64 LSN

Figure 9: Technical description of NAT64

10.2.3 Feature Synopsis

NAT64 features are based mainly on software requirements. As there is no specific CPE in a NAT64 architecture, the intelligence has to come from either client or a separate entity to provide separation between ordinary IPv6 native traffic and NAT64 traffic. This is achieved through DNS initialization and features that send a AAAA record in response to an A record request. The main features on the LSN are represented by ALG requirements. ALGs are providing the convergence layer between the different packet formats and protocol messages of IPv4 and IPv6.

10.2.3.1 NAT64 Hardware Features/Topology

- Role / location
  In some cases the NAT64 device performs the role of a 6PE or a 6VPE. In this case, the NAT64 device is placed at the network edge as close to the external peering points as possible. If the operator's topology includes multiple exit points, a balanced path should be implemented to all exit points (NAT64 devices will only be added based on capacity requirements).
  In case the NAT64 device does not perform the role of a 6PE or a 6VPE, it is placed at the network edge as close to the internal PEs as possible.

- Node type
  The NAT64 device should be chassis-based to guarantee scalability and hot-swap capability. Features and functionality should be distributed on blade or port level.

- Physical ports / traffic balance
  All physical ports should be nominally capable to handle 10 Gbit/s of traffic. The actual capacity requirement may depend on aggregation of node capacity.

- Memory
  Memory capacity should be sufficient to hold the full IPv6 and IPv4 BGP routing table. The minimum requirement is based upon the operator's PE standards.

- Integrated forwarding and NAT64 function
  The ICXF and NAT64 mechanisms should be integrated on a blade or ingress port.

- Forwarding Architecture
  Forwarding should be implemented in hardware with a minimum number of 4 million sessions per blade. Node latency is expected to be below 1 ms.

10.2.3.2 NAT64 Specific Software Requirements

- Client-Customer ID
  For data retention purposes, tunnel identifiers should be uniquely associated with a single CPE.

- NAT Timers
  The appropriate timers should be implemented to ensure proper operation of UDP, TCP, ICMP.

- Softwire Initialization (NAT64 Timers Extended)
  Quick drop quick pick-up approach is preferred.

- Port Block Allocation
  Port block allocation is configurable to allow for any ratio assignment.

- Static / Dynamic Port Allocation (STDPA)
  NAT64 requires efficiency and scalability not only with regard to throughput and node latency but also in terms of IP address ratios and customer port allocation (CPA - the amount of ports within the block assigned to a specific IPv6 home subnet or address (CBSA)). The NAT64 device needs to take into account the prefix-length of subscribers as all IPv6 addresses within this prefix need to be considered as belonging to a single subscriber. The port allocation limitations described above will apply on the aggregate of all home IPv6 devices. Setting the prefix length to 128 would imply that each IPv6 device in the home can independently use the maximum number of allocated ports.

- NDP / ECMP Integration
  Due to the nature of a NAT64 load balancing on links is extremely important. ECMP should be seen as a requirement and be integrated into NDP to allow for effecting events and ECMP use of best path matches.

- MTU Size / TCP MSS
  TCP MSS support should be mandatory for NAT64 due to the removal of an end-to-end MTU sizing functionality. This will avoid the need for excessive fragmentation.

- Fragmentation
  Fragmentation should be placed on the ASIC running in the linecard and should be pre-tunnel SI. NAT64 is expected support fragmentation and re-assembly. IPv6 mandates end points to discover the MTU of the communication path such that it is not required to fragment packets. However, NAT64 devices may receive already fragmented datagrams from the IPv4 domain or may need to fragment packets, if after conversion to IPv6 the packet exceeds the maximum IPv6 MTU (and the DF bit was not set in the IPv4 packet). In order to apply NAPT, the translator needs to re-assemble fragmented packets to the extent where the port numbers are visible. Other than the specific handling of the translation between IPv4 and IPv6 headers (including fragments), the same fundamental approach to fragmentation and re-assembly as adopted by other existing flavours of NAT should be followed. Some specific issues related to fragmentation and reassembly for NAT64 should be considered. If after translation to IPv6 the downstream IPv4 packet exceeds the IPv6 MTU (configured value or default value), an ICMP error "packet too big" will be sent back to the IPv4 source of the packet if the DF bit is set in the original IPv4 packet. Exceeding the IPv6 MTU is not expected to be common as the IPv6 initiating end-point will perform path MTU discovery (and should accordingly indicate TCP MSS to the peer). It is possible to receive already fragmented packets from IPv4 side. In this case, the translation rules for fragments have to be followed. If the translated fragment itself exceeds the IPv6 MTU, the fragment is discarded and an ICMP error is sent back to the IPv4 source.

- Support of 127-bit Prefixes on Inter-router Links
  Using 127-Bit IPv6 prefixes on inter-router links as described in [i.44] has to be supported natively.

- Load Balancing
  Load balancing across aggregated interfaces should be supported. This includes IP address and port allocation, ECMP, PVST load balancing, LDP hashing, etc.

- IPv4 Private Subnet Segmentation
  NAT64 should be able to segment IP address blocks into smaller blocks for the local interfaces.

- Non-ALG Deployment
  NAT64 should have a non-ALG approach for decapsulation where possible. However, for data applications such FTP, SIP and RTSP ALGs have to be supported.

- Traffic Prioritization
  Traffic prioritization should be possible to e.g. support B2B customers running video or voice to the translation interface. NAT64 should be able to copy the Traffic Class of the inbound IPv6 packet into the DSCP field of the outbound IPv4 packet. The DSCP of the inbound IPv4 packet should be copied into the Traffic Class field of the outbound IPv6 packet.

- NAT64 Logging
  IPFIX/Netflow is to be supported for flow logging. RADIUS/SYSLOG should be supported for port block logging.

- Clustering
  NAT64 clustering allows redundancy between two NAT64 devices. It should optimize capacity for NAT cache and configuration matching. Clustering should use the maximum efficiency for IP allocation as possible.

- Shared Resource (Single NAT64 GW Address)
  The architecture should support anycast addressing to simplify network dependencies.

- Physical Redundancy
  Physical redundancy in the service NPUs is required allowing for stand-by or active allocation during failure of a linecard or NPU.

- Thresholds / Watermarks
  The NAT64 device should provide thresholds and watermarks at system, pool and subscriber level to manage capacity assignments.

- NAT64 Address Withdrawal
  When the NAT64 device goes out of service the IPv6 addressing and IPv4 public addressing configured on the NAT64 device should be withdrawn.

- Transitional Co-existence
  NAT64 should be able to co-exist and share resources with other transitional methodologies such as NAT44 and DS-Lite.

- SI Monitoring
  The NAT64 device provides the necessary management interfaces and traps to monitor its operation.

- Data Retention (DR) RADIUS
  For DR purposes two potential methodologies exist, NAT caching with RADIUS or Netflow. The requirement is to hold the following parameters for each flow:

  - Source IPv6 address
  - Source IPv6 port
  - Remote destination port (external IPv4)
  - Time stamp
  - Remote IPv4 address
  - Remote IPv6 address (this will be the NAT64 address) (optional)
  - Incoming interface (optional)

  - Outgoing interface (optional)
  - Client ID (if different from the IPv6 address)

- UPnP/PMP/Port forwarding/PCP
  The NAT64 device has to allow UPnP forwarding and port mapping.

NAT64 software requirements depend on the deployment topology. In most cases, two topologies are required. In the "integrated topology", the NAT64 device functions as a full MPLS 6PE router. Alternatively, the NAT64 device as a L3 router, hairpinning connections through an external 6PE router. This topology is called "hairpin topology".

10.2.4 Performance Requirements and Comparison

NAT64 performance evaluated against lowest acceptable benchmarks and the actual delivery of throughput, convergence, failover and latency for all aspects and features within the chassis. All performance requirements are based on peak capacity and average throughput which are considered for the capacity the platform is designed for. Capacity estimates are determined by expected growth of subscribers and increase of throughput per subscriber.

- Throughput interfaces
  All interfaces are required to operate at line rate with 10 Gbit/s. The interface should be compatible with the operator's PE router connectivity.

- Node latency
  The required maximum node latency is 100 µs.

- Flow throughput
  The flow throughput is defined by three main performance figures:

  - CPE initialization
    This is the initialization of the port allocation per subscriber (i.e. the ports that are allocated when a CPE comes up for the first time) with a single external IP address per CPE.

  - Primary flow initialization
    Primary flow initialization is performed after the CPE has already been granted a port allocation. The flow is established as a "new" flow in the NAT cache. New flows are characterized by the lack of any entry except a source IPv6 address already in the cache. The whole flow needs to be allocated and registered into the NAT cache. The requirement is to be able to handle 800 000 flows per 40 Gbit/s chassis throughput capacity.

  - Secondary flow initialization
    The requirement is 1 million flows per 40 Gbit/s chassis throughput capacity.

- Convergence
  Convergence of routing and link failure should be well within 10 ms.

10.2.5 Development and Deployment Status

Product solutions for NAT64/DNS64 are widely available. But deployments are limited due to the fact that the technology is only applicable to IPv6-only devices. There are a number of applications which do not operate via DNS which further limits the applicability of NAT64/DNS64.

10.2.6 Failings/Issues of NAT64

There is a number of well-known issues with NAT64 which present topics for further development.

- The main issue with this technology is the fact that it does not operate on IPv4-only devices and that applications need to run via DNS. Applications like P2P do not consult the DNS and, thus, fail to operate since the NAT64 GW address cannot be discovered.

- NAT64 requires ALGs (ALPs to be exact in most cases) which require the NAT64 device to do some form of intelligent vicissitudes to the transit packets bearing the risk of drop in performance and limits in functionality. The risk of service-deprecation is minimized if all functionality is included natively in the NAT64 device. In addition, any application that communicates an IPv4 address in its downstream payload or a non-local IPv6 address in the upstream payload will not work flawlessly with NAT64.

- With NAT64, all functional intelligence is located in the NAT64 device. Thus, functionality that requires a public address in the local network has to be performed in the NAT64 device itself. An example of this is PCP. In NAT64 is placed on the LSN due to UPnP 1 and 2 requirements.

- To predict scaling requirements comparative to IPv6 utilization can be misleading. This is not specific to NAT64 but occurs similarly in other transition technologies. The issue is that it is largely unknown what portion of the traffic originated by a certain client will be IPv6. Thus, the system has to be scaled for a constant maximum of all clients and designs cannot adapt based on exact traffic predictions.

- NAT64 is a tunnel technology and, thus, suffers from MTU requirements beyond the norm. This means that MSS clamping, IPv4 fragmentation and IPv6 fragmentation are required in order to avoid ICMP blocking. Fragmentation resends and general PMTU control can have performance effects on customer services. Therefore, there is a major requirement to optimize implementations such that they do not cause a large impact on current RTTs and node latency.

- The requirement for Data Retention capacity is rather large and, thus, making the capacity available may become a major issue for the receiving logging server. Normally, on public addresses there is a single entry or two per week depending on lease times and utilization. Within NAT64, each time an individual subscriber receives a port and IP address assignment, the event should be logged. This can amount to 200 assignments a day with a start and stop time resulting in 400 logs per user per day. Assuming that an MSO could have 10 % of its customer base online at one time, that would mean a huge amount of logs.

10.2.7 Summary Assessment

NAT64 is less preferred as transition technology as only DNS based application will work at present. The CPE requirements, however, are limited. The only indispensable requirement is the support of IPv6. The NAT64 device itself, if deployed centrally, does not represent a major investment. Example figures indicate the expected range to be around 200 000 EUR per 1 million customers, which translates into a cost of 20 cents per subscriber.

Technically, the solution has been continuously improved in recent years with many of the known issues resolver and general performance improved. Current implementations are expected to provide 40 µs node latency on average and still less than 1 ms even with logging, fragmentation and ALGs.

10.3 Teredo

Although Teredo is not a transition technology for deployment from an MSO's point of view but the industry might be forced into considering its effects on its own transition technologies and native IPv6 deployment and thus in this clause we analyse the technology. There have been examples recently of customers bypassing the functionality of their own transition technology by switching on Teredo on their clients.

10.3.1 Technical Summary

Teredo is a transition technology that gives full IPv6 connectivity for IPv6-capable hosts which have no direct native connection to an IPv6 network. Compared to other similar technologies Teredo's distinguishing feature is that it is able to perform its function even from behind Network Address Translation (NAT) devices such as CPE Routers. Teredo operates using a platform independent tunneling protocol designed to provide IPv6 connectivity by encapsulating IPv6 packets within IPv4 User Datagram Protocol (UDP) packets. These datagrams can be routed on the IPv4 Internet and through NAT devices. Other Teredo nodes sometimes also called Teredo Relays that have access to the IPv6 network receive the packets, decapsulate them and route them to their final destinations. Teredo is intended to be a temporary measure since in the long term all IPv6 hosts should use native IPv6 connectivity. Teredo should, therefore, be disabled when native IPv6 connectivity becomes available. Teredo is standardized by the IETF in [i.33]. The Teredo Server listens on UDP port 3544.

A potential standardization of Teredo implementations in Cable Networks would need to address the following objectives:

- To define the logical and physical parameters allowing for customers to access the public Internet across an IPv6 network using automatic tunneling.
- To define the specific features to use.

Teredo can be summarised to provide the following:

- IPv6 connectivity for IPv6-capable hosts which are connected to an IPv4 network but which have no direct native connection to an IPv6 network
- Need to maintain IPv4 Access Network
- A lightweight solution

The Teredo protocol performs several functions:

- Diagnoses UDP over IPv4 (UDPv4) connectivity and discovers the type of NAT if any in the communication path (using a simplified replacement to the STUN protocol)
- Assigns a globally routable unique IPv6 address to each host using the Teredo service
- Encapsulates IPv6 packets inside UDPv4 datagrams for transmission over an IPv4 network (this includes NAT traversal)
- Routes traffic between Teredo hosts and native (or otherwise non-Teredo) IPv6 hosts

Teredo defines several different kinds of nodes:

- Teredo client: A host which has IPv4 connectivity to the Internet from behind a NAT and uses the Teredo tunneling protocol to access the IPv6 Internet. Teredo clients are assigned an IPv6 address that starts with the Teredo prefix (2001:0::/32).
- Teredo server: A well-known host which is used for initial configuration of a Teredo tunnel. A Teredo server never forwards any traffic for the client (apart from IPv6 pings) and has, therefore, very modest bandwidth requirements (a few hundred bits per second per client at most). This allows a single server to support large numbers of clients. Additionally, a Teredo server can be implemented in a fully stateless manner, thus using the same amount of memory regardless of how many clients it supports.
- Teredo relay: The remote end of a Teredo tunnel. A Teredo relay should forward all of the data on behalf of the Teredo clients it serves, with the exception of direct Teredo client to Teredo client exchanges. Therefore, a relay requires a lot of bandwidth and can only support a limited number of simultaneously active clients. Each Teredo relay serves a range of IPv6 hosts (e.g. a single campus/company, an ISP or a whole operator network or even the whole IPv6 Internet). It forwards traffic between any Teredo clients and any host within its network realm.
- Teredo host-specific relay: A Teredo relay whose range of service is limited to the individual host it runs on. It has no particular bandwidth or routing requirements. A computer with a host-specific relay will use Teredo to communicate with Teredo clients, but it will stick to its main IPv6 connectivity to reach the rest of the IPv6 Internet.

Each Teredo client is assigned a public IPv6 address which is constructed as follows:

- Bits 0 to 31 are set to the Teredo prefix (normally 2001::/32).
- Bits 32 to 63 embed the primary IPv4 address of the Teredo server that is used.

- Bits 64 to 79 can be used to define some flags. Currently only the higher order bit is used. It is set to 1 if the Teredo client is located behind a cone NAT, it is 0 otherwise. Some implementations use more bits with proprietary definitions. Usually, a defined bit mask identifies the functionality of the flags. An example is the format "CRAAAAUG AAAAAAAA" for the 16 bits, MSB first. In this case, "C" remains the "Cone" flag. The "R" bit is reserved for future use. The "U" bit is for the Universal/Local flag (set to 0). The "G" bit is Individual/Group flag (set to 0). The A bits are set to a 12-bit randomly generated number chosen by the Teredo client to introduce additional protection against IPv6-based scanning attacks.

- Bits 80 to 95 contain the obfuscated UDP port number. This is the port number that is mapped by the NAT to the Teredo client with all bits inverted.

- Bits 96 to 127 contains the obfuscated IPv4 address. This is the public IPv4 address of the NAT with all bits inverted.

Teredo servers are used by Teredo clients to autodetect the type of NAT behind which they are located (if any) through a simplified STUN-like qualification procedure. Teredo clients also maintain a binding on their NAT toward their Teredo server by sending a UDP packet at regular time intervals. That ensures that the server can always contact any of its clients which is required for mechanisms like hole punching to work properly.

If a Teredo relay (or another Teredo client) has to send an IPv6 packet to a Teredo client, it will first send a Teredo bubble packet to the client's Teredo server whose IP address can be inferred from the Teredo IPv6 address of the Teredo client. The server can then forward the bubble to the client, so the Teredo client software knows that hole punching should be done toward the Teredo relay.

Teredo servers can also transmit ICMPv6 packet from Teredo clients toward the IPv6 Internet. In practice, when a Teredo client wants to contact a native IPv6 node it should find out where the corresponding Teredo relay is (i.e. which public IPv4 and UDP port number to send encapsulated IPv6 packets to). To do that, the client crafts an ICMPv6 Echo Request (ping) toward the IPv6 node and sends it through its configured Teredo server. The Teredo server decapsulates the ping and forwards it to the IPv6 Internet. The ping eventually reaches the IPv6 node. The IPv6 node should then reply with an ICMPv6 Echo Reply. This reply packet will be routed to the closest Teredo relay which will finally try to contact the Teredo client.

Maintaining a Teredo server requires little bandwidth because it is not involved into the actual transmission and reception of IPv6 traffic packets. Also, it does not involve any access to the Internet routing protocols. The only requirements for a Teredo server are:

- the ability to emit ICMPv6 packets with a source address belonging to the Teredo prefix

- assignment of two distinct public IPv4 addresses (in some Teredo implementations both addresses are expected to be consecutive); the second IPv4 address is needed for the purpose of NAT detection

A Teredo relay potentially requires a lot of network bandwidth. Also, it should export (advertise) a route towards the Teredo IPv6 prefix (2001:0::/32) to other IPv6 hosts. In that way, the Teredo relay will receive traffic from the IPv6 hosts addressed to any Teredo client and forward it over UDP/IPv4. Symmetrically, it will receive packets from Teredo clients addressed to native IPv6 hosts over UDP/IPv4 and inject those into the native IPv6 network.

In practice, network administrators can set up a private Teredo relay for their company or campus. This will provide a short path between their IPv6 network and any Teredo client. However, setting up a Teredo relay on a scale beyond that of a single network requires the ability to export BGP IPv6 routes to other Autonomous Systems (ASs).

Unlike 6to4 where the two legs of a connection can use different relays, traffic between a native IPv6 host and a Teredo host will use the same Teredo relay in both directions, namely the one that is closest to the native IPv6 host with regard to the network topology. The Teredo host cannot localize a relay by itself since it cannot send IPv6 packets by itself. If it needs to initiate a connection to a native IPv6 host, it will send the first packet through the Teredo server which sends a packet to the native IPv6 host using the client's Teredo IPv6 address. The IPv6 host responds as usual to the client's Teredo IPv6 address which will eventually cause the packet to find a Teredo relay. The relay is able to initiate a connection to the client (possibly using the Teredo server for NAT piercing). Subsequently, the relay is used for communication between the two hosts for as long as is needed. This design means that neither the Teredo server nor the client needs to know the IPv4 address of any Teredo relays; a suitable one is automatically found by means of the global IPv6 routing table. This is enabled by the requirement on all Teredo relays to advertise the network 2001:0::/32.

A Teredo relay encapsulation causes a change in MTU size as the IPv4 header with 20 Bytes and the UDP header with 8 Bytes are added. If after encapsulation into UDP the IPv4 link MTU is exceeded, it is recommended to fragment the IPv6 packet before it is encapsulated in IPv4 and UDP and to set the Max Outside MTU on the Teredo relay accordingly:

Max Outside MTU = IPv4 MTU - 20 Bytes IPv4 header - 8 Bytes UDP header

For QoS consistency, the Traffic Class of the inbound IPv6 packet should be copied into the DSCP field of the outbound IPv4 packet. The DSCP of the inbound IPv4 packet should be copied into the Traffic Class field of the outbound IPv6 packet.

10.3.2 Technical Requirements Summary

Two network components are involved in the end-to-end Teredo communication, the Teredo client and the Teredo relay. A Teredo relay is a device that is placed at the edge of the network (IPE connected to the internal MPLS PE/LER) as the IPv6/IPv4 gateway to perform Teredo tunnel termination and IPv6 forwarding.

Requirement considerations for Teredo deployment are:

- Hardware topology
- Logical topology
- Software/hardware features
- Scalability
- Resilience and redundancy
- IP address allocation & DHCP specific features
- Forwarding/convergence performance
- Monitoring, management, reporting and access
- DR specific technical requirements

10.3.3 Technical Definition

Teredo uses IPv4-only links between the provider and the customer to carry IPv6 over IPv4. The home gateway (CPE) is provisioned with only an IPv4 address on its WAN interface. At the LAN-side interface, the CPE operates its own DHCP server, handing out IPv4 addresses to home devices. The CPE performs native IPv4 packet forwarding. The Teredo relay is a device in the provider's network, which performs tunnel termination and IPv6 forwarding.

The IPv6overIPv4 packet from the home device to an external IPv6 destination is tunnelled to a Teredo relay. Tunnel terminations are highly scalable and are required to offer N+1 redundancy. IPv4 destined packets are natively forwarded to a public IPv4 address.

If deployed as described above, the following benefits for the service provider apply:

- Ensures service provider IPv6 connectivity over an IPv4 infrastructure
- Allows IPv4 services to continue and evolve during the migration to IPv6 services

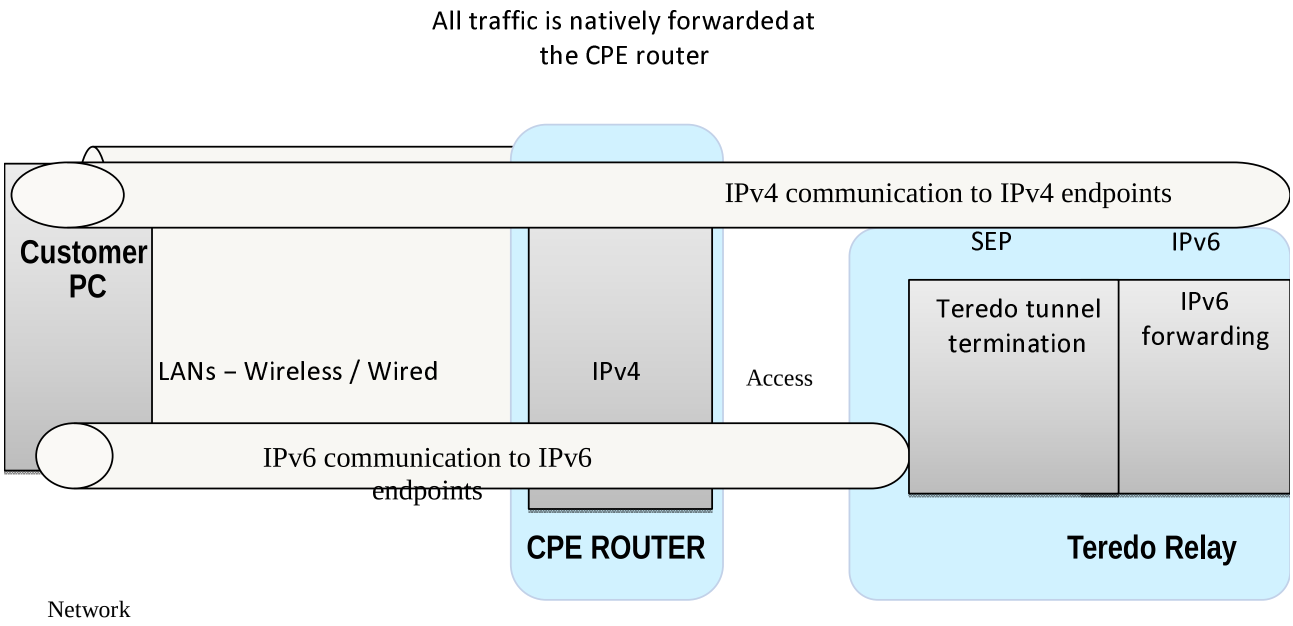
All traffic is natively forwarded at the CPE router
IPv4 communication to IPv4 endpoints
Customer PC
SEP
IPv6
Teredo tunnel termination
IPv6 forwarding
LANs – Wireless / Wired
IPv4
Access
IPv6 communication to IPv6 endpoints
CPE ROUTER
Teredo Relay
Network

Figure 10: Technical definition of Teredo

10.3.4 Performance Requirements and Comparison

Teredo performance is evaluated against lowest acceptable benchmarks and the actual delivery of throughput, convergence, failover and latency for all aspects and features within the chassis. All performance requirements are based on peak capacity and average throughput which are considered for the capacity the platform is designed for. Capacity estimates are determined by expected growth of subscriber numbers and increase of throughput per subscriber.

- Throughput interfaces
  All interfaces are required to operate at line rate with 10 Gbit/s. The interface should be compatible with the operator's PE router connectivity.

- Node latency
  The required node latency is 100 µs.

- Flow throughput
  The flow throughput is defined by three main performance figures:

  - CPE initialization
    This is the initialization of the port allocation per subscriber (i.e. the ports that are allocated when a CPE comes up for the first time) with a single external IP address per CPE.

  - Secondary flow initialization
    The requirement is 1 million flows per 40 Gbit/s chassis throughput capacity.

- Convergence
  Convergence of routing and link failure should be well within 10 ms.

10.3.5 Development and Deployment Status

Teredo is a solution that is specifically implemented for integration into the Microsoft Windows Operating System. It is not widely adopted in devices beyond those running this Operating System.

10.3.6 Failings/Issues of Teredo

There is a number of well-known issues with Teredo which present topics for further development.

- The main issue with this technology is the fact that it does not mitigate the depletion of IPv4 addresses and does not provide native IPv6 forwarding. As such, this technology can support fast IPv6 introduction but operators need to consider introducing another technology such as NAT64 or DS-Lite for native IPv6 forwarding.
- To predict scaling requirements comparative to IPv6 utilization can be misleading. This is not specific to Teredo but occurs similarly in other transition technologies. The issue is that it is largely unknown what portion of the traffic originated by a certain client will be IPv6. Thus, the system has to be scaled for a constant maximum of all clients and designs cannot adapt based on exact traffic predictions.
- Teredo is a tunnel technology and, thus, suffers from MTU requirements beyond the norm. This means that MSS clamping and IPv6 fragmentation are required in order to avoid ICMP blocking. Fragmentation resends and general PMTU control can have performance effects on customer services. Therefore, there is a major requirement to optimize implementations such that they do not cause a large impact on current RTTs and node latency.
- Teredo relies on ICMPv6 ping between a Teredo client and an IPv6 host on the Internet. The ICMPv6 ping is used for the Teredo client to discover his IPv6 address. As not all IPv6 hosts on the internet respond to ICMPv6 ping requests, the discovery might fail which prevents Teredo from operating.
- Teredo uses NAT44 UDP bindings. These bindings should be maintained using periodical messages. If the binding is lost, the overlaying IPv6 TCP/UDP/ICMP sessions fail.

10.3.7 Summary Assessment

Teredo is a transition technology that supports fast IPv6 introduction. It does not require any update of the CPEs. The Teredo relay and Teredo server do not have a high cost. As such, Teredo can be considered a simple and low-cost transition mechanism.

The main weakness of Teredo is its reliance on ICMPv6. When traversing NAT devices, care should be taken with UDP NAT bindings that may time out prematurely.

But the fundamental limitation of Teredo is that it does not mitigate IPv4 address depletion. NAT44 is required as a complement to Teredo to address IPv4 address depletion. Therefore, the total cost of transition based on Teredo may be less preferable.

10.4 dIVI

10.4.1 Technical Summary

dIVI is a 4464 solution focusing on the CPE to enforce the transitional functionality. Its concepts are in part reused and enhanced by the future MAP-T/E technologies. It is a stateless IPv6 transition technology. Dual IVI (dIVI) enables a service provider to offer IPv4 services to IPv6 enabled hosts. This approach utilizes stateless IPv4 to IPv6 translation (i.e. NAT64) to transit the IPv6-enabled network infrastructure. The Access Network can operate IPv6 only, while customers see IPv6 and IPv4 services simultaneously. dIVI keeps the stateful NAT44 on CPE to handle IPv4 address exhaustion. In addition, it introduces stateless NAT64 on CPE and Border Router to enable cross-connectivity between IPv4 and IPv6.

10.4.2 Address Family Translation (AFT)

Address Family Translation (AFT) refers to the mapping of an IP address from one address family to another IP address from another address family, for instance mapping an IPv4 address into an IPv6 address or vice versa. This translation is sometimes denoted as NAT46 (when the initiator is on the IPv4 side) or NAT64 (when the initiator is on the IPv6 side). AFT can be stateful or stateless. Stateless AFT is also known as IVI (in Roman numerals IV = 4 and VI = 6). IVI can be IPv4 or IPv6 initiated. An example architecture is shown in Figure 11.

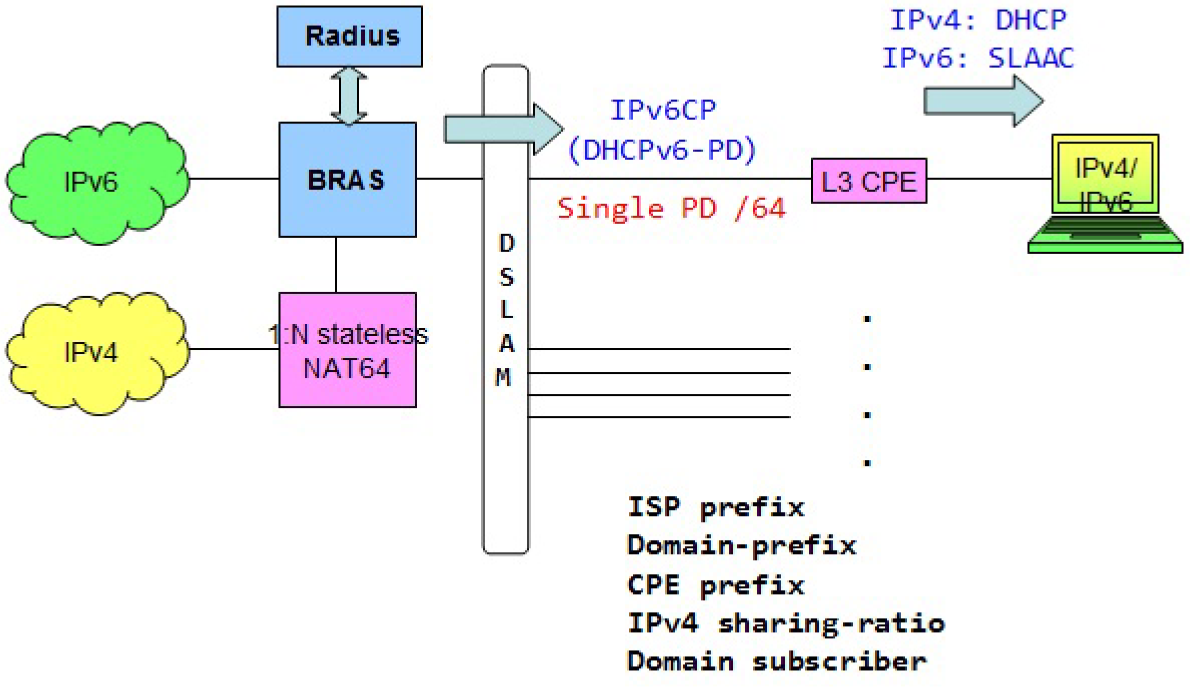

Figure 11: Example architecture of IVI

There are several variants of IVI that are offering different capabilities and requiring different architectures.

- Stateless 1:1 IVI
  - Full function on server and client (bi-directional initiation of connections)
  - Restricted IPv6 addresses (few, 1:1mapping)
  - Stateless and scalable
  - DNS64 and ALG required
- Stateless 1:N IVI (dIVI)
  - Limited function on server and client (bi-directional initiation of connections)
  - Restricted IPv6 addresses (many, 1:N mapping)
  - Stateless and scalable
  - DNS64 and ALG not required
- Stateful 1:N (NAT64)
  - Limited function on client (IPv6 initiation of connections only)
  - Unrestricted IPv6 addresses
  - Stateful and less scalable
  - DNS64 and ALG required

IPv6 traffic uses a separate path not involving the dIVI gateway.

10.4.3 Technical Definition

There is a significant difference between stateful NAT464 and stateless NAT464 also called dIVI or dIVI-pd if prefix delegation is required. dIVI-pd is used by A+P to share an IPv4 public address among multiple users/CPE without the help of NAT or double NAT.

Stateless NAT464 (dIVI, dIVI-pd) is intended to enable network operators (ISPs) to effectively share public IPv4 addresses among a set of customers. In parallel, it leverages IPv6 in the network in a manner that makes traffic originated by IPv4 customers looks like native IPv6 traffic resulting in simplified operations. More importantly, dIVI/dIVI-pd does not require any stateful NAT, DNS64 and ALG in the network. This avoids the need for the network operator to deal with any NAT logging, etc. dIVI maintains end-to-end address transparency and bidirectionally initiated communications.

With dIVI, address sharing is enabled by dual stateless IPv4/IPv6 translation. The 1:N XLAT (first translator) is the IPv4/IPv6 translator performing stateless 1:N translation between the IPv4 Internet and an IPv6 network. The CPE to CPE XLAT (second translator) is a 1:1 IPv6/IPv4 translator typically embedded with an IPv6 router that also performs port mapping functions for a host when it is communicating with the IPv4 Internet from a shared IPv4 address.

The IPv6 network routes the more specific IPv4-translatable addresses (longer than /64) to the corresponding CPEs. In the case of prefix delegation (shorter than /64), the dual stateless IPv4/IPv6 translation with prefix delegation (dIVI-pd) defined in [i.18] can be used.

Stateless translation requires that IPv6 nodes generate source port numbers in the range defined by the extended IPv4-translatable address. If this condition does not hold, the partial-state translation algorithm can be used which can map a random source port number to the port number range defined by the extended IPv4-translatable address. Related to the requirement to maintain state in the translator for the port mapping (no state is required for the address mapping), the algorithm refers to a partial state. The technical details of the mapping, state maintaining algorithm are determined by the implementation.

Prefix extension is only applied to dIVI-pd. It constructs different delegated prefixes for each CPE.

dIVI is an extension of stateless translation (IVI), which defines the following features:

- Sharing of public IPv4 addresses among several IPv6 hosts
- Support of bidirectionally initiated communications
- No requirement for ALG
- No requirement for DNS64
- Specifically considering operational aspects of commercial service provider networks:
    - When placing the dIVI XLAT in a centralized model:
        - It can replace a core router or be positioned at the interface with an interconnection partner
        - It can reduce cost, simplify management and is prepared for a stateless solution
    - When placing the dIVI XLAT in a distributed model:
        - It can be integrated with BRAS
        - Dual stack has to be maintained in the Core Network limiting the potential for decrease in operational cost
- Support of the identification of IPv6-only/dIVI subscribers in AAA
- Specifically considering aspects of network management:
    - Management of IPv6-only networks, including IPv6 MIB modules, Netflow, etc.
    - Manage the translation process, including dIVI device management, dIVI traffic monitoring, etc.

- Specifically considering aspects of CPE:
  - Support of IVI address assignment
  - Integration of user authentication function
  - Support of IPv4-to-IPv6 DNS proxy
- Introduction of port restricted NAT44 and stateless NAT46 in order to allow IPv4-only as well as IPv6-only hosts to access the IPv4 Internet
- Algorithmic mapping of IPv4 ports to/from IPv6 addresses (based on configured or well-known schema)
- Encapsulation employs IPv4-embedded IPv6 addresses
- Support of stateless NAT64 which can also be enabled in stateful mode for other IPv6-only clients
- IPv6 hosts use native addressing and IPv6 routing to public IPv6 Internet
- Specifically considering aspects of service continuity
  - Only upgrade Core Network to dual stack
  - No changes to existing IPv4 Access Network
- Minimal customer impact
  - Deploy IPv6-only data center with 1:1 IVI to move content to IPv6 without losing the IPv4 users
  - Deploy new IPv6-only Access Network with 1:N double IVI for new customers using shared IPv4 addresses
- Enabling of scalability and incremental deployment
- Support of a majority of current IPv4 applications
- Support of a variety of Operating Systems
- Avoid that port restriction have evident impact on user experience by reasonable port configuration

If deployed as described above, the following benefits apply:

- Mitigates service provider IPv4 address exhaustion
- Allows IPv4 services to continue and evolve during the migration to IPv6 services
- Layer 2-aware NAT removes requirement for unique private IP addresses per subscriber
- Allows for unique NAT subscriber policies such as PPP, PPPoE, L2TP LNS, and IPoE (DHCP)

10.4.4 dIVI Hardware Features/Topology

- Role/Location
  In some cases the AFT device performs the role of a 6PE. The AFT device should be placed at the network edge as close to the external peering points as possible. In case the network realm has multiple exit points, a balanced path should be implemented to all exit points.

  In case the AFT device does not perform the role of a 6PE, the device is placed at the network edge as close to the internal PEs as possible.

- Node Type
  The type of the node should be chassis-based to guarantee later scalability and hot-swap capability. Features and functionality should be distributed on blade or port level.

- Physical Ports / Traffic Balance
  Physical ports should have sufficient capacity to sustain the expected traffic requirements. Ports capable of 10 Gbit/s are recommended. The actual design may depend on aggregation of node capacity.

- Memory
  Memory capacity should be sufficient to hold the full IPv6 and IPv4 BGP routing table; minimum requirements are based on the operator's PE requirements.

- Integrated Forwarding
  The ICXF and AFT mechanisms should be integrated on blade or ingress port.

- Forwarding Architecture
  Forwarding should be implemented in hardware with a minimum number of 4 million sessions per blade. Node latency is expected to be below 1 ms.

10.4.5 AFTR dIVI Specific Software Requirements

- Tunnel Identifiers / Client-Customer ID
  For data retention purposes, tunnel identifiers should be uniquely associated with a single CPE.

- dIVI Timers
  Timers have to be configurable.

- Softwire Initialization (dIVI Timers Extended)
  Quick drop quick pickup approach is preferred.

- Port Block Allocation per IP Address
  Port block allocation is configurable to allow for any ratio assignment.

- Static / Dynamic Port Allocation (STDPA)
  dIVI requires efficiency and scalability not only with regard to throughput and node latency but also in terms of for IP address ratios and customer port allocation (CPA - the amount of ports within the block assigned to a specific CPE (CBSA).

- NDP / ECMP Integration
  Due to the nature of an AFTR load balancing on links is extremely important. ECMP should be seen as a requirement and be integrated into NDP to allow for effecting events and ECMP use of best path matches.

- MTU Size / TCP MSS
  TCP MSS support should be mandatory for the AFTR due to the removal of an end-to-end MTU sizing functionality. This will avoid the need for excessive fragmentation.

- Fragmentation
  Fragmentation should be placed on the ASIC running in the linecard and should be pre-tunnel SI.

- Load balancing
  Load balancing across aggregated interfaces should be supported. This includes IP address and port allocation, ECMP, PVST load balancing, LDP hashing, etc.

- IPv4 Private Subnet Segmentation
  The AFTR should be able to segment the IP address blocks into smaller blocks for the local interfaces.

- Non-ALG Deployment
  The AFTR should have a non-ALG approach for decapsulation where possible.

- Traffic Prioritization
  Traffic prioritization should be possible to e.g. support B2B customers running video or voice to the translation interface.

- dIVI IPFIX/Netflow
  Netflow v9 or IPFIX should be a common template allowing for both IPv4 and IPv6 to be demonstrated in a single entry. Note that 1:1 ratios should be applicable and both a dIVI and standard traffic flow template should be configurable.

- Clustering
  AFTR clustering allows redundancy between two AFTRs and should optimize capacity for NAT cache and configuration matching. Clustering should use maximum efficiency for IP address allocation.

- Shared Resource (Single AFTR GW Address)
  The AFTR should be a shared resource as much as possible with a single AFTR address. Backup AFTR addresses can be configured but one node should only use one address at a time.

- Physical Redundancy
  Physical redundancy in the service NPUs is required allowing for stand-by or active allocation during failure of a linecard or NPU.

- Thresholds / Watermarks
  Thresholds should be aligned to resources and be present in watermarks leading to events and changes in resource allocation.

- Relative Buffering
  Buffering should allow for multiple entry points into single or multiple NPU for NAT entries, packet buffers and logging.

- AFTR Address Withdrawal
  The AFTR should have at least five points of AFTR GW address withdrawal occurrence. The list includes loss of route-out, loss of all BGP/IGP sessions, loss of forwarding, loss of NPU capacity and certain errors in the NAT caching. Any of the failures should be detectable based on configurable timers with 15 seconds being the default setting.

- Transitional Co-existence
  dIVI should be able to co-exist and share resources with other transitional methodologies such as NAT44.

- IPv4 Private SI-ID for Block Resource Allocation
  At times there will be more than five Internet users behind a Cable Modem. This will mean that standard port block allocation may not be enough. As a solution IPv4 Private SI-IDs are proposed to be associated with heavily used B4 configurations. This feature would allow extra blocks to be dynamically allocated to those devices, thus, allowing higher levels of port allocation.

- SI Monitoring
  In order to supervise SI activity either for a return-path connection or for PCP, the tunnel establishment should be monitored and have an event.

- Anycast
  Anycast AFTR gateway addresses are a requirement to allow simplicity of deployment for a single address across multiple AFTRs.

- Data Retention (DR) RADIUS
  For DR purposes two potential methodologies exist, NAT caching with RADIUS or Netflow. The requirement is to hold the following parameters for each flow:

  - Source IPv4 address

  - Source IPv6 address

  - Source port (internal IPv4)

  - Remote destination port (external IPv4)

  - Time stamp

  - Remote IPv4 address

  - Remote IPv6 address (this will be the AFTR address) (optional)

  - Incoming interface (optional)

  - Outgoing interface (optional)

  - Client ID (if different from the IPv6 address)

- UPnP/PMP/Port forwarding/PCP
  The AFTR has to allow UPnP forwarding and port mapping through the tunnel.

The AFTR software requirements depend on the deployment topology. In most cases, topologies are required. In the "integrated topology", the AFTR functions as a full MPLS 6PE router. Alternatively, the AFTR functions as a L3 router, hairpinning connections through an external 6PE router. This topology is called "hairpin topology".

10.4.6 Feature Development Requirements

Table 7 presents features that have been identified to be in development aiming at enhancement of the AFTR capabilities. They are generalized and more analysis needs to be done to validate their scope and their necessity for deployment.

Table 7: dIVI features under development

FEATURE	DESCRIPTION
IPv6 Filtering	IPv6 filtering has been released for testing.
PCP	PCP is in development and a first release is expected in June. Further development will be required.
AFTR Clustering	AFTR clustering allows for the synchronization of the NAT cache states/entries across the network with limited additional network traffic. First release is expected by end 2012.
IPv4 Private SI-ID	The SI-ID is based on the IPv6 B4 address. It will eventually change to using both the IPv6 source address and the IPv4 private source address. There are issues with this feature due to performance changes and resource utilization. If available, port allocation can be carried out per host device and not per CPE.
RADIUS	RADIUS is tested for dIVI DR requirements. However, it will also serve its purpose for monitoring. RADIUS only effects performance by about 5 % for transit node latency. It is, therefore, the preferred method where legally available.

10.4.7 Performance Requirements and Comparison

The AFTR performance is evaluated against lowest acceptable benchmarks and the actual delivery of throughput, convergence, failover and latency for all aspects and features within the chassis. All performance requirements are based on peak capacity and average throughput which are considered for the capacity the platform is designed for. Capacity estimates are determined by expected growth of subscriber numbers and increase of throughput per subscriber.

- Throughput interfaces
  All interfaces are required to operate at line rate with 10 Gbit/s. The interface should be compatible with the operator's PE router connectivity.

- Node latency
  The required node latency is 100 µs.

- Flow throughput
  The flow throughput is defined by three main performance figures:

  - CPE initialization
    This is the initialization of the port allocation per subscriber (i.e. the ports that are allocated when a CPE comes up for the first time) with a single external IP address per CPE.

  - Primary flow initialization
    Primary flow initialization is performed after the CPE has already been granted a port allocation. The flow is established as a "new" flow in the NAT cache. New flows are characterized by the lack of any no entry except a source IPv6 address already in the cache. The whole flow needs to be allocated and registered into the NAT cache. The requirement is to be able to handle 800 000 flows per 40 Gbit/s chassis throughput capacity.

- Convergence
  Convergence of routing and link failure should be well within 10 ms.

10.4.8 dIVI CPE Requirements

dIVI is a dual stateless IPv4 / IPv6 translation mechanism. It is an extension of the 1:1 stateless IPv4 / IPv6 translation (IVI) mechanism. dIVI uses IPv4 to IPv6 address mapping at the CPE to translate an IPv4 header into IPv6. This packet is sent across an IPv6 access network to the XLAT / CGN. A translation back from IPv6 to IPv4 is achieved on the XLAT using the stateful NAT64 translation as defined in [i.42] and the IP/ICMP translation algorithm [i.41].

dIVI is defined in [i.17] with an extension defined in [i.18]. This adds support for prefix delegation at the CPE, improves the IPv4 address usage efficiency and removes the need for DNS64 and ALGs.

On initial connection, the CPE Router should have native IPv6 connectivity on its interface facing the Access Network. It is provisioned with:

- A global IPv6 prefix for native IPv6 communication
- A global IPv6 prefix for dIVI
- A public IPv4 address
- A port range (which is fixed by the provisioning system; the entire range can be used where no address sharing is required)
- One or more IPv6 DNS servers
- An XLAT IPv6 address

The following diagram shows the structure for a dIVI IPv6 mapped address.

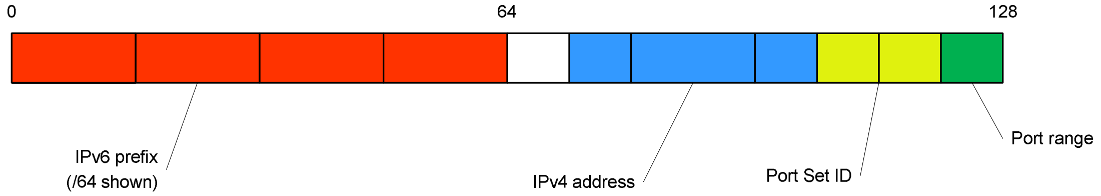

Figure 12: Address mapping in dIVI

The Home Network is required to be dual stack, with both global IPv6 addressing and private IPv4 addressing. Each host can be IPv4-only, IPv6-only or dual stack. All IPv6 connectivity is native end-to-end.

For an IPv4 client to reach an IPv4 server the following process is followed:

- The DNS query is forwarded to the CPE Router which acts as a DNS proxy. This is to avoid unnecessary port utilization by encapsulating DNS queries being sent to external IPv4 DNS servers.
- The CPE Router forwards the query to the specified IPv6 DNS server of the service provider and returns the response to the client.
- The client forwards the IPv4 packet to its default gateway (the CPE Router). The CPE Router creates an IPv6 header for the packet. The IPv6 source address is calculated using the mapping information as provisioned. This includes the IPv6 prefix, the public IPv4 address and a TCP or UDP port from the port range provided.
- The destination IPv6 address for the packet is the provisioned XLAT address.
- The packet formatted in the previous steps can be forwarded across the IPv6 Access Network to the XLAT.
- The XLAT is responsible for the IPv6 / IPv4 translation, restoring the public IPv4 address from the IPv6 source.

For inbound packets, the CPE Router will identify the dIVI packet by its destination address. It will perform IPv6 / IPv4 header translation using the mechanism described in [i.41]. After translation, the IPv4 packet can be forwarded on the Home Network to its intended recipient.

The dIVI specification [i.17] mentions that special MTU and fragmentation actions should be taken in the case of dual translation. However, it does not describe any details.

10.4.8.1 Feature Synopsis for dIVI in the Access Network

Requirements on the host are determined by the available connectivity in the Home Network. The Home Network may be dual stack or IPv4 only or IPv6 only. In the case that IPv4 is used, the Home Network has to support private addressing [i.23].

The following requirements apply to the CPE Router:

- WAN interface facing the Access Network
  - Request IPv6 global address via DHCPv6
  - Request IPv6 prefix via DHCPv6-PD
  - Request IPv4 address via DHCP
- LAN interface facing the Home Network
  - IPv4 private addressing
  - DHCPv4 service for LAN addressing
  - IPv4 DNS Proxy
  - IPv4 MTU set to 1 460 Bytes
  - TCP MSS clamping to 1 420 Bytes
- Router function
  - Receive XLAT IPv6 address via DHCPv6
  - Forward DNS queries to IPv6 server learnt from DHCPv6
  - Perform NAT44 between private IPv4 address in the Home Network and public IPv4 address
  - Translate ports to assigned port range
  - Perform stateless mapping [i.39]
  - Decapsulate IPv6 packet for IPv4 Home Network
  - Fragment encapsulated IPv6 packets
  - Reassemble received IPv6 fragments
  - Support UPnP and static port mapping
  - PCP

Assuming that the CPE Router is embedded in a Cable Modem, additional requirements apply:

- Should support bridging of IPv6 packets
- May support IPv4 or IPv6 management

10.4.9 Failings/Issues of dIVI

The main limitation of dIVI is the lack of focus it has received within the IETF. It is an attractive technology for IPv6 transition. Still, the MTU and fragmentation issues remain undefined.

The ability to provide static port mapping is limited as the NAT44 functionality on the CPE Router can only utilize the port range specified from provisioning. Additional communication between the CPE Router and the XLAT would be required to allow the selection of a specific port. Such a functionality is not available in current implementations.

The port randomization feature being introduced on many CGN devices to enhance security cannot be used with dIVI since a specific port range is provisioned for each customer.

10.5 Technical Comparison

As a result from the analysis above, a technical comparison of the introduced transition technologies takes into account the functionality of the technologies, their applicability for the environment of a Cable Network and the current state of development for vendor products. In Table 8 each transition technology is assessed with regard to several aspects, and a score from 1 to 10 with 1 being the lowest and 10 being the highest against the specific function for each architectural and deployment dependency is assigned.

Table 8: Technical Comparison

Feature/Function	DS-Lite	Teredo	NAT64	dIVI
Cost	2	10	9	1
Scalability	9	3	7	9
Redundancy	9	4	9	9
Technical Challenge	6	7	6	4
Service Deprecation	9	10	2	8
Standard Throughput	9	6	7	8
Logging Requirement/Ability	8	2	4	5
Logging Performance	9	1	9	9
Manageability by MSO	9	0	2	9
Future Toward full IPv6	9	0	9	9

The most significant cost in any subscriber network is bound in the CPE. A transition technology that requires both new Cable Modems and some kind of address translation functionality (e.g. AFTR) will have an impact on increasing CAPEX or OPEX of an MSO. Example technologies include DS-Lite and dIVI. Second to the cost, the ability to continue to provide all services to all customers is of highest priority. With the other two technologies (Teredo and NAT64) some serious service deprecation can be expected as the technologies do not accommodate an IPv6 transit network or fail in other areas. This represents the fundamental trade-off when determining the choice of transition technology.

Those technologies that use an IPv6 network to carry the IPv4 traffic, i.e. DS-Lite and dIVI, align with a strategic objective to smoothly migrate via the transitional step to an IPv6-only network.

11 IPv4 to IPv6 Future/Developing Transition Technologies

In this clause transition technologies are analysed that are more in a conceptual stage compared to those technologies found in clause 10. This may be due to the lack of a completed or stable specification as well as due to limited support by implementations and feature availability in products. However, similar to those technologies in clause 10 the technologies presented in this clause have a high potential to support the strategic goals for Cable Operators to implements IPv6 while addressing some of the shortfalls identified in current technologies.

11.1 464XLAT

11.1.1 Technical Summary

464XLAT is one of the future transition technologies. The main reason for considering 464XLAT for further analysis is its ability to operate without inter-communication between IPv4 and IPv6. This technology will allow customers to access services natively over IPv6 and through translation over IPv4. As IPv6 is not "backward compatible" to IPv4, i.e. the two protocols exist independently of each other without any interaction in most topologies, 464XLAT is introduced to allow for a smooth transition towards IPv6 once no more IPv4 addresses are available and a co-habitation of IPv4 and IPv6 on the same network infrastructure is required. 464XLAT also provides the ability for IPv4-only home devices, applications and OSs to continue to access the Internet with minimal use of IPv4 public addresses from the operator's remaining pool. It also avoids dual or multiple layers of address translation which otherwise introduces its own set of problems.

A potential standardization of 464XLAT technology would need to address the following objectives:

- To define the logical and physical parameters allowing customers to access the public Internet across an IPv6 network using IPv4/IPv6 translation techniques.
- To define the specific features to use.

464XLAT can be summarised to provide the following:

- IPv4 connectivity to IPv4 hosts over home routers (CPE) and Access Networks that are provisioned with IPv6 addresses only
- Dual stack connectivity for hosts connected to IPv6-only Access Networks
- Less need to maintain an IPv4 or dual stack Access Network
- A lightweight solution
- Single NAT, i.e. no need to have multiple layers of NAT
- Sharing of a limited number of public IPv4 addresses among a large number of customers using port translation
- Port forwarding capability on a PLAT using technologies such as: Web-UI, NAT-PMP, UPnP, A+P

464XLAT does not require DNS64 [i.43] since an IPv4 host may simply send IPv4 packets - including packets to an IPv4 DNS server - that are translated on the CLAT to IPv6 and back to IPv4 on the PLAT.

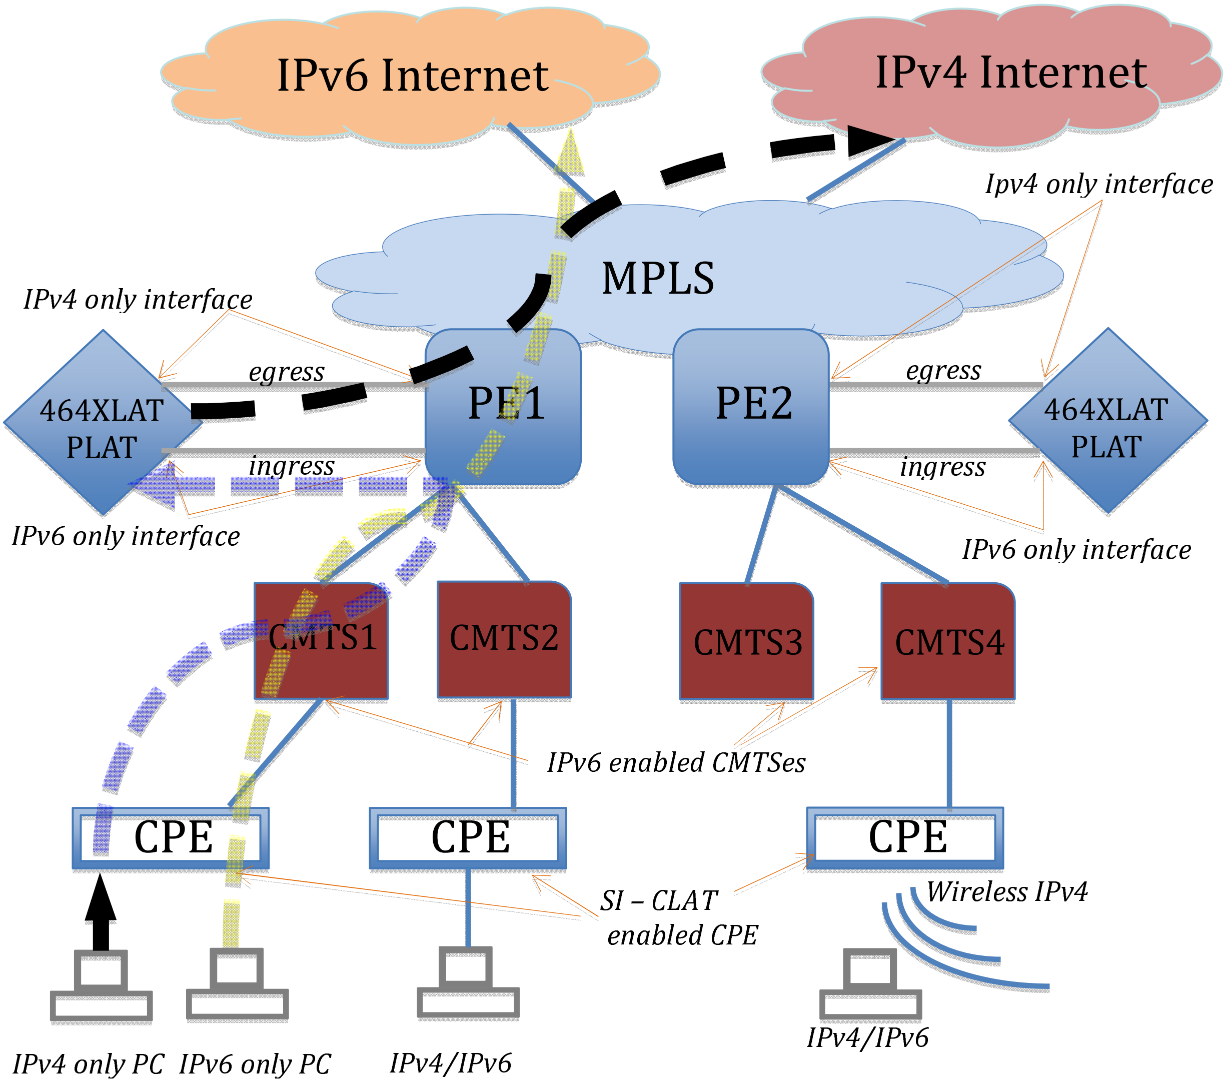

Figure 13: Packet flow in 464XLAT

The IPv6 address format in 464XLAT is defined in [i.39]. Prefix delegation mechanism such as DHCPv6-PD [i.30] are available to assign a dedicated translation prefix to the CLAT. From the delegated DHCPv6 prefix, a /64 prefix is dedicated for addressing outgoing and incoming IPv6 packets associated with the stateless translation as defined in [i.41]. The CPE (CLAT) may discover the related prefix of the PLAT via a variety of specified methods such as by means of a DHCPv6 option, using TR-069 [i.48], DNS APL RR [i.28] or the mechanism defined in [i.12].

In the case that DHCPv6-PD [i.30] is not available the CLAT does not have a dedicated IPv6 prefix for translation. It, alternatively, performs NAT44 and stateless translation [i.41]. For IPv4 packets coming from the Home Network and containing a source address from the private IPv4 address space [i.23] NAT44 is used to forward the packets to the CLAT IPv4 host address. The CLAT will execute a stateless translation [i.41] such that the IPv4 packets received at the CLAT IPv4 host interface are translated to the CLAT WAN IPv6 address as described in [i.39]. The IPv6 prefix is constructed of the delegated prefix which is completed if needed to form a /64 prefix by adding a subnet ID of 0.

The addressing scheme defined in [i.39] allows subnet lengths for the XLAT prefix to be 32, 40, 48, 56, 64 or 96 bits. Depending on the prefix length, the IPv6 address with embedded IPv4 address is formatted according to Table 9.

Table 9: Embedding IPv4 addresses in IPv6 addresses with different prefix lengths

0-15	16-31	32-47	48-63	64-79	80-96	96-111	112-128	
Prefix Prefix	Prefix Prefix	Prefix Prefix	Prefix Prefix	0 Prefix	Prefix Prefix	IPv4 IPv4	IPv4 IPv4	/96
Prefix Prefix	Prefix Prefix	Prefix Prefix	Prefix Prefix	u IPv4	IPv4 IPv4	IPv4 Suffix	Suffix Suffix	/64
Prefix Prefix	Prefix Prefix	Prefix Prefix	Prefix IPv4	u IPv4	IPv4 IPv4	Suffix Suffix	Suffix Suffix	/56
Prefix Prefix	Prefix Prefix	Prefix Prefix	IPv4 IPv4	u IPv4	IPv4 Suffix	Suffix Suffix	Suffix Suffix	/48
Prefix Prefix	Prefix Prefix	Prefix IPv4	IPv4 IPv4	u IPv4	Suffix Suffix	Suffix Suffix	Suffix Suffix	/40
Prefix Prefix	Prefix Prefix	IPv4 IPv4	IPv4 IPv4	u Suffix	Suffix Suffix	Suffix Suffix	Suffix Suffix	/32

Bits 64 to 71 (u) should always be set to zero even when using a /96 prefix.

A 464XLAT translation causes a change in MTU size. In addition to the minimum length of 40 Bytes for the IPv6 header, 20 Byte length of the IPv4 header have to be taken into account. If after IPv4 to IPv6 translation the IPv6 link MTU is exceeded, it is recommended to fragment the IPv4 packets before they enter the NAT and to set the Max Outside MTU of the NAT accordingly.

Max Outside MTU = IPv6 MTU - 40 Bytes IPv6 header - 8 Bytes IPv6 fragmentation header + 20 Bytes IPv4 header

For QoS consistency, the Traffic Class of the inbound IPv6 packet should be copied into the DSCP field of the outbound IPv4 packet. The DSCP of the inbound IPv4 packet should be copied into the Traffic Class field of the outbound IPv6 packet.

The CPE (CLAT) should implement a DNS proxy as defined in [i.37]. The case of an IPv4-only node behind the CLAT querying an IPv4 DNS server is undesirable since it requires both stateful and stateless translation for each DNS lookup. The CLAT itself should offer a DNS server via DHCP or other means and proxy DNS queries for IPv4 and IPv6 clients. Gateway functions are commonly implemented in CLAT-enabled home routers such that the operation of a DNS proxy is straight forward. The ultimate goal is to simplify traffic flows such that only IPv6 native queries are forwarded across the Access Network. The CLAT should allow for a client to query any DNS server of its choice and bypass the proxy.

Table 10 shows by means of an example the IP addresses that are used in different parts of the network during the IPv4/IPv6 address translation mechanism in the 464XLAT architecture. Note that the PLAT might also translate source ports in case NAPT is used. In the example the following IP addresses apply:

• IPv4 client	192.168.1.2
• IPv4 server	198.51.100.1
• CLAT IPv6 prefix	2001:DB8:AAAA::/96
• PLAT IPv6 prefix	2001:DB8:1234::/96
• PLAT IPv4 pool	192.0.2.1/24

Table 10: Addressing scheme in 464XLAT

IPv4 Home Network		IPv6 Access Network		IPv4 Internet	
IPv4 src	IPv4 dest	IPv6 src	IPv6 dest	IPv4 src	IPv4 dest
192.168.1.2	198.51.100.1	2001:DB8:AAAA: :192.168.1.2	2001:DB8:1234: :198.51.100.1	192.0.2.1	198.51.100.1

11.1.2 Technical Requirements Summary

Two network components are involved in the end-to-end 464XLAT communication, the PLAT and the CPE (CLAT). The PLAT is the LSN device that is placed at the edge of the network (IPE connected to the internal MPLS PE/LER) as the IPv6/IPv4 gateway. For egress packets, the PLAT performs address family translation from an IPv6 packet to a IPv4 packet with a public address.

Requirement considerations for PLAT deployment are:

- Hardware topology

- Logical topology
- Software/hardware features
- Scalability
- Resilience and redundancy
- IP address allocation and DHCP specific features (v4 and v6)
- Forwarding/convergence performance
- Monitoring, management, reporting and access
- DR specific technical requirements

11.1.3 Technical Definition

464XLAT uses IPv6-only links between the provider and the customer to carry IPv4 privately addressed packets. The 464XLAT home gateway (CPE) is provisioned with only an IPv6 address on its WAN interface. At the LAN-side interface, the CPE operates its own DHCP server handing out RFC 1918 private IPv4 addresses [i.23] to home devices. The CPE does not perform NAT, instead algorithmically translates private IPv4 addresses to global IPv6 addresses and vice versa. The NAT function is located on a carrier-grade NAT device in the provider's network, which is also a translator for global IPv6 addresses to global IPv4 addresses. This device is called PLAT.

The IPv4 packet from the home device to an external destination is translated in an IPv6 packet by the CPE and transported on the provider network. The packet is translated again to IPv4 at the PLAT, and NAT44 is performed to map the hosts' private IPv4 address to a public IPv4 address. The IPv6 source address created at the CPE is added to the NAT table. The IPv6 source address includes the CPE IPv6 prefix and the IPv4 source address of the host. As such it uniquely identifies the subscriber (CPE) and the host. If a home device needs to access an IPv6 service, packets are transported "as-is" and routed to an Internet server. With 464XLAT technology, the communications between endpoints remain within their address family without requiring protocol family translation.

NAT services on the PLAT allow service providers to conserve IPv4 addresses and maintain IPv4 Internet access while migrating to IPv6. If implemented in components of the Core Network (such as on the PLAT), an optimized software implementation provides scale, improved transaction rates and complete logging and accounting.

NAT services operate in two modes: Network Address and Port Translation (NAPT) optimized to provide scale and a subscriber-aware, Layer 2- aware NAT. In NAPT mode, an operator can deploy centralized devices to provide IPv4 service continuity with minimal changes to the Access Network and host devices. Layer 2-aware NAT enhances NAPT to create a virtual NAT table per subscriber. This allows for customized NAT policies and integrated RADIUS accounting; and it uniquely permits all subscribers to share a common internal IP address to simplify IPv4 address assignment and administration.

Example features of a PLAT include:

- NAT44 with high scalability, high transaction rate and N+1 redundancy
- Two modes of operation:
  - NAPT with traditional inside and outside addresses
  - Layer 2-aware NAT allowing for overlap and reuse of subscriber (inside) IP addresses
- Assigned port range blocks for each host to improve scalability and reduce logging requirements
- Per-subscriber port limits with application priority based on traffic forwarding class
- Configurable high/low prioritization of applications traversing the NAT based on traffic forwarding class

If deployed as described above, the following benefits for the service provider apply:

- Mitigates service provider IPv4 address exhaustion
- Allows IPv4 services to continue and evolve during the migration to IPv6 services

- Layer 2-aware NAT removes requirement for unique private IP addresses per subscriber
- Allows for unique NAT subscriber policies such as PPP, PPPoE, L2TP LNS and IPoE (DHCP)
- Does not require new protocols and, thus, enables quick deployment
- Leverage simpler IPv6 addressing structure to decrease operating costs

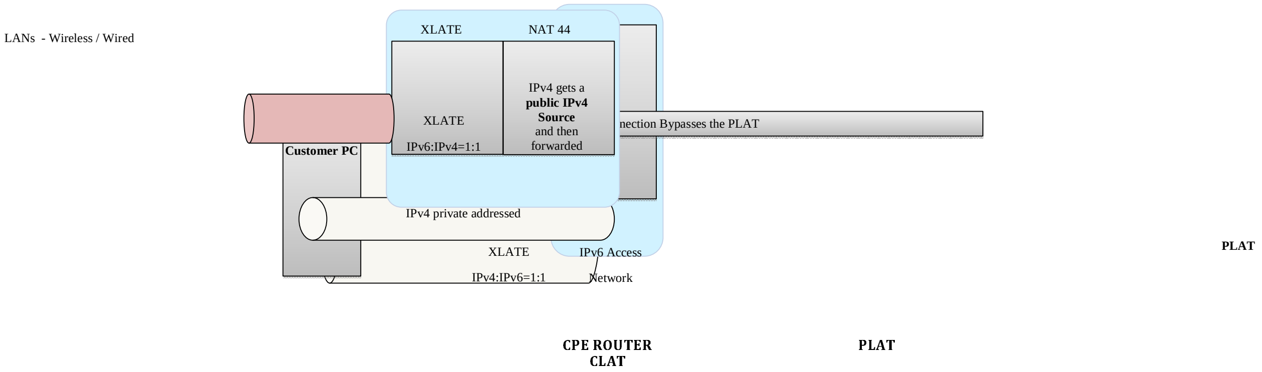

Figure 14: Technical definition of 464XLAT

11.1.4 PLAT Hardware Features/Topology

- Role / Location
  In some cases the PLAT performs the role of a 6PE or 6VPE. The PLAT should be placed at the network edge as close to the external peering points as possible. In case the network realm has multiple exit points, a balanced path should be implemented to all exit points (PLAT nodes will only be added based on capacity requirements).

  In case the PLAT does not perform the role of a 6PE, the PLAT is placed at the network edge as close to the internal PEs as possible.
- Node Type
  The type of the node should be chassis-based to guarantee later scalability and hot-swap capability. Features and functionality should be distributed on blade or port level.
- Physical Ports / Traffic Balance
  Physical ports should have sufficient capacity to sustain the expected traffic requirements. Ports capable of 10 Gbit/s are recommended. The actual design may be depend on aggregation of node capacity.
- Memory
  Memory capacity should be sufficient to hold the full IPv6 and IPv4 BGP routing table; minimum requirements are based on the operator's PE requirements.
- Integrated Forwarding
  The ICXF and PLAT mechanisms should be integrated on blade or ingress port.
- Forwarding Architecture
  Forwarding should be implemented in hardware with a minimum number of 4 million sessions per blade. Node latency is expected to be below 1 ms.

11.1.5 PLAT 464XLAT Specific Software Requirements

- NAT Timers
  Timers have to be configurable.
- Port Block Allocation per IP Address
  Port block allocation is configurable to allow for any ratio assignment.

- Static / Dynamic Port Allocation (STDPA)
  464XLAT requires efficiency and scalability not only with regard to throughput and node latency but also in terms of IP address ratios and customer port allocation (CPA - the amount of ports within the block assigned to a specific IPv6 source address (CBSA)).

- NDP / ECMP Integration
  Due to the nature of a PLAT load balancing on links is extremely important. ECMP should be seen as a requirement and be integrated into NDP to allow for effecting events and ECMP use of best path matches.

- MTU Size / TCP MSS
  TCP MSS support is mandatory for the PLAT due to the removal of an end-to-end MTU sizing functionality. This will avoid the need for excessive fragmentation.

- Fragmentation
  Fragmentation should be placed on the ASIC running in the line card. Fragmentation of IPv6 packets under certain requirements should be supported.

- Load Balancing
  Load balancing across aggregated interfaces should be supported. This includes IP address and port allocation ECMP, PVST load balancing, LDP hashing, etc.

- IPv4 Private Subnet Segmentation
  The PLAT should be able to segment IP address blocks into smaller blocks for the local interfaces.

- Non-ALG Deployment
  The PLAT should have a non-ALG approach for decapsulation where possible. However, ALGs should be a requirement with FTP, RSTP and SIP.

- Traffic Prioritization
  Traffic prioritization should be possible to e.g. support B2B customers running video or voice to the translation interface.

- 464XLAT IPFIX/Netflow
  Netflow v9 or IPFIX should be a common template allowing for both IPv4 and IPv6 to be demonstrated in a single entry. Note that 1:1 ratios should be applicable and both a 464XLAT and standard traffic flow template should be configurable.

- Clustering
  PLAT clustering allows redundancy between two PLATs and should optimize capacity for NAT cache and configuration matching. Clustering should use the maximum efficiency for IP address allocation.

- Shared Resource (Single PLAT GW Address)
  The PLAT should be a shared resource as much as possible with a single PLAT address. Backup PLAT addresses can be configured but one node should only use one address at a time.

- Physical Redundancy
  Physical redundancy in the service NPUs is required allowing for stand-by or active allocation during failure of a linecard or NPU.

- Thresholds / Watermarks
  Thresholds should be aligned to resources and be present in watermarks leading to events and changes in resource allocation.

- Relative Buffering
  Buffering should allow for multiple entry points into single or multiple NPU for NAT entries, packet buffers and logging.

- PLAT Address Withdrawal
  The PLAT should have at least five points of PLAT GW address withdrawal occurrence. The list includes loss of route-out, loss of all BGP/IGP sessions, loss of forwarding, loss of NPU capacity and certain errors in the NAT caching. Any of the failures should be detectable based on configurable timers with 15 seconds being the default setting.

- Transitional Co-existence
  464XLAT should be able to co-exist and share the resources with other transitional methodologies such as NAT44, NAT64 and DS-Lite.

- Anycast
  Anycast PLAT gateway addresses are a requirement to allow simplicity of deployment for a single address across multiple PLATs.

- Data Retention (DR) RADIUS
  For DR purposes two potential methodologies exist, NAT caching with RADIUS or Netflow. The requirement is to hold the following parameters for each flow:

  - Source IPv4 address
  - Source IPv6 address
  - Source port (internal IPv4)
  - Remote destination port (external IPv4)
  - Time stamp
  - Remote IPv4 address
  - Remote IPv6 address (this will be a PLAT address) (optional)
  - Incoming interface (optional)
  - Outgoing interface (optional)
  - Client-ID (if different from the IPv6 address)

The PLAT software requirements depend on the deployment topology. In most cases, two topologies are required. In the "integrated topology", the PLAT functions as a full MPLS 6PE router. Alternatively, the PLAT functions as a L3 router, hairpinning connections through an external 6PE router. This topology is called "hairpin topology".

11.1.6 Integrated Topology/Transport Requirements

The following list of features has to be supported by the 464XLAT architecture for both IPv4 and IPv6 / L2 and L3 in an integrated topology solution. Features marked with an asterisk (*) are topology dependent. Generally, all forwarding functionality should be implemented in hardware to prevent performance issues. However, there are some features where software based forwarding can compete with the equivalent hardware architecture.

- MP-BGP (including 6PE and 6VPE)
- BGP community / 32 bit AS
- MPLS LDP (currently only IPv4 is supported natively but the requirement for native IPv6 in MPLS is becoming indispensable)
- ECMP
- QoS (IPv4/IPv6) - classification, priority queuing, etc.
- QPPB/SCU/DCU
- SNMP v1/v2/v3 (transport over IPv4 and IPv6)
- ACLs/prefix lists/filtering (both IPv4/IPv6)
- TACACS/RADIUS (IPv4/IPv6)
- SYSLOG (event reporting for IPv4 and IPv6 as well as transport over both protocols)
- CoPP (IPv4/IPv6)

- Netflow v9 (potentially also previous versions will be required depending on the state of the implementation of Netflow)
- XML* (IPv4 and IPv6 reporting and transport)
- MAC accounting
- IEEE 802.1Q [i.52]
- EtherChannel
- Ethernet OAM
- NSF/GR (IPv4/IPv6)
- Policy Based Routing (IPv4/IPv6)
- IS-IS (potentially MT* for IS-IS as well if MPLS IPv6 LDP allows for dual stack) (IPv4/IPv6)
- Static Routing (IPv4/IPv6)
- OSPFv2/v3
- CDP/LLDP (IPv4/IPv6)
- VRRP/HSRP (IPv4/IPv6)
- VLAN mapping/Double Tagging
- L3 multicasting/MFIB (IPv4/IPv6)
- IPv6 forwarding (hardware)
- IPv4 forwarding (hardware)
- Ethernet technologies
- Virtual interfaces (IPv4/IPv6)
- AAA* (IPv4/IPv6) (extended beyond RADIUS)
- BFD (IPv4/IPv6)
- MLD/L2 multicast
- Full NDP (ICMPv6, DAD, NUD...)
- PIM/IGMPv2/v3
- CEF/dCEF
- Anycast
- Route reflection (IPv4/IPv6)
- Standard IPv4 VPN
- ISSU/SSO technologies
- NTP (IPv4/IPv6)
- SEND*
- IPsec
- DNS (IPv4/IPv6 server and client)

- DHCP relay (IPv4/IPv6)
- Graphical traffic and threshold monitor
- NMS
- Jump-off Services / Terminal console services

11.1.7 Failings/Issues of 464XLAT

There is a number of well-known issues with 464XLAT which present topics for further development.

- The cost, particularly of the CPE (CLAT), is expected to be increased due to the extension of the required functionality. MSOs that are not deploying WiFi solutions and stick to bridging CPE devices will face the issue of having to exchange CPE just for the reason of introducing 464XLAT.
- 464XLAT requires ALGs (ALPs to be exact in most cases) which require the PLAT to do some form of intelligent vicissitudes to the transit packets bearing the risk of drop in performance and limits in functionality.
- With 464XLAT, all functional intelligence is located in the PLAT. Thus, functionality that requires a public address in the local network has to be performed in the PLAT itself.
- To predict scaling requirements comparative to IPv6 utilization can be misleading. This is not specific to 464XLAT but occurs similarly in other transition technologies. The issue is that it is largely unknown what portion of the traffic originated by a certain client will be IPv6. Thus, the system has to be scaled for a constant maximum of all clients and designs cannot adapt based on exact traffic predictions.
- 464XLAT is an address family translation technology and, thus, suffers from MTU requirements beyond the norm. This means that MSS clamping, IPv4 fragmentation and IPv6 fragmentation are required in order to avoid ICMP blocking. Fragmentation resends and general PMTU control can have performance effects on customer services. Therefore, there is a major requirement to optimize implementations such that they do not cause a large impact on current RTTs and node latency.
- The requirement for Data Retention capacity is rather large and, thus, making the capacity available may become a major issue for the receiving logging server. Normally, on public addresses there is a single entry or two per week depending on lease times and utilization. Within 464XLAT, every time an individual subscriber receives a port and IP address assignment the event should be logged. This can amount to 200 assignments a day with a start and stop timer resulting in 400 logs per user per day. Assuming that an MSO could have 10 % of its customer base online at one time, that would mean a huge amount of logs.

11.1.8 Summary Assessment

464XLAT is extremely usable as a transition technology where the biggest concern remaining to be the additional cost of the CPE. The PLAT itself, if deployed centrally, does not represent a major investment.

464XLAT saves on public IPv4 addresses and allows for an IPv6-only Access Network with the potential of reducing operational costs.

Technically, the solution faces similar issues like NAT44 and DS-Lite and may leverage the continuous improvements that have been achieved on those technologies in recent years.

With the cost issue resolved, 464XLAT represents a fully functional solution.

11.2 Stateless DS-Lite

11.2.1 Technical Summary

Stateless DS-Lite is one of the future transition technologies. The main reason for considering Stateless DS-Lite for further analysis is its ability to operate without inter-communication between IPv4 and IPv6. This technology will allow customers to access services natively over IPv6 and through translation over IPv4. As IPv6 is not "backward compatible" to IPv4, i.e. the two protocols exist independently of each other without any interaction in most topologies, Stateless DS-Lite is introduced to allow for a smooth transition towards IPv6 once no more IPv4 addresses are available and co-existence of IPv4 and IPv6 on the same network infrastructure is required. Stateless DS-Lite also provides the ability for IPv4-only home devices, applications and OSs to continue to access the Internet with minimal use of IPv4 public addresses from the operator's remaining pool without requiring dual or multiple layers of address translation.

A potential standardization of Stateless DS-Lite would need to address the following objectives:

- To define the logical and physical parameters allowing customers to access the public Internet across an IPv6 network using IPv4/IPv6 translation techniques.
- To define the specific features to use.

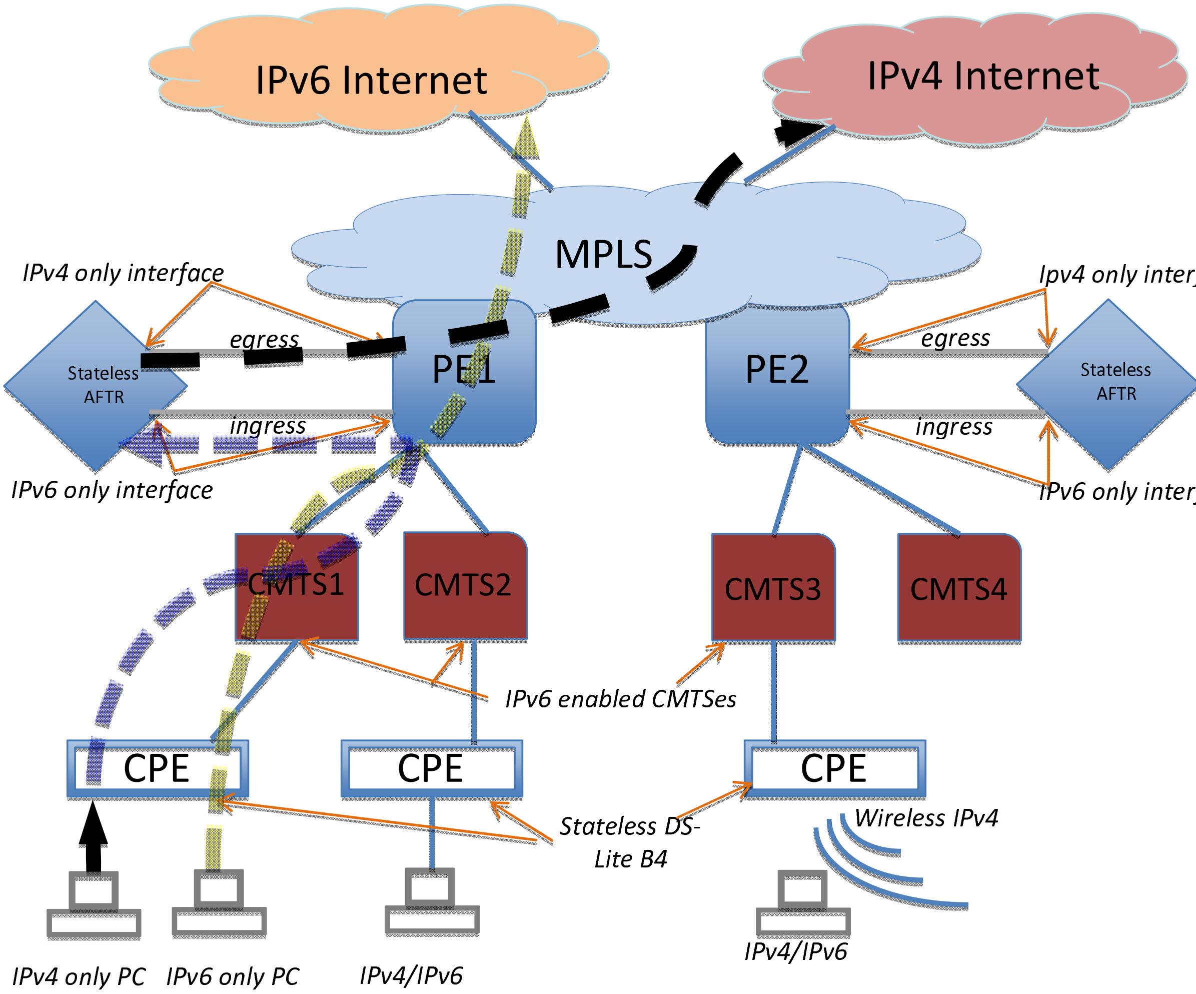

Figure 15: Packet flow in Stateless DS-Lite

Stateless DS-Lite can be summarised to provide the following:

- IPv4 connectivity to IPv4 hosts over home routers (CPEs) and Access Networks that are provisioned with IPv6 addresses only
- Dual stack connectivity for hosts connected to IPv6-only Access Networks

- Less need to maintain an IPv4 or dual stack Access Network
- A lightweight solution
- Single NAT on the CPE - i.e. no need to have multiple layers of NAT
- Sharing of a limited number of public IPv4 addresses among a large number of customers using port translation on the CPE
- Port forwarding capability on a CPE

Stateless DS-Lite behaves similar to stateful DS-Lite as described in clause 10.1 using a B4 component in the CPE Router and an AFTR as concentrator in the Core Network. The main difference is that the NAT function is provided by the CPE Router instead of the AFTR. A port range is provided to the CPE Router using DHCPv4 or DHCPv6.

Figure 16 shows an example of IPv4 connectivity on a Stateless DS-Lite architecture. The CPE Router learns the port range through DHCPv4 or DHCPv6 and the AFTR learns the port mapping for upstream traffic validation and downstream traffic encapsulation.

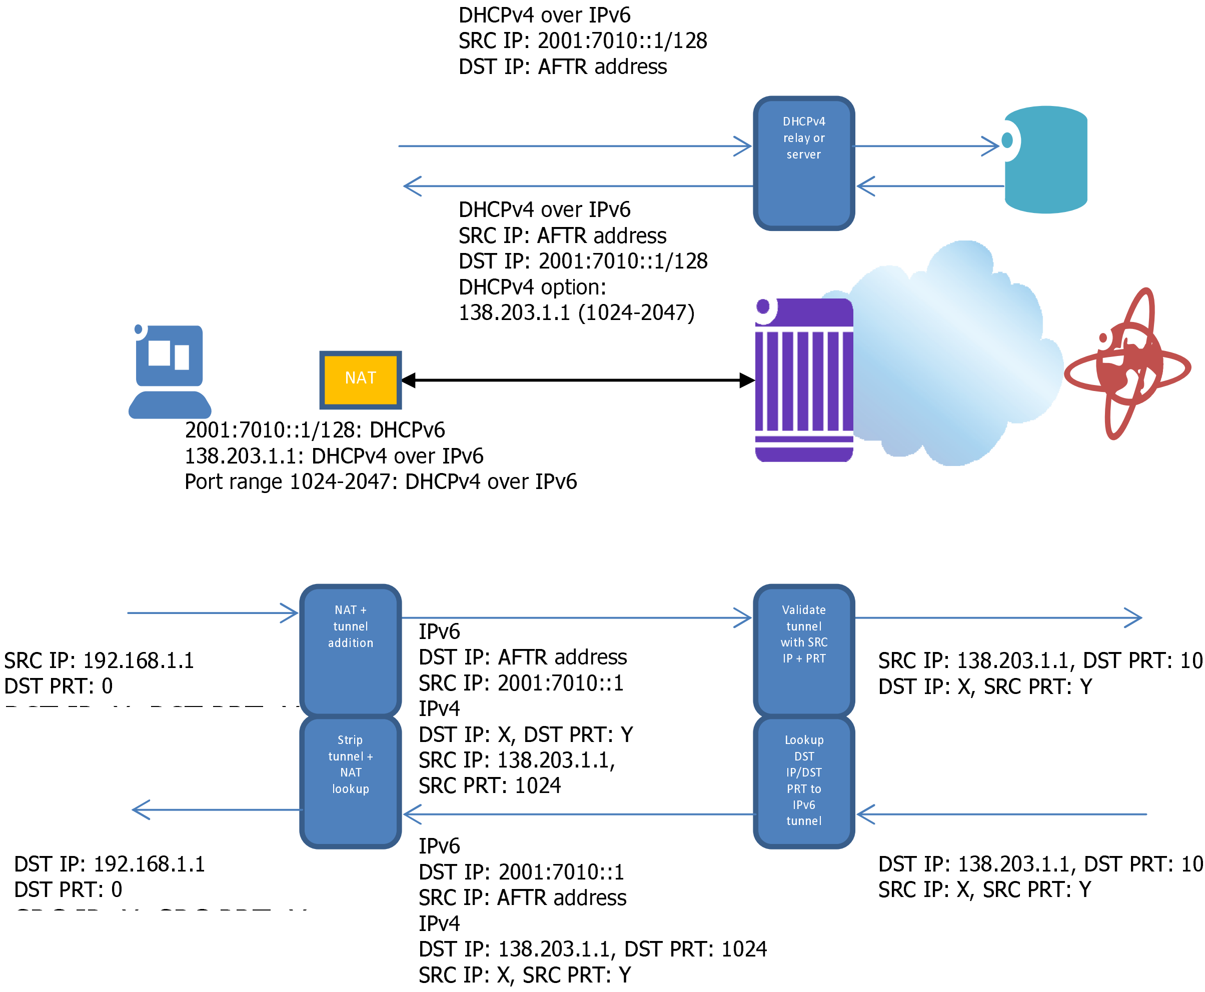

Figure 16: Addressing scheme in Stateless DS-Lite

11.2.2 Technical Requirements Summary

Two network components are involved in the end-to-end Stateless DS-Lite communication, the AFTR and the CPE Router (B4). The AFTR is the device that is placed at the edge of the network (IPE connected to the internal MPLS PE/LER) as the IPv6/IPv4 gateway to perform IPv4 decapsulation upstream and IPv4 over IPv6 encapsulation downstream based on the port mapping table.

Requirement considerations for AFTR deployment are:

- Hardware topology
- Logical topology
- Software/hardware features
- Scalability
- Resilience and redundancy
- IP address allocation & DHCP specific features (v4 and v6)
- Forwarding/convergence performance
- Monitoring, management, reporting and access
- DR specific technical requirements

11.2.3 Technical Definition

Stateless DS-Lite uses IPv6-only links between the provider and the customer to carry IPv4 packets. The B4 home gateway (CPE) is provisioned with only an IPv6 address on its WAN interface. At the LAN-side interface, the CPE operates its own DHCP server handing out RFC 1918 private IPv4 addresses [i.23] to home devices. The CPE performs IPv4 network address translation (NAT) within the port range allocated using DHCP. IPv4 packets are encapsulated on the CPE to be sent to the AFTR over IPv6 like with DS-Lite.

The privately addressed IPv4 packet from the home device to an external IPv4 destination is first translated to a packet with a public IPv4 address and afterwards tunnelled in an IPv6 packet towards the AFTR. If a home device needs to access an IPv6 service, packets are transported "as-is" and routed to an Internet server.

Example features of a Stateless DS-Lite architecture include:

- NAT44 within the CPE using the port range allocated by the provider
- DS-Lite encapsulation/decapsulation

If deployed as described above, the following benefits for the service provider apply:

- Mitigates service provider IPv4 address exhaustion
- Allows IPv4 services to continue and evolve during the migration to IPv6 services
- Requires extensions of DHCPv4 or DHCPv6
- Leverage simpler IPv6 addressing structure to decrease operating costs

11.2.4 AFTR Hardware Features/Topology

- Role/Location
  In some cases the AFTR performs the role of a 6PE or 6VPE. The AFTR should be placed at the network edge as close to the external peering points as possible. In case the network realm has multiple exit points, a balanced path should be implemented to all exit points (AFTR nodes will only be added based on capacity requirements).

  In case the AFTR does not perform the role of a 6PE, the AFTR is placed at the network edge as close to the internal PEs as possible.

- Node Type
  The type of the node should be chassis-based to guarantee later scalability and hot-swap capability. Features and functionality should be distributed on blade or port level.

- Physical Ports / Traffic Balance
  Physical ports should have sufficient capacity to sustain the expected traffic requirements. Ports capable of 10 Gbit/s are recommended. The actual design may be depend on aggregation of node capacity.

- Memory
  Memory capacity should be sufficient to hold the full IPv6 and IPv4 BGP routing table; minimum requirements are based on the operator's PE requirements adding the memory required for port forwarding tables.

- Forwarding Architecture
  Forwarding should be implemented in hardware with a minimum number of 4 million sessions per blade. Node latency should be below 1 ms.

11.2.5 B4/AFTR Specific Software Requirements

- NAT Timers
  Timers have to be configurable.

- Port Block Allocation per IP Address
  Port block allocation is configurable to allow for any ratio assignment.

- MTU Size / TCP MSS
  TCP MSS support is mandatory for the B4 due to the removal and end-to-end MTU sizing functionality. This will avoid the need for excessive fragmentation.

- Fragmentation
  Fragmentation should be placed on the ASIC running in the line card and should be pre-tunnel SI. Downstream fragmented packets will be buffered on the AFTR for downstream forwarding. The fragmentation/reassembly of IPv6 packets under certain requirements should be supported. A capability to handle large size packets (larger than the MTU) should be added.

- Load Balancing
  Load balancing across aggregated interfaces should be supported. This includes IP address and port allocation ECMP, PVST load balancing, LDP hashing, etc.

- Non-ALG Deployment
  The B4 should have a non-ALG approach for decapsulation where possible. However, ALGs should be a requirement with FTP, RSTP and SIP.

- Traffic Prioritization
  Traffic prioritization should be possible to e.g. support B2B customers running video or voice to the translation interface.

- B4 IPFIX/Netflow
  Netflow v9 or IPFIX should be a common template allowing for both the IPv4 and IPv6 to be demonstrated in a single entry.

- Clustering
  AFTR clustering allows redundancy between two AFTRs. Simultaneously operated AFTRs should maintain synchronization on the port block allocations.

- Physical Redundancy
  Physical redundancy in the service NPUs is required allowing for stand-by or active allocation during failure of a linecard or NPU.

- Thresholds / Watermarks
  Thresholds should be aligned to resources and be present in watermarks leading to events and changes in resource allocation.

- Relative Buffering
  Buffering should allow for multiple entry points into single or multiple NPU for NAT entries, packet buffers and logging.

- Transitional Co-existence
  Stateless DS-Lite should be able to co-exist and share the resources with other transitional methodologies such as NAT44, NAT64 and DS-Lite.

- Anycast
  Anycast Stateless AFTR gateway addresses are a requirement. If multiple AFTRs are identified by the same IP address, the port block tables need to be synchronized.

- UPnP/PMP/Port forwarding/PCP
  The B4 should allow UPnP forwarding and port mapping.

The AFTR software requirements depend on the deployment topology. In most cases two topologies are required. In the "integrated topology" the AFTR functions as a full MPLS 6PE router. Alternatively, the AFTR functions as a L3 router, hairpinning connections through an external 6PE router. This topology is called "hairpin topology".

11.2.6 Failings/Issues of Stateless DS-Lite

There is a number of well-known issues with Stateless DS-Lite which present topics for further development.

- Per Flow logging needs to be provided by an external DPI platform.

- Users which require to run servers or other traffic sources on well-known ports need special address/port range allocation considerations as the port range normally does not contain the appropriate ports.

- Dynamic port blocks are not possible and, hence, the allocation of the port blocks needs to manage carefully.

- Applications without ports do not work across Stateless DS-Lite. An example is. PPTP.

- The AFTR needs to buffer fragmented packets for proper downstream forwarding.

11.2.7 Summary Assessment

Stateless DS-Lite is usable as a transition technology where the biggest concern appearing to be the cost/complexities of the CPE. The AFTR itself, if deployed centrally, does not represent a major investment.

Stateless DS-Lite saves on public IPv4 addresses and allows for an IPv6-only Access Network with the potential of reducing operational costs.

Technically, the solution faces similar issues like DS-Lite where the NAT function is provided by the AFTR. It may leverage the continuous improvements that have been achieved on those technologies in recent years.

With the cost issue resolved, Stateless DS-Lite represents a fully functional solution. However, the notes in clause 10.1 apply to Stateless DS-Lite as well.

11.3 MAP-T

11.3.1 Technical Definition

MAP-T uses an IPv4/IPv6 address and port mapping technique as defined in [i.16] to encapsulate the IPv4 address suffix bits and port set ID within the IPv6 prefix.

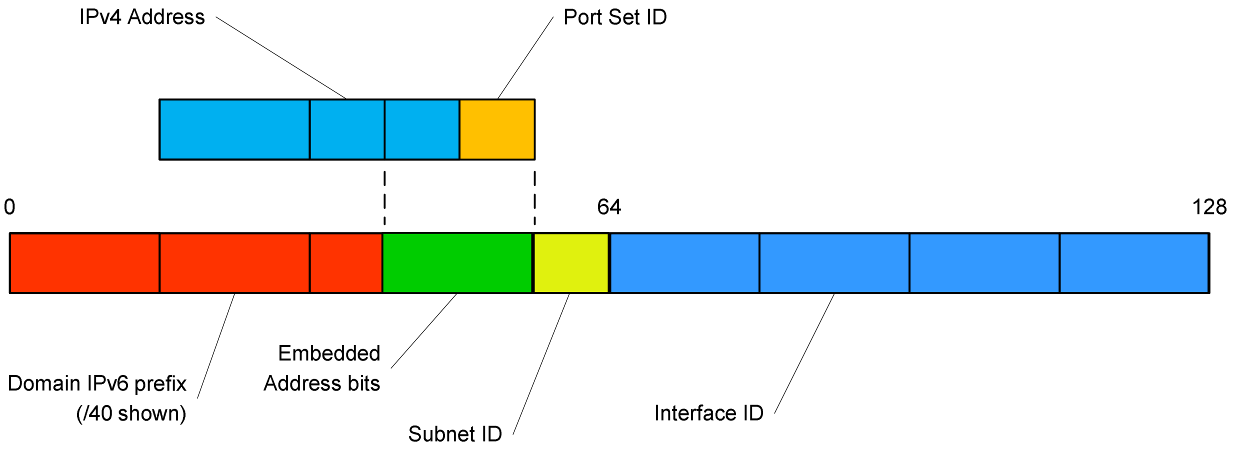

Figure 17: IPv6 address for MAP address and port mapping

Provisioning of the CPE Router requires the following attributes:

- IPv6 prefix - note that the IPv6 prefix is not dedicated for use with MAP-T and can be part of the customer-assigned IPv6 prefix used for native IPv6 traffic
- Embedded address bits
- MAP-T Border Relay IPv6 prefix
- Public IPv4 address (and prefix)
- Port set range
- DHCPv6 options defined in [i.15]

The Border Relay has to be configured with the public IPv4 address and length and the IPv6 prefix for translation in addition to its own IPv6 prefix.

For each outbound packet, the CPE Router is required to perform NAT44 translation from the private IPv4 address [i.23] to an address provided by the service provider. Once translated to a public IPv4 address, the CPE Router creates the mapped IPv6 packet header in accordance with the stateless IPv4/IPv6 translation mechanism defined in [i.41] and using the MAP header mapping rules. The translated packet is sent across the IPv6 network to the MAP-T Border Relay.

The Border Relay is responsible for translating the received packet back to IPv4, recovering the original IPv4 address from the IPv6 header using the MAP header mapping and translation technique defined in [i.41]. Finally, the packet is forwarded to its destination.

Inbound IPv4 packets are received by a Border Relay. The Border Relay translates the IPv4 header into IPv6 using the same translation technique and MAP header mapping to produce the IPv6 packet. The IPv6 packet, in turn, is forwarded to the CPE Router. The CPE Router uses once more the same translation mechanism to retrieve the IPv4 header and performs regular NAT44 to pass the packet on to the destination. Where an IPv6 server is located between the CPE Router and Border Relay, it remains accessible.

As the packet size varies between IPv4 and IPv6, MTU size and fragmentation are handled by the IP/ICMP translation. The use of IPv4 PMTUD will result in ICMP "packet too big" messages. TCP MSS Clamping is additionally required to ensure the overall IPv4 packet size does not exceed the IPv6 path MTU.

Data retention and lawful intercept systems will be required to capture the mapped IP address and port range for each customer.

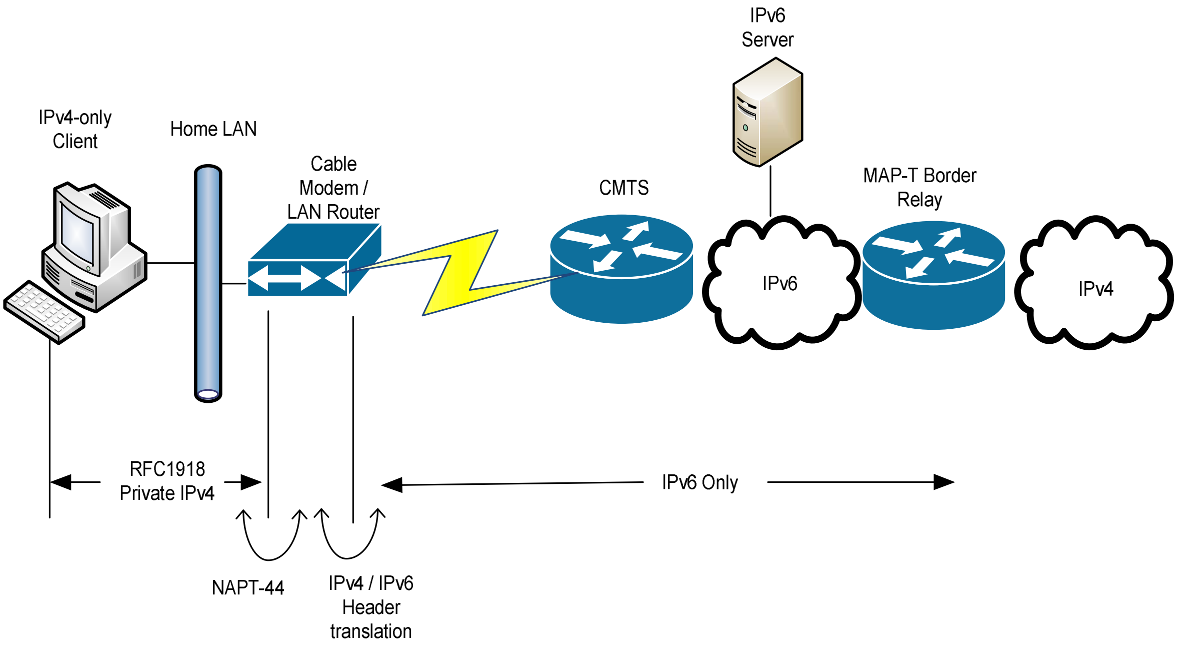

Figure 18: Technical definition of MAP-T

11.3.2 Feature Synopsis of MAP-T

The following requirements apply to the CPE Router:

- Native IPv6 on the WAN interface facing the Access Network
- Dual stack IPv4 / IPv6 on the LAN interface facing the Home Network
- IPv4 to IPv6 header translation in accordance with [i.41]
- IPv6 to IPv4 header translation
- NAPT44 to translate private IPv4 addresses to public IPv4 address and port range
- Provisioning of IPv6 prefix via DHCPv6
- DHCPv6 MAP-T options as defined in [i.15]
- Port forwarding and mapping for NAT translation
- UPnP NAT Traversal
- MSS Clamping for TCP/IPv4 connection negotiation
- PMTUD for both IPv4 and IPv6

Assuming that the CPE Router is embedded in a Cable Modem, the following requirements apply to the Access Network:

- IPv6 as defined in [i.5] or [i.49]
- Device management over IPv6 or IPv4

The MAP-T Border Relay has to support the following features:

- IPv6 prefix assigned for IPv6/IPv4 translation (customer side)
- IPv4 address for forwarding to/from IPv4 Internet
- IPv4 to IPv6 header translation in accordance with [i.41]

- IPv6 to IPv4 header translation
- MSS Clamping for TCP/IPv4 connection negotiation
- PMTUD for both IPv4 and IPv6 support

MAP-T requires dedicated functionality on the CPE Router which is similar to other transition technologies. However, due to the dual translation the CPE Router can still provide access to IPv6 servers and services from IPv4 hosts in the Home Network.

One of the unique benefits of MAP-T is that there is no requirement for Port Control Protocol (PCP) to configure the NAT within the network. Regular NAPT44 translation is maintained on the CPE Router.

For MAP-T, there are no dependencies on the hosts. The hosts can all be IPv4-only, IPv6-only or dual stack. This allows multiple hosts irrespective of configuration / IP stack availability to operate seamlessly with either protocol.

As a stateless technology, scaling and resilience can be achieved with the addition of common hardware. The address and port mapping technique provide a mechanism to store port and address information for DR/LI purposes.

11.3.3 Development and Deployment Status

MAP-T is currently in active development and - if adopted by the IETF - is likely to replace dIVI. There are currently no Border Relays or CPE Routers available for deployment.

11.3.4 Summary Assessment

MAP-T provides a good method of dual translation between IPv4 and IPv6. It has benefits over existing technologies since it allows for address sharing as well as fixed IPv4 address assignments.

The stateless design adds improvements for scalability and reliability within the service provider network.

A MAP-T CPE Router is backwards compatible with stateful NAT64, stateless NAT64 and DS-Lite technologies. A development effort of the CPE can be used to integrate MAP-T.

11.4 MAP-E

11.4.1 Technical Definition

MAP-E uses an IPv4/IPv6 address and port mapping technique as defined in [i.16] to associate an IPv4 address with an IPv6 address. This allows the automatic creation of tunnels between the Border Relay and the CPE Router. The format of the IPv6 address is the same as for MAP-T and shown in Figure 17. In order to perform the mapping, MAP-E allows the provisioning of an IPv4 address and port range to the CPE Router.

Provisioning of the CPE Router requires the following attributes:

- IPv6 prefix - note that the IPv6 prefix is not dedicated for use with MAP-E and can be part of the customer-assigned IPv6 prefix used for native IPv6 traffic
- Embedded address bits - these indicate the public IPv4 address and port range
- MAP-E Border Relay IPv6 prefix
- DHCPv6 options defined in [i.15]

The Border Relay has to be configured with a public IPv4 address and the IPv6 prefix for translation in addition to its own IPv6 prefix.

For each outbound packet, the CPE Router is required to perform NAT44 translation from the private IPv4 address [i.23] to the address provided by the Service provider and a port within the range specified. Once translated to a public IPv4 address, the CPE Router creates the mapped IPv6 packet header in accordance with the stateless IPv4/IPv6 translation mechanism defined in [i.41] and using the MAP header mapping rules. The translated packet is sent across the IPv6 network to the MAP-E Border Relay.

The Border Relay is responsible for translating the received packet back to IPv4, recovering the original IPv4 addresses from the IPv6 header using the MAP header mapping and translation technique defined in [i.41]. Finally, the packet is forwarded to its destination.

Inbound IPv4 packets are received by a Border Relay. The Border Relay translates the IPv4 header into IPv6 using the same translation technique and MAP header mapping to produce the IPv6 packet. The IPv6 packet, in turn, is forwarded to the CPE Router. The CPE Router uses once more the same translation mechanism to retrieve the IPv4 header and performs regular NAT44 to pass the packet on to the destination. Where an IPv6 server is located between the CPE Router and Border Router, it remains accessible.

As the packet size varies between IPv4 and IPv6, MTU size and fragmentation are handled by the IP/ICMP translation. The use of IPv4 PMTUD will result in ICMP "packet too big" messages. TCP MSS Clamping is additionally required to ensure the overall IPv4 packet size does not exceed the IPv6 path MTU.

The MAP-E Border Relay is responsible for fragmentation and reassembly of any IPv4 packet received prior to the encapsulation to IPv6.

Data retention and lawful intercept systems will be required to capture the mapped IP address and port range for each customer.

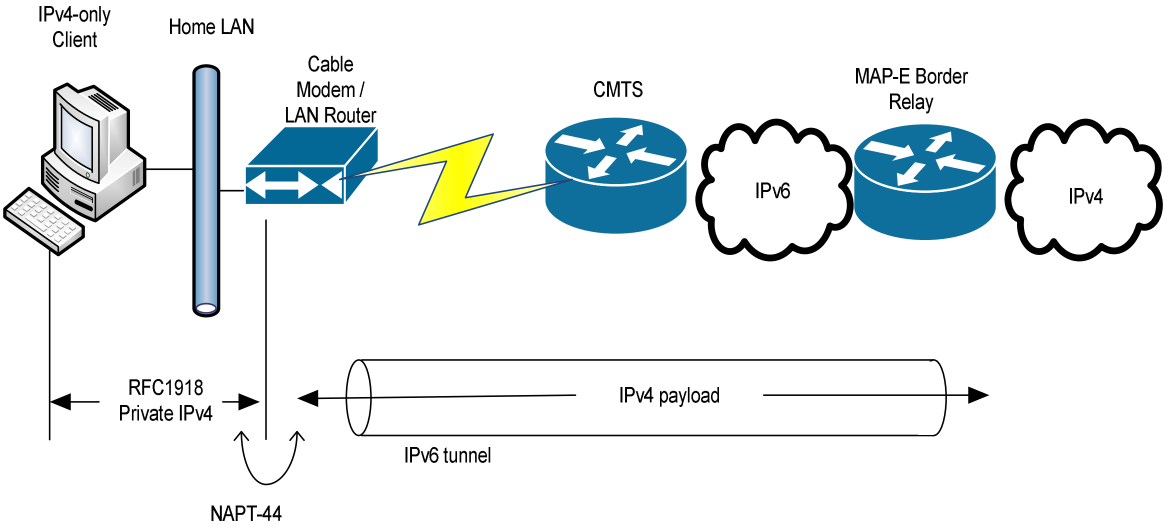

Figure 19: Technical definition of MAP-E

11.4.2 Feature Synopsis of MAP-E

The following requirements apply to the CPE Router:

- Native IPv6 on the WAN interface facing the Access Network
- Dual stack IPv4 / IPv6 on the LAN interface facing the Home Network
- IPv4 to IPv6 header translation in accordance with [i.41]
- IPv6 to IPv4 header translation
- NAPT44 to translate private IPv4 addresses to public IPv4 address and port range
- Provisioning of IPv6 prefix via DHCPv6
- DHCPv6 MAP-E options as defined in [i.15]
- Port forwarding and mapping for NAT translation
- UPnP NAT Traversal
- MSS Clamping for TCP/IPv4 connection negotiation

- PMTUD for both IPv4 and IPv6
- Fragmentation/reassembly for IPv6 packets

Assuming that the CPE Router is embedded in a Cable Modem, the following requirements apply to the Access Network:

- IPv6 as defined in [i.5] or [i.49]
- Device management over IPv6 or IPv4

The MAP-E Border Relay has to support the following features:

- IPv6 prefix assigned for IPv6/IPv4 translation (customer side)
- IPv4 address for forwarding to/from IPv4 Internet
- IPv4 to IPv6 header translation in accordance with [i.41]
- IPv6 to IPv4 header translation
- MSS Clamping for TCP/IPv4 connection negotiation
- PMTUD for both IPv4 and IPv6 support
- Fragmentation/reassembly for IPv6 packets

The primary difference between MAP-T and MAP-E is that MAP-E provides an encapsulation of the IPv4 packet whereas MAP-T translates the existing header between IPv4 and IPv6.

MAP-E has a slightly larger overhead in the packet header compared to MAP-T. However, as the original packet is encapsulated it has less risk of IP or ICMP header translation failure.

On a Cable Network, it is not possible to increase the overall MTU size due to the limitations specified for DOCSIS, so the maximum packet size has to be reduced to accommodate overhead for any transition technology. The additional overhead requirement for MAP-E is therefore less significant in this environment.

11.4.3 Development and Deployment Status

MAP-E is currently in active development and - if adopted by the IETF - is likely to replace 4RD. There are currently no Border Relays or CPE Routers available for deployment.

11.4.4 Summary Assessment

MAP-E provides a good method of dual translation between IPv4 and IPv6. It has benefits over existing technologies since it allows for address sharing as well as fixed IPv4 address assignments.

The stateless design adds improvements for scalability and reliability within the service provider network.

12 Technical Conclusion

The time is rapidly approaching when the last of the IPv4 addresses will be allocated worldwide. Consequently, the Cable Industry has to meet the new challenges in adopting IPv6 into its end-to-end networks including the home networks and end-user equipment. The most immediate investment being made is to transition a Cable Network from its current deployed base of equipment towards IPv6. There is actually no chicken and egg situation; deploying transition technologies merely enables a smooth migration from IPv4 to IPv6 but it will not mean that the industry will not move towards all IPv6 network topology. The timing of the migration will be governed by the operator's business requirements. Customers will be looking to have the ability to access services on IPv6 in addition to the IPv4 services that they are used to. To this end the Cable Operator will be challenged to provide an IPv6 address unless the Cable Operator decides to manage the IPv6 access via a tunneling technology such as Teredo. So, scenarios become more apparent where Cable ISPs will need to provide IPv6 in the very near future whether or not they do or do not have available IPv4 addresses that for example may suffice for their business needs beyond that. Those MSOs who do have the pool of addresses to last for their customer business continuity, the recommendation is to deploy IPv4/IPv6 dual stack. There are still concerns with dual stack - as with IPv6 as a whole - due to security, DNS prioritization issues and general lack of full validation/testing of the 70 000+ applications certified for Windows and 50 000+ applications certified for Unix based operating systems. This is already an immense challenge for the industry let alone similar issues with the mobile and tablet markets. The issues related to testing and verification are, by the way, irrespective of deploying either a dual stack or a transition methodology.

In the case, where a service provider's IPv4 address pool is being depleted and considering an eight year expected transitional life cycle of IPv4 to IPv6 to a point of 90 % IPv6 penetration, the service provider will need to invest into a solution for a transition method as an alternative to use NAT44/444. NAT44 holds its own issues and problems and, although it can be used along with dual stack, it does have its drawbacks. The present document described the possible transitional methods for Cable Networks with their pros and cons. Most of the conclusions are applicable to any IP-based network transitioning from IPv4 to IPv6. There are only a limited number of available choices at present, with little development going into many of the technologies and vendor market decisions looking at potential new technologies that will not fit the MSO network. There are reasonable recommendations revealed by this analysis that will allow, if the cost is not too high, a Cable ISP to encounter virtually no service deprecation.

Strategy is the main concern. How does a transitional strategy fit into an MSO's present environment and how does the vendor development of the transitional techniques within their products affect that strategy? Relevant proposals for answers to these questions have been developed in the present document taking into account various network and service scenarios.

The first potential strategy is dual stack with NAT44 or NAT444. Nested NAT in itself does not seem to be an optimal solution. But compared against many of the other transition technologies that were evaluated in preparation of the present document, the amount of IPv4 applications that have to be abandoned is about the same or even slightly lower while the implementation complexity and feature requirements are clearly in favour of NAT44. But for many networks this means the need to add private IPv4 space into the access layer up to the location of the CGN. This in itself is unwanted generally and has its own administrative and scalability issues.

For a clean and reasonably efficient transition to IPv6 dedicated technologies seem to be a more suitable option. Therefore, actually four types of transition technologies can be categorized; tunnel or translator and CPE-based or host-based. A tunnelled host-based example would be Teredo, whereas a tunnelled CPE-based example for a transition technology would be DS-Lite. NAT64 would be a representative of the host-based translator category mainly due to its ALTS (Application Layer Translation Service) as opposed to the proxy-defined ALGs known as ALPs. In both cases, packets stay in the same form end-to-end although they might get encapsulated by another protocol while in transit. The difference between the two approaches to ALGs is that one requires intelligent proxy methodology for functionality (ALP) and the other requires translation between the protocol functionality (ALTS). No technology from the CPE-based translator category was part of this analysis.

Clause 10 provides a detailed comparison between various technologies and assigns them to one of the above categories. Concluding from this comparison, there are only three technologies for recommendation at present when disregarding the potential of future developments.

DS-Lite seems to be by far the best option for most of the transitional requirements of MSOs identified in the course of this analysis. The most important reasons for this assessment are:

- the lack of the need to deprecate services that can only be delivered on native IPv4;
- the advanced development of relevant features in products currently available on the market;

- the CPE-based invocation of the functionality enabling operators to directly manage the behaviour; and
- the level of adoption by ISPs.

The throughput figures of DS-Lite packet delivery have been verified to be equal or better to a similar network without the CGN. With some network configurations, a node latency of 30 µs even for primary flows can be achieved. Logging is now limited to 2 Mbit/s write speed on average for a fully functioning set up with 100 000 customers. Although the peak rate can still be high, most implementations include RADIUS buffers to accommodate the peaks. Especially with potentially 1 million primary flow block allocations per second being available per NPU, a device with 9 NPUs could theoretically establish 9 million primary connections per second and, therefore, as a theoretical maximum require two times 9 million log entries if all connections are started and stopped within the same second. DS-Lite has significant benefits with most implementations but has a considerable cost. Although wireless and Layer 3 implementations on the CPEs being delivered to customers and the number of deployed devices reaching levels that are affecting the cost, there is still a large cost difference between a standard CPE and a DS-Lite CPE. This cost would likely need to be absorbed by the MSO with no real recovery options due to market competition and the lack of added value obvious to the customer. But if this cost premium can be handled, DS-Lite provides the market with a superior transitional mechanism.

There are other options that can be considered such as leaving it to Teredo to handle customer IPv6 needs or considering NAT64, which with its ALTS requirements and DNS initialization leaves significant risks for service deprecation from native IPv4.

In summary, transitional technologies are available that would allow to address all service dependencies and, thus, provide an ideal transition mechanism if the cost can be sustained.

Annex A: List of Applicable IP Standards

Date of Publication	***RFC***	***Title***	***URL***
IPv4			
September 1981	*791*	*Internet Protocol, Darpa Internet Program, Protocol Specification*	*http://www.ietf.org/rfc/rfc791.txt*
September 1981	*792*	*Internet Control Message Protocol, Darpa Internet Program, Protocol Specification*	*http://tools.ietf.org/html/rfc792*
September 1981	*793*	*Transmission Control Protocol, Darpa Internet Program, Protocol Specification*	*http://www.ietf.org/rfc/rfc793.txt*
August 1985	*950*	*Internet Standard Subnetting Procedure*	*http://tools.ietf.org/html/rfc950*
November 1987	*1034*	*Domain names - concepts and facilities*	*http://www.ietf.org/rfc/rfc1034.txt*
November 1987	*1035*	*Domain names - implementation and specification*	*http://www.ietf.org/rfc/rfc1035.txt*
October 1989	*1122*	*Requirements for Internet Hosts -- Communication Layers*	*http://tools.ietf.org/html/rfc1122*
December 1990	*1195*	*Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*	*http://tools.ietf.org/rfc/rfc1195.txt*
June 1992	*1338*	*Supernetting: an Address Assignment and Aggregation Strategy*	*http://tools.ietf.org/html/rfc1338*
November 1992	*1380*	*IESG Deliberations on Routing and Addressing*	*http://tools.ietf.org/rfc/rfc1380.txt*
September 1993	*1519*	*Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*	*http://www.ietf.org/rfc/rfc1519.txt*
IPv6			
August 1996	*1972*	*A Method for the Transmission of IPv6 Packets over Ethernet Networks*	*http://tools.ietf.org/html/rfc1972*
August 1996	*1981*	*Path MTU Discovery for IP version 6*	*http://www.ietf.org/rfc/rfc1981.txt*
January 1997	*2080*	*RIPng for IPv6*	*http://tools.ietf.org/html/rfc2080*
March 1997	*2119*	*Key words for use in RFCs to Indicate Requirement Levels*	*http://www.ietf.org/rfc/rfc2119.txt*
March 1997	*2132*	*DHCP Options and BOOTP Vendor Extensions*	*http://www.ietf.org/rfc/rfc2132.txt*
September 1997	*2185*	*Routing Aspects Of IPv6 Transition*	*http://tools.ietf.org/html/rfc2185*
December 1998	*2460*	*Internet Protocol, Version 6 (IPv6), Specification*	*http://www.ietf.org/rfc/rfc2460.txt*
December 1998	*2461*	*Neighbor Discovery for IP Version 6 (IPv6)*	*http://www.ietf.org/rfc/rfc2461.txt*
December 1998	*2462*	*IPv6 Stateless Address Autoconfiguration*	*http://tools.ietf.org/html/rfc2462*
December 1998	*2464*	*Transmission of IPv6 Packets over Ethernet Networks*	*http://tools.ietf.org/html/rfc2464*
January 1991	*2491*	*IPv6 over Non-Broadcast Multiple Access (NBMA) networks*	*http://tools.ietf.org/html/rfc2491*
January 1999	*2492*	*IPv6 over ATM Networks*	*http://tools.ietf.org/html/rfc2492*
May 1999	*2590*	*Transmission of IPv6 Packets over Frame Relay Networks Specification*	*http://tools.ietf.org/html/rfc2590*
August 1999	*2675*	*IPv6 Jumbograms*	*http://tools.ietf.org/html/rfc2675*
October 1999	*2710*	*Multicast Listener Discovery (MLD) for IPv6*	*http://www.ietf.org/rfc/rfc2710.txt*
May 2000	*2827*	*Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*	*http://www.ietf.org/rfc/rfc2827.txt*
October 2001	*3146*	*Transmission of IPv6 Packets over IEEE 1394 Networks*	*http://tools.ietf.org/html/rfc3146*
August 2001	*3152*	*Delegation of IP6.ARPA*	*http://tools.ietf.org/html/rfc3152*
July 2003	*3315*	*Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*	*http://www.ietf.org/rfc/rfc3315.txt*
July 2003	*3316*	*Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*	*http://www.ietf.org/rfc/rfc3315.txt*
August 2002	*3363*	*Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)*	*http://tools.ietf.org/html/rfc3363*

Date of Publication	*RFC*	*Title*	*URL*
February 2003	*3484*	*Default Address Selection for Internet Protocol version 6 (IPv6)*	*http://www.ietf.org/rfc/rfc3484.txt*
February 2003	*3493*	*Basic Socket Interface Extensions for IPv6*	*http://tools.ietf.org/html/rfc3493*
May 2003	*3542*	*Advanced Sockets Application Program Interface (API) for IPv6*	*http://www.ietf.org/rfc/rfc3542.txt*
September 2003	*3590*	*Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*	*http://tools.ietf.org/html/rfc3590*
October 2003	*3596*	*DNS Extensions to Support IP Version 6*	*http://www.ietf.org/rfc/rfc3596.txt*
December 2003	*3633*	*IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*	*http://www.ietf.org/rfc/rfc3633.txt*
December 2003	*3646*	*DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*	*http://www.ietf.org/rfc/rfc3646.txt*
January 2004	*3678*	*Socket Interface Extensions for Multicast Source Filters*	*http://www.ietf.org/rfc/rfc3678.txt*
April 2004	*3736*	*Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*	*http://tools.ietf.org/html/rfc3736*
June 2004	*3776*	*Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents*	*http://www.ietf.org/rfc/rfc3776.txt*
June 2004	*3810*	*Multicast Listener Discovery Version 2 (MLDv2) for IPv6*	*http://tools.ietf.org/html/rfc3810*
March 2005	*3971*	*SEcure Neighbor Discovery (SEND)*	*http://www.ietf.org/rfc/rfc3971.txt*
March 2005	*3972*	*Cryptographically Generated Addresses (CGA)*	*http://www.ietf.org/rfc/rfc3972.txt*
March 2005	*4029*	*Scenarios and Analysis for Introducing IPv6 into ISP Networks*	*http://www.ietf.org/rfc/rfc4029.txt*
May 2005	*4075*	*Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6*	*http://tools.ietf.org/rfc/rfc4075.txt*
November 2005	*4191*	*Default Router Preferences and More-Specific Routes*	*http://tools.ietf.org/html/rfc4191*
October 2005	*4193*	*Unique Local IPv6 Unicast Addresses*	*http://tools.ietf.org/html/rfc4193*
December 2005	*4301*	*Security Architecture for the Internet Protocol*	*http://tools.ietf.org/html/rfc4301*
November 2005	*4311*	*IPv6 Host-to-Router Load Sharing*	*http://tools.ietf.org/html/rfc4311*
January 2006	*4338*	*Transmission of IPv6, IPv4, and Address Resolution Protocol (ARP) Packets over Fibre Channel*	*http://tools.ietf.org/html/rfc4338*
February 2006	*4291*	*IP Version 6 Addressing Architecture*	*http://tools.ietf.org/html/rfc4291*
February 2006	*4364*	*BGP/MPLS IP Virtual Private Networks (VPNs)*	*http://www.ietf.org/rfc/rfc4364.txt*
April 2006	*4429*	*Optimistic Duplicate Address Detection (DAD) for IPv6*	*http://www.ietf.org/rfc/rfc4429.txt*
March 2006	*4443*	*Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*	*http://tools.ietf.org/html/rfc4443*
July 2006	*4584*	*Extension to Sockets API for Mobile IPv6*	*http://www.ietf.org/rfc/rfc4584.txt*
August 2006	*4605*	*Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")*	*http://tools.ietf.org/html/rfc4605*
August 2006	*4607*	*Source-Specific Multicast for IP*	*http://tools.ietf.org/html/rfc4607*
August 2006	*4632*	*Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*	*http://tools.ietf.org/html/rfc4632*
January 2007	*4760*	*Multiprotocol Extensions for BGP-4*	*http://tools.ietf.org/html/rfc4760*
March 2007	*4821*	*Packetization Layer Path MTU Discovery*	*http://www.ietf.org/rfc/rfc4821.txt*
September 2007	*4861*	*Neighbor Discovery for IP version 6 (IPv6)*	*http://tools.ietf.org/html/rfc4861*
September 2007	*4862*	*IPv6 Stateless Address Autoconfiguration*	*http://www.ietf.org/rfc/rfc4862.txt*
April 2007	*4877*	*Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture*	*http://tools.ietf.org/html/rfc4877*
April 2007	*4884*	*Extended ICMP to Support Multi-Part Messages*	*http://tools.ietf.org/html/rfc4884*
September 2007	*4941*	*Privacy Extensions for Stateless Address Autoconfiguration in IPv6*	*http://tools.ietf.org/html/rfc4941*

Date of Publication	***RFC***	***Title***	***URL***
September 2007	*4944*	*Transmission of IPv6 Packets over IEEE 802.15.4 Networks*	*http://tools.ietf.org/html/rfc4944*
September 2007	*5014*	*IPv6 Socket API for Source Address Selection*	*http://www.ietf.org/rfc/rfc5014.txt*
September 2007	*5006*	*IPv6 Router Advertisement Option for DNS Configuration*	*http://tools.ietf.org/html/rfc5006*
September 2007	*5072*	*IP Version 6 over PPP*	*http://tools.ietf.org/html/rfc5072*
December 2007	*5095*	*Deprecation of Type 0 Routing Headers in IPv6*	*http://www.ietf.org/rfc/rfc5095.txt*
March 2008	*5175*	*IPv6 Router Advertisement Flags Option*	*http://tools.ietf.org/html/rfc5175*
February 2008	*5121*	*Transmission of IPv6 via the IPv6 Convergence Sublayer over IEEE 802.16 Networks*	*http://tools.ietf.org/html/rfc5121*
October 2008	*5305*	*IS-IS Extensions for Traffic Engineering*	*http://tools.ietf.org/html/rfc5305*
October 2008	*5308*	*Routing IPv6 with IS-IS*	*http://tools.ietf.org/html/rfc5308*
February 2009	*5453*	*Reserved IPv6 Interface Identifiers*	*http://tools.ietf.org/html/rfc5453*
June 2009	*5555*	*Mobile IPv6 Support for Dual Stack Hosts and Routers*	*http://tools.ietf.org/html/rfc5555*
October 2009	*5722*	*Handling of Overlapping IPv6 Fragments*	*http://tools.ietf.org/html/rfc5722*
February 2010	*5790*	*Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols*	*http://tools.ietf.org/html/rfc5790*
July 2010	*5942*	*IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes*	*http://tools.ietf.org/html/rfc5942*
August 2010	*5952*	*A Recommendation for IPv6 Address Text Representation*	*http://tools.ietf.org/html/rfc5952*
September 2010	*5996*	*Internet Key Exchange Protocol Version 2 (IKEv2)*	*http://tools.ietf.org/html/rfc5996*
November 2010	*6106*	*IPv6 Router Advertisement Options for DNS Configuration*	*http://tools.ietf.org/html/rfc6106*
April 2011	*6204*	*Basic Requirements for IPv6 Customer Edge Routers*	*http://tools.ietf.org/html/rfc6204*
July 2011	*6275*	*Mobility Support in IPv6*	*http://tools.ietf.org/html/rfc6275*
December 2011	*6434*	*IPv6 Node Requirements*	*http://tools.ietf.org/html/rfc6434*
November 2011	*6437*	*IPv6 Flow Label Specification*	*http://tools.ietf.org/html/rfc6437*
IPv4 - IPv6 transition			
August 2000	*2893*	*Transition Mechanisms for IPv6 Hosts and Routers*	*http://www.ietf.org/rfc/rfc2893.txt*
January 2001	*3053*	*IPv6 Tunnel Broker*	*http://www.ietf.org/rfc/rfc3053.txt*
February 2001	*3056*	*Connection of IPv6 Domains via IPv4 Clouds*	*http://www.ietf.org/rfc/rfc3056.txt*
October 2005	*4213*	*Basic Transition Mechanisms for IPv6 Hosts and Routers*	*http://tools.ietf.org/html/rfc4213*
October 2005	*4213*	*Basic Transition Mechanisms for IPv6 Hosts and Routers*	*http://tools.ietf.org/html/rfc4213*
February 2006	*4380*	*Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)*	*http://www.ietf.org/rfc/rfc4380.txt*
September 2006	*4659*	*BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*	*http://www.ietf.org/rfc/rfc4659.txt*
February 2007	*4798*	*Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*	*http://tools.ietf.org/html/rfc4798*

Annex B:
IPv4/IPv6 Transition Standards

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
RFC 2460 [i.57]	**IPv6 architecture**	IP version 6 (IPv6) is a new version of the Internet Protocol designed as the successor to IP version 4 (IPv4). The changes from IPv4 to IPv6 fall primarily into the following categories: Expanded Addressing Capabilities IPv6 increases the IP address size from 32 bits to 128 bits, to support more levels of addressing hierarchy, a much greater number of addressable nodes, and simpler auto-configuration of addresses. The scalability of multicast routing is improved by adding a "scope" field to multicast addresses. And a new type of address called an "anycast address" is defined, used to send a packet to any one of a group of nodes. Header Format Simplification Some IPv4 header fields have been dropped or made optional, to reduce the common-case processing cost of packet handling and to limit the bandwidth cost of the IPv6 header. Improved Support for Extensions and Options Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing new options in the future. Flow Labeling Capability A new capability is added to enable the labeling of packets belonging to particular traffic "flows" for which the sender requests special handling, such as non-default quality of service or "real-time" service.
RFC 4193 [i.58]	**ULA**	This RFC defines an IPv6 unicast address format that is globally unique and is intended for local communications. These addresses are called Unique Local IPv6 Unicast Addresses and are abbreviated in this document as Local IPv6 addresses. They are not expected to be routable on the global Internet. They are routable inside of a more limited area such as a site. They may also be routed between a limited set of sites. Local IPv6 unicast addresses have the following characteristics: - Globally unique prefix (with high probability of uniqueness). - Well-known prefix to allow for easy filtering at site boundaries. - Allow sites to be combined or privately interconnected without creating any address conflicts or requiring renumbering of interfaces that use these prefixes. - Internet Service Provider independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity. If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses. In practice, applications may treat these addresses like global scoped addresses. This document defines the format of Local IPv6 addresses, how to allocate them, and usage considerations including routing, site border routers, DNS, application support, VPN usage, and guidelines for how to use for local communication inside a site.
RFC 5942 [i.59]	**IPv6 subnet model**	IPv4 implementations typically associate a netmask with an address when an IPv4 address is assigned to an interface. That netmask together with the IPv4 address designates an on-link prefix. Nodes consider addresses covered by an on-link prefix to be directly attached to the same link as the sending node, i.e. they send traffic for such addresses directly rather than to a router. Prior to the development of subnetting and Classless Inter-Domain Routing (CIDR) an address's netmask could be derived directly from the address simply by determining whether it was a Class A, B, or C address. Today, assigning an address to an interface also requires specifying a netmask to use. In the absence of specifying a specific netmask when assigning an address, some implementations would fall back to deriving the netmask from the class of the address. The behavior of IPv6 as specified in Neighbor Discovery (ND) is quite different. The on-link determination is separate from the address assignment. A host can have IPv6 addresses without any related on-link prefixes or can have on-link prefixes that are not related to any IPv6 addresses that are assigned to the host. Any assigned address on an interface should initially be considered as having no internal structure In IPv6, by default, a host treats only the link-local prefix as on-link.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		The reception of a Prefix Information Option (PIO) with the L-bit set and a non-zero valid lifetime creates (or updates) an entry in the Prefix List. All prefixes on a host's Prefix List (i.e. those prefixes that have not yet timed out) are considered to be on-link by that host. The on-link definition modified by this document defines the complete list of cases in which a host considers an address to be on-link. Individual address entries can be expired by the Neighbor Unreachability Detection mechanism. IPv6 packets sent using the Conceptual Sending Algorithm only trigger address resolution for IPv6 addresses that the sender considers to be on-link. Packets to any other address are sent to a default router. If there is no default router, then the node should send an ICMPv6 Destination Unreachable indication. In the old version of Neighbor Discovery, if the Default Router List is empty, rather than sending the ICMPv6 Destination Unreachable indication, the node assumed that the destination was on-link. Note that ND is scoped to a single link. All Neighbor Solicitation (NS) responses are assumed to be sent out the same interface on which the corresponding query was received without using the Conceptual Sending Algorithm. Failure of host implementations to correctly implement the IPv6 subnet model can result in lack of IPv6 connectivity.
RFC 4861 [i.34]	**Neighbor discovery**	This specification defines the Neighbor Discovery (ND) protocol for Internet Protocol Version 6 (IPv6). Nodes (hosts and routers) use Neighbor Discovery to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid. Hosts also use Neighbor Discovery to find neighboring routers that are willing to forward packets on their behalf. Finally, nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses. When a router or the path to a router fails, a host actively searches for functioning alternates. Unless specified otherwise (in a document that covers operating IP over a particular link type) this document applies to all link types. However, because ND uses link-layer multicast for some of its services, it is possible that on some link types (e.g. Non-Broadcast Multi-Access (NBMA) links), alternative protocols or mechanisms to implement those services will be specified (in the appropriate document covering the operation of IP over a particular link type). The services described in this document that are not directly dependent on multicast, such as Redirects, Next-hop determination, Neighbor Unreachability Detection, etc., are expected to be provided as specified in this document.
RFC 4862 [i.35]	**SLAAC**	This document specifies the steps a host takes in deciding how to autoconfigure its interfaces in IP version 6 (IPv6). The auto-configuration process includes generating a link-local address, generating global addresses via stateless address auto-configuration, and the Duplicate Address Detection procedure to verify the uniqueness of the addresses on a link. The IPv6 stateless auto-configuration mechanism requires no manual configuration of hosts, minimal (if any) configuration of routers, and no additional servers. The stateless mechanism allows a host to generate its own addresses using a combination of locally available information and information advertised by routers. Routers advertise prefixes that identify the subnet(s) associated with a link, while hosts generate an "interface identifier" that uniquely identifies an interface on a subnet. An address is formed by combining the two. In the absence of routers, a host can only generate link-local addresses. However, link-local addresses are sufficient for allowing communication among nodes attached to the same link. The stateless approach is used when a site is not particularly concerned with the exact addresses hosts use, so long as they are unique and properly routable. On the other hand, Dynamic Host Configuration Protocol for IPv6 (DHCPv6) (RFC 3315 [i.68]) is used when a site requires tighter control over exact address assignments. Both stateless address auto-configuration and DHCPv6 may be used simultaneously. IPv6 addresses are leased to an interface for a fixed (possibly infinite) length of time. Each address has an associated lifetime that indicates how long the address is bound to an interface. When a lifetime expires, the binding (and address) become invalid and the address may be reassigned to another interface elsewhere in the Internet. To handle the expiration of address bindings gracefully, an address goes through two distinct phases while assigned to an interface. Initially, an address is "preferred", meaning that its use in arbitrary communication is unrestricted. Later, an address becomes "deprecated" in anticipation that its current interface binding will become invalid. While an address is in a deprecated state, its use is discouraged, but not strictly forbidden. New communication (e.g. the opening of a new TCP connection) should use a preferred address when possible. A deprecated address should be used only by applications that have been using it and would have difficulty switching to another address without a service disruption.
RFC 3971 [i.60]	**SEND**	IPv6 defines the Neighbor Discovery Protocol (NDP). Nodes on the same link use NDP to discover each other's presence and link-layer addresses, to find routers, and to maintain reachability information about the paths to active neighbors. NDP is used by both hosts and routers. Its functions include Neighbor Discovery (ND), Router Discovery (RD), Address Auto configuration, Address Resolution, Neighbor Unreachability Detection (NUD), Duplicate Address Detection (DAD), and Redirection.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		The original NDP specifications called for the use of IPsec to protect NDP messages. However, the RFCs do not give detailed instructions for using IPsec to do this. In this particular application, IPsec can only be used with a manual configuration of security associations, due to bootstrapping problems in using IKE. Furthermore, the number of manually configured security associations needed for protecting NDP can be very large, making that approach impractical for most purposes. The SEND protocol is designed to counter the threats to NDP. SEND is applicable in environments where physical security on the link is not assured (such as over wireless) and attacks on NDP are a concern. This document is organized as follows. Sections 2 and 3 define some terminology and present a brief review of NDP, respectively. Section 4 describes the overall approach to securing NDP. This approach involves the use of new NDP options to carry public key - based signatures. A zero-configuration mechanism is used for showing address ownership on individual nodes; routers are certified by a trust anchor. The formats, procedures, and cryptographic mechanisms for the zero-configuration mechanism are described in a related specification. The required new NDP options are discussed in section 5. Section 6 describes the mechanism for distributing certification paths to establish an authorization delegation chain to a trust anchor. Finally, section 8 discusses the co-existence of secured and unsecured NDP on the same link, and section 9 discusses security considerations for SEcure Neighbor Discovery (SEND). The use of identity certificates provisioned on end hosts for authorizing address use is out of the scope for this document, as is the security of NDP when the entity defending an address is not the same as the entity claiming that address (also known as "proxy ND"). These are extensions of SEND that may be treated in separate documents, should the need arise.
RFC 4443 [i.61]	**ICMPv6**	The Internet Protocol version 6 (IPv6) uses the Internet Control Message Protocol (ICMP) as defined for IPv4, with a number of changes. The resulting protocol is called ICMPv6 and has an IPv6 Next Header value of 58. This document describes the format of a set of control messages used in ICMPv6. It does not describe the procedures for using these messages to achieve functions like Path MTU discovery; these procedures are described in other documents. Other documents may also introduce additional ICMPv6 message types, such as Neighbor Discovery messages, subject to the general rules for ICMPv6 messages given in section 2 of this document. Terminology defined in the IPv6 specification and the IPv6 Routing and Addressing specification applies to this document as well.
RFC 4605 [i.62]	**IGMP/MLD proxy**	This document applies spanning tree multicast routing to an Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD)-only environment. The topology is limited to a tree, since no protocol to build a spanning tree over a more complex topology is specified. The root of the tree is assumed to be connected to a wider multicast infrastructure. In a simple tree topology, it is not necessary to run a multicast routing protocol. It is sufficient to learn and proxy group membership information and simply forward multicast packets based upon that information. One typical example of such a tree topology can be found on an edge aggregation box such as a Digital Subscriber Line Access Multiplexer (DSLAM). In most deployment scenarios, an edge box has only one connection to the core network side and has many connections to the customer side. Using IGMP/MLD-based forwarding to replicate multicast traffic on devices such as the edge boxes can greatly simplify the design and implementation of those devices. By not supporting more complicated multicast routing protocols such as Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP), it reduces not only the cost of the devices but also the operational overhead. Another advantage is that it makes the proxy devices independent of the multicast routing protocol used by the core network routers. Hence, proxy devices can be easily deployed in any multicast network. Robustness in an edge box is usually achieved by using a hot spare connection to the core network. When the first connection fails, the edge box fails over to the second connection. IGMP/MLD-based forwarding can benefit from such a mechanism and use the spare connection for its redundant or backup connection to multicast routers. When an edge box fails over to the second connection, the proxy upstream connection can also be updated to the new connection.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
RFC 4632 [i.63]	**CIDR**	What is now known as the Internet started as a research project in the 1970s to design and develop a set of protocols that could be used with many different network technologies to provide a seamless, end- to-end facility for interconnecting a diverse set of end systems. When it was determined how the 32-bit address space would be used, certain assumptions were made about the number of organizations to be connected, the number of end systems per organization, and total number of end systems on the network. The end result was the establishment of three classes of networks: Class A (most significant address bits '00'), with 128 possible networks each and 16 777 216 end systems (minus special bit values reserved for network/broadcast addresses); Class B (MSB '10'), with 16 384 possible networks each with 65 536 end systems (less reserved values); and Class C (MSB '110'), and 2097152 possible networks each and 254 end systems (256 bit combinations minus the reserved all-zeros and all-ones patterns). The set of addresses with MSB '111' was reserved for future use; parts of this were eventually defined (MSB '1110') for use with IPv4 multicast and parts are still reserved as of the writing of this document. In the late 1980s, the expansion and commercialization of the former research network resulted in the connection of many new organizations to the rapidly growing Internet, and each new organization required an address assignment according to the Class A/B/C addressing plan. As demand for new network numbers (particularly in the Class B space) took what appeared to be an exponential growth rate, some members of the operations and engineering community started to have concerns over the long-term scaling properties of the class A/B/C system and began thinking about how to modify network number assignment policy and routing protocols to accommodate the growth. In November 1991, the Internet Engineering Task Force (IETF) created the ROAD (Routing and Addressing) group to examine the situation. This group met in January 1992 and identified three major problems: Exhaustion of the Class B network address space. One fundamental cause of this problem is the lack of a network class of a size that is appropriate for mid-sized organization. Class C, with a maximum of 254 host addresses, is too small, whereas Class B, which allows up to 65 534 host addresses, is too large for most organizations but was the best fit available for use with subnetting. Growth of routing tables in Internet routers beyond the ability of current software, hardware, and people to effectively manage. Eventual exhaustion of the 32-bit IPv4 address space. It was clear that then-current rates of Internet growth would cause the first two problems to become critical sometime between 1993 and 1995. Work already in progress on topological assignment of addressing for Connectionless Network Service (CLNS), which was presented to the community at the Boulder IETF in December of 1990, led to thoughts on how to re-structure the 32-bit IPv4 address space to increase its lifespan. The design and deployment of CIDR was intended to solve these problems by providing a mechanism to slow the growth of global routing tables and to reduce the rate of consumption of IPv4 address space. It did not and does not attempt to solve the third problem, which is of a more long-term nature; instead, it endeavors to ease enough of the short- to mid-term difficulties to allow the Internet to continue to function efficiently while progress is made on a longer-term solution. The solution that the community created was to deprecate the Class A/B/C network address assignment system in favor of using "classless", hierarchical blocks of IP addresses (referred to as prefixes). The assignment of prefixes is intended to roughly follow the underlying Internet topology so that aggregation can be used to facilitate scaling of the global routing system. One implication of this strategy is that prefix assignment and aggregation is generally done according to provider-subscriber relationships, since that is how the Internet topology is determined. This addressing plan was intended to be a relatively short-term response, lasting approximately three to five years, during which a more permanent addressing and routing architecture would be designed and implemented. As can be inferred from the dates on the original documents, CIDR has far outlasted its anticipated lifespan and has become the mid-term solution to the problems described above. Note that in the following text a description of the current policies and procedures that have been put in place to implement the allocation architecture discussed here. This description is not intended to be interpreted as direction to IANA. Coupled with address management strategies implemented by the Regional Internet Registries, the deployment of CIDR-style addressing has also reduced the rate at which IPv4 address space has been consumed, thus providing short- to medium-term relief to problem #3, described above.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		Note that, as defined, this plan neither requires nor assumes the re-assignment of those parts of the legacy "Class C" space that are not amenable to aggregation (sometimes called "the swamp"). Doing so would somewhat reduce routing table sizes (current estimate is that "the swamp" contains approximately 15,000 entries), though at a significant renumbering cost. Similarly, there is no hard requirement that any end site renumber when changing transit service provider, but end sites are encouraged do so to eliminate the need for explicit advertisement of their prefixes into the global routing system.
RFC 5308 [i.64]	**IS-IS routing IPv6**	IS-IS is an extendible intra-domain routing protocol. Each router in the routing domain issues an Link State Protocol Data Unit (LSP) that contains information pertaining to that router. The LSP contains typed variable-length data, often referred to as TLVs (type-length-values). The protocol is extended with two new TLVs to carry information required to perform IPv6 routing. A method to route both OSI and IPv4 is used with some minor changes to allow for IPv6. To do so, two new TLVs, namely "IPv6 Reachability" and "IPv6 Interface Address", and a new IPv6 protocol identifier should be defined.
RFC 2740 [i.65]	**OSPFv3**	This document describes the modifications to OSPF to support version 6 of the Internet Protocol (IPv6). The fundamental mechanisms of OSPF (flooding, DR election, area support, SPF calculations, etc.) remain unchanged. However, some changes have been necessary, either due to changes in protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6. This document is organized as follows. Section 2 describes the differences between OSPF for IPv4 and OSPF for IPv6 in detail. Section 3 provides implementation details for the changes. Appendix A gives the OSPF for IPv6 packet and LSA formats. Appendix B lists the OSPF architectural constants. Appendix C describes configuration parameters.
RFC 4798 [i.66]	**6PE**	There are several approaches for providing IPv6 connectivity over an MPLS core network including (i) requiring that MPLS networks support setting up IPv6-signaled Label Switched Paths (LSPs) and establish IPv6 connectivity by using those LSPs, (ii) use configured tunneling over IPv4-signaled LSPs, or (iii) use the IPv6 Provider Edge (6PE) approach defined in this document. The 6PE approach is required as an alternative to the use of standard tunnels. It provides a solution for an MPLS environment where all tunnels are established dynamically, thereby addressing environments where the effort to configure and maintain explicitly configured tunnels is not acceptable. This document specifies operations of the 6PE approach for interconnection of IPv6 islands over an IPv4 MPLS cloud. The approach requires that the edge routers connected to IPv6 islands be Dual Stack Multiprotocol-BGP-speaking routers, while the core routers are only required to run IPv4 MPLS. The approach uses MP-BGP over IPv4, relies on identification of the 6PE routers by their IPv4 address, and uses IPv4-signaled MPLS LSPs that do not require any explicit tunnel configuration. In this document an 'IPv6 island' is a network running native IPv6. A typical example of an IPv6 island would be a customer's IPv6 site connected via its IPv6 Customer Edge (CE) router to one (or more) Dual Stack Provider Edge router(s) of a Service Provider. These IPv6 Provider Edge routers (6PE) are connected to an IPv4 MPLS core network. The interconnection method described in this document typically applies to an Internet Service Provider (ISP) that has an IPv4 MPLS network, that is familiar with BGP (possibly already offering BGP/MPLS VPN services), and that wants to offer IPv6 services to some of its customers. However, the ISP may not (yet) want to upgrade its network core to IPv6, nor use only IPv6-over-IPv4 tunneling. With the 6PE approach described here, the provider only has to upgrade some Provider Edge (PE) routers to Dual Stack operations so that they behave as 6PE routers (and route reflectors if those are used for the exchange of IPv6 reachability among 6PE routers) while leaving the IPv4 MPLS core routers untouched. These 6PE routers provide connectivity to IPv6 islands. They may also provide other services simultaneously (IPv4 connectivity, IPv4 L3VPN services, L2VPN services, etc.). Also with the 6PE approach, no tunnels need to be explicitly configured, and no IPv4 headers need to be inserted in front of the IPv6 packets between the customer and provider edge. The ISP obtains IPv6 connectivity to its peers and upstreams using means outside of the scope of this document, and its 6PE routers readvertise it over the IPv4 MPLS core with MP-BGP. The interface between the edge router of the IPv6 island (Customer Edge (CE) router) and the 6PE router is a native IPv6 interface which can be physical or logical. A routing protocol (IGP or EGP) may run between the CE router and the 6PE router for the distribution of IPv6 reachability information. Alternatively, static routes and/or a default route may be used on the 6PE router and the CE router to control reachability. An IPv6 island may connect to the provider network over more than one interface. The 6PE approach described in this document can be used for customers that already have an IPv4 service from the network provider and additionally require an IPv6 service, as well as for customers that require only IPv6 connectivity.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		Note that the 6PE approach specified in this document provides global IPv6 reachability. Support of IPv6 VPNs is not within the scope of this document. Deployment of the 6PE approach over an existing IPv4 MPLS cloud does not require an introduction of new mechanisms in the core (other than potentially those described at the end of section 3 for dealing with dynamic MTU discovery). Configuration and operations of the 6PE approach have a lot of similarities with the configuration and operations of an IPv4 VPN service or IPv6 VPN service over an IPv4 MPLS core because they all use MP-BGP to distribute non-IPv4 reachability information for transport over an IPv4 MPLS Core. However, the configuration and operations of the 6PE approach is somewhat simpler, since it does not involve all the VPN concepts such as Virtual Routing and Forwarding (VRFs) tables.
RFC 4659 [i.67]	**6VPE**	This document describes a method by which a Service Provider may use its packet-switched backbone to provide Virtual Private Network services for its IPv6 customers. This method reuses, and extends where necessary, the "BGP/MPLS IP VPN" method for support of IPv6. In particular, this method uses the same "peer model", in which the customers' edge routers ("CE routers") send their IPv6 routes to the Service Provider's edge routers ("PE routers"). BGP ("Border Gateway Protocol") is then used by the Service Provider to exchange the routes of a particular IPv6 VPN among the PE routers that are attached to that IPv6 VPN. Eventually, the PE routers distribute, to the CE routers in a particular VPN, the IPv6 routes from other CE routers in that VPN. As with IPv4 VPNs, a key characteristic of this "peer model" is that the (IPv6) CE routers within an (IPv6) VPN do not peer with each other; there is no "overlay" visible to the (IPv6) VPN's routing algorithm. A VPN is said to be an IPv6 VPN when each site of this VPN is IPv6 capable and is natively connected over an IPv6 interface or sub- interface to the Service Provider (SP) backbone via a Provider Edge device (PE). A site may be both IPv4 capable and IPv6 capable. The logical interface on which packets arrive at the PE may determine the IP version. Alternatively, the same logical interface may be used for both IPv4 and IPv6, in which case a per-packet lookup at the Version field of the IP packet header determines the IP version. This document only concerns itself with handling of IPv6 communication between IPv6 hosts located on IPv6-capable sites. Handling of IPv4 communication between IPv4 hosts located on IPv4-capable sites is outside the scope of this document. Communication between an IPv4 host located in an IPv4- capable site and an IPv6 host located in an IPv6-capable site is outside the scope of this document. In a similar manner to how IPv4 VPN routes are distributed, BGP and its extensions are used to distribute routes from an IPv6 VPN site to all the other PE routers connected to a site of the same IPv6 VPN. PEs use "VPN Routing and Forwarding tables" (VRFs) to maintain the reachability information and forwarding information of each IPv6 VPN separately. As is done for IPv4 VPNs, each IPv6 VPN is allowed to have its own IPv6 address space, which means that a given address may denote different systems in different VPNs. This is achieved via a new address family, the VPN-IPv6 Address Family, in a fashion similar to that of the VPN-IPv4 address family which prepends a Route Distinguisher to the IP address. In addition to its operation over MPLS Label Switched Paths (LSPs), the IPv4 BGP/MPLS VPN solution has been extended to allow operation over other tunneling techniques, including GRE tunnels, IP-in-IP tunnels, L2TPv3 tunnels, and IPsec protected tunnels. In a similar manner, this document allows support of an IPv6 VPN service over MPLS LSPs, as well as over other tunneling techniques. This document allows support for an IPv6 VPN service over an IPv4 backbone, as well as over an IPv6 backbone. The IPv6 VPN service supported is identical in both cases. The IPv6 VPN solution defined in this document offers the following benefits: From both the Service Provider perspective and the customer perspective, the VPN service that can be supported for IPv6 sites is identical to the one that can be supported for IPv4 sites. From the Service Provider perspective, operations of the IPv6 VPN service require the exact same skills, procedures, and mechanisms as those for the IPv4 VPN service. Where both IPv4 VPNs and IPv6 VPN services are supported over an IPv4 core, the same single set of MP-BGP peering relationships and the same single PE-PE tunnel mesh MAY be used for both. The IPv6 VPN service is independent of whether the core runs IPv4 or IPv6. This is so that the IPv6 VPN service supported before and after a migration of the core from IPv4 to IPv6 is undistinguishable to the VPN customer. Note that supporting IPv4 VPN services over an IPv6 core is not covered by this document.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		The BGP Multiprotocol Extensions allow BGP to carry routes from multiple "address families". The notion of the "VPN-IPv6 address family", which is similar to the VPN-IPv4 address family is introduced. A VPN-IPv6 address is a 24-octet quantity, beginning with an 8-octet "Route Distinguisher" (RD) and ending with a 16-octet IPv6 address. The purpose of the RD is solely to allow one to create distinct routes to a common IPv6 address prefix, which is similar to the purpose of the RD. The RD can be used to create multiple different routes to the very same system. This can be achieved by creating two different VPN-IPv6 routes that have the same IPv6 part but different RDs. This allows the provider's BGP to install multiple different routes to the same system and allows policy to be used to decide which packets use which route.
RFC 3315 [i.68]	**DHCPv6**	This document describes DHCP for IPv6 (DHCP), a client/server protocol that provides managed configuration of devices. DHCP can provide a device with addresses assigned by a DHCP server and other configuration information, which are carried in options. DHCP can be extended through the definition of new options to carry configuration information not specified in this document. DHCP is the "stateful address auto-configuration protocol" and the "stateful auto-configuration protocol" referred to in "IPv6 Stateless Address Auto configuration". The operational models and relevant configuration information for DHCPv4 and DHCPv6 are sufficiently different that integration between the two services is not included in this document. If there is sufficient interest and demand, integration can be specified in a document that extends DHCPv6 to carry IPv4 addresses and configuration information. The remainder of this introduction summarizes DHCP, explaining the message exchange mechanisms and example message flows. The message flows in sections 1.2 and 1.3 are intended as illustrations of DHCP operation rather than an exhaustive list of all possible client-server interactions. Sections 17, 18, and 19 explain client and server operation in detail. Clients and servers exchange DHCP messages using UDP. The client uses a link-local address or addresses determined through other mechanisms for transmitting and receiving DHCP messages. DHCP servers receive messages from clients using a reserved, link-scoped multicast address. A DHCP client transmits most messages to this reserved multicast address, so that the client need not be configured with the address or addresses of DHCP servers. To allow a DHCP client to send a message to a DHCP server that is not attached to the same link, a DHCP relay agent on the client's link will relay messages between the client and server. The operation of the relay agent is transparent to the client and the discussion of message exchanges in the remainder of this section will omit the description of message relaying by relay agents. Once the client has determined the address of a server, it may under some circumstances send messages directly to the server using unicast.
RFC 3633 [i.30]	**DHCPv6 prefix delegation**	This document describes new options for Dynamic Host Configuration Protocol (DHCP) that provide a mechanism for the delegation of IPv6 prefixes. Through these options, a delegating router can delegate prefixes to authorized requesting routers. The prefix delegation mechanism described in this document is intended for simple delegation of prefixes from a delegating router to requesting routers. It is appropriate for situations in which the delegating router does not have knowledge about the topology of the networks to which the requesting router is attached, and the delegating router does not require other information aside from the identity of the requesting router to choose a prefix for delegation. For example, these options would be used by a service provider to assign a prefix to a Customer Premise Equipment (CPE) device acting as a router between the subscriber's internal network and the service provider's core network. Many applications expect stable addresses. Even though this mechanism makes automatic renumbering easier, it is expected that prefixes have a long lifespan. During renumbering it is expected that the old and the new prefix co-exist for some time. The design of this prefix delegation mechanism meets the requirements for prefix delegation in Requirements for IPv6 prefix delegation. Note that this use of DHCP is not bound to the assignment of IP addresses or other configuration information to hosts, and that no mechanism is currently available to communicate delegated prefixes to a DHCP server that serves such a function. This may be an item of future work, should usage warrant.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		This document describes new DHCPv6 options for IPv6 prefix delegation. This document should be read in conjunction with the DHCPv6 specification for a complete specification of the Prefix Delegation options and mechanism.
RFC 3646 [i.69]	**DNS for IPv6**	The DNS Recursive Name Server option provides a list of one or more IPv6 addresses of DNS recursive name servers to which a client's DNS resolver MAY send DNS queries. The DNS servers are listed in the order of preference for use by the client resolver. The Domain Search List option specifies the domain search list the client is to use when resolving hostnames with DNS. This option does not apply to other name resolution mechanisms. The DNS Recursive Name Server option SHOULD NOT appear in any other than the following messages: Solicit, Advertise, Request, Renew, Rebind, Information-Request, and Reply. The Domain Search List option SHOULD NOT appear in any other than the following messages: Solicit, Advertise, Request, Renew, Rebind, Information-Request, and Reply. The DNS Recursive Name Server option may be used by an intruder DHCP server to cause DHCP clients to send DNS queries to an intruder DNS recursive name server. The results of these misdirected DNS queries may be used to spoof DNS names. To avoid attacks through the DNS Recursive Name Server option, the DHCP client SHOULD require DHCP authentication before installing a list of DNS recursive name servers obtained through authenticated DHCP.
RFC 3736 [i.70]	**stateless DHCPv6**	This document assumes that a node using stateless DHCP configuration is not using DHCP for address assignment, and that a node has determined at least a link-local address. To obtain configuration parameters through stateless DHCP, a node uses the DHCP Information-request message. DHCP servers respond to the node's message with a Reply message that carries configuration parameters for the node. The Reply message from the server can carry configuration information, such as a list of DNS recursive name servers and SIP servers. This document does not apply to the function of DHCP relay agents. A network element can provide both DHCP server and DHCP relay service. For example, a network element can provide stateless DHCP service to hosts requesting stateless DHCP service, while relaying messages from hosts requesting address assignment through DHCP to another DHCP server. Several sections of the DHCP specification provide background information or define parts of the specification that are common to all implementations: 1-4: give an introduction to DHCP and an overview of DHCP message flows 5: defines constants used throughout the protocol specification 6, 7: illustrate the format of DHCP messages 8: describes the representation of Domain Names 9: defines the "DHCP unique identifier" (DUID) 13-16: describe DHCP message transmission, retransmission, and validation 21: describes authentication for DHCP The client indicates that it is requesting configuration information by sending an Information-request message that includes an Option Request option specifying the options that it wishes to receive from the DHCP server. For example, if the client is attempting to obtain a list of DNS recursive name servers, it identifies the DNS Recursive Name Server option in the Information-request message. The server determines the appropriate configuration parameters for the client based on its configuration policies and responds with a Reply message containing the requested parameters. In this example, the server would respond with DNS configuration parameters. A node may include a Client Identifier option in the Information-request message to identify itself to a server, because the server administrator may want to customize the server's response to each node, based on the node's identity.
RFC 4191 [i.71]	**default router preferences**	Neighbor Discovery specifies a conceptual model for hosts that includes a Default Router List and a Prefix List. Hosts send Router Solicitation messages and receive Router Advertisement messages from routers. Hosts populate their Default Router List and Prefix List based on information in the Router Advertisement messages. A conceptual sending algorithm uses the Prefix List to determine if a destination address is on-link and uses the Default Router List to select a router for off-link destinations. In some network topologies where the host has multiple routers on its Default Router List, the choice of router for an off-link destination is important. In some situations, one router may provide much better performance than another for a destination. In other situations, choosing the wrong router may result in a failure to communicate. This document describes an optional extension to Neighbor Discovery Router Advertisement messages for communicating default router preferences and more-specific routes from routers to hosts. This improves the ability of hosts to pick an appropriate router for an off-link destination.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		Note that since these procedures are applicable to hosts only, the forwarding algorithm used by the routers (including hosts with enabled IP forwarding) is not affected. Neighbor Discovery provides a Redirect message that routers can use to correct a host's choice of router. A router can send a Redirect message to a host, telling it to use a different router for a specific destination. However, the Redirect functionality is limited to a single link. A router on one link cannot redirect a host to a router on another link. Hence, Redirect messages do not help multi-homed (through multiple interfaces) hosts select an appropriate router. Multi-homed hosts are an increasingly important scenario, especially with IPv6. In addition to a wired network connection, like Ethernet, hosts may have one or more wireless connections, like 802.11 or Bluetooth. In addition to physical network connections, hosts may have virtual or tunnel network connections. For example, in addition to a direct connection to the public Internet, a host may have a tunnel into a private corporate network. Some IPv6 transition scenarios can add additional tunnels. For example, hosts may have 6to4 or configured tunnel network connections. This document requires that the preference values and specific routes advertised to hosts require explicit administrative configuration. They are not automatically derived from routing tables. In particular, the preference values are not routing metrics and it is not recommended that routers "dump out" their entire routing tables to hosts. Router Advertisement messages are used, instead of some other protocol like RIP, because Router Advertisements are an existing standard, stable protocol for router-to-host communication. Piggybacking this information on existing message traffic from routers to hosts reduces network overhead. Neighbor Discovery shares with Multicast Listener Discovery the property that they both define host-to-router interactions, while shielding the host from having to participate in more general router-to-router interactions. In addition, RIP is unsuitable because it does not carry route lifetimes so it requires frequent message traffic with greater processing overheads. The mechanisms specified here are backwards-compatible, so that hosts that do not implement them continue to function as well as they did previously.
RFC 4075 [i.72]	**sNTP**	The Simple Network Time Protocol servers option provides a list of one or more IPv6 addresses of SNTP servers available to the client for synchronization. The clients use these SNTP servers to synchronize their system time to that of the standard time servers. Clients SHOULD treat the list of SNTP servers as an ordered list. The server MAY list the SNTP servers in decreasing order of preference. The option defined in this document can only be used to configure information about SNTP servers that can be reached using IPv6. Mechanisms for configuring IPv4/IPv6 dual-stack applications are being considered, but are not specified in this document. The SNTP servers option SHOULD NOT appear in messages other than the following: Solicit, Advertise, Request, Renew, Rebind, Information-Request, and Reply. If this option appears in messages other than those specified above, the receiver SHOULD ignore it. The option number for this option MAY appear in the Option Request Option in the following messages: Solicit, Request, Renew, Rebind, Information-Request, and Reconfigure. If this option number appears in the Option Request Option in messages other than those specified above, the receiver SHOULD ignore it.
RFC 3596 [i.73]	**DNS IPv6**	Current support for the storage of Internet addresses in the Domain Name System (DNS) cannot easily be extended to support IPv6 addresses since applications assume that address queries return 32-bit IPv4 addresses only. To support the storage of IPv6 addresses in the DNS, this document defines the following extensions: - A resource record type is defined to map a domain name to an IPv6 address. - A domain is defined to support lookups based on address. - Existing queries that perform additional section processing to locate IPv4 addresses are redefined to perform additional section processing on both IPv4 and IPv6 addresses. The changes are designed to be compatible with existing software. The existing support for IPv4 addresses is retained. Transition issues related to the co-existence of both IPv4 and IPv6 addresses in the DNS are discussed in. The IP protocol version used for querying resource records is independent of the protocol version of the resource records; e.g. IPv4 transport can be used to query IPv6 records and vice versa. A record type is defined to store a host's IPv6 address. A host that has more than one IPv6 address should have more than one such record. The AAAA resource record type is a record specific to the Internet class that stores a single IPv6 address.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		The IANA assigned value of the type is 28 (decimal). A 128 bit IPv6 address is encoded in the data portion of an AAAA resource record in network byte order (high-order byte first). An AAAA query for a specified domain name in the Internet class returns all associated AAAA resource records in the answer section of a response. A type AAAA query does not trigger additional section processing. The textual representation of the data portion of the AAAA resource record used in a master database file is the textual representation of an IPv6 address. All existing query types that perform type A additional section processing, i.e. name server (NS), location of services (SRV) and mail exchange (MX) query types, should be redefined to perform both type A and type AAAA additional section processing. These definitions mean that a name server should add any relevant IPv4 addresses and any relevant IPv6 addresses available locally to the additional section of a response when processing any one of the above queries.
RFC 6434 [i.74]	**IPv6 node requirements**	This document defines common functionality required from both IPv6 hosts and routers. Many IPv6 nodes will implement optional or additional features, but this document collects and summarizes requirements from other published Standards Track documents in one place. This document tries to avoid discussion of protocol details and references RFCs for this purpose. This document is intended to be an applicability statement and to provide guidance as to which IPv6 specifications should be implemented in the general case and which specifications may be of interest to specific deployment scenarios. This document does not update any individual protocol document RFCs. Although this document points to different specifications, it should be noted that in many cases, the granularity of a particular requirement will be smaller than a single specification, as many specifications define multiple, independent pieces, some of which may not be mandatory. In addition, most specifications define both client and server behavior in the same specification, while many implementations will be focused on only one of those roles. This document defines a minimal level of requirement needed for a device to provide useful internet service and considers a broad range of device types and deployment scenarios. Because of the wide range of deployment scenarios, the minimal requirements specified in this document may not be sufficient for all deployment scenarios. It is perfectly reasonable (and indeed expected) for other profiles to define additional or stricter requirements appropriate for specific usage and deployment environments. For example, this document does not mandate that all clients support DHCP, but some deployment scenarios may deem it appropriate to make such a requirement. For example, government agencies in the USA have defined profiles for specialized requirements for IPv6 in target environments. As it is not always possible for an implementer to know the exact usage of IPv6 in a node, an overriding requirement for IPv6 nodes is that they should adhere to Jon Postel's Robustness Principle: "Be conservative in what you do, be liberal in what you accept from others".
RFC 2464 [i.75]	**IPv6 over Ethernet**	This document specifies the frame format for transmission of IPv6 packets and the method of forming IPv6 link-local addresses and statelessly autoconfigured addresses on Ethernet networks. It also specifies the content of the Source/Target Link-layer Address option used in Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement and Redirect messages when those messages are transmitted on an Ethernet. The default MTU size for IPv6 packets on an Ethernet is 1 500 octets. This size may be reduced by a Router Advertisement containing an MTU option which specifies a smaller MTU, or by manual configuration of each node. If a Router Advertisement received on an Ethernet interface has an MTU option specifying an MTU larger than 1 500, or larger than a manually configured value, that MTU option may be logged to system management but should be otherwise ignored. For purposes of this document, information received from DHCP is considered "manually configured" and the term Ethernet includes CSMA/CD and full-duplex subnetworks based on ISO/IEC 8802-3 [i.54], with various data rates. IPv6 packets are transmitted in standard Ethernet frames. The Ethernet header contains the Destination and Source Ethernet addresses and the Ethernet type code, which should contain the value 86DD hexadecimal. The data field contains the IPv6 header followed immediately by the payload, and possibly padding octets to meet the minimum frame size for the Ethernet link. The Interface Identifier for an Ethernet interface is based on the EUI-64 identifier derived from the interface's built-in 48-bit IEEE 802 address. The EUI-64 is formed as follows. (Canonical bit order is assumed throughout.) The OUI of the Ethernet address (the first three octets) becomes the company_id of the EUI-64 (the first three octets). The fourth and fifth octets of the EUI are set to the fixed value FFFE hexadecimal. The last three octets of the Ethernet address become the last three octets of the EUI-64.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		The Interface Identifier is then formed from the EUI-64 by complementing the "Universal/Local" (U/L) bit, which is the next-to-lowest order bit of the first octet of the EUI-64. Complementing this bit will generally change a 0 value to a 1, since an interface's built-in address is expected to be from a universally administered address space and hence have a globally unique value. A universally administered IEEE 802 address or an EUI-64 is signified by a 0 in the U/L bit position, while a globally unique IPv6 Interface Identifier is signified by a 1 in the corresponding position. A different MAC address set manually or by software should not be used to derive the Interface Identifier. If such a MAC address should be used, its global uniqueness property should be reflected in the value of the U/L bit. An IPv6 address prefix used for stateless auto-configuration of an Ethernet interface should have a length of 64 bits. The IPv6 link-local address for an Ethernet interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.
RFC 2827 [i.76]	**Network ingress filtering**	A resurgence of Denial of Service Attacks aimed at various targets in the Internet have produced new challenges within the Internet Service Provider (ISP) and network security communities to find new and innovative methods to mitigate these types of attacks. The difficulties in reaching this goal are numerous; some simple tools already exist to limit the effectiveness and scope of these attacks, but they have not been widely implemented. This method of attack has been known for some time. Defending against it, however, has been an ongoing concern. Bill Cheswick is quoted as saying that he pulled a chapter from his book, "Firewalls and Internet Security" at the last minute because there was no way for an administrator of the system under attack to effectively defend the system. By mentioning the method, he was concerned about encouraging its use. While the filtering method discussed in this document does absolutely nothing to protect against flooding attacks which originate from valid prefixes (IP addresses), it will prohibit an attacker within the originating network from launching an attack of this nature using forged source addresses that do not conform to ingress filtering rules. All providers of Internet connectivity are urged to implement filtering described in this document to prohibit attackers from using forged source addresses which do not reside within a range of legitimately advertised prefixes. In other words, if an ISP is aggregating routing announcements for multiple downstream networks, strict traffic filtering should be used to prohibit traffic which claims to have originated from outside of these aggregated announcements. An additional benefit of implementing this type of filtering is that it enables the originator to be easily traced to its true source, since the attacker would have to use a valid, and legitimately reachable, source address. Also worthy of mention is a case wherein the source address is forged to appear to have originated from within another legitimate network which appears in the global routing table(s). For example, an attacker using a valid network address could wreak havoc by making the attack appear to come from an organization which did not, in fact, originate the attack and was completely innocent. In such cases, the administrator of a system under attack may be inclined to filter all traffic coming from the apparent attack source. Adding such a filter would then result in a denial of service to legitimate, non-hostile end-systems. In this case, the administrator of the system under attack unwittingly becomes an accomplice of the attacker. Further complicating matters, TCP SYN flood attacks will result in SYN-ACK packets being sent to one or many hosts which have no involvement in the attack, but which become secondary victims. This allows the attacker to abuse two or more systems at once. Similar attacks have been attempted using UDP and ICMP flooding. The former attack (UDP flooding) uses forged packets to try and connect the changed UDP service to the echo UDP service at another site. Systems administrators should NEVER allow UDP packets destined for system diagnostic ports from outside of their administrative domain to reach their systems. The latter attack (ICMP flooding), uses an insidious feature in IP subnet broadcast replication mechanics. This attack relies on a router serving a large multi- access broadcast network to frame an IP broadcast address (such as one destined for 10.255.255.255) into a Layer 2 broadcast frame (for ethernet, FF:FF:FF:FF:FF:FF). Ethernet NIC hardware (MAC-layer hardware, specifically) will only listen to a select number of addresses in normal operation. The one MAC address that all devices share in common in normal operation is the media broadcast, or FF:FF:FF:FF:FF:FF. In this case, a device will take the packet and send an interrupt for processing. Thus, a flood of these broadcast frames will consume all available resources on an end-system. It is perhaps prudent that system administrators should consider ensuring that their border routers do not allow directed broadcast packets to be forwarded through their routers as a default. When an TCP SYN attack is launched using unreachable source address, the target host attempts to reserve resources waiting for a response. The attacker repeatedly changes the bogus source address on each new packet sent, thus exhausting additional host resources.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		Alternatively, if the attacker uses someone else's valid host address as the source address, the system under attack will send a large number of SYN/ACK packets to what it believes is the originator of the connection establishment sequence. In this fashion, the attacker does damage to two systems: the destination target system, as well as the system which is actually using the spoofed address in the global routing system. The result of both attack methods is extremely degraded performance, or worse, a system crash.
RFC 2460 [i.57]	**IPv6 Addressing Architecture**	This specification defines the addressing architecture of the IP Version 6 (IPv6) protocol. It includes the basic formats for the various types of IPv6 addresses (unicast, anycast, and multicast). IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces. There are three types of addresses: Unicast: An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. Anycast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance). Multicast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address. There are no broadcast addresses in IPv6, their function being superseded by multicast addresses. In this document, fields in addresses are given a specific name, for example "subnet". When this name is used with the term "ID" for identifier after the name (e.g. "subnet ID"), it refers to the contents of the named field. When it is used with the term "prefix" (e.g. "subnet prefix") it refers to all of the address from the left up to and including this field. In IPv6, all zeros and all ones are legal values for any field, unless specifically excluded. Specifically, prefixes may contain, or end with, zero-valued fields. IPv6 addresses of all types are assigned to interfaces, not nodes. An IPv6 unicast address refers to a single interface. Since each interface belongs to a single node, any of that node's interfaces' unicast addresses may be used as an identifier for the node. All interfaces are required to have at least one link-local unicast address. A single interface may also have multiple IPv6 addresses of any type (unicast, anycast, and multicast) or scope. Unicast addresses with scope greater than link-scope are not needed for interfaces that are not used as the origin or destination of any IPv6 packets to or from non-neighbors. This is sometimes convenient for point-to-point interfaces. There is one exception to this addressing model: A unicast address or a set of unicast addresses may be assigned to multiple physical interfaces if the implementation treats the multiple physical interfaces as one interface when presenting it to the internet layer. This is useful for load-sharing over multiple physical interfaces. Currently IPv6 continues the IPv4 model that a subnet prefix is associated with one link. Multiple subnet prefixes may be assigned to the same link.
RFC 2185 [i.55]	**Routing aspects of IPv6 Transition**	This document gives an overview of the routing aspects of IPv4 to IPv6 transition. The approach outlined here is designed to be compatible with the existing mechanisms for IPv6 transition. During an extended IPv4-to-IPv6 transition period, IPv6-based systems should coexist with the installed base of IPv4 systems. In such a dual internetworking protocol environment, both IPv4 and IPv6 routing infrastructure will be present. Initially, deployed IPv6-capable domains might not be globally interconnected via IPv6-capable internet infrastructure and therefore may need to communicate across IPv4-only routing regions. In order to achieve dynamic routing in such a mixed environment, there need to be mechanisms to globally distribute IPv6 network layer reachability information between dispersed IPv6 routing regions. The same techniques can be used in later stages of IPv4-to-IPv6 transition to route IPv4 packets between isolated IPv4-only routing region over IPv6 infrastructure. The IPng transition provides a dual-IP-layer transition, augmented by use of encapsulation where necessary and appropriate. Routing issues related to this transition include: (1) Routing for IPv4 packets (2) Routing for IPv6 packets (2a) IPv6 packets with IPv6-native addresses (2b) IPv6 packets with IPv4-compatible addresses

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		(3) Operation of manually configured static tunnels (4) Operation of automatic encapsulation (4a) Locating encapsulators (4b) Ensuring that routing is consist with encapsulation Basic mechanisms required to accomplish these goals include: (i) Dual-IP-layer Route Computation; (ii) Manual configuration of point- to-point tunnels; and (iii) Route leaking to support automatic encapsulation. The basic mechanism for routing of IPv4 and IPv6 involves dual-IP-layer routing. This implies that routes are separately calculated for IPv4 addresses and for IPv6 addressing. Tunnels (either IPv4 over IPv6, or IPv6 over IPv4) may be manually configured. For example, in the early stages of transition this may be used to allow two IPv6 domains to interact over an IPv4 infrastructure. Manually configured static tunnels are treated as if they were a normal data link. Use of automatic encapsulation, where the IPv4 tunnel endpoint address is determined from the IPv4 address embedded in the IPv4-compatible destination address of IPv6 packet, requires consistency of routes between IPv4 and IPv6 routing domains for destinations using IPv4-compatible addresses. For example, consider a packet which starts off as an IPv6 packet, but then is encapsulated in an IPv4 packet in the middle of its path from source to destination. This packet should locate an encapsulator at the correct part of its path. Also, this packet has to follow a consistent route for the entire path from source to destination.
RFC 3053 [i.56]	**IPv6 Tunnel Broker**	The growth of IPv6 networks started mainly using the transport facilities offered by the current Internet. This led to the development of several techniques to manage IPv6 over IPv4 tunnels. At present most of the 6bone network is built using manually configured tunnels over the Internet. The main drawback of this approach is the overwhelming management load for network administrators, who have to perform extensive manual configuration for each tunnel. Several attempts to reduce this management overhead have already been proposed and each of them presents interesting advantages but also solves different problems than the Tunnel Broker, or poses drawbacks not present in the Tunnel Broker: - the use of automatic tunnels with IPv4 compatible addresses is a simple mechanism to establish early IPv6 connectivity among isolated dual-stack hosts and/or routers. The problem with this approach is that it does not solve the address exhaustion problem of IPv4. Also there is a great fear to include the complete IPv4 routing table into the IPv6 world because this would worsen the routing table size problem multiplying it by 5; - 6over4 is a site local transition mechanism based on the use of IPv4 multicast as a virtual link layer. It does not solve the problem of connecting an isolated user to the global IPv6 Internet; - 6to4 has been designed to allow isolated IPv6 domains attached to a wide area network with no native IPv6 support (e.g. the IPv4 Internet), to communicate with other such IPv6 domains with minimal manual configuration. The idea is to embed IPv4 tunnel addresses into the IPv6 prefixes so that any domain border router can automatically discover tunnel endpoints for outbound IPv6 traffic. The Tunnel Broker idea is an alternative approach based on the provision of dedicated servers, called Tunnel Brokers, to automatically manage tunnel requests coming from the users. This approach is expected to be useful to stimulate the growth of IPv6 interconnected hosts and to allow early IPv6 network providers to provide easy access to their IPv6 networks. The main difference between the Tunnel Broker and the 6to4 mechanisms is that the they serve a different segment of the IPv6 community: - the Tunnel Broker fits well for small isolated IPv6 sites, and especially isolated IPv6 hosts on the IPv4 Internet, that want to easily connect to an existing IPv6 network; - the 6to4 approach has been designed to allow isolated IPv6 sites to easily connect together without having to wait for their IPv4 ISPs to deliver native IPv6 services. This is very well suited for extranet and virtual private networks. Using 6to4 relays, 6to4 sites can also reach sites on the IPv6 Internet. In addition, the Tunnel Broker approach allows IPv6 ISPs to easily perform access control on the users enforcing their own policies on network resources utilization. This document is intended to present a framework describing the guidelines for the provision of a Tunnel Broker service within the Internet. It does not specify any protocol but details the general architecture of the proposed approach. It also outlines a set of viable alternatives for implementing it. Section 2 provides an overall description of the Tunnel Broker model; section 3 reports known limitations to the model; section 4 briefly outlines other possible applications of the Tunnel Broker approach; section 5 addresses security issues.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		Tunnel brokers can be seen as virtual IPv6 ISPs, providing IPv6 connectivity to users already connected to the IPv4 Internet. In the emerging IPv6 Internet it is expected that many tunnel brokers will be available so that the user will just have to pick one. The list of the tunnel brokers should be referenced on a "well known" web page to allow users to choose the "closest" one, the "cheapest" one, or any other one. The TB is the place where the user connects to register and activate tunnels. The TB manages tunnel creation, modification and deletion on behalf of the user. For scalability reasons the tunnel broker can share the load of network side tunnel end-points among several tunnel servers. It sends configuration orders to the relevant tunnel server whenever a tunnel has to be created, modified or deleted. The TB may also register the user IPv6 address and name in the DNS. A TB should be IPv4 addressable. It may also be IPv6 addressable, but this is not mandatory. Communications between the broker and the servers can take place either with IPv4 or IPv6.
RFC 4213 [i.31]	**Basic Transition Mechanisms for IPv6 Hosts and Routers**	The key to a successful IPv6 transition is compatibility with the large installed base of IPv4 hosts and routers. Maintaining compatibility with IPv4 while deploying IPv6 will streamline the task of transitioning the Internet to IPv6. This specification defines two mechanisms that IPv6 hosts and routers may implement in order to be compatible with IPv4 hosts and routers. The mechanisms in this document are designed to be employed by IPv6 hosts and routers that need to interoperate with IPv4 hosts and utilize IPv4 routing infrastructures. It can be expected that most nodes in the Internet will need such compatibility for a long time to come, and perhaps even indefinitely. The mechanisms specified here are: - Dual IP layer (also known as dual stack): A technique for providing complete support for both Internet protocols -- IPv4 and IPv6 -- in hosts and routers. - Configured tunneling of IPv6 over IPv4: A technique for establishing point-to-point tunnels by encapsulating IPv6 packets within IPv4 headers to carry them over IPv4 routing infrastructures. The mechanisms defined here are intended to be the core of a "transition toolbox" -- a growing collection of techniques that implementations and users may employ to ease the transition. The tools may be used as needed. Implementations and sites decide which techniques are appropriate to their specific needs. This document defines the basic set of transition mechanisms, but these are not the only tools available. Additional transition and compatibility mechanisms are specified in other documents. Types of Nodes IPv4-only node: A host or router that implements only IPv4. An IPv4-only node does not understand IPv6. The installed base of IPv4 hosts and routers existing before the transition begins are IPv4-only nodes. IPv6/IPv4 node: A host or router that implements both IPv4 and IPv6. IPv6-only node: A host or router that implements IPv6 and does not implement IPv4. The operation of IPv6-only nodes is not addressed in this memo. IPv6 node: Any host or router that implements IPv6. IPv6/IPv4 and IPv6-only nodes are both IPv6 nodes. IPv4 node: Any host or router that implements IPv4. IPv6/IPv4 and IPv4-only nodes are both IPv4 nodes. Techniques Used in the Transition IPv6-over-IPv4 tunneling: The technique of encapsulating IPv6 packets within IPv4 so that they can be carried across IPv4 routing infrastructures. Configured tunneling: IPv6-over-IPv4 tunneling where the IPv4 tunnel endpoint address(es) are determined by configuration information on tunnel endpoints. All tunnels are assumed to be bidirectional. The tunnel provides a (virtual) point-to-point link to the IPv6 layer, using the configured IPv4 addresses as the lower-layer endpoint addresses. Other transition mechanisms, including other tunneling mechanisms, are outside the scope of this document.

RFC	RFC Name	RFC Description (partially generated from the RFC introductions)
		The most straightforward way for IPv6 nodes to remain compatible with IPv4-only nodes is by providing a complete IPv4 implementation. IPv6 nodes that provide complete IPv4 and IPv6 implementations are called "IPv6/IPv4 nodes". IPv6/IPv4 nodes have the ability to send and receive both IPv4 and IPv6 packets. They can directly interoperate with IPv4 nodes using IPv4 packets, and also directly interoperate with IPv6 nodes using IPv6 packets. Even though a node may be equipped to support both protocols, one or the other stack may be disabled for operational reasons. Here a rather loose notion of "stack" is used. A stack being enabled has IP addresses assigned, but whether or not any particular application is available on the stacks is explicitly not defined. Thus, IPv6/IPv4 nodes may be operated in one of three modes: - With their IPv4 stack enabled and their IPv6 stack disabled. - With their IPv6 stack enabled and their IPv4 stack disabled. - With both stacks enabled. IPv6/IPv4 nodes with their IPv6 stack disabled will operate like IPv4-only nodes. Similarly, IPv6/IPv4 nodes with their IPv4 stacks disabled will operate like IPv6-only nodes. IPv6/IPv4 nodes MAY provide a configuration switch to disable either their IPv4 or IPv6 stack. The configured tunneling technique, which is described in section 3, may or may not be used in addition to the dual IP layer operation. Because the nodes support both protocols, IPv6/IPv4 nodes may be configured with both IPv4 and IPv6 addresses. IPv6/IPv4 nodes use IPv4 mechanisms (e.g. DHCP) to acquire their IPv4 addresses, and IPv6 protocol mechanisms (e.g. stateless address auto-configuration and/or DHCPv6) to acquire their IPv6 addresses.

Annex C: List of Transition Technologies

C.1 Current Transition Technologies

The following currently available IPv4/IPv6 transition technologies are analyzed in clause 10:

- DS-Lite
- NAT64/DNS64
- Teredo
- dIVI

C.2 Future Transition Technologies

The following IPv4/IPv6 transition technologies that are currently under development are analyzed in clause 11:

- 464XLAT
- Stateless DS-Lite
- MAP-T
- MAP-E

C.3 Other Transition Technologies

The following IPv4/IPv6 transition technologies were also considered in the course of the analysis presented in the present document:

- 4in6
- 4rd
- 6in4
- 6over4
- 6rd
- 6to4
- AYIYA
- ISATAP
- IVI
- NAPT-PT
- NAT-PT
- SA46T
- SIIT

- TRT
- TSP

None of the above were analyzed in detail for one or more of the following reasons:

- The fundamental mechanisms defined by the transition technology are re-used in more advanced technologies.
- The connectivity issues that are addressed by the transition technology are also or better resolved by an alternative technologies.
- The functionality of the transition technology does not align with strategic goal of introducing and, finally, transitioning to IPv6.
- The transition technology lacks a mature technical description, implementation or market acceptance or has been deprecated from the standardization process.

History

Document history		
V1.1.1	December 2012	Publication