ETSI TS 103 443-2 V1.1.1 (2016-08)

TECHNICAL SPECIFICATION

Integrated broadband cable telecommunication networks (CABLE); IPv6 Transition Technology Engineering and Operational Aspects; Part 2: NAT64

Reference

DTS/CABLE-00018-2

Keywords

cable, 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

The present document can be downloaded from:
http://www.etsi.org/standards-search

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the only prevailing document is the print of the Portable Document Format (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
https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx

If you find errors in the present document, please send your comment to one of the following services:
https://portal.etsi.org/People/CommiteeSupportStaff.aspx

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 (https://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 Specification (TS) has been produced by ETSI Technical Committee Integrated broadband cable telecommunication networks (CABLE).

The present document is part 2 of a multi-part deliverable. Full details of the entire series can be found in part 1 [23].

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Introduction

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.

The present document accommodates an urgent need in the industry to implement and integrate the IPv6 transition technologies as specified by ETSI TS 101 569-1 [1] into their cable networks. The choice of the technology implemented depends on factors such as the business needs, current deployed architectures and plans for cost effectively transition from IPv4 to IPv6.

Current global IPv4 address space was projected to be depleted around the middle of 2012; depletion for the operator was estimated around end 2012. As part of the resulting roll-out of IPv6 in the operator's network, specific measures had to be taken to allow a smooth transition and coexistence between IPv4 and IPv6. ETSI developed requirements to address transition from IPv4 to IPv6 specifying six transition technologies as given by ETSI TS 101 569-1 [1] that were at the time considered to be the most appropriate to assist cable operators to transition there cable networks to IPv6.

Since then the industry has acquired more experience with the technology options settling in the main for DS-Lite across the cable network market and NAT64 IPv6 transition technologies across the mobile market.

The objective of the present document is to define the operational and engineering requirements to enable engineers to implement a seamless transition of the cable networks to IPv6 with the application of the 6RD transition technology.

The present document is the final part of a companion of ETSI standards developed in 4 phases to provide the cable sector in particular cable operators engineering and operational staff a standardized approach when integrating one of the five IPv6 transition technologies, NAT64, DS-Lite, 464XLAT, 6RD and MAP-E.

The first phase assessed the different IPv6 transition technology options being defined by industry with recommendation for the most appropriate with consideration of current network architectures, ensuring adequate scale and a cost effective transition approach from IPv4 to IPv6 as the IPv4 addresses deplete. The objective being to examine the pros and cons of the IPv6 transition technologies and recommend the most cost effective solution that would enable the cable operators to minimize the cost of upgrades to their existing network plant whilst maintain continuity of services to their present and new added customers. The details of the study are given by ETSI TR 101 569 [2].

In the second phase an ETSI technical specification was developed to specify technical requirements for six transition technologies that industry were considering for use by Cable Operators depending on the current state of their deployed cable network architecture, service model requirements and their IPv6 transition strategy as the IPv4 addresses depleted. These six IPv6 transition technologies are specified by ETSI TS 101 569-1 [1], covering NAT64, DSLite, 6RD, NAT44, 464XLAT and MAP-E.

In the third phase ETSI developed a series of conformance test specifications to enable the compliance verification of the five IPv6 transition technologies, NAT64, DS-Lite, 464XLAT, 6RD and MAP-E that were specified during phase 2 standardization. The conformance tests are developed against the requirements given by the ETSI TS 101 569-1 [1]. The series of conformance tests developed for each of the four transition technologies, are as given by, ETSI TS 103 238 parts 1 [3] to 3 [5] respectively for NAT64; ETSI TS 103 239 parts 1 [6] to 3 [8] respectively for MAP-E; ETSI TS 103 241 parts 1 [9] to 3 [11] respectively for DS-Lite; ETSI TS 103 242 parts 1 [12] to 3 [14] respectively for XLAT and ETSI TS 103 243 parts 1 [15] to 3 [17] respectively for 6RD.

Phase 4 is the present project phase for development of technical specifications covering the operational and engineering requirements with the present document being part 2 of a multi-part series covering the IPv6 transition technology NAT64.

DOCSIS® is a registered Trade Mark of Cable Television Laboratories, Inc., and is used in the present document with permission.

1 Scope

The present document presents the engineering and operational requirements for the application of the IPv6 transition technology NAT64 as defined by ETSI TS 101 569-1 [1] (IPv6 Transition Requirements) implemented within an integrated broadband cable network end to end across its network domains.

The present document is part 2 of a multi-part series and presents the operational aspects of the IPv6 transition technology NAT64 across the cable network domains.

Only those elements of the network that have to be engineered to operate the IPv6 transition technology NAT64 are presented. Descriptions and interface details of network elements that do not change are already addressed by the relevant equipment cable standards and therefore this information is not included in the present document.

The conformity of the NAT64 implementation is relevant when assessing its implementation and operational requirements across the cable network to ensure the implementation is correctly engineered to conform to the requirements of the base standard ETSI TS 101 569-1 [1]. These conformance tests are not specified in the present document as they are already specified by ETSI TS 103 243 parts 1 [15] to 3 [17].

The operational aspects for the IPv6 transition technology NAT64 are considered when engineered end to end across the cable network domains:

- CPE Home Networking Domain
- Access Network Domain
- Core Network Domain
- Data Centre Domain
- DMZ Service Domain
- Transit and Peering Domain
- Management and Monitoring Domain
- Security Domain

The present document specifies the requirements to be considered when the defined IPv6 transition technology NAT64 is engineered across the cable network domains.

2 References

2.1 Normative 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.

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

[1] ETSI TS 101 569-1: "Integrated Broadband Cable Telecommunication Networks (CABLE); Cable Network Transition to IPv6; Part 1: IPv6 Transition Requirements".

[2] ETSI TR 101 569: "Access, Terminals, Transmission and Multiplexing (ATTM); Integrated Broadband Cable and Television Networks; Cable Network Transition to IPv6".

[3] ETSI TS 103 238-1: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for NAT64 technology; Part 1: Protocol Implementation Conformance Statement (PICS) proforma".

[4] ETSI TS 103 238-2: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for NAT64 technology; Part 2: Test Suite Structure and Test Purposes (TSS&TP)".

[5] ETSI TS 103 238-3: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for NAT64 technology; Part 3: Abstract Test Suite (ATS) and Protocol Implementation eXtra Information for Testing (PIXIT)".

[6] ETSI TS 103 239-1: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for MAP-E technology; Part 1: Protocol Implementation Conformance Statement (PICS) proforma".

[7] ETSI TS 103 239-2: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for MAP-E technology; Part 2: Test Suite Structure and Test Purposes (TSS&TP)".

[8] ETSI TS 103 239-3: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for MAP-E technology; Part 3: Abstract Test Suite (ATS) and Protocol Implementation eXtra Information for Testing (PIXIT)".

[9] ETSI TS 103 241-1: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for DS-Lite technology; Part 1: Protocol Implementation Conformance Statement (PICS) proforma".

[10] ETSI TS 103 241-2: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for DS-Lite technology; Part 2: Test Suite Structure and Test Purposes (TSS&TP)".

[11] ETSI TS 103 241-3: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for DS-Lite technology; Part 3: Abstract Test Suite (ATS) and Protocol Implementation eXtra Information for Testing (PIXIT)".

[12] ETSI TS 103 242-1: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for 464XLAT technology; Part 1: Protocol Implementation Conformance Statement (PICS) proforma".

[13] ETSI TS 103 242-2: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for 464XLAT technology; Part 2:Test Suite Structure and Test Purposes (TSS&TP)".

[14] ETSI TS 103 242-3: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for 464XLAT technology; Part 3: Abstract Test Suite (ATS) and Protocol Implementation eXtra Information for Testing (PIXIT)".

[15] ETSI TS 103 243-1: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for 6rd technology; Part 1: Protocol Implementation Conformance Statement (PICS) proforma".

[16] ETSI TS 103 243-2: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for 6rd technology; Part 2: Test Suite Structure and Test Purposes (TSS&TP)".

[17] ETSI TS 103 243-3: "Integrated broadband cable telecommunication networks (CABLE); Testing; Conformance test specifications for 6rd technology; Part 3: Abstract Test Suite (ATS) and Protocol Implementation eXtra Information for Testing (PIXIT)".

[18] IETF RFC 6147 (April 2011): "DNS64: DNS Extension for Network Address Translation from IPv6 Clients to IPv4 Servers".

[19] IETF RFC 6146 (April 2011): "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers".

[20] IETF RFC 7269 (June 2014): "NAT64 Deployment Options and Experience".

[21] IETF RFC 6334 (August 2011): "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite".

[22] IETF RFC 6519 (February 2012): "RADIUS Extensions for Dual-Stack Lite".

[23] ETSI TS 103 443-1: "Integrated broadband cable telecommunication networks (CABLE); IPv6 Transition Technology Engineering and Operational Aspects; Part 1: General".

[24] IETF RFC 4787: "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP".

[25] IETF RFC 5382: "NAT Behavioral Requirements for TCP".

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

[27] IETF RFC 6145: "IP/ICMP Translation Algorithm".

2.2 Informative 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.

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

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] CableLabs.

NOTE: Available at http://www.cablelabs.com/specs/.

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

4in6: encapsulation of IPv4 packets within IPv6 packet format

NAT44: network address translation from an IPv4 address to another IPv4 address

P Router: label switching router acting as a transit router in the core network of an MPLS network

3.2 Abbreviations

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

6PE	IPv6 Provider Edge
6RD	IPv6 Rapid Deployment
6VPE	IPv6 VPN Provider Edge
AAA	Authentication, Authorization and Accounting
AAAA	Quad-A Resource Record
ALG	Application Layer Gateway
ALTS	Application-Level Transport Service
AMPS	Amplifiers
AS	Autonomous System

ASCII	American Standard Code for Information Interchange
ASIC	Application Specific Integrated Circuit
B4	Basic Bridging BroadBand element
BFD	Bidirectional Forwarding Detection
BGP	Boarder Gateway Protocol
BNG	Broadband Network Gateway
CDP	Cisco Discovery Protocol
CEF	Cisco Express Forwarding
CLAT	Customer-side TransLATor
CMTS	Cable Modem Termination System
CoPP	Control Plane Policing
CPE	Customer Premises Equipment
DAD	Duplicate Address Detection
dCEF	distributed Cisco Express Forwarding
DCU	Destination Class Usage
DHCP	Dynamic Host Configuration
DMZ	DeMilitarised Zone
DNS	Domain Name System
DR	Data Retention
DSCP	Differentiated Services Code Point
DS-Lite	Dual Stack-Lite
ECMP	Equal-Cost-Multi-Path
FTP	File Transfer Protocol
GRT	Global Routing Table
GW	Gateway
HA	High Availability
HA	High Availability
HFC	Hybrid Fibre Coax
HSRP	Hot Standby Router Protocol
ICMP	Internet Control Message Protocol
ID	IDentifier
IGP	Interior Gateway Protocol
IMIX	Internet MIX
IP	Internet Protocol
IPFIX	IP Flow Information Export PPTP Point-to-Point Tunnelling Protocol
IPv4	IP version 4
IPv6	IP version 6
IRB	Integrated Routing and Bridging
ISIS	Intermediate System to Intermediate System
ISSU	In-Service Software Upgrade
IXPE	Internet Exchange Provider Edge
L2	Layer 2
LDP	Label Distribution Protocol
LI	Lawful Intercept
LLDP	Link Layer Discovery Protocol
LP	Application-Level Proxy
LSN	Large Scale NAT
MAP-E	Mapping of Address and Port - Encapsulation mode
MFIB	Multicast Forwarding Information Base
MLD/L2	Multicast Listener Discovery/Layer 2
MP BGP	MultiProtocol Boarder Gateway Protocol
MP	MultiProtocol
MPLS	MultiProtocol Label Switching
MSO	Multiple-System Operator
MSS	Maximum Segment Size
MSTP	Multiple Spanning Tree Protocol
MT	Multi-Topology
MTU	Maximum Transmission Unit
MVNO	Mobile Virtual Network Operator
NAPT	Network Address and Port Translation
NAT	Network Address Translation
NAT44	Network Address Translation IPv4 to IPv4

NAT64	Network Address Translation IPv6 to IPv4
NDP	Neighbour Discovery Protocol
NFv9	NetFlow version 9
NPU	Network Processing Unit
NSF/GR	Non-Stop Forwarding/Graceful Restart
NTP	Network Time Protocol
NUD	Neighbour Unreachability Detection
OAM	Operation, Administration and Maintenance
PCP	Port Control Protocol
PE	Provider Edge
PIM	Protocol Independent Multicasting
PMTU	Path Maximum Transport Unit
PPTP	Point-to-Point Tunnelling Protocol
PS-BGP	Pretty Secure Boarder Gateway Protocol
QoS	Quality of Service
QPPB	QoS Policy Propagation via Boarder Gateway Protocol
RADIUS	Remote Authentication Dial-In User Service
RDT	Reliable Data Transfer
RP	Route Processor
RSTP	Rapid Spanning Tree Protocol
RTCP	Real-Time Transmission Control Protocol
RTP	Real-Time Protocol
RTSP	Real-Time Streaming Protocol
SCU	Source Class Usage
SEND	Secure Neighbour Discovery
SIP	Session Initiated Protocol
SNMP	Simple Network Management Protocol
SSH	Secure SHell
SSO	Stateful Switchover
SVI	Switched Virtual Interface
SYSLOG	Syslog Protocol
TACACS	Terminal Access Controller Access Control System
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
UDP	User Datagram Protocol
VLAN	Virtual LAN
VPLS	Virtual Protocol Local Area Network Service
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VRRP	Virtual Router Redundancy Protocol
XLAT	transLATor
XML	eXtensible Markup Language

4 General Considerations

4.1 Background

The present document is part of a multi-part series specifying requirements to engineer and operate the NAT64 transition technology end to end across a cable operator's network. Its implementation would ensure the network provider can continue to provide business continuity throughout the depletion of publicly routable IPv4 addresses and the subsequent rollout and migration to IPv6 in the operator's network.

To aid this transition some sectors of industry are currently deploying NAT64 in particular the mobile sector as a suitable transition technology whilst the cable sector are evaluating NAT64 deployments to mitigate the gap and lack of integration and compatibility between IPv4 and IPv6.

4.2 General Overview

The present document uses network address translation IPv6 to IPv4 (NAT64) technology to provide a seamless internet experience to users accessing IPv4 internet services from an IPv6 only client through a cable network enabling service providers to transparently deliver and enable new and existing services to IPv6 internet users with little or no change in their existing network infrastructure.

An IPv4 client cannot operate using NAT64.

It should be noted that Cable broadband access networks may vary in build and design with characteristics that may be vendor equipment specific. Consequently there may be aspects to the engineering and operation of the IPv6 transition technology NAT 64 that are dependent on the network build and vendor specific equipment deployed.

The present document does not offer information that may be vendor and network build specific since such information may be confidential to the network operator and/or based on proprietary data.

The present document assumes the reader is familiar with the cable network architecture requirements since the description of the various elements within a cable network across its domains are already defined by ETSI standards and standards developed by CableLabs [i.1]. The present document details only the changes to the network aspects when operating the transition technology NAT64.

The network elements required to implement the IPv6 transition technology NAT64 across the cable network domains is as illustrated by figure 1.

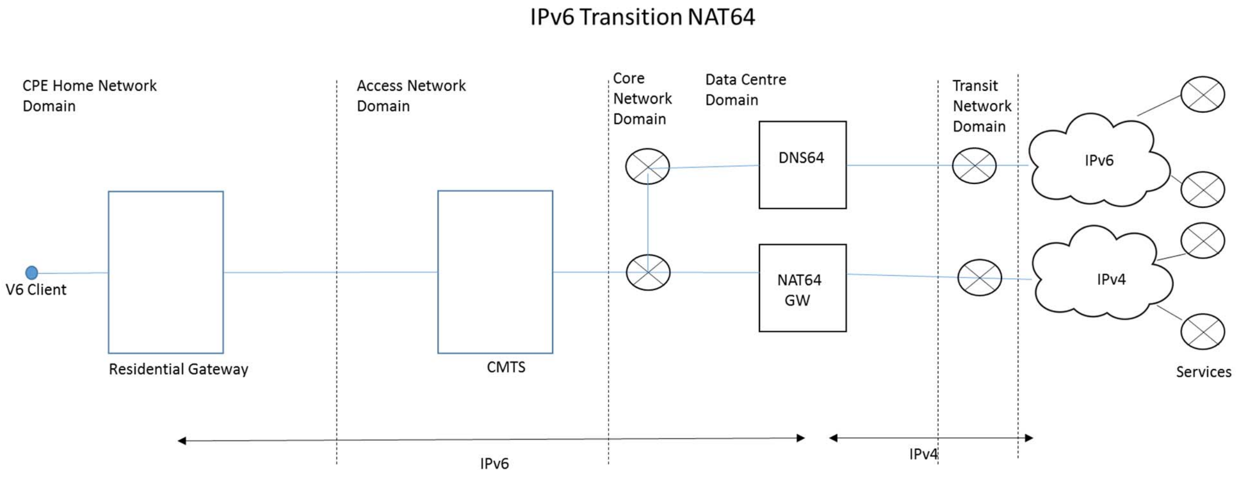

Figure 1: Illustration of network elements to support IPv6 transition technology NAT64 across Cable Network Domains

The specific aspects are given in the subsequent clauses for each network domain.

When engineering IPv6 transition technology IPv6 needs to be implemented on all of the network elements.

The engineered network elements to enable NAT64 in each cable network domain is integrated with existing network elements and shall be validated by network integration testing. The conformance of the implementation for NAT64 would need to be verified before operation as given by ETSI TS 103 238 parts 1 [3] to 3 [5].

4.3 Vendor Considerations

Generally there are many equipment manufacturers (vendors) in the market place delivering NAT64 as it is used heavily in the mobile market. However, its utilization in the MSO market has been limited even in the cable operators MVNO arena. NAT64 within a cable network requires the engineering of several ALGs to function and consequently it is considered to not be a technically viable proposal as an IPv6 service enabler. Each protocol requires more than a single stream or ports to function and requires consideration of the control plane to engineer a functioning ALG whether ALP or ALTS.

5 Gap Analysis

5.1 Consideration

The engineering and operational requirements applying ETSI TS 101 569-1 [1] need to be specified with the following design objective to define the logical and physical parameters to allow for IPv6 customers to access IPv4 services across the public IPv6 internet network using NAT64.

NAT64 is to be deployed such that it will provide a seamless experience to IPv6 users accessing IPv4 network services through new IPv6 only networks and to enable current and new content to be delivered seamlessly to IPv6 users.

Deployment of DNS64 is essential for actual internet usage and NAT64 will not function under normal requirements without it.

ALGs are an expensive consideration when using NAT64 as a transition technology as they require specific and extensive engineering to implement specific use cases that is considered undesirable from an economic approach and also regional networks would in general be dimensioned to require four or five ALGs to be implemented.

Another engineering challenge with NAT64 is using applications/protocols that embed IPv4 or IPv6 addresses in the payloads, and expect a gateway to translate those without use of NAT traversal techniques. This can add complexity and additional engineering requirements to extend the ALGs to support this capability.

5.2 Overview

There are two main methodologies within NAT64, stateful and stateless, explained below. The cable operators internal requirements would govern which option they decide to deploy.

In a stateful NAT64 configuration, an LSN shall be configured as a NAT64 gateway, engineered to enable IPv6-only clients and IPv4 resources to communicate with each other by way of address and packet translation. This translation operation shall be performed on the NAT64 gateway using stateful sessions.

When configured as a stateless NAT64 gateway an LSN shall be engineered to enable IPv6-only clients to communicate with IPv4-only resources by means of address and packet translation. In a stateless NAT64 configuration, LSN shall use a mapping table to match IPv6 request packets sent from the IPv6 client to an IPv4 destination address of an IPv4 resource.

6 Domain Functionality

6.1 End to End Network Domains

In order to operate the IPv6 transition technology it has to be engineered and verified end to end across the cable broadband network addressing each of the domains as illustrated in figure 2.

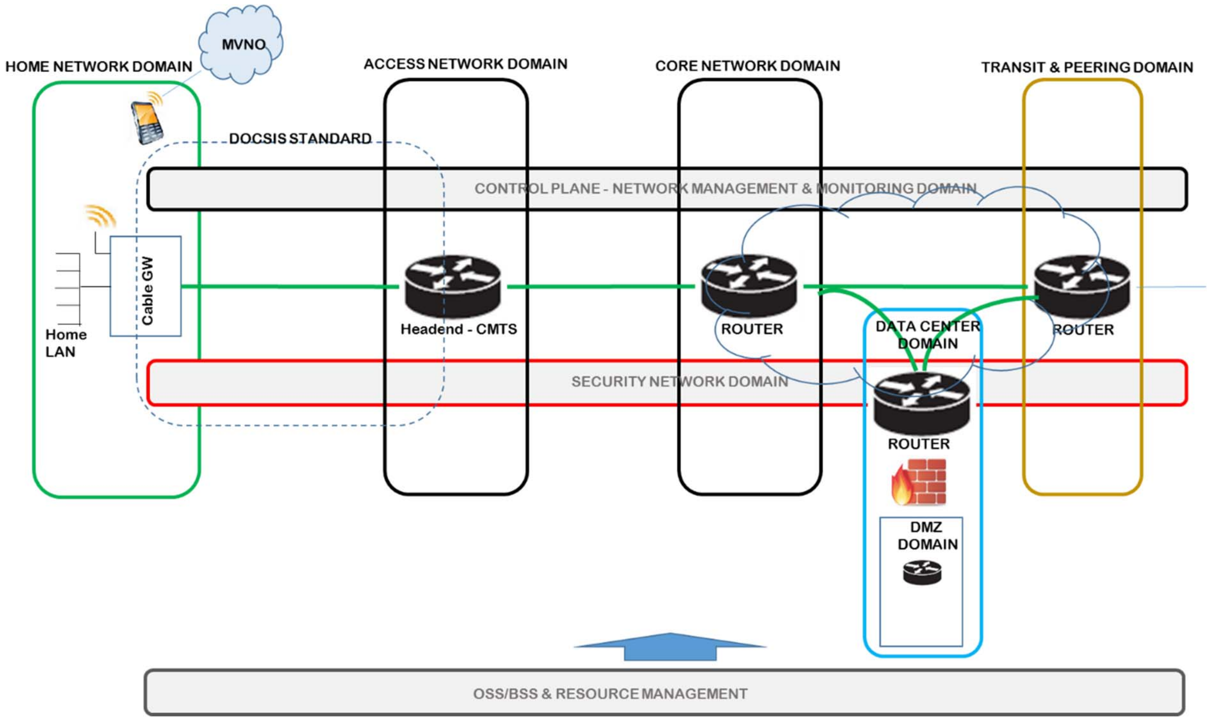

Figure 2: Illustration of the Cable Broadband Network Domains

6.2 CPE Home Network Domain

Customer's devices shall be V6 only, and customer device that is V4 is not supported with this technology.

Functionality to be engineered in the cable network:

- The cable residential gateway device shall be capable of IPv6 connectivity only.
- The DOCSIS management between the Cable Modem and CMTS may be IPv4 or IPv6.

6.3 Access Network Domain

Functionality to be engineered in the cable network:

- The cable headend CMTS shall be capable of IPv6 connectivity for customer traffic.
- The DOCSIS management between the Cable Modem and CMTS may be IPv4 or IPv6.

NOTE: There are no additional requirements on the HFC distribution network i.e. AMPS, taps, etc.

6.4 Core Network Domain

Functionality to be engineered in the cable network:

- The core network routers shall support IPv6 routing and forwarding capabilities.
- NAT64 gateway shall be implemented as specified by IETF RFC 6146 [19] for further information on its operation refer also to IETF RFC 7269 [20].

Operation:

A v6 client communicating with v6 service has direct end to end connectivity and does not utilize the NAT64 resources, however a v6 client communicating with a v4 service will utilize the functions of the NAT64 gateway such that the V6 packets are routed to the NAT64 GW which translates the address family from IPv6 to a shared IPv4 address which is then forwarded to the v4 server.

The communication between the NAT64 and V4 server for the session will be using v4 packets, whereas the communication between the v6 client and NAT64 gateway will be using v6 packets so this enables the V6 client to communicate to a V4 service with the NAT64 gateway maintaining the state of the session.

6.5 Data Centre Domain

Functionality to be engineered in the cable network:

- The DNS64 shall be implemented when operating IPv6 transition technology NAT64. The requirements for DNS64 shall be in accordance with IETF RFC 6147 [18].
- The DSN64 functionality replaces the exiting DNS caching server within the cable network.

Operation:

For an IPv6 request from the client to the DNS64 server if a valid IPv6 response is received by the DSN64 server then this is returned to the client. In this case the v6 client shall have a direct connectivity with the V6 service. The NAT64 server is not involved.

If it does not receive an IPv6 record then a synthetic AAAA record is created from the IPv4 record as specified by IETF RFC 6147 [18] to direct the IP packets to the NAT64 gateway. The NAT64 gateway then translates the address family as defined in clause 6.4 of the present document.

The DHCP server is inherently IPv6 only. The DHCP server shall include the DNS64 IPv6 addresses in its reply options to the CPE.

6.6 DMZ Service Domain

Functionality to be engineered in the cable network:

- There are no specific operational requirements for the DMZ however it is strongly recommended that all services within the DMZ are IPv6 capable either by dual stack configuration or the use of V6/V4 load balancing.
- V6/V4 load balancers allow V6 access to existing v4 services such as a web server avoiding the need to engineer v6 only content on these servers.
- To minimize the impact on the NAT64 server from the requirement of network address translation for the network providers own services then all services within the DMZ shall be IPv6 capable.

6.7 Transit and Peering Service Domain

Functionality to be engineered in the cable network:

There are no changes to be engineered when considering the transit and peering domain since the transmit and peering links are dual stack and support both IPv6 and IPv4 packets simultaneously and therefore no specific additional requirements are needed to be defined for this domain.

6.8 Management and Monitoring Domain

Functionality to be engineered in the cable network:

- DOCSIS management between the CMTS and Cable Modem may remain deployed either using IPv4 or IPv6 addresses.
- There is the additional functionality to monitor and manage the NAT64 gateway and DNS64 server in order to provide sufficient capacity to scale with the traffic throughput. Also the available V4 address pool shall need to be managed based on the number of customers and the address sharing ratio between the V6 and V4 ports e.g. 16 customers sharing 1 IPv4 address.

6.9 Security Domain

Functionality to be engineered in the cable network:

- The cable modem packet classifiers shall be updated to support IPv6 filtering.
- The continuity of the security of the end to end service shall be maintained when operating NAT64, however the operational implications from NAT64 shall be minimized as specified in IETF RFC 7269 [20], section 9, IETF RFC 6147 [18], section 8 and IETF RFC 6146 [19], section 5.
- The logging of IPv4 addresses for LI and DR purposes shall additionally contain the IP port and the customers IPv6 address.

7 Topologies

This clause explains the potential NAT64 domain topologies. At the time of writing, two topologies are possible, the integrated topology and the hairpin topology; however, depending on future hardware and software developments, additional topologies may be introduced.

Use of the integrated topology would provide the most flexibility, as in the hairpin topology the use of an IGP is required and will not allow to connect the LSN anywhere on the core to any PE without running BGP in the core.

To support the hairpin topology P-routers may need to be converted into PE-routers or extra PE-router devices may need to be deployed for one to distribute the LSN's as required.

In the integrated topology the LSN would function as a full MPLS 6PE router:

In the hairpin topology the LSN functions as L3 router, hair-pinning connections through an external 6PE router. This topology would only be used when application of the integrated topology is not possible.

8 Technical Requirements

8.1 General

NAT64 is one of the technologies that has seen some adoption in the mobile markets but little in the cable network markets. 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.

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 3.

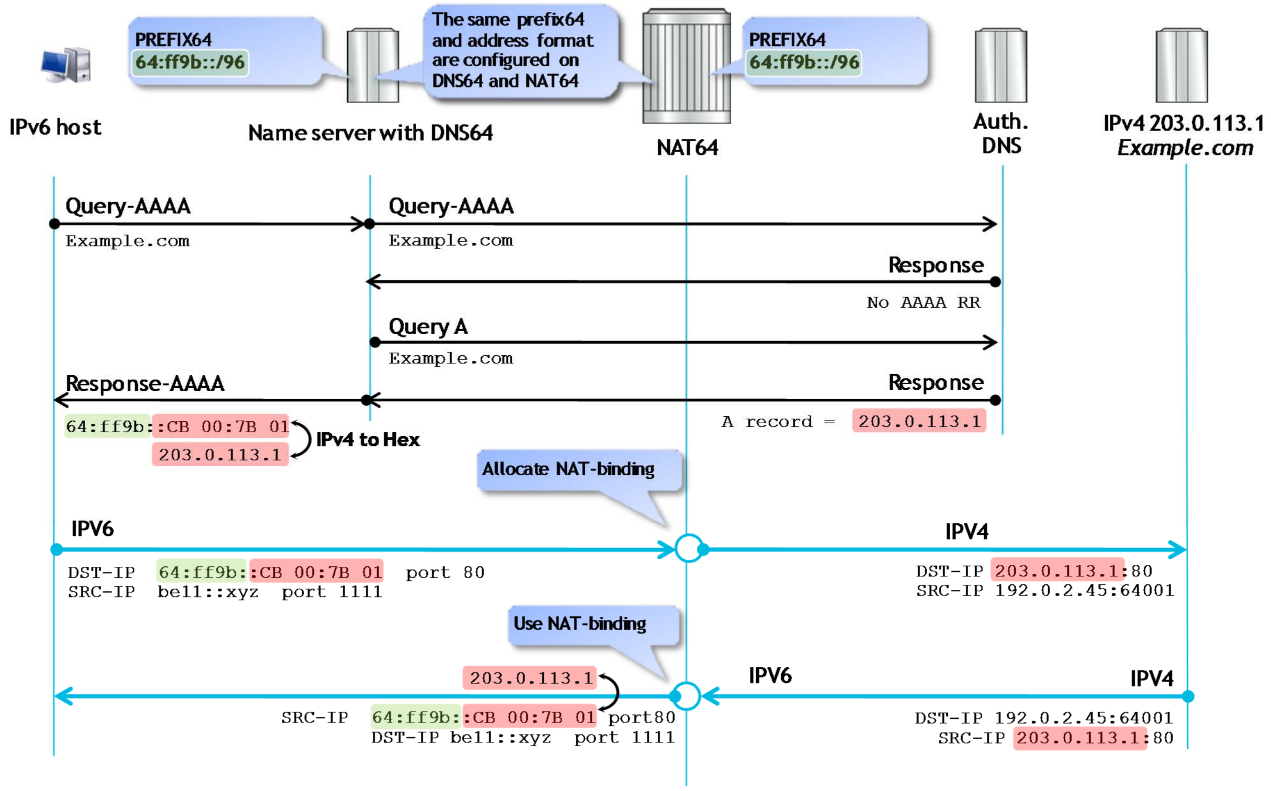

Figure 3: Addressing scheme in NAT64

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 bit).
  This is best practice.

- A network specific prefix.
  The addressing scheme defined for NAT64 [26] allows subnet lengths for the NAT64 prefix to be 32, 40, 48, 56, 64 or 96 bit.

Depending on the prefix length, the IPv6 address with embedded IPv4 address is formatted according to table 1.

Table 1: 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

NOTE: 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 Byte 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 Byte IPv6 header - 8 Byte IPv6 fragmentation header + 20 Byte 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.

8.2 High Level Requirements

8.2.1 General

The requirements of the NAT64 implementation are based on enabling seamless NAT64 connections without degradation in service, access, functionality or speed.

Two network components are involved in the end-to-end NAT64 approach; the LSN and the CPE or client. Requirement considerations for both are listed below.

8.2.2 LSN

The LSN device placed in the edge of the network (IXPE) as the IPv4 gateway to perform de-capsulation on the egress from a 4in6 packet to a pure IPv4 packet. Requirement considerations for LSN are:

- Hardware Topology
- Logical Topology
- Software/Hardware Features
- Scalability
- Resilience and Redundancy
- IP Allocation & DHCP specific features (v4 & v6)
- Forwarding/Convergence Performance
- Monitoring, Management, Reporting & Access
- DR Specifics

8.2.3 CPE

The CPE is a device in customer's home to encapsulate the traffic on egress from a pure IPv4 packet to a 4in6 packet. Requirement considerations for NAT64 CPE are:

- Hardware Topology
- Logical Topology
- Software/Hardware Features
- Scalability
- Stability
- IP Allocation/DHCP (v4 & v6)
- Forwarding performance

8.3 NAT64 technology feature requirements

The NAT64 technology feature is summarized in table 2 detailing for each function the requirement as required or optional with a brief description of each of the named functions.

Table 2: Summary of NAT64 Features Requirements

Functional Name	Requirement	Description
NAT64 - IETF RFC 4787 [24] (UDP)		Compliance with NAT behaviour according to IETF RFC 4787 [24] for UDP.
NAT64 - IETF RFC 5382 [25] (TCP)		Compliance with NAT behaviour according to IETF RFC 5382 [25] for TCP.
IETF RFC 6052 [26]	Required	Compliance with IETF RFC 6052 [26]: IPv6 Addressing of IPv4/IPv6 Translators.
IETF RFC 6145 [27]	Required	Compliance with IETF RFC 6145 [27]: IP/ICMP Translation Algorithm.
IETF RFC 6146	Required	Compliance with IETF RFC 6146 [19]: Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers.
Redundancy	Required	All critical components shall be redundant in such a manner that they can fail-over without impact to customer or management traffic greater than 1 ms.
Shared Resource	Required	Single NAT64 GW prefix. The NAT64 IPv6 prefix should be able to be shared amongst different NPU's in the NAT64. A hashing mechanism should be in place to hash all upstream packets based on the source IPv6 address.
Si-ID	Optional	Tunnel/customer identifier based on IPv6 CPE address.
NAT64 Addressing & Virtual Interfaces	Required	NAT64 shall be able to assign a single virtual interface with up to 8 NAT64 GW prefixes for any given NAT64 instance on the node.
Anycast	Required	Anycast NAT64 gateway prefixes are a requirement to allow simplicity of deployment for a single prefix across multiple NAT64's.
NAT64 Address withdrawal	Required	The NAT64 should have at least five points of NAT GW prefix 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.

Functional Name	Requirement	Description
NPU Load Balancing (hashing)	Required	NAT64 Inside to Outside hashing performed on the Source IPv6 (128 bits) address of the CPE device. NAT64 Outside to Inside hashing performed on the Destination IPv4 (lower-order 2 bits) address assigned from the pool prefix.
Tunnel MTU	Required	The Maximum Transmission Unit for NAT64 shall be configurable. Expected value will be 1 480 bytes.
MSS Clamping	Required	TCP MSS support is mandatory for the NAT64 due to the removal of an end-to-end MTU sizing functionality. This will avoid the need for excessive fragmentation.
Fragmentation	Required	Fragmentation should be done on the IPv4 packet. Fragmentation should be placed on the ASIC running in the line card. Optionally, fragmentation can be done on the IPv6 packet.
NAT - Network Address and Port Mapping - Endpoint Independent Mapping	Required	For two flows for a common inside source IPv6 address and port, the external address/port translation is independent of the destination IPv4 address and port and when the flows exist simultaneously in the NAT state table they will use the same translation.
NAT - Translation Filtering - Endpoint Independent Filtering	Required	A flow initiated externally can use the existing External/Inside IPv4 address/port mapping and it is independent of the source IPv4 address/port of the senders.
NAT - Paired IP Address Assignment	Required	Translation to External IPv4 address is done in a paired fashion. A given Inside IPv6 address is always translated to the same External IPv4 address.
NAT - Hair-pinning	Required	Different internal addresses on the same internal interface shall be able reach each other using external address/port translations.
NAT - 1:1 IP Mapping	Required	Ability to configure a one to one type of mapping for particular inside-VRFs: every public IP will be mapped to one and only one private IP (multiple ports are allowed).
NAT44 - Outside-Service-App mapping for inside-VRF	Required	Ability in the inside-vrf to provide the explicit outside serviceapp to be paired.
NAT - Port Limit configuration	Required	A maximum amount of ports can be configured for every IPv6 source address.
NAT - Per-Protocol Timeout configuration	Required	Timeout of mappings is critical and removing a NAT mapping at the appropriate time maximizes the shared port resource.
NAT - Dynamic Port Range start configuration	Required	The start port for dynamic port ranges should be configurable to allow for a range of ports for static port mappings.
NAT - Software Load Balancing	Required	NAT Inside to Outside hashing performed on the Source IPv6 CLAT Prefix. NAT44 Outside to Inside hashing performed on the Destination IPv4 (lower-order 2 bits) address assigned from the pool prefix.
Port Allocation	Required	In order to reduce the volume of data generated by the NAT device (logging creation and deletion data), bulk port allocation can be enabled. When bulk port allocation is enabled and when a subscriber creates the first session, a number of contiguous outside ports are pre-allocated. A bulk allocation message is logged indicating this allocation.
Deterministic NAT64	Optional	Algorithmically maps a customer's private IPv6 address to a set of public IPv4 address ports, allowing a significant reduction in logging.
FTP ALG (Active and Passive)	Required	FTP clients are supported with inside (private) address and servers with outside (public) addresses. Passive FTP is provided by the basic NAT function. Active FTP is used with the ALG.
RTSP ALG	Required	Remote control protocol for streamers (which use RTP/RTCP or RDT). Our implementation considers the server is located "outside" and clients are "inside". RTSP is used in many streamers like QuickTime or RealNetworks. Both "SETUP" and "SETUP response" are analysed. Based on this information, the translation entry will be created in the table.

Functional Name	Requirement	Description
ICMP ALG	Required	
SIP ALG	Required	
TFTP		
PPTP ALG	Required	
Use of Application SVI's, connecting multiple routing entities (inside/private VRF, outside/public VRF)	Required	NAT should be possible from any routing context (VRF, GRT) to any routing context (VRF, GRT).
Stateful ICMP	Required	Stateful ICMP mappings between inside and outside ICMP identifiers should be supported.
Thresholds	Required	Configurable thresholds using watermarks should be supported to monitor the resources on the NAT64.
QoS translation	Required	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.
Chassis NAT Clustering	Optional	Clustering of NAT64's to allow for inter-chassis resiliency.
PCP	Required	Support for PCP to allocate static port bindings.
Logging via Netflow V9/IPFIX	Required	Netflow v9 support for logging of the translation records. Logging of the translation records can be mandated by for Lawful Intercept. The Netflow uses binary format and hence requires software to parse and present the translation records.
Logging via Syslog	Required	Syslog support for logging of the translation records as an alternative to Netflow. Syslog uses ASCII format and hence can be read by users. However, the log data volume is higher in Syslog than Netflow.
Destination based Logging	Required	Destination Based Logging will generate one record per session, including the destination port and address. NFv9 templates including destination IP address and destination port will be used.
Base Logging Fields	Required	NAT64 logs the following information when a translation entry is created: - Inside instance ID. - Outside instance ID. - Inside IPv6 Address. - Inside Port. - Outside IPv4 Address. - Outside Port. - Protocol. - Start Time. - Stop Time.
Radius Logging	Required	Logging using Radius accounting messages on a per block allocation requirement.
XML l	Optional	Logging using XML files.
Static port forwarding (up to 6K static forward entries per npu)	Required	Static port forwarding configures a fixed, private (internal) IPv6 address and port that are associated with a particular subscriber while the NAT64 allocates a free public IP address and port. Therefore, the inside IPv6 address and port are associated to a free outside IP address and port.
Static port forwarding 1:1 active/standby	Required	Static Port Forwarding mapping will be kept constant with two npu cards in Active/Standby mode.
64 NAT64 instances per npu Card	Required	
20 M+ Translations (per npu)	Required	
Minimum Gbps throughput per npu (Inside-to-outside VRF + Outside-to-inside VRF) with IMIX traffic	Required	
1 M+ connections per second setup rate	Required	
1 M users	Required	
Latency	Required	Latency is between 390 and 580 micro seconds (µs).
6 npu Cards per chassis	Required	

Functional Name	Requirement	Description
IRB support	Required	Integrated Routing and Bridging (L3 interface for Bridge Domain).
Broadband Network Gateway (BNG) support	Optional	32 k BNG sessions and up to 256 k NAT users at the same time.

8.4 Detailed LSN Requirements

8.4.1 LSN Hardware Feature/Topology

NAT64 is an implementation of an IPv4/IPv6 transition protocol, and all CPE/CEs running NAT64 LSN connect to the CORE node. It is the translation from the IPv6 to IPv4 and the originating IPv4 NAT device egressing traffic towards the destination on IPv4 after translation.

An IPv6 client, whether CPE or an edge customer device like a mobile phone, does DNS64 lookup if needed, then embeds the IPv4 address into an IPv6 address in the last 32 bits. The client then sends packets to the resulting address. The NAT64 gateway creates a mapping between the IPv6 and the IPv4 addresses, which may be statically configured or dynamic in nature. The source address is replaced with the LSN northbound IPv4 interface so all packets can return to their origin.

The hardware topology is based on two distinct containers, CPE/Client and the LSN, and so both shall be present for NAT64 to function.

8.4.2 LSN General software requirements

8.4.2.1 Integrated topology requirements

List of features required for both IPv4 and IPv6/L2 and L3 for an integrated topology solution:

- MP BGP (as well 6PE and 6VPE)
- BGP Community/32 bit AS
- MPLS LDP (potentially only v4 native at present but the requirement for v6 native MPLS may become an absolute)
- ECMP
- QoS (v4/v6) - classification, priority queuing, etc.
- QPPB/SCU/DCU
- SNMP (transport over v4 and v6) v1/v2/v3
- ACLs/Prefix Listing (both v4/v6)
- TACACS/RADIUS (v4/v6)
- Syslog (event reporting for v4 & v6 as well as transporting over both protocols)
- CoPP (v4/v6)
- Netflowv9 (potentially previous versions will be required depending on the state of the NA4 implementation of Netflow)
- XML (v4 & v6 reporting and transport)
- Mac Accounting
- 802.1q

- Ether-channel
- Ether OAM
- NSF/GR (v4/v6)
- Policy Based Routing (v4/v6)
- ISIS (Potentially MT for ISIS as well if the MPLS IPv6 LDP allows for dual stacking) (v4/v6)
- Static Routing (v4/v6)
- OSPFv2/v3
- CDP/LLDP (v4/v6)
- VRRP/HSRP (v4/v6)
- VLAN Mapping/Double Tagging
- L3 Multicasting/MFIB (v4/v6)
- IPv6 Forwarding
- IPv4 Forwarding
- Ethernet technologies
- Hardware forwarding functionality, instead of software, across all priority flows (v4/v6)
- Virtual Interfaces (v4/v6)
- AAA (v4/v6)
- BFD (v4/v6)
- MLD/L2 Multicast
- Full NDP (ICMPv6, DAD, NUD, etc.)
- PIM/IGMPv2/v3
- CEF/dCEF
- Anycast
- Route Reflection (v4/v6)
- Standard v4 VPN
- ISSU/SSO
- NTP (v4/v6)
- SEND
- IP-Sec
- DNS (v4/v6 server & client)
- DHCP relay (v4/v6)
- MSTP/RSTP
- L2/L3 Load Balancing (v4/v6)
- VPLS (v4/v6)

- L2 Bridging
- NAT64
- NAT64
- NAT Cache Writing
- SSH/Telnet (v4/v6)
- Authentication across most protocols such as SNMPv3/BGP/LDP/ISIS/, etc.
- PS-BGP

8.4.2.2 Hairpin topology requirements

List of features required for both IPv4 and IPv6/L2 and L3 for a "hairpin topology" solution:

- SNMP (transport over v4 and v6) v1/v2/v3
- ACLs/Prefix Listing (both v4/v6)
- TACACS/RADIUS (v4/v6)
- Syslog (event reporting for v4 & v6 as well as transporting over both protocols)
- CoPP (v4/v6)
- Netflowv9 (potentially previous versions will be required depending on the state of the NA4 implementation of Netflow)
- XML (v4 & v6 reporting and transport)
- 802.1q
- Ether-channel
- Ether OAM
- Static Routing (v4/v6)
- CDP/LLDP (v4/v6)
- VRRP/HSRP (v4/v6)
- IPv6 Forwarding
- IPv4 Forwarding
- Ethernet technologies
- Hardware forwarding functionality, instead of software, across all priority flows (v4/v6)
- Virtual Interfaces (v4/v6)
- Full NDP (ICMPv6, DAD, NUD, etc.)
- CEF/dCEF
- Anycast
- ISSU
- NTP (v4/v6)
- SEND

- BFD
- IP-Sec
- DNS (v4/v6 server & client)
- MSTP/RSTP
- L2/L3 Load Balancing (v4/v6)
- NAT64
- NAT64
- NAT Cache Writing
- SSH/Telnet (v4/v6)
- Authentication across most protocols such as SNMPv3/BGP/LDP/ISIS/, etc.
- PS-BGP
- ECMP

8.4.3 Scalability

A NAT64 solution exhibits lower scaling and performance capabilities than other transition technologies due to the address embedding and ALG requirement. It can be difficult to engineer scaling for a minimum of 100 000 CPE addresses. Consequently latency is a challenge to engineer for a carrier grade solution, especially on primary flows and ALG functions, and this has a direct correlation to the processing requirement.

9 Technical Considerations

9.1 Hardware

Carrier-grade NAT64 solutions exist with a high level of capacity. The customer end point, based on the maximum number of flows per client and mapping, but dependent on the ALG process requirements. As NAT64 is quite dynamic in processing requirements hardware shall match and scale to fit, in other words the hardware cannot be scaled with just a 20 % or 30 % variable but with the potential of peaks well exceeding these numbers and thus can be more expensive in implementation.

The LSN build would be specific to the operator's service requirements. When engineering the build to meet the specific service needs considerations are given as an example guide:

- Engineered as a carrier grade router allowing for all requirements within the determined PE requirements.
- Allow for Route-Processor redundancy, with e.g. 4 to 10 open blade slots allowing for a single Integrated Services Adapter and a port based NPU placed within it.
- Each port card/module - support a 10 Gb/s interface which is matched at the PE router.
- Two 10 Gb/s connections into the LSN; one primary ingress and one primary egress. The primary egress supports the public IPv4 interface, dual stacked for redundancy in case the primary ingress, holding the IPv6 interface, goes down, and vice versa on the primary ingress.
- Priority of the traffic - have two LSNs, in HA mode to allow a client to access either LSN giving chassis redundancy.

- Growth as NAT64 requirement increases in capacity and customer base, by holding two chassis' with a single RP, one blade holding for example 20 Gb/s NPU and a port card, so it is lightly loaded to give for example a maximum of 40 Gb/s capacity with failover of 20 Gb/s. Allowing for redundancy comparative to a single chassis holding two NPUs and ISAs considering the chassis accounts for only approximately 3 % of the actual overall cost.

9.2 MTU and fragmentation

NAT64 is a translation technology and, thus, suffers less MTU requirements however there is the issue that fragmentation resends and general PMTU control can have performance effects on customer services. A further issue is that PMTU functions well in IPv6 but translating from an IPv6 packet to an IPv4 packet it is necessary to engineer functional capabilities to render each flow relative. This may be controlled with some success with protocol based MTU settings calculated to match one another residing within the remit of the transit networks themselves.

9.3 Scaling

The details of the scalability are structured based on the service requirements for delivery in the cable operator's network.

An example of the scalability requirements for a network are given below.

Engineer the network for:

1) Minimum number of flows per NPU e.g. 4-6 million.

2) The throughput on a single direction e.g. 60 Gb/s.

3) Load balancing on ingress ports across the chassis.

4) Two LSNs per cluster as a minimum.

5) Two global LSNs for failover of traffic.

6) Number of bindings per second NPU e.g. 500 k.

7) Initial bindings per second per NPU e.g. 500 k.

8) Scalability configurable for sessions per IP.

9) Minimum throughput per chassis e.g. 120 Gb/s.

10) Minimum and maximum number of slots to support the minimum throughput e.g. minimum 4 slots, maximum 10 port slots, allowing for a 120 Gb/s backplane.

11) Matching blades using both a NAT NPU and a port card in a single blade.

9.4 Reliability

To increase reliability and robustness or to distribute the load, operators could opt to deploy multiple LSNs, to share load across a large footprint. In the case of having multiple LSN's deployed, the IPv6 destination prefixes should be Anycast based if stateless and the client should hold a primary and secondary if stateful.

9.5 Quality of Service

The QoS policies defined by the cable operator for their network shall be engineered to operate properly with the new NAT64 environment.

A NAT64 stream can be viewed as a particular case of uniform conceptual tunnel model. This uniform model views an IP tunnel only as a necessary mechanism to forward traffic to its destination: the tunnel has no significant impact on traffic conditioning. In this model, any packet has exactly one DSCP field (in the outermost IP header) that is used for traffic conditioning at any point. In the NAT64 model, this is the Traffic Class field in the IPv6 header.

Implementations of this model copy the DSCP value to the outer IP header at encapsulation and copy the outer header's DSCP value to the inner IP header at de-capsulation.

Operators using this model shall provision the network such that the LSN copies the DSCP value in the IPv4 header to the Traffic Class field in the IPv6 header or the other way around, after the translation for the downstream or upstream traffic. Similarly, the B4 shall copy the Traffic Class field value in the IPv6 header to the DSCP to the IPv4 header. Traffic identification and classification can be implemented by examining the outer IPv6 header in the IPv6 access network and IPv4 in the corresponding CORE network after the LSN.

9.6 B4 deployment

In order to configure the IPv4-in-IPv6 tunnel, the B4 needs the IPv6 address of the LSN. This IPv6 address can be configured using a variety of methods ranging from an out-of-band mechanism, manual configuration, and DHCPv6 option to RADIUS. If a network operator decides to use DHCPv6 to provision the B4, the B4 shall implement the DHCPv6 option defined in IETF RFC 6334 [21]. If an operator decides to use RADIUS to provision the B4, the B4 shall implement IETF RFC 6519 [22].

Annex A (informative): Bibliography

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

IETF RFC 4459 (April 2006): "MTU and Fragmentation Issues with In-the-Network Tunneling".

IETF RFC 2983 (October 2000): "Differentiated Services and Tunnels".

draft-ietf-softwire-dual-stack-Lite-11: "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion".

History

Document history		
V1.1.1	August 2016	Publication