

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 16, 2012

W. Dec
R. Asati
Cisco
C. Congxiao
CERNET Center/Tsinghua
University
H. Deng
China Mobile
M. Boucadair
France Telecom
October 14, 2011

Stateless 4Via6 Address Sharing
draft-dec-stateless-4v6-04

Abstract

This document presents an overview of the characteristics of stateless 4V6 solutions, alongside a assessment of the issues attributes. The impact of translated or mapped tunnel transport modes is also presented in the broader context of other industry standard reference architectures and existing deployments.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 16, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Terminology	4
3. Stateless 4V6 Technical and Architectural Overview	5
3.1. IPv4 address and algorithmic port indexing	7
3.2. 4V6 CE IPv6 Address and domain info	7
3.3. IPv6 Adaptation Function	8
3.3.1. 4V6 Stateless Tunneling Mode	8
3.3.2. 4V6 Stateless Translation mode	9
4. Comparison of 4V6 transport modes	9
4.1. General Characteristics of 4V6 modes	9
4.2. Mobile SP Architecture and 4V6 Applicability	12
4.2.1. 3GPP overview	13
4.2.2. 3GPP and 4V6 modes	15
4.3. Cable SP Architectures & 4V6 Applicability	18
4.3.1. PacketCable Introduction	18
4.3.2. PacketCable Construct - Classifier	20
4.3.3. 4V6 Modes Impact on PacketCable	20
5. Overview of potential issues and discussion	21
5.1. Notion of Unicast Address	21
5.1.1. Overview	21
5.1.2. Discussion	22
5.2. Implementation on hosts	22
5.2.1. Overview	22
5.2.2. Discussion	23
5.3. 4V6 address and impact on other IPv6 hosts	23
5.3.1. Overview	23
5.3.2. Discussion	23
5.4. Impact on 4V6 CE based applications	24
5.4.1. Overview	24
5.4.2. Discussion	24
5.5. 4V6 interface	24
5.5.1. Overview	24

5.5.2. Discussion	24
5.6. Non TCP/UDP port based IP protocols - ICMP)	25
5.6.1. Overview	25
5.6.2. Discussion	25
5.7. Provisioning and Operational Systems	25
5.7.1. Overview	25
5.7.2. Discussion	25
5.8. Training & Education	27
5.8.1. Overview	27
5.8.2. Discussion	27
5.9. Security and Port Randomization	28
5.9.1. Overview	28
5.9.2. Discussion	28
5.10. Unknown Failure Modes	28
5.10.1. Overview	28
5.10.2. Discussion	28
5.11. Possible Impact on NAT66 use & design	29
5.11.1. Overview	29
5.11.2. Discussion	29
5.12. Port statistical multiplexing and monetization of port space	29
5.12.1. Overview	29
5.12.2. Discussion	29
5.13. Readdressing	30
5.13.1. Overview	30
5.13.2. Discussion	30
5.14. Ambiguity about communication between devices sharing an IP address.	31
5.14.1. Overview	31
5.14.2. Discussion	31
5.15. Other	32
5.15.1. Abuse Claims	32
5.15.2. Fragmentation and Traffic Asymmetry	32
5.15.3. Multicast Services	33
6. Conclusion	33
7. IANA Considerations	33
8. Security Considerations	33
9. Contributors and Acknowledgements	34
10. References	34
10.1. Normative References	34
10.2. Informative References	34
Authors' Addresses	36

1. Introduction

As network service providers move towards deploying IPv6 and IPv4 dual stack networks, and further on towards IPv6 only networks, a problem arises in terms of supporting residual IPv4 services, over an infrastructure geared for IPv6-only operations, and doing so in the context of IPv4 address depletion. This class of problem is referred to by the draft as the 4via6 problem, for which a stateless solution is desired driven by motivation as documented in [I-D.operators-softwire-stateless-4v6-motivation]. Solutions such as a 4rd [I-D.despres-softwire-4rd], [I-D.murakami-softwire-4v6-translation], and [I-D.xli-behave-divi-pd], as well as dIVI [I-D.xli-behave-divi] offer such stateless solutions, by using fully distributed NAT44 functionality located on end user CPEs, which allows the network operators' core to remain effectively stateless in terms of NAT44. The solutions, collectively called Stateless4V6, rely on the same IPv4 address being used by multiple CPEs, each with a different TCP/UDP port range, and are derived from the Address+Port (A+P) solution space [I-D.ymbk-aplusp]. Differences between the solutions come down to the mode of transport (translation or mapped tunneling), and the mapping algorithm used. This document looks at the issues that have been claimed as applying to A+P technology, in the specific context of the referenced solutions, and also analyzes the two modes of transport.

2. Terminology

Stateless4V6 domain: A domain is composed out of an arbitrary number of 4V6 CE and Gateway nodes that share a mapping relationship between an operator assigned IPv6 prefix and one or more IPv4 subnets along with all the applicable TCP/UDP ports, all mapped into the IPv6 address space. An 4V6 system can have multiple domains.

Stateless4V6 CE: A CPE node that implements 4V6 functionality including NAT44 which is provisioned by means of 4V6. The device interfaces to the SP network using native IPv6 and a IPv4-IPv6 adaptation service.

Stateless4V6 Gateway A Service Provider node that implements the stateless 46 adaptation functionality for interfacing between the SP's IPv6 domain and an IPv4 domain in delivering end user IPv4 connectivity beyond the domain.

IPv4 Address sharing The notion of attributing the same IPv4 address by multiple CEs in an 4V6 domain.

Port-set: A set composed of unique TCP/UDP ports (ranges) associated to a IPv4 address. A single 4V6 CE is expected to have a single port-set for each IPv4 address.

Port-set-id: A numeric identifier of a given port set that is unique in a given 4V6 domain. A port-set-id is used to algorithmically determine the port-set members. The port-set-id is conveyed to CEs as part the CE's IPv6 addressing information, ie it is part of IPv6 subnet or address of a given CE, and its format places no restriction on the use of SLAAC or DHCP addressing.

CE-index: A numeric value, composed of a full or partial IPv4 address and optionally a port-set-id, which uniquely identifies a given CE in an 4V6 domain.

3. Stateless 4V6 Technical and Architectural Overview

This section presents the architectural and technical overview of a stateless 4v6 solution, and evidenced in whole or in part by various stateless 4via6 solution proposals such as 4rd, dIVI. Figure 1 depicts the overall architecture with two IPv4 user networks connected via 4via6 CPEs that share an IPv4 address. The goal of the system is to allow IPv4 user connectivity to the Public IPv4 network, across an operator's IPv6 network.

A key characteristic of the system, and a major differentiator with respect to previous solutions, is that translation state is only (ever) present on the CE, with the rest of the system performing stateless transport. This stateless transport applies to both the mapped-tunnel and translated modes, as described in the dedicated sections.

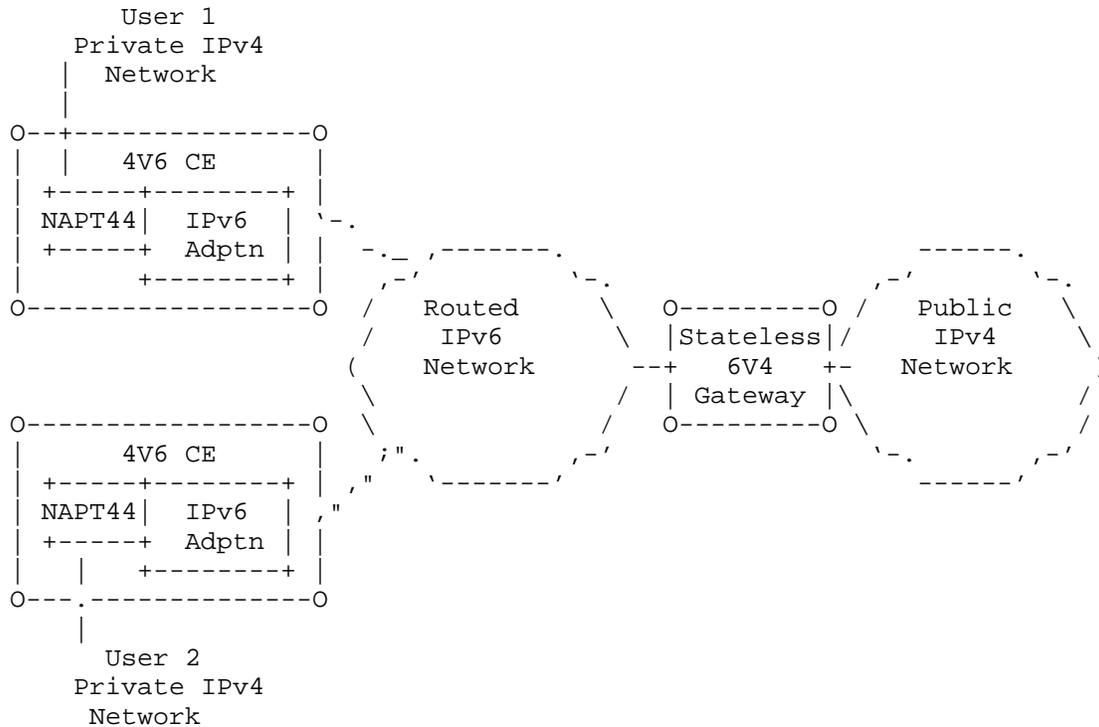


Figure 1 - Generalized Stateless 4V6 system

On IPv4 network user side, the routed IPv6 service provider network is demarcated with a 4V6 CE. The CPE externally has only a native IPv6 interface to the SP network, and a native IPv4 interface towards the end user network.

The IPv4 Internet is demarcated from the operator IPv6 network with one or more operator managed stateless 6V4 gateways that contain an IPv6 adaptation function (not detailed in the diagram) matching the one in the CE. Note: The stateless 6v4 gateway can be integrated into any existing network element (eg a core router, or an IP Edge).

Internally, the 4V6 CE is modelled as having a port restricted NAPT44 function coupled with a stateless IPv6 adaptation function that is able to ferry the end-user's IPv4 traffic across the IPv6 network, besides deriving 4V6 provisioning info from it. The NAPT44 function derives its IPv4 address, which may be shared with that of other users, and its unique Layer 4 (TCP/UDP) port range from the IPv6 address/prefix by means of an 4V6 algorithm and a port indexing schema. Any IPv4 ALG functionality that the CPE may support, remain unaffected. The CPE is expected to act as a DNS resolver proxy, using native DNS over IPv6 to the SP network.

Two forms of the IPv6 adaptation function are: i) 4v6 stateless tunneling ii) 4v6 stateless translation, each described in further in this document.

The service provider is assumed to be operating all the necessary provisioning and accounting infrastructure to support a regular IPv6 deployment. Similarly, the network operator is assumed to have the ability to assign an IPv6 prefix or IPv6 address to a CPE, and log such an address assignment.

End user host's DO NOT implement any of the 4V6, or other address sharing technologies, nor are they addressed directly with a shared IPv4 address. End user IPv4 hosts connected to the CPE receive unique private addresses assigned by the CPE, and it is the CPE that is directly addressed by the shared IPv4 address.

Although tangential to the discussion of stateless 4V6, it is useful to note that the CPE is expected to have a native IPv6 interface to the end user network, with any of the end user IPv6 hosts (single or dual stack) receiving IPv6 addresses from an IPv6 delegated prefix issued to the CPE.

3.1. IPv4 address and algorithmic port indexing

At the heart of the 4V6 solution, irrespective of mode of transport, lies the algorithm described in the specific solution drafts that allows the mapping of a shared IPv4 address and a TCP/UDP given port-set to a single IPv6 prefix or address. Notably, the 4V6 system allows both the shared IPv4 address use, as well as full non-shared IPv4 address use, all subject to the 4V6 domain configuration.

The S46 domain information required to compute the IPv4 address and correct port set is retrieved from the 4V6 prefix advertised to the CE, and pre-configured or statelessly acquired domain information.

3.2. 4V6 CE IPv6 Address and domain info

As presented in Section 2, IPv6 address of an 4V6 CE is composed out of the SP advertised IPv6 4V6 prefix, containing the CE-index, and an algorithmically computed appendix to complete the 128-bit address. This IPv6 address is *in addition* to any other IPv6 interface address that the CE configures or is configured with, including a SLAAC address from the 4V6 prefix or any IPv6 address source. One characteristics of the resulting IPv6 prefix or address is that it is for all intents and purposes a regular IPv6 prefix address that can be assigned to any regular IPv6 host.

The IPv6 4V6 interface is reserved for the 4V6 application and the

4V6 IPv6 adaptation function will exclusively use this IPv6 address. This is because the 4V6 system supports stateless communication between the 4V6 CE and the 4V6 gateway only by means of packets sent to/from this address.

3.3. IPv6 Adaptation Function

The IPv6 adaptation function plays a key role in the 4V6 system, in statelessly allowing the IPv4 user payload to be transported across an IPv6 (only) network. Two modes of such a function are currently proposed and presented in the following subsections

3.3.1. 4V6 Stateless Tunneling Mode

This type of IPv6 adaptation function is adopted and described in [I-D.despres-softwire-4rd].

The 4V6 gateway operates in the IPv4->IPv6 direction by mapping all or part of the IPv4 destination address and the port Index derived from the UDP/TCP payload into an IPv6 CE destination address. The resulting packet is sent using IPv4inIPv6 encapsulation to the CE, sourced from the 4V6's gateway IPv6 address, where the original IPv4 packet is extracted and passed to the stateful NAPT44 function.

The 4V6 CE operates in the IPv4->IPv6 direction, for traffic bound to the IPv4 internet, by encapsulating the IPv4 packet in an IPv6 header using IPv4inIPv6 encapsulation, and sending the resulting packet to the (well known) unicast address of the 4V6 gateway. There the IPv4 packet is extracted and forwarded.

The the original IPv4 packet addressing information is only partially visible on the IPv6 data plane, and the original Layer 4 information is only visible as part of the encapsulated IPv4 payload packet.

The figure below illustrates the CE model of a 4v6 Mapped Tunnel mode.

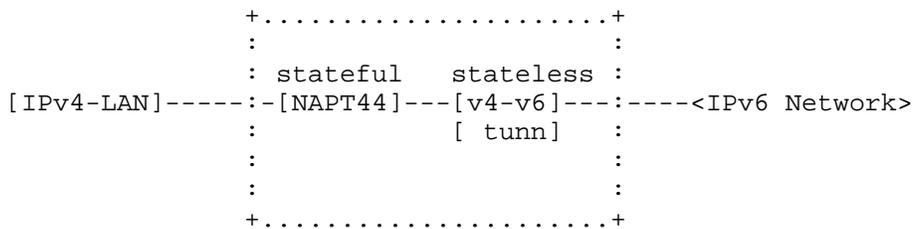


Figure 2 - 4v6 CE model with Tunnel mode

3.3.2. 4V6 Stateless Translation mode

This type of IPv6 adaptation function is adopted and described in [I-D.murakami-softwire-4v6-translation], I-D.xli-behave-divi-pd, and[I-D.xli-behave-divi] The 4V6 translation mode transport operates by means of stateless NAT46 [RFC6145] extended to map the the TCP/UDP port index algorithmically derived from received IPv4 packets into an IPv6 address suffix, in the IPv6 header, besides the full IPv4 mapped representation of the original IPv4 address information. The resulting packet is then sent across the IPv6 domain as an IPv6 packet - this IPv6 packet, besides mapping the original original IPv4 address information into a determinate IPv6 format, also places the Layer 4 and packet content directly after the IPv6 header, as any regular IPv6 with TCP/UDP packet. This IPv6 packet is thus capable of being processed by regular IPv6 network elements or servers in the IPv6 domain. At either end of the IPv6 domain, the IPv4 packet header is statelessly recreated, by the 4v6 CE or gateway, again using exactly the same NAT64 process as in [RFC6145].

The figure below illustrates the IPv6 4v6 Stateless Translation model of a 4v6 CE.

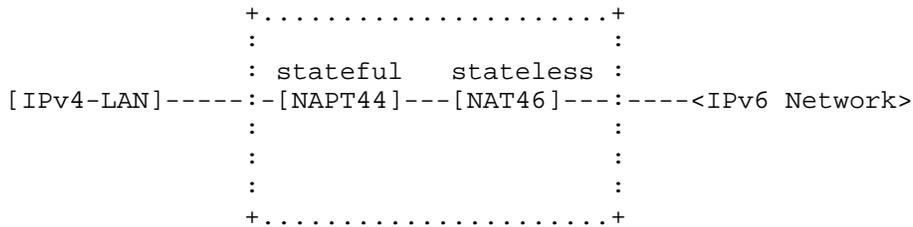


Figure 3 - 4v6 CE model with stateless NAT64

4. Comparison of 4V6 transport modes

This section presents the an overview of the similarities and differences between an IPv4-IPv6 translation based 4V6 transport mode and one that utilizes IPv4-in-IPv6 tunnelling. The comparison takes into consideration a wider deployment view composed of functionality that is known to be in common use today.

4.1. General Characteristics of 4V6 modes

The following table presents a comparison of the 4V6 transport modes, in terms of the base technology, and constrains, including also IPv4.

Item	4V6 Translation mode	4V6 Tunnel Mode
Base Technology	Port restricted NAPT44 with modified stateless NAT64 on CPE and Gateway	Port restricted NAPT44 with IPv4 in IPv6 mapped tunneling on CPE and Gateway
Location of stateful NAPT44 function	CPE	CPE
IPv4 Forwarding paradigm	L3 + L4 lookup	L3 + L4 lookup
IPv6 Addressing Constraints	CE uses 4V6 suffix.	CE uses 4V6 suffix.
Type of IPv6 prefix/address announcement method supported	ICMPv6 (SLAAC), DHCPv6 (both IA_NA and IA_PD)	ICMPv6 (SLAAC), DHCPv6 (both IA_NA and IA_PD)
Can the 4V6 IPv6 prefix be used by non 4V6 devices?	Yes	Yes
IPv4 addressing constraints	Fixed sharing ratio per IPv4 address.	Fixed sharing ratio per IPv4 address.
TCP/UDP Port range constraint	Ports are statically allocated	Ports are statically allocated
Requires ALG64 or DNS64	No	No
Requires IPv6 DNS on CPE	Recommended	Recommended
4V6 CE Parameter provisioning methods (assuming suitable protocol extensions)	ICMPv6, Stateless DHCPv6, TR69	ICMPv6, Stateless DHCPv6, TR69.

IPv6 Domain Routing to CE based on:	Regular closest IP match to CE-IPv6 subnet	Regular closest IP match to CE-IPv6 subnet
-----	-----	-----
IPv6 Domain Routing to 4V6 Gateway based on	IPv6 4V6 domain aggregate route	4V6 Gateway unicast/anycast address
-----	-----	-----
IPv4 Header Checksum recalculation required	Yes	No
-----	-----	-----
Supports non TCP/UDP Protocols	No*	No*
-----	-----	-----
ICMPv4 Limitations	No ICMPv4 from "outside the domain". Internal ICMPv4-v6 translation as per [RFC6145]	No ICMPv4 from "outside the domain".
-----	-----	-----
ICMPv5 identifier NAT/Markup needed	Yes	Yes
-----	-----	-----
Supports IPv4 fragmentation (without additional state)	No	No
-----	-----	-----
Requires IPv6 PMTU discovery/configuration	Yes	Yes
-----	-----	-----
Supports IPv4 Header Options	No - as per NAT64 [RFC6145]	Yes (use of source route option is constrained)
-----	-----	-----
TCP/UDP Checksum recalculation	Yes - depending on suffix, as per NAT64 [RFC6145]	No
-----	-----	-----
Supports UDP null checksum	Yes/Configurable - as per NAT64 [RFC6145]	Yes
-----	-----	-----
Transparency to DF bit	Yes, configurable as per [RFC6145]	Yes
-----	-----	-----

Supports IPv4 Fragmentation	Partial (no fragments from outside the domain)	Partial (no fragments from outside the domain)
-----	-----	-----
Transparency to IPv4 TOS	Yes, configurable as per [RFC6145]	Yes
-----	-----	-----
Overhead in relation to original average payload on IPv6 of a) ~550 bytes b) 1400 bytes).	a) 0% b) 0%	a) 4.36% b) 1.71%
-----	-----	-----
Supports non-shared IPv4 usage (ie whole IPv4 address assignment to a single device)	Yes	Yes
-----	-----	-----
Can support IPv4 to IPv6 host communication (for traffic not requiring ALGs)	Yes - As per [RFC6145] stateless NAT64 specification	No
-----	-----	-----
Changes to network element provisioning tool(s)**	Yes - Mapping IPv4 to IPv6 addresses	Yes - Enabling IPv4inIPv6 functionality
-----	-----	-----

* Without specific ALGs. Non UDP/TCP protocols, like ICMP, can be supported with specific ALGs.

**Network (feature) provisioning tools/applications need to be 4V6 aware. With the translation technique, the tool needs to enable the operator to map IPv4 addresses to IPv6 addresses as dictated by the 4V6 domain. With the tunneling technique, the tool needs to allow the operator to enable IPv4 (inIPv6) functionality and modify its characteristics.

4.2. Mobile SP Architecture and 4V6 Applicability

This section presents the applicability and comparison of the 4V6 modes to current 3GPP architectures used by Mobile SP for delivering all sorts of mobile services.

4.2.1. 3GPP overview

The 3rd Generation Partnership Project (3GPP) is a collaboration between groups of telecommunications associations, whose scope is to develop a globally applicable mobile phone systems and architectures based on service requirements. 3GPP standards are structured as Releases, each of which incorporates numerous individual standard documents. Currently, 3GPP Release 7 is the latest release in common practical deployment, with Release 8 being readied for deployment. Releases 9 and 10 are finalized, and work is underway on Release 11.

One of the major service requirement drivers of recent and ongoing 3GPP releases is the realization of services that deliver specific QoS, or user charging goals, all based on a policy system (eg tiered data rate or volume plans). Technically this translates to the Policy and Charging Control (PCC) framework, which in turn attributes specific functionality to nodes in the 3GPP architecture, such as the PDN-Gw and the PCRF. This functionality comprises both data-plane features (eg IP flow classification) as well as the interfaces/protocols between nodes (eg Diameter, and its specific 3GPP applications).

The 3GPP specifications allow both IPv4 and IPv6 traffic to be handled, and subject to operator defined handling and charging policies by means of applying suitable user traffic filters. Such filters are currently defined to be either IPv4 or IPv6, are applicable to user plane traffic, and are used in a variety of critical roles including the signalling of PDP contexts/EPC Bearers, as well as PCC signalling and interaction with applications.

The following table illustrates the impact of the 4V6 translation and tunnel transport modes respectively on the 3GPP architecture including PCC interfaces. In assessing the impact of these 4V6 transport modes a number of additional assumptions are taken:

- o The 3GPP system supports native IPv6 user traffic, as say per either of the E-UTRAN Release 8 or 9 specifications, using the relevant EPS bearer or PDP functionality.
- o The 4V6 gateway functionality is not part of the 3GPP core architecture (given that currently it is not scoped by a 3GPP Release). Instead, the 4V6 gateway is taken to be a stand alone component in the 3GPP network operator's core reachable via the SGi interface.

The above system, in the context of 3GPPs E-UTRAN architecture as defined in [E-UTRAN] is shown in Figure 2

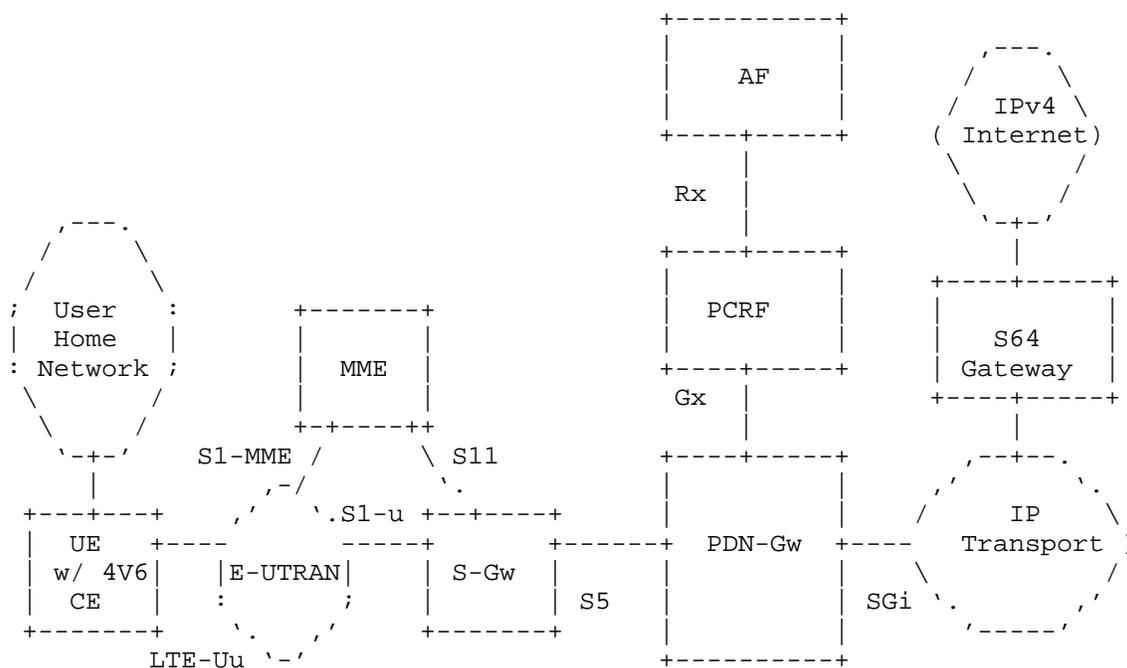


Figure 2 - 3GPP Architecture with 4V6

The main 3GPP system components, and terms are summarized as follows (the reader is referred to [E-UTRAN for a more detailed definition]:

- UE The User Equipment, typically a phone or a 3G/4G capable Home Router (shown to incorporate 4V6 functionality)
- E-UTRAN Evolved Universal Terrestrial Radio Access Network. The Radio Access network, composed on E-NodeB elements.
- MME Mobility Management Entity. Responsible for user authentication, PDN/SGw selection. Does not interact with the user data plane
- S-Gw Serving Gateway (function). Responsible for handling local mobility, (some) traffic accounting, traffic forwarding, bearer establishment.
- PDN-Gw Packet Data Network Gateway (function). Responsible for per user IP traffic handling, incl. address assignment, filtering, QoS, accounting.

PCRF Policy And Charging Rules Function. Responsible for authorizing and applying policy rules, as well as binding them to user bearers.

Bearer The bearer represents a virtual connection, typically that between a UE and a PDN-Gw. The bearer is specified as an IP Fliter (in terms of IP address, port numbers) and is the object of policy rules. 3GPP, depending on Release and document, defines many terms that are used to refer to the same notion: PDP context, EPS Bearer.

AF Application Function. A functional element offering (higher level) applications that require dynamic policy and/or charging control over the user plane (bearer) behaviour. The AF can be seen as bridging the gap between applications and how they affect the IP data plane of a user.

S5 It provides user plane tunnelling and tunnel management between SGW and PDN-GW, using GTP or PMIPv6 as the network based mobility management protocol.

S1-u Provides user plane tunnelling and inter eNodeB path switching during handover between eNodeB and SGW, using the GTP-U protocol

SGi It is the interface between the PDN-GW and the packet data network. Packet data network may be an operator external public or private packet data network or an intra operator packet data network.

Gx Bearer and flow control interface between the user data-plane element (PDN-Gw) and the Policy System. A Diameter based interface with a suite of 3GPP applications

4.2.2. 3GPP and 4V6 modes

4V6 translated traffic appears for all intents and purposes as regular IPv6-user traffic to the 3GPP system and packet processing functions (eg the PDN-Gw). Hence, and based on the stated assumptions, any such 4V6 traffic can be handled using existing native IPv6 functionality defined by the core 3GPP specifications.

In contrast, 4V6 tunneled traffic requires additional data plane processing to get to the "real" user IPv4 payload and apply the desired functions. Such additional processing is currently not part of the functionality covered by the 3GPP specifications. In view of this, and solely in relation to the 4V6 tunnel transport mode, two alternative hypotheses need to be placed in order to complete the comparison

i) that such IPv4 in IPv6 processing functionality will be supported as part of the existing EPS bearer functionality defined in E-UTRAN, perhaps as a dedicated EPS bearer (ie an additional virtual interface per subscriber). Or, that;

ii) a new 46 EPS bearer type (ie interface type) identification and signalling will be defined by the 3GPP architecture, which formalizes the v4inv6 relationship between the IPv4-user payload and the v6-user layers.

An apparent benefit of approach (ii) would be in allowing the system to clearly distinguish and expose to other systems v4-user traffic versus v6-user traffic, which is composed of v4inv6 and regular v6 traffic that a UE may generate. The former approach (i) is more convoluted given the ambiguity in distinguishing, and representing such a combination of v6-user and v6-user-bearer and v4-user traffic, all while keeping coherence in terms of the policy system. These two options are designated with ** in the table below.

Item	4V6 Translation Mode	4V6 Mapped Tunnel Mode
User Data Plane at the PDN-Gw (as per section 5.1.2 in [EUTRAN])	IPv6 over GTP-U over UDP over IP	IPv4 over IPv6 over GTP-U over UDP over IP
Gx (Diameter)	No discernible impact	Impacted: no way to express v4 over v6 in TFT Filter and Flow Descriptors
Rx (Diameter)	No discernible impact	Impacted: no way to express v4 over v6 in Media-Component-Description and, Flow-Description-AVP
S5 (GTP)	No impact	Impacted with new PDP/EPS Bearer type*
New 46 Bearer definition	Not required	Possibly required**

Secondary interface (dedicated bearer or secondary PDP) for 4G traffic	Not required	Possibly required**
-----	-----	-----
PDN-Gw	No impact	New TFT capability, IP Gate functionality, changes to Gx, and likely changes to S5/S7 related to signalling the new bearer
-----	-----	-----
SGW	No Impact	No discernible impact
-----	-----	-----
PCRF	No impact for IPv6. Feature to map IPv4-IPv6 addresses needed only in case of IPv4-only applications.	Impacted for both IPv6 and IPv4-only applications and Gx applications utilizing flow control/charging
-----	-----	-----
AF Application Function	No discernible impact	Flow based application control impacted
-----	-----	-----
UE	4V6 application	4V6 application
-----	-----	-----
LTE-Uu	No discernible impact	Likely changes required to support signalling of EPS bearer or PDP type
-----	-----	-----
Lawful Intercept	No discernible impact	New rules for tunnel support
-----	-----	-----

*A new PDP Type or EPS bearer signalling has a broader 3GPP system wide impact not fully covered here.

As the table illustrates, the 4V6 tunnel transport model appears to affect a significant number of 3GPP elements, when the intent is to realize a full suite of services. This observation appears to apply to any other carrier inserted tunneling technology (eg DS-lite). Hence, a substantial investment in 3GPP standard terms and in the evolution of deployed systems appears to be required.

In contrast the 4V6 translation mode bears none to no discernible impact on existing 3GPP Release 8/9 specifications and their deployments, while allowing the operator to realize the full set of services on 4V6, alongside any native IPv6 traffic, allowed for by these architecture. Hence, little beyond the addition of 4V6 components operating using translation mode appears to be required.

4.3. Cable SP Architectures & 4V6 Applicability

Cable SPs (commonly referred to as Multi System Operators (MSOs)) usually deliver video, data, and voice service over the cable and fiber access to residential and commercial customers. Many MSOs offer SLAs with various services by exploiting QoS not only in their IP/MPLS network, but also their access network.

The cable access network (now synonymous with Hybrid Fiber Coax (HFC)) is commonly enabled with Data Over Cable Service Interface Specifications (DOCSIS, a CableLabs standard) to facilitate the implementation of packet based services. In this paradigm, the HFC/DOCSIS access bandwidth is typically shared among a number of customers, hence, ensuring optimal service quality & experience per customer becomes extremely important for MSOs' success.

Cable SPs/MSOs ensure the optimal service quality of various advanced & real-time multimedia services (such as IP telephony, multimedia conferencing, interactive gaming etc.) by utilizing "PacketCable" framework to enforce QoS on the HFC/DOCSIS access.

The next sub-section 4.3.1 provides a brief introduction to PacketCable, section 4.3.2 explains a key PacketCable construct - Classifier, and section 4.3.3 tabulates the impact of 4V6 modes to PacketCable enabled DOCSIS/IP services.

4.3.1. PacketCable Introduction

PacketCable, a CableLabs standard, defines a framework for ensuring the Quality of Service (QoS) on the HFC/DOCSIS Access. PacketCable specifications (e.g. PacketCable 1.0, PacketCable Multi Media [PCMM], PacketCable Dynamic QoS [PC-DQOS], PacketCable 2.0) specify interoperable interface specifications for executing QoS, Admission Control, Accounting, Policy, and Security functions on Cable Modem (CM) and Cable Modem Termination System (CMTS), as/when needed. They all require DOCSIS 1.1 or later versions.

The PacketCable framework is also critically important for MSOs to comply with government regulations for things such as E911 when they offer voice/telephony services, Lawful Intercept (LI) etc.

Dec, et al.

Expires April 16, 2012

[Page 19]

4.3.2. PacketCable Construct - Classifier

PacketCable framework fundamentally relies on Cable Modem (CM) and Cable Modem Termination System (CMTS) to first qualify and then classify the appropriate IP traffic between them, for effective QoS enforcement. The framework requires the usage of "Classifier" for both qualification (in control plane) and classification (in data plane).

Taking PCMM specification [PCMM] again as an example, PCMM mandates the usage of classifier in the control plane (i.e. 'Upstream Packet Classification Encoding' in pkt-mm-1 interface (DOCSIS) , whereas 'Multimedia Classifier Object' in pkt-mm-2 and pkt-mm-3 interfaces (COPS)) for conveying the attributes of an IP flow belonging to an application (telephony, say), and subsequently its usage in the data plane i.e. filter matching on the IP packets' layer2/3/4 headers prior to QoS treatment.

The PCMM specification mandates the 'classifier' to include Source and Destination IP addresses, DSCP/TOS, IP Protocol, Source and Destination ports for an IPv4 traffic flow received by the CMTS, and similarly, Source and Destination IP addresses, TC, Next Header, Source and Destination ports for an IPv6 traffic flow received by the CMTS.

Similar to PCMM, PacketCable DQOS specification [PC-DQOS] also mandates the usage of classifier in the control plane (DSx messaging). In particular, PC-DQOS mandates the classifier definition to have 'protocol' (or next header) in IP header to be 17 (=UDP) along with specific Source and Destination ports (and Source and Destination IP addresses, optionally) so as to accommodate voice RTPoUDPoIP traffic.

In summary, the CMTS (and CM) construct their data-plane filter based on the 'classifier' information.

4.3.3. 4V6 Modes Impact on PacketCable

In 4V6 Tunnel mode, the 4V6 tunneled traffic requires additional data plane processing to get to the "real" user IPv4 payload and apply the desired functions. Such additional processing is currently not part of the functionality covered by the PacketCable specifications, nor part of compliant implementations.

In 4V6 Translation Mode, the 4V6 translated traffic appears for all intents and purposes as regular IPv6-user traffic to the PacketCable framework (both control plane and data plane). Hence, it is likely that any such 4V6 traffic can be handled using native IPv6

functionality e.g. classifier as defined by the PacketCable specifications and supported by CMTS and CM.

Taking PCMM specification as an example, it is worth noting that PCMM already allows for (and mandates) a minimum of four classifiers to be included in Gate-set. Hence, a Policy Server can communicate (via pkt-mm-2) both IPv4 and IPv6 classifier to the CMTS, which can use IPv6 classifier for constructing its data-plane filters (for DownStream processing), and convey IPv4 classifier to the CM via DOCSIS messages (pkt-mm-1) for any Upstream Processing. So, the 4V6 Translation Mode would work out in current implementations/deployment reasonably well.

Separately, it is likely that the CPE Router would be engaged in serving IPv4 multicast content to IPv6 receivers (and vice versa) in future, requiring 'translation' function.

In summary, while 4V6 Translation mode can work with the existing PacketCable framework, 4V6 Tunnel mode can not.

5. Overview of potential issues and discussion

This section summarizes the issues attributed to an A+P, or port restricted scheme, along with a discussion of applicability to the assumed system and possible resolutions. The summary of issues stem from [I-D.thaler-port-restricted-ip-issues] and associated discussions.

5.1. Notion of Unicast Address

5.1.1. Overview

The issue, referred to as the "definition of a unicast address", relates to the notion that in a shared IPv4 address system, multiple hosts will be visible as having a single IPv4 address outside of the system. This issue is a general characteristic of any NAT44 based solution [I-D.ietf-intarea-shared-addressing-issues], including DS-Lite. However, a more specific aspect of this issue in the context of an address sharing system is the possibility that a single host having multiple interfaces will be assigned the same IPv4 address (with different port ranges) on each of its interfaces. It may also be that multiple hosts sharing an address find themselves on the same Layer 2 segment. Either would impede hosts from working within the notion of known host IP stack and protocol implementations.

5.1.2. Discussion

A number of the characteristics of the 4via6 solution architecture cause the issues not to be applicable, key of which is that there is no expectation for any kind of end hosts to be part of the shared IPv4 address system.

In the stateless 4via6 system, CPE nodes are assigned with a shared IPv4 address+port range by means of the unique IPv6 address, containing the embedded IPv4 address + port index, of that CPE node. The CPE node is in addition enabled to be running the port restricted NAPT44 function from the IPv6 derived address, a key characteristic of the solution. On the IPv6 plane, the IPv6 address of the CPE is practically indistinguishable from any "regular" IPv6 address, and in fact any host that is not aware of it conveying an embedded IPv4 address would be able to use this just like any other regular IPv6 address, ie the 4via6 solution uses standard IPv6 addressing. In terms of the IPv4 dimension, since the shared address and port index are never used to address native IPv4 nodes or hosts, but instead uniquely assigned to a single NAPT44 function that is part of the CPEs, all legacy or other IPv4 hosts are not exposed to the presented issues.

Going beyond the ascribed issue however, it appears desirable to have the 4via6 CPEs that are to be part of the shared system to be able to provide a hint to the network operator in terms of their special capability. Such a hint can be a DHCPv6 Option Request Option, which would be useful to allow the DHCPv6 sub-system to also inform the CPE of any other stateless 4via6 system parameters. A largely similar ORO option is currently being defined as part of [I-D.ietf-software-ds-lite-tunnel-option]

Recommendation: Define a suitable DHCPv6 ORO for conveying the 4via6 capability of a CPE.

5.2. Implementation on hosts

5.2.1. Overview

The issue, as presented, relates to the need for modifications on end hosts or devices to support a port constrained mechanism and the overall impossibility of realizing such modifications. Furthermore, host applications that attempt to bind to specific ports that are not part of the allowed port range will fail to do so and may also require modifications.

5.2.2. Discussion

As presented in Section 3 the solution assumes the use of a dedicated CPE implementing the 4via6 functionality within a port constrained mode and NAPT44. Granted, CPE nodes will require to implement new functionality such as the IPv6 adaptation function, that is likely alongside introducing native IPv6 support. However, any and all existing end user IPv4 devices (eg PCs, etc) will not be affected. Nor are such devices expected to behave in any way different from that of today, where they typically obtain a private rfc1918 address and multiplexed by a CPE using a NAPT44 function.

In summary, the assumed 4via6 solution requires a specific 4via6 CPE but does not require any IPv4 host stack changes.

5.3. 4V6 address and impact on other IPv6 hosts

5.3.1. Overview

The issue relates to the question of whether the operation of a regular IPv6, non 4V6 capable, host would be adversely impacted should it be assigned or auto-configured with an address from an S64 address or prefix pool.

5.3.2. Discussion

The 4V6 prefix is for all intents and purposes a regular IPv6 prefix, and as such can be announced/assigned to any IPv6 host which in turn can use derived addresses like any other IPv6 address. Thus, an 4V6 IPv6 domain can address non-4V6 devices, leaving such devices to operate as native IPv6.

There is however a restriction on the 4V6 CE devices. As described in Section 2, a 4V6 CE constructs itself the full 128 bit address from the concatenation of the IPv6 prefix, 4V6 domain information acquired statelessly, and a pre-determined or algorithmic interface-id. By definition, only one 4V6 CE can use the same IPv4 address and port index. Hence, while there is no exact limitation on the number of non 4V6 hosts that can be addressed from an 4V6 prefix, there is a limit of one 4V6 CE per 4V6 prefix. Using a 4V6 prefix to address network segments without 4V6 devices does diminish the efficiency of the IPv4 address sharing mechanism, in terms of using up port ranges onto segments that will not use them. This is naturally a deployment consideration which an operator can optimize.

5.4. Impact on 4V6 CE based applications

5.4.1. Overview

It has been claimed that applications implemented on the CE itself, eg a DNS resolver-client, may be impacted by the 4V6 functionality. In particular, a concern is that such applications would either need to be specially engineered to issue socket calls or extensive IP stack modifications made to support them.

5.4.2. Discussion

By definition the 4V6 CE is an IPv6 capable device, and any IPv6 capable applications will be able to use the native IPv6 stack (note: IPv6 interface selection, is discussed in section 5.5). As such, the concern raised does not apply to applications that can be expected to support IPv6, and instead only to IPv4-only applications running on the 4V6 CE.

The shared IPv4 address is intended to be used only by the 4V6 CE function. This shared IPv4 address does not need to be assigned to an interface on the 4V6 CE and thus a target for potential applications. Any such applications running on the 4V6 will use any of the other (likely private) IPv4 address on the CE, which then will be routed to the 4V6 function this is applied post routing for the packets generated by these applications.

5.5. 4V6 interface

5.5.1. Overview

A 4V6 CE will have a "self configured" 4V6 IPv6 interface address, alongside any other SLAAC or DHCPv6 derived addresses, potentially from the same prefix. This particular 4V6 address may be subject to specific filtering rules or restrictions by the operator, besides usage and filtering restrictions on the 4V6 CE. Also, for the 4V6 system to operate as intended, the 4V6 application on the CE must be restricted to using the specific 4V6 address when sourcing 4V6 packets. Also, the 4V6 CE needs to be set-up to correctly forward IPv4 traffic to the 4V6 application.

5.5.2. Discussion

While the method of creating the interface is implementation specific, the generic operating model that is envisaged is for the 4V6 application to create the 4V6 interface as a virtual interface with an IPv4 unnumbered address. The application would then install a default IPv4 route pointing to this virtual interface, which would

be effectively see the 4V6 application acting as a network appliance on the forwarded traffic. In terms of IPv6 behaviour, the 4V6 application is expected to be set up to specify the use (binding) to the 4v6 IPv6 virtual interface.

5.6. Non TCP/UDP port based IP protocols - ICMP)

5.6.1. Overview

This issue relates to the inability of using regular ICMP messages to "ping" an end-host that has been addressed with a shared IPv4 address. The issue can be generalized one applicable to any IP protocol that is not TCP/UDP port based, and also in terms of the ability of using such protocols from end hosts that are assigned a shared IPv4 address.

5.6.2. Discussion

The inability to ping a CPE from the IPv4 Internet is shared by other IPv4 address sharing mechanisms such as DS-Lite. Thus, the issue is no better or worse in the case of the stateless 4via6 solution. The same can be said of end user hosts using other non UDP/TCP port based protocols from behind a NAT44 function, ie they will not function irrespective of address sharing or not.

As discussed in [I-D.ietf-intarea-shared-addressing-issues], all IP address sharing solutions break protocols which do not use transport numbers. A mitigation solution is to utilize specific ALGs. For ICMP in particular, a mitigation solution would be to rewrite the "Identifier" and perhaps "Sequence Number" fields in the ICMP request, treating them as if they were port numbers.

As a conclusion, this issue can be partially mitigated, likely at par to centralized NAT solutions.

5.7. Provisioning and Operational Systems

5.7.1. Overview

The general claim of this issue is that a service providers' provisioning and accounting systems would need to [radically] evolve to deal with the notions of shared IPv4 addresses and port range constrains.

5.7.2. Discussion

The stateless 4via6 solution relies on a fully operational IPv6 network, which on the IPv6 plane fundamentally does not differ from a

regular IPv6 network, and the stateless 4via6 solution may be seen as an IPv6 application - devices connecting to the network, need unique IPv6 addresses which the network is able to provide. In the 4via6 solution it happens that these unique IPv6 addresses embed an IPv4 address. Hence, additional system enhancements that the stateless 4via6 solution requires, over and above those simply needed to deploy and operate an IPv6 network, lie in the domain of supporting the provisioning of the IPv6 adaptation functionality of the CPEs. This may require the operator to use DHCPv6, or other provisioning methods such as IPv6CP, TR-69, in order to configure any relevant 4via6 service parameters to a CPE.

From an IPv4 perspective, an operator will likely want to have a management system capable of the assignment of IPv4 addresses to the shared pool, and tuning the re-use factor. In this, the solution exhibits no grossly different characteristics than those of any system with an operator managed NAT44 function where similar management capabilities need to be introduced.

One additional aspect of the stateless 4via6 solution needs to be highlighted. On a par basis this solution requires less per subscriber management, accounting and logging capabilities than centralized NAPT44 alternatives such as DS-Lite, due to the following:

- o The assignment of an IPv6 address that embeds a deterministic IPv4 address and port range removes the need for the operator to perform any NAPT44 binding logging, ie the task of determining which user had a given IPv4 address and port at a given time is simply a matter of determining who had the corresponding IPv6 address, rather than collecting large amounts of dynamic binding data.
- o There is no need for the operator to manage NAPT44 binding data access and retention.
- o Given the stateless nature of the 4via6 solution, all subscriber CPEs in an operator's domain can share exactly the same 4via6 service configuration, i.e. The operator does not need to be concerned with managing on a per user basis specific AFTR assignment and/or load balancing such users and throughout ensuring symmetric traffic flows throughout.
- o The location of the NAPT44 function on the user's CPE, allows easy and direct management of the port mappings by the end user removing a need for the operator to introduce PCP [I-D.wing-software-port-control-protocol] (or similar) protocols in on AFTRs, and on CPE devices. In effect the end user can

retains control of any bindings, which could be via today's GUI, or UPnP IGDv2, or even PCP.

- o As and when needed, a stateless 4via6 solution readily supports the assignment of an unshared IPv4 address, and full port control by the end user. A similar capability with centralised NAPT44 solutions involve onerous management of per subscriber configurations on the operator's AFTR.

5.8. Training & Education

5.8.1. Overview

The issue claims a concern with the need for developers and support staff to be trained & educated in dealing with a port constrained systems.

5.8.2. Discussion

There appear to be at least two levels of looking at this issue in the stateless 4via6 context. On one level, it is perfectly true that developers and support staff will need to be trained with running/supporting a native IPv6 network, that is now a basis of the solution. This however is an inherent aspect of deploying an IPv6 network and applications. On another level, support and developers need to be familiarized with the NAPT44 characteristics of the system, that are not different from those already known about such systems today. More specifically, there appears to be no such thing as a port unconstrained carrier grade NAPT44 system, in either tomorrow's stateless 4via6 or AFTR guises, or today's residential CPE NAPT44 implementations that have an inherent hard set translation limit (often 1024 translation, corresponding to a usage of 1024 ports). That application developers should be trained to be reasonably conservative in the usage of ports is thus not an issue of the stateless 4via6 solution, but pretty much of any NAPT44 based solution, even those in use today.

Another useful observation here is that the stateless 4via6 solution, actually allows an operator to retain existing troubleshooting procedures, given which today encompass CPE based NAPT44, rather than changing them radically to an AFTR. Furthermore, it is possible to alleviate any port-range constraints for users by allocating more generous port ranges without the need to manage such users configuration on active core network devices (eg AFTR).

5.9. Security and Port Randomization

5.9.1. Overview

Preserving port randomization [RFC6056] may be more or less difficult depending on the address sharing ratio (i.e., the size of the port space assigned to a CPE). Port randomization may be more difficult to achieve with a stateless solution than stateful solution. The CPE can only randomize the ports inside be assigned a fixed port range.

5.9.2. Discussion

The difference in the random port selection range may be significant in practice and using port-restricted systems without any measures (like random port selection in draft-bajko-pripaddrassign-03) is one of the trade-offs of the mechanism. It should be however noted that even full port unrestricted systems, today, rarely implement random port selection from the full port range, as such the difference is largely theoretical, again viewed from today's perspective. Only with a longer term prospect of devices/hosts adopting random port selection according to RFC 6056 the NAT-based port-restricted mechanisms, will degrade security to a certain extent.

5.10. Unknown Failure Modes

5.10.1. Overview

The issue purports that a system with a port constraints introduces new unknown failure modes, not known with NAT44 or NAPT44 systems, and in general is more complex than such a system.

5.10.2. Discussion

This claim does not appear to have objective technical arguments that can be discussed. A restricted port range system, such as the one assumed in this document, does not appear to have any more or less complexity than any of the other NAPT44 solutions against which the same issue has not been levelled. That is a statement that can be made in consideration of each of those alternative solution network design (eg elaborate routing rules or topologies) and feature implementation complexities, which appear to be no better than that of a stateless 4via6 address port range system. Ultimately, system complexity is something best left adjudicated by the operators choosing to deploy one or the other of these IP based transition solutions.

5.11. Possible Impact on NAT66 use & design

5.11.1. Overview

The notion of a shared address with a constrained port range is seen as possibly bearing influence on use in future schemes involving NAT66, where IPv6 address sharing is in general deemed not to be desired (ie there is good reason to avoid PAT66).

5.11.2. Discussion

The authors do not propose, nor expect to see the IP address sharing characteristic applying to future NAT66/PAT66 discussions and specification. However, having said that it is useful to take a humble step back and consider the general aspect of causality in this context. The direct cause that brought about IPv4 shared address solutions to the fore was a shortage/exhaustion of a limited IPv4 address resource, alongside a failure of the community to migrate IPv4 networks to IPv6 in a timely manner. At the time of writing it is hard to imagine the same occurring with respect to IPv6 address resources, and hopefully the same set of causes will not be allowed to re-occur. This appears to be the only way to ensure that IPv6 address sharing effect does not come to be, as opposed to precluding such notions within the context of today's IPv4 world where the causality is rather clear.

5.12. Port statistical multiplexing and monetization of port space

5.12.1. Overview

An issue attributed to 4V6 solutions is that due to their characteristic of assigning a fixed amount of ports to participating system nodes, the overall pool of ports cannot be dynamically/statistically multiplexed.

A corollary of this claimed issue is the claim that port range constraints will lead to monetization by service providers of such port ranges, for example by charging users based on the number of ports assigned or creating some bronze, silver, gold type of port based service categories.

5.12.2. Discussion

The 4via6 address shared solution indeed limits the ability to "overload" ie statistically multiplex amongst users, the ports available of a given public IPv4 address. This can be seen as a trade off vs dynamic allocation and the need to log (large amounts) of NAT bindings. Furthermore, the solution is meant to be

fundamentally a transitional one for supporting legacy IPv4 users till full migration to IPv6 can occur. As an example, even with a static allocation of ~1000 ports per shared IP user, it allows an operator to effectively multiply by ~64 the current IPv4 unrealizable address space. To put it into a network growth perspective, it allows an operator to support for some 10 years a steady 50% annual increase in users, without requiring new IPv4 addresses. This is likely an alluring (if unlikely) prospect for most, but it demonstrates the fact that even with static port allocations, IPv4 address sharing can go a long way for many operators.

CGN-based solutions, because they can dynamically assign ports, provide better IPv4 address sharing ratio than stateless solutions (i.e., can share the same IP address among a larger number of customers). For Service Providers who desire an aggressive IPv4 address sharing, a CGN-based solution is more suitable than the stateless. However, in case a CGN pre-allocates port ranges, for instance to alleviate traceability complexity it also reduces its port utilization efficiency.

5.13. Readdressing

5.13.1. Overview

Due to the port range encoding being part of the CPE's IPv6 address, any change in the range requires a re-configuration of the CPEs 4via6 address. This is said to be an issue given the impact that IP address changes have on existing traffic flows, as well as general IPv6 network routing

5.13.2. Discussion

It is true that under the assumed notions of the stateless 4via6 solution, IPv6 re-addressing is required to effect a change in terms of the shared IPv4 address or ports. Such changes can and are likely best done using dynamic address configuration methods such as DHCPv6, or alternatively out of band management tools, eg TR-69, especially when the 4via6 address can be derived from a delegated prefix. Using these, the impact of the address change does not translate to a neither a classic IPv6 host renumbering problem nor an unmanageable network renumbering problem. On the CPE, the change only affects the 4via6 address of the CPE and not any end user IPv6 hosts behind the CPE (which would likely continue to derive their IPv6 addresses from an unchanged delegated prefix). On the service provider network side, the change, if any, represents a network renumbering case which the operator can be reasonably expected to handle within their network numbering plan, especially given that the IPv6-prefix of the an IPv4-in-IPv6 address is summarizable.

An addressing change will impact any existing IPv4 flows that are being NAT'ed by the CPE. This is also analogous to the today's practice of IPv4 address changes espoused by some operators, which while not being commendable, is established in the market. Nevertheless, as a means of alleviating such an impact it appears desirable for the solutions to investigate the viability of mechanisms that could allow for more graceful addressing changes.

To facilitate IPv6 summarization and operator appears to have two 4V6 deployment choices. When encoding IPv4 addresses in lower order address space bits that are subject to summarization, the operator would need to assign a modest dedicated IPv6 prefix (such as a /64) as an 4V6 IPv6 addressing sub-domain. Alternatively, without resorting to a separate 4V6 addressing sub-domain, an operator could allow for the IPv4 address embedding to be embedded in a high-order portion of the IPv6 domain address space, one that closely follows the IPv6 domain prefix. These two valid address subnetting and deployment options deserve better description in the solution specifications.

5.14. Ambiguity about communication between devices sharing an IP address.

5.14.1. Overview

A regular IPv4 destination based routed system inherently does not allow two devices to communicate while sharing the same IPv4 address, even if with different ports. Similarly, such a system does not allow on the basis of a IPv4 source address alone to perform address spoofing prevention. These two issues naturally render regular IPv4 based routed networks incapable of supporting a shared address solution.

5.14.2. Discussion

In terms of the IPv4 data plane of the 4via6 solution, the CPE and the stateless gateway components need to be modified in terms of their IPv4 forwarding behaviour. The CPE's NAPT44 function, must be capable of sending traffic towards the IPv6 adaptation function when the traffic is addressed to its (shared) IPv4 address but a different port than the one assigned to the CPE. Similarly, the CPE's NAPT44 function must be capable of receiving traffic addressed from its (shared) IPv4 address but a different port than the one assigned to it.

On the IPv6 data plane the stateless 4via6 solution does not suffer from the issue by the nature of relying on regular IPv6 forwarding. Address-spoofing security can be realized on regular IPv6 devices

plane, in a way which effectively does not allow a CPE to send IPv6 traffic from a source IPv6 address that it has not been assigned. The spoofing of IPv4 addresses can be prevented in this manner in 4via6 solution relying on translation (dIVI). Tunneling 4via6 solutions (4rd) require IPv6+IPv4 source address validation to be performed at the CPE and stateless gateway, by the IPv6 adaptation function.

The conceptual IPv6 adaptation function has many of its core principles already defined either as part of IPinIP tunneling or stateless NAT64 drafts. However additional work, such as defining the port indexing schemes, is needed and is at the heart of what needs to be covered in the individual solution drafts that fall under the stateless 4via6 family. Throughout, no legacy IPv4 end-systems are expected to implement these techniques.

5.15. Other

5.15.1. Abuse Claims

Because the IPv4 address is shared between several customers, and in order to meet the traceability requirement discussed in Section 12 of [I-D.ietf-intarea-shared-addressing-issues], Service Providers must store the assigned ports in addition to the IPv4 address.

If the remote server does not implement the recommendation detailed in [I-D.ietf-intarea-server-logging-recommendations], the Service Provider may be obliged to reveal the identity of all customers sharing the same IP address at a given time.

5.15.2. Fragmentation and Traffic Asymmetry

In order to deliver a fragmented IPv4 packet to its final destination, among those having the same IPv4 address, a dedicated procedure similar to the one defined in Section 3.5 of [RFC6146] is required to reassemble the fragments in order to look at the destination port number.

When several stateless IPv4/IPv6 interconnection nodes are deployed, and because of traffic asymmetry, situations where fragments are not handled by the same stateless IPv4/IPv6 interconnection node may occur. Such context would lead to session breakdowns. As a mitigation, a solution would be to redirect fragments towards a given node which will be responsible for implementing the procedure documented in [RFC6146]. The redirection procedure is stateless.

As a conclusion, this issue can be mitigated.

5.15.3. Multicast Services

IPv4 service continuity must be guaranteed during the transition period, including the delivery of multicast-based services such as IPTV. Because only an IPv6 prefix will be provided to a CPE, dedicated functions are required to be enabled for the delivery of legacy multicast services to IPv4 receivers. This is critical since many of the current IPTV contents are likely to remain IPv4-formatted and there will remain legacy receivers (e.g., IPv4-only Set Top Boxes (STB)) that can't be upgraded or be easily replaced.

This issue is similar to the one encountered in the stateful case, and the same solution can be used to mitigate the issue (e.g., [I-D.qin-software-dslite-multicast]).

As a conclusion, this issue can be solved.

6. Conclusion

As per the discussion in this document, the authors believe that the set of issues specifically attributed to A+P based such as the stateless 4via6 solution with characteristics as per Section 3, either do not apply, or can be mitigated. In several aspects, a stateless 4V6 solution represents a reasonable trade off compared to alternatives in areas such as NAT logging, ease as of deployment and operations, all of which are actually facilitated by such a solution.

In terms of the 4V6 transport mode, both translation and mapped tunnel appear to be share the same key characteristics, but applicable to different contexts. The mapped tunnel mode appears desirable when the operator has no expectations of applying any more elaborate traffic based services, and/or concerned about the loss of IP Options or the use of NAT64 technology. The translation based approach appears particularly attractive to operators who are concerned about integrating traffic into a more elaborate suite of services based on regular IPv6 data-plane functionality, as opposed to specific IPinIP data plane functionality.

7. IANA Considerations

This document does not raise any IANA considerations.

8. Security Considerations

This document does not introduce any security considerations over and

above those already covered by the referenced solution drafts.

9. Contributors and Acknowledgements

The authors thank Dan Wing, Nejc Skoberne, Remi Depres, Xing Li, Jan Zorz, Satoru Matsushima, Mohamed Boucadair, Qiong Sun, and Arkadiusz Kaliwoda for their reviews and draft input.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informative References

[I-D.bajko-pripaddrassign]

Bajko, G., Savolainen, T., Boucadair, M., and P. Levis, "Port Restricted IP Address Assignment", draft-bajko-pripaddrassign-03 (work in progress), September 2010.

[I-D.despres-softwire-4rd]

Despres, R., "IPv4 Residual Deployment across IPv6-Service networks (4rd) A NAT-less solution", draft-despres-softwire-4rd-00 (work in progress), October 2010.

[I-D.ietf-intarea-server-logging-recommendations]

Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging recommendations for Internet facing servers", draft-ietf-intarea-server-logging-recommendations-04 (work in progress), April 2011.

[I-D.ietf-intarea-shared-addressing-issues]

Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", draft-ietf-intarea-shared-addressing-issues-05 (work in progress), March 2011.

[I-D.ietf-softwire-ds-lite-tunnel-option]

Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual- Stack Lite", draft-ietf-softwire-ds-lite-tunnel-option-10 (work in progress), March 2011.

- [I-D.ietf-softwire-dual-stack-lite]
Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", draft-ietf-softwire-dual-stack-lite-11 (work in progress), May 2011.
- [I-D.murakami-softwire-4v6-translation]
Murakami, T., Chen, G., Deng, H., Dec, W., and S. Matsushima, "4via6 Stateless Translation", draft-murakami-softwire-4v6-translation-00 (work in progress), July 2011.
- [I-D.operators-softwire-stateless-4v6-motivation]
Boucadair, M., Matsushima, S., Lee, Y., Bonness, O., Borges, I., and G. Chen, "Motivations for Stateless IPv4 over IPv6 Migration Solutions", draft-operators-softwire-stateless-4v6-motivation-02 (work in progress), June 2011.
- [I-D.qin-softwire-dslite-multicast]
Wang, Q., Qin, J., Boucadair, M., Jacquenet, C., and Y. Lee, "Multicast Extensions to DS-Lite Technique in Broadband Deployments", draft-qin-softwire-dslite-multicast-04 (work in progress), June 2011.
- [I-D.thaler-port-restricted-ip-issues]
Thaler, D., "Issues With Port-Restricted IP Addresses", draft-thaler-port-restricted-ip-issues-00 (work in progress), February 2010.
- [I-D.vixie-dnsexst-dns0x20]
Vixie, P. and D. Dagon, "Use of Bit 0x20 in DNS Labels to Improve Transaction Identity", draft-vixie-dnsexst-dns0x20-00 (work in progress), March 2008.
- [I-D.wing-softwire-port-control-protocol]
Wing, D., Penno, R., and M. Boucadair, "Pinhole Control Protocol (PCP)", draft-wing-softwire-port-control-protocol-02 (work in progress), July 2010.
- [I-D.xli-behave-divi]
Bao, C., Li, X., Zhai, Y., and W. Shang, "dIVI: Dual-Stateless IPv4/IPv6 Translation", draft-xli-behave-divi-03 (work in progress), July 2011.

- [I-D.xli-behave-divi-pd]
Li, X., Bao, C., Dec, W., Asati, R., Xie, C., and Q. Sun,
"dIVI-pd: Dual-Stateless IPv4/IPv6 Translation with Prefix
Delegation", draft-xli-behave-divi-pd-01 (work in
progress), September 2011.
- [I-D.ymbk-aplusp]
Bush, R., "The A+P Approach to the IPv4 Address Shortage",
draft-ymbk-aplusp-10 (work in progress), May 2011.
- [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's
Robustness to Blind In-Window Attacks", RFC 5961,
August 2010.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X.
Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052,
October 2010.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-
Protocol Port Randomization", BCP 156, RFC 6056,
January 2011.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation
Algorithm", RFC 6145, April 2011.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful
NAT64: Network Address and Protocol Translation from IPv6
Clients to IPv4 Servers", RFC 6146, April 2011.

Authors' Addresses

Wojciech Dec
Cisco
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands

Email: wdec@cisco.com

Rajiv Asati
Cisco
Raleigh, NC
USA

Phone:
Fax:
Email: rajiva@cisco.com
URI:

Congxiao Bao
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing, 100084
CN

Phone: +86 10-62785983
Fax:
Email: congxiao@cernet.edu.cn
URI:

Hui Deng
China Mobile
Beijing,
CN

Phone:
Fax:
Email: denghui@chinamobile.com
URI:

Mohamed Boucadair
France Telecom
France

Phone:
Fax:
Email: mohamed.boucadair@orange-ftgroup.com
URI:

