Individual Submission
Internet Draft
Intended status: Informational
Expires: April 2011

E. Jankiewicz (Ed.)
SRI International, Inc.

October 25, 2010

An Annotated Bibliography for IPv4-IPv6 Transition and Coexistence
draft-jankiewicz-v6ops-v4v6biblio-03.txt

Abstract

The Internet is in the early stages of what may be a protracted period of coexistence of IPv4 and IPv6. Network operators are challenged with the task of activating IPv6 without negative impact on operating IPv4 networks and their customers. This draft is an informational “annotated bibliography” compiled to help in the analysis and development of basic guidelines and recommendations for network operators. The goal of this document is to survey the current state of RFCs, Internet-Drafts and external reference materials that define the use cases, problem statements, protocols, transition mechanisms and coexistence tools that will be of interest to a network operator planning to turn on IPv6.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on April 25, 2009.

Copyright Notice

Copyright (c) 2010 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.

Table of Contents

1. Introduction

Since the IPv6 protocol was defined in 1995 as RFC 1883 (replaced in 1998 by RFC 2460) the Internet has been in a long transition from IPv4 to IPv6. In reality, we are still in the early stages of what is likely to be a protracted period of coexistence, where IPv6 penetration in hosts (both servers and clients) will gradually ramp up as networks make IPv6 available through their infrastructures.

Network operators face a daunting task to design and implement plans to activate IPv6 without negative impact on large (in some cases very large) operating IPv4 networks with many live customers. Some basic guidelines and recommendations for network operators are being developed (http://tools.ietf.org/html/draft-lee-v4v6tran-problem) and this draft is an informational companion to that effort. The goal of this document is to survey the current state of RFCs, active (and expired but still relevant) Internet-Drafts and external reference materials that define the use cases, problem statements, protocols, transition mechanisms and coexistence tools that will be of interest to a network operator planning to turn on IPv6.

This is a dynamic and evolving marketplace of ideas. At best, this draft is a blurry snapshot of the landscape near to the time of its publication. The editor intends this compendium to be merely the starting point for an active database or wiki available for community contribution including feedback on the real-world experience of network operators as they turn on IPv6. Note that the links to RFCs and drafts are based on the IETF Tools view of the repository at http://tools.ietf.org/html/. The links for active drafts are not for a specific revision but should link to the last or latest version.

The following sections comprise an annotated bibliography of the currently available documentation to knowledge of the editor. It is provided as informational guidance only, and any network operator contemplating an IPv6 implementation will of course exercise due

diligence in researching all the issues, standards and recommendations and analyze applicability to the particular network operation.

Note that as the body of this text includes full reference information for the bibliography entries these are not included in the normal Reference section.

[Editor's note to be removed before publication:

While this draft is circulating, the editor is interested in any and all pointers to additional useful references. Contributions of capsule summaries and applicability for any of the listed entries would also be appreciated and will be graciously acknowledged. If I have missed anyone who already chipped in, this will be cheerfully rectified upon your reminder via a private e-mail.]

1.1. The Three Laws of IPv4/IPv6 Coexistence Mechanisms

The Editor of this draft thought it might be helpful to briefly explore the motivations driving the current profusion of coexistence mechanisms. In the not so distant past little or no discussion of this topic was going on in the IETF, as many felt the case was closed. A discussion in the Intarea meeting at IETF 71 in Dublin and a presentation at the plenary at that meeting led to a reawakening of interest in coexistence and transition tools. This discussion continued at a special meeting in Montreal in October 2008, and has occupied substantial time on the mailing lists and meetings of several Working Groups since then. The Internet Area, IPv6 Operation (v6ops), Softwires and Behave WGs have generated many contributions, and an ad-hoc discussion mailing list has been established at https://www.ietf.org/mailman/listinfo/v4tov6transition.

Early in the life of IPv6, the assumption was made that IPv6 deployment, based on dual-stack implementations, would be ubiquitous long before the IPv4 address pool would run out. For special cases, tunneling through dissimilar networks or use of an external translation box such as NAT-PT would allow interim operation of legacy equipment. At present, this has not yet come to pass. The impending exhaustion of IPv4 address space renders dual-stack impossible in some deployments and issues have resulted in NAT-PT being deprecated to Historic status.

Nature (and your average Internet-Draft author) abhors a vacuum. With the demise of NAT-PT and the increasing urgency to get moving on IPv6 transition, we are now in a period of "Let 1000 Flowers Bloom" where many ideas are being advanced, and a lot of IETF brainpower is

being spent debating the relative merits and evilness of various approaches. The spectrum of opinion on coexistence mechanisms has two extremes:

IPv4 is so Over: Concentrate on deploying native IPv6 and managing it effectively, rather than spinning more complex webs of IPv4 accommodation. Deploying anything that delays IPv6 and enables more IPv4 usage at this point is irresponsible.

Where’s the Business Case: Real customers need IPv4, there is no IPv6 content, no demand for IPv6. Scale up NAT to keep IPv4 viable, provide some sort of artificial IPv6 access, if and when customers ask. No plans for native IPv6 in the foreseeable future.

A reasonable position recognizes the valid motivation on both sides. An ISP may not be able to dictate updates to customer computers and routers, and must provide access to all legacy customers, not just eager IPv6 adopters, so an interim mechanism that minimizes their inconvenience is needed. One size will never fit all, so some solutions may be a good fit for one ISP, and not for others. While evaluating all the alternative documented here, the principle to keep in mind is that the IETF should provide good engineering opinions on all these alternatives, to permit things that will help, and prevent things that will cause problems.

This can be summed up in the “Three Laws of IPv4/IPv6 Coexistence Mechanisms”:

1. First, do no harm.

2. Keep it simple.

3. Keep moving towards more native IPv6.

“No harm” in this case means that a good solution will not unduly interfere with good experience for the legacy IPv4 customer, nor will it impede the eager IPv6 adopter. The solution must not cause problems for peer or backbone networks or for the Internet community at large.

“Simple” means to solve particular problems with specific solutions focused to the point of need rather than attempting broad and complex methods that impinge on all traffic. However, do not simplify any more than necessary to avoid harm.

The compulsion to move towards native IPv6 follows from the first two laws. Over time, even minimal harm and complexity that even a good

mechanism presents can and should be reduced over time by continuing to enable, promote and encourage transition to native IPv6. Design and deploy your interim solution(s) with a clear migration path that will eventually render them redundant. Set a date after which you will not deploy any new equipment that does not support IPv6. Set a date to sunset IPv4 access, giving legacy customers plenty of time (and incentive) to upgrade their old equipment.

In summary, it seems that the Robustness Principle (Postel's Law) would apply, as it does in many situations:

> "Be conservative in what you do, be liberal in what you accept from others." [RFC 793]

Following the Robustness Principle and the Three Laws should allow an operator complete freedom to manage their own network and to chose and operate any coexistence mechanism as long as they need to for supporting their customers, except where those choices cause harm to someone else. Of course, there is no universal definition of "harm" so reasonable people can disagree, e.g. if a mechanism in use on the access side causes additional delay, content providers may see that as "harming" their users' experience. That's why Working Group mailing lists and IETF meetings are just so much fun.

Oh, and by the way, the Fourth Law should be "Don't reinvent the wheel" so please explore the RFCs, drafts and other citations to see if someone has already proposed something similar to your idea. Your contributions are needed, but time and energy is better spent exploring novel approaches and building on what has already been proposed.

2. IPv6 and related Protocol Specifications

"IPv6 Node Requirements" J. Loughney, Ed. April 2006
http://tools.ietf.org/html/rfc4294

"IPv6 Node Requirements RFC 4294-bis" E. Jankiewicz, J. Loughney, T. Narten
http://tools.ietf.org/html/draft-ietf-6man-node-req-bis

RFC 4294 and its update draft are included by reference. These provide a comprehensive overview of the IPv6 baseline specifications and the reader is directed to them to avoid a redundant listing here.

3. Problem Statements and Use Cases

"Problem Statements of IPv6 Transition of ISP" Y. Lee, Ed.
http://tools.ietf.org/html/draft-lee-v4v6tran-problem

This draft is being developed by an ad-hoc group interested in providing guidance to network operators on the IPv6 transition. It will include high level use cases (as contributed by IETF participants with network operator experience) and a problem statement documenting what additional work IETF could do to provide sufficient tools and guidance for the network operators

"Mobile Networks Considerations for IPv6 Deployment" R. Koodli
http://tools.ietf.org/html/draft-ietf-v6ops-v6-in-mobile-networks

Mobile Internet access from smartphones and other mobile devices is accelerating the exhaustion of IPv4 addresses. IPv6 is widely seen as crucial for the continued operation and growth of the Internet, and in particular, it is critical in mobile networks. This document discusses the issues that arise when deploying IPv6 in mobile networks. Hence, this document can be a useful reference for service providers and network designers.

"Routing Loop Attack using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", G. Nakibly and F. Templin
http://tools.ietf.org/html/draft-ietf-v6ops-tunnel-loops

This document is concerned with security vulnerabilities in IPv6-in-IPv4 automatic tunnels. These vulnerabilities allow an attacker to take advantage of inconsistencies between the IPv4 routing state and the IPv6 routing state. The attack forms a routing loop which can be abused as a vehicle for traffic amplification to facilitate DoS attacks. If automatic tunnels are used in a deployment the warnings and mitigations in this draft should be considered.

"Use Case for IPv6 Transition for a Large-Scale Broadband Network" CC. Huang (Ed.), XY. Li and LM. Hu
http://tools.ietf.org/html/draft-huang-v6ops-v4v6tran-bb-usecase

"IPv6 Transition Cable Access Network Use Cases" Y. Lee and V. Kuarsingh
http://tools.ietf.org/html/draft-lee-v4v6tran-usecase-cable

"IPv6 Transition Use Case for a Large Mobile Network" C. Zhou (Ed.) and T. Taylor
http://tools.ietf.org/html/draft-zhou-v6ops-mobile-use-case-00

Each of these use case drafts is focused on a particular deployment model for a specific market segment. While each may be based on a singular operator's experience or planning, the intention is to develop the set of use cases drafts to be of interest to any network operator in the segment.

"Considerations for Stateless Translation (IVI/dIVI) in Large SP Network" Q. Sun et al.
http://tools.ietf.org/html/draft-sunq-v6ops-ivi-sp

"dIVI" is a prefix-specific and stateless address mapping method based on IVI which can directly translate IPv4 packet to IPv6 packet. This document describes the challenges and requirements for large Service Provider to deploy IPv6 in an operational network and specifically considerations for dIVI deployment.

4. Transition and Coexistence Scenarios and Architectures

RFC 5211 "An Internet Transition Plan." J. Curran, July 2008
http://tools.ietf.org/html/rfc5211

While the abstract for this RFC humbly describes it as just "one possible plan" for the IPv6 transition, it provides very good context and a common language to use when talking about transition plans, and can be seen as a call to action. It describes three phases of the transition, and proposes a timeline based on predictions of the imminent exhaustion of the IPv4 address space. The phases are:

1. Preparation, where IPv4 predominates while service providers trial and experiment with IPv6, and end-users prepare to provide Internet-facing IPv6 services in the future. The timeline in the RFC described this phase as in progress, and optimally this phase would have ended already.

2. Transition, where both IPv4 and IPv6 services are offered and used, with production level support for IPv6, although this may be via transition mechanisms rather than native IPv6. The RFC targeted this phase to end in 2011.

3. Post-Transition, where native IPv6 services should be offered while IPv4 services may still be supported.

"Guidelines for Using Transition Mechanisms During IPv6 Deployment" J. Arkko and F. Baker
http://tools.ietf.org/html/draft-arkko-ipv6-transition-guidelines

IPv6 deployment requires some effort, resources, and expertise. The availability of many different deployment models is one reason why expertise is required. This draft discusses the IPv6 deployment models and migration tools, and recommends ones that have been found to work well in operational networks in many common situations.

“IPv6 Transition Guide For A Large ISP Providing Broadband Access”, G. Yang (Ed.), L. Hu and J. Lin
http://tools.ietf.org/html/draft-yang-v6ops-v4v6tran-bb-transition-guide

This draft is a product of the current v4tov6transition effort and it examines IPv6 migration solutions for each part of the Large-scale broadband infrastructure with a layer 2 access network. The analysis is based on the requirements for providing existing broadband services in v4v6-coexisting or IPv6-only situations. The draft describes the suitable scenarios for each solution.

“IPv6 Transition Guide for a Large Mobile Operator” T. Tsou (Ed.) and T. Taylor
http://tools.ietf.org/html/draft-tsou-v6ops-mobile-transition-guide

Similarly, this draft examines IPv6 migration solutions for a large mobile network.

RFC 6036 “Emerging Service Provider Scenarios for IPv6 Deployment”, B. Carpenter, S. Jiang
http://www.rfc-editor.org/rfc/rfc6036.txt

This document describes practices and plans that are emerging among Internet Service Providers for the deployment of IPv6 services. They are based on practical experience so far, as well as current plans and requirements, reported in a survey of a number of ISPs carried out in early 2010. The document identifies a number of technology gaps, but does not make recommendations.

“Framework for IP Version Transition Scenarios”, B. Carpenter, S. Jiang and V. Kuarasingh
http://tools.ietf.org/html/draft-carpenter-v4v6tran-framework

This document sets out a framework for the presentation of scenarios and recommendations for a variety of approaches to the transition from IPv4 to IPv6, given the necessity for a long period of co-existence of the two protocols.

5. Transition/Coexistence Tools

As network operators and end-users independently proceed with transition to IPv6 while others continue to use IPv4, a potentially long period of coexistence will ensue. Variations on terminology have been used since the specification of IPv6; transition implies a process whereby the star of IPv6 rises and the star of IPv4 sets; coexistence implies that both will operate together. Due to thoroughly discussed limits to the growth of an Internet using only IPv4, IPv6 is a necessary technology for the future of the Internet. However, nothing compels the elimination of IPv4; no protocol police will forbid its use in the foreseeable future. IPv4 may disappear due to irrelevance when IPv6 is so pervasive to make it redundant, but network operators should be prepared to operate IPv4 and IPv6 in a mixed deployment for some time. However, the techniques and mechanisms supported by a network operator can be expected to evolve and change over time as a rational goal would be to gradually shift coexistence costs (real operational expense as well as convenience) from “early adopters” of IPv6 to the shrinking pool of IPv4 maintainers.

Various techniques are required for coexistence, roughly divided into three categories:

1. Address Mapping: Many situations will require the use of address mapping to maintain scalability in the face of dwindling IPv4 global address space and to support translation and tunneling approaches.

2. Tunneling: A method for the encapsulation and transport of one protocol over or through the infrastructure that favors the other, e.g. IPv6 traffic via an IPv4 infrastructure

3. Protocol Translation: A mechanism for rewriting packets from one protocol to the other so they can be delivered as native (non-encapsulated) packets typically due to incompatible end nodes, e.g. an IPv6 client to an IPv4 server.

These categories are not mutually exclusive, as some scenarios and solutions incorporate aspects of multiple approaches.

RFC 4213 “Basic Transition Mechanisms for IPv6 Hosts and Routers” E. Nordmark and R. Gilligan October 2005
http://tools.ietf.org/html/rfc4213

5.1. Address Mapping

The introduction of address family translation presents challenges similar to those experienced with Network Address Translation (NAT) as it has evolved in the IPv4 Internet. The depletion of IPv4 global address space conspires with the continuing need for routable IPv4 address in some coexistence approaches to further press proliferation and scale of NAT. While alternatives exist, some network operators will continue to see the various flavors of NAT as a necessary evil, so it remains important to understand the impact on network operations, on the end-user and on applications.

Dual-Stack Lite (DS-lite) is one of the alternatives to providing dual-stack support to end-users in the face of limited global IPv4 address space.

RFC 2663 “IP Network Address Translator (NAT) Terminology and Considerations” P. Srisuresh and M. Holdrege August 1999 http://tools.ietf.org/html/rfc2663

This document attempts to describe the operation of NAT devices and the associated considerations in general, and to define the terminology used to identify various flavors of NAT.

5.1.1. Address Translation in Network Operations

“Common Requirements for IP Address Sharing Schemes” I. Yamagati et al. http://tools.ietf.org/html/draft-ietf-behave-lsn-requirements

This document defines common requirements of multiple types of Large Scale Network Address Translation (NAT) that handles Unicast UDP, TCP and ICMP.

“Issues with IP Address Sharing” M. Ford (Ed.) et al. http://tools.ietf.org/html/draft-ietf-intarea-shared-addressing-issues

The completion of IPv4 address allocations from IANA and the RIRs is causing service providers around the world to question how they will continue providing IPv4 connectivity service to their subscribers when there are no longer sufficient IPv4 addresses to allocate them one per subscriber. Several possible solutions to this problem are now emerging based around the idea of shared IPv4 addressing. These solutions give rise to a number of issues and this memo identifies those common to all such address sharing approaches.

“An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition", Sheng Jiang, Dayong Guo, Brian Carpenter
http://tools.ietf.org/html/draft-ietf-v6ops-incremental-cgn

Carrier-Grade NAT (CGN) devices with integrated transition mechanisms can reduce the operational change required during the IPv4 to IPv6 migration or coexistence period. This document proposes an incremental CGN approach for IPv6 transition. It can provide IPv6 access services for IPv6-enabled hosts and IPv4 access services for IPv4 hosts while leaving much of a legacy IPv4 ISP network unchanged. It is suitable for the initial stage of IPv4 to IPv6 migration. Unlike NAT444 based CGN alone, Incremental CGN also supports and encourages transition towards dual-stack or IPv6-only ISP networks. A smooth transition to IPv6 deployment is also described in this document.

“Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers” Bagnulo, Matthews, van Beijnum
http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate-stateful

This document describes stateful NAT64 translation, which allows IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP. The public IPv4 address can be shared among several IPv6-only clients. When the stateful NAT64 is used in conjunction with DNS64 no changes are usually required in the IPv6 client or the IPv4 server.

“NAT64-CPE Mode Operation for Opening Residential Service” G. Chen and H. Deng
http://tools.ietf.org/html/draft-chen-v6ops-nat64-cpe

The authors of this draft describe the application of fundamental NAT64 functionality in CPE deployment scenarios. The approach is intended to eliminate the need for CPE to cooperate with DNS64, and to be compatible with legacy residential servers without changes to DNS requirements.

“Flexible IPv6 Migration Scenarios in the Context of IPv4 Address Shortage” M. Boucadair (Ed.) et al, October 20, 2009 (expired)
http://tools.ietf.org/html/draft-boucadair-behave-ipv6-portrange-04

This memo presents a solution to solve IPv4 address shortage and ease IPv4-IPv6 interconnection. The document presents a set of incremental steps for the deployment of IPv6 as a means to solve IPv4 address exhaustion. Stateless IPv4/IPv6 address mapping functions are introduced and IPv4-IPv6 interconnection scenarios presented.

This memo advocates for a more proactive approach for the deployment of IPv6 into operational networks. This memo specifies the IPv6 variant of the A+P. Both encapsulation and translation scheme are covered. Moreover, two modes are elaborated: the binding mode (compatible mode with DS-lite) and the stateless mode.

“A Note on NAT64 Interaction with Mobile IPv6” W. Haddad and C. Perkins
http://tools.ietf.org/html/draft-haddad-mext-nat64-mobility-harmful

This memo discusses potential NAT64 technology repercussions for mobile nodes using Mobile IPv6. An ambiguity is identified related to the use of DNS during bootstrapping, which is likely to inhibit proper signaling between mobile node and home agent.

“NAT64 for Dual Stack Mobile IPv6” B. Sarikaya and F. Xia
http://tools.ietf.org/html/draft-sarikaya-behave-mext-nat64-dsmip

This memo specifies how IPv6 only mobile nodes (MN) receiving host-based mobility management using Dual Stack Mobile IPv6 (DSMIPv6) can communicate with IPv4 only servers. The protocol is based on home agents maintaining a table similar to NAT64 and linking it to the binding cache. This technique avoids the problems encountered when NAT64 is used for mobile nodes in Dual Stack Mobile IPv6. How IPv6 only mobile nodes can receive multicast data from IPv4 only content providers is also explained.

“NAT64 for Proxy Mobile IPv6” B. Sarikaya and F. Xia
http://tools.ietf.org/html/draft-sarikaya-behave-netext-nat64-pmip

Similarly, this memo specifies how IPv6 only mobile nodes (MN) receiving network-based mobility management using Proxy Mobile IPv6 (PMIPv6) can communicate with IPv4 only servers.

5.1.2. Application and End-User Considerations With NAT

“Problem Statement for Referrals” B. Carpenter, S. Jiang and B. Zhou
http://tools.ietf.org/html/draft-carpenter-referral-ps

The purpose of a referral is to enable a given entity in a multiparty Internet application to pass information to another party. It enables a communication initiator to be aware of relevant information of its destination entity before launching the communication. This memo discusses the problems involved in referral scenarios.

“Referrals Across an IPv6/IPv4 Translator” D. Wing, October 19, 2009
http://tools.ietf.org/html/draft-wing-behave-nat64-referrals-01

While this draft is expired, this issue remains a topic of conversation, including a Bar-BoF at IETF 78. Referrals across disparate address domains may be needed for provision of services such as SIP during transition.

“Legacy NAT Traversal for IPv6: Simple Address Mapping for Premises Legacy Equipment (SAMPLE)”
http://tools.ietf.org/html/draft-carpenter-softwire-sample

IPv6 deployment is delayed by the existence of millions of subscriber network address translators (NATs) that cannot be upgraded to support IPv6. This document specifies a mechanism for traversal of such NATs. It is based on an address mapping and on a mechanism whereby suitably upgraded hosts behind a NAT may obtain IPv6 connectivity via a stateless server, known as a SAMPLE server, operated by their Internet Service Provider. SAMPLE is an alternative to the Teredo protocol.

“Some Considerations on the Load-Balancer for NAT64” D. Zhang et al.
http://tools.ietf.org/html/draft-wang-behave-nat64-load-balancer

This draft investigates issues with deploying load-balancers with NAT64 devices.

“An FTP ALG for IPv6-to-IPv4 Translation” I. van Beijnum
http://tools.ietf.org/html/draft-ietf-behave-ftp64

The File Transfer Protocol (FTP) has a very long history, and despite the fact that today, other options exist to perform file transfers, FTP is still in common use. As such, it is important that in the situation where some client computers are IPv6-only while many servers are still IPv4-only and IPv6-to-IPv4 translators are used to bridge that gap, FTP is made to work through these translators as best it can. This document specifies a middlebox that enables legacy usage of FTP with translation.

“Assessing the Impact of NAT444 on Network Applications” C. Donley et al. http://tools.ietf.org/html/draft-donley-nat444-impacts

NAT444 is an IPv4 extension technology being considered by Service Providers to continue offering IPv4 service to customers while transitioning to IPv6. This technology adds an extra Large-Scale NAT ("LSN") in the Service Provider network, thereby resulting in two NATs. CableLabs, Time Warner Cable, and Rogers Communications independently tested the impacts of NAT444 on many popular Internet services using a variety of test scenarios, network topologies, and vendor equipment. This document identifies areas where adding a

second layer of NAT disrupts the communication channel for common Internet applications.

5.1.3. Dual-Stack Lite (DS-lite)

“Understanding Dual-Stack Lite” Jeff Doyle, Network World October 22, 2009 http://www.networkworld.com/community/node/46600

This article provides a good introduction to DSlite, at the time of its publication. Please see the following drafts for details and more current work.

“Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion” A. Durand et al.
http://tools.ietf.org/html/draft-ietf-softwire-dual-stack-lite

This document revisits the dual-stack model and introduces the dual-stack lite technology aimed at better aligning the costs and benefits of deploying IPv6 in service provider networks. Dual-stack lite enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT).

“Dual-stack Lite Mobility Solutions” B. Sarikaya and F. Xia October 11, 2009 (expired)
http://tools.ietf.org/html/draft-sarikaya-softwire-dslitemobility-01

Two solutions are presented to show how to use Dual-Stack Lite transition technique in mobile networks: one for Proxy Mobile IPv6 and the other for Dual-Stack Mobile IPv6. Proxy Mobile IPv6 allows IPv4 nodes to receive mobility services using an IPv4 home address. In case of client based mobility using DSMIPv6, mobile node is a dual-stack node and it can receive an IPv4 home address from the home agent which is co-located with DS-lite carrier-grade NAT.

“Scalable Operation of Address Translators with Per-Interface Bindings” J. Arkko and L. Eggert February 9, 2009 (expired)
http://tools.ietf.org/html/draft-arkko-dual-stack-extra-lite-00

This document explains how to employ address translation in networks that serve a large number of individual customers without requiring a correspondingly large amount of private IPv4 address space.

“Gateway Initiated Dual-Stack Lite Deployment” F. Brockners et al.
http://tools.ietf.org/html/draft-ietf-softwire-gateway-init-ds-lite

Gateway-Initiated Dual-Stack lite (GI-DS-lite) is a modified approach to the original Dual-Stack lite (DS-lite) applicable to certain tunnel-based access architectures. GI-DS-lite extends existing access tunnels beyond the access gateway to an IPv4-IPv4 NAT using softwires with an embedded context identifier, that uniquely identifies the end-system the tunneled packets belong to. The access gateway determines which portion of the traffic requires NAT using local policies and sends/receives this portion to/from this softwire tunnel.

“Deployment DS-lite in Point-to-Point Access Network” Y. Lee (Ed.) et al. http://tools.ietf.org/html/draft-zhou-softwire-ds-lite-p2p

Gateway-Initiated Dual-Stack lite (GI-DS-lite) is a proposal to logically extend existing access tunnels beyond the access gateway to DS-Lite Address Family Transition Router element (AFTR) using softwires with an embedded context identifier. This memo describes a deployment model using GI-DS-lite in Point-to-Point access network.

“Deploying Dual-Stack Lite in IPv6 Network” M. Boucadair (Ed.) et al. http://tools.ietf.org/html/draft-boucadair-dslite-interco-v4v6

Dual-Stack lite requires that the AFTR must have IPv4 connectivity. This forbids a service provider who wants to deploy AFTR in an IPv6-only network. This memo proposes an extension to implement a stateless IPv4-in-IPv6 encapsulation in the AFTR so that AFTR can be deployed in an IPv6-only network.

“IPv6 RA Option for DS-lite AFTR Element” Y. Lee, M. Boucadair and X. Xu http://tools.ietf.org/html/draft-lee-6man-ra-dslite

This document specifies a new optional extension to IPv6 Router Advertisement messages to allow IPv6 routers to advertise DS-Lite AFTR addresses to IPv6 hosts (i.e., a default IPv6 route for DS-Lite traffic). The provisioning of the AFTR address is crucial to access IPv4 connectivity services in a DS-Lite context. Means to ensure reliable delivery of this information to connecting hosts is a must.

Furthermore, this RA option can be used as a means to distribute DS-Lite serviced customers among a set of deployed AFTRs without requiring a central knowledge of the underlying topology and deployed AFTRs.

5.2. Tunneling Mechanisms

RFC 2473 “Generic Packet Tunneling in IPv6 Specification.” A. Conta and S. Deering, December 1998
http://tools.ietf.org/html/rfc2473

This document defines the model and generic mechanisms for IPv6 encapsulation of Internet packets, such as IPv6 and IPv4. The model and mechanisms can be applied to other protocol packets as well, such as AppleTalk, IPX, CLNP, or others.

RFC 2529 “Transmission of IPv6 over IPv4 Domains without Explicit Tunnels” B. Carpenter and C. Jung March 1999.
http://tools.ietf.org/html/rfc2529

This memo specifies the frame format for transmission of IPv6 packets and the method of forming IPv6 link-local addresses over IPv4 domains. The motivation for this method is to allow isolated IPv6 hosts, located on a physical link which has no directly connected IPv6 router, to become fully functional IPv6 hosts by using an IPv4 domain that supports IPv4 multicast as their virtual local link.

RFC 3056 “Connection of IPv6 Domains via IPv4 Clouds” B. Carpenter and K. Moore February 2001
http://tools.ietf.org/html/rfc3056

This memo specifies an optional interim mechanism for IPv6 sites to communicate with each other over the IPv4 network without explicit tunnel setup, and for them to communicate with native IPv6 domains via relay routers.

RFC 3053 “IPv6 Tunnel Broker” A. Durand, I. Guardini and D. Lento January 2001
http://tools.ietf.org/html/rfc3053

The IPv6 global Internet as of today uses a lot of tunnels over the existing IPv4 infrastructure. Those tunnels are difficult to configure and maintain in a large scale environment, and the process is too complex for the isolated end user. The motivation for the development of the tunnel broker model is to help early IPv6 adopters to hook up to an existing IPv6 network with stable, permanent IPv6 addresses and DNS names.

5.2.1. Teredo

RFC 4380 “Teredo: Tunneling IPv6 over UDP” C. Huitema February 2006
http://tools.ietf.org/html/rfc4380

This RFC defined a service that enables nodes located behind one or more IPv4 Network Address Translations (NATs) to obtain IPv6 connectivity by tunneling packets over UDP; we call this the Teredo service. Running the service requires the help of "Teredo servers" and "Teredo relays". The Teredo servers are stateless, and only have to manage a small fraction of the traffic between Teredo clients; the Teredo relays act as IPv6 routers between the Teredo service and the "native" IPv6 Internet. The relays can also provide interoperability with hosts using other transition mechanisms such as "6to4". Teredo client capability has been included in Windows operating systems since Windows XP and public servers are available.

RFC 5991 “Teredo Security Extensions” D. Thaler, S. Krishnan and J. Hoagland September 2010
http://tools.ietf.org/html/rfc5991

The Teredo protocol defines a set of flags that are embedded in every Teredo IPv6 address. This document specifies a set of security updates that modify the use of this flags field, but are backward compatible.

“Teredo Extensions”, D. Thaler
http://tools.ietf.org/html/draft-thaler-v6ops-teredo-extensions

This document specifies a set of extensions to the Teredo protocol. These extensions provide additional capabilities to Teredo, including support for more types of Network Address Translations (NATs), and support for more efficient communication.

5.2.2. IPv6 Rapid Deployment (6rd)and Extensions

IPv6 Rapid Deployment (6rd) is an approach that allows a service provider to quickly roll out an IPv6 service offering. Free, a large French ISP, successfully deployed a 6rd offering in 5 weeks. It is also being used in a current IPv6 trial offered by Comcast in the USA.

“How 6rd Eases the Transition to IPv6” Mike Capuano on Cisco SP360 blog, August 5, 2010
http://blogs.cisco.com/sp/how_6rd_eases_the_transition_to_ipv6/

This article provides a quick overview of 6rd. The fundamental protocol specification and initial implementation experience can be found in RFC 5969 and 5569.

RFC 5969 “IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)– Protocol Specification” W. Townsley and O. Troan August 2010 http://tools.ietf.org/html/rfc5969

RFC 5569 “IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)” R. Despres January 2010 http://tools.ietf.org/html/rfc5569

“IPv6 Across NAT44 CPEs (6a44)” R. Despres, B. Carpenter and S. Jiang http://tools.ietf.org/html/draft-despres-softwire-6a44

IPv6 Across NAT44 CPEs (6a44) 6a44 is based on an address mapping and on a mechanism whereby suitably upgraded hosts behind a NAT may obtain IPv6 connectivity via a stateless 6a44 server function operated by their Internet Service Provider. With it, traffic between two 6a44 hosts in a single site remains within the site. Except for IANA numbers that remain to be assigned, the specification is intended to be complete enough for running codes to be independently written and interwork.

[Note that this draft converges and supersedes work started in two separate drafts, which are no longer relevant: http://tools.ietf.org/html/draft-despres-softwire-6rdplus-00 http://tools.ietf.org/html/draft-carpenter-softwire-sample-00]

“UDP Encapsulation of 6rd” Y. Lee and P. Kapoor http://tools.ietf.org/html/draft-lee-softwire-6rd-udp-02

This memo specifies the UDP encapsulation to IPv6 Rapid Deployment (6rd) protocol which enables hosts behind unmodified Home Gateway device to access 6rd service. One variation (Server Model) avoids host modification by offloading the implementation to a small server (relay) on the home LAN.

“Gateway Initiated 6rd” T. Tsou et al. http://tools.ietf.org/html/draft-tsou-softwire-gwinit-6rd

This document proposes an alternative to the deployment model defined in RFC 5969 for 6rd. This model extends existing access tunnels beyond an operator-owned gateway collocated with the operator's IPv4 network edge to the Border Router. This modification makes it unnecessary to provide IPv4 routes to IPv6 UEs. The gateway serves as an aggregation point for IPv4 routing.

5.2.3. Tunnel Support Protocol (TSP)

RFC 5572 “IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP)” M. Blanchet and F. Parent, February 2010
http://tools.ietf.org/html/rfc5572

TSP is an Experimental RFC defining a method for a tunnel client to negotiate tunnel characteristics with a tunnel broker. It enables tunnels in various deployment architectures including NAT traversal and mobility, and for user authentication it utilizes:

RFC 4422 “Simple Authentication and Security Layer (SASL)” A. Melikov ad K. Zeilenga(Eds.) June 2006
http://tools.ietf.org/html/rfc4422

5.2.4. Residual IPv4 Deployment over IPv6-only Infrastructure

Further down the transition road, operators may desire to retire IPv4 routing support and move their backbone networks to IPv6-only. There may be residual IPv4 legacy customers (clients and servers) still requiring the delivery of IPv4 packets. While the previously proposed Dual-Stack Transition Mechanism (DSTM) approach attempted to satisfy this use case, it was complex and stateful. A stateless approach to IPv4 residual deployment (4rd) is defined in section 3.2 of the Stateless Address Mapping (SAM) draft. At the time of this publication, several network operators in Japan are planning implementation to support residual IPv4 customers.

“Stateless Address Mapping (SAM) - a Simplified Mesh-Softwire Model” Despres, R. July 12, 2010
http://tools.ietf.org/html/draft-despres-softwire-sam

“IPv4 Residual Deployment across IPv6-Service networks (4rd): A NAT-less Solution” R. Despres
http://tools.ietf.org/html/draft-despres-softwire-4rd

5.2.5. Address Plus Port (AplusP)

“The A+P Approach to the IPv4 Address Shortage” R. Bush (Ed.) October 27, 2009 (expired, but authors indicate a new draft is coming)
http://tools.ietf.org/html/draft-ymbk-aplusp

This draft discusses the possibility of address sharing by treating some of the port number bits as part of an extended IPv4 address (Address plus Port, or A+P). Instead of assigning a single IPv4

address to a customer device, we propose to extended the address by "stealing" bits from the port number in the TCP/UDP header, leaving the applications a reduced range of ports. This means assigning the same IPv4 address to multiple clients (e.g., CPE, mobile phones), each with its assigned port-range. In the face of IPv4 address exhaustion, the need for addresses is stronger than the need to be able to address thousands of applications on a single host. If address translation is needed, the end-user should be in control of the translation process - not some smart boxes in the core.

“Aplusp Lite – A light weight aplusp approach” Z. Xiaoyu http://tools.ietf.org/html/draft-xiaoyu-aplusp-lite

This document proposes a solution aimed at providing IPv4 continuity in IPv6 environment. The proposed solution is expected to alleviate the public IPv4 depletion problem while maximize the benefits from IPv6 deployment, and meet the desired service availability and reliability with affordable cost.

5.2.6. IRON-RANGER and ISATAP Solutions

A body of RFCs and drafts in progress provide an alternative approach to IPv4/IPv6 coexistence. This approach utilizes tunneling techniques to create “overlay” networks. While currently considered “Experimental” it may be of interest to network operators as an alternative network architecture.

RFC 5214 “Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)” F. Templin et al. March 2008 http://tools.ietf.org/html/rfc5214

RFC 5579 “Transmission of IPv4 Packets over Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) Interfaces” F. Templin (Ed.) February 2010 http://tools.ietf.org/html/rfc5579

RFC 5320 “The Subnetwork Encapsulation and Adaptation Layer (SEAL)” F. Templin (Ed.) February 2010 http://tools.ietf.org/html/rfc5320

Fred Templin originally published SEAL as an Experimental RFC, and is currently updating with the intention to publish as Standards Track: http://tools.ietf.org/html/draft-templin-intarea-seal

RFC 5558 “Virtual Enterprise Traversal (VET)” F. Templin (Ed.) February 2010 http://tools.ietf.org/html/rfc5558

Fred Templin originally published VET as an Informational RFC, and is currently updating with the intention to publish as Standards Track: http://tools.ietf.org/html/draft-templin-intarea-vet

RFC 5720 “Routing and Addressing in Networks with Global Enterprise Recursion (RANGER)” F. Templin (Ed.) February 2010 http://tools.ietf.org/html/rfc5720

“The Internet Routing Overlay Network (IRON)” F. Templin (Ed.) http://tools.ietf.org/html/draft-templin-iron

5.2.7. Softwires Hub and Spoke with L2TP

RFC 5571 “Softwire Hub and Spoke Deployment Framework with Layer Two Tunneling Protocol Version 2 (L2TPv2)” B. Storer et al. June 2009 http://tools.ietf.org/html/rfc5571

This document describes the framework of the Softwire "Hub and Spoke" solution with the Layer Two Tunneling Protocol version 2 (L2TPv2). The implementation details specified in this document should be followed to achieve interoperability among different vendor implementations.

5.3. Translation

From the earliest specification of IPv6 IETF contributors have recognized that translation would be a necessary tool for transition and coexistence, as IPv6 was designed as an incompatible replacement rather than an extension of IPv4. The original approach to stateless translation defined in RFC 2765 and its implementation as NA(P)T-PT as described in RFC 2766 had a number of issues that resulting in the approach being deprecated by RFC 4966. Recently the Behave WG has taken on the work of defining a set of scenarios covering the use cases for translation, prioritizing the work and defining new solutions that overcome the deficiencies of the historic approach.

5.3.1. Historic Approach

RFC 2765 “Stateless IP/ICMP Translation (SIIT).” E. Nordmark, February 2000 http://tools.ietf.org/html/rfc2765

This document specifies a transition mechanism algorithm in addition to the mechanisms already specified in RFC 1933 (note that this reference was subsequently obsoleted by RFC 2893 which in turn was obsoleted by RFC 4213). The algorithm translates between IPv4 and IPv6 packet headers (including ICMP headers) in separate translator "boxes" in the network without requiring any per-connection state in those "boxes". This new algorithm can be used as part of a solution that allows IPv6 hosts, which do not have a permanently assigned IPv4 addresses, to communicate with IPv4-only hosts. The document neither

specifies address assignment nor routing to and from the IPv6 hosts when they communicate with the IPv4-only hosts.

SIIT has been applied in several translation implementations, including the historic NAT-PT specified in RFC 2766 and deprecated by RFC 4966. SIIT is currently being revised in “IP/ICMP Translation Algorithm” X. Li, C. Bao and F. Baker
http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate

RFC 2766 “Network Address Translation – Protocol Translation (NAT-PT).” G. Tsirtsis and P. Srisresh, February 2000
http://tools.ietf.org/html/rfc2766

This solution attempted to provide transparent routing to end-nodes in an IPv6 realm trying to communicate with end-nodes in an IPv4 realm and vice versa. This combined Network Address Translation and Protocol Translation. While it did mandate dual-stack support or special purpose routing requirements (such as requiring tunneling support) on end nodes, it did introduce issues that were considered harmful enough to lead to its deprecation in July 2007 by RFC 4966 “Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status” http://tools.ietf.org/html/rfc4966.

RFC 2767 “Dual-Stack Hosts Using ‘Bump in the Stack’ Technique (BIS)” K. Tsuchiay, H. Higuchi and Y. Atarashi February 2000

RFC 3338 “Dual-Stack Hosts Using ‘Bump in the API’ (BIA)” S. Lee, et al. October 2002
http://tools.ietf.org/html/rfc3338

These two RFCs are proposed for obsolescence by a draft that combines both:

“Dual-Stack Hosts Using ‘Bump in the Host’(BIH)” B. Huang, H. Deng and T. Savolainen
http://tools.ietf.org/html/draft-ietf-behave-v4v6-bih

5.3.2. Current Translation Approaches

A renewed effort to define new translation mechanisms started with discussions in the Internet Area (intarea) meeting and the Technical Plenary at IETF 71 in Dublin, and continued at a special meeting in Montreal in October 2008. This led to a commitment by contributors in the Behave WG to take on the work. A set of scenarios were defined along with a framework for the translation solutions.

"IPv4 Run-Out and IPv4-IPv6 Co-Existence Scenarios" J. Arkko and M. Townsley
http://tools.ietf.org/html/draft-arkko-townsley-coexistence

When IPv6 was designed, it was expected that the transition from IPv4 to IPv6 would occur more smoothly and expeditiously than experience has revealed. The growth of the IPv4 Internet and predicted depletion of the free pool of IPv4 address blocks on a foreseeable horizon has highlighted an urgent need to revisit IPv6 deployment models. This document provides an overview of deployment scenarios with the goal of helping to understand what types of additional tools the industry needs to assist in IPv4 and IPv6 co-existence and transition.

This document was originally created as input to the Montreal co-existence interim meeting in October 2008, which led to the rechartering of the Behave and Softwire working groups to take on new IPv4 and IPv6 coexistence work. This document is published as a historical record of the thinking at the time.

"A Framework for IPv4/IPv6 Translation" F. Baker et al.
http://tools.ietf.org/html/draft-ietf-behave-v6v4-framework

This draft (Framework) is the place to start to understand the historic context for translation, the definition and rationale for the set of translation scenarios and canonical definitions for some of the terminology that arises when talking about translation and coexistence in general.

The 4 deployment modes for these scenarios are:

1. Connecting between the IPv4 Internet and the IPv6 Internet

2. Connecting an IPv6 network to the IPv4 Internet

3. Connecting an IPv4 network to the IPv6 Internet

4. Connecting between an IPv4 network and an IPv6 network

As solutions may differ with respect to the initiating end of the conversation, 8 scenarios are defined in the Framework draft, as recapped in the following sections along with specifications that fit each scenario.

Some general specifications that are cited in the various solution specifications (or may be in subsequent revisions) are:

"IPv6 Addressing of IPv4/IPv6 Translators" C. Bao et al. August 16, 2010 http://tools.ietf.org/html/draft-ietf-behave-address-format-10

"DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers" M. Bagnulo et al. July 5, 2010 http://tools.ietf.org/html/draft-ietf-behave-dns64-10

"Analysis of 64 Translation" R. Penno, T. Saxena and D. Wing http://tools.ietf.org/html/draft-penno-behave-64-analysis

Due to specific problems, NAT-PT was deprecated by the IETF as a mechanism to perform IPv6-IPv4 translation. Since then, new effort has been undertaken within IETF to standardize alternative mechanisms to perform IPv6-IPv4 translation. This document evaluates how the new translation mechanisms avoid the problems that caused the IETF to deprecate NAT-PT.

5.3.2.1. An IPv6 network to the IPv4 Internet

The Framework defines Scenario 1 for an early adopter (end user or network operator) which establishes an IPv6 network and needs to maintain access to the global IPv4 Internet, preferably without assigning IPv4 addresses to the nodes of the IPv6 network. Either the Stateful or Stateless solutions proposed may satisfy this deployment scenario.

"Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers" M. Bagnulo, P. Matthews and I. van Beijnum http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate-stateful

"IP/ICMP Translation Algorithm" X. Li, C. Bao and F. Baker http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate

5.3.2.2. The IPv4 Internet to an IPv6 network

The Framework defines Scenario 2 for a node on the IPv4 Internet initiating a transmission to a node on an IPv6 network. The original approach to this deployment was the NAT-PT implementation of SIIT (as defined in RFC 2766) which has been deprecated (by RFC 4966). The Stateless Translation solution for Scenario 1 also would work for this case as it does support IPv4-initiated communication with a subset of IPv6 addresses.

5.3.2.3. The IPv6 Internet to an IPv4 network

The Framework defines Scenario 3 where a legacy IPv4 network has a requirement to provide services to users in the IPv6 Internet.

Stateful Translation with static AAAA records in DNS to represent the IPv4-only hosts will work.

“Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers” M. Bagnulo, P. Matthews and I. van Beijnum http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate-stateful

“DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers” M. Bagnulo et al. http://tools.ietf.org/html/draft-ietf-behave-dns64

Alternatively, host-based translation (BIH) or tightly-coupled translators may be considered.

5.3.2.4. An IPv4 network to the IPv6 Internet

Scenario 4 is not easy to solve but fortunately will not arise until significant IPv6 uptake. In-network translation is not viable, and other techniques should be considered including host-based translation (BIH) or tightly-coupled translators that adapt legacy hosts or networks to the IPv6 Internet.

5.3.2.5. An IPv6 network to an IPv4 network

Scenario 5 describes a configuration where both the IPv6 network and IPv4 network are within the administrative control of the same organization. It appears amenable to the same solutions proposed for Scenario 1.

5.3.2.6. An IPv4 network to an IPv6 network

Scenario 6 is the mirror image of Scenario 5, with communication initiated from the IPv4 side. It appears amenable to the same solution proposed for Scenario 2.

5.3.2.7. The IPv6 Internet to the IPv4 Internet

The Framework indicates that Scenario 7, the interconnection of the IPv4 Internet with the IPv6 Internet may appear to be an ideal case for an in-network translator (such as the deprecated NAT-PT), but there is no viable way to map the immense IPv6 address space onto IPv4. This situation would not entail until significant IPv6 adoption, and has not been a priority for solution.

5.3.2.8. The IPv4 Internet to the IPv6 Internet

Scenario 8 presents a challenge similar to Scenario 7.

5.4. Connectivity Checking and Delay Avoidance

One important issue that arises in a coexistence environment is negative impact on the initiation of peer-to-peer connections, such as VoIP, video, etc. The initiator doesn't know a priori whether the peer is using the same address family incurring a possible delay as the first attempt may fail. There is also ambiguity, as the IPv6 path may be temporarily broken.

"IPv6 Connectivity Check and Redirection by HTTP Servers" E. Vyncke http://tools.ietf.org/html/draft-vyncke-http-server-64aware

Rather than forcing the client to decide whether IPv4 or IPv6 is more convenient to reach a web server; this document proposes to let the web server check whether there is IPv6 connectivity to the client; then the web server can do a HTTP redirect to the force the client to use IPv6.

This is done easily by a script within the server HTML pages and does not require any change in the client applications or configuration. The client still can control whether he/she wants to enabled IPv6.

"Happy Eyeballs: Trending Towards Success (IPv6 and SCTP)", D. Wing, A. Yourtchenko, P. Natarajan.
http://tools.ietf.org/html/draft-wing-http-new-tech

This draft makes several recommendations to ensure user satisfaction and a smooth transition from HTTP's pervasive IPv4 to IPv6 and from TCP to SCTP. While the target audience is app developers and content providers, network operators should be aware of techniques needed to maintain peaceful coexistence without negative impact on end-user perception of service level.

"Migrating SIP to IPv6 Media Without Connectivity Checks" D. Wing, A. Yourtchenko
http://tools.ietf.org/html/draft-wing-dispatch-v6-migration

During the migration from IPv4 to IPv6, it is anticipated that an IPv6 path might be broken for a variety of reasons, causing endpoints to not receive RTP data. Connectivity checks would detect and avoid the user noticing such a problem, but there is industry reluctance to implement connectivity checks.

This document describes a mechanism allowing dual-stack SIP endpoints to attempt communications over IPv6 and fall back to IPv4 if the IPv6 path is not working. The mechanism does not require connectivity checks.

6. Prefix and Address Assignment and Distribution

RFC 4291 “IP Version 6 Addressing Architecture.” R. Hinden, S. Deering. February 2006.
http://tools.ietf.org/html/rfc4291

RFC 5952 “A Recommendation for IPv6 Text Representation” S. Kawamura and M. Kawashima, August 2010
http://tools.ietf.org/html/rfc5952

RFC 4291 defines the addressing architecture of the IP Version 6 (IPv6) protocol. The document includes the IPv6 addressing model, text representations of IPv6 addresses, definition of IPv6 unicast addresses, anycast addresses, and multicast addresses, and an IPv6 node's required addresses. RFC 5952 updates RFC 4291 with a recommended method for rendering IPv6 addresses in a standard form for user interfaces, logging and reporting.

“IPv6 Addressing of IPv4/IPv6 Translators” C. Bao et al. (Status: Standards Track, in RFC Editor will update RFC 4291)
http://tools.ietf.org/html/draft-ietf-behave-address-format

This document discusses the algorithmic translation of an IPv6 address to a corresponding IPv4 address, and vice versa, using only statically configured information. It defines a well-known prefix for use in algorithmic translations, while allowing organizations to also use network-specific prefixes when appropriate. Algorithmic translation is used in IPv4/IPv6 translators, as well as other types of proxies and gateways (e.g., for DNS) used in IPv4/IPv6 scenarios.

RFC 3177 “IAB/IESG Recommendations on IPv6 Address Allocations to Sites.” IAB, IESG. September 2001.
http://tools.ietf.org/html/rfc3177

RFC 3177 provides recommendations to the addressing registries (APNIC, ARIN and RIPE-NCC) on policies for assigning IPv6 address blocks to end sites. In particular, it recommends the assignment of /48 in the general case, /64 when it is known that one and only one subnet is needed and /128 when it is absolutely known that one and only one device is connecting.

“IPv6 Address Assignment to End Sites", T. Narten, G. Huston, R. Roberts, 12-Jul-10
http://tools.ietf.org/html/draft-ietf-v6ops-3177bis-end-sites

The proposed update to RFC 3177 revises the recommendation to leave the exact choice to the operational community. The role of the IETF

is limited to providing guidance on IPv6 architectural and operational considerations. This document reviews the architectural and operational considerations of end site assignments as well as the motivations behind the original 3177 recommendations. Moreover, the document clarifies that a one-size-fits-all recommendation of /48 is not nuanced enough for the broad range of end sites and is no longer recommended as a single default.

RFC 4192 “Procedures for Renumbering an IPv6 Network without a Flag Day” F. Baker, E. Lear and R. Droms http://www.ietf.org/rfc/rfc4192.txt

RFC 5942 “IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes.” H. Singh, W. Beebee, E. Nordmark. July 2010. http://tools.ietf.org/html/rfc5942

IPv6 specifies a model of a subnet that is different than the IPv4 subnet model. The subtlety of the differences has resulted in incorrect implementations that do not interoperate. This document spells out the most important difference: that an IPv6 address isn't automatically associated with an IPv6 on-link prefix. This document also updates (partially due to security concerns caused by incorrect implementations) a part of the definition of "on-link" from RFC 4861.

RFC 4862 “IPv6 Stateless Address Autoconfiguration.” S. Thomson, T. Narten, T. Jinmei. September 2007. http://tools.ietf.org/html/rfc4862

RFC 4941 “Privacy Extensions for Stateless Address Autoconfiguration in IPv6.” T. Narten, R. Draves, S. Krishnan. September 2007. http://tools.ietf.org/html/rfc4941

The IPv6 addressing architecture presumes that the remaining 64 bits are an endpoint interface identifier. This could be the MAC Address (EUI-64 Address) in an appropriate encoding, or it could be what is called a "privacy address", which is a random number. You will find the most common approach to that, for hosts, in this RFC.

RFC 3315 “Dynamic Host Configuration Protocol for IPv6 (DHCPv6).” R. Droms (Ed.), J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney. July 2003. http://tools.ietf.org/html/rfc3315

“Analysis of Solution Proposals for hosts to learn NAT64 Prefixes” J. Korhonen (Ed.) and T. Savolainen http://tools.ietf.org/html/draft-korhonen-behave-nat64-learn-analysis

Hosts and applications may benefit from the knowledge if an IPv6 address is synthesized, which would mean a NAT64 is used to reach the IPv4 network or Internet. This document analyses number of proposed solutions for communicating if the synthesis is taking place, used address format, and the IPv6 prefix used by the NAT64 and DNS64. This enables both NAT64 avoidance and intentional utilization by allowing local IPv6 address synthesis.

7. How-to, Whitepapers and FAQs

“IPv6 Rollout: Where do we start?” O. Crepin-Leblond
http://www.slideshare.net/ocl999/suggestion-for-an-ipv6-roll-out

“Everything Sysadmin” T. Limoncelli
http://everythingsysadmin.com/2009/01/google-enables-ipv6-for-most-s.html
http://everythingsysadmin.com/2010/08/methods-of-converting-to-ipv6.html

"IPv6 Deployment in Internet Exchange Points (IXPs)", Roque Gagliano
http://tools.ietf.org/html/draft-ietf-v6ops-v6inixp

This draft suggests that in an Internet Exchange Point one might use an address that helps in debugging routing exchanges. One could also look at what other folks do, embedding identifying marks in addresses. For example, Facebook includes “face:b00c” in the IID portion of their address.

8. Experiments, Trials and Prototypes

6bone (concluded)
http://go6.net/ipv6-6bone/

Hurricane Electric (ongoing)
http://www.he.net/

T-Mobile USA (ongoing)
http://groups.google.com/group/tmoipv6beta

Comcast (ongoing)
http://www.comcast6.net/

Internode ADSL (Ongoing)
http://ipv6.internode.on.net/access/adsl/

Verizon FiOS (small scale test – concluded)
http://newscenter.verizon.com/press-releases/verizon/2010/verizon-begins-testing-ipv6.html

“Considerations for Stateless Translation (IVI/dIVI) in Large SP Network” Q. Sun et al.
http://tools.ietf.org/html/draft-sunq-v6ops-ivi-sp

In addition to the deployment use case this draft describes, the draft documents an experimental use of the translation in a research network.

Measurements of IPv6 Path MTU Discovery Behavior
http://www.ripe.net/ripe/meetings/ripe-60/presentations/Stasiewicz-Measurements_of_IPv6_Path_MTU_Discovery_Behaviour.pdf

9. Implementation Reports

“A Basic Guideline for Listing ISPs that Run IPv6” S. Kawamura
http://tools.ietf.org/html/draft-kawamura-ipv6-isp-listings

This draft attempts to gather information about currently known sites that rate ISP readiness for IPv6 and to look at their evaluation methods. This document also summarizes basic guidelines that these listings may consider when checking an ISPs IPv6 readiness. As the draft says, there are many opinions about what it means to be ready for IPv6, and it would be helpful to evaluate ISPs based on some common criteria.

IPv6 Rapid Deployment
http://tools.ietf.org/html/rfc5569

Google has hosted a meeting of IPv6 Implementers in 2009 and 2010, several presentations covered experimental or live transition experience.
https://sites.google.com/site/ipv6implementors/2009/agenda
https://sites.google.com/site/ipv6implementors/2010/agenda

10. Books on IPv6

Blanchet, Marc. “Migrating to IPv6: a Practical Guide to Implementing IPv6 in Mobile and Fixed Networks.” Chichester, England: J. Wiley & Sons, 2006. Print.

Hagen, Silvia. “IPv6 Essentials – Second Edition” Sebastapol, CA: O’Reilly Media, Inc, 2006. Print.

Loshin, Peter. “IPv6, Second Edition: Theory, Protocol and Practice” Morgan Kaufmann Publishing, 2003

Popoviciu, Ciprian, Eric Levy-Abengoli and Patrick Grossetete “Deploying IPv6 Networks” Indianapolis, IN: Cisco Press, 2006. Print.

Siil, Karl A. “IPv6 Mandates: Choosing a Transition Strategy, Preparing Transition Plans, and Executing the Migration of a Network to IPv6.” Indianapolis, IN: Wiley, 2008. Print.

11. Miscellaneous

See the Dancing Turtle, but only if you have native IPv6!
http://www.kame.net/

A little more detail than a Dancing Turtle, on your IPv6 readiness can be obtained by using this site put up by Jason Fesler:
http://test-ipv6.com/

There is an extension for Firefox (and perhaps other browsers) that displays the IP address of web pages you visit, clearly indicating when you are connected via IPv4 or IPv6. In Firefox, click on Tools..Add-ons..Extensions and search for ShowIP.

Eric Vyncke is collecting some statistics on IPv6 penetration.
http://www.vyncke.org/ipv6status/

A reasonable estimation of how fast the sky is falling.
http://www.potaroo.net/tools/ipv4/

A graphical representation of IPv4 depletion.
http://www.ipv4depletion.com/old.html

“IPv6 Adoption Remains Slow, Survey Says” W. Jackson, GCN Sept. 5, 2101
http://gcn.com/articles/2010/09/14/adoption-of-ipv6-is-slow.aspx
http://www.nro.net/documents/GlobalIPv6SurveySummaryv2.pdf

Some troubling, yet interesting news about what operators and end-user organizations are thinking about IPv6 adoption at this time.

A study of some of the brokenness around Path MTU Discovery
http://www.ripe.net/ripe/meetings/ripe-60/presentations/Stasiewicz-Measurements_of_IPv6_Path_MTU_Discovery_Behaviour.pdf

Cluenet hosts a mailing list with IPv6 operator participation. Various transition-related topics are brought up there from time to time.http://lists.cluenet.de/mailman/listinfo/ipv6-ops

“IPv6 for Dummies, Part 1: It’s Time!” http://www.xtranormal.com/watch/7201125/

“IPv6 for Dummies, Part 2: Comparing IPv4 and IPv6” http://www.xtranormal.com/watch/7210035/

12. Security Considerations

This draft does not introduce any security considerations.

13. IANA Considerations

This draft does not require any action from IANA.

[Note to RFC Editor: this section may be removed.]

14. Conclusions

This draft is merely the starting point for a network operator planning an IPv6 rollout. The intention of the editor was to document the great work that is already available that can help in the process and to perhaps save a few hours of redundant effort for someone to find this information. Of course, this will be out of date before it is published as active research continues in coexistence and transition tools. The editor hopes it is at least a useful “You Are Here” map to help navigate the thrill rides available in the IPv6 theme park.

This compendium could serve as an initial set of data to populate an active database or wiki. This would allow continuing community contribution including feedback on the real-world experience of network operators as they turn on IPv6.

15. References

15.1. Normative References

None.

15.2. Informative References

Complete reference information is included in the body of the draft.

16. Acknowledgments

This bibliography is a recapitulation of the contributions of the authors of the cited RFCs, drafts, websites and other publications and many folks on the v6ops and v4v6transition mailing lists, the editor has freely borrowed abstract and summary text from the cited works and e-mail postings. In addition, the editor wishes to acknowledge significant contributions and suggestions from Fred Baker, Brian Carpenter, Remi Despres, Suresh Krishnan, Tina Tsou, Yiu Lee, Marc Blanchet, Med Boucadair, Fred Templin, Andrew Yourtchenko and many contributors on the v4v6trans mailing list. All credit is due to those contributors while the editor takes responsibility for any errors, omissions or mischaracterization of the work in the process of abstracting and summarizing it here.

The IPv4-IPv6 Transition mailing list archive can be found at: https://www.ietf.org/mailman/listinfo/v4tov6transition and the readers are also directed to the mailing list archives of the various IETF Working Groups mentioned for the history of the cited drafts and RFCs.

This document was prepared using 2-Word-v2.0.template.dot.

Author’s Address

Edward J. Jankiewicz
SRI International, Inc.
333 Ravenswood Ave
Menlo Park, CA USA

Phone: 732-389-1003 or 650-859-2000
Email: edward.jankiewicz@sri.com