

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.,
Petitioner,

v.

MONARCH NETWORKING SOLUTIONS LLC,
Patent Owner.

U.S. Patent No. 8,451,845

“Method of receiving a data packet in an IPv6 domain, an associated device and an associated home gateway”

**DECLARATION OF RON BONICA IN SUPPORT OF PETITION FOR
INTER PARTES REVIEW OF U.S. PATENT NO. 8,451,845**


I, Ron Bonica, declare and state as follows:

1. I am a Distinguished Engineer at Juniper Networks, Inc. (“Petitioner”).

I am an active participant in the Internet Engineering Task Force (“IETF”) and have been since 1998. For six years, I served as the co-director of the IETF Operations and Management Area. I am currently co-chairing the IETF working groups for V6OPS and OPSEC. A capture of the IETF website showing the various roles I have held over these years is attached as **Appendix A**.

2. I understand that much of the analysis done in an *inter partes* review must be evaluated from the perspective of a person of ordinary skill in the art (“POSITA”). I understand that a POSITA is a hypothetical person who is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity—not an automaton. I further understand that Dr. Bhattacharjee has provided a declaration in connection with this proceeding and has stated that “a POSITA at the time of the invention (*i.e.*, in June 2008) would have had a Bachelor of Science degree in electrical engineering, computer engineering, computer science, or a related subject, and three years of work experience in computer networking, including experience with the IPv4 and IPv6 protocols, or a related field. Lack of professional experience can be remedied by additional education, and vice versa.” EX1003, ¶55. I meet these criteria.

3. I am—and was, as of June 30, 2008—familiar with the process of submitting documents, including Internet Drafts, to the IETF. Indeed, I have authored and co-authored many IETF Internet Drafts. The first Internet Draft that I submitted was called “ICMP Extensions for MultiProtocol Label Switching.” **Appendix B** is a capture of the IETF website’s Datatracker for this Internet Draft, which shows the draft was initially submitted on February 1999, and that subsequent versions of the draft were uploaded in the months that followed. Since that time I have authored or co-authored over 100 additional Internet Drafts, of which twenty have issued as a Request For Comments (“RFC”), seven are active Internet Drafts, and seventy-five are expired Internet Drafts. A capture of the IETF website listing the IETF drafts I have authored or co-authored can be found at **Appendix A**. I have submitted at least one IETF draft every year since 1999, including multiple drafts in 2008. **Appendix A** illustrates a timeline with my Internet Draft activity throughout the years, and I have reproduced below a more detailed version of that chart:



4. I understand that Petitioner has submitted as Exhibit 1038 the file titled “draft-xli-ivi-00-7431-cancelled.txt.” (“Li-2”). The staging folder lists the last

modified date of Li-2 as “2008-06-21.” EX1044, 13. This indicates that Li-2 was available in the staging folder by June 21, 2008. Moreover, although the words “cancelled” appear in the title of the file, this merely indicates that the authors chose not to pursue Li-2 as an RFC. It does not mean the submission was removed from the staging folder. To the contrary, as I discuss further below, Li-2 would have remained publicly accessible in the staging folder. Participants of the IETF were aware of the staging folder at the time and could have accessed Li-2 freely and without a password (so too could general members of the public). Indeed, Li-2 remained accessible in the staging folder until at least January 2024. I note that the IETF recently updated the staging folder so that it only shows recent submissions and not prior submitted IETF drafts such as Li-2. However, this update occurred only recently and prior to this update Li-2 was available in the staging folder, as reflected in the capture of the staging folder that Petitioner has filed as Exhibit 1044. My discussion of the staging folder in the paragraphs below relates to how it existed and operated during the relevant time period in 2008 and until January of 2024—prior to the IETF update that removed prior drafts that were in the staging folder.

5. The IETF provides a submission tool for uploading IETF drafts, which was available in June 2008. The submission tool included an “Upload” button that a user could use to select and upload an IETF draft in multiple file formats. When an author uploaded a draft to the IETF, the draft was stored in the staging folder.

The IETF then sent the author (or authors) an e-mail that linked to the document in the staging folder asking them to confirm the submission. See EX1055, 1-2. The staging folder could be accessed at the following link: <https://www.ietf.org/staging/>. The staging folder also included a “last modified” column that reflects the date on which the submissions were last modified.

6. If the authors confirmed the submission, it moved from the staging folder into the IETF Internet Drafts repository, and the draft was announced to the community. If the author “canceled” the submission, the submission was not moved to the IETF Internet Drafts repository. However, the draft remained available in the IETF staging folder. The staging folder was not password protected and could be accessed by any IETF participant or member of the public.

7. I am familiar with the staging folder and its operation, including during the relevant time period in 2008. I have submitted drafts that I later “cancelled,” which were uploaded to the staging folder.

8. For instance, the staging folder includes the file titled “draft-kumar-mpls-fec-to-nhlfe-mib-00-16143-cancelled.txt.” EX1044, 46. This is an Internet Draft that I submitted with another IETF participant (Subodh Kumar) on July 6, 2009, as is reflected on the first page of the IETF Draft (reproduced below):

Network Working Group
Internet-Draft
Obsoletes: 3814 (if approved)
Expires: January 7, 2010

S. Kumar, Ed.
R. Bonica
Juniper Networks
July 6, 2009

`draft-kumar-mpls-fec-to-nhlfe-mib-00`

EX1045, 1; EX1044, 46. On July 6, 2009, the file was uploaded with the file name: draft-kumar-mpls-fec-to-nhlfe-mib-00. Because we decided not to confirm the submission, it was renamed to add the word “cancelled,” such that the staging folder now displays the title “draft-kumar-mpls-fec-to-nhlfe-mib-00-16139-cancelled.txt.” The addition of the word “cancelled” to the file name does not alter the substance or text of the original disclosure on July 6, 2009, as reflected by the “last modified” date of the file: “2009-07-06.” The file continued to be publicly available in the staging folder in the same form that it was originally uploaded, and could be accessed without a password and by any participant of the IETF or general public. The same would have been true for Li-2 (*i.e.*, it would have remained accessible to the public in the IETF staging folder as of June 21, 2008).

9. A POSITA would have been aware of the IETF’s staging folder. The IETF is a preeminent standards organization that establishes technical standards for the Internet. A POSITA would have known about the IETF and kept up with technical developments submitted to the IETF. Each year, including in 2008, thousands of people—many of whom are POSITAs—participate in the IETF. See *Introduction to the IETF*, IETF, available at <https://www.ietf.org/about/introduction/>

(last accessed Jan. 11, 2024) (“In any one year, over 7000 people actively participate in the IETF either by authoring a document, engaging in a mailing list discussion, or attending a meeting.”). The July 2008 IETF 72 meeting alone, for example, had 1,182 participants. See *Meeting Statistics*, IETF, available at <https://datatracker.ietf.org/stats/meeting/overview/> (last accessed Jan. 11, 2024). The staging folder was known among POSITAs. Indeed, when POSITAs, as authors, uploaded documents to the IETF, the IETF would send a confirmation e-mail that linked to the staging folder. See EX1055, 1-2. Moreover, the staging folder was mentioned multiple times in e-mails that POSITAs sent to each other, for example e-mails that POSITAs sent to IETF mailing lists and working groups comprising other POSITAs. The record is replete with such e-mails. EXS1046-1055. In fact, I was a POSITA and IETF participant at the time and was aware of and had accessed the IETF staging folder by June of 2008.

10. Thus, POSITAs would have either monitored the IETF’s staging folder for new document uploads or would have browsed the staging folder for documents relevant to the specific field of networking that the POSITA was interested in. For example, Li-2 is titled “draft-xli-ivi-00-7431-cancelled.txt.” A POSITA would have recognized that “xli” represents the primary author’s name, X. Li or Dr. Xing Li. A POSITA would have recognized Dr. Li’s name, as he had authored numerous influential networking papers at the time, including papers on IPv4 and IPv6 and the

transition between the two protocols. A POSITA would have further recognized that “ivi” refers to the IPv4 to IPv6 transition technology, IVI, where IV represents 4, VI represents 6, and the combination IVI represents the transition from 4 to 6. Indeed, I was aware of Dr. Li and IETF submissions that he had made at that time. For example, attached as **Appendix C** is a presentation that I gave in 2008 regarding the transition from IPv4 to IPv6. On slide 19, I refer to a draft by Dr. Li regarding the IVI transition method. Thus, a POSITA would have been aware from the document’s name in the staging folder that Li-2 is pertinent to IPv4 to IPv6 transition mechanisms and other networking developments.

11. Moreover, once a draft is uploaded to the IETF—even before it is officially converted into an Internet Draft—it is added to the IETF Datatracker, which is a sophisticated tool that allows users to search drafts by several criteria, including working group, author, and technical area. EX1056. This process is shown in Exhibit 1055, which is a capture of an email from the IETF Mailing Archive. On February 22, 2008, a group of authors submitted a draft with the title “draft-jiang-p2psip-sep-01.txt.” EX1055, 1-2. The IETF sent the authors an e-mail with a link to the draft in the staging folder as well as a link to a page on the IETF Datatracker that tracked the submission. *Id.* If the author cancels the submission the Datatracker notes that the submission was withdrawn (*i.e.*, cancelled). After a given period, it is removed from the Datatracker, but remains in the staging folder.

Thus, in addition to being able to locate Li-2 in the staging folder, Li-2 would have also been indexed on the IETF Datatracker for a given period of time, and a POSITA would have been able to find Li-2 using the IETF Datatracker's search function.

12. In conclusion, in my opinion a POSITA would have been aware of the IETF staging folder in June 2008, would have been able to find Li-2 in the staging folder as of June 21, 2008, and would have also been able to find Li-2 using the IETF Datatracker's search function at that time.

13. I have personal knowledge of the facts stated in this declaration and could testify competently to them if asked to do so. My employment with Juniper and compensation are not contingent upon the substance of my declaration, any statements made, or any outcome of this matter.

14. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

* * *

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true.

I further declare that these statements were made with the knowledge that willful false statements and the like are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

Respectfully submitted,

Date: February 9, 2024

/s/ Ron Bonica
Ron Bonica



Ron Bonica

Ron Bonica has participated in the IETF since 1998. His areas of expertise include IP Routing, IPv6, Operational Security and network operations.

Ron lives in Northern Virginia. He enjoys yoga, and drawing.



[Ron Bonica](#)

Roles

[RFCs \(20\)](#)

[Active Internet-Drafts \(7\)](#)

[Expired Internet-Drafts \(75\)](#)

[Internet-Draft Activity](#)

Roles

Role ▾	Group ▾	Email ▾
Chair	IPv6 Operations (v6ops)	rbonica@juniper.net
Chair	Operational Security Capabilities for IP Network Infrastructure (opsec)	rbonica@juniper.net
Materials Manager	IPv6 Operations (v6ops)	rbonica@juniper.net
Reviewer	Internet Area Directorate (intdir)	rbonica@juniper.net
Reviewer	Ops Directorate (opsdir)	rbonica@juniper.net
Reviewer	Routing Area Directorate (rtgdir)	rbonica@juniper.net
Tech Advisor	Network Virtualization Overlays (nvo3)	rbonica@juniper.net

RFCs (20)

RFC ▾	Date ▾	Title ▾	Cited by ▾
RFC 3609	Oct 2003	Tracing Requirements for Generic Tunnels	
RFC 4684	Nov 2006	Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)	17 RFCs
RFC 4797	Jan 2007	Use of Provider Edge to Provider Edge (PE-PE) Generic Routing Encapsulation (GRE) or IP in BGP/MPLS IP Virtual Private Networks	2 RFCs
RFC 4884	Apr 2007	Extended ICMP to Support Multi-Part Messages	19 RFCs
RFC 4950	Aug 2007	ICMP Extensions for Multiprotocol Label Switching	6 RFCs
RFC 5837	Apr 2010	Extending ICMP for Interface and Next-Hop Identification	11 RFCs
RFC 5925	Jun 2010	The TCP Authentication Option	72 RFCs

RFC 6291	Jun 2011	Guidelines for the Use of the "OAM" Acronym in the IETF	38 RFCs
RFC 6890	Apr 2013	Special-Purpose IP Address Registries	19 RFCs
RFC 7112	Jan 2014	Implications of Oversized IPv6 Header Chains	8 RFCs
RFC 7543	May 2015	Covering Prefixes Outbound Route Filter for BGP-4	1 RFC
RFC 7588	Jul 2015	A Widely Deployed Solution to the Generic Routing Encapsulation (GRE) Fragmentation Problem	2 RFCs
RFC 7676	Oct 2015	IPv6 Support for Generic Routing Encapsulation (GRE)	10 RFCs
RFC 7746	Jan 2016	Label Switched Path (LSP) Self-Ping	
RFC 8161	May 2017	Benchmarking the Neighbor Discovery Protocol	
RFC 8190	Jun 2017	Updates to the Special-Purpose IP Address Registries	2 RFCs
RFC 8335	Feb 2018	PROBE: A Utility for Probing Interfaces	1 RFC
RFC 8642	Aug 2019	Policy Behavior for Well-Known BGP Communities	
RFC 8900	Sep 2020	IP Fragmentation Considered Fragile	13 RFCs
RFC 9502	Nov 2023	IGP Flexible Algorithm in IP Networks	

Active Internet-Drafts (7)

- [draft-ietf-6man-comp-rtg-hdr](#)
- [draft-ietf-mpls-lspping-norao](#)
- [draft-li-ivy-power](#)
- [draft-bonica-6man-vpn-dest-opt](#)
- [draft-bonica-spring-srv6-end-dtm](#)
- [draft-mirsky-mpls-bfd-bootstrap-clarify](#)
- [draft-salih-spring-srv6-inter-domain-sids](#)


Expired Internet-Drafts (75)

- [draft-ietf-6man-hbh-header-handling](#)
- [draft-ietf-6man-nd-opt-validation](#)
- [draft-ietf-spring-compression-analysis](#)
- [draft-ietf-spring-compression-requirement](#)
- [draft-hegde-lsr-asla-any-app](#)
- [draft-zzhang-intarea-generic-delivery-functions](#)
- [draft-bonica-lsr-crh-isis-extensions](#)
- [draft-bonica-6man-ext-hdr-update](#)
- [draft-ietf-opsec-ipv6-nd-security](#)
- [draft-bonica-6man-deprecate-router-alert](#)
- [draft-bonica-6man-crh-helper-opt](#)
- [draft-bonica-spring-sr-mapped-six](#)
- [draft-bonica-intarea-func-path](#)
- [draft-bonica-6man-seg-end-opt](#)
- [draft-ssangli-bess-bgp-vpn-srm6](#)
- [draft-smith-6man-in-flight-eh-insertion-harmful](#)
- [draft-nayak-intarea-probe-by-vfi](#)
- [draft-alston-idr-crh-bgp-signalling](#)
- [draft-bonica-intarea-lossless-pmtud](#)
- [draft-bonica-6man-oam](#)
- [draft-aghule-intarea-oam](#)
- [draft-kompella-spring-dhcp](#)
- [draft-leddy-6man-truncate](#)
- [draft-bonica-6man-unrecognized-opt](#)
- [draft-rosen-bess-secure-l3vpn](#)
- [draft-spiriyath-ipsecme-dynamic-ipsec-pmtu](#)
- [draft-gont-6man-ipv6-opt-transmit](#)
- [draft-wkumari-long-headers](#)
- [draft-bonica-tcp-auth](#)
- [draft-ietf-l3vpn-auth](#)
- [draft-liu-yfeng-dhcpv6-close-guidance](#)

- [draft-liu-v6ops-dhcpv6-slaac-guidance](#)
- [draft-gont-6man-lla-opt-validation](#)
- [draft-liu-bonica-dhcpv6-slaac-problem](#)
- [draft-bonica-6man-frag-deprecate](#)
- [draft-mpls-ipv6-only-gap](#)
- [draft-bb-2544like-production-tests](#)
- [draft-bonica-v6-multihome](#)
- [draft-kumar-mpls-fec-to-nhlfe-mib](#)
- [draft-berger-l3vpn-ip-tunnels](#)
- [draft-ietf-opsec-nmasc](#)
- [draft-ietf-l3vpn-l3vpn-auth](#)
- [draft-ietf-ccamp-tunproto](#)
- [draft-marques-ppvpn-rt-constrain](#)
- [draft-bonica-tunproto](#)
- [draft-bonica-tunneltrace](#)
- [draft-bonica-l3vpn-auth](#)
- [draft-pan-lsp-ping](#)
- [draft-kb-ppvpn-l2vpn-motiv](#)
- [draft-worster-ppvpn-layers](#)
- [draft-bonica-icmp-mpls](#)

(Excluding replaced Internet-Drafts.)

Internet-Draft Activity



[Abstract](#)[Authors](#)

Extensions for MultiProtocol Label Switching

602

[Email expansions](#)[History](#)

Type Expired Internet-Draft (individual) Expired & archived

Authors [Ron Bonica](#) , [Dan Tappan](#) , [Der-Hwa Gan](#)

Last updated 2000-11-13

RFC stream (None)

Intended (None)

RFC status

Formats [txt](#) [htmlized](#) [pdf](#) [bibtex](#) [bibxml](#)

Stream state (No stream defined)

Consensus Unknown
boilerplate

RFC Editor (None)

Note

IESG

[IESG state](#) Expired

Telechat (None)
date

Responsible (None)
AD

Send notices (None)
to

[Email authors](#)[IPR](#)[References](#)[Referenced by](#)[Nits](#)[Search lists](#)

This Internet-Draft is no longer active. A copy of the expired Internet-Draft is available in these formats:

[txt](#) [htmlized](#) [pdf](#) [bibtex](#) [bibxml](#)

Abstract



Datatracker

Sign in

... through the MPLS user plane. Although these extensions are not being proposed as Internet standards, they are
... here because they have been implemented by several vendors and deployed by several operators.

S

[ica](#)

[pan](#)

[i Gan](#)

... the e-mail addresses provided for the authors of this Internet-Draft may no longer be valid.)

[IETF](#)

[IESG](#)

[IAB](#)

[IRTF](#)

[IETF LLC](#)

[IETF Trust](#)

[RFC Editor](#)


[IANA](#)

[Privacy Statement](#)

[About IETF Datatracker](#)

[Version 11.16.0 \(release - 7d5afaf\)](#)

Report a bug: [GitHub](#) [Email](#)



IPv6 Transition: Gradual Steps to a Big Change

Robert Raszuk
Ron Bonica

Juniper Networks' IPv6 View

The story so far....

- **IPv6 is not an Optional IP Feature**
 - IPv6 is an essential part of any serious IP support
 - Performance, features must be on par with IPv4
- **IPv6 support has been available for many years in:**
 - JUNOS for CPE, Edge and Core Routers
 - JUNOSe for Subscriber Management
 - ScreenOS for Security products
- **Drive has been Government led:**
 - APAC:
 - Japan
 - South Korea
 - Taiwan
 - People's Republic of China
 - US – Federal Government Programs


Juniper Networks IPv6 deployments include:

- **US: Mainly Federal Government and Research Networks**
 - Department of Energy Sciences Network (ESnet)
 - Department of Defense
 - Defense Research and Engineering Network (DREN)
 - Internet 2
 - Canadian Network for the Advancement of Research, Industry, and Education (CANARIE)
 - Pacific Northwest
- **EMEA: Research Networks**
 - GEANT & Delivery of Advanced Network Technology to Europe (DANTE)
 - **NORDUnet, the IP infrastructure for research and education networks in the Nordics**
- **APAC**
 - Japan-European IP research backbone
 - Chunghwa Telecom's nationwide broadband network (HiNet)
 - NTT Communications
- **International Networks:**
 - Telefonica International
 - NTT/Verio Global IP Network backbone

Juniper Networks' IPv6 View It's serious this time.....



The IPv4 Consumption Model



- From <http://www.potaroo.net/presentations/2007-10-23-ipv4-depletion.pdf>

Juniper Networks' IPv6 View

Pressure is increasing

- **Fixed line voice, video and data are converged on IP**
 - But services are not yet converged – video is mainly inside a walled garden
- **Mobile Voice (and Video) is not yet converged on IP**
- **Mobile IP will require millions or billions of care-of addresses**
 - Potential user base larger than entire IPv4 address space
- **Mobile IPv4 solution: Foreign Agent**
 - Allows multiple mobile nodes to share a care-of address
 - NAT-like function, NAT-like problems
- **Mobile IPv6: No Foreign Agent**
 - Inheriting IPv6 addresses happens automatically
 - Smooth transition from one foreign network to another

Juniper Networks' IPv6 View

Current IPv6 Support - Summary

- **Common and Standard models supported in JUNOS and JUNOSe for:**
 - Business Internet Access
 - Fixed and Dynamic L3 VPN access
 - Residential Internet Access and Wholesale
 - IGP, BGP and MPLS Routing
 - Dual Stack and Transition Models (e.g. 6PE) across an IP and/or MPLS core
- **Due to limited deployment esp in EMEA, some of the customer specific “bells and whistles” that we have developed over the last 10 years for IPv4 are not complete for IPv6**
 - Some may not be necessary for IPv6 or new deployments

Possible Responses to IPv4 Address Space Exhaustion

- **Rapid, universal deployment of IPv6**
 - Wishful thinking
- **Continued use of the current, IPv4 infrastructure**
 - Painful
 - Possible
 - More NATs !!!
- **IPv4 ↔ IPv6 Coexistence**
 - An attractive alternative

Rapid, Universal Deployment of IPv6

- During a short transition period, most hosts deploy “dual stacks”
 - Some tunneling may be required
- Some say that after transition period, IPv4 stacks slowly disappear (**I do not agree too much with it**)
- For years, this transition plan has been recommended by IETF
- This transition plan has years of non-deployment to speak for it
 - Enterprises and ISPs are not economically motivated to be the first to transition
 - v6 www services, wikis, blogs are not there yet in critical mass !

Continued Use of IPv4 infrastructure

- **Address space trading**
 - Holders of large, sparsely populated address blocks are motivated to renumber and sell unused space
 - May lead to significant increase of routing table size
 - May change the economics of the Internet
- **More IPv4-to-IPv4 NAT**
 - Breaks some things, but ex: Skype can get around it
 - But these things are well-known as NAT is widely deployed
- **Buys the community some time**
 - But sooner or later, address space scarcity becomes an issue again

IPv4↔IPv6 Coexistence

- **In the short term, most traffic is IPv4**
 - Adapt business and operational procedures to compensate for shortage of IPv4 address space
- **Community develops better transition mechanisms**
 - IPv4 hosts can communicate with IPv6 hosts
 - These mechanisms don't exist yet
 - When they do, enterprises and ISPs will be motivated to migrate to IPv6
- **IPv6 traffic levels increase slowly**
- **More content becomes available over IPv6**

How Should the Community Prepare for Co-existence?

- As a community, adopt a plan similar to the one that is described in **RFC 5211**
- RFC 5211 describes a preparation phase, a transition phase, and a post transition phase
- Each phase details responsibilities for enterprises and network service providers
 - To provide IPv6-based services
 - To provide IPv6-based connectivity
 - To provide IPv6 support
- The plan described in RFC 5211 is particularly well suited to small, closed communities (e.g., R&E networks)

To NAT in the core or not to NAT

<http://ran.psg.com/~randy/080820.alt-to-cgn.pdf>

A Better Approach than Carrier-Grade-NAT

Olaf Maennel	Randy Bush	Luca Cittadini	Steven M. Bellovin
T-Labs Berlin, Germany olaf@maennel.net	IIIJ Tokyo, Japan randy@psg.com	Universita' Roma Tre Rome, Italy luca.cittadini@gmail.com	Columbia University New York, USA smb@cs.columbia.edu

„Some prominent ones target towards carrier-grade-NATs (CGN), e.g. [1], which we feel is a bad idea.”

„1.1 Why Carrier-Grade-NATs are harmful”

With NAT in the core, if a customer wants a new application protocol which requires cooperation from the NAT, they get to beg help from the broadband providers' engineers and lawyers, and all other users of carrier grade NATs. This is the ultimate horror the NAT-haters fear, and they are not all that wrong.

What Else Can a Community Do To Prepare for Co-existence

- Offer services over dual stacks
- Offer incentives to organizations providing IPv6 services
- Monitor emerging transition standards
 - **address-port-borrowing-protocol (APBP)**
 - Dual-Stack Lite
 - IVI
 - NAT6
 - NAT64
 - sNAT-PT
- **Experiment with emerging transition mechanisms**
 - More on these later....

Is Juniper Prepared to Support IPv6 Operations?

- **All JUNOS platforms support IPv6**
 - M-Series
 - T-Series
 - MX-Series
- **JUNOSe supports IPv6**
 - E-Series
- **J-series access router supports IPv6**
- **(Waiting to hear about BX)**

Will All JUNOS Features Work Under IPv6?

- **Most IP routing features will work**
 - Notable exceptions are MPLS related routing protocols (LDP, RSVP)
- **All QoS features will work**
- **A limited set of firewall filtering functionality will work**
 - Firewall filtering limited to fields contained by the IPv6 header and the first “next-header”

So, Tell Me More About Transition Mechanisms

- **Initially, the IETF saw a need for transition mechanisms**
- **And a thousand flowers blossomed**
 - Configured tunnels (6bone)
 - Automatic tunnels (requires derived IPv6 addressing)
 - 6to4
 - ISATAP
 - Tunnel Broker
 - NAT-PT with SIIT
 - Transport relay translators
 - Torpedo

Were Any Transition Mechanisms Widely Deployed?

- No. There was no economic incentive for anyone to move to IPv6. So, there was no incentive to experiment with transition mechanisms.
- NAT-PT with SIIT was reclassified to HISTORIC
- V6OPS WG stopped entertaining new transition proposals
- IETF adopts dual-stack as recommended transition mechanism
- Years of non-deployment follow

What Has Changed?

- **IPv4 address space exhaustion approaches rapidly**
- **The V6OPS WG produces a transition requirements document**
 - draft-bagnulo-behave-nat64
- **IETF Internet, Transport, and O&M areas co-operate to create a new transition mechanism**
- **Many new proposals now on the table**

New Proposals

- **draft-durand-dual-stack-lite**
- **draft-droms-sofwires-snat**
- **draft-baker-behave-ivi**
- **draft-xli-behave-ivi**
- **draft-jennings-behave-nat6**
- **draft-miyata-v6ops-snatpt**

Dual Stack Lite

- Assumes that each host has an IPv6 address and a non-routable IPv4 address
- Hosts communicate with IPv4 ↔ IPv4 NAT boxes over IPv6 tunnels
 - Possibly cascading NAT boxes
- No need for IPv4 ↔ IPv6 NAT

Softwires

- Also assumes that each host has an IPv6 address and a non-routable IPv4 address
- At each endpoint, a gateway performs IPv4↔IPv4 translation, as well as encapsulation over IPv6
- No need for IPv4↔IPv6 NAT

IVI

- **Most flexible proposal**
- **Enable systems that have no IPv6 address to access servers and peers with IPv4-derived IPv6 addresses (IVI addresses) in the IPv6 domain.**
 - Stateless NAT
- **Enable systems that have no IPv4 address to access servers and peers in the IPv4 domain.**
 - For systems with IPv4- derived IPv6 addresses (IVI addresses), this is solved by a SIIT extension, given the appropriate AAAA record by IVI DNS ALG. This operation is also stateless.
 - Other systems with non-IVI IPv6 addresses require some form of stateful translation.

Observations on Co-existence solutions

- **Prefer native mode communications, when available**
 - IPv4 ↔ IPv4 over IPv4 network
- **When native mode communication is not possible, resort to tunneling**
 - IPv4 ↔ IPv4 through IPv4 over IPv6 tunnel
- **When neither of the above are possible, restore to stateless NAT**
 - Over tunnels, if necessary
- **When all else fails, resort to stateful NAT**
 - Over tunnels if necessary

Stateless versus Stateful NAT

■ Stateless

- More robust and better scaling
- No extra configuration required to support incoming connection requests (P2P applications)
- More restrictive addressing scheme required

■ Stateful

- Less robust and does not scale as well
- Extra port forwarding configuration required to support incoming calls (P2P applications)

IPv4↔IPv6 NAT Deployment

- **At Customer Premise**
 - Scales well
 - Does not introduce a single point of failure
 - Does little to conserve IPv4 address space
- **In Network Core**
 - Does not scale well
 - Introduces single point of failure
 - Conserves IPv4 address space
- **At Provider Aggregation Point**
 - A likely compromise between the two extremes

Conclusion

- **There is much uncertainty in the community**
 - Will IPv6 finally enjoy universal deployment?
 - If so, which co-existence strategy will facilitate deployment?
- **Monitor the following IETF groups for details**
 - Behave
 - INT Area
 - V6OPS

Major Junos IPv6 features

Basic IPv6 Support

- **Link/Global Address Configuration**
 - Dual Stack
- **ICMPv6**
- **Stateless Address Autoconfiguration**
- **Neighbor Discovery Protocol**
- **VRRP for IPv6**
- **Secure Neighbor Discovery (9.3)**

Operations

- **SNMP Agents and Data over IPv6**
- **IPv6 MIBs**
- **IPv6 Packet Filtering**
 - Source/Destination Address
 - ICMP Values
 - Next Header field
 - Packet Length
 - Traffic Class field
 - TCP Flags
- **cFlowdV9 for IPv6 (9.4)**
- **Support for IPv6 NTP server (tbd)**
- **IPv6 Ping**
- **SSH**
- **Telnet**
- **Syslog**

Packet Forwarding

- **Hardware Forwarding**
 - M and T
- **DSCP-IPv6 Classification**
 - BA Classifiers
- **Filter-Based Classification**
 - MF Classifiers DSCP-IPv6 Rewrite Markers
- **Two-Rate Tri-Color Marking**
- **Unicast RPF**
- **IPv6 Aggregate Policing**
- **Filter-Based Forwarding**
- **Class-Based Forwarding**
- **Destination/Source Class Usage**
- **Route Class Policy**
- **Forwarding Table Filtering**
- **filtering tcp/udp in the presence of ipv6 extension headers (9.4)**

Routing Protocols

- **Static Routes**
 - BFD
- **RIPng**
 - Graceful Restart
- **OSPFv3**
 - IPv6 EH authentication
 - IPSec encryption (ESP Header)
 - Overloading
 - Graceful Restart
 - BFD (9.3)
- **IS-IS**
 - Authentication
 - Unicast Mesh Groups
 - Multicast Mesh Groups
 - Graceful Restart
 - BFD
- **Multitopology IS-IS**
 - Unicast
 - Multicast
- **MBGP**
 - Authentication
 - BGP peering to IPv6 endpoints
 - IPv6 routes over IPv4 peering
 - IPv6 Prefix Limits
 - Interface counters
- **Routing Policy**
 - IPv6 multicast scoping
 - IPv6 address family
 - IPv6 prefixes
 - IPv6 route destination address

MPLS

- IPv6 Tunneling over MPLS LSPs
- RSVP-TE for IPv6 (not scheduled yet)
- LDP for IPv6 (not scheduled yet)
- Little demand for pure IPv6/MPLS Cores
 - IPv4 MPLS infrastructure in place
 - Supporting multiple services
 - Public Internet (IPv4 and IPv6)
 - L3 VPN (IPv4 and IPv6)
 - L2 VPNs and VPLS
 - Multicast & Multicast VPNs

MPLS VPNs

- **Layer 3 (RFC 2547) BGP/MPLS VPNs**
 - Unicast
 - Multicast (9.3)
- **VRF Table Label for IPv6 VPNs**

IPv6 Multicast - Native

- **Multicast Address Support**
- **PIMv2**
- **PIM Anycast RP**
- **Statically Defined RP**
- **Embedded RP Addresses**
- **Source-Specific Multicast (SSM)**
- **Multicast Listener Discovery (v1 and v2)**
- **Bootstrap Router (BSR) for IPv6**

Transition

- Dual Stacking
- IPv6-in-IPv4 Configured Tunnels
- IPv6 over MPLS LSPs

Security & Services

■ Security

- Unicast RPF
- Discard Interface
- IPSec for IPv6
 - Requires MS PIC

■ Services

- IPv6 support in VoIP BGF (9.2)
 - IPv6 <-> IPv6 VoIP BGF
 - IPv4 <-> IPv6 VoIP BGF
 - Requires MS PIC

Juniper *your* **Net™**