

v6ops  
Internet-Draft  
Intended status: Best Current Practice  
Expires: January 1, 2018

J. Brzozowski  
Comcast Cable  
D. Schinazi  
S. Cheshire  
Apple Inc.  
L. Colitti  
E. Kline  
J. Linkova  
Google  
M. Keane  
Microsoft  
P. Saab  
Facebook  
June 30, 2017

**Incremental Deployment of IPv6-only Wi-Fi for IETF Meetings**  
**draft-jjmb-v6ops-ietf-ipv6-only-incremental-00**

Abstract

The purpose of this document is to provide a blueprint and guidance for deploying IPv6-only Wi-Fi at IETF meetings. This document outlines infrastructure and operational guidance that operators should consider when deploying IPv6-only networks using NAT64 and DNS64 to support communication to legacy IPv4-only services.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on January 1, 2018.

## Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

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

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.1.</a>	Requirements Language . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Design Principles . . . . .	<a href="#">4</a>
<a href="#">2.1.</a>	Network Infrastructure . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Network Services . . . . .	<a href="#">5</a>
<a href="#">3.1.</a>	DNS64 . . . . .	<a href="#">5</a>
<a href="#">3.2.</a>	NAT64 . . . . .	<a href="#">6</a>
<a href="#">3.3.</a>	DHCPv6 . . . . .	<a href="#">6</a>
<a href="#">4.</a>	User Equipment . . . . .	<a href="#">7</a>
<a href="#">4.1.</a>	Host Address Assignment and Configuration . . . . .	<a href="#">7</a>
<a href="#">4.2.</a>	IPv4 support . . . . .	<a href="#">7</a>
<a href="#">5.</a>	Network Management . . . . .	<a href="#">8</a>
<a href="#">6.</a>	Telemetry and Monitoring . . . . .	<a href="#">9</a>
<a href="#">7.</a>	Support for User Applications and Services . . . . .	<a href="#">10</a>
<a href="#">8.</a>	Support and Operations . . . . .	<a href="#">10</a>
<a href="#">8.1.</a>	Reporting Issues (Ticketing) . . . . .	<a href="#">10</a>
<a href="#">8.2.</a>	Interactive Support . . . . .	<a href="#">11</a>
<a href="#">9.</a>	Known Client-side Issues . . . . .	<a href="#">11</a>
<a href="#">10.</a>	IANA Considerations . . . . .	<a href="#">11</a>
<a href="#">10.1.</a>	Security Considerations . . . . .	<a href="#">12</a>
<a href="#">11.</a>	Future Work . . . . .	<a href="#">12</a>
<a href="#">12.</a>	Related Industry Efforts . . . . .	<a href="#">12</a>
<a href="#">13.</a>	References . . . . .	<a href="#">13</a>
<a href="#">13.1.</a>	Normative References . . . . .	<a href="#">13</a>
<a href="#">13.2.</a>	Informative References . . . . .	<a href="#">14</a>
	Authors' Addresses . . . . .	<a href="#">14</a>



## **1. Introduction**

The purpose of this document is to provide a blueprint and guidance for deploying IPv6-only Wi-Fi at IETF meetings. This document outlines infrastructure and operational guidance that operators should consider when deploying IPv6-only networks using NAT64 and DNS64 to support communication to legacy IPv4-only services.

One of the main strengths of the IETF has always been an insistence on running code. As such, IETF meetings were one of the first deployments of a dual-stack network to help test the first implementations of IPv6. Many years later, as several networks are shifting towards IPv6-only, it is the responsibility of the IETF to lead the trend and make their main network IPv6-only.

This document outlines the requirements and design principles for an IPv6-only network infrastructure that includes support for IPv4-only content. It also discusses techniques and requirements for network management, telemetry, and the operations and support for the IPv6-only network. Recommendations and best practices for operations and support will be provided, however, alternate approaches may be utilized. Disabling or removal of IPv4 stacks is out of scope for this document. This document focuses on the explicit provisioning of IPv6-only using NAT64 [[RFC6146](#)] and DNS64 [[RFC6147](#)] to access IPv4-only content and services.

### **1.1. Requirements Language**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in "Key words for use in RFCs to Indicate Requirement Levels" [[RFC2119](#)].



## **2. Design Principles**

### **2.1. Network Infrastructure**

The following are specific network design details that are minimally required to support an IPv6-only network that utilize NAT64 and DNS64. The following have been drawn from real deployment scenarios for large scale uses of IPv6-only with NAT64 and DNS64. The parameters specified here are specific to providing IPv6-only connectivity. It is assumed that IPv6-only is provisioned and that IPv4 stacks remain active on network and host interfaces. The disabling or removal of IPv4 stacks from hosts or routers is out of scope for this document. As such, it is important to note that link local IPv4 [[RFC3927](#)] will likely remain active and will appear on hosts and network infrastructure.

The following section outlines the requirement to provisioning IPv6-only. We minimally assume that SLAAC will be utilized, however, for completeness the parameters required for DHCPv6 [[RFC3315](#)] and [[RFC3736](#)] are also provided:

- o IPv6-only hosts are expected to be provisioned with IPv6-only connectivity, however, link local IPv4 is likely to be present.
- o RA interval is RECOMMENDED to be minimally set to 600 seconds per the guidance outlined in [[RFC7772](#)].
- o Support for solicited unicast router advertisements are also recommended per [[RFC7772](#)]
- o At least one prefix information option (PIO) MUST be included in router advertisements, the transmitted PIO MUST correspond to the IPv6 prefix that is valid for a given IPv6 link.
- o The use of SLAAC [[RFC4862](#)] MUST be signalled by the network, specifically for each transmitted PIO the A bit MUST be set to one.
- o DHCPv6 support SHOULD be included to support legacy operating systems that do not support DNS RA options but is not required. Whether stateless or stateful DHCPv6 is used, both the DNS Server IPv6 address and DNS Search List options [[RFC8106](#)] MUST minimally be included. The DNS server IPv6 address(es) MUST be those used for DNS64. It is RECOMMENDED that these values be identical to those used in the IPv6 router advertisements that include the DNS options [[RFC8106](#)]. If DHCPv6 support is deployed, stateless DHCPv6 MUST minimally be available.



- o IPv6 router advertisements MUST include the DNS options [[RFC8106](#)]. Both the DNS Server IPv6 address(es) and DNS Search List are REQUIRED. If DHCPv6 support is deployed the values sent here for DNS RA options are RECOMMENDED to match those sent via DHCPv6.

To ensure seamless and to support an incremental deployment of IPv6-only access to legacy dual stack infrastructure should remain available. The following are recommended approaches that may be considered to achieve the same.

The deployment of IPv6-only with NAT64 and DNS64 may very well help to identify applications, services, or use cases that are not entirely compatible with the same. It is therefore important to ensure that users of IP networks, whether wired or wireless, have access to legacy dual stack infrastructure as a fallback. For wireless network it is recommend to have a secondary SSID labelled accordingly, e.g. example-ssid-dual-stack or example-ssid-legacy. For wired network connectivity having secondary ports that are dual stack enabled is also recommended. Note that while it is recommended to ensure the presence of a fallback network, the goal remains to make the IPv6-only network the primary network.

This document assumes that dual stack connectivity is available by default and that IPv4-only connectivity is no longer supported. As such, it is out of scope for this document to outline fallback or access to legacy connectivity that is IPv4-only.

### **[3.](#) Network Services**

The following network services are required for an IPv6-only where support for and access to IPv4 content, services, and applications are required.

#### **[3.1.](#) DNS64**

The following recommendations apply to the use and deployment of DNS64:

- o Use of the well known DNS64 prefix per [[RFC6052](#)]
- o It is also recommended that query logging be enabled for DNS64, performance impacts of query logging must be noted but are largely out of scope for this document. Query logging is essential to determine the volume and make up of DNS queries and replies that are are specific to DNS64 and IPv4-only content, services, and applications.





### **3.2. NAT64**

The following recommendations apply to the use and deployment of NAT64:

- o DNS64 is a critical aspect to direct requests from IPv6-only hosts to a NAT64 service.
- o NAT64 configurations vary widely, port allocation techniques are largely out of scope for this document. One-to-one (1:1) mappings can be used to allocate an IPv4 address per connected device or alternatively blocks of IPv4 ports can also be assigned per device, each has different properties. It is generally recommended to allocate IPv4 ports per device in an effort to maximize IPv4 utilization for NAT64.

### **3.3. DHCPv6**

Support for DHCPv6 may be required in some deployments. If required, parameters pertaining to IPv6 router discovery may require adjustment. The following outlines the guidance specific to the use of DHCPv6:

- o Stateless DHCPv6 SHOULD be supported to facilitate the transmission of DNS servers IPv6 address(es) and DNS search lists to legacy hosts that do not support DNS RA options.
- o Stateful DHCPv6 for address assignment MAY be supported, but is not required. If stateful DHCPv6 is used the DNS parameters mentioned above MUST be included.
- o If, at some future date, support for IPv6 prefix delegation becomes necessary, stateful DHCPv6 will likely be mandatory (Future Work ([Section 11](#))). The details of IPv6 prefix delegation are out of scope for this document.



## **4. User Equipment**

### **4.1. Host Address Assignment and Configuration**

- o Hosts MUST support SLAAC.
- o Hosts SHOULD support DNS RA options [[RFC8106](#)] for the acquisition of DNS server IPv6 addresses and a DNS Search List.
- o Hosts MAY support DHCPv6 for address acquisition, the use of DHCPv6 for address acquisition is not prohibited.
- o DHCPv6 option to configure DNS server option 23 and domain search list option 24 [[RFC3646](#)] address MUST be implemented if DHCPv6 is to be utilized.

### **4.2. IPv4 support**

The IPv4 stacks of hosts MAY remain enabled, which means that Link Local IPv4 [[RFC3927](#)] (169.254/16) addresses MAY continue to be present and in use. Disabling of the IPv4 stack of hosts is out of scope for this document.

Host operating systems SHOULD provide a means for applications to easily connect to IPv4-only servers by using the NAT64/DNS64. While modern applications simply need to make AAAA queries and connect to the resulting IPv6 address, operating systems SHOULD provide simple ways for applications to do so or even connect to IPv4 literals in the absence of host names. Possible solutions include 464XLAT [[RFC6877](#)], "Bump-in-the-Host" [[RFC6535](#)] and Happy Eyeballs v2 [[HEv2](#)].

Finally, it is RECOMMENDED that support for DHCPv4 be explicitly suppressed in particular to prevent the inadvertent assignment of IPv4 addresses on networks that do not have a valid IPv4 egress. DHCPv4 servers, rogue or otherwise, could adversely impact the experience of end users of the IPv6-only network.



## 5. Network Management

The focus of this document is user equipment and hosts. The network and network service requirements are oriented around providing IPv6-only connectivity that allows for the use of NAT64 and DNS64 to maintain reachability to IPv4-only content, applications, and services. Operations and management of the underlying network is technically out of scope for this document, however, given the relevance of the same to the focus of this draft some guidance is being provided.

Strictly speaking the primary requirement for the underlying network is that IPv6 is supported along with the services required to enable the use of NAT64 and DNS64. This suggests that the underlying network could in fact be dual stack for management and operations. It is required that the provisioning of IPv4 for user equipment and host connectivity not be supported. User equipment or host facing interfaces **MUST NOT** acquire non-link-local IPv4 addresses or IPv4 DNS server addresses. Additionally, the network **MUST NOT** respond to DHCPv4 requests or DNS queries sent over IPv4.

Given the above, within a given VLAN it is possible and likely that IPv4 may be observed, present, and possibly used. It is out of scope for this document to prevent the use of IPv4 entirely.

Depending on the level of readiness IPv6-only network management may or may not be possible. Network management and operations includes but is not limited to the following:

- o Remote access to network infrastructure via SSH or telnet
- o Remote SNMP communications
- o Remote NETCONF communications
- o Remote Syslog communications

While it is strongly recommended that all network management and operations be performed over IPv6-only it is not strictly required. However, it is important to note that the presence and use of IPv4 for network management and operations must not impede or impact the use of IPv6-only with NAT64 and DNS64.



## 6. Telemetry and Monitoring

At this point in time, IPv6-only networks with no IPv4 support at all are still not widespread and may expose issues in host operating systems or applications. It is therefore recommended that telemetry summarizing how hosts are being provisioned and accessing the Internet be collected and analyzed. In order to preserve the privacy of users of the network, it is paramount that connectivity information (e.g. DNS64 records) cannot be correlated with individual client nodes.

We can measure how hosts:

- o Configure IPv6 addresses (SLAAC, DHCPv6) and which ones they use
- o Configure DNS server addresses (DNS RA options vs DHCPv6)

We can measure what percentage of the traffic:

- o Uses native IPv6
- o Uses NAT64

Recording the most common hostnames that require the DNS64 would also allow operators to establish a list of the most prominent IPv4-only services.

Observing the TCP/UDP ports used by applications that still leverage IPv4 link-local on an IPv6-only network will also help prepare for the time when routers stop supporting IPv4 communications altogether.

Given that some users may have devices running legacy IPv4-only software, the network should provide a different fallback network that is dual-stack. It is worth measuring the number of users that switch to this network, and possibly use an anonymous survey asking users what software failure caused them to switch. Additionally, the fallback network SHOULD use different authentication credentials per meeting (such as SSID) to make sure a failure causing a user to switch does not mean they will stay on the fallback network forever.





## **7. Support for User Applications and Services**

Following is a list of commonly used applications and services that are expected to operate, without incident, when used in an IPv6-only environment that utilizes NAT64 and DNS64. The list below is not exhaustive.

- o VPN
- o Chat
- o Email
- o SSH/Telnet
- o Git
- o Voice

## **8. Support and Operations**

Most every network has customers or end users of some sort, therefore it essential to ensure that end users or consumers of the have means to do the following while transitions are occurring in networks and related infrastructure. One key item referenced earlier is the availability of temporary fallback networks that support legacy communications.

The following outline additional items that end users must have available to communicate with network operators. All of the items below must be available via dual stack connectivity.

### **8.1. Reporting Issues (Ticketing)**

Tools and systems that can be used to report issues with applications, services, or content must be available for end-users. Network and systems operators are responsible for acknowledging and classifying issues and ultimately ensuring that the same are properly addressed. Specifically to this document "fixed" is meant to imply that proper support for IPv6 is available. In some cases network and system operators may need to implement temporary workarounds to ensure that end users can access the desired content, application, or service.



In order for users experiencing IPv6-specific issues to be able to report them, the ticketing system **MUST** also be reachable over the dual-stack fallback network. The existence of the fallback network **SHOULD** also be made clear to users ahead of time. In order to help narrow down issues, the ticketing system **SHOULD** ask the user whether the issue is specific to IPv6-only and whether they have experienced the issue or a different outcome on the fallback network.

## **8.2. Interactive Support**

Interactive support is often desired in lieu or in conjunction with traditional support models like trouble ticket creation. It is recommended that interactive support be available via real time and near real time mechanisms like Slack or electronic mail (e-mail).

## **9. Known Client-side Issues**

Following are known client side issues that are specific to the deployment of IPv6-only networks and/or the use of NAT64/DNS64:

- o Use of literal IPv4 addresses - the use of literal IPv4 addresses is a known issue given the approach that is documented in this I-D. Addressing the use of literal IPv4 addresses is out of scope for this document.
- o Applications that explicitly require IPv4 by only performing AAAA queries or restricting the type of underlying socket they use.
- o Unreachable but valid AAAA RR in the DNS - in some cases a valid AAAA RR is returned by the DNS, however, if the same is unreachable or is not configured the presence of the same will prevent a DNS64 query which in turn prevents the use of the NAT64 to reach the target host references by the address in the AAAA DNS RR.

## **10. IANA Considerations**

This memo includes no request to IANA.



### **10.1. Security Considerations**

The vastness of the IPv6 address space often makes it more difficult to scan the same unlike legacy IPv4-only or dual stack IP networks. It is conceivable that IPv6-only network represent a reduction in attack surface area which in turn could be viewed a security improvement compared to IPv4-only or dual stack IP networks.

Given the criticality of the DNS64 for reachability to the NAT64, poisoning of one or both could represent a vector for the attack of the DNS64 and NAT64 which could in turn impact the end user experience. Worse poisoning of the DNS64 and/or NAT64 could result in redirection of end use devices to malicious hosts. It is likely that this vulnerability is no greater in IPv6-only networks utilizing DNS64 and NAT64 compared to traditional IPv4-only or dual stack networks.

### **11. Future Work**

The following items are out of scope for this document, however, the following are listed as future work items specific to incremental IPv6-only deployments:

- o Support for IPv6 prefix delegation
- o Disabling IPv4 stacks at some point in the future
- o Fully deprecating the fallback legacy IPv4 network

### **12. Related Industry Efforts**

- o Comcast new building and IPv6-only (John Jason Brzozowski <john\_brzozowski@comcast.com>)
- o Microsoft corporate IT IPv6-only (Marcus Keane <marcus.keane@microsoft.com>)
- o Google (Jen Linkova <furry@google.com>)



## **13. References**

### **13.1. Normative References**

- [HEV2] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2", Work in Progress, [draft-ietf-v6ops-rfc6555bis](#), June 2017.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3646](#), DOI 10.17487/RFC3646, December 2003, <<http://www.rfc-editor.org/info/rfc3646>>.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", [RFC 3736](#), DOI 10.17487/RFC3736, April 2004, <<http://www.rfc-editor.org/info/rfc3736>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), DOI 10.17487/RFC6052, October 2010, <<http://www.rfc-editor.org/info/rfc6052>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), DOI 10.17487/RFC6146, April 2011, <<http://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [RFC 6147](#), DOI 10.17487/RFC6147, April 2011, <<http://www.rfc-editor.org/info/rfc6147>>.





- [RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", [BCP 202](#), [RFC 7772](#), DOI 10.17487/RFC7772, February 2016, <<http://www.rfc-editor.org/info/rfc7772>>.
- [RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", [RFC 8106](#), DOI 10.17487/RFC8106, March 2017, <<http://www.rfc-editor.org/info/rfc8106>>.

### **13.2. Informative References**

- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", [RFC 3927](#), DOI 10.17487/RFC3927, May 2005, <<http://www.rfc-editor.org/info/rfc3927>>.
- [RFC6535] Huang, B., Deng, H., and T. Savolainen, "Dual-Stack Hosts Using "Bump-in-the-Host" (BIH)", [RFC 6535](#), DOI 10.17487/RFC6535, February 2012, <<http://www.rfc-editor.org/info/rfc6535>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", [RFC 6877](#), DOI 10.17487/RFC6877, April 2013, <<http://www.rfc-editor.org/info/rfc6877>>.

#### Authors' Addresses

John Jason Brzozowski  
Comcast Cable  
1701 John F. Kennedy Blvd.  
Philadelphia, PA  
USA

Email: [john\\_brzozowski@cable.comcast.com](mailto:john_brzozowski@cable.comcast.com)

David Schinazi  
Apple Inc.  
1 Infinite Loop  
Cupertino, California 95014  
US

Email: [dschinazi@apple.com](mailto:dschinazi@apple.com)



Stuart Cheshire  
Apple Inc.  
1 Infinite Loop  
Cupertino, California 95014  
USA

Email: [cheshire@apple.com](mailto:cheshire@apple.com)

Lorenzo Colitti  
Google

Email: [lorenzo@google.com](mailto:lorenzo@google.com)

Erik Kline  
Google

Email: [ek@google.com](mailto:ek@google.com)

Jen Linkova  
Google

Email: [furry@google.com](mailto:furry@google.com)

Marcus Keane  
Microsoft

Email: [marcus.keane@microsoft.com](mailto:marcus.keane@microsoft.com)

Paul Saab  
Facebook

Email: [ps@fb.com](mailto:ps@fb.com)

