IPv6 Operations Working Group (v6ops)   F. Gont
Internet-Draft   G. Gont
Intended status: Informational   SI6 Networks
Expires: August 25, 2021   February 21, 2021

IPv6 Addressing Considerations
draft-gont-v6ops-ipv6-addressing-considerations-01

Abstract

IPv6 addresses can differ in a number of properties, such as scope, stability, and intended usage type. This document analyzes the impact of these properties on aspects such as security, privacy, interoperability, and network operations. Additionally, it identifies challenges and gaps that currently prevent systems and applications from leveraging the increased flexibility and availability of IPv6 addresses.

Table of Contents

1. Introduction

IPv6 addresses can differ in a number of properties, such as address scope (e.g. link-local vs. global), stability (e.g. stable addresses vs. temporary addresses), and intended usage type (outgoing communications vs. incoming communications). While often overlooked, these properties have direct impact on areas such as security, privacy, interoperability, and network operations.

IPv6 hosts typically configure addresses based on local system policy, which tends to be static and irrespective of the specific network the host attaches to. For example, most IPv6 host implementations configure one link-local address for each network interface, and one stable and one (or more) temporary addresses per each Stateless Address Auto-configuration (SLAAC) [RFC4862] prefix for each network interface. However, this static policy for address configuration might be inappropriate. For example, mobile nodes might benefit from employing only temporary addresses, which generally offer better privacy properties than stable addresses. On the other hand, an enterprise network might prefer that local hosts employ only stable addresses, which might be more convenient when enforcing access control, performing network trouble-shooting, or identifying hosts that might have been infected by malware.

On the other hand, Each application on a given host could have its own set of requirements or expectations for the underlying IPv6 addresses. For example, an application meaning to offer a public service might expect to employ addresses that are both globally-reachable [RFC8190] and stable [RFC7721] [RFC8064], while a privacy-sensible client application might prefer short-lived temporary addresses [I-D.ietf-6man-rfc4941bis], or might even expect to employ single-use ("ephemeral") IPv6 addresses when connecting to public servers. However, the subtleties associated with IPv6 addresses are often ignored or overlooked by application programmers. This means that applications could fail to signal their requirements and preferences to the underlying host, or that the addresses configured by the underlying host might be inappropriate to satisfy the requirements of the corresponding applications.

Finally, a number of limitations in components that range from network devices to Application Programming Interfaces (APIs) could also prevent hosts and applications from leveraging the increased flexibility of IPv6 addressing.

This document identifies a set of properties that can be associated with IPv6 addresses (such as scope and stability), and analyzes the impact of these properties on areas ranging from security and privacy to network operations, with the goal of providing guidance about IPv6 address usage. Additionally, it identifies challenges and gaps that currently prevent systems and applications from leveraging the increased flexibility and availability of IPv6 addresses.

2. Terminology

This document employs the definitions of "public address", "stable address", and "temporary address" from Section 2 of [RFC7721].

This document employs the definition of "globally reachable" from Section 2.1 of [RFC8190].

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Conventions

3.1. Legacy Specifications and Schemes

IPv6 SLAAC has traditionally employed schemes for generating Interface Identifiers (IIDs) that have negatively affected the security and privacy properties of IPv6 addresses. For example, IPv6 SLAAC originally generated stable addresses by embedding the underlying link-layer address in the IPv6 Interface Identifier (IID), thus negatively affecting the security and privacy properties of IPv6 addresses [RFC7721] [RFC7707]. Similarly, IPv6 temporary addresses [RFC4941] reused the same randomized IID for different auto-configuration prefixes [RFC4941], thus allowing for network activity correlation across different addresses of the same host.

These schemes have become formally superseded by other schemes, such as [RFC7217] and [I-D.ietf-6man-rfc4941bis], that mitigate the aforementioned issues. Therefore, this document does not discuss issues arising from legacy IID generation algorithms.

NOTE:
The security and privacy implications of such schemes are discussed in [RFC7721], [RFC7707], and [RFC7217].

3.2. Address Scope

[RFC4007] defines the scope of an address as:

> "[the] topological span within which the address may be used as a unique identifier for an interface or set of interfaces"

And defines the "global scope" to be used for:

> "uniquely identifying interfaces anywhere in the Internet"

However, the term "scope" is employed in conflicting ways in different specifications (see [I-D.gont-6man-ipv6-ula-scope]). Throughout this document, we employ the notion of "scope" defined in [RFC4007]. As a result, addresses that do not uniquely identify interfaces Internet-wide are considered to have "non-global" or "limited" scope. Grouping addresses in such a way is simply useful for the purpose of discussing address properties.

4. IPv6 Address Properties

There are, at least, four properties that can be associated with every IPv6 address:

- Scope

- Reachability

- Stability

- Provider Dependency

The address scope essentially represents the topological span where an address can be expected to uniquely identify an interface; i.e., the topological span where an given address is meaningful. For example, link-local addresses are only meaningful within a given network link, and are expected to be unique only within such network link.

Address reachability represents the topological span where an address can be expected to be used for receiving and transmitting packets. Reachability is implicitly constrained by the address scope, and may also be affected by network devices: for example, Customer Edge Routers (CE Routers) that enforce a filtering policy of "only allowing outgoing communications" can render otherwise globally reachable addresses as "unreachable from the public Internet, unless communication is initiated from the customer’s network".

The stability of an address is associated with the invariance of an address over time. For example, a manually-configured address will typically remain stable while the node remains attached to the same subnet, while a temporary address will, by definition, change over time. While address stability does depend on the inherent properties of a given address (e.g. stable vs. temporary), it also depends on other factors, such as provider dependency: if a network employs a prefix that is assigned/leased by an upstream provider, then the overall stability an address will also depend on the stability corresponding network prefix.

Provider-dependency is typically discussed in the context of Global Unicast Addresses, where the address space may be allocated by an Internet Service Provider (ISP) (and hence "provider aggregatable") or by a Regional Internet Registry (RIR) (and hence "provider independent"). However, this document considers "provider dependency" in a more general way: "provider aggregatable" address space is assigned or leased by an upstream provider and carved out from the provider's address space, and thus is topologically-related to the upstream provider's address space; on the other hand, "provider independent" address space is "owned" by the network in question and thus is not necessarily topologically-related to the upstream provider.

4.1. Address Scope Considerations

The IPv6 address scope [RFC4007] has a direct implication on address reachability: the address scope essentially constrains address reachability. For example, addresses that have a non-global/limited scope are not, in principle, globally reachable.

NOTE:
This assumption becomes invalid if technologies such as Network Prefix Translation (NPT) [RFC6296] are employed, though. However, strictly speaking, in these scenarios the non-global addresses are still not globally reachable, but rather the middle-box acts as an interface with the "external realm" via globally-reachable addresses (i.e., the middle-box provides an interface between two topological spans).

The IPv6 address scope can, in some scenarios, limit the attack exposure of a node as a result of the implicit isolation provided by a non-global/limited address scopes. For example, a node that only employs link-local addresses will, in principle, only be exposed to attacks from other nodes on the same local link.

The potential protection provided by a non-global-scope addresses should not be regarded as a complete security strategy, but rather as a form of "prophylactic" security (see [I-D.gont-opsawg-firewalls-analysis]).

We note that non-global scope addresses are normally only of use for a limited number of applications/protocols that operate on a limited scope (e.g., mDNS), or deployments where the intended participants are known to operate in a limited domain [RFC8799] (e.g., OpenSSH client and server attached to the same link and employing link-local addresses, or mDNS hosts employing link-local addresses).

The address scope can at times be somewhat related with the provider dependency property. For example, link-local addresses are, by definition, provider independent. In the same light, a locally-generated ULA prefix will be, by definition, provider independent. However, a router might also employ a ULA prefix leased by an upstream router, in which case this prefix would be "provider dependent". The possible implications of the address scope on "provider dependency" may also affect address stability: for example, a locally-generated ULA prefix is "provider independent", and will not be subject to renumbering events triggered by the upstream provider. However, a router (e.g. CE Router) might, in some circumstances, be unable to guarantee prefix stability -- as in the case where the locally-generated ULA prefix is not recorded on stable storage, and thus cannot be guaranteed to remain stable across power outages.

4.2. Provider Dependency

Provider-dependency is typically discussed in the context of Global Unicast Addresses, where the address space may be allocated by an Internet Service Provider (ISP) (and hence "provider agreggatable") or by a Regional Internet Registry (RIR) (and hence "provider independent"). However, this document considers "provider dependency" in a more general way: "provider aggregatable" address space is assigned or leased by an upstream provider and carved out from the provider's address space, and thus is topologically-related to the upstream provider's address space; on the other hand, "provider independent" address space is "owned" by the network in question and thus is not necessarily topologically-related to the upstream provider.

An implicit consequence of PA address space is that its use is tied to the specific provider/upstream provider that provides the address space. This has a number of consequences, including:

- Multi-homing (employing local address space with multiple upstream providers) is not possible.

- A renumbering event at the upstream provider will typically cause the local network to be renumbered.

Since PA space has a topological relationship with the upstream provider, it will prevent multi-homing. This has led some organizations to employ NPT [RFC6296] such that:

- The local network is isolated of renumbering events caused by the upstream provider.

- The local network employs the same address space regardless of the upstream provider employed to communicate with the external realm.

While PA space may impact address stability, PI address space generally has better stability properties. For example, a home network could internally employ both ULAs and GUAs, where a ULA prefix is locally generated by the CE Router (and hence resulting in PI space), and a global prefix is leased by the ISP via DHCPv6 Prefix Delegation [RFC8415] (hence PA space). If for some reason there was an outage involving the connection with the upstream ISP, the lease time for the GUA prefix would eventually expire, and therefore addresses configured for such prefix would need to be invalidated. Similarly, if upon prefix lease expiration the ISP were to lease a new GUA prefix (rather than renew the current prefix), the network would need to be renumbered. On the other hand, locally-generated ULA prefixes can be employed independently from the upstream ISP.

Similarly, an organizational network that employs PI global address space obtained from a RIR would be able to employ the same address space irrespective of renumbering events or outages involving the upstream provider.

4.3. Address Reachability

Address reachability represents the area of the network (and the associated conditions), where an address can be used for receiving and transmitting packets. As noted in Section 4.1, the address scope has a direct implication on address reachability, since it constrains the network span where the address is reachable.

In addition to the reachability semantics of each address type, network filtering policies may also affect address reachability. For example, there is widespread deployment of Customer Edge Routers that implement a (stateful) filtering policy of "only allowing outgoing communications" -- mimicking the filtering policy enforced (as a side-effect) by IPv4 NATs. In such scenarios, even otherwise globally-reachable addresses become unreachable, unless:

- o communication is initiated from the internal network, or,

- o the CE Router is manually configured override the default filtering policy, or,

- o a technology to dynamically override the filtering policy (such as UPnP [UPnP] or PCP [RFC6887]) is employed.

Address reachability is what ultimately determines the application architecture that may be successfully employed by an IPv6 node.

NOTE:
Ironically, an IPv6-only host (with global-scope addresses) attached to a home network where the CE Router "only allows outgoing communications" and does not implement protocols such as UPnP [UPnP] or PCP [RFC6887], will normally have a harder time using peer-to-peer (P2P) applications than an IPv4-only host (with a private address) attached to a home network where the CE Router employs NAT but implements a protocols such as UPnP or PCP.

Address reachability has a direct impact on security, since the ability to attack a system normally relies on the ability of the attacker to reach the system in the first place. Firewalls [I-D.gont-opsawg-firewalls-analysis] are, indeed, devices that are specifically devoted to administer address reachability.

4.4. Address Stability Considerations

Address stability typically depends on two factors:

- o Stability of the network prefix

- o Inherent stability of address type

Depending on whether the local prefix is PI or PA (see Section 4.2) and whether the prefix is stable or dynamic (see [I-D.ietf-v6ops-slaac-renum]), the resulting addresses will have different stability properties. Additionally, even in the presence of stable prefixes, a host may configure stable addresses [RFC8064] and/or temporary addresses [RFC4941].

The stability of an address has two associated security/privacy implications:

- o Ability of an attacker to correlate network activity

- o Exposure to attack

For obvious reasons, an address that is employed for multiple communication instances allows the aforementioned network activities to be correlated. The longer an address is employed (i.e., the more stable it is), the longer such correlation will be possible. In the worst-case scenario, a stable address that is employed for multiple communication instances over time will allow all such activities to be correlated. On the other hand, if a host were to generate (and eventually remove) one new address for each communication instance (e.g., TCP connection), network activity correlation would be mitigated.

NOTE:
The security and privacy implications of predictable addresses are discussed in [RFC7721] and [RFC7707].

Typically, the longer an address is employed the longer the window of exposure of a host as being accessible via an address that becomes revealed as a result of active communication. While such exposure is traditionally associated with the stability of the address, the usage type of the address may also have an impact on attack exposure (see Section 5.2).

A popular approach to mitigate network activity correlation is the use of "temporary addresses" [RFC4941]. Temporary addresses are typically auto-configured and employed along with stable addresses, with the temporary addresses employed for outgoing communications, and the stable addresses employed for incoming communications.

NOTE:
Ongoing work [I-D.ietf-6man-rfc4941bis] aims at updating [RFC4941] such that temporary addresses can be employed without the need to configure stable addresses.

We note that the extent to which temporary addresses provide improved mitigation of network activity correlation and/or reduced attack exposure may be questionable and/or limited in some scenarios. For example, a temporary address that is reachable for, say, a few hours has a questionable "reduced exposure" (particularly when automated attack tools do not typically require such a long period of time to complete their task). Similarly, if network activity can be correlated for the life of such address (e.g., on the order of several hours), such period of time might be long enough for the attacker to correlate all the network activity of interest. However, they temporary addresses do limit the attack window and the amount of time during which address-based network activity correlation can be performed.

In order to better mitigate network activity correlation and/or possibly reduce host exposure, an implementation might want to either reduce the preferred lifetime of temporary addresses or, even better, generate one new IPv6 address for each application or new transport protocol instance (sometimes referred to as "ephemeral addresses"). However, reduced address lifetimes and the use of multiple IPv6 addresses may have a negative impact on the network (please see Section 6.3).

Additionally, enforcing a maximum lifetime on IPv6 addresses may cause long-lived TCP connections to fail. For example, an address becoming "Invalid" (after transitioning through the "Preferred" and "Deprecated" states) would cause the TCP connections employing them to break, which would in turn cause e.g. long-lived SSH sessions to break/fail. Traditionally, many application protocols have assumed or expected address stability. However, in the light of mobile roaming nodes that may frequently switch among different connections (e.g. Wi-Fi, 4G, etc.) or that may be subject to renumbering events (see [I-D.ietf-v6ops-slaac-renum]), robust applications should assume and expect "ephemeral" IPv6 addresses (i.e., gracefully handle the case where the underlying IPv6 addresses change over short periods of time).

In some scenarios, attack exposure may be further mitigated by limiting the usage of temporary addresses to outgoing connections, and prevent such addresses from being used for incoming connections (please see Section 5.2).

Finally, we note that on different single-use (i.e., "ephemeral") IPv6 address is employed for each transport protocol instance, the possibility of an attacker successfully performing off-path attacks (such as the TCP reset attacks discussed in [RFC4953]) is reduced, since the ephemeral IPv6 address will typically be unknown and unpredictable to the off-path attacker.

5. IPv6 Address Usage

5.1. Default IPv6 Address Selection

Applications use system API's to implicitly or explicitly select the IPv6 addresses that will be used for incoming and outgoing connections. These choices have consequences in terms of privacy, security, stability and performance.

Default Address Selection for IPv6 is specified in [RFC6724]. The selection starts with a set of potential destination addresses, such as returned by getaddrinfo(3), and the set of potential source addresses currently configured for the selected interfaces. For each potential destination address, the algorithm will select the source address that provides the best route to the destination, while choosing the appropriate scope and preferring temporary addresses. The algorithm will then select the destination address, while giving a preference to reachable addresses with the smallest scope. The selection may be affected by system settings. We note that [RFC6724] only applies for outgoing connections, such as those made by clients trying to use services offered by other hosts.

We note that [RFC6724] selects IPv6 addresses from all the currently available addresses on the host, and there is currently no way for an application to indicate expected or desirable properties for the IPv6 source addresses employed for such outgoing communications. For example, a privacy-sensitive application might want that each outgoing communication instance employs a new, single-use IPv6 address, or to employ a new reusable address that is not employed or reusable by any other application on the host. Reuse of an IPv6 address by an application would allow the correlation of all network activities corresponding to such application as being performed by the same host, while reuse of an IPv6 address by multiple different applications would allow the correlation of all such network activities as being performed by the host with such IPv6 address (see Section 4.4 for further details).

When a host provides a service, the common pattern is to just wait for incoming connections over all configured addresses. For example, applications using the BSD Sockets API will commonly bind(2) the listening socket to the undefined address. This long-established behavior is appropriate for hosts providing public services, but can have unexpected results for hosts providing semi-private services, such as various forms of peer-to-peer or local-only applications (e.g. mDNS).

This behavior leads to three problems: host tracking, discussed in Section 6.2.2; unexpected address discovery, discussed in Section 6.2.3; and availability outside the expected scope, discussed in Section 6.2.4. These problems are caused in part by the limitations of available address selection API, discussed in Section 7.4.

5.2. Usage Type Considerations

IPv6 hosts may configure stable [RFC8064] and/or temporary [RFC4941] addresses, where stable addresses are typically employed for incoming (server-like) communications, and temporary addresses are employed for outgoing (client-like) communications. That is, the stability properties of an address have an implicitly associated usage type.

A host that employs one of its addresses to communicate with a remote server (i.e., that performs an "outgoing connection") will expose that address to the target server (and to on-path nodes). Once the remote server receives an incoming connection, it could readily launch an attack against the host via the exposed address. A real-world instance of this type of scenario has been documented in [Hein].

However, we note that employing an IPv6 address for outgoing communications need not increase the exposure of local services to other parties. For example, nodes could employ temporary addresses only for outgoing communications, and disallow their use for incoming communications. Thus, nodes that learn about a client's addresses could not really leverage such addresses for actively contacting clients. Unfortunately, current APIs represent a challenge when trying to leverage IPv6 addresses in this way (please see Section 5.2.1 and Section 7.4 for further details).

The following subsections possible techniques that could be employed by applications to better leverage IPv6 addresses for both incoming and outgoing communications

5.2.1. Incoming communications

There are a number of ways in which a system or network may affect which addresses (and how) may be employed for different services and cases. Namely,

- o TCP/IP stack address filtering
- o Application-based address filtering
- o Firewall-based address filtering

Clearly, the most elegant approach for address selection would be for applications to be able to specify the properties of the addresses they are willing to employ by means of an API, such the TCP/IP stack itself could "filter" which addresses are allowed for the given service/application. For example, an application could specify the stability and scope properties of the addresses on which incoming communications should be accepted, such that the application can be relieved from dealing with low-level networking details, portability is improved, and duplicate code in applications is avoided. However, constraints in the current APIs (see Section 7.4) prevent application programmers from leveraging this technique. Alternatively, services could be bound to specific (explicit) addresses, rather than to all locally-configured addresses. However, there are a number of short-comings associated with this approach. Firstly, an application would need to be able to learn all of its addresses and associated properties, something that tends to be non-trivial and non-portable, and that also makes applications protocol-dependent, unnecessarily. Secondly, the BSD Sockets API does not allow a socket to be bound to a subset of the node's addresses. That is, sockets can be bound to a single address or to all available addresses (wildcard), but not to a subset of all the configured addresses.

Another possible approach would be for applications to e.g. bind services to all available addresses, and perform the associated selection/filtering at the application level. While possible, this would have a number of drawbacks. Firstly, it would require applications to deal with low-level networking details, lead to duplicated code in all applications, and also negatively affect portability. Secondly, performing address/selection filtering at the application level may not mitigate some possible attacks. For example, port scanning would still be possible, since the aforementioned filtering would be performed once UDP packets have been received or TCP connections have been established.

A client could simply run a host-based firewall that only allows incoming connections on the stable addresses. This would be more of an operational approach for achieving the desired functionality, and would require good firewall/host integration (e.g., the firewall should be able to tell stable vs. temporary addresses), would require the client to run additional firewall software for this specific purpose, etc. In other scenarios, a network-based firewall could be configured to allow outgoing communications from all internal addresses, but only allow incoming communications to stable addresses (either via manual configuration or via a helper protocol such as [UPnP] or PCP [RFC6887]). For obvious reasons, this is generally only applicable to networks where incoming communications are allowed to a limited number of hosts/servers.

5.2.2. Outgoing communications

An application might be able to obtain the list of currently-configured addresses, and subsequently select an address with desired properties, and explicitly "bind" the address to the socket, to override the default source address selection.

However, this approach is problematic for a number of reasons. Firstly, there is no portable way of obtaining the list of currently-configured addresses on the local node, and even less to check for address properties such "valid lifetime". Secondly, as discussed in Section 5.2.1, it would require application programmers to understand all the subtleties associated with IPv6 addressing, and would also lead to duplicate code on all applications. Finally, applications would be limited to use already-configured addresses and unable to trigger the generation of new addresses where desirable (e.g. the generation of a new single-use address for this application instance or communication instance).

6. Current Issues Associated with IPv6 Addressing

The following subsections discuss current problems associated with IPv6 addresses, namely:

- o Sub-optimal Address Configuration (Section 6.1)
- o Sub-optimal IPv6 Address Usage (Section 6.2)
- o Operational Problems (Section 6.3)

6.1. Sub-optimal Address Configuration

6.1.1. Number of Addresses

Two mechanisms exist for automatic network configuration: SLAAC [RFC4862] and DHCPv6 [RFC8415]. DHCPv6 centralizes network configuration and address assignment, and may thus prevent hosts from leveraging the increased flexibility and availability of IPv6 addresses. On the other hand, SLAAC may result in network configuration anarchy, where hosts may e.g. configure and use addresses in a way that may negatively affect the network (please see Section 6.3.1).

Most of the challenges associated with the use of multiple addresses can be addressed by allocating one /64 per host via mechanisms such as DHCPv6-PD [RFC8415]. However, support for such mechanisms in host implementations and e.g. the LAN-side of CE Routers is rather uncommon. On the other hand, SLAAC lacks the means for conveying information about e.g., the number of addresses per host that the network is able or willing to support.

NOTE:
Use of a /64 prefix per host could also render techniques such as temporary addresses [RFC4941] ineffective, since hosts would become identified by corresponding /64 prefix.

6.1.2. SLAAC/DHCPv6 Interaction

Many CE Routers offer address configuration via both SLAAC and DHCPv6, by including Prefix Information Options (PIOs) with the "A" flag set in Router Advertisement messages, and also setting the "M" flag in such RA messages. This has a number of implications:

- o The outcome of the configuration process is non-deterministic, difficulting network troubleshooting (see [I-D.ietf-v6ops-dhcpv6-slaac-problem]).

- o Nodes end up configuring more addresses than needed (or even used), normally configuring multiple stable addresses for each autoconfiguration prefix, with at least one address for each configuration mechanism (SLAAC and DHCPv6).

- o A host may end up employing predictable addresses resulting from DHCPv6, thus thwarting the security and privacy improvements of SLAAC-configured addresses (i.e., [RFC7217] and [RFC4941]).

6.2. Sub-optimal IPv6 Address Usage

6.2.1. Correlation of Network Activity

As discussed in [RFC7721], a node that reuses an IPv6 address for multiple communication instances will enable the correlation of such network activities. This could be the case when the same IPv6 address is employed by several instances of the same application (e.g., a browser in "privacy" mode and a browser in "normal" mode), or when the same IPv6 address is employed by two different applications on the same node (e.g., a browser in "privacy" mode, and an email client).

Particularly in the case of privacy-sensitive applications, an application or system might want to limit the usage of a given IPv6 address to a single communication instance, a single application, a single user on the system, etc. However, as discussed in Section 5, given current APIs, this is practically impossible.

6.2.2. Host Tracking

The stable addresses recommended in [RFC8064] use stable IIDs defined in [RFC7217]. One key part of that algorithm is that if a device connects to a given network at different times, it will always configure the same IPv6 addresses on that network. If the device hosts a service ready to accept connections on that stable address, adversaries can test the presence of the device on the network by attempting connections to that stable address. Stable addresses will thus enable testing whether a specific device is returning to a particular network, which in a number of cases might be considered a privacy issue.

6.2.3. Unintended Service Disclosure

Systems like DNS-Based Service Discovery [RFC6763] allow clients to discover services within a limited domain (e.g. a local link). These services are not advertised outside of that domain, and thus do not expect to be discovered by random parties on the Internet. However, such services may be easily discoverable if they allow incoming connections on IPv6 addresses that client processes also use when connecting to remote servers.

NOTE:

An example of such service disclosure is described in [Hein]. A network manager observed port scanning traffic directed at the temporary addresses of local host. The analysis in [Hein] shows that the scanners learned the addresses by observing the device contact an NTP service ([RFC5905]). The remote scanning was possible because the local services were accepting connections on all configured addresses, including temporary addresses.

It is obvious from this example that local services are disclosed because they are bond to the same IPv6 addresses that are also used by clients for outgoing communications with remote systems. But the overlap between "client" and "server" addresses is only one part of the problem. Suppose that a host operates both a video game server and a home automation application server. The video game users will be able to discover the IPv6 address of the game server; if the home automation server listens to the same IPv6 addresses, its address will be revealed to all these users, thus increasing the exposure of the home automation server.

We note that a host or network that wants to limit access to local services should filter incoming connection attempts by affecting address reachability (see Section 4.3) via firewalls [I-D.gont-opsawg-firewalls-analysis] and/or the use of IPv6 addresses of appropriate scope (see Section 4.1). However, it is also prudent to avoid unintended service disclosure by avoiding the scenarios discussed in this section.

6.2.4. Availability of Service Outside the Expected Domain

IPv6 defines [RFC4291] [RFC4007] multiple address scopes, with hosts typically configuring Global Unicast Addresses (GUAs), link local addresses, and Unique Local IPv6 Unicast Addresses (ULAs) [RFC4193]. Availability of a service outside the expected scope happens when a service is expected to be available only in some limited domain, but inadvertently becomes available from outside of that domain. This could happen, for example, if a service is meant to be accessible only within a given link, but becomes reachable from outside that link via ULAs or GUAs, or if a service is meant to be accessible only within some organization's perimeter but becomes accessible from the public Internet via GUAs. This will commonly happen if a service intended for a limited domain is implemented by bind()ing the listening socket to the "unspecified" addresses (please see Section 7.4).

6.3. Operational Problems

6.3.1. Implications of Employing Multiple Addresses

Network deployments are currently recommended to provide multiple IPv6 addresses to general-purpose hosts [RFC7934]. However, in some scenarios, use of a large number of IPv6 addresses may have negative implications on network devices that need to maintain entries for each IPv6 address in network data structures (e.g., [RFC7039]). Additionally, concurrent active use of multiple IPv6 addresses will normally increase neighbour discovery traffic if Neighbour Caches in network devices are not large enough to store all addresses on the link. This can impact performance and energy efficiency on networks on which multicast is expensive (e.g. [I-D.ietf-mboned-ieee802-mcast-problems]). Finally, network devices may interpret the use of a number of addresses above a certain threshold as a security event, and block the offending device from using the network.

6.3.2. Legitimate Network Activity Correlation

The desires of protecting individual privacy versus the desire to effectively maintain and debug a network can conflict with each other. For example, having clients use addresses that change over time will make it more difficult to track down and isolate operational problems. When looking at packet traces, it could become more difficult to determine whether one is seeing behavior caused by a single errant machine, or by a number of them.

6.3.3. Routing in Multi-Prefix/Multi-Router Networks

If the network is provided with multiple upstream connections via different providers and different local routers, each of them will typically provide its own PA address space (see Section 4.2) and thus local hosts will typically configure addresses for each of PA address spaces. In this scenario, packets sourced from a given PA space should only employ the local router of the corresponding upstream provider, since otherwise packets might be dropped as a result of ingress/egress filtering [RFC2827]. Unfortunately, traditional Neighbor Discovery [RFC4861] can advertise routes only with a per-destination granularity, irrespective of the source address/prefix.

[RFC8028] addresses the most important challenges associated with these scenarios. However, [RFC8028] is not yet widely implemented. As a result, operating a multi-prefix/multi-router IPv6 network represents a major challenge -- if at all possible.

6.3.4. Renumbering

The challenges posed by network renumbering have been known for a very long time [RFC5887], with network renumbering typically being assumed to be performed in a planned manner.

However, in scenarios where a host is moved to a different network without the host detecting the network re-attachment event, or where the network a host attaches to is moved to a different point of the network topology (i.e., the network itself is migrated/"moved"), the aforementioned host will also experience a renumbering event. In an era in which migrating virtual machines, containers, and networks around a network topology is commonplace, and where mobile systems changing network connectivity to and from e.g. WiFi and 4G is also commonplace, renumbering events are anything but rare.

One of the challenges represented by network renumbering is how hosts can infer that an existing network prefix and associated address(es) have become stale (such that stale prefixes and addresses can be removed and replaced by new prefixes and addresses). In scenarios where the network topology does not change and the network is renumbered, network elements may be aware of the renumbering event and signal this condition to attached systems (i.e., signal that existing network configuration information should be removed and replaced). However, in scenarios where it is the host, virtual machine, container or network that move around the network topology, the network might not be able to signal the "renumbering event", and these events might be harder to infer and react to.

Unfortunately, both SLAAC and DHCPv6 assume that network configuration information is somewhat stable. SLAAC has traditionally employed long lifetimes for network configuration information, meaning that stale information could be employed for an unacceptably long period of time. DCHPv6 operates on the same premise, and lacks widespread support for RECONFIGURE messages -- so even if the network were in a position to signal a renumbering event, hosts will normally rely on expiration of lease times for stale information to be cleared up.

Some of these problems have been discussed in detail in [I-D.ietf-v6ops-slaac-renum], and there is ongoing work [I-D.ietf-6man-slaac-renum] [I-D.ietf-v6ops-cpe-slaac-renum] to mitigate this issue.

7. Current Gaps that Prevent Leveraging IPv6 Addressing

The following subsections identify and discuss areas where further work is needed. Namely,

- Profile-based IPv6 Address Configuration (see Section 7.1)

- Advice on IPv6 Address Usage (see Section 7.2)

- Protocol Improvements to Deal with Many Addresses (see Section 7.3)

- Improved Address Selection APIs (see Section 7.4)

- Universal Support of RFC 8028 (see Section 7.5)

- Support for Firewall Traversal in CE Routers (see Section 7.6)

7.1. Profile-based IPv6 Address Configuration

Most operating systems configure the same type of addresses regardless of the current "operating mode" or "profile" of the device (e.g., device connected to an enterprise network vs. roaming across untrusted networks). For example, many operating systems configure both stable [RFC8064] and temporary [RFC4941] addresses for all network types. However, this "one size fits all" approach tends to be sub-optimal or even inappropriate for some scenarios. For example, enterprise networks typically prefer the use of only stable addresses, thus requiring the network administrator to configure each host to disable the use of temporary addresses. On the other hand, mobile devices typically configure both stable and temporary addresses, even when their operating mode (client-like operation) would allow for the more privacy-sensible option of configuring only temporary addresses.

The lack of fine-grained address configuration policies forces nodes to rely on a "one size fits all" approach that, as noted, usually leads to suboptimal results. Advice in this area might help achieve profile-based address configuration policies such that IPv6 addressing capabilities are fully leveraged.

NOTE:
One might envision a document that provides advice regarding IPv6 address generation for different typical scenarios (e.g., when to configure stable-only, temporary-only, or stable+temporary). In the most simple analysis, one might expect nodes in a typical enterprise network to employ only stable addresses. General-purpose nodes in a home or "trusted" network might want to employ

both stable and temporary addresses. Finally, mobile nodes (e.g. when roaming across non-trusted networks) might want to employ only temporary addresses).

7.2. Advice on IPv6 Address Usage

An application programmers typically rely to the Default Source IPv6 Address Selection for IPv6 (see Section 5.1) for outgoing communications, and to accepting incoming communications on all configured addresses. As discussed throughout this document, this leads to sub-optimal or undesirable results. Applications on a node share the same pool of configured addresses, and currently available APIs prevent applications from requesting the generation of new addresses (e.g. to be employed for a particular application or communication instance).

Guidance in this area is warranted such that applications and systems can fully leverage IPv6 addressing.

NOTE:

Such guidance would elaborate, among other things, both on the usage of IPv6 addresses when offering network services and when performing client-like communications. For example, for incoming communications, hosts might want to employ only the smallest-scope applicable addresses (if available) and, if stable addresses are available only accept incoming connections on such addresses. For client-like communications, hosts might prefer temporary addresses, unless the corresponding communication instances are expected to be long-lived (e.g., SSH sessions).

7.3. Protocol Improvements to Deal with Many Addresses

Possible improvements to IPv6 SLAAC should be evaluated, including:

- o Enabling IPv6 routers to convey information about network constraints such as maximum number of addressees per node.

- o Enabling hosts to register/de-register configured addresses, such that e.g. routers need not tie resources to addresses that are no longer used.

On the other hand, in order for DHCPv6-PD (or some alternative protocol) to be employed to support the "one /64 per node" paradigm, widespread support for DHCPv6-PD (or an alternative protocol) would be necessary.

7.4. Improved Address Selection APIs

Application developers using the BSD Sockets API can "bind()" a listening socket to a specific address, and ensure that the application is only reachable through that address. In theory, careful selection of the binding address could mitigate the problems described in Section 6.2. Binding services to temporary addresses could mitigate the ability of an attacker from testing for the presence of the node in the network. Binding different services to different addresses could mitigate unexpected discovery. Binding services to non-globally-reachable addresses (e.g. link-local addresses or ULAs) could mitigate availability outside the expected domain. However, explicitly managing addresses adds significant complexity to application development. It requires that application developers master IPv6 addressing architecture subtleties, and implement logic that reacts adequately to connectivity events and address changes. Experience shows that application developers would probably prefer some much simpler solution.

In addition, we note that many application developers use high level APIs that listen to TLS, HTTP, or some other application protocol. These high level APIs seldomly provide detailed access to specific IPv6 addresses, and typically default to listening to all available addresses.

A more advanced API could allow application programmers to select desired properties in an address (scope, stability, etc.), such that the best-suitable addresses are selected, while relieving the application from low-level IPv6 addressing details. Such API could also trigger the generation of new IPv6 addresses if/when the specified properties require so.

7.5. Universal Support of RFC 8028

To put it bluntly, multi-prefix/multi-router networks cannot possibly work properly without implementation of [RFC8028]. Unfortunately, [RFC8028] is not widely implemented yet. On the protocol standardization side, the IETF should consider elevating the requirement to support RFC8028 in the IPv6 Node Requirements RFC [RFC8504] from "SHOULD" to "MUST".

7.6. Support for Firewall Traversal in CE Routers

Customer Edge Routers that implement a default filtering policy of "only allowing outgoing communications" need to support helper protocols such as [UPnP] or PCP [RFC6887], so that applications can open holes in the CE Router firewall to be able to receive incoming communications. Otherwise, P2P applications that currently work in IPv4 networks might not function in IPv6-only networks.

Support for these protocols is particularly important for IPv6 deployments since, as hosts will normally employ "provider aggregatable" addresses (see Section 4.2), renumbering events will result in host address changes, and thus static firewall rules will be harder to implement than for the IPv4 networks. Similarly, use of temporary addresses [RFC4941] will also lead to changing IPv6 addresses, which will require that the associated firewall rules be updated.

8. IANA Considerations

This document has no IANA actions.

9. Security Considerations

The security and privacy implications associated with the predictability and lifetime of IPv6 addresses has been analyzed in [RFC7217] [RFC7721], and [RFC7707]. This document complements and extends the aforementioned analysis by also considering other IPv6 properties such as address scope and address reachability, and the associated trade-offs.

10. Acknowledgements

The authors would like to thank (in alphabetical order) Mikael Abrahamsson, Fred Baker, Brian Carpenter, Owen DeLong, Francis Dupont, Tatuya Jinmei, Ted Lemon, and Dave Thaler for providing valuable comments on earlier versions of this document.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, DOI 10.17487/RFC4007, March 2005, <https://www.rfc-editor.org/info/rfc4007>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <https://www.rfc-editor.org/info/rfc4193>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <https://www.rfc-editor.org/info/rfc4941>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/info/rfc5905>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <https://www.rfc-editor.org/info/rfc6724>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <https://www.rfc-editor.org/info/rfc6763>.

[RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <https://www.rfc-editor.org/info/rfc6887>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <https://www.rfc-editor.org/info/rfc7934>.

[RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016, <https://www.rfc-editor.org/info/rfc8028>.

[RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <https://www.rfc-editor.org/info/rfc8064>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

[RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <https://www.rfc-editor.org/info/rfc8504>.

11.2. Informative References

[Hein]   Hein, B., "The Rising Sophistication of Network Scanning", January 2016, <http://netpatterns.blogspot.be/2016/01/the-rising-sophistication-of-network.html>.

[I-D.gont-6man-ipv6-ula-scope]
Gont, F., "Scope of Unique Local IPv6 Unicast Addresses", draft-gont-6man-ipv6-ula-scope-00 (work in progress), January 2021.

[I-D.gont-opsawg-firewalls-analysis]
Gont, F. and F. Baker, "On Firewalls in Network Security", draft-gont-opsawg-firewalls-analysis-02 (work in progress), February 2016.

[I-D.ietf-6man-rfc4941bis]
Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", draft-ietf-6man-rfc4941bis-12 (work in progress), November 2020.

[I-D.ietf-6man-slaac-renum]
Gont, F., Zorz, J., and R. Patterson, "Improving the Robustness of Stateless Address Autoconfiguration (SLAAC) to Flash Renumbering Events", draft-ietf-6man-slaac-renum-02 (work in progress), January 2021.

[I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-12 (work in progress), October 2020.

[I-D.ietf-v6ops-cpe-slaac-renum]
Gont, F., Zorz, J., Patterson, R., and B. Volz, "Improving the Reaction of Customer Edge Routers to Renumbering Events", draft-ietf-v6ops-cpe-slaac-renum-06 (work in progress), December 2020.

[I-D.ietf-v6ops-dhcpv6-slaac-problem]
Liu, B., Jiang, S., Gong, X., Wang, W., and E. Rey, "DHCPv6/SLAAC Interaction Problems on Address and DNS Configuration", draft-ietf-v6ops-dhcpv6-slaac-problem-07 (work in progress), August 2016.

[I-D.ietf-v6ops-slaac-renum]
Gont, F., Zorz, J., and R. Patterson, "Reaction of Stateless Address Autoconfiguration (SLAAC) to Flash-Renumbering Events", draft-ietf-v6ops-slaac-renum-05 (work in progress), November 2020.

[RFC4953] Touch, J., "Defending TCP Against Spoofing Attacks", RFC 4953, DOI 10.17487/RFC4953, July 2007, <https://www.rfc-editor.org/info/rfc4953>.

[RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, DOI 10.17487/RFC5887, May 2010, <https://www.rfc-editor.org/info/rfc5887>.

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <https://www.rfc-editor.org/info/rfc6296>.

[RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <https://www.rfc-editor.org/info/rfc7039>.

[RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", RFC 7707, DOI 10.17487/RFC7707, March 2016, <https://www.rfc-editor.org/info/rfc7707>.

[RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <https://www.rfc-editor.org/info/rfc7721>.

[RFC8190] Bonica, R., Cotton, M., Haberman, B., and L. Vegoda, "Updates to the Special-Purpose IP Address Registries", BCP 153, RFC 8190, DOI 10.17487/RFC8190, June 2017, <https://www.rfc-editor.org/info/rfc8190>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/info/rfc8799>.

[UPnP]   UPnP, "UPnP Device Architecture 2.0", April 17, 2020, <https://openconnectivity.org/upnp-specs/UPnP-arch-DeviceArchitecture-v2.0-20200417.pdf>.

Authors' Addresses

Fernando Gont
SI6 Networks
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Email: fgont@si6networks.com
URI: https://www.si6networks.com

Guillermo Gont
SI6 Networks
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Email: ggont@si6networks.com
URI: https://www.si6networks.com

Network Working Group S. Peng
Internet-Draft Z. Li
Intended status: Informational Huawei Technologies
Expires: July 26, 2021 C. Xie
China Telecom
Z. Qin
China Unicom
G. Mishra
Verizon Inc.
January 22, 2021

Processing of the Hop-by-Hop Options Header
draft-peng-v6ops-hbh-03

Abstract

This document describes the processing of the Hop-by-Hop Options Header (HBH) in today's routers in the aspects of standards specification, common implementations, and default operations. This document outlines the reasons why the Hop-by-Hop Options Header is rarely utilized in current networks. In addition, this document describes how the HBH could be used as a powerful mechanism allowing deployment and operations of new services requiring a more optimized way to leverage network resources of an infrastructure. The Hop-by-Hop Options Header is taken into consideration by several network operators as a valuable container for carrying the information facilitating the introduction of new services. The desired, and proposed, processing behavior of the HBH and the migration strategies towards it are also suggested.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Due to historical reasons, such as incapable ASICs, limited IPv6 deployments, and few service requirements, the most common Hop-by-Hop Options header (HBH) processing implementation is that the node sends the IPv6 packets with the Hop-by-Hop Options header to the slow path (i.e., the control plane) of the node. The option type of each option carried within the Hop-by-Hop Options header will not even be examined before the packet is sent to the slow path. Very often, such processing behavior is the default configuration or, even worse, is the only behavior of the ipv6 implementation of the node.

Such default processing behavior of the Hop-by-Hop Options header could result in various unpleasant effects such as a risk of Denial of Service (DoS) attack on the router control plane and inconsistent packet drops due to rate limiting on the interface between the router control plane and forwarding plane, which will impact the normal end-to-end IP forwarding of the network services.

This actually introduced a circular problem:

-> An implementation problem caused HBH to become a DoS vector.

-> Because HBH is a DoS vector, network operators deployed ACLs that discard packets containing HBH.

-> Because network operators deployed ACLs that discard packets containing HBH, network designers stopped defining new HBH Options.

-> Because network designers stopped defining new HBH Options, the community was not motivated to fix the implementation problem that cause HBH to become a DoS vector.

The purpose of this draft is to break the cycle described above, fixing the problem that caused HBH not actually being utilized in operators' networks so to allow a better leverage of the HBH capability.

Driven by the wide deployments of IPv6 and ever-emerging new services, the Hop-by-Hop Options Header is taken as a valuable container for carrying the information to facilitate these new services.

This document suggests the desired processing behavior and the migration strategies towards it.

2. Modern Router Architecture

Modern router architecture design maintains a strict separation of the router control plane and its forwarding plane [RFC6192], as shown in Figure 1. Either the control plane or the forwarding plane is composed of both software and hardware, but each plane is responsible for different functions.

```
                          +----------------+
                          | Router Control |
                          |     Plane      |
                          +------+-+-------+
                                 | |
                            Interface Z
                                 | |
                          +------+-+-------+
                          |   Forwarding   |
          Interface X ==[     Plane      ]== Interface Y
                          +----------------+
```

Figure 1. Modern Router Architecture

The router control plane supports routing and management functions, handling packets destined to the device as well as building and sending packets originated locally on the device, and also drives the programming of the forwarding plane. The router control plane is generally realized in software on general-purpose processors, and its hardware is usually not optimized for high-speed packet handling. Because of the wide range of functionality, it is more susceptible to security vulnerabilities and a more likely a target for a DoS attack.

The forwarding plane is typically responsible for receiving a packet on an incoming interface, performing a lookup to identify the packet's next hop and determine the outgoing interface towards the destination, and forwarding the packet out through the appropriate outgoing interface. Typically, forwarding plane functionality is realized in high-performance Application Specific Integrated Circuits (ASICs) or Network Processors (NPs) that are capable of handling very high packet rates.

The router control plane interfaces with its forwarding plane through the Interface Z, as shown in the Figure 1, and the forwarding plane connects to other network devices via Interfaces such as X and Y. Since the router control plane is vulnerable to the DoS attack, usually a traffic filtering mechanism is implemented on Interface Z in order to block unwanted traffic. In order to protect the router control plane, a rate-limiting mechanism is always implemented on this interface. However, such rate limiting mechanism will always cause inconsistent packet drops, which will impact the normal IP forwarding.

Semiconductor chip technology has advanced significantly in the last decade, and as such the widely used network processing and forwarding process can now not only forward packets at line speed, but also easily support other feature processing such as QoS for DiffServ/

MPLS, Access Control List (ACL), Firewall, and Deep Packet Inspection (DPI).

A Network Processing Unit (NPU) is a non-ASIC based Integrated Circuit (IC) that is programmable through software. It performs all packet header operations between the physical layer interface and the switching fabric such as packet parsing and forwarding, modification, and forwarding. Many equipment vendors implement these functions in fixed function ASICs rather than using "off-the-shelf" NPUs, because of proprietary algorithms. Classification Co-processor is a specialized processor that can be used to lighten the processing load on an NPU by handling the parsing and classification of incoming packets such as IPv6 extended header HBH options processing. This advancement enables network processors to do the general process to handle simple control messages for traffic management, such as signaling for hardware programming, congestion state report, OAM, etc. Industry trend is for intelligent multi-core CPU fast path hardware using modern NPUs for forwarding packets at line rate while still being able to perform other complex tasks such as HBH forwarding options processing without having to punt to slow path.

Many of the fast-path packet-processing devices employed in modern switch and router designs are fixed-function ASICs to handle proprietary functions. While these devices can be very efficient for the set of functions they are designed for, they can be very inflexible. There is a tradeoff of price, performance and flexibility when vendors make a choice to use a fixed function ASIC as opposed to NPU. Due to the inflexibility of the fixed function ASIC, tasks that require additional processing such as IPv6 HBH header processing must be punted to the slow path. This problem is still a challenge today and is the reason why operators to protect against control plane DOS attack vector must drop or ignore HBH options. As industry shifts to Merchant Silicon based NPU evolution from fixed function ASIC, the gap will continue to close increasing the viability ubiquitous HBH use cases due to now processing in the fast path.

Most modern routers maintain a strict separation between forwarding plane and control plane hardware. Forwarding plane "fast path" bandwidth and resources are plentiful, while control plane "slow path" bandwidth and resources are constrained. In order to protect scarce control plane resources, routers enforce policies that restrict access from the forwarding plane to the control plane. Effective policies address packets containing the HBH Options Extension header, because HBH control options require access from the forwarding plane to the control plane. Many network operators perceive HBH Options to be a breach of the separation between the forwarding and control planes. In this case HBH control options would be required to be punted to slow path by fixed function ASICs as well as NPUs.

The maximum length of an HBH Options header is 2,048 bytes. A source node can encode hundreds of options in 2,048 bytes. With today's technology it would be cost prohibitive to be able to process hundreds of options with either NPU or proprietary fixed function ASIC.

While [RFC8200] required that all nodes must examine and process the Hop-by-Hop Options header, it is now expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so. This can be beneficial in cases where transit nodes are legacy hardware and the destination endpoint PE is newer NPU based hardware that can process HBH in the fast path.

IPv6 Extended Header limitations that need to be addressed to make HBH processing more efficient and viable in the fast path:

[RFC8504] defines the IPv6 node requirements and how to protect a node from excessive header chain and excessive header options with various limitations that can be defined on a node. [RFC8883] defines ICMPv6 Errors for discarding packets due to processing limits. Per [RFC8200] HBH options must be processed serially. However, an implementation of options processing can be made to be done with more parallelism in serial processing grouping of similar options to be processed in parallel.

The IPv6 standard does not currently limit the header chain length or number of options that can be encoded.

Each Option is encoded in a TLV and so processing of a long list of TLVs is expensive. Zero data length encoded options TLVs are a valid option. A DOS vector could be easily generated by encoding 1000 HBH options (Zero data length) in a standard 1500 MTU packet. So now imagine if you have a Christmas tree long header chain to parse each with many options.

3. Specification of RFC 8200

[RFC8200] defines several IPv6 extension header types, including the Hop-by-Hop (HBH) Options header. As specified in [RFC8200], the Hop-by-Hop (HBH) Options header is used to carry optional information that will be examined and processed by every node along a packet's delivery path, and it is identified by a Next Header value of zero in the IPv6 header.

The Hop-by-Hop (HBH) Options header contains the following fields:

-- Next Header: 8-bit selector, identifies the type of header immediately following the Hop-by-Hop Options header.

-- Hdr Ext Len: 8-bit unsigned integer, the length of the Hop-by-Hop Options header in 8-octet units, not including the first 8 octets.

-- Options: Variable-length field, of length such that the complete Hop-by-Hop Options header is an integer multiple of 8 octets long.

The Hop-by-Hop (HBH) Options header carries a variable number of "options" that are encoded in the format of type-length-value (TLV).

The highest-order two bits (i.e., the ACT bits) of the Option Type specify the action that must be taken if the processing IPv6 node does not recognize the Option Type. The third-highest-order bit (i.e., the CHG bit) of the Option Type specifies whether or not the Option Data of that option can change en route to the packet's final destination.

While [RFC2460] required that all nodes must examine and process the Hop-by-Hop Options header, with [RFC8200] it is expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so. It means that the HBH processing behavior in a node depends on its configuration.

However, in the current [RFC8200], there is no explicit specification of the possible configurations. Therefore, the nodes may be configured to ignore the Hop-by-Hop Options header, drop packets containing a Hop-by-Hop Options header, or assign packets containing a Hop-by-Hop Options header to a slow processing path [RFC8200]. Because of these likely uncertain processing behaviors, new hop-by-hop options are not recommended.

4. Common Implementations

In the current common implementations, once an IPv6 packet, with its Next Header field set to 0, arrives at a node, it will be directly sent to the slow path (i.e., the control plane) of the node. With such implementations, the value of the Next Header field in the IPv6 header is the only trigger for the default processing behavior. The option type of each option carried within the Hop-by-Hop Options header will not even be examined before the packet is sent to the slow path.

Very often, such processing behavior is the default configuration on the node, which is embedded in the implementation and cannot be changed or reconfigured.

Another critical component of IPv6 HBH processing which is in some cases is overlooked is the operator core network which can be designed to use the global Internet routing table for internet traffic and in other cases use an overlay MPLS VPN to carry Internet traffic. In the global Internet routing table scenario where only an underlay global routing table exists, and no VPN overlay carrying customer Internet traffic, the IPv6 HBH options can be used as a DOS attack vector for both the operator nodes, adjacent inter-as peer nodes as well as customer nodes along a path. In a case where the Internet routing table is carried in a MPLS VPN overlay payload, the HBH options header does not impact the operator underlay framework and only impacts the VPN overlay payload and thus the operator underlay topmost label global table routing FEC LSP instantiation is not impacted as the operator underlay is within the operators closed domain. However HBH options DOS attack vector in the VPN overlay can still impact the customer CE destination end nodes as well as other adjacent inter-as operators that only use underlay global Internet routing table. In an operator closed domain where MPLS VPN overlay is utilized to carry internet traffic, the operator has full control of the underlay and IPv6 Extended header chain length as well as the number of HBH options encoded. However in contrast, in the global routing table scenario for Internet traffic there is no way to control the IPv6 Extended header chain lenghth as well as the number of HBH forward or HBH control options encoded.

4.1. Historical Reasons

When IPv6 was first implemented on high-speed routers, HBH options were not yet well-understood and ASICs were not as capable as they are today. So, early IPv6 implementations dispatched all packets that contain HBH options to their slow path.

4.2. Consequences

Such implementation introduces a risk of a DoS attack on the control plane of the node, and a large flow of IPv6 packets could congest the slow path, causing other critical functions (including routing and network management) that are executed on the control plane to fail. Rate limiting mechanisms will cause inconsistent packet drops and impact the normal end-to-end IP forwarding of the network services.

5. Operators' Typical Processing

To mitigate this DoS vulnerability, many operators deployed Access Control Lists (ACLs) that discard all packets containing HBH Options.

[RFC6564] shows the Reports from the field indicating that some IP routers deployed within the global Internet are configured either to ignore or to drop packets having a hop-by-hop header. As stated in [RFC7872], many network operators perceive HBH Options to be a breach of the separation between the forwarding and control planes. Therefore, several network operators configured their nodes so as to discard all packets containing the HBH Options Extension Header, while others configured nodes to forward the packet but to ignore the HBH Options. [RFC7045] also states that hop-by-hop options are not handled by many high-speed routers or are processed only on a slow path.

Due to such behaviors observed and described in these specifications, new hop-by-hop options are not recommended in [RFC8200] hence the usability of HBH options is severely limited.

6. New Services

As IPv6 is being rapidly and widely deployed worldwide, more and more applications and network services are migrating to or directly adopting IPv6. More and more new services that require HBH are emerging and the HBH Options header is going to be utilized by the new services in various scenarios.

In-situ OAM (IOAM) with IPv6 encapsulation [I-D.ietf-ippm-ioam-ipv6-options] is one of the examples. IOAM in IPv6 is used to enhance diagnostics of IPv6 networks and complements other mechanisms, such as the IPv6 Performance and Diagnostic Metrics Destination Option described in [RFC8250]. The IOAM data fields are encapsulated in "option data" fields of the Hop-by-Hop Options header if Pre-allocated Tracing Option, Incremental Tracing Option, or Proof of Transit Option are carried [I-D.ietf-ippm-ioam-data], that is, the IOAM performs per hop.

Alternate Marking Method can be used as the passive performance measurement tool in an IPv6 domain. The AltMark Option is defined as a new IPv6 extension header option to encode alternate marking technique and Hop-by-Hop Options Header is considered [I-D.ietf-6man-ipv6-alt-mark].

The Minimum Path MTU Hop-by-Hop Option is defined in [I-D.ietf-6man-mtu-option] to record the minimum Path MTU along the forward path between a source host to a destination host. This Hop-by-Hop option is intended to be used in environments like Data Centers and on paths between Data Centers as well as other environments including the general Internet. It provides a useful tool for allowing to better take advantage of paths able to support a large Path MTU.

As more services start utilizing the HBH Options header, more packets containing HBH Options are going to be injected into the networks. According to the current common configuration in most network deployments, all the packets of the new services are going to be sent to the control plane of the nodes, with the possible consequence of causing a DoS on the control plane. The packets will be dropped and the normal IP forwarding may be severely impacted. The deployment of new network services involving multi-vendor interoperability will become impossible.

7. The Desired Processing Behavior

The following requirements SHOULD be met:

o The control plane SHOULD be protected from undesired traffic.

-* The HBH options header SHOULD NOT be directly sent to the control plane once the packets are received since these options may not aim for the control plane.

-* The HBH options that are not supposed to be processed by the control plane SHOULD NOT be sent to the control plane, potentially causing the DoS attack.

o Since generally the two types of HBH options (control plane (e.g., Route Alert Option [RFC2711]) and forwarding plane (e.g., AltMark Option [I-D.ietf-6man-ipv6-alt-mark])) serve different purposes and require different processing procedures on a node, they should be encoded separately and carried in different packets.

Note: More details on the two types of HBH options can be found in [I-D.li-6man-hbh-fwd-hdr].

o The packets carrying the HBH Forwarding Options are supposed to be maintained in the forwarding plane rather than being directly sent up to the control plane. While the packets carrying the HBH Control Options are supposed to be sent to the control plane.

o The source node SHOULD NOT encode the HBH Options that exceed the maximum length of an HBH Options header i.e. 2,048 bytes.

- o The source node SHOULD NOT encode the number of HBH Options that exceeds the lowest processing capability of the nodes along the path.

- o The source node SHOULD NOT encode the HBH Options that exceed the maximum overall length of the IPv6 extensions header chain.

- o The options aimed for the control plane are better if they do not consume the forwarding plane resources.

- o A simple and efficient way to discriminate the two types of HBH options is required.

- o The new deployments should be compatible with the existing deployments, since default configuration of some devices running in the networks cannot be changed or reconfigured. The update of the networks in operation will usually take time.

- o If the IPv6 extension header including the HBH options header of a packet cannot be recognized by the node, or the option in the HBH header is unknown to the node, and the node is not the destination of the packet, the packet SHOULD NOT be dropped or sent to the control plane, rather this unrecognized extension header should be skipped and the rest of the packet should be processed.

8. Migration Strategies

In order to achieve the desired processing behavior of the HBH options header and facilitate the ever-emerging new services to be deployed in operators' networks across multiple vendors' devices, the migration can happen in three parts as described below:

1. The source of the HBH options header encapsulation.

The information to be carried in the HBH options header needs to be first categorized and encapsulated into either control options or forwarding options, and then encapsulated in different packets.

2. The nodes within the network.

The nodes within the network are updated to the proposed behavior introduced in the previous section.

3. The edge nodes of the network.

The edge nodes should check whether the packet contains an HBH header with control or forwarding option. Packets with a control option may still be filtered and dropped while packets with forwarding option SHOULD be allowed by the ACL.

If it is certain that there is no harm that can be introduced by the HBH control options to the nodes and the services, they can also be allowed.

Note: During the migration stage, the nodes that are not yet updated will stay with their existing configurations.

9. Security Considerations

The same as the Security Considerations apply as in [RFC8200] for the part related with the HBH Options header.

10. IANA Considerations

This document does not include an IANA request.

11. Acknowledgements

The authors would like to acknowledge Ron Bonica, Fred Baker, Bob Hinden, Stefano Previdi, and Donald Eastlake for their valuable review and comments.

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.

[RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <https://www.rfc-editor.org/info/rfc6192>.

[RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <https://www.rfc-editor.org/info/rfc7045>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <https://www.rfc-editor.org/info/rfc7872>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

12.2. Informative References

[I-D.ietf-6man-ipv6-alt-mark] Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate Marking Method", draft-ietf-6man-ipv6-alt-mark-02 (work in progress), October 2020.

[I-D.ietf-6man-mtu-option] Hinden, R. and G. Fairhurst, "IPv6 Minimum Path MTU Hop-by-Hop Option", draft-ietf-6man-mtu-option-04 (work in progress), October 2020.

[I-D.ietf-ippm-ioam-data] Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-11 (work in progress), November 2020.

[I-D.ietf-ippm-ioam-ipv6-options] Bhandari, S., Brockners, F., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., Spiegel, M., Krishnan, S., Asati, R., and M. Smith, "In-situ OAM IPv6 Options", draft-ietf-ippm-ioam-ipv6-options-04 (work in progress), November 2020.

[I-D.li-6man-hbh-fwd-hdr] Li, Z. and S. Peng, "Hop-by-Hop Forwarding Options Header", draft-li-6man-hbh-fwd-hdr-00 (work in progress), July 2020.

[RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, DOI 10.17487/RFC2711, October 1999, <https://www.rfc-editor.org/info/rfc2711>.

[RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <https://www.rfc-editor.org/info/rfc8250>.

[RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <https://www.rfc-editor.org/info/rfc8504>.

[RFC8883] Herbert, T., "ICMPv6 Errors for Discarding Packets Due to Processing Limits", RFC 8883, DOI 10.17487/RFC8883, September 2020, <https://www.rfc-editor.org/info/rfc8883>.

Authors' Addresses

Shuping Peng
Huawei Technologies
Beijing
China

Email: pengshuping@huawei.com

Zhenbin Li
Huawei Technologies
Beijing
China

Email: lizhenbin@huawei.com

Chongfeng Xie
China Telecom
China

Email: xiechf@chinatelecom.cn

Zhuangzhuang Qin
China Unicom
Beijing
China

Email: qinzhuangzhuang@chinaunicom.cn

Gyan Mishra
Verizon Inc.
USA

Email: gyan.s.mishra@verizon.com

V6OPS G. Fioccola
Internet-Draft P. Volpato
Intended status: Informational Huawei Technologies
Expires: August 26, 2021 N. Elkins
Inside Products
S. Lourdez
Post Luxembourg
February 22, 2021

IPv6 Deployment Status
draft-vf-v6ops-ipv6-deployment-02

Abstract

Looking globally, IPv6 is growing faster than IPv4 and this means that the collective wisdom of the networking industry has selected IPv6 for the future. This document provides an overview of IPv6 transition deployment status and a view on how the transition to IPv6 is progressing among network operators and enterprises that are introducing IPv6 or have already adopted an IPv6-only solution. It also aims to analyze the transition challenges and therefore encourage actions and more investigations on some areas that are still under discussion. The overall IPv6 incentives are also examined.

Requirements Language

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

Copyright Notice

Table of Contents

1. Introduction

The focus of this document is to provide a survey of the deployed IPv6 transition technologies and to highlight the difficulties in the transition. This process helps to understand what is missing and how to improve the current IPv6 deployment strategies of network operators, enterprises, content and cloud service providers. The objective is to give an updated view of the practices and plans already described in [RFC6036]. The scope is to report the current IPv6 status and encourage actions and more investigations on some areas that are still under discussion as well as the main incentives for the IPv6 adoption.

[RFC6180] discussed the IPv6 deployment models and migration tools. [RFC6036] described the Service Provider Scenarios for IPv6 Deployment, [RFC7381] introduced the guidelines of the IPv6 deployment for Enterprise and [RFC6883] provided guidance and suggestions for Internet Content Providers and Application Service Providers. On the other hand, this document focuses on the end-to-end services and in particular on the device - network - content communication chain.

[ETSI-IP6-WhitePaper] reported the IPv6 Best Practices, Benefits, Transition Challenges and the Way Forward. IPv6 is becoming a priority again and a new wave of IPv6 deployment is expected, due the exhaustion of the IPv4 address space since 2010, in addition technologies like 5G, cloud, IoT require its use, governments and standard bodies (including IETF) demand it, and the device - network - content communication chain is calling for its adoption. In this regard it is possible to mention the IAB Statement on IPv6 stating that "IETF will stop requiring IPv4 compatibility in new or extended protocols".

The following sections go through the issue of IPv4 address exhaustion and give the global picture of IPv6 to show how IPv6 is growing faster than IPv4 worldwide in all measures including number of users, percentage of content, and amount of traffic. This testifies that the key Internet industry players have decided strategically to invest and deploy IPv6 in large-scale to sustain the Internet growth.

Then it is presented the survey among network operators as well as considerations and observations for enterprises and content and cloud service providers about the IPv6 deployment and the considerations that have come out. IPv6 transition solutions for Mobile BroadBand (MBB), Fixed BroadBand (FBB) and enterprise services are ready. Dual-Stack is the most deployed solution for IPv6 introduction, while 464XLAT and Dual Stack Lite (DS-Lite) seem the most suitable for IPv6-only service delivery.

Finally, The IPv6 incentives are presented but the general IPv6 challenges are also reported in particular in relation to Architecture, Operations, Performance and Security issues. These considerations aim to start a call for action on the areas of improvement, that are often mentioned as reason for not deploying IP6.

2. IPv4 Adress Exhaustion

According to [CAIR] there will be 29.3 billion networked devices by 2023, up from 18.4 billion in 2018. This poses the question on whether the IPv4 address space can sustain such a number of allocations and, consequently, if this is affecting the process of its exhaustion. The answer is not straightforward as many aspects have to be considered.

On the one hand, the RIRs are reporting scarcity of available and still reserved addresses. Table 3 of [POTAROO1] shows that the available pool of the five RIRs counts a little more than 6 million IPv4 address, while the reserved pool includes another 12 million, for a total of "usable" addresses equal to 18.3 million. The same reference, in table 1, shows that the total IPv4 allocated pool equals 3.684 billion addresses. The ratio between the "usable" addresses and the total allocated brings to 0.005% of remaining space.

On the other, [POTAROO1] again highlights the role of both NAT and the address transfer to counter the IPv4 exhaustion. NAT systems well fit in the current client/server model used by most of the available Internet applications, with this phenomenon amplified by the general shift to cloud. The transfer of IPv4 addresses also contributes to mitigate the the need of addresses. As an example, [IGP-GT] shows the amount of transfers to recipient organizations in the ARIN region in 2018. Cloud Service Providers (CSPs) appear to be the most active is buying available addresses to satisfy their need of providing IPv4 connectivity to their tenants.

3. The global picture of IPv6

The utilization of IPv6 has been monitored by many agencies and institutions worldwide. Different analytics have been made available, ranging from the number of IPv6 users, its relative utilization over the Internet, to the number of carriers able to route IPv6 network prefixes. [ETSI-IP6-WhitePaper] provided several of those analytics. The scope of this section then is to summarize the status of the IPv6 adoption, so to get an indication of the relevance of IPv6 today. For the analytics listed here, the trend over the past five years is given, expressed as the Compound Annual Growth Rate (CAGR). In general, this shows how IPv6 has grown in the past few years, and that is growing faster than IPv4.

3.1. IPv6 users

[ETSI-IP6-WhitePaper] provided the main statistics about the utilization of IPv6 worldwide and references the organizations that make their measurement publicly available through their web sites. To give a rough estimation of the relative growth of IPv6, the next table shows the total number of estimated IPv6 users at December 2020 as measured by [POTAROO2], [APNIC1].

	Dec 2016	Dec 2017	Dec 2018	Dec 2019	Dec 2020	CAGR
World	300.85	473.14	543.04	990.19	1,201.09	41%

Figure 1: IPv6 users worldwide (in millions)

3.2. IPv6 allocations and networks

Regional Internet Registries (RIRs) are responsible for assigning an IPv6 address block to ISPs or enterprises. An ISP will use the assigned block to provide addresses to their end users. For example, a mobile carrier will assign one or several /64 prefixes to the end users. Several analytics are available for the RIRs. The next table shows the amount of individual allocations, per RIR, in the time period 2016-2020 [APNIC2].

Registry	Dec 2016	Dec 2017	Dec 2018	Dec 2019	Dec 2020	Cumulated	CAGR
AFRINIC	116	112	110	115	109	562	48%
APNIC	1,681	1,369	1,474	1,484	1,498	7,506	45%
ARIN	646	684	659	605	644	3,238	50%
LACNIC	1,009	1,549	1,448	1,614	1,801	7,421	65%
RIPE NCC	2,141	2,051	2,620	3,104	1,403	11,319	52%
Total	5,593	5,765	6,311	6,922	5,455	30,046	52%

Figure 2: IPv6 allocations worldwide

Note that the decline in 2020 of IPv6 allocations from the RIPE NCC could be explained with the COVID-19 measures that affect many European countries. Anyway countries all over the world have been similarly affected, but the decline in IPv6 allocation activity in 2020 is only seen in the data from the RIPE NCC.

[APNIC2] also compares the number of allocations for both address families, and the result is in favor of IPv6. The average yearly growth is 52% for IPv6 in the period 2016-2020 versus 49% for IPv4, a sign that IPv6 is growing bigger than IPv4. This is described in the next table.

Address family	Dec 2016	Dec 2017	Dec 2018	Dec 2019	Dec 2020	Cumulated	CAGR
IPv6	5,593	5,765	6,311	6,922	5,455	30,046	52%
IPv4	10,515	9,437	10,192	14,019	7,437	51,600	49%

Figure 3: Allocations per address family

The next table is based on [APNIC3], [APNIC4] and shows the percentage of ASes supporting IPv6 compared to the total ASes worldwide. The number of IPv6-capable ASes increases from 22.6% in January 2017 to 30.4% in January 2021. This equals to 14% CAGR for IPv6 enabled networks. This also shows that the number of networks supporting IPv6 is growing faster than the ones supporting IPv4, since the total (IPv6 and IPv4) networks grow at 6% CAGR.

Advertised ASN	Jan 2017	Jan 2018	Jan 2019	Jan 2020	Jan 2021	CAGR
IPv6-capable	12,700	14,500	16,470	18,600	21,400	14%
Total ASN	56,100	59,700	63,100	66,800	70,400	6%
Ratio	22.6%	24.3%	26.1%	27.8%	30.4%	

Figure 4: Percentage of IPv6-capable ASes

4. Survey among Network Operators

It was started an IPv6 poll to more than 50 network operators about the status of IPv6 deployment. This poll reveals that more than 30 operators will migrate fixed and mobile users to IPv6 in next 2 years. The IPv6 Poll has been submitted in particular to network operators considering that, as showed by the previous section, both user devices and contents seem more ready for IPv6. The answers to the questionnaire can be found in Appendix.

The main Questions asked are:

* Do you plan to move more fixed or mobile or enterprise users to IPv6 (e.g. Dual-Stack) or IPv6-only in the next 2 years? What are the reasons to do so? Which transition solution will you use, Dual-Stack, DS-Lite, 464XLAT, MAP-T/E?

* Do you need to change network devices for the above goal? Will you migrate your metro or backbone or backhaul network to support IPv6?

The result of this questionnaire highlights that major IPv6 migration will happen in next 2 years. Dual Stack is always the most adopted solution and the transition to IPv6-only is motivated in particular by business reasons like the 5G and IoT requirements. In addition it is worth mentioning that the migration of transport network (metro and backbone) is not considered a priority today for many network operators and the focus is in particular on the end to end IPv6 services.

More details about the answers received can be found in the Appendix.

5. Considerations for Enterprises

As described in [RFC7381], enterprises face different challenges than operators. The overall problem for many enterprises is to handle IPv6-based connectivity to the upstream providers, while supporting a mixed IPv4/IPv6 domain in the internal network.

The business reasons for IPv6 is unique to each enterprise especially for the internal network. But the most common drivers are on the external network due to the fact that when Internet service providers, run out of IPv4 addresses, they will provide native IPv6 and non-native IPv4. So for client networks trying to reach enterprise networks, the IPv6 experience will be better than the transitional IPv4 if the enterprise deploys IPv6 in its public-facing services. Enterprise that is or will be expanding into emerging markets or that partners with other companies who use IPv6 (larger enterprise, governments, service providers) has to deploy IPv6 or plan to do in the near term to support the long term goals. As an example it is possible to mention the emerging energy market and in partiuclar SmartGrid where high density of IP-enabled endpoints are needed and IPv6 is a key technology.

6. Observations on Content and Cloud Service Providers

The number of addresses required to connect all of the virtual and physical elements in a Data Center and the necessity to overcome the limitation posed by [RFC1918] has been the driver to adopt IPv6 in several Content and Cloud Service Provider (CSP) networks.

Several public references discuss how most of the major players find themselves at different stages in the transition to IPv6-only in their DC infrastructure. In some cases, the transition already happened and the DC infrastructure of these hyperscalers is completely based on IPv6. This can be considered a good sign because the end-to-end connectivity between a client (e.g. an application on a smartphone) and a server (a Virtual Machine in a DC) may be based on IPv6.

7. Industrial Internet application

There are potential advantages for implementing IPv6 for IIoT (Industrial Internet of Things) applications, in particular the large IPv6 address space, the automatic IPv6 configuration and resource discovery.

However, there are still many obstacles that prevent its pervasive use. The key problems identified are the incomplete or immature tool support, the dependency on manual configuration and the poor knowledge of the IPv6 protocols among insiders. To advance and ease the use of IPv6 for smart manufacturing systems and IIoT applications in general, a generic approach to remove these pain points is therefore highly desirable.

8. IPv6 deployments worldwide

This section reports the most deployed approaches for the IPv6 migration in MBB, FBB and enterprise.

8.1. IPv6 service design for Mobile, Fixed broadband and enterprises

The consolidated strategy, as also described in [ETSI-IP6-WhitePaper], is based on two stages, namely: (1) IPv6 introduction, and (2) IPv6-only. The first stage aims at delivering the service in a controlled manner, where the traffic volume of IPv6-based services is minimal. When the service conditions change, e.g. when the traffic grows beyond a certain threshold, then the move to the second stage may occur. In this latter case, the service is delivered solely on IPv6.

8.1.1. IPv6 introduction

In order to enable the deployment of an IPv6 service over an underlay IPv4 architecture, there are two possible approaches:

- o Enabling Dual-Stack at the CPE

- o Tunneling IPv6 traffic over IPv4, e.g. with 6rd.

So, from a technical perspective, the first stage is based on Dual-Stack [RFC4213] or tunnel-based mechanisms such as Generic Routing Encapsulation (GRE), IPv6 Rapid Deployment (6rd), Connection of IPv6 Domains via IPv4 Clouds (6to4), and others.

Dual-Stack [RFC4213] is more robust, and easier to troubleshoot and support. Based on information provided by operators with the answers to the poll (see Appendix A), it can be stated that Dual-Stack is currently the most widely deployed IPv6 solution, for MBB, FBB and enterprises, accounting for about 50% of all IPv6 deployments, see both Appendix A and the statistics reported in [ETSI-IP6-WhitePaper]. Therefore, for operators that are willing to introduce IPv6 the most common approach is to apply the Dual-Stack transition solution.

With Dual-Stack, IPv6 can be introduced together with other network upgrade and many parts of network management and IT systems can still work in IPv4. This avoids major upgrade of such systems to support IPv6, which is possibly the most difficult task in IPv6 transition.

In other words, the cost and effort on the network management and IT system upgrade are moderate. The benefits are to start to accommodate future services and save the NAT costs.

The CPE has only an IPv6 address at the WAN side and uses an IPv6 connection to the operator gateway, e.g. Broadband Network Gateway (BNG) or Packet Gateway (PGW) / User Plane Function (UPF). However, the hosts and content servers can still be IPv4 and/or IPv6. For example, NAT64 can enable IPv6 hosts to access IPv4 servers. The backbone network underlay can also be IPv4 or IPv6.

Although the Dual-Stack IPv6 transition is a good solution to be followed in the IPv6 introduction stage, it does have few disadvantages in the long run, like the duplication of the network resources and states, as well as other limitations for network operation. For this reason, when IPv6 increases to a certain limit, it would be better to switch to the IPv6-only stage.

8.1.2. IPv6-only service delivery

The second stage, named here IPv6-only, can be a complex decision that depends on several factors, such as economic factors, policy and government regulation.

[I-D.lmhp-v6ops-transition-comparison] discusses and compares the technical merits of the most common transition solutions for IPv6-only service delivery, 464XLAT, DS-lite, Lightweight 4over6 (lw4o6), MAP-E, and MAP-T, but without providing an explicit recommendation. As the poll highlights, the most widely deployed IPv6 transition solution for MBB is 464XLAT and for FBB is DS-Lite.

Based on the survey among network operators in Appendix A it is possible to analyze the IPv6 transition technologies that are already deployed or that will be deployed. The different answers to the questionnaire and in particular [ETSI-IP6-WhitePaper] reported detailed statistics on that and it can be stated that, besides Dual-Stack, the most widely deployed IPv6 transition solution for MBB is 464XLAT [RFC6877], and for FBB is DS-Lite [RFC6333], both of which are IPv6-only solutions.

Looking at the different feedback from network operators, in some cases, even when using private addresses, such as 10.0.0.0/8 space [RFC1918], the address pool is not large enough, e.g. for large mobile operators or large Data Centers (DCs), Dual-Stack is not enough, because it still requires IPv4 addresses to be assigned. Also, Dual-Stack will likely lead to duplication of several network operations both in IPv6 and IPv4 and this increases the amount of state information in the network with a waste of resources. For this reason, in some scenarios (e.g. MBB or DCs) IPv6-only stage could be more efficient from the start since the IPv6 introduction phase with Dual-Stack may consume more resources (for example CGNAT costs).

So, in general, it is possible to state that, when the Dual-Stack disadvantages outweigh the IPv6-only complexity, it makes sense to migrate to IPv6-only. Some network operators already started this process, while others are still waiting.

9. Findings of the IPv6 Survey

Global IPv4 address depletion is reported by most network operators as the important driver for IPv6 deployment. Indeed, the main reason for IPv6 deployment given is related to the run out of private 10.0.0.0/8 space [RFC1918]. 5G and IoT service deployment is another incentive not only for business reasons but also for the need of more addresses.

The answers in Appendix shows that the IPv6 deployment strategy is based mainly on Dual Stack architecture and most of the network operators are migrating or plan to migrate in the next few years. The main motivation is related to the depletion of IPv4 addresses and to save the NAT costs.

It is interesting to see that most of the network operators have no big plans to migrate transport network (metro and backbone) soon, since they do not see business reasons. It seems that there is no pressure to migrate to native IPv6 forwarding in the short term, anyway the future benefit of IPv6 may justify in the long term a migration to native IPv6. Some network operators also said that a software upgrade can be enough to support IPv6 where it is needed for now.

This survey demonstrates that full replacement of IPv4 will take long time. Indeed the transition to IPv6 has different impacts and requirements depending on the network segment:

- It is possible to say that almost all mobile devices are already IPv6 capable while for fixed access most of the CPEs are Dual Stack. Data Centers are also evolving and deploying IPv6 to cope with the increasing demand of cloud services.

- While the access network seems not strongly impacted because it is mainly based on layer 2 traffic, regarding Edge and BNG, most network operators that provide IPv6 connectivity runs BNG devices in Dual Stack in order to distribute both IPv4 and IPv6.

- o For Metro and Backbone, the trend is to keep MPLS Data Plane and run IPv6/IPv4 over PE devices at the border. All MPLS services can be guaranteed in IPv6 as well through 6PE/6VPE protocols.

In this scenario it is clear that the complete deployment of a full IPv6 data plane will take more time. If we look at the long term evolution, IPv6 can bring other advantages like introducing advanced protocols developed only on IPv6 (e.g. SRv6) to implement all the controlled SLA services aimed by the 5G technology and beyond.

10. IPv6 incentives

It is possible to state that IPv6 adoption is no longer optional, indeed there are several incentives for the IPv6 deployment:

Technical incentives: all Internet technical standard bodies and network equipment vendors have endorsed IPv6 and view it as the standards-based solution to the IPv4 address shortage. The IETF, as well as other SDOs, need to ensure that their standards do not assume IPv4. The IAB expects that the IETF will stop requiring IPv4 compatibility in new or extended protocols. Future IETF protocol work will then optimize for and depend on IPv6. It is recommended that all networking standards assume the use of IPv6 and be written so they do not require IPv4 ([RFC6540]). In addition, every Internet registry worldwide strongly recommends immediate IPv6 adoption.

Business incentives: with the emergence of new digital technologies, such as 5G, IOT and Cloud, new use cases have come into being and posed more new requirements for IPv6 deployment. Over time, numerous technical and economic stop-gap measures have been developed in an attempt to extend the lifetime of IPv4, but all of these measures add cost and complexity to network infrastructure and raise significant barriers to innovation. It is widely recognized that full transition to IPv6 is the only viable option to ensure future growth and innovation in Internet technology and services. Several large networks and Data Centers have already evolved their internal infrastructures to be IPv6-only. Forward looking large corporations are also working toward migrating their enterprise networks to IPv6-only environments.

Governments incentives: governments have a huge responsibility in promoting IPv6 deployment within their countries. There are example of governments already adopting policies to encourage IPv6 utilization or enforce increased security on IPv4. So, even without funding the IPv6 transition, governments can recommend to add IPv6 compatibility for every connectivity, service or products bid. This will encourage the network operators and vendors who don't want to miss out on government related bids to evolve their infrastructure to be IPv6 capable. Any public incentives for technical evolution will be bonded to IPv6 capabilities of the technology itself. In this regard, in the United States, the Office of Management and Budget is calling for an implementation plan to have 80% of the IP-enabled resources on Federal networks be IPv6-only by 2025. If resources cannot be converted, then the Federal agency is required to have a plan to retire them. The Call for Comment is at [US-FR] and [US-CIO].

11. Call for action

There are some areas of improvement, that are often mentioned in the literature and during the discussions on IPv6 deployment. This section lists these topics and wants to start a call for action to encourage more investigations on these aspects.

11.1. Transition choices

From an architectural perspective, a service provider or an enterprise may perceive quite a complex task the transition to IPv6, due to the many technical alternatives available and the changes required in management and operations. Moreover, the choice of the method to support the transition may depend on factors specific to the operator's or the enterprise's context, such as the IPv6 network design that fits the service requirements, the deployment strategy, and the service and network operations.

This section briefly highlights the basic approaches that service providers and enterprises may take. The scope is to raise the discussion whether actions may be taken that allow to overcome the issues highlighted and further push the adoption of IPv6.

11.1.1. Service providers

For a service provider, the IPv6 transition often refers to the service architecture (also referred to as overlay) and not to the network architecture (underlay). IPv6 is introduced at the service layer when a service requiring IPv6-based connectivity is deployed in an IPv4-based network. In this case, as already mentioned in the previous sections, a strategy is based on two stages: IPv6 introduction and IPv6-only.

For fixed operators, the massive CPE software upgrade to support Dual Stack started in most of service providers network and the traffic percentage is currently between 30% and 40% of IPv6, looking at the global statistics. This is valid for a network operator that provides Dual Stack and gives the same opportunity for end terminal applications to choose freely the path that they want and assuming a normal internet usage. Anyway, it is interesting to see that in the latest years all major content providers have already implemented dual stack access to their services and most of them have implemented IPv6-only in their Data Centers. This aspect could affect the decision on the IPv6 adoption for an operator, but there are also other aspects like the current IPv4 addressing status, CPE costs, CGNAT costs and so on. Most operators already understood the need to adopt IPv6 in their networks and services, and also to promote the diffusion into their clients, while others are still at the edge of a massive implementation decision. Indeed, two situations are possible:

Operators that have already employed CGNAT and have introduced IPv6 in their networks, so they remain attached to a Dual Stack architecture. Although IPv6 brought them to a more technological advanced state, CGNAT, on the other end, boosts for some time their ability to supply CPE IPv4 connectivity.

Operators with a Dual Stack architecture that have introduced IPv6 both in the backbone and for the CPEs, but when reaching the limit in terms of number of IPv4 addresses available, they need to start defining and start to apply a new strategy that can be through CGNAT or with an IPv6-only approach.

For mobile operators, the situation is different since they are stretching their IPv4 address space since CGNAT translation levels have been reached and no more IPv4 public pool addresses are available. The new requirements from IoT services, 5G 3GPP release implementations, Voice over Long-Term Evolution (VoLTE) together with the constraints of national regulator lawful interception are seen as major drivers for IPv6. For these reasons, two situations are possible:

Some mobile operators choose to implement Dual-Stack as first and immediate mitigation solution.

Other mobile operators prefer to move to IPv6-only solution(e.g. 464XLAT) since Dual-Stack only mitigates and does not solve completely the IPv4 number scarcity issue.

11.1.2. Enterprises

The dual stage approach described in the previous sections can be still applicable for enterprises, even if the priorities to apply either stage are different since they have to consider both the internal and external network:

It is possible to start with Dual-Stack on hosts/OS and then in client network distribution layer. This allows the IPv6 introduction independently since both hosts/OS and client networks belong to the domain of the enterprise.

Dual-Stack can be further extended to WAN/campus core/edge routers. Also, as temporary solution, the use of NAT64 is recommended for servers/apps only capable of IPv4. Enterprise Data Center is also to be considered for the IPv6 transition. In this regard the application support needs to be taken into account, even if virtualization should make DCs simpler and more flexible.

There are additional challenges also related to the campus network and the cloud interconnection, indeed the networking may be not homogeneous. IPv6 could help to build a flat network by leveraging SD-WAN integration. The perspective of IPv6-only could also ensure better end-to-end performance.

Enterprises (private, managed networks) in US and Europe have failed to adopt IPv6, especially on internal networks. Other countries, in particular in Asia, who faced a shortage of IPv4 addresses, have moved somewhat more quickly. But, even there, the large "brick-and-mortar" enterprises find no business reason to adopt IPv6.

The enterprise engineers and technicians also don't know how IPv6 works. The technicians want to get trained yet the management does not feel that they do not want to pay for such training because they do not see a business need for adoption. This creates an unfortunate cycle where misinformation about the complexity of the IPv6 protocol and unreasonable fears about security and manageability combine with the perceived lack of urgent business needs to prevent adoption of IPv6.

In 2019 and 2020, there has been a concerted effort by some grass roots non-profits working with ARIN and APNIC to provide training [ARIN-CG] [ISIF-ASIA-G].

Having said that, some problems such as the problem of application conversion from IPv6 are quite difficult. The reliance of the economic, governmental, and military enterprise organizations on computer applications is great; the number of legacy systems, and ossification at such organizations, is also great. A number of mission-critical computer applications were written in the 1970's. While they have the source code, no one at the enterprise may be familiar with the application nor do they have the funds for external resources. So, transitioning to IPv6 is quite difficult.

The problem may be that of "First Mover Disadvantage". Understandably, corporations, having responsibility to their stockholders, have upgraded to new technologies and architectures, such as IPv6, only if it gains them revenue. Thus, legacy programs and technical debt accumulate.

11.1.3. Cloud and Data Centers

It was already highlighted how CSPs have adopted IPv6 in their internal infrastructure but are also active in gathering IPv4 addresses on the transfer market to serve the current business needs of IPv4 connectivity. This is primarily directed to serve the transition to cloud of enterprise's applications.

As noted in the previous section, most enterprises do not consider the transition to IPv6 as a priority. To this extent, the use of IPv4-based network services by the CSPs will last. Yet, CSPs are struggling to buy IPv4 addresses. If, in the next years, the scarcity of IPv4 addresses becomes more evident, it is likely that the cost of buying an IPv4 address by a CSP will be charged to an enterprise as a fee. From a financial standpoint this effect might be taken into consideration when evaluating the decision of moving to IPv6.

11.1.4. Industrial Internet

As the most promising protocol for network applications, IPv6 is frequently mentioned in relation to Internet of Things and Industry 4.0. However, its industrial adoption, in particular in smart manufacturing systems, has been much slower than expected. Indeed, it is important to provide an easy way to familiarize system architects and software developers with the IPv6 protocol and its role in the application development life cycle in order to limit the dependency on manual configuration and improve the tool support.

It is possible to differentiate types of data and access to understand how and where the IPv6 transition can happen. In the control network, determinism is required with full operational visibility and control, as well as reliability and availability. In monitoring IoT, best effort can be acceptable and low OPEX, zero-touch functions autoconfiguration, zero-configuration. For diagnostics and alerts, trust and transmissions that do not impact the control network are needed. For safety, guarantees in terms of redundancy, latency similar to the control network but with total assurance, is necessary.

For IIoT applications, it would be desirable to be able to implement a truly distributed system without dependencies to central components like a DHCP server. In this regard the distributed IIoT applications can leverage the configuration-less characteristic of IPv6 and in this regard all the possible problems and compatibility issues with IPv6 link local addresses, SLAAC (StateLess Address Auto Configuration) needs to be investigated.

In addition, it could be interesting to have the ability to use IP based communication and standard application protocols at every point in the production process and further reduce the use of specialized communication systems like PLCs (Programmable Logic Controllers) and fieldbuses for real-time control to subsystems where this is absolutely necessary.

11.1.5. Government and Regulators

The slogan should be "stimulate if you can, regulate if you must". The global picture shows that the deployment of IPv6 worldwide is not uniform at all [G_stats], [APNIC1]. Countries where either market conditions or local regulators have stimulated the adoption of IPv6 show clear sign of growth.

As an example, zooming into the European Union area, countries such as Belgium, France and Germany are well ahead in terms of IPv6 adoption. The French National Regulator, Arcep, can be considered a good reference of National support to IPv6. [ARCEP] introduced an obligation for the operators awarded with a license to use 5G frequencies (3.4-3.8GHz) in Metropolitan France to be IPv6 compatible. As stated, "the goal is to ensure that services are interoperable and to remove obstacles to using services that are only available in IPv6, as the number of devices in use continues to soar, and because the RIPE NCC has run out of IPv4 addresses". A slow adoption of IPv6 could prevent new Internet services to widespread or create a barrier to entry for newcomers to the market. "IPv6 can help to increase competition in the telecom industry, and help to industrialize a country for specific vertical sectors".

A renewed industrial policy might be advocated in other countries and regions to stimulate IPv6 adoption. As an example, in the United States, the Office of Management and Budget is also calling for IPv6 adoption [US-FR], [US-CIO].

11.2. Network Operations

An important factor is represented by the need for training the network operations workforce. Deploying IPv6 requires it as policies and procedures have to be adjusted in order to successfully plan and complete an IPv6 migration. Staff has to be aware of the best practices for managing IPv4 and IPv6 assets. In addition to network nodes, network management applications and equipment need to be properly configured and in same cases also replaced. This may introduce more complexity and costs for the migration.

11.3. Performance

Despite their relative differences, people tend to compare the performance of IPv6 versus IPv4, even if these differences are not so important for applications. In some cases, IPv6 behaving "worse" than IPv4 tends to re-enforce the justification of not moving towards the full adoption of IPv6. This position is supported when looking at available analytics on two critical parameters: packet loss and latency. These parameters have been constantly monitored over time, but only a few extensive researches and measurement campaigns are currently providing up-to-date information. This paragraph will look briefly at both of them, considering the available measurements. Operators are invited to bring in their experience and enrich the information reported below.

11.3.1. IPv6 latency

[APNIC5] constantly compares the latency of both address families. Currently, the worldwide average is still in favor of IPv4. Zooming at the country or even at the operator level, it is possible to get more detailed information and appreciate that cases exist where IPv6 is faster than IPv4. [APRICOT] highlights how when a difference in performance exists it is often related to asymmetric routing issues. Other possible explanations for a relative latency difference lays on the specificity of the IPv6 header which allows packet fragmentation. In turn, this means that hardware needs to spend cycles to analyze all of the header sections and when it is not capable of handling one of them it drops the packet. Even considering this, a difference in latency stands and sometimes it is perceived as a limiting factor for IPv6. A few measurement campaigns on the behavior of IPv6 in Content Delivery Networks (CDN) are also available [MAPRG-IETF99], [INFOCOM]. The TCP connect time is still higher for IPv6 in both cases, even if the gap has reduced over the analysis time window.

11.3.2. IPv6 packet loss

[APNIC5] also provides the failure rate of IPv6. Two reports, namely [RIPE1] and [APRICOT], discussed the associated trend, showing how the average worldwide failure rate of IPv6 worsened from around 1.5% in 2016 to a value exceeding 2% in 2020. Reasons for this effect may be found in endpoints with an unreachable IPv6 address, routing instability or firewall behaviours. Yet, this worsening effect may appeae as disturbing for a plain transition to IPv6. Operators are once again invited to share their experience and discuss the performance of IPv6 in their network scenarios.

11.3.3. Router's performance

It is worth mentioning the aspect of Router's performance too. IPv6 is 4 times longer than IPv4 and it is possible to do a simple calculation: the same memory on routers could permit to have 1/4 of different tables (routing, filtering, next hop). Anyway most of the routers showed a remarkably similar throughput and latency for IPv4 and IPv6. For smaller software switching platforms, some tests reported a lower throughput for IPv6 compared to IPv4 only in case of smaller packet sizes, while for larger hardware switching platforms there was no throughput variance between IPv6 and IPv4 both at larger frame sizes and at the smaller packet size.

11.4. IPv6 security

IPv6 presents a number of exciting possibilities for the expanding global Internet, however, there are also noted security challenges associated with the transition to IPv6. [I-D.ietf-opsec-v6] analyzes the operational security issues in several places of a network (enterprises, service providers and residential users).

The security aspects have to be considered to keep the same level of security as it exists nowadays in an IPv4-only network environment. The autoconfiguration features of IPv6 will require some more attention for the things going on at the network level. Router discovery and address autoconfiguration may produce unexpected results and security holes. The IPsec protocol implementation has initially been set as mandatory in every node of the network, but then relaxed to recommendation due to extremely constrained hardware deployed in some devices e.g., sensors, Internet of Things (IoT).

There are some concerns in terms of the security but, on the other hand, IPv6 offers increased efficiency. There are measurable benefits to IPv6 to notice, like more transparency, improved mobility, and also end to end security (if implemented).

As reported in [ISOC], comparing IPv6 and IPv4 at the protocol level, one may probably conclude that the increased complexity of IPv6 results in an increased number of attack vectors, that imply more possible ways to perform different types attacks. However, a more interesting and practical question is how IPv6 deployments compare to IPv4 deployments in terms of security. In that sense, there are a number of aspects to consider.

Most security vulnerabilities related to network protocols are based on implementation flaws. Typically, security researchers find vulnerabilities in protocol implementations, which eventually are "patched" to mitigate such vulnerabilities. Over time, this process of finding and patching vulnerabilities results in more robust implementations. For obvious reasons, the IPv4 protocols have benefited from the work of security researchers for much longer, and thus IPv4 implementations are generally more robust than IPv6.

Besides the intrinsic properties of the protocols, the security level of the resulting deployments is closely related to the level of expertise of network and security engineers. In that sense, there is obviously much more experience and confidence with deploying and operating IPv4 networks than with deploying and operating IPv6 networks.

Finally, implementation of IPv6 security controls obviously depends on the availability of features in security devices and tools. Whilst there have been improvements in this area, there is a lack of parity in terms of features and/or performance when considering IPv4 and IPv6 support in security devices and tools.

11.4.1. Protocols security issues

It is important to say that IPv6 is not more or less secure than IPv4 and the knowledge of the protocol is the best security measure.

In general there are security concerns related to IPv6 that can be classified as follows:

- o Basic IPv6 protocol (Basic header, Extension Headers, Addressing)

- o IPv6 associated protocols (ICMPv6, NDP, MLD, DNS, DHCPv6)

- o Internet-wide IPv6 security (Filtering, DDoS, Transition Mechanisms)

ICMPv6 is an integral part of IPv6 and performs error reporting and diagnostic functions. Since it is used in many IPv6 related protocols, ICMPv6 packet with multicast address should be filtered carefully to avoid attacks. Neighbor Discovery Protocol (NDP) is a node discovery protocol in IPv6 which replaces and enhances functions of ARP. Multicast Listener Discovery (MLD) is used by IPv6 routers for discovering multicast listeners on a directly attached link, much like Internet Group Management Protocol (IGMP) is used in IPv4.

These IPv6 associated protocols like ICMPv6, NDP and MLD are something new compared to IPv4, so they adds new security threats and the related solutions are still under discussion today. NDP has vulnerabilities [RFC3756] [RFC6583]. The specification says to use IPsec but it is impractical and not used, on the other hand, SEND (SEcure Neighbour Discovery) [RFC3971] is not widely available.

[RIPE2] describes the most important threats and solutions regarding IPv6 security.

11.4.2. IPv6 Extension Headers and Fragmentation

IPv6 Extension Headers imply some issues, in particular their flexibility also means an increased complexity, indeed security devices and software must process the full chain of headers while firewalls must be able to filter based on Extension Headers. Additionally, packets with IPv6 Extension Headers may be dropped in the public Internet.

There are some possible attacks through EHs, for example RH0 can be used for traffic amplification over a remote path and it is deprecated. Other attacks based on Extension Headers are based on IPv6 Header Chains and Fragmentation that could be used to bypass filtering, but, to mitigate this effect, Header chain should go only in the first fragment and the use of the IPv6 Fragmentation Header is forbidden in all Neighbor Discovery messages.

Fragment Header is used by IPv6 source node to send a packet bigger than path MTU and the Destination host processes fragment headers. There are several threats related to fragmentation to pay attention to e.g. overlapping fragments (not allowed) resource consumption while waiting for last fragment (to discard), atomic fragments (to be isolated).

11.4.3. Oversized IPv6 packets

A lot of additional functionality has been added to IPv6 primarily by adding Extension Headers and/or using overlay encapsulation. All of the these expand the packet size and this could lead to oversized packets that would be dropped on some links.

It is better to investigate the potential problems with oversized packets in the first place. Fragmentation must not be done in transit and a better solution needs to be found, e.g. upgrade all links to bigger MTU or follow specific recommendations at the source node.

12. Security Considerations

This document has no impact on the security properties of specific IPv6 protocols or transition tools. The security considerations relating to the protocols and transition tools are described in the relevant documents.

13. Contributors

TBC

14. Acknowledgements

TBC

15. IANA Considerations

This document has no actions for IANA.

16. References

16.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

16.2. Informative References

[APNIC1] APNIC, "IPv6 Capable Rate by country (%)", 2020, <https://stats.labs.apnic.net/ipv6>.

[APNIC2] APNIC2, "Addressing 2020", 2021, <https://labs.apnic.net/?p=1400>.

[APNIC3] APNIC, "BGP in 2019 - The BGP Table", 2020, <https://blog.apnic.net/2020/01/14/bgp-in-2019-the-bgp-table/>.

[APNIC4] APNIC, "IPv6 in 2020", 2021, <https://blog.apnic.net/2021/02/08/ipv6-in-2020/>.

[APNIC5] APNIC, "Average RTT Difference (ms) (V6 - V4) for World (XA)", 2020, <https://stats.labs.apnic.net/v6perf/XA>.

[APRICOT] Huston, G., "Average RTT Difference (ms) (V6 - V4) for World (XA)", 2020, <https://2020.apricot.net/assets/files/APAE432/ipv6-performance-measurement.pdf>.

[ARCEP] ARCEP, "Arcep Decision no 2019-1386, Decision on the terms and conditions for awarding licences to use frequencies in the 3.4-3.8GHz band", 2019, <https://www.arcep.fr/uploads/tx_gsavis/19-1386.pdf>.

[ARIN-CG] ARIN, "Community Grant Program: IPv6 Security, Applications, and Training for Enterprises", 2020, <https://www.arin.net/about/community_grants/recipients/>.

[CAIR] Cisco, "Cisco Annual Internet Report (2018-2023) White Paper", 2020, <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>.

[ETSI-IP6-WhitePaper] ETSI, "ETSI White Paper No. 35: IPv6 Best Practices, Benefits, Transition Challenges and the Way Forward", ISBN 979-10-92620-31-1, 2020.

[G_stats] Google, "Google IPv6 Per-Country IPv6 adoption", 2021, <https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption>.

[I-D.ietf-opsec-v6] Vyncke, E., Kk, C., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", draft-ietf-opsec-v6-21 (work in progress), November 2019.

[I-D.lmhp-v6ops-transition-comparison] Lencse, G., Martinez, J., Howard, L., Patterson, R., and I. Farrer, "Pros and Cons of IPv6 Transition Technologies for IPv4aaS", draft-lmhp-v6ops-transition-comparison-06 (work in progress), January 2021.

[IGP-GT] Internet Governance Project, Georgia Tech, "The hidden standards war: economic factors affecting IPv6 deployment", 2019, <https://via.hypothes.is/https://www.internetgovernance.org/wp-content/uploads/IPv6-Migration-Study-final-report.pdf>.

[INFOCOM] Doan, T., "A Longitudinal View of Netflix: Content Delivery over IPv6 and Content Cache Deployments", 2020, <https://dl.acm.org/doi/abs/10.1109/INFOCOM41043.2020.9155367>.

[ISIF-ASIA-G] ISIF Asia, "Internet Operations Research Grant: IPv6 Deployment at Enterprises. IIESoc. India", 2020, <https://isif.asia/2020-grantees/>.

[ISOC] Internet Society, "IPv6 Security FAQ", 2019, <https://www.internetsociety.org/wp-content/uploads/2019/02/Deploy360-IPv6-Security-FAQ.pdf>.

[MAPRG-IETF99] Bajpai, V., "Measuring YouTube Content Delivery over IPv6", 2017, <https://www.ietf.org/proceedings/99/slides/slides-99-maprg-measuring-youtube-content-delivery-over-ipv6-00.pdf>.

[POTAROO1] POTAROO, "Addressing 2020", 2020, <https://www.potaroo.net/ispcol/2021-01/addr2020.html>.

[POTAROO2] POTAROO, "IPv6 Resource Distribution Reports", 2021, <https://resources.potaroo.net/iso3166/archive/>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

[RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <https://www.rfc-editor.org/info/rfc3756>.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <https://www.rfc-editor.org/info/rfc3971>.

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <https://www.rfc-editor.org/info/rfc4213>.

[RFC6036] Carpenter, B. and S. Jiang, "Emerging Service Provider Scenarios for IPv6 Deployment", RFC 6036, DOI 10.17487/RFC6036, October 2010, <https://www.rfc-editor.org/info/rfc6036>.

[RFC6180] Arkko, J. and F. Baker, "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment", RFC 6180, DOI 10.17487/RFC6180, May 2011, <https://www.rfc-editor.org/info/rfc6180>.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <https://www.rfc-editor.org/info/rfc6333>.

[RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, DOI 10.17487/RFC6540, April 2012, <https://www.rfc-editor.org/info/rfc6540>.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, DOI 10.17487/RFC6583, March 2012, <https://www.rfc-editor.org/info/rfc6583>.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <https://www.rfc-editor.org/info/rfc6877>.

[RFC6883] Carpenter, B. and S. Jiang, "IPv6 Guidance for Internet Content Providers and Application Service Providers", RFC 6883, DOI 10.17487/RFC6883, March 2013, <https://www.rfc-editor.org/info/rfc6883>.

[RFC7381] Chittimaneni, K., Chown, T., Howard, L., Kuarsingh, V., Pouffary, Y., and E. Vyncke, "Enterprise IPv6 Deployment Guidelines", RFC 7381, DOI 10.17487/RFC7381, October 2014, <https://www.rfc-editor.org/info/rfc7381>.

[RIPE1] Huston, G., "Measuring IPv6 Performance", 2016, <https://ripe73.ripe.net/wp-content/uploads/presentations/35-2016-10-24-v6-performance.pdf>.

[RIPE2] RIPE, "IPv6 Security", 2019, <https://www.ripe.net/support/training/material/ipv6-security/ipv6security-slides.pdf>.

[US-CIO] The CIO Council, "Memorandum for Heads of Executive Departments and Agencies. Completing the Transition to Internet Protocol Version 6 (IPv6)", 2020, <https://www.cio.gov/assets/resources/internet-protocol-version6-draft.pdf>.

[US-FR] Federal Register, "Request for Comments on Updated Guidance for Completing the Transition to the Next Generation Internet Protocol, Internet Protocol Version 6 (IPv6)", 2020, <https://www.federalregister.gov/documents/2020/03/02/2020-04202/request-for-comments-on-updated-guidance-for-completing-the-transition-to-the-next-generation>.

Appendix A. Summary of Questionnaire and Replies

This Appendix summarizes the questionnaire and the replies received.

1. Do you have plan to move more fixed or mobile or enterprise users to IPv6 in the next 2 years?

a. If yes, fixed, or mobile, or enterprise?

b. What're the reasons to do so?

c. When to start: already on going, in 12 months, after 12 months?

d. Which transition solution will you use, Dual-Stack, DS-Lite, 464XLAT, MAP-T/E?

2. Do you need to change network devices for the above goal?

a. If yes, what kind of devices: CPE, or BNG/mobile core, or NAT?

b. Will you migrate your metro or backbone or backhaul network to support IPv6?

Some answers below:

Answer 1: (1) Yes, IPv6 migration strategy relies upon the deployment of Dual Stack architecture. IPv4 service continuity designs is based on DS-Lite for fixed environments and 464XLAT for mobile environments. No plans to move towards MAP-E or MAP-T solutions for the time being. (2) Yes, it's a matter of upgrading CPE, routers (including BNGs), etc. Tunneling options (ISATAP, TEREDO, 6rd) will also be used for migration.

Answer 2: (1) Yes, at this moment we widely use IPv6 for mobile services while we are using DS-Lite for fixed services (FTTH and DSL). (2) We have no pressure to migrate to native IPv6 forwarding in the short term and it would represent a significant work without clear immediate benefit or business rationale. However we may see a future benefit with SRv6 which may justify in the long term a migration to native IPv6.

Answer 3: (1) Yes, fixed. The IP depletion topic is crucial, so we need to speed up the DS-Lite deployment and also Carrier Grade NAT introduction. (2) Yes, CGNAT introduction.

Answer 4: (1) No, we are rolling IPv6 users back to IPv4. DS-Lite. (2) No, it was already done. IPv6 works worse than IPv4. it is immature.

Answer 5: (1) Yes, all 3. Target is Dual-stack for fixed, mobile and enterprise. (2) Yes, we are adding specific services cards inside our FTTH equipment for dealing with CGNAT. Metro and backbone are already Dual Stack.

Answer 6: (1) Yes, Enterprises customer demand is high and the transition is on going through Dual-Stack. (2) No big plan for transport network.

Answer 7: No such requirements

Answer 8: (1) Yes, mobile. The Internet APN is not yet enabled for IPv6, this will be done soon. 464XLAT will be used to save on RFC1918 address space. (2) Yes, PGW; Metro is already IPv6 and Backbone is currently IPv4/MPLS. No native IPv6 planned as for now.

Answer 9: (1) Yes, Dual-Stack for all 3. Not all services are available on IPv6. IPv6 adoption has been stated from many years but still not finished. Dual-Stack is used. (2) No, at the moment it is 6PE solution. No plan to migrate on native IPv6.

Answer 10: (1) Yes, all 3. Ongoing transition with Dual-stack and 464XLAT. (2) No plan for Metro and Backbone.

Answer 11: No such requirements.

Answer 12: (1) Yes, mobile and fixed. To mitigate IPv4 exhaustion in 12 months, Dual-Stack is used. (2) No (hopefully). Managed by software upgrade.

Answer 13: (1) Yes, on Mobile and Fixed. Mobile: IPv4 exhaustion for the RAN transport and IPv6 roll out ongoing. Fixed: Enterprises are requesting IPv6 and also competitors are offering it. Mobile: dual stack and 6VPE; Enterprise: Dual Stack and 6VPE. (2) No, maybe only a software upgrade.

Answer 14: (1) Yes, fixed. IPv4 address depletion, on going, Dual-Stack with NAT444. (2) No.

Answer 15: (1) Yes, Mobile. Running out of private IPv4 address space and do not want to overlap addresses. Transition on going through 464XLAT. (2) Not yet, this is not the most pressing concern at the moment but it is planned.

Answer 16: No, already on Dual-Stack for many years. Discussing IPv6-only.

Answer 17: (1) Yes, all 3, strategy on going, Dual-Stack, MAP-T. (2) Yes, CPE, BR Dual-Stack.

Answer 18: (1) Yes, Mobile, due to address deficit. It would be very likely 464XLAT. (2) It is not clear at the moment. Still under investigation. CPE, Mobile Core, NAT. For IPv6 native support no plans for today.

Answer 19: No. Difficult to do it for enterprises, and don't really care for residential customers.

Answer 20: (1) Yes, fixed, mobile. IP space depletion. Mobile and Backbone are already done, Fixed is becoming Dual-Stack. (2) Yes, ordinary CPE and small routers. Some of them needs just software upgrade. Backbone done, no plan for metro and backhaul.

Answer 21: No such requirements

Answer 22: (1) Yes, mobile, we have few enterprise requests for IPv6; fixed already Dual-Stack. We are in the exhaustion point in public IPv4 usage in mobile so we need to move to IPv6 in the terminals. Dual-Stack deployment is ongoing. (2) No, all devices already support dual-stack mode. No migration needed. We already support IPv6 forwarding in our backbone.

Answer 23: No, already Dual-Stack

Answer 24: (1) Yes, fixed. DS-Lite. (2) Yes, BNG supporting CGNAT.

Answer 25: (1) Yes, fixed. DS-Lite will be deployed. (2) Yes.

Answer 26: (1) Yes, Mobile (Fixed already Dual-Stack). IPv4 depletion and Business customers are asking for it. Dual-Stack will be deployed. (2) No.

Answer 27: (1) Yes, Mobile. Dual-Stack is on going. (2) Yes, MBH, mobile core.

Answer 28: No such requirements.

Answer 29: (1) Yes, fixed and mobile, enterprise is not certain. IPv4 addressing is not enough, fixed and mobile should be started in 12 months. (2) Telco Cloud, BNG and PEs already support IPv6.

Answer 30: (1) Yes, all 3. Government has pushed. Dual-Stack for FBB in 12 months. (2) Yes, RGs have not good readiness, but not much could be done about it. PPPoE access does not create problem in access and aggregation. BNG should only change configuration.

Answer 31: (1) Yes, mobile for 5G sites. Plan to use IPv6 soon. 6VPE in the beginning, then migrate to Dual-stack. (2) IP BH devices already support IPv6.

Answer 32: No.

Answer 33: Yes, Enterprises. We are running short of IPV4 addresses. In our Internet Core IPV4/IPV6 Dual Stack was already introduced. The rollout of IPV6 services is slow and we started with business services. From customer perspective Dual Stack is still a "must have" and this will be true for many years to come. Another thought is related to regulatory obligations. Anyway a total switch from IPv4 to IPv6 will not be possible for many more years.

Answer 34: No, we have no plans to introduce new wave of IPv6 in our network.

Answer 35: (1) Yes. Fixed, Enterprise. IPv4 addressing is not enough. Dual Stack deployment is ongoing. (2) Yes, CPE for metro and backbone.

Answer 36: (1) Yes, Fixed, Enterprise. Dual-Stack. (2) Yes, CPE for IPv6 service delivery support.

Answer 37: Yes, mobile and enterprise. 6PE is deployed on the PEs, and dual-stack. The PE supports IPv6 by modifying the live network configuration or upgrading the software.

Answer 38: Yes, both home broadband and enterprise services support IPv6. IPv6 services are basic capabilities of communication networks. Currently 6RD, dual stack (native IPv6) in the future. The dual-stack feature does not require device changes. The home gateway is connected to the switch and the BNG. The Dual Stack can be supported through configuration changes. Both the metro and backbone networks use MPLS to provide bearer services and do not require IPv6 capabilities. IPv6 is not enabled on both the metro and backbone networks. IPv6 services are implemented through 6VPE.

Answer 39: (1) Yes, Enterprises B2B needs more IP addresses. Dual-Stack is already on going. (2) No, BNG/mobile core and NAT. Metro and Backbone already support today.

Answer 40: Not for now.

Authors' Addresses

Giuseppe Fioccola
Huawei Technologies
Riesstrasse, 25
Munich 80992
Germany

Email: giuseppe.fioccola@huawei.com

Paolo Volpato
Huawei Technologies
Via Lorenteggio, 240
Milan 20147
Italy

Email: paolo.volpato@huawei.com

Nalini Elkins
Inside Products
36A Upper Circle
Carmel Valley CA 93924
United States of America

Email: nalini.elkins@insidethestack.com

Sebastien Lourdez
Post Luxembourg

Email: sebastien.lourdez@post.lu