IPv6-Only vs IPv6-Mostly: Appropriate Use Cases

"Birds of a feather flock together"

A Matter of Terminology

IETF terminology is not always clearly defined, which can lead to confusion and, in some cases, the use of mechanisms outside the scenarios for which they were originally intended, potentially affecting the quality of service offered to users.

The terms IPv6-Only and IPv6-Mostly are a clear example of this. But despite the difficulty of reaching consensus on definitions, v6ops is making progress on this front through draft-palet-v6ops-ipv6-only.

One of the key points emerging from this work is that terms such as IPv6-Only and IPv6-Mostly only make sense when their scope is clearly defined. For example, saying that a network is IPv6-Only is probably incorrect, because it would imply that IPv4 is neither configured nor used anywhere in the network.

This is not a common scenario today, as services, applications, and internal or external content may still require IPv4 connectivity. If instead we say "IPv6-Only Access Network", we make it clear that only the access portion of the network is IPv6-only. The CPE LAN and other parts of the operator's network may still operate in dual-stack mode.

Understanding where these terms apply in practice helps explain why IPv6-Only and IPv6-Mostly should not be treated as interchangeable.

IPv6-Only+IPv4aaS

This is the most common scenario today in mobile access networks, using 464XLAT, and is a growing trend in residential access networks. It is one of the five IPv6-Only + IPv4aaS ("IPv4 as a Service") transition mechanisms for **access networks**. It allows the WAN portion to be IPv6-Only while maintaining dual-stack functionality on the user network, preventing any device or application from ceasing to function. Relevant documents include: RFC6877, RFC8585, RFC8683, RFC9313, draft-ietf-v6ops-rfc6146-bis, and draft-ietf-v6ops-rfc7084bis.

These five mechanisms, and especially **464XLAT** (which is the only one supported in mobile networks), have a clear use case: **unmanaged networks**. Users of these networks do not need to perform special configurations, regardless of whether the devices they have on their home networks or tethering networks (when the mobile phone becomes a router to provide access to other devices) use IPv4, IPv6, or both, since the CPE incorporates a stateless NAT46 mechanism (CLAT or

equivalent) to facilitate communication between IPv4 and IPv6, **without local communication traffic having to leave the local network or be translated**.

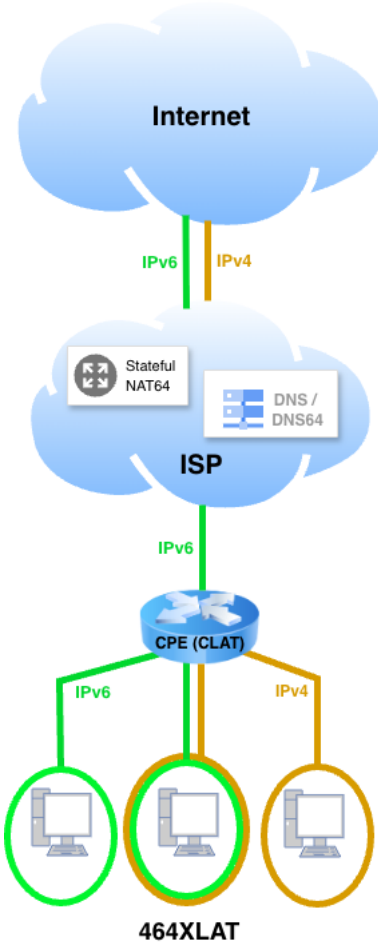

In recent years, IPv6 deployment has progressed globally, especially in mobile and residential networks. This is reflected in statistics from Akamai, APNIC, Facebook, Google, among others, which indicate IPv6 adoption rates exceeding 50%. This figure doesn't even include traffic from China (which has legal IPv6-Only obligations for ISPs), likely placing that percentage closer to 70-75% and growing rapidly in that country. However, while major content providers have deployed IPv6-Only data centres, organizations (both public and private) generally have low levels of IPv6 adoption.

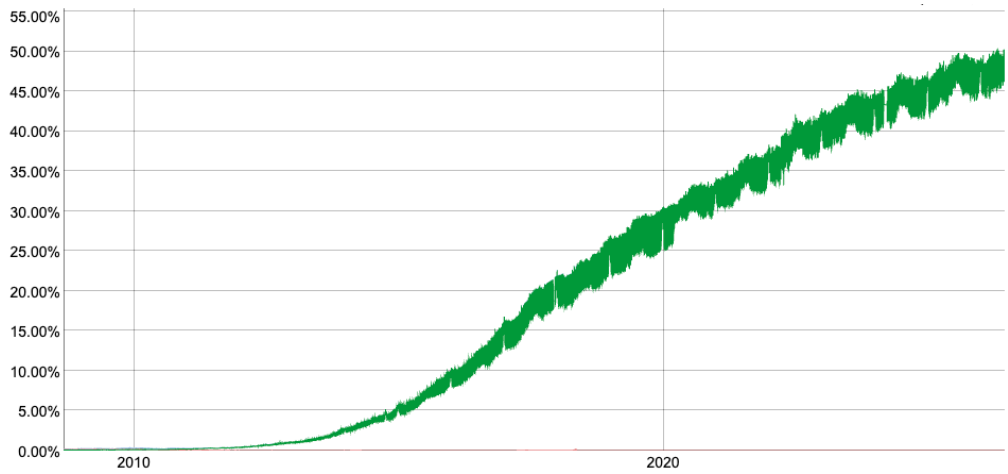

IPv6-Mostly

Therefore, the IETF has taken steps to facilitate deployment in **managed networks**, i.e., corporate networks, with mechanisms to achieve the "IPv6-Mostly" model (draft-ietf-v6ops-6mops).

IPv6-Mostly, compared to the dual-stack model in corporate networks, has several objectives:

- Reduce the consumption of private IPv4 addresses and DHCPv4 resources.
- Avoid masking IPv6 problems (due to HappyEyeballs' fallback to IPv4), facilitating the debugging of IPv6 errors.
- Enable a gradual, host-by-host transition to IPv6-Only, eliminating the need for dedicated IPv6-Only segments and simplifying operations, thus increasing scalability.
- Discover and prevent incompatibilities in a future IPv6-Only environment.
- Reduce NAT4 usage.

IPv6-Mostly is based on the premise of RFC8925: a dual-stack host can signal, using a new DHCPv4 option (108), that it does not need to use IPv4 on its interface. The goal is to avoid deploying dedicated IPv6-Only network segments, allowing the host to behave as such on a dual-stack segment, provided the network administrator deems it appropriate for that segment. To achieve this, the DHCPv4 server configuration on that segment must be enabled with support for option 108, enabling granular control over which hosts can disable IPv4.

However, disabling IPv4 on a host means it will not be able to access IPv4-Only destinations. Therefore, the network administrator should only enable option 108 if they have implemented support for stateful NAT64 (PLAT in 464XLAT terminology).

If support for literal IPv4 addresses, applications lacking IPv6 support, and similar situations are also required, hosts supporting option 108 must also incorporate CLAT. In fact, CLAT (draft-ietf-v6ops-claton) support in desktop operating systems

has been progressively implemented since the publication of RFC8925, first in macOS and Linux, and more recently (Public Preview) in Windows 11.

Mobile operating systems (iOS, Android, Harmony) have incorporated CLAT or equivalent features into their cellular interfaces since the inception of 464XLAT, and have recently added it to other interfaces as well (such as Wi-Fi). However, IPv6-Mostly not supported on the cellular interface because DHCPv4 is generally not implemented on that interface, either on the device itself or on mobile network infrastructure devices.

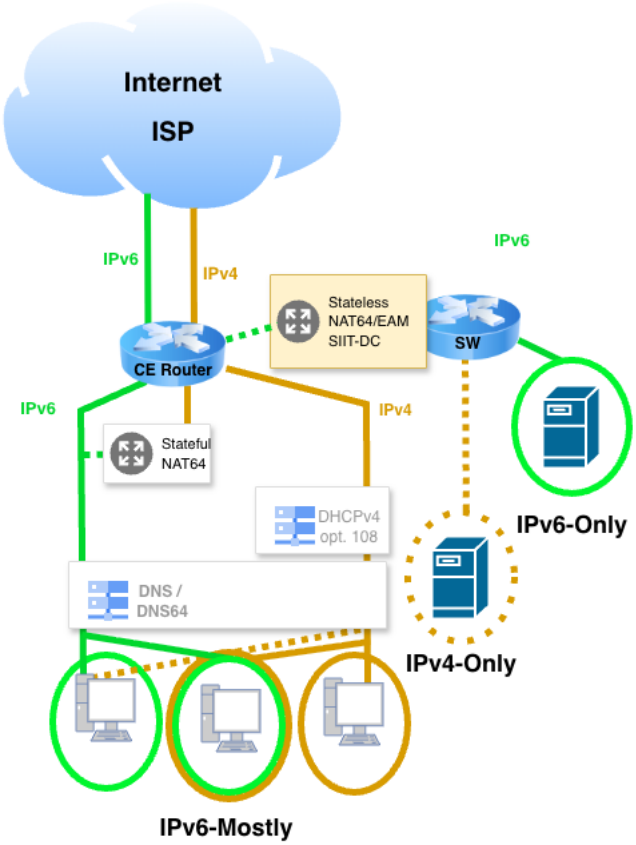

It's important to note that using IPv6-Only on a host also prevents communication with IPv4-Only devices on the same or other segments of that network (including external access via IPv4 in the case of servers). In other words, printers, cameras, and other IPv4-Only devices cannot be used by IPv6-Only devices. Obviously, in a corporate network, the administrator will explicitly configure solutions to allow such communication, such as through the internal PLAT, or stateless NAT64 mechanisms with EAM (Explicit Address Mappings, RFC7757/STD103), SIIT-DC (RFC7755), etc.

Commonalities

In both 464XLAT and IPv6-Mostly, various mechanisms have been used for discovering the NAT64 prefix, ranging from manual configuration to 3GPP's own mechanisms for mobile networks, and heuristic systems such as RFC7050 (updated by RFC8880, but generally not correctly implemented in most networks). Recently, a better alternative has been adopted, based on Router Advertisement (RA) messages, according to RFC8781.

Likewise, the use of DNS64 (RFC6147) is generally recommended, as it avoids translations in NAT46, reduces the load on the DNS server, and improves connection establishment times. Whenever possible, although it is not a common problem if DNS64 is correctly configured, to avoid DNSSEC validation failures, operating systems should perform self-synthesis.

"Birds of a feather flock together"

IPv6-Only+IPv4aaS (e.g., 464XLAT) is a suitable transition mechanism for **access networks**. Networks connected using these mechanisms are typically unmanaged internally, generally mobile and residential, and even small and medium-sized corporate networks when they don't need to publish internal services on the Internet (or have few services and manual configuration is worthwhile).

IPv6-Mostly was designed **exclusively** for use **within managed corporate networks**, whose access network should be dual-stack. Applying it to residential (unmanaged) or mobile networks would imply:

- The access network must be maintained in a dual-stack configuration with public IPv4 addresses or incorporate CLAT into the CPE.
- Communication between IPv6-Mostly and IPv6-Only devices is not possible unless mechanisms such as PLAT, EAM, or SIIT-DC are manually configured, possibly on the CPE, in which case the network becomes managed. Alternatively, these mechanisms could be configured virtually by the ISP in the cloud or on the operator's infrastructure, considering the implications of many residential users potentially sharing the same private IPv4 addresses. This would force that traffic, instead of remaining local, to "upload" and "download" to the cloud or the operator's infrastructure, with the consequent latency and additional bandwidth consumption.

Residential CPEs often fail to incorporate CLAT, usually citing its complexity as an excuse, which is untrue. Incorporating option 108 in DHCPv4 is even more complex, requiring additional code, as well as the implementation of PLAT, EAM, or

SIIT-DC, and a user-friendly interface to facilitate the management to non-experts. Furthermore, in all cases, numerous open-source implementations exist, and the greatest cost actually lies in a deep understanding of these protocols, including IPv6 and transition mechanisms.

464XLAT vs IPv6-Mostly		
	464XLAT	**IPv6-Mostly**
Access Network	IPv6-Only	Dual-Stack
Mechanism Scope	Access Network	Network Segments
Managed	NO	YES
CPE Requirements	CLAT	Dual-Stack
ISP Network Requirements	PLAT DNS64 (recommended)	
Local Network Requirements		PLAT DHCPv6 Option 108
Local IPv4-IPv6 Communication	YES	Manually Configured: PLAT, EAM, SIIT-DC, ...
Mobile Networks Support	YES	NO
Advantages	Saving Public IPv4	Saving Private IPv4 Easy IPv6 error debugging Gradual IPv6-Only Discover Incompatibilities Reduced usage of NAT44
Mandatory RFCs	RFC6146/draft-ietf-v6ops-rfc6146-bis RFC6147 RFC6877 RFC7084/draft-ietf-v6ops-rfc7084bis	RFC6146/draft-ietf-v6ops-rfc6146-bis RFC6147 RFC8925 draft-ietf-v6ops-6mops RFC7757/STD103 RFC7755
Recommended RFCs	RFC8585 RFC8683 RFC8781 draft-ietf-v6ops-claton	RFC8683 RFC8781 draft-ietf-v6ops-claton

Note: The RFCs listed in the table are references in various parts of the network; their specific use may vary depending on the details of the deployment model.

Conclusion

The most important thing is user satisfaction, and this compels us, on one hand, to use the appropriate protocols for each case and, on the other hand, to call on CPE manufacturers to incorporate CLAT into their devices (and, for corporate use, the relevant IPv6-Mostly protocols). We also urge ISPs to demand this support. There are alternatives on the market, and we must choose not only based on price (often, "you get what you pay for"), but also what is best for users and our networks. Only the market itself (or explicit regulation) will force manufacturers to implement what we need. Let's not rule out highly competitive OEM options with OpenWrt, which is even suitable for avoiding replacing existing CPEs with custom firmware for our network.