PHAT MONKEY IT

SECURITY BRIEFING

Your company is probably running IPv6 — without knowing it

A plain-English briefing for business leaders on the network you didn't plan, the risks it quietly creates, and the audit that closes the gap.

Vishal Vashisht
July 2026

EXECUTIVE SUMMARY

IPv6 is already switched on inside most organisations – built into Windows, macOS, Linux and every modern phone, and often enabled automatically. Many companies believe they "don't use IPv6" when they simply haven't looked. The result is a second network running alongside the managed IPv4 one: the same laptops, the same switches, carrying the same data – but frequently unmonitored, unfiltered and unaudited. This briefing explains, in plain terms, why that gap exists, the specific risks it creates, and the practical audit that closes it – with no large migration project required.

01 The corridor you can't see

Most business leaders have never heard of IPv6. That is understandable – internet addressing feels like something for network engineers, sitting behind the Wi-Fi, the laptops and the cloud services people use every day. Yet IPv6 can create a serious gap between the network a company *believes* it is running and the network that actually exists.

Many organisations say they are not using IPv6. What they usually mean is that they have not planned an IPv6 project, bought equipment for it, or presented a migration plan to the board. But that does not mean IPv6 is absent. It is already built into Windows, macOS, Linux, iPhones and Android devices, and is often switched on automatically. A company can have IPv6 traffic moving around its offices even when nobody has formally approved it.

> *This is a little like discovering that your office has a second set of corridors behind the walls. The main corridors have security cameras and access controls. The hidden ones do not.*

02 What is IPv6?

Every device on a network needs an address, so data can find the right laptop, server or phone – just as a postal address lets a letter reach the right house. Most companies still think mainly in IPv4: addresses such as `192.168.1.25`. IPv4 has been used for decades, but the world eventually began to run out of available addresses.

IPv6 was created to solve that. Its addresses are much longer, and look more like `fe80::a123:45ff:fe67:890b`. Device makers have already added it to their products, so an employee's laptop may give itself an IPv6 address as soon as it joins a network. It does not need a company server to issue one, or an administrator to configure anything. The organisation might not have decided to use IPv6. The laptop already has.

03 Why should a CEO care?

A CEO does not need to understand network packets. The business issue is simpler: **you cannot protect something you do not know exists.** A company may have spent heavily protecting its IPv4 network – firewalls configured, security teams watching alerts, activity recorded for investigations. IPv6 traffic may not receive the same attention.

Some monitoring tools support it but have not been configured correctly. Some security rules recognise only the older IPv4 format. Flow logs may be collected for IPv4 and ignored for IPv6. An attacker does not need IPv6 to be better than IPv4 – only to be watched less carefully. In effect, the organisation may be monitoring the front door while leaving a side door unobserved.

04 The network may choose IPv6 automatically

Modern computers are designed to make connections quickly. When both IPv4 and IPv6 are available, a device tests the possible routes and uses whichever works best – all in the background. The user never sees a message saying "this website is now using IPv6". The browser simply opens the page.

That is useful for staff, because applications keep working. It is unhelpful for security teams, because people may not notice which type of connection was used. A company can therefore have IPv6 traffic without users, managers or even some technical teams being aware of it. The network is doing exactly what it was designed to do. The organisation has simply failed to look.

05 The false comfort of NAT

Many companies have long believed that a technology called Network Address Translation, or NAT, protects them. NAT lets many internal devices share a smaller number of public addresses; it became common because IPv4 addresses were scarce, and it gave businesses a clear dividing line between the internal network and the outside world. That felt like security.

The real protection normally came from the firewall sitting next to it. The firewall decided which connections were allowed; NAT merely changed the addresses. These are different jobs – changing an address does not make dangerous traffic safe.

Some companies try to make IPv6 look exactly like their old IPv4 network, adding extra address translation because it feels familiar. That can make the network harder to understand without improving security. A better approach uses proper firewall controls, sensible network separation and clear monitoring – the measures that actually provide protection. Familiar technology can feel safe simply because people recognise it. That does not make it the right answer.

06 When a laptop becomes a router

One of the clearest IPv6 risks begins with something called a Router Advertisement. The name sounds complicated; the idea is simple. A router is a device that tells computers where to send their traffic. In IPv6, routers can announce themselves to other devices on the local network – a computer hears the announcement and may decide to use that router as its way out.

The problem is that these announcements are not normally authenticated. A device can claim to be a router, and other computers may believe it. Sometimes this is deliberate: an attacker uses a fake announcement to move traffic through a device they control. More often it happens by accident – an employee switches on internet sharing from a laptop, a developer creates a badly configured virtual network, someone connects a small travel router to an office socket, or a virtual server host is set up incorrectly.

> *Suddenly, an ordinary device begins telling other computers: "send your traffic through me". Some computers may obey.*

This can cause unstable connections. It can also create an opportunity to observe or interfere with traffic. The business may experience strange outages while technical teams spend days looking in the wrong place.

07 The rogue router problem

Imagine an office with a proper road leading to the internet. Everyone normally follows the official road. Now imagine that someone puts up a new sign – "internet this way" – pointing through their laptop. Some drivers ignore it; others follow it.

That is roughly what can happen with a rogue Router Advertisement. The laptop does not have to break into the firewall – it is already inside the office network, speaking directly to nearby devices.

Network switches can help prevent this. A control known as RA Guard stops ordinary office ports from sending router announcements. In simple terms, the network can be told: "only the real routers are allowed to say they are routers". This is called first-hop security – it protects the first connection between a device and the network, and should not be treated as an optional extra.

08 Why blocking everything causes trouble

When technical teams discover unfamiliar network traffic, the natural response is often to block it. That can be sensible – but with IPv6, blocking all of a certain type of traffic can break the network in confusing ways. IPv6 uses a system called ICMPv6 for important messages: these help devices find routers, report errors, and let devices discover one another. They also tell a computer when the data it is sending is too large for part of the network.

Think of a delivery van reaching a bridge with a height restriction. The driver needs a warning: "your vehicle is too tall – use a smaller one". In IPv6, a message called *Packet Too Big* performs a similar job, telling the sender to use smaller packets. If a firewall blocks that warning, the connection may only partly work.

A website might begin to load and then stop. A login page may appear while the main application hangs. Small pieces of data get through but larger ones fail. This is difficult to diagnose because nothing appears completely broken: the application team blames the network, the network team sees that the connection was allowed, and support staff cannot reproduce the problem reliably. Everyone chases the wrong cause. The answer is not to allow every type of IPv6 traffic – it is to understand which messages are necessary and permit them safely.

09 Security teams may be missing the evidence

IPv6 can also make investigations harder. Modern devices may regularly change part of their IPv6 address to improve privacy, which makes it harder for websites to track the same device over long periods. That privacy feature has a reasonable purpose – but it also means a security team cannot always look at an IPv6 address and assume it belongs to the same person forever.

The company needs other records to connect network activity to a user or device: login records, switch information and network access control data. Many organisations do not collect enough of it. Their security tools may also contain rules designed only for IPv4 – an alert that looks for an old-style address may completely ignore a longer IPv6 one, and traffic records may not include IPv6 at all. This creates a visibility gap. A malicious person using IPv6 may appear less clearly in

the company's monitoring; their activity is not necessarily invisible, but it may be harder to detect and investigate. For an attacker, that may be enough.

10 This is not a huge migration project

Discovering this problem does not mean the company must move every system to IPv6 immediately. The first step is much smaller: find out what is already happening. Technical teams should check whether company devices have IPv6 addresses, look for unexpected router announcements, and identify which devices are acting as routers. They should review office networks, wireless services, development environments and temporary internet connections.

The company should then check whether its network equipment can stop ordinary devices from pretending to be routers. Firewall policies need to allow the IPv6 messages required for the network to function, and block those that appear where they do not belong. Security monitoring must include IPv6 traffic, alerting rules should recognise IPv6 addresses, and the organisation should be able to connect an address seen at a particular time to the device and person using it. New technology purchases should also include IPv6 requirements – a firewall that only provides excellent visibility for IPv4 is no longer providing complete visibility.

11 The questions leaders should ask

A board does not need a lesson in network engineering. It does need assurance. Leaders can ask a few straightforward questions:

- Do we know where IPv6 is active?
- Can an employee's device accidentally become a router?
- Would our security team see suspicious IPv6 traffic?
- Have our firewall rules been designed for IPv6?
- Can we connect an IPv6 address to a real device and user?

These do not require the CEO to understand the technical answer in detail. They require the organisation to *have* an answer. "We are not using IPv6" is not enough – it may only mean that nobody has checked.

12 The real risk is the network you forgot

IPv6 is not automatically dangerous. The danger comes from treating it as somebody else's future problem when it is already present in today's devices. A business may have a carefully managed IPv4 network and a poorly understood IPv6 network running beside it. Both use the same laptops. Both use the same switches. Both can carry company data. One is being watched. The other may not be.

A PHAT MONKEY PERSPECTIVE

The right response is not panic – it is a proper audit. Find the traffic. Check the controls. Fix the monitoring. Then make an informed decision about the future. As an independent, London-based consultancy with senior engineers across network security, cloud and on-premises estates, we run exactly this kind of IPv6 discovery and hardening exercise – mapping where IPv6 is live, closing first-hop and firewall gaps, and making sure your monitoring sees both networks, not just the one you planned.

Your company may not have chosen to deploy IPv6. That does not mean IPv6 waited for permission.