Traditional VPNs, while still vital for enterprise security, struggle to meet the demands of modern, latency-sensitive applications such as voice, video conferencing, and augmented and extended reality (AR/XR). VPNs reliance on tunnels (along with optional encryption) obscures traffic metadata, thus breaking network QoS mechanisms and can introduce significant additional latency due to the routing inefficiencies associated with typical hub-and-spoke VPN topologies. These constraints can significantly degrade the quality of real time communications and can adversely impact the performance of dynamic, cloud-native workloads. UM Labs' Unicus™ platform addresses these challenges by delivering low-latency, post-quantum-ready encryption without the use of tunnels, thus preserving QoS and retaining optimal routing. It offers a future-proof alternative to legacy VPNs for securing real-time communications without sacrificing performance or manageability.

Executive Summary

Site-to-site Virtual Private Networks (VPNs) have long been foundational to enterprise IT, offering secure communication over untrusted networks. Initially designed to emulate the security of private WANs using public infrastructure, VPNs have since evolved into core components of modern distributed networks.

However, as real-time and latency-sensitive applications (voice, video, AR/VR, telepresence) become mainstream, legacy VPN architectures and their overlay-based tunnelling models are proving to be a performance bottleneck. These limitations stem from the obfuscation of critical network metadata and suboptimal routing topologies—both of which undermine the effectiveness of Quality of Service (QoS) mechanisms and negatively impact user experience.

This paper examines the architectural and operational limitations of traditional VPNs in modern enterprise environments, particularly when supporting real-time applications, and highlights the challenges posed by static tunnel configurations, packet encryption overhead, and inefficient traffic flows.

The Role of Site-to-Site VPNs in Enterprise IT

Site-to-site Virtual Private Networks (VPNs) remain a fundamental part of enterprise IT architecture. They are widely used to establish a dedicated network topology over untrusted networks, particularly the public internet. This is achieved by building "tunnels" that encapsulate traffic between remote sites, users, and cloud environments, effectively simulating a private network. It is also possible to encrypt the data relayed in these tunnels (i.e. the whole of the original packet, including both data and headers) thus making the information relayed in the VPN tunnels unreadable to outsiders.

By ensuring confidentiality, integrity, and controlled access, VPNs support business continuity and security compliance across distributed organisations.

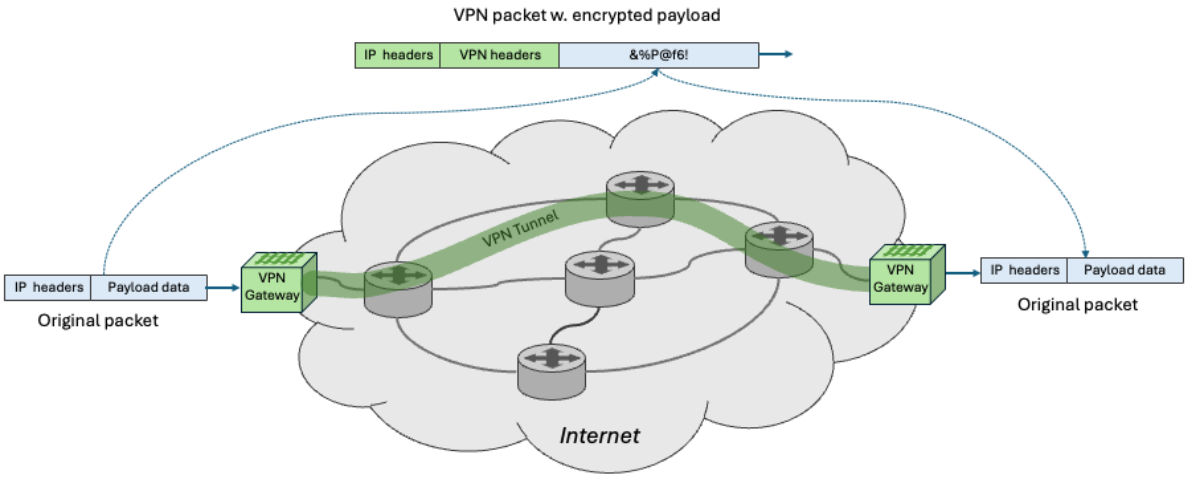

Figure 1 - An Overview of Site-to-site VPNs Architecture

Historical Context and Evolution

The concept of VPNs emerged in the mid-1990s as a cost-effective alternative to leased lines. Early protocols like Point-to-Point Tunnelling Protocol (PPTP) and IPsec enabled enterprises to securely connect branch offices and remote users without the cost and rigidity of dedicated WAN links.

Over time, VPNs have adapted to new use cases, extending beyond traditional applications such as email, file sharing, and terminal access. Today, they must support a diverse array of modern workloads—including real-time voice/video, containerized microservices, IoT, and hybrid cloud architectures. Technologies like OpenVPN and WireGuard have improved VPN performance and security, while software-defined approaches like SD-WAN have built on VPN tunnelling as a foundation for intelligent, policy-driven traffic routing.

Technical Constraints of Tunnelling and Encapsulation

Despite their evolution, VPNs still rely on tunnelling, which inherently encrypts not only payloads but also original protocol headers (IP, TCP/UDP). These headers carry essential metadata, including:

- Source/destination addresses
- Application-specific port numbers
- QoS-related markings such as Differentiated Services Code Points (DSCP)

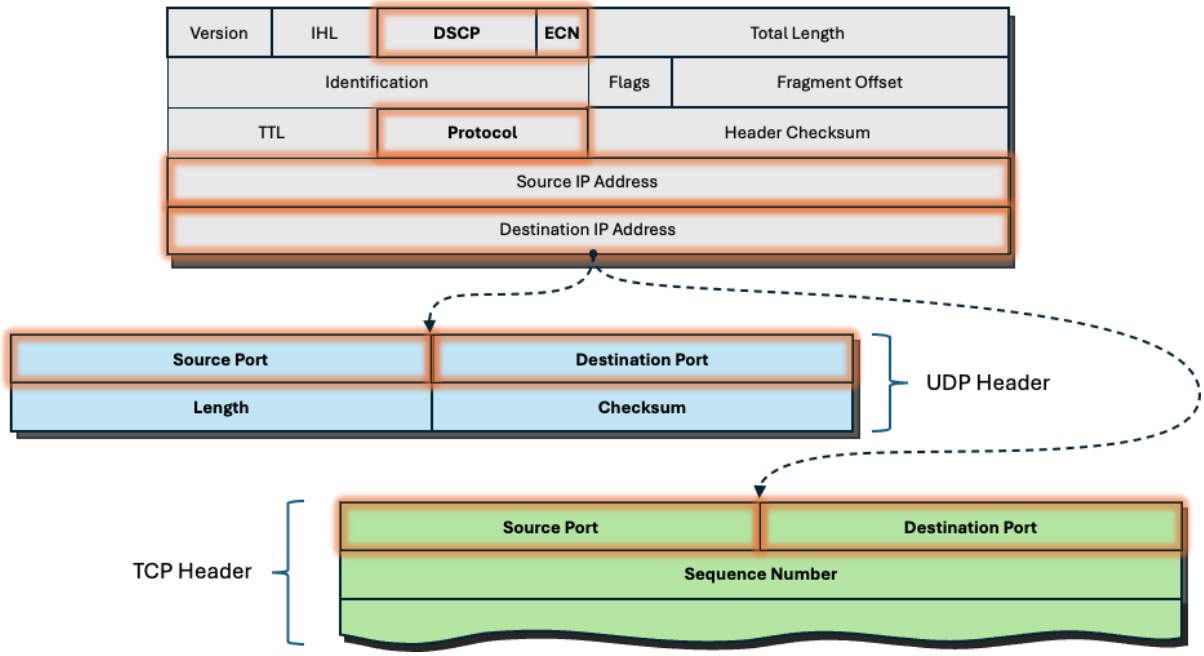

Figure 2 - Key metadata (highlighted) in IP, UDP & TCP headers used to determine traffic priority

This encryption effectively blinds intermediate network devices, such as routers and switches, from reading or acting on this metadata. As a result, enterprise QoS mechanisms that prioritize latency-sensitive traffic (e.g. voice, video, AR/VR) are rendered ineffective across the encrypted tunnel.

Consequently, all traffic is treated equally en route to the VPN gateway, which can lead to increased delay, jitter, and packet loss for critical real-time applications.

Real-Time Application Requirements

Modern applications are increasingly latency sensitive. For context:

- The **ITU-T** recommends a maximum one-way latency of **150ms** for acceptable voice quality.

- **Cisco Telepresence** adheres to the same threshold for video conferencing.

- **NVIDIA CloudXR** specifies latency between **60ms and 120ms** for extended reality (XR) workloads.

When traffic is forced through inefficient or overloaded VPN tunnels, without the benefit of QoS enforcement, these thresholds are easily breached, degrading the end-user experience.

Overlay Network Topologies and Routing Inefficiencies

Enterprise VPN deployments often rely on overlay networks formed by tunnels between gateways at headquarters, branch offices, data centres, or cloud services. While these topologies simplify routing and policy enforcement, they introduce rigidity.

A common approach is a **hub-and-spoke** model, where all branch traffic is routed through a central location (e.g., a New York HQ). In such scenarios, traffic between two European offices (e.g., London to Paris) must traverse the US, adding **140ms–180ms of unnecessary latency**. This suboptimal routing undermines application performance, especially for real-time collaboration tools and immersive experiences.

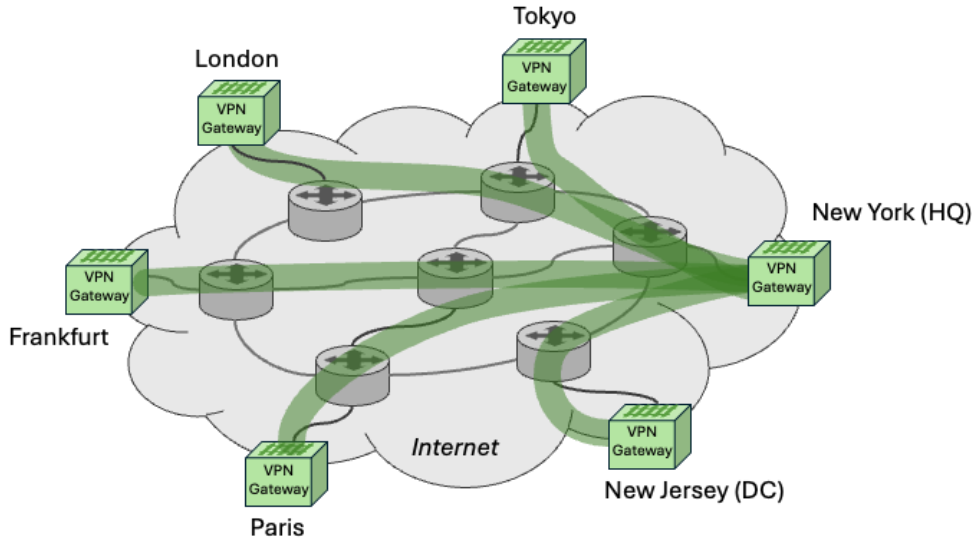

Figure 3 - Typical "Hub & Spoke" site-to-site VPN overlay topology

Although split tunnelling and Policy-Based Routing (PBR) are sometimes used to mitigate this issue, these approaches require complex Access Control List (ACL) based configurations and are poorly suited for dynamic environments, such as containerized workloads or ephemeral cloud resources, where endpoints change frequently and are location-agnostic.

Conclusion

VPNs remain essential for secure connectivity in modern enterprises, but their legacy architecture poses challenges for latency-sensitive and dynamic workloads:

- **Encrypted tunnels obfuscate key protocol metadata**, making QoS enforcement impossible along the path.

- **Overlay topologies introduce unnecessary latency** due to rigid, centralized routing models.

- **Static configurations and complex routing policies** are incompatible with cloud-native and hybrid workloads that demand agility and automation.

As enterprises increasingly depend on real-time applications and dynamic infrastructure, a reassessment of the role of VPNs is needed and alternative mechanisms need to be implemented to secure real-time communications. Any alternative solution needs to ensure all real-time traffic can benefit from network-based quality of service and optimal routing whilst delivering security equivalent or better than the VPNs solution they are replacing.

Fortunately, a practical solution to these challenges exists today in the form of **UM Labs' Unicus® platform**. Unicus™ is specifically engineered to address the limitations of traditional VPN architectures in real-time communication environments. It delivers **low-latency, high-assurance encryption**, including support for **Post-Quantum Cryptography (PQC)**, across a broad range of real-time protocols, all while preserving critical **QoS markings** to maintain traffic prioritization.

In addition to its encryption capabilities, Unicus™ offers **native NAT and proxy services**, which significantly reduce the complexity typically associated with **Policy-Based Routing (PBR)** and **split tunnelling**. By minimizing the number of IP addresses requiring manual classification for non-VPN routing, Unicus™ streamlines network operations and reduces administrative overhead.

In essence, Unicus® offers a secure, scalable, and future-proof alternative to legacy VPN solutions—optimizing performance without compromising security, control or visibility.

Gareth Flook- Lead at UM-Labs R&D