Securing OT Protocols: Addressing Vulnerabilities with Host Identity Protocol

This paper examines the inherent cybersecurity vulnerabilities in common Operational Technology (OT) protocols when transmitting over public cloud infrastructure and presents Host Identity Protocol (HIP) as a more secure alternative to traditional VPN solutions. OT protocols such as MQTT, BACnet, and Modbus were designed primarily for efficiency and interoperability within isolated networks, with security considerations often secondary. As industrial systems increasingly connect to cloud services, these protocols' security limitations create significant risks. This paper demonstrates how HIP's cryptographic identity approach addresses these vulnerabilities more effectively than traditional VPN solutions, providing consulting engineers with a comprehensive understanding of both the security challenges and potential solutions.

The Vulnerability Landscape of OT Protocols

Industrial and building automation systems have historically relied on protocols designed for isolated environments. As these systems connect to broader networks and cloud infrastructure, their inherent security weaknesses become critical vulnerabilities.

MQTT Security Vulnerabilities

MQTT (Message Queuing Telemetry Transport) has gained widespread adoption for IoT applications due to its lightweight design and publish-subscribe model. However, this protocol exhibits significant security vulnerabilities when exposed to public networks.

Recent research revealed approximately 49,000 exposed MQTT servers accessible from the Internet, with over 32,000 having no password protection whatsoever[1]. This widespread misconfiguration creates an enormous attack surface. Even when encryption is implemented, MQTT suffers from architectural weaknesses that undermine security.

The protocols publish-subscribe model inherently exposes topic structures that reveal significant information about system architecture and operations. Attackers can exploit MQTT's wildcard subscription feature (using "#") to listen for all messages being passed through the broker, gaining complete visibility of all communications[1]. This information exposure occurs even when the payload itself is encrypted, as the topic structure remains visible.

Furthermore, MQTT exposes broker addresses when traversing networks, creating entry points for attackers. This visibility issue persists regardless of whether TLS encryption is implemented, as the protocol design fundamentally requires visible broker addressing[2][3]. When attackers gain access to an MQTT broker, they can not only subscribe to sensitive information but also publish false messages to topics, potentially manipulating building systems, security features, or industrial processes.

BACnet Security Exposures

BACnet (Building Automation and Control Networks) presents even more severe security challenges. This protocol, ubiquitous in building management systems, was designed without fundamental security mechanisms and becomes extremely vulnerable when connected to public networks.

BACnet lacks authentication, encryption, and authorisation controls at the protocol level[4]. Security researchers at Dragos concluded that "freely available tools can connect to Internet-connected building systems which use this protocol and modify their operation without the need for an exploit"[5]. This assessment highlights that attackers don't need sophisticated techniques to compromise BACnet systems-the protocol itself provides the access.

The security deficiencies in BACnet include:

1. Complete lack of authentication, allowing device spoofing and unauthorised commands[4]
2. No encryption, with all data transmitted in plaintext[6]
3. Absence of authorisation controls, permitting any connected device to issue commands[6]
4. Vulnerability to denial-of-service attacks through network flooding[4]
5. Susceptibility to network connection disruption via routing table corruption[4]

These vulnerabilities become particularly dangerous when BACnet systems connect to cloud services or allow remote access, as the protocol offers no inherent protection against unauthorised manipulation[5][4].

Modbus Protocol Weaknesses

Modbus, one of the oldest and most widely deployed industrial control protocols, presents similar security concerns. Originally designed for serial communications in isolated environments, Modbus TCP/IP adaptations have carried forward fundamental security deficiencies into networked environments.

The most critical Modbus vulnerabilities include:

1. Cleartext data transmission with no encryption[7][8]
2. Complete absence of authentication mechanisms[7][8]
3. No integrity checks to verify data hasn't been altered in transit[7][8]
4. Vulnerability to command injection and replay attacks[7]
5. Simplistic framing and lack of session structure[8]

These vulnerabilities create significant opportunities for attackers. As documented by Red Bot Security, "When the lack of encryption is combined with non-existent authentication and session handling, opportunities to intercept and modify messages between the client and server become available"[9]. An attacker with network access can perform man-in-the-middle attacks to manipulate industrial processes while reporting normal operations to monitoring systems-a scenario reminiscent of the infamous Stuxnet attack[9].

Traditional Security Approach: Virtual Private Networks

Organisations have traditionally relied on Virtual Private Networks (VPNs) to secure OT communications traversing public networks. While VPNs provide a basic security layer, they introduce significant limitations and vulnerabilities when applied to OT environments.

VPN Security Model and Implementation

VPNs operate on a perimeter-based security model, creating encrypted tunnels between network endpoints at the network layer (Layer 3). They provide encryption and tunnelling capabilities that can protect data in transit from eavesdropping and basic tampering.

The primary security features of VPNs include:

1. Data encryption between endpoints
2. Authentication at connection establishment
3. Tunnelling of traffic through untrusted networks
4. IP address concealment from the public internet

These capabilities address some basic security concerns when transmitting OT protocols over public infrastructure, but they fail to address the fundamental vulnerabilities inherent in the protocols themselves.

Limitations of VPNs for OT Security

VPNs create several significant security challenges when applied to OT environments, potentially introducing more risk than they mitigate.

The most critical limitation is that VPNs typically create flat networks with broad access after authentication[10]. Once a user or device authenticates to a VPN, they often gain broad network access rather than being limited to specific resources. This access model contradicts the principle of least privilege fundamental to secure system design.

In industrial environments, this flat network model means VPNs defeat IT/OT air-gaps through bidirectional tunnels[10]. As Zscaler notes, "By design, VPNs create bidirectional tunnels between two networks, but inbound traffic flows are the source of all things bad"[10]. This bidirectional access creates pathways for attackers to move laterally between IT and OT systems once they've compromised a VPN connection.

VPNs also create significant new attack surfaces. Dragos research indicates that "adversaries use IT systems as an entry point to gain visibility and access to OT environments" and identified widespread brute-force attacks targeting VPN appliances across critical infrastructure sectors[11]. Once compromised, these VPNs become the gateway to industrial systems.

The operational challenges of VPNs also undermine security. They require re-authentication when changing networks, resulting in reconnection times of 31-100 seconds. This authentication model makes VPNs ill-suited for mobile or unreliable network environments common in industrial settings. Configuration complexity, especially for IPsec implementations, often leads to security misconfigurations that create additional vulnerabilities.

Performance degradation represents another significant limitation, with VPNs reducing throughput by up to 90% in some implementations. This performance impact can disrupt time-sensitive industrial processes and monitoring systems.

Host Identity Protocol: A Zero-Trust Alternative

Host Identity Protocol (HIP) offers a fundamentally different approach to securing OT communications that addresses many of the limitations of traditional VPN solutions.

HIP Fundamentals and Architecture

HIP is an IETF-approved security protocol that operates between the network and transport layers (Layer 3.5)[12][13]. Unlike traditional TCP/IP networking, which uses IP addresses as both locators and identifiers, HIP separates these functions by introducing cryptographic host identities[12][13].

This architectural difference is fundamental: rather than identifying hosts by their network locations (IP addresses), HIP establishes cryptographic identities independent of network topology[13]. As explained in RFC 9063, "The Host Identity is referred to by its public component, the public key. Thus, the name representing a Host Identity in the Host Identity namespace, i.e., the Host Identifier, is the public key"[13].

HIP establishes secure communications through a four-way handshake called the Base Exchange, which authenticates both parties using their cryptographic identities and establishes encrypted communications[14][15]. This exchange includes computational puzzles that protect against denial-of-service attacks by requiring initiators to demonstrate computational commitment before a responder allocates significant resources[15].

HIP Security Benefits for OT Systems

HIP implements a true zero-trust security model that fundamentally transforms how OT networks are secured. Rather than creating a trusted perimeter, HIP verifies every communication attempt based on cryptographic identity[16][17].

The key security benefits include:

1. Complete invisibility of protected devices to unauthorised users[16]
2. Significant reduction in attack surface through cryptographic cloaking[16]
3. Continuous verification rather than single-point authentication[17]
4. Micro segmentation that limits access to specific resources rather than entire networks[17]
5. Protection against address spoofing through cryptographic verification[13]

These capabilities address the fundamental vulnerability of OT protocols: their lack of built-in security. As Tempered Networks explains, "Rather than probing an IP device for vulnerabilities, attackers will never know of the existence of the host on the network. To them, the device is completely cloaked"[16].

Securing Specific OT Protocols with HIP

HIP addresses the specific vulnerabilities of common OT protocols through its cryptographic identity model.

For MQTT, HIP hides broker addresses and secures topics through cryptographic identities. This approach prevents unauthorised discovery of MQTT brokers and ensures that only authenticated devices can subscribe to topics or publish messages. The protocol weaknesses remain, but they're contained within a secure overlay network only accessible to authenticated devices.

For BACnet, HIP creates a secure overlay network that prevents unauthorised access. While BACnet still lacks authentication and encryption internally, the HIP layer ensures that only verified devices can participate in BACnet communications, effectively adding the missing security layer.

For Modbus, HIP prevents address spoofing and secures the communication channel. This protection mitigates the risk of man-in-the-middle attacks and unauthorised commands that could manipulate industrial processes. The cryptographic verification ensures that only authorised devices can send Modbus commands, even though the protocol itself lacks authentication.

Performance Comparison: VPN vs HIP

While security benefits represent the primary consideration for protecting critical OT systems, performance characteristics significantly impact implementation decisions. Both VPN and HIP solutions introduce overhead that affects network performance.

Latency and Throughput Impact

Performance testing reveals significant differences between VPN and HIP implementations in terms of network impact.

VPN solutions, particularly IPsec, demonstrate high impact on throughput, reducing it by 70-90% in testing scenarios[18]. This dramatic reduction occurs primarily due to the intensive cryptographic operations and encapsulation overhead. OpenVPN typically shows even greater performance impacts than IPsec[18].

In contrast, HIP introduces a more moderate impact on network performance. Measurements indicate a 35-45% latency increase compared to plain IP traffic[19]. While still significant, this performance impact is substantially lower than that of most VPN implementations and generally remains acceptable for most industrial applications.

Connection establishment times also differ significantly. IPsec VPNs typically require 0.7-2.0 seconds for connection initialisation, with reconnection times of 31-100 seconds when network changes occur. HIP achieves faster connection initialisation, requiring approximately 200-300ms with 768-bit Diffie-Hellman key exchange[19]. More importantly, HIP provides seamless reconnection during network transitions, maintaining communications without requiring re-authentication[19].

Resource Utilisation and System Requirements

Resource utilisation presents another significant consideration, particularly for industrial environments with resource-constrained devices.

VPNs, especially IPsec and OpenVPN, demonstrate high CPU utilisation in testing scenarios[18]. This resource consumption can impact the performance of industrial controllers and edge devices that may already operate near their computational limits.

HIP demonstrates more moderate resource requirements, with testing showing acceptable performance even on lightweight hardware such as mobile devices[19]. Measurements on a Nokia 770 Internet Tablet showed that HIP could achieve reasonable performance by adjusting cryptographic parameters, such as using 768-bit Diffie-Hellman groups to balance security and performance[19].

For resource-constrained environments, HIP offers greater flexibility in balancing security and performance. As one study notes, "Using different DH Groups makes it possible to affect the generation time of the DH session key and as a result the total duration of the HIP base exchange. In reality, this means an opportunity for a server to offer smaller DH public values to lightweight clients that are not powerful enough or if the security is not of critical importance"[19].

Implementation and Operational Considerations

Operational aspects of security solutions significantly impact their effectiveness and sustainability in industrial environments.

VPNs typically require complex configuration, especially for IPsec implementations. This complexity increases the likelihood of security misconfigurations and creates ongoing management challenges. VPNs also often require significant infrastructure changes to implement properly, particularly for site-to-site connections.

HIP requires infrastructure updates but generally presents simpler configuration than IPsec VPN. While adoption requires changes to networking architecture, the centralised management of cryptographic identities can simplify ongoing operations compared to managing complex VPN access controls.

Mobility support represents a significant operational advantage for HIP. The protocol maintains the same cryptographic identity across networks, providing seamless mobility support[20]. This capability proves particularly valuable in industrial environments where devices may connect through multiple network paths or change connectivity frequently.

Comparative Analysis: VPN vs HIP for OT Security

Direct comparison of VPN and HIP approaches reveals significant differences in how they address OT protocol vulnerabilities and overall security posture.

Security Model Comparison

The fundamental security models of VPN and HIP differ significantly, with direct implications for OT security.

VPNs implement a perimeter-based security model that creates a trusted zone after authentication[17]. This model assumes that authenticated users and devices should have broad network access, conflicting with zero-trust principles that have become essential in modern cybersecurity.

HIP implements a true zero-trust security model based on cryptographic identities[16][17]. Each communication attempt requires verification, and access is granted only to specific resources rather than entire network segments. This approach aligns with current security best practices that assume networks are already compromised and require continuous verification.

Research comparing HIP with mobile VPNs found that HIP-based solutions provide better security by 3.58% and better mobility by 5.26%[21]. While these percentage improvements may seem modest, they represent significant enhancements in critical security capabilities.

Access Control and Network Visibility

The approaches to access control and network visibility fundamentally differ between VPN and HIP implementations.

VPNs typically create flat networks with broad access once authentication occurs[10][17]. This approach allows lateral movement within the protected network and potentially enables attackers to access critical systems once they breach the VPN. Network resources remain visible to all authenticated users, creating unnecessarily broad attack surfaces.

HIP implements micro segmentation and least privilege access, granting users and devices access only to specific resources they require[17]. This granular access control prevents lateral movement and contains potential breaches. Additionally, devices protected by HIP remain completely invisible to unauthorised users, preventing reconnaissance and dramatically reducing the attack surface[16].

Protocol-Specific Protection

The effectiveness of VPN and HIP solutions in addressing specific OT protocol vulnerabilities differs significantly.

VPNs provide basic encryption but don't address protocol-specific vulnerabilities. While they create encrypted tunnels for MQTT, BACnet, and Modbus traffic, they don't mitigate the fundamental security weaknesses in these protocols. Once an attacker gains VPN access, they can exploit these vulnerabilities just as if they had direct network access.

HIP addresses fundamental protocol flaws by hiding vulnerable systems from unauthorised access. For MQTT, it hides broker addresses and secures topics through cryptographic identities. For BACnet, it creates secure overlay networks that prevent unauthorised access despite the protocol's lack of authentication. For Modbus, it prevents address spoofing and secures communication channels, mitigating the risk of unauthorised commands.

Conclusion and Recommendations

The security challenges facing OT protocols in cloud-connected environments require robust solutions that address fundamental vulnerabilities. While VPNs have traditionally been used to secure these communications, their perimeter-based security model introduces significant limitations and risks.

Host Identity Protocol offers a more comprehensive security approach for protecting vulnerable OT protocols when transmitted over public cloud infrastructure. Its zero-trust architecture with cryptographic identities provides stronger protection against modern threats while offering operational benefits like seamless mobility and simplified management.

For consulting engineers designing secure OT networks that must traverse public infrastructure, we recommend:

1. Implementing HIP-based solutions for critical OT systems that require connectivity across public networks or cloud infrastructure
2. Considering the performance implications of security solutions, with HIP generally offering better balance between security and performance than VPN alternatives
3. Adopting zero-trust principles regardless of the specific technology implemented, ensuring least privilege access and continuous verification
4. Maintaining proper network segmentation between IT and OT environments, even when using secure communication protocols
5. Regularly assessing security posture against evolving threats, recognising that no single technology provides complete protection

The evolution of industrial systems toward greater connectivity necessitates a corresponding evolution in security approaches. Host Identity Protocol represents a significant advancement that addresses the fundamental security challenges of traditional OT protocols while providing the flexibility and performance needed for modern industrial applications.

Future Considerations

As industrial systems continue to evolve, security solutions must adapt to address emerging challenges. Future developments in HIP implementations may further improve performance on resource-constrained devices and enhance integration with cloud platforms. Organisations should maintain awareness of these developments and regularly reassess their security architecture to ensure it continues to provide adequate protection for critical industrial systems.

The transition to more secure architectures requires both technical implementation and organisational change. Consulting engineers play a crucial role in educating clients about security risks and guiding them toward appropriate solutions. By understanding both the vulnerabilities of OT protocols and the capabilities of modern security approaches like HIP, engineers can help organisations build truly secure industrial systems that reliably operate across modern network infrastructures.

References

1. https://blog.paessler.com/why-mqtt-is-everywhere-and-the-security-issues-it-faces
2. https://sen.news/mqtt-cyber-security-risks/
3. https://www.txone.com/blog/mqtt-series-2-potential-risks-of-exposed-mqtt-brokers/
4. https://www.infosecinstitute.com/resources/scada-ics-security/bacnet/
5. https://www.dragos.com/blog/industry-news/assessing-the-bacnet-control-system-vulnerability/
6. https://www.veridify.com/bacnet-security-issues-and-how-to-mitigate-cyber-risks/
7. https://www.veridify.com/modbus-security-issues-and-how-to-mitigate-cyber-risks/
8. http://dione.lib.unipi.gr/xmlui/bitstream/handle/unipi/11394/Evangeliou_1508.pdf?sequence=1&isAllowed=y
9. https://redbotsecurity.com/examining-the-modbus-protocol/
10. https://www.sscaler.com/blogs/product-insights/vpns-biggest-threat-your-industrial-control-system
11. https://www.dragos.com/blog/why-adversaries-target-vpn-appliances-the-pathway-from-it-to-ot-cyber-attack/
12. https://www.ida.liu.se/~TDDE21/info/primer-host-identity-protocol-whitepaper.pdf
13. https://www.rfc-editor.org/rfc/rfc9063.pdf
14. https://www.linuxjournal.com/content/experimenting-python-implementation-host-identity-protocol
15. https://github.com/nihalpasham/rustdhipv2
16. https://www.tempered.io/airwall/host-identity-protocol/
17. https://www.veridify.com/zero-trust-vs-remote-access-vpn-for-building-control-systems/
18. https://core.ac.uk/download/pdf/322886318.pdf
19. https://conferences.sigcomm.org/sigcomm/2007/mobiarch/Khurri_HIP.pdf
20. http://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1870069
21. https://oulurepo.oulu.fi/bitstream/handle/10024/23168/nbnfi-fe2018073133162.pdf?sequence=1&isAllowed=y