ARQIT

White paper

Post-quantum security for the **next-generation battlespace**

How today's militaries can build and maintain their information advantage in the face of Cryptographically Relevant Quantum Computers (CRQCs)

Contents

Information in warfighting	3
Network-centric warfare	4
Warfare in the cyber domain	4
Designing tomorrow's military networks	5
Why militaries are investing in Zero Trust	6
The quantum threat to encryption	7
Current military cryptography isn't future-proof	8
Cryptographic must-haves for modern militaries	9
Post-quantum security for military applications	10
Testing our quantum-safe MACP architecture	11
Functional testing results	11
Performance testing results	11
Quantum-safe, CSfC-compliant, and Zero Trust-ready	12
Quantum-safe solutions with full sovereignty	13
Appendix A: Hardware, software & network configuration for quantum-safe MACP testing	13

Information in warfighting

The US Department of Defense (DoD) expects to spend more than $15 billion on cyberspace activities in 2026, aiming to protect its information advantage and disrupt those of its adversaries.

Why? Because information is crucial in warfighting.

It was true when Sun Tzu wrote The Art of War in 5th-century China. It was true when Miyomoto Musashi wrote his Book of Five Rings in 17th-century Japan. And it's true today.

Modern militaries need current and accurate information in the right place and at the right time to aid decision-making. The ability to collect, share, and use information to make real-time decisions is the true differentiator in many military engagements, particularly when hardware and manpower are evenly matched.

Military forces of the US, UK, and EU nations must constantly monitor the capabilities of adversarial nations including Russia, China, North Korea, and Iran. These nations invest heavily in defense, and it's unlikely that a lasting advantage will be found in the development of new hardware. Instead, today's militaries aim to secure an advantage in the information domain.

This means the ability to:

- Collect more **high quality information** than adversaries
- **Analyze information** more quickly and effectively
- Share information instantaneously and safely across remote operations
- **Use information in real-time** to inform decisions

This information advantage speeds up the ability to Observe, Orient, Decide, and Act (the "OODA loop") on the battlefield. The shorter the OODA loop, the more quickly a military unit can act. All of this in the pursuit of outmaneuvering adversaries before a shot is fired, or better yet, avoiding conflict altogether by demonstrating information dominance.

"Military services use network-centric warfare and information dominance to create a tight OODA loop to enable real time decision-making, much faster than the opposing forces," *says Sean Carnew, Senior Director: Government, Defence & CNI for Arqit.* **"There's always going to be uncertainty on the battlefield, the fog of war. When you have superior situational awareness, you have an advantage. It speeds up your decision-making and helps you understand the truth of what's happening around you."**

But how can militaries protect and maintain an information advantage in the face of ever-evolving cyber threats and network-centric warfare?

This paper highlights the information, network, and security issues facing today's militaries, and outlines a new approach to secure, quantum-safe networking that's suitable for military applications.

Key takeaways

- Cyber warfare strategies are now common and pose a threat to all military networks.
- US and allied militaries are committed to adopting Zero Trust and future-proof encryption.
- Cryptographically relevant quantum computers (CRQCs) will break asymmetric encryption and PKI within 3-5 years.
- Today's militaries need a quantum-safe encryption solution that's software-based, compatible with Zero Trust principles, usable across untrusted networks, and hyper-scalable.
- Arqit, ECS, and Intel have developed a quantum-safe architecture that fulfills these needs.
- Based on the US National Security Agency (NSA) Mobile Access Capability Package (MACP), the architecture addresses the CRQC threat and the logistics issue at the heart of symmetric key management.

Military services **use network-centric warfare** and **information dominance to create a tight OODA loop** to enable **real time decision-making,** much faster than the opposing forces...

Network-centric warfare

Today's militaries handle an unprecedented amount of data, information, and intelligence. For context, a single F-35 fighter jet can produce up to 24 terabytes of data per flight hour.

Militaries collect, analyze, and operate on this data via:

- Sensors, wearables, devices, weapon systems, and vehicles — everything from tiny cameras to Triton unmanned aerial vehicles (UAVs) — to acquire, share, and receive information.
- Cloud and edge computing to store, send, and receive information, and to enable more effective, efficient, and precise warfighting.
- AI and Machine Learning (ML) tools for rapid analysis of huge datasets (big data).

Still, this only produces an information advantage if militaries have continuous access to fast and secure networking capabilities that function across vast and disparate regions. Again, information is only an advantage if it's available at the right time and in the right place; warfare is rarely fought on home turf or in an ideal environment. This is where modern military networks come in.

In addition to traditional network infrastructure, the US and its allies make frequent use of satellite-based networking and radio-based Mobile Ad-hoc Networking (MANET) meshes to maintain data connections around the world. These networks extend beyond front-end warfighting to improve support operations such as personnel management, communications, training, simulation, logistics, maintenance, and medical services.

However, the problem of sharing information extends beyond connectivity. Having an information advantage requires that adversaries don't have the same information, or at a minimum, they don't know you have it.

Basic information requirements for military networks

- Supports diverse devices and endpoints that collect, share, and receive information
- Widespread networking to share information across remote regions
- High-speed information analysis and sharing
- High security and total confidentiality

Herein lies the problem: making information available immediately, wherever it's needed, is hard. Doing so while protecting against theft and disruption by advanced adversaries is even harder.

Done well, network-centric warfare increases combat effectiveness through improved situational awareness, reduced sensor-to-shooter time, enhanced force effectiveness, efficient decentralized operations, and a host of other benefits. The DoD recognizes this value as well as the challenge of operationalizing network-centric warfare both effectively and securely.

"Operational success in the cyberspace domain demands speed, agility and unity of effort," *said U.S. Army Gen. Paul Nakasone (Ret.), former commander of U.S. Cyber Command*, in testimony before the House Armed Services Committee, March 2023. **"Defending the nation is paramount among our missions. It means defending our military systems, networks and the critical infrastructure that enable national security."**

Warfare in the cyber domain

Attacking military communications infrastructure and attempting to intercept messages has a long and storied history. From **intercepted telegrams** during the American Civil War to **cracking the Enigma code**, compromising communications has long been part of the warfighting playbook.

More recently, Russia launched widespread cyberattacks ahead of its land invasion of Ukraine - including disrupting the ViaSat satellite network - to hamper Ukrainian military capabilities. These attacks reflect what we've already covered: disrupting an adversary's flow of information can make it difficult for them to coordinate operations.

Today, the most prevalent examples of these attacks occur in the cyber domain. Network-based attacks take many forms, such as espionage, data breaches, malware, and denial of service attacks. They can also include electronic warfare attacks like jamming or spoofing network signals, and even physical attacks on network hardware including cables, radios, and satellites.

Adversary nation-states such as Russia, China, Iran, and North Korea have made huge investments in their cyber warfare capabilities. There is also evidence to suggest cooperation between these adversary nations, particularly when it comes to sharing intelligence. The military must also be prepared for attacks from non-state actors such as hacktivists and organized criminal groups.

"We are at a defining time for our Nation and our military. Near-peer competitors are posturing themselves, and threats to the United States' global advantage are growing. Nowhere is this challenge more manifest than in cyberspace."

— Gen. Paul Nakasone, 3rd Commander of United States Cyber Command (ret)

Put simply, cyber warfare has become an entire field of warfighting in its own right, and all mature militaries are investing in both offensive and defensive cyber capabilities. Case in point: all NATO countries now have cyber commands with dedicated budgets, with the US DoD's cyber budget skyrocketing more than 5X from $2.8 billion in 2012 to a requested $15.1 billion for FY 2026.

In response, DoD published its Joint All-Domain Command and Control Strategy (JADC2) in March 2022. The strategy notes:

> **Rapid changes in the global** security environment are presenting significant **new challenges to the U.S. military and the ability of the Joint Force to seize, maintain, and protect our information and decision advantage over our adversaries. [...]** We must anticipate that future military operations will be conducted in degraded and contested electromagnetic spectrum environments. **These challenges require a coherent and focused Departmental** effort to modernize how we develop, **implement, and manage our [command & control] capabilities to prevail in all operational domains, across echelons, and with our mission partners.**

Similarly, the UK's Ministry of Defence (MoD) has issued its Cyber Resilience Strategy for Defence where it lays out seven strategic priorities, along with its stated aim: "...for Defence's critical functions to be significantly hardened to cyber attack by 2026, with all Defence organisations resilient to known vulnerabilities and attack methods no later than 2030."

Designing tomorrow's military networks

The rapidly changing technology landscape poses significant challenges for network design. Modern militaries must deploy networks that can dynamically adapt to changing tactics and strategies, and changes are occurring faster than ever in response to hardware and software advancements. Forces that lock themselves into stovepiped designs risk being unable to adapt quickly and will likely find themselves at a disadvantage.

"If I make a decision based on what I'm seeing faster than your ability to respond, I'm constantly ahead of you," Carnew says. "If I design my networks with the ability to dynamically change, that's how I beat you in the battlefield. Not by just getting bigger, stronger, and going head-to-head with you. By being faster than you, more dynamic, and able to adapt."

Military networks are also increasingly adopting network architectures that bring critical computing capabilities to the edge, meaning data processing and storage are moved closer to the data's source.

This allows forces to make more effective use of different networking capabilities such as MANET, satellite-based networking, and mobile technologies while using edge computing in rugged and denied environments. If an adversary is jamming or otherwise attacking one system, users can switch to another system dynamically, ensuring lines of communication remain open and forces maintain their information advantage.

Beyond this, what exactly should military networks look like? If you read through the strategy documents released by the DoD and MoD, there are plenty of descriptive statements detailing the outcomes of an effective and secure military network.

They include phrases such as:

- Layered defence
- Interoperability
- Scalability
- Future proofed
- Security best practices
- Cost management
- Common data practices
- Conformant IT standards

While light on detail, this makes sense: militaries should be flexible while predominantly using standardized and accepted practices and standards, because not doing so would make it difficult to ensure consistency and interact seamlessly with allies.

Beyond this, there are two additional requirements that practically every modern military aims to build into its networks:

1. Zero Trust and Zero Trust Network Architecture (ZTNA)
2. Some form of standardized and future-proofed encryption methodology

While traditional cybersecurity measures like firewalls won't go away anytime soon, today's militaries recognize that more fundamental strategies are required to ensure granular control over where specific data is transmitted and who can access it, without slowing down the flow of information.

Why militaries are investing in Zero Trust

How can modern militaries share massive amounts of data in real time, all with different levels of authorization, while protecting against the myriad ways bad actors may try to intercept or interrupt communications? Zero Trust is an attempt to answer this question.

> **Zero Trust is an information security model that denies access to applications and data by default. Threat prevention is achieved by only granting access to networks and workloads utilizing policy informed by continuous, contextual, risk-based verification across users and their associated devices. Zero Trust advocates these three core principles: All entities are untrusted by default; least privilege access is enforced; and comprehensive security monitoring is implemented.**
>
> *— Forrester,*
> *The Definition of Modern Zero Trust (emphasis ours)*

In 2020, NIST published SP 800-207, Zero Trust Architecture, which fleshed out the original concept with seven "tenets":

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7. The enterprise collects as much information as possible on the current state of assets, network infrastructure, and communications, and uses it to improve its security posture.

The above definition and tenets describe the "philosophy" of Zero Trust.

In practice, actual implementation usually combines a variety of security and network capabilities, including:

- **Multi-Factor Authentication (MFA)**
- **Network segmentation**
- **Advanced encryption**
- **Endpoint security**
- **Identity & Access Management (IAM)**
- **Network analytics**
- **Data Loss Prevention (DLP)**
- **Networking monitoring**
- **Security monitoring**

It's easy to see why Zero Trust is attractive to modern militaries. The concept requires least privilege access, along with continuous reauthentication to access resources. This is a clear improvement over traditional perimeter security, where users and devices would authenticate just once and then be granted access until the end of their session.

Modern militaries are continuously targeted with sophisticated cyberattacks, but it's much more difficult to "dwell" within a target network that requires continuous reauthentication, particularly if authentication requires unbroken knowledge of changing credentials.

Implemented correctly, this perfectly meets a modern military organization's need for data security and confidentiality. This is why both DoD and MoD have committed to implementing Zero Trust:

- In 2022, DoD published its "**DoD Zero Trust Strategy**," which lays out a plan to implement zero trust capabilities and activities across the department by 2027.
- Also in 2022, MoD published its "**Cyber Resilience Strategy for Defence**", committing to implementing "a zero-trust architecture [...] based on a data-centric security model".

The quantum threat to encryption

Zero Trust and ZTNA aim to protect military networks. But what about the data those networks transmit? This also requires protection, and that's where encryption comes into play.

Secure encryption of data in transit is fundamental to a military's ability to share information between assets in a manner that can't be intercepted by adversaries. Until now, asymmetric key encryption has proven the most scalable and usable form of high-grade encryption.

Asymmetric encryption uses two mathematically linked keys: a public key and a private key. The public key can be distributed freely, and anyone can use it to encrypt a message. However, only someone with the corresponding private key can decrypt that message. This allows military assets anywhere in the world to encrypt information so that only the intended recipient can access it.

Asymmetric encryption works because up until now, the mathematics required to "crack" private keys is too difficult and time-intensive for standard computers (even supercomputers) to process. However, CRQCs are expected to be available within the next 3-5 years, and they threaten to render asymmetric encryption insufficient. As a result, today's militaries need quantum-secure encryption solutions that can be thoroughly tested and implemented before CRQCs become available.

Harvest now, decrypt later

Anticipating the availability of CRQCs, adversarial nations are already aiming to steal and store encrypted data until it can be decrypted. Some military data needs to remain secret for a long time, so this tactic (known as "harvest now, decrypt later") is a serious threat. In practice, it means militaries and governments ideally need a quantum-safe encryption solution significantly in advance of when CRQCs become available.

The U.S. Government has issued a **National Security Memorandum** mandating the adoption of quantum-secure encryption schemes, **supported by the NSA's Commercial Solutions for Classified (CSfC) Symmetric Key Management Requirements Annex V2.0**, and a 2025 **Executive Order** amending EO 14144 to include the following requirements, due by December 1, 2025:

- The Secretary of Homeland Security [...] shall release and thereafter regularly update a list of product categories in which products that support post-quantum cryptography (PQC) are widely available.
- The Director of the National Security Agency [...] and the Director of OMB shall each issue requirements for agencies to support, as soon as practicable, but not later than January 2, 2030, Transport Layer Security protocol version 1.3 or a successor version."

Unsurprisingly, the military is the first area where most countries aim to implement Post-Quantum Cryptography (PQC). But what will it look like?

Symmetric key encryption is more resistant to quantum computing because it doesn't rely on mathematical hardness problems, and is therefore not nearly as susceptible to the same algorithmic "shortcuts" to compromising. Unlike asymmetric, symmetric key encryption requires both sender and receiver to have the same key.

While trivial to implement in simple ecosystems, symmetric key encryption has proven challenging to scale. Both parties must have the same key for a communication to be successful, which raises the question of how the same key is made available to both in a secure way. This is difficult when dealing with a large number of devices and services, particularly when keys must be refreshed frequently to adhere to Zero Trust principles.

The challenge isn't the encryption/decryption process itself, rather the ability to generate, share, and refresh keys:

- Across complex and distributed networks, including denied and distributed environments.
- Between a wide range of endpoints, including autonomous devices.
- While being "bearer agnostic" to the communication channel in use.

In short, symmetric key encryption is simple and, when using an appropriate encryption standard, highly resistant to CRCQs. Key management for symmetric key encryption, on the other hand, has proven a significant challenge.

Current military cryptography isn't future-proof

Historically, military organizations have relied heavily on hardware devices for the encryption of classified data in transit. The NSA classifies cryptographic products or algorithms into four product types. Type 1 products are defined as:

"Cryptographic equipment, assembly or component classified or certified by NSA for encrypting and decrypting classified and sensitive national security information when appropriately keyed. Developed using established NSA business processes and containing NSA-approved algorithms. Used to protect systems requiring the most stringent protection mechanisms."

Type 2 products are for "sensitive national security information", Type 3 for "unclassified sensitive U.S. Government or commercial information", and Type 4 were not certified for government usage.

Naturally, the US military has made extensive use of Type 1 hardware devices for classified data. However, this has come with a variety of issues. **Notably, cryptography solutions based on small form-factor devices are:**

- Expensive to set up and maintain
- Restrictive of scalability and flexibility
- Difficult to operate in remote or denied locations
- Hard to extend to autonomous devices

Naturally, modern militaries would prefer a software encryption solution, as software is inherently cheaper, more flexible, and more scalable. So the question is: what does a software encryption solution that can be delivered across complex military networks look like?

The answer can be found in the NSA's **Mobile Access Capability Package (MACP)** developed for its **Commercial Solutions for Classified (CSfC)** program. The MACP aims to protect classified data in transit across untrusted networks to and from mobile endpoints.

This can be achieved using another NSA solution: **Enterprise Gray architecture.**

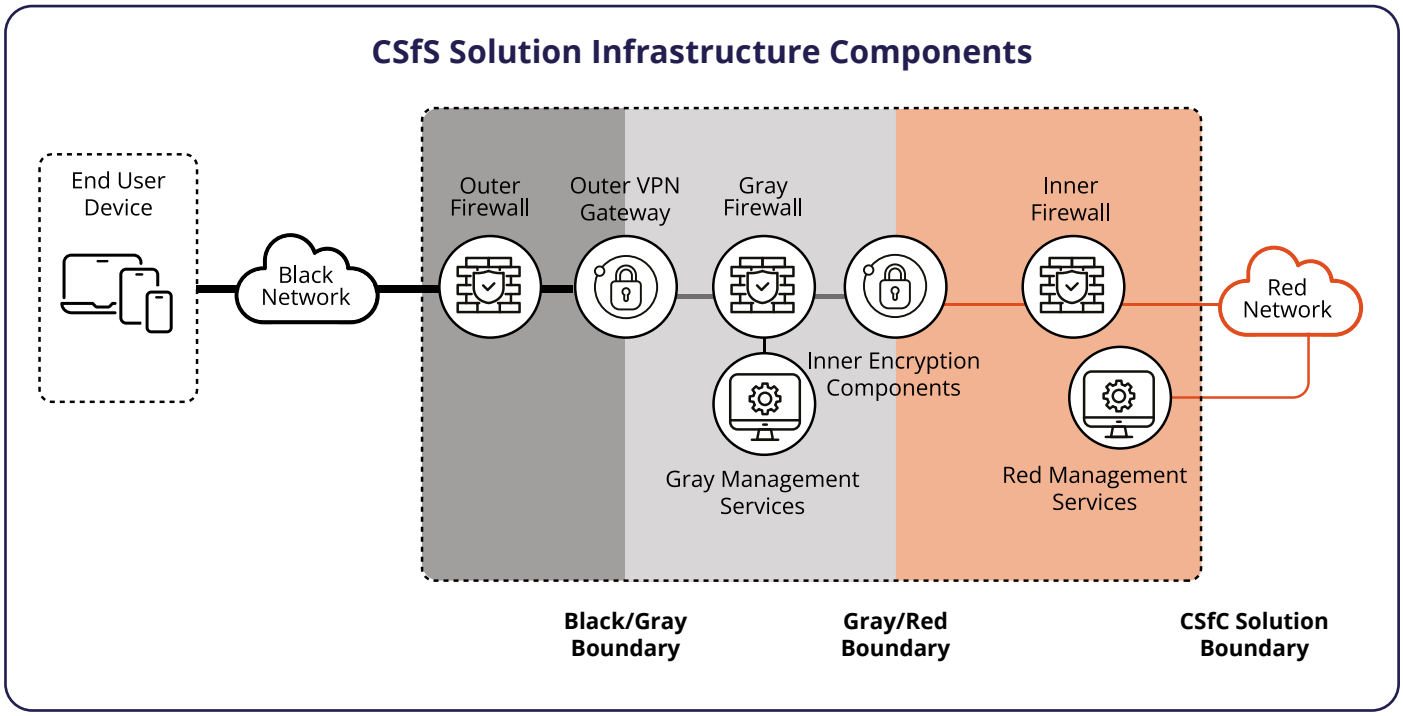

Figure 1. *Image depicts the Enterprise Gray architecture enabling end user device access to a secure network. Source: CSfC Mobile Access Capability Package*

Using Enterprise Gray architecture, sensitive data transported across unsecured internet space (black) from one secure network enclave (red) to another must go through two layers of encrypted tunnels. These are built using commercial off-the-shelf encryption components, such as VPN clients and gateways operating in an intermediate service network layer (gray).

MACP solutions via Enterprise Gray are intended to protect classified data in transit across untrusted networks, both to and from mobile devices. At the same time, this architecture eliminates the need for Type 1 encryption products, resulting in:

- Significant cost savings
- Reduced size, weight, and power (SWaP) requirements
- Eliminating technical support requirements
- Superior flexibility and scalability

In short, this approach is vastly superior to hardware-based encryption for military applications in all but one way: it still retains the vulnerability to CRQCs. The MACP solution uses Internet Protocol Security (IPsec) for the outer tunnel, and the key exchange method used to establish IPsec connections can be broken using a CRQC.

Using **Enterprise Gray architecture**, sensitive data transported across unsecured internet space (black) from one **secure network enclave (red) to another must go through two layers** of encrypted tunnels.

Cryptographic must-haves for modern militaries

At this point, it's worth taking stock of what today's militaries need from a cryptographic solution. While resistance to CRQCs is essential, it's not the only requirement.

To maintain an information advantage in modern warfighting, a military organization needs:

1. Quantum-safe encryption for data in transit (to resist CRQCs),
2. ...that doesn't rely on Type 1 hardware encryption products,
3. ...is compatible with Zero Trust principles (e.g., continuous authentication),
4. ...works for remote mobile and IoMT devices, including on untrusted networks,
5. ... and is flexible, lightweight, and hyper-scalable enough for dynamic, real-world use.

And this can be done. Italian telecommunication company, Sparkle, in collaboration with Arqit, has already deployed post-quantum VPNs across terrestrial links, seamlessly integrating with cloud-based infrastructure. Sparkle's **Quantum Safe over Internet (QSI)** initiative, which combines Arqit's symmetric key agreement technology with Sparkle's global backbone, can now enable Zero Trust-compliant connectivity between distributed defense nodes, cloud infrastructures, and command centers using symmetric key encryption that is fully quantum-safe.

These deployments show how symmetric key agreement can be orchestrated at scale, providing the foundation for secure data transport in CSfC-compliant and Zero Trust aligned architectures.

An additional solution is to use CSfC Enterprise Gray architecture, but with symmetric key encryption that doesn't rely on hardware devices, and doesn't require significantly more network and device resources compared to current implementations.

This would enable a military to continue evolving its warfighting capabilities, maximize the utility of its current and future hardware, and retain its information advantage over adversaries.

Post-quantum security for military applications

Arqit has teamed up with Equus Compute Solutions (ECS) and Intel to develop a quantum-safe MACP architecture. Deployed between ECS and Intel technical labs in California and Oregon, the solution addresses the CRQC threat and the logistics issue at the heart of symmetric key management.

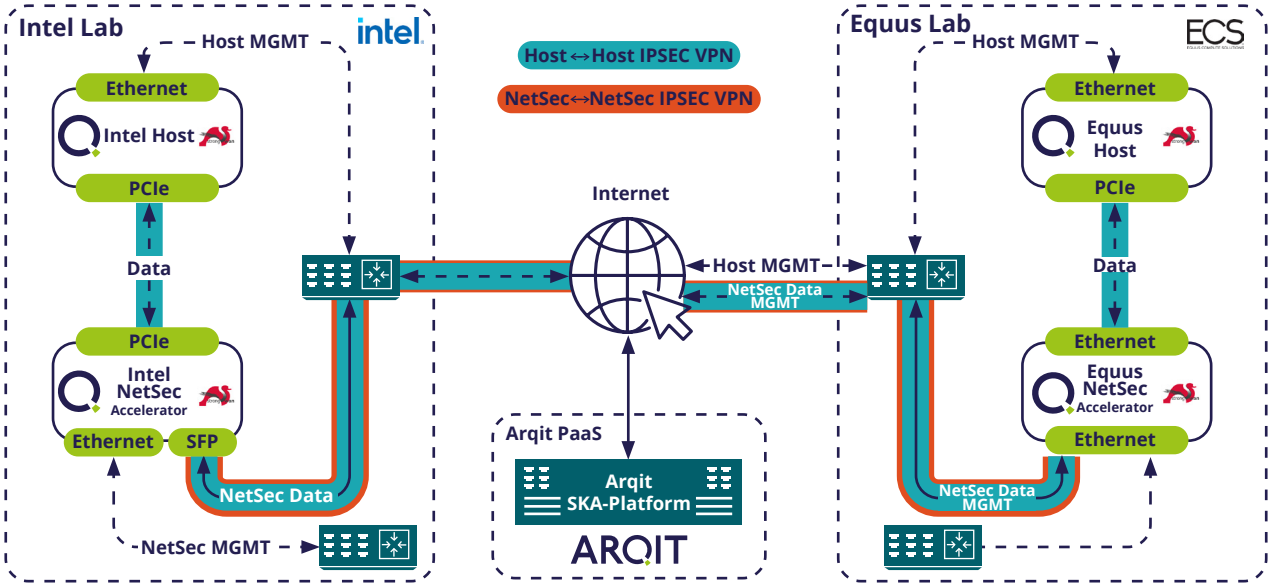

Figure 2. *Highly secure and performant MACP architecture achieved by ECS, Arqit and Intel. Note that standard firewall components were deployed at the edge of each lab and on the Intel hosts to ensure MACP compliance.*

The diagram above shows the architecture deployed between the labs. For the red and gray (inner and outer) VPN components, the open-source IPsec-based strongSwan application was chosen for its implementation of RFC 8784[3], a post-quantum standard that complies with CSfC requirements.

Note that standard firewall components were deployed at the edge of each lab and on the Intel hosts to ensure MACP compliance. Other key components were:

Arqit's SKA-Platform

SKA-Platform allows IPsec tunnel endpoints to generate quantum-safe symmetric keys. Keys can be refreshed multiple times per second, enabling dynamic rekeying of IPsec tunnels. In line with Zero Trust principles, this minimizes the lifespan of keys, preventing device spoofing and impersonation.

The platform also eliminates the need for manual key generation, couriering, loading, auditing, accounting, and other manpower-intensive and unscalable key management operations.

Security Hardware Accelerators

Intel® Xeon® Scalable processor-based hosts with Intel® NetSec Accelerator Reference Design network security accelerator cards. These cards combine an Intel® Ethernet Controller with an Intel® Xeon® D processor, packaged in a PCIe add-in card form factor. They deliver the data plane and cryptography performance needed, and their form factor allows deployment of additional network security optimized computers in space- and power-constrained locations.

RFC-8784-Compliant VPN

strongSwan, a widely used open-source VPN library, creates an out-of-the-box quantum-safe VPN. With SKA-Platform, this passes post-quantum, symmetric pre-shared keys (PSK) into the strongSwan configuration, ensuring RFC-8784 compliance. Keys can be refreshed as often as required.

strongSwan VPN nodes are monitored to ensure continuous verification, and symmetric session keys are rotated every 30 seconds to ensure perfect forward secrecy.

Testing our quantum-safe MACP architecture

The architecture was tested in two phases:

1. **Functional validation** between the two distant lab environments to demonstrate end-to-end feasibility in a real-world environment.
2. **Performance benchmarking** in a single lab to quantify the impact of introducing nested encryption schemes into high-throughput network testing.

Hardware and software specifications and network configuration are in **Appendix A**. The test environment was left unoptimized to reflect out-of-the-box behavior across all layers of the stack.

The configuration represents a realistic deployment scenario, allowing evaluation of IPsec tunnel performance in a typical production environment without advanced tuning.

Functional testing results

The functional test established a nested, quantum secure VPN connection across the open internet.

Initial "outer" tunnels were established between the NetSec accelerator card installed on the bare metal hosts' PCIe interfaces. These tunnels were dynamically keyed using Arqit SKA-Platform API integrated directly into the strongSwan daemon, with a remotely hosted SKA-Platform instance acting as the broker for key agreement between the cards.

The hosts and NetSec accelerator cards were configured so the cards acted as the default data ingress/egress route for non-management network traffic. By default, strongSwan does not support nesting of IPsec security associations and tunnels, requiring the bypassing of port-based XFRM traps and default traffic policies on the initiator NIC.

Once plugin and policy modifications were made to support nesting, a second "inner" tunnel was established between the bare metal hosts. This inner tunnel was also keyed using SKA-Platform.

Successful nesting of the IPsec tunnels was validated through analysis of native strongSwan security association metrics, XFRM policy states, and TCP/IP stack traces showing encrypted Encapsulated Security Payload (ESP) packets over the PCIe network interfaces to and from the host systems.

Test outcome
The first evaluation successfully demonstrated the MACP concept's viability.

Performance testing results

While successful, the first evaluation was inappropriate for performance testing due to the non-deterministic uncertainties imposed by the open internet backhaul. Instead, both systems were migrated to the same physical and logical network enclave, displayed below.

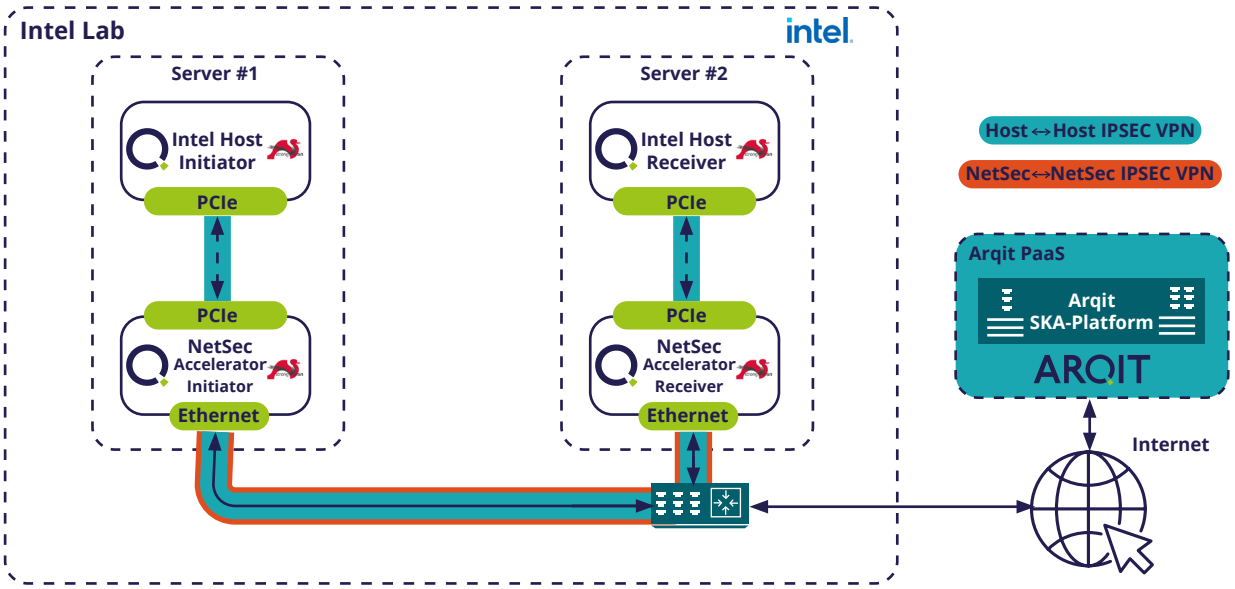

Figure 3. *Modified lab laydown for performance evaluation*

While SKA-Platform required open internet connectivity to execute the symmetric key agreement process, the strongSwan tunnels were configured to be re-keyed every 30 seconds, "make before break." This meant re-keying wouldn't result in network performance degradation, and the key agreement process did not limit throughput.

The evaluation consisted of three sequential iperf3 test scenarios, each designed to incrementally increase the complexity of the traffic path while maintaining consistency in hardware, software, and environmental conditions.

Unencrypted Baseline (No Tunnel) — Traffic was exchanged directly between the hosts via their NetSec accelerator cards, which acted as default gateways. No IPsec tunnels were in place. The measured throughput in this configuration averaged **937 Mbps**, establishing the raw performance ceiling of the setup in its unencrypted form.

Single Tunnel Configuration (NIC-to-NIC IPsec Tunnel) — Two NetSec accelerators acting as NICs established the single tunnel. The average measured throughput for a single tunnel was **904 Mbps**, representing a **3.5% decrease** from the unencrypted baseline. Despite the additional overhead of encapsulation and cryptographic processing, the impact on performance was minor, and throughput remained stable and consistent.

Nested Tunnel Configuration (Host-to-Host Tunnel Encapsulated in NIC-to-NIC Tunnel) — Under the nested configuration, the average throughput measured was **852 Mbps**, a **9.1% decrease** from the unencrypted baseline. Despite the compounded encryption and additional protocol encapsulation, the system continued to deliver stable and reliable throughput.

In all three tests, no adverse behaviors (e.g., significant retransmission, fragmentation, or erratic TCP behavior) were observed. Note that the performance impact is additive, not multiplicative, and each layer introduces a predictable, bounded penalty.

Test outcome

Quantum-secure, nested IPsec tunnels can be deployed without significant performance loss, even without tuning.

Quantum-safe, CSfC-compliant, and Zero Trust-ready

The CSfC-compliant MACP solution described here is quantum-safe, operationally scalable, and suitable to secure confidential military data in transit anywhere in the world.

Complementary capabilities (e.g., Zscaler Branch Connector and Client Connector) can extend these protections to cloud environments, hybrid networks, and end-user devices. By enabling Zero Trust enforcement at both the edge and device level, this approach implements tightly controlled, identity-driven access policies that safeguard users and data.

Further, it delivers on the specific requirements of modern military organizations by:

1. Removing the need for Type 1 hardware encryption products
2. Enabling continuous rekeying and reauthentication (in line with Zero rust principles)
3. Supporting remote mobile and connected devices, including on untrusted networks
4. Being lightweight and performant enough for real-world military applications

The solution described here overcomes the limitations of traditional symmetric key distribution. These results confirm that high-assurance mobile access to classified networks is possible using commercial, software-defined components that scale to dynamic mission environments.

This architecture provides a secure, scalable, and future-ready encryption solution to help modern military organizations maintain their information advantage — even in the face of CRQCs.

Quantum-safe solutions with full sovereignty

Naturally, militaries need to be able to implement quantum-safe solutions while retaining sovereignty over all aspects of their infrastructure. To see how Arqit could help your military organization protect against the imminent threat of CRQCs, book a demo today.

Appendix A: Hardware, software & network configuration for quantum-safe MACP testing

This appendix provides details of the systems used as host platforms and network settings for the functional and performance evaluations.

Hardware and software specifications for test systems:

Specification	Initiator Host	Receiver Host
Platform	Dell R750	2U Supermicro SYS-521C-NR|
Operating System	Ubuntu 22.04|	Ubuntu 22.04.5 LTS|
Kernel Version|	6.8.0-57-generic	5.15.0-131-generic
Network Driver	ice	ice
Driver Version	6.8.0-57-generic	5.15.0-131-generic
Firmware Version	4.40 0x8001c98b 1.3534.0|	4.40 0x8001c98b 1.3534.0|
RAM	512 GB DDR4	32GB DDR5 4800mhz
CPU	Dual socket Intel Xeon 6338N|	Intel® Xeon® Gold 6444Y

Table 1. Host system specifications

Both systems were configured with ample compute and memory headroom to prevent host-level bottlenecks from influencing throughput. Each system ran a different kernel version to validate compatibility across software revisions. Both platforms used the same Intel network interface hardware and firmware, and each system was equipped with a dual-socket Intel® Xeon® 6338N processor configuration. The test environment was left unoptimized to reflect default, out-of-the-box behavior across all layers of the network stack.

The solution described here **overcomes the limitations of traditional symmetric** key distribution.

Network settings:

- All systems operated with a standard Ethernet MTU of 1500 bytes, and no modifications were made to MSS values, socket buffer sizes, or TCP tuning parameters.
- Traffic was generated using iperf3 with default settings, resulting in standard TCP flows without application-level packet size constraints or protocol enhancements.
- IPsec tunnels used strongSwan, configured with the AES256-GCM12 cipher suite for authenticated encryption and x25519 as the key exchange mechanism, with no additional parameters set beyond what was necessary to establish the Security Associations.
- Fragmentation control, MSS clamping, and PMTU discovery behavior were all system defaults.

ARQIT

arqitgroup.com