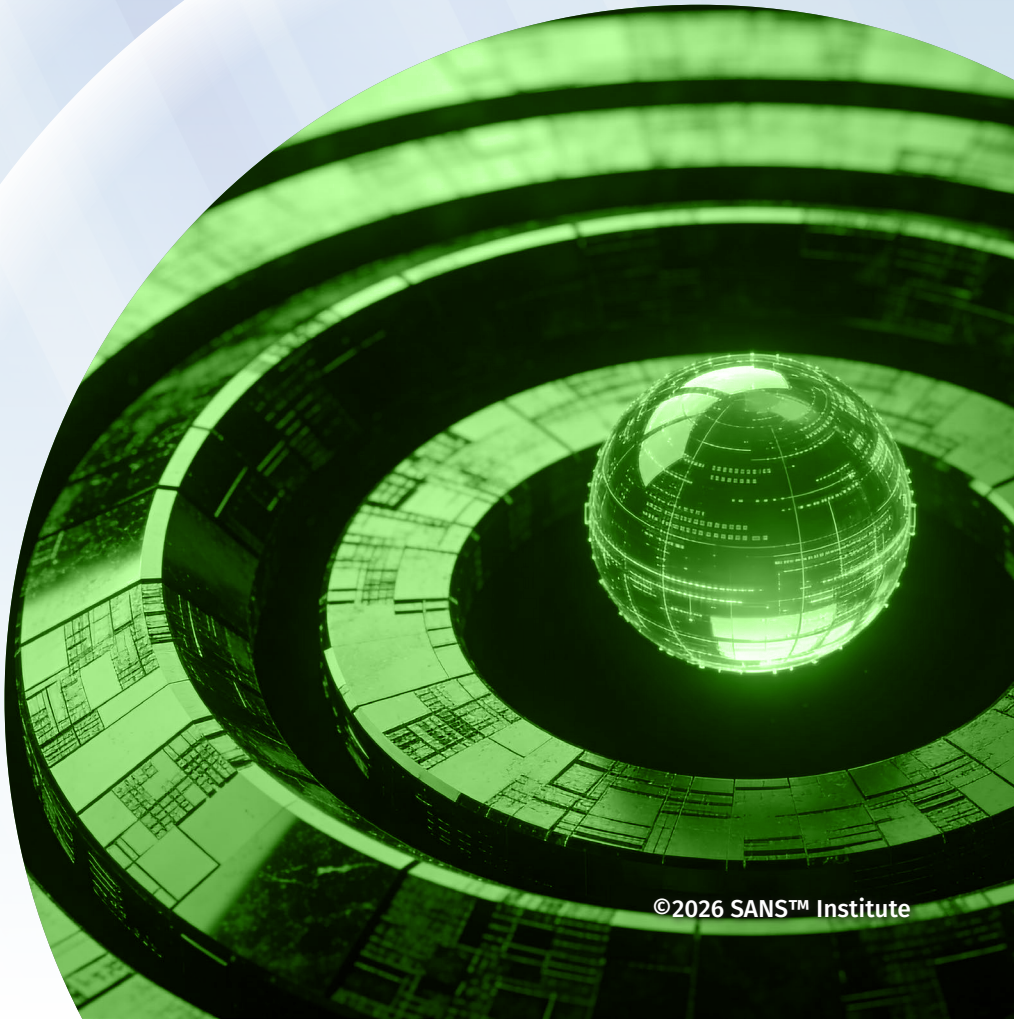
SANS | Research Program

WHITE PAPER

Air-Gapped: Isolation in an Always-Connected World

Written by **My-Ngoc Nguyen**
with ICS Contribution from
Larry Grate, ICS Expert
May 2026



BROADCOM

©2026 SANS™ Institute

Evolution of the Air Gap

For decades, the air gap—physically isolating a network from all external connections—has been the gold standard for protecting the nation's most sensitive information. From national nuclear security systems to intelligence community networks, the premise has been simple: Systems that are physically protected and cannot logically be reached by outside networks are significantly harder to compromise. For the systems that matter most (e.g., classified weapons programs, intelligence collection infrastructure, safety-critical industrial controls), that premise remains not just valid but mandatory. Today, however, sustaining it has never been more difficult, because the relentless demand for data-driven connectivity requires a more considered approach rather than the default complete separation.

The air gap was never meant to be a strategy. It was simply a fact. In the early decades of classified computing, physical separation between sensitive networks and the outside world was the natural state of things. There was no internet to connect to, no cloud to migrate workloads into, no contractor laptop sitting a few feet from a sensitive compartmented information facility (SCIF) workstation. Security came from geography and physics. That world is gone.

Today, the agencies and enterprises responsible for the nation's most sensitive systems operate in an environment of relentless connectivity pressure, with classified workloads migrating to cloud regions, operational technology networks sharing telemetry with enterprise IT platforms, and mission partners across agencies and allied nations demanding real-time data exchange. Each of these forces is legitimate, and each creates a new connectivity point to what was once an impermeable boundary. To address those issues, this white paper examines key risks facing modern air-gapped environments, including malware introduced via removable media, insider threats, supply chain compromise, governance challenges, and IT/OT convergence. Despite the risks air gap shares with any other security solutions, it is one tool within a needed defense-in-depth strategy and a rigorous, mandatory security architecture for some of the most sensitive systems. Thus, although it is not to be thought of as a silver bullet, it is an effective tool in the security toolkit. Isolation itself can be one component of a well-rounded security strategy that includes antivirus (AV), application control, device control, file integrity monitoring (FIM), data leak prevention (DLP), and incident response (IR).

Authoritative Definition and Historical Origins

The term "air gap" derives from the physical gap of air between two conductors, a space that electricity, light, radio frequency—and by analogy, data—cannot cross. The Committee on National Security Systems (CNSS) and the National Institute of Standards and Technology (NIST) define an air-gapped system as one that is "physically isolated from unsecured networks, such as the public internet or an unclassified network."[1,2] This concept predates the internet. Early government mainframes processed classified data in rooms with no external wiring at all.

A true air gap requires the absence of any network interface, wired or wireless, between protected system(s) and any external network(s). This isolation extends to indirect pathways such as shared power circuits, acoustic channels, optical links, and electromagnetic emanations vectors, which are addressed by Telecommunications Electronics Materials Protected from Emanating Spurious Transmissions (TEMPEST) standards. Physical access controls, strict media management, and procedural controls are prerequisites, not enhancements.

When Is an Air Gap Really a Virtual Gap?

In its truest form and aligned to the authoritative definition, an air gap is truly a physical gap; hence, the word "air." A common and dangerous misconception equates network or logical segmentation such as virtual local area networks (VLANs), virtual machines (VMs), and other forms of logical isolation, firewalls, or even dedicated fiber links to air-gapping. These are valuable controls, but are not air gaps. Instead, they are virtual gaps. A firewall can be misconfigured, a VLAN can be traversed through a misconfigured trunk port, and a "dedicated" fiber still represents a physical path that can be tapped or compromised. Organizations that treat logical isolation as physical isolation systematically underestimate their attack surface. However, those who understand that they are applying a virtual gap, while applying strict restricted access and formal documentation as well as enforcing the virtual gap in policy, can better manage risks and apply the needed layered defense-in-depth approach to the security posture of the sensitive system with the recommendations in upcoming sections.

[1] Committee on National Security Systems, CNSS Instruction No. 4009, 2015, https://nsarchive.gwu.edu/document/22385-document-08-committee-national-security

[2] "Guide to Operational Technology (OT) Security: NIST Publishes SP 800-82, Revision 3," September 2023, https://csrc.nist.gov/News/2023/nist-publishes-sp-800-82-revision-3

The Isolation Spectrum: Air to Virtual Gap

The earliest classified computing environments, Semi-Automatic Ground Environment (SAGE), Automatic Digital Network (AUTODIN), and the predecessor networks to today's Secret Internet Protocol Router Network (SIPRNet) and Joint Worldwide Intelligence Communications System (JWICS), used physical air gaps by necessity. There was no internet to connect to back then. Data moved by courier, punched tape, and magnetic reel. Physical separation was not a security strategy. It was simply how computing worked in the 1950s and 1960s.

Industrial control systems (ICS) and operational technology (OT) environments followed a parallel path. Power grids, water treatment facilities, and manufacturing plants ran proprietary protocols (e.g., Modbus, Distributed Network Protocol 3, Profibus) over isolated serial links. The OT world and the classified IT world shared a common architecture without knowing it; both relied on physical isolation as a primary control.

By the late 1990s and up to present day, cost pressure and mission demands forced a rethink. Agencies began aggregating classified workloads onto shared infrastructure by using VLANs and hardware security modules to enforce separation. A cross-domain solution (CDS), which includes hardware/software appliances designed to pass specific inspected data between classification levels, formalized the concept of a "virtual gap." The application of the air gap started to be interpreted by many as no longer being binary. Instead, air-gapping was seen as existing on a spectrum from strict physical isolation to policy-enforced logical separation. The evolution of a defined physical air gap to what really is a virtual gap needed to be considered with the varying degrees of logical restricted access.

The Intelligence Community's Commercial Cloud Services (IC C2S) contract, and later the Joint Warfighting Cloud Capability (JWCC), demonstrated that classified workloads could run on cloud infrastructure if the cloud itself was air-gapped from commercial internet, operated in purpose-built regions, and was governed by stringent access controls. Zero-Trust Architecture (ZTA), as defined in NIST SP 800-207, shifts the framing from perimeter defense to continuous verification. However, it does not eliminate the need for physical isolation where national security demands it. The air gap becomes a control objective and property to be verified and maintained rather than an architectural assumption.

Scaling Air Gaps Across Agencies and Enterprises

What began as isolated mainframe rooms has evolved into SIPRNet (Secret) and JWICS (Top Secret/Sensitive Compartmentalized Information/SCI), which are global wide-area networks that carry classified traffic across thousands of endpoints. The network is still "air-gapped" from the internet in the sense that no routing path exists between the two, but the internal attack surface is vast. A compromise on one SIPRNet node can traverse the same WAN that was supposed to protect it. This is why ZTA and its principles of encrypting and securing all traffic are essential to the expanded connection across global areas.

Modern intelligence and defense missions require real-time sharing of classified data across agencies with different security policies, accreditation frameworks, and risk tolerances. Federated trust models where Agency A grants Agency B limited access to specific data objects introduce complexity that scales poorly. As such, each federated connection is a potential seam in the fabric of isolation.

Each classified enclave, each CDS deployment, and each cloud extension represents a node in a governance graph that must be continuously monitored, patched, audited, and accredited. Organizations that scale infrastructure faster than their security and compliance programs create "operational debt," which is a growing backlog of unpatched systems, expired accreditations, and undocumented data flows that attackers can exploit before defenders can remediate.

Expansion of Classified Environments

Escalating strategic competition with China and Russia has significantly expanded the scope and volume of classified mission work, thus driving a surge in sensitive operations spanning intelligence collection and signal intelligence (SIGINT), cyber defense, and advanced research such as AI. This expansion has necessitated the rapid growth of classified environments to support critical research and development initiatives, including next-generation defense systems, low-yield nuclear weapons programs, and advanced deterrence technologies designed to maintain strategic advantage against near-peer adversaries.

The legacy model, a single classified enclave per facility managed by a single Information System Security Manager or Officer (ISSM or ISSO), has given way to a constellation of interconnected classified islands. A defense contractor might operate a classified lab enclave, a SCIF with a CDS connection to a government customer's network, and a classified cloud tenant in JWCC, all simultaneously. Each island has its own boundary, but the bridges between them create a composite attack surface that no single security team fully owns.

Mission partner environments (MPEs) and coalition networks require classified data to flow across not just agency boundaries but national boundaries. The Five Eyes intelligence partnership, for example, involves classified data sharing among the US, UK, Canada, Australia, and New Zealand. Each has its own legal authorities, technical standards, and risk frameworks. Managing these flows while preserving the integrity of air-gap controls requires mature CDS governance, rigorous data tagging, and continuous auditing of what crosses each boundary.

Classified Cloud Implications

The escalating demands of modern classified missions require unprecedented computational power and analytical capabilities that traditional on-premises infrastructure simply cannot scale to meet. This has driven agencies to leverage cloud environments for elastic computing resources, advanced AI/ML processing, and rapid data analytics. Cloud adoption enables intelligence and defense agencies to dynamically surge capacity during critical operations, reduce hardware procurement cycles, and tap into cutting-edge technologies. This allows analysts and researchers to focus on mission outcomes rather than infrastructure management, all while achieving significant cost efficiencies through a pay-for-use model.

As a solution for this growing need, Amazon's Commercial Cloud Services (C2S) and Microsoft's Azure Government Secret regions, along with the JWCC and other similar types of contracts—like those for the intelligence community and other agencies—enable agencies to run classified workloads on shared physical infrastructure. The "air gap" is enforced through hypervisor isolation, dedicated hardware regions, and access control rather than physical network separation. This is a materially different threat model: The attack surface includes hypervisor vulnerabilities, identity and access management misconfigurations, and the security posture of the cloud provider itself. NIST SP 800-144 and Cybersecurity and Infrastructure Security Agency's (CISA's) Cloud Security Technical Reference Architecture provide frameworks for managing these risks. From this model, the air gap concept has morphed into something it was not originally intended to be. As mentioned previously, the use of virtual and logical isolation should not be considered air-gapping an environment. Therefore, this contradicts the authoritative definition of an air gap and is really a virtual gap. For this reason, there are systems in mission environments (e.g., JWICS enclaves, SCIFs, nuclear command and control, and OT/ICS systems) that will still have continuous use of a true air-gapped architecture leveraging solutions that are fully on premises with no cloud dependency. This approach supports true air-gap deployments with manual threat intelligence import and local policy management, as recommended in following sections.

IT and OT/ICS Convergence and the Isolation Spectrum

Because ICS environments followed a parallel path to the classified environment with air-gapping, the two share a common architecture that relies on physical isolation as a primary control. However, the evolution of the need for data from the OT and ICS, to manage and monitor them, has created the convergence of IT and OT/ICS. As such, the air gap has morphed.

ICS engineers have recognized similar challenges of managing classified IT. A nuclear plant's control network, which has Modbus over serial and running programmable logic controllers (PLCs) with 20-year life cycles and no patch mechanism, is the OT equivalent of a legacy classified enclave. Both environments were designed for reliability and safety, not for security. The lessons from OT are directly applicable to classified IT and vice versa. However, the types of risk are inherently different. It is often said that "in IT, we protect data, whereas in operational environments, we are concerned with the process and how it could impact the life, safety, and health of those impacted by it."

Risks Introduced into Previously Isolated Systems

Over the years, there have been pressures and movements for enterprises to converge OT with IT. This is because of legitimate operational needs, such as remote monitoring of critical infrastructure, predictive maintenance using enterprise analytics platforms, and supply chain integration. For example, a power utility that can see real-time turbine telemetry in its enterprise data lake can optimize dispatch and reduce downtime. However, every data connection between the OT control network and the enterprise IT network is a potential lateral movement path for an attacker.

In addition, most legacy OT environments, and many newer ones, were built with the assumption that they would never be connected to external systems. They were, and often still are, designed to solve specific problems in the physical world while frequently using protocols that do not support any form of encryption or authentication. As such, when operating environments are connected to external systems without appropriate security controls in place, they become easy targets for adversarial activity or fallout from broad attacks. One of the major risks to operational environments recently has been ransomware. According to a recent report, there were 392 recorded ransomware incidents that impacted operational environments in Q3 of 2025.[3]

[3] "Dragos Industrial Ransomware Analysis: Q3 2025," December 2025, www.dragos.com/blog/dragos-industrial-ransomware-analysis-q3-2025

Why Classic Air Gaps Are Stressed but Unavoidable—The Reality

Organizations cannot function as pure islands. Mission demands, maintenance requirements, and the economics of shared infrastructure all push toward connectivity. In practice, OT environments described as "air-gapped" are frequently found to have multiple remote access connections—often installed by third parties and often unknown to the enterprise IT teams responsible for supporting them. This problem accelerated sharply during 2020 and 2021, when pandemic-driven remote access was deployed without adequate risk analysis.

Simultaneously, programs under labels like "Digital Transformation" and "Industry 4.0" are pulling operational data into enterprise analytics platforms and AI/ML pipelines, creating enormous pressure on control teams to connect systems they would otherwise keep isolated. The answer is not to abandon air-gap principles, but rather to enforce them more precisely by:

- Defining what data must cross which boundary
- Enforcing it with hardware controls
- Monitoring every crossing
- Strongly securing every endpoint and monitoring for anomalous behavior
- Auditing continuously

The air gap becomes a virtual gap and a property of specific data flows rather than a property of entire networks.

Security and Operational Challenges in Air-Gapped Environments

The two most persistent threats to air-gapped environments still share a common root: trust placed in people and equipment that enters the enclave. The 2010 Stuxnet attack on Iran's Natanz enrichment facility demonstrated that a true air gap can be defeated through supply chain compromise and removable media when a weaponized USB drive, almost certainly introduced by an unwitting insider, crossed the gap that no network connection could. But the more common defeat mechanism is far less sophisticated: a well-intentioned employee who plugs in a personal drive to "just copy one file" or a contractor who connects a laptop to a classified workstation to avoid a cumbersome transfer procedure. Insider risk in air-gapped environments is amplified by the low monitoring density typical of these spaces and by the operational pressure to work around inconvenient controls. Therefore, security culture, supported by mandatory training, enforced policy, and visible leadership accountability, is as important as any technical control.

Supply chain risk compounds the insider risk. Every device entering a classified enclave (e.g., servers, networking gear, storage media, even keyboards and KVM switches) is a potential attack vector. NSA guidance, aligned with NIST SP 800-161, emphasizes trusted supplier programs, firmware verification, and tamper-evident packaging as baseline requirements. Removable media policy should default to prohibition, with narrow, tightly controlled exceptions. Where media transfer is operationally necessary, write-once removable media with hardware-enforced encryption and mandatory hash logging is the minimum acceptable standard. In both cases, insider behavior and supply chain integrity are controls that are only as strong as the organization's willingness to enforce them consistently, including against senior personnel and high-value contractors.

Vulnerability and Risk Mitigation

Patch management is a key vulnerability mitigation. In air-gapped environments, it is where discipline meets physics. In a connected enterprise, patches flow automatically from vendor repositories to endpoints via solutions like Windows Server Update Services (WSUS) or a patch management platform. In a classified network, every patch must travel by hand—it must be downloaded on an unclassified system, verified by hash, scanned for malware, approved by change management, physically transferred via write-once-read-many (WORM) media, rescanned on the classified side, and applied after testing in an isolated staging environment. Table 1 illustrates this workflow.

Secure Patch Transfer Workflow (Unclassified → Classified)				
UNCLASSIFIED/INTERNET-FACING ZONE				
1. Vendor Patch Download (Internet)	**2.** Hash Verify (SHA-256/GPG)	**3.** AV/Malware Scan	**4.** Change Management Approval (ITSM)	**5.** Staging on Unclassified Server
▼ PHYSICAL TRANSFER BOUNDARY ▼				
Controlled Transfer: WORM DVD/encrypted USB → Two-person integrity (TPI) required → Media label, hash log, and chain-of-custody form				
▼ CLASSIFIED/AIR-GAPPED ZONE ▼				
6. Receive and Log Media	**7.** Rescan on Classified AV	**8.** Test Patch in Isolated Lab	**9.** ISSO/ISSM Sign-Off	**10.** Deploy to Prod + Log
Media physically destroyed or sanitized after deployment | All steps logged in a Workorder Management system (classified side) | RMF evidence package updated				

Table 1. Secure Patch Transfer Workflow (Unclassified to Classified)

This is very manually driven, so it's not as efficient as it could be if automated and streamlined with technical solutions commonly used in the unclassified side. This is the same for the OT environment. With the increase in patching needs from the systems to the protective devices such as firewalls, this process demands the automation and orchestration done for vulnerability management on the IT side. In addition, because the manually driven process has a patching cadence that is inherently slow by nature, application allowlisting is needed to provide an architectural approach in which only explicitly trusted executables are permitted to run. As a result, malware introduced through removable media cannot execute regardless of whether it bypassed the transfer workflow.

Accreditation, Compliance, and Change Management

Every change to a classified system—a new application, a hardware upgrade, a network topology modification—requires a formal change management review and, in most cases, a re-accreditation under the Risk Management Framework (RMF). This governance overhead is not bureaucratic excess; it is the mechanism by which the security baseline is preserved against configuration drift. Organizations that treat accreditation as a one-time event rather than a continuous process are, in effect, operating on borrowed time.

Incident Response (IR) and Forensics

Despite the protective and preventative measures, IR is much needed. Most standard IR playbooks assume network-based containment: isolate the host, pull logs from SIEM, revoke credentials remotely. In an air-gapped environment, none of these actions may be possible through normal channels. Physical access is required. Chain-of-custody procedures for forensic evidence collection must be defined in advance. Offline forensic toolkits (e.g., bootable drives with known-good tools) must be pre-staged inside the classified enclave. Without pre-planned procedures, IR teams discover their limitations at the worst possible moment.

Solutions, Controls, and Recommendations

Securing an air-gapped environment requires a layered controls architecture that operates independently of external connectivity. Unlike conventional enterprise security, where many controls assume internet access for licensing, telemetry, updates, and orchestration, air-gapped environments demand solutions that function entirely on premises, within the enclave boundary, on indefinite timescales. Cloud-based security services can complement this architecture at the boundary (e.g., threat intelligence feeds delivered via secure transfer), but the core defensive stack must be on premises by design.

The controls described in Table 2 span hardware, software, process, and governance layers, with on-premises viability as a baseline requirement. These seven recommendations are prioritized by risk reduction impact and practical ability to implement for government classified environments and their contractors.

#	Recommendation	Rationale	Implementation Note
1	Formalize and enforce patch transfer workflows and monitor the environment	Removable media remains the primary vector for air-gap defeats (e.g., Stuxnet). Documented, TPI-enforced workflows with hash verification close this gap.	Adopt NIST SP 800-53 CM-3/SI-2 controls; implement write-once media policy; log all transfers in SIEM. Deploy on-premises SIEM platforms within the enclave and aggregate logs from all classified endpoints without requiring external connectivity for licensing or telemetry.
2	Deploy hardware-based unidirectional gateways (data diodes) for high-side log export	Software firewalls can be misconfigured; data diodes offer physics-enforced one-way flow ideal for exporting telemetry without creating a return path.	Vendors such as Waterfall Security and Owl Cyber Defense offer commercial solutions; validate placement at enclave boundaries.
3	Enact endpoint protection	Beyond boundary enforcement, on-premises endpoint protection platforms (EPPs) like endpoint detection and response (EDR) provide defense-in-depth without requiring any external connectivity. These solutions operate entirely within the air-gapped perimeter and represent the primary security layer available to enclaves where cloud integration and internet connection is prohibited by policy or architecture. In addition, process execution, file writes, and registry changes are not observable through network monitoring. In environments where the threat model includes insider action or media-introduced malware, endpoint visibility inside the enclave boundary is necessary—not optional—alongside network monitoring.	Apply on-premises EPP (e.g., AV, EDR) and implement a default deny policy to include application allowlisting (aka whitelisting) as guided by NIST SP 800-167. Implement secure baseline configurations from the Center of Internet Security (CIS) or Defense Information Security Agency Security Technical Implementation Guides (DISA STIGs). Implement privileged access workstations (PAWs) with hardware-enforced isolation to reduce the attack surface for administrative access. Conduct vulnerability scanning with locally hosted signature databases.
4	Implement passive, air gap-aware monitoring using OT/classified-specific IDS	You cannot defend what you cannot see. Passive terminal access points (TAPs) on classified and OT segments provide visibility without disrupting real-time operations or safety systems.	Deploy passive sensors in-line with network TAPs. Aggregate events in an on-premises SIEM deployed within or adjacent to the classified enclave; forward summarized alerts outbound via data diode only if an enterprise SOC integration is required. Use out-of-band management networks for sensor health monitoring. On-premises SIEM deployments should not require internet connectivity for licensing or signature updates.
5	Establish a zero-trust (ZT) identity posture for human access and apply ZT principles to air-gapped zones	Jump servers and privileged workstations are high-value targets. ZT principles (e.g., least privilege, continuous verification, securing all traffic with authentication) reduce lateral movement risk.	Implement privileged access management (PAM) solutions with session recording. Enforce multifactor authentication (MFA) on all access paths, review privileged accounts quarterly per STIG guidance, and implement ZT principles.
6	Develop and exercise an air-gap incident response playbook	Standard IR playbooks assume network-based containment tools. Air-gapped environments require physical procedures like media seizure, manual isolation, and offline forensics.	Define playbooks for malware introduction via media, insider data exfiltration, and firmware tampering. Conduct tabletop exercises annually and maintain an offline forensic toolkit.
7	Integrate classified cloud (JWCC/C2S) as a controlled extension, not a replacement	For environments operating at the highest classification levels or under the most restrictive physical security requirements, on-premises infrastructure remains the authoritative deployment model. Cloud extends the architecture where policy permits, but does not replace it. If cloud adoption is necessary, it must be architected as a controlled enclave with strict CAC/PIV authentication, FIPS 140-2/3 encryption, post-quantum encryption FIPS 203/204/205, and continuous compliance monitoring rather than a bypass of air-gap principles.	Leverage FedRAMP High/IL5/IL6 cloud offerings or gain agency or enclave authorization to operation (ATO) from the respective approving official with the recommendation provided and solution for any private cloud/hybrid configurations when necessary to maintain data sovereignty and privacy while securing the environment. Implement cloud access security broker (CASB) and cloud security posture management (CSPM) tooling. Audit IAM policies monthly and treat cloud as another “island” subject to air-gap governance.

Table 2. Top Seven Prioritized Recommendations for Air-Gapped Network Security

Conclusion

The air gap is not dead. However, it can be impractical for many environments in its literal sense and might no longer be sufficient on its own. Many of the environments that the air gap was designed to protect have grown more complex, interconnected, and targeted than their architects anticipated. Today's classified networks span data centers, cloud regions, contractor facilities, and coalition partners. In addition, the control systems that run critical infrastructure share telemetry with enterprise IT platforms that share nothing of their operational resilience.

The organizations that navigate this reality best are those that move deliberately toward a virtual gap. They treat isolation not as a binary state but as a control set to be continuously measured, enforced, and defended. They invest in hardware-based controls at the highest-risk boundaries, build monitoring architectures that provide visibility without creating return channels, and maintain the operational disciplines such as patch workflows, endpoint protection, media controls, configuration control, and change management—all to keep the gap from becoming a fiction. The use of the virtual gap honors the true intent and authoritative definition of an air gap in that it must still be considered a binary state. Where a true physical air gap is required, organizations preserve this principle without compromise. Where a virtual gap is the operational reality, they govern it with the same rigor. A virtual gap must be controlled within a spectrum of isolation; the sensitivity of the system determines which controls are sufficient within that spectrum.

Sustaining either an air gap or a virtual gap demands a defense-in-depth and layered security architecture that does not just depend on network connectivity to function. This is because the invisible perimeter is only as strong as the least-disciplined person who crosses it. Technical controls set the conditions while security culture and governance sustain them. This means a defense-in-depth model and ZT principles are needed to have the security culture and governance necessary to enforce and ensure processes and technical controls are maintained. The technical controls would include endpoint protection platforms deployed and managed entirely on premises, including:

- Application allowlisting to prevent unauthorized code execution on systems that may go months between patch cycles
- Antivirus and anti-malware with locally managed signature updates delivered through the same verified transfer process used for patches
- Host-based intrusion detection operating continuously within the enclave, regardless of whether any boundary crossing is occurring

These controls are not supplementary in an air-gapped or virtual-gap environment, but instead are the last line of defense when the boundary is crossed. These are aspects that are established as part of the seven recommended controls in this paper. The adversary community understands this. Stuxnet, the SolarWinds supply chain attack, and the Volt Typhoon campaign targeting US critical infrastructure all exploited the boundary between connected and isolated systems. Defending that boundary is not a project. It is a continuous program that must be sustained across budget cycles, personnel rotations, and technology generations, whether it is air gapped or virtually gapped.

SPONSORED BY

