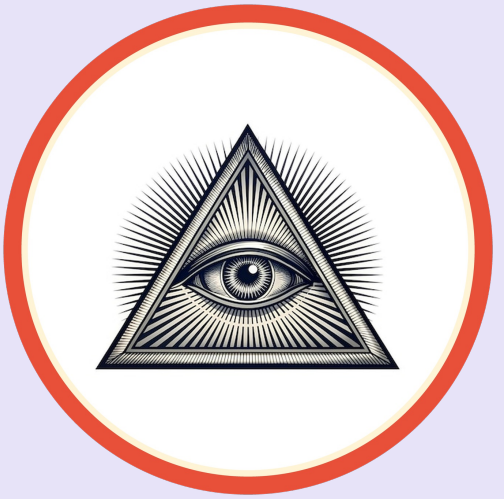
Cyber Security Report

Athens International Airport (ATH)

Landside network & public-access risk framework — Wi-Fi, remote access, perimeter hygiene, and the line between passenger and operational systems.



Version: ATH Ground Systems — Illustrative Edition, 10 September 2026

Executive Summary

About This Report

This document is a general risk-framework overview, written for briefing and training purposes. It uses Athens International Airport as a named example of a busy, multi-terminal European hub to make the discussion concrete, but it does not draw on, describe, or claim knowledge of ATH's actual network architecture, device inventory, or configuration. Every control weakness discussed here is a pattern documented across the aviation sector generally, not a finding about a live system.

Scope Note

No scanning, testing, or reconnaissance of any real airport system was performed to produce this report. Content is drawn from publicly known industry patterns, regulatory frameworks, and one verified historical incident (Section 10). Treat every figure and matrix as illustrative.

Why Landside Networks Deserve Attention

Airports sit at an unusual intersection: thousands of unmanaged passenger devices, dozens of ground-handling and retail vendors, and safety-critical operational systems all sharing the same physical campus. The interesting risk questions are rarely exotic — they're about whether the boring stuff (Wi-Fi captive portals, VPN configs, a NAT box someone forgot about) is quietly weaker than the systems everyone worries about.

The goal of good airport network hygiene isn't a perfect fortress — it's making sure a compromised passenger laptop and a compromised runway light are never on speaking terms.

Table of Contents

Executive Summary 02
- About This Report 02
- Why Landside Networks Deserve Attention 02

Table of Contents 03

Landside & Passenger Network Overview 04
- Typical Network Zones 04

Public & Transit Wi-Fi 05
- Common Weaknesses 05
- Baseline Controls 05

VPN & Remote Access 06
- Typical Uses on an Airport Campus 06
- Baseline Controls 06

NAT & Perimeter Hygiene 07
- What NAT Does (and Doesn't) for Security 07
- Baseline Controls 07

Network Segmentation 08
- Passenger vs Operational Zones 08
- Why the Line Must Hold 08

Regulatory & Governance Landscape 09
- EU & International Frameworks 09

Platform & Data Sovereignty 10
- Reducing Platform Monoculture 10
- Where Passenger Data Actually Sits 10

Threat Landscape 11
- Who Targets Airports, and Why 11

Physical & Social Engineering Risk 12
- Badge & Access Control 12
- Phishing & Pretexting 12

Incident Response & Continuity 13

Recommendations Summary 14

Risk Overview Matrix 15

Glossary 16

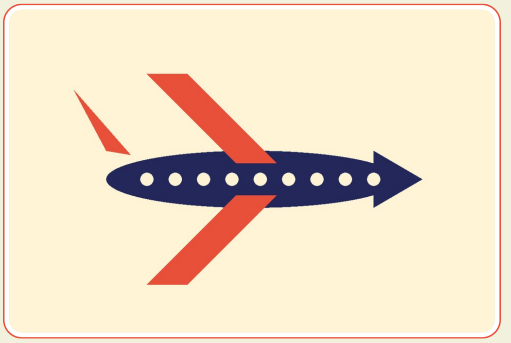
Document Control 17

Landside & Passenger Network Overview

Typical Network Zones

Large hub airports generally run several logically distinct network zones side by side on the same campus. None of the specifics below are ATH-specific — they describe the shape most major airports share:

- **Public passenger Wi-Fi** — terminals, lounges, and often shuttle buses between terminals or to remote parking.
- **Retail & concession networks** — point-of-sale, card payment, and vendor management systems.
- **Corporate / administrative IT** — email, HR, finance, ticketing back-office.
- **Ground operations** — baggage handling, gate allocation, check-in kiosks.
- **Safety-critical operational technology** — air traffic control, navigation aids, runway lighting, radar feeds. In a well-run airport this zone is logically and physically isolated from everything above it.



Illustrative only — not a diagram of any specific airport's infrastructure.

A Note on Air Traffic Control

ATC and navigation systems are not, and should never be, reachable from passenger-facing networks. Every credible framework in this report treats "no path from landside Wi-Fi to ATC" as a baseline requirement, not an aspiration — see Section 8.

Public & Transit Wi-Fi

RISK PROFILE

Likelihood	**Medium–High**
Impact	**Low–Medium**
Primary control	**Client isolation**

Common Weaknesses

Rogue & Twin Access Points

1. Free public Wi-Fi is an easy target for "evil twin" access points using a near-identical network name.
2. Shuttle-bus and inter-terminal transit Wi-Fi is a known soft spot: it's often run by a third-party transit contractor on cheaper consumer-grade hardware, with weaker oversight than terminal Wi-Fi.

Lack of Client Isolation

1. Without client isolation, one passenger's device can sometimes see and probe another's on the same access point.



Transit Wi-Fi (shuttle buses, inter-terminal links) is frequently the least-monitored segment of a public network.

Documented Incident

In 2024, Australian Federal Police charged a man who had been running "evil twin" Wi-Fi hotspots at Perth, Melbourne, and Adelaide airports, and on domestic flights — fake networks with airport-like names pushing a fake login page to anyone who connected. He was sentenced in late 2025 to over seven years. The tell was never technical: airline staff simply noticed a Wi-Fi name that didn't belong.

BASELINE CONTROLS

Client/AP isolation by default, captive portal certificate pinning where feasible, clearly branded SSIDs with no lookalike variants left unmonitored, and treating transit-vehicle Wi-Fi as part of the same security program as terminal Wi-Fi — not an afterthought run by a separate contractor.

VPN & Remote Access

RISK PROFILE

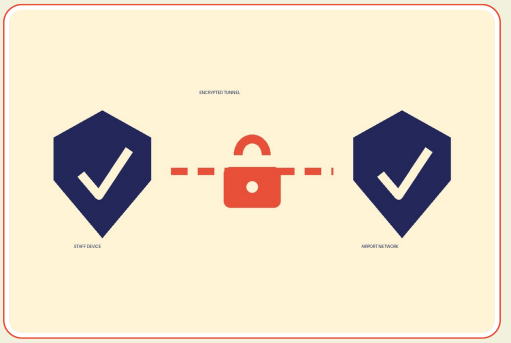
Likelihood	**Medium**
Impact	**Medium–High**
Primary control	**MFA + no split-tunnel**

Typical Uses on an Airport Campus

Airports lean on VPNs more than most sites, because so much of the workforce is contracted: ground handlers, maintenance crews, catering, cleaning, IT vendors, and airline staff who need remote access to scheduling or maintenance systems without being physically on campus.

Common Weaknesses

- **Split tunnelling** left enabled by default, so a compromised remote device keeps one foot on the public internet while connected to internal systems.
- **Shared or vendor-generic credentials** instead of per-technician accounts, making it hard to revoke access when a contractor's staff changes.
- **Flat routing** — a VPN that drops a user straight onto a broad internal subnet rather than a narrowly scoped one.
- Stale accounts for contractors whose engagement ended months earlier.



A VPN protects the tunnel — not what's allowed to happen at either end of it.

Baseline Controls

Certificate- or token-based MFA for every remote session, split-tunnelling disabled by policy, per-contractor identities tied to a joiner/mover/leaver process, and routing scoped to only the systems that specific vendor actually needs.

NAT & Perimeter Hygiene

RISK PROFILE

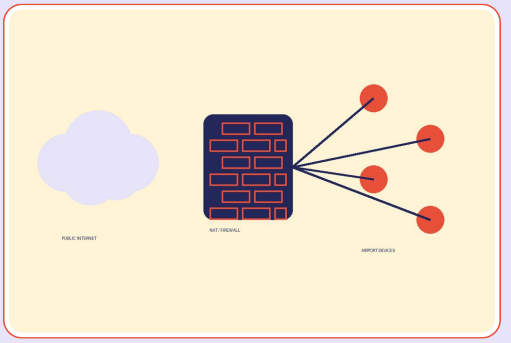
Likelihood	**Medium**
Impact	**Medium**
Primary control	**Config audits**

What NAT Does (and Doesn't) for Security

Network Address Translation lets many internal devices share a small number of public IP addresses — genuinely useful for address conservation, and it does make internal hosts harder to address directly from outside. But NAT was never designed as a security boundary, and treating "it's behind NAT" as equivalent to "it's protected" is one of the most common misreadings in network operations.

Common Misconfigurations

- Port-forwarding rules added for a one-off vendor need and never removed.
- Default credentials left on the NAT/firewall appliance's management interface.
- Firmware patch cadence lagging behind the vendor's advisories, especially on hardware installed by a contractor who has since moved on.
- UPnP left enabled, letting devices open ports automatically without review.



NAT hides addressing — it doesn't inspect or restrict what crosses the boundary.

Documented Incident

In 2018, researchers uncovered "GhostDNS" — malware that had quietly hijacked well over 100,000 home routers, mostly in Brazil, by guessing weak default admin passwords. It rewrote each router's DNS settings to point at cloned banking websites, harvesting logins from customers of several major banks who never left their own homes. The entry point was never exotic — it was the router's own default password.

BASELINE CONTROLS

Quarterly rule-set review with a "why does this exist" owner attached to every forwarded port, management interfaces never exposed to the public side, disabled UPnP, and NAT paired with a real stateful firewall policy rather than standing in for one.

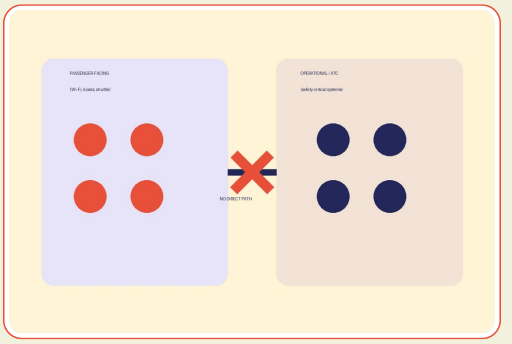
Network Segmentation

RISK PROFILE

Likelihood	**Low**
Impact	**Severe**
Primary control	**Physical/logical air-gap**

Passenger vs Operational Zones

This is the section that matters most, and it should be the shortest, because the answer is simple: passenger-facing networks — Wi-Fi, kiosks, shuttle connectivity, retail — should have **no direct network path** to operational technology or air traffic systems. Not a firewalled path, not a "monitored" path. No path.



The separation between passenger and operational zones is the single control everything else in this report supports.

Why the Line Must Hold

Every other section of this report — Wi-Fi hardening, VPN scoping, NAT hygiene — is really in service of this one boundary. A weakness in transit Wi-Fi is a landside inconvenience. A weakness that lets landside traffic reach operational systems is a different category of problem entirely, which is why aviation frameworks treat segmentation as non-negotiable rather than risk-managed.

Verification, Not Assumption

Segmentation should be periodically tested from the passenger side outward — not just documented on a network diagram and assumed to hold as vendors, contractors, and equipment change over the years.

Documented: Malware Found Near ATC

In January 2016, malware was detected inside Kyiv Boryspil Airport's IT network, in a segment that also carried its air traffic control systems, traced to a server based in Russia. Officials said no damage was done — it was caught before it could act — but it remains a rare, publicly confirmed case of state-linked activity reaching the kind of segment this section argues should never be reachable in the first place.

Regulatory & Governance Landscape

EU & International Frameworks

Airports operating in the EU increasingly sit at the overlap of several regulatory regimes, in broad terms:

- **Aviation-specific information security rules** from EASA, which push airports and their suppliers toward formal information-security management systems rather than ad hoc IT policy.
- **EU-wide critical-infrastructure cyber rules** (the NIS2 framework), which bring large airports into scope as essential-service operators with incident-reporting duties.
- **ICAO-level guidance** encouraging member states to fold cybersecurity into existing aviation-security programmes rather than treating it as a separate discipline.
- **National data-protection law**, relevant wherever passenger or staff personal data touches these networks (payment systems, Wi-Fi captive portals, HR).

Reading This Section

Regulatory detail changes frequently and varies by jurisdiction and airport ownership structure. Treat the above as a map of the terrain, not a compliance checklist — an airport's own legal and security teams are the authority on which specific obligations apply.

Governance Expectations, in Practice

Across these frameworks, the recurring themes are the same ones this report keeps returning to: named accountability for security decisions, regular independent testing, incident-reporting timelines measured in hours rather than weeks, and hard segmentation between passenger-facing and safety-critical systems.

Typical Role	Governance Responsibility
CISO / Head of IT Security	Overall accountability for segmentation and incident-response readiness
Airport Duty Manager	Escalation authority to isolate systems during an active incident
Vendor / Contractor Security Lead	Oversight of third-party remote access, badges, and VPN scoping
Data Protection Officer	Personal data handling across Wi-Fi, payment, and HR systems

Overlap, Not Redundancy

These regimes are frequently described as overlapping bureaucracy. In practice they check different things — safety-management culture, critical-infrastructure resilience, and personal-data handling rarely fail for the same reason, so a genuinely resilient airport tends to satisfy all three rather than picking one.

Platform & Data Sovereignty

Reducing Platform Monoculture

A single operating system across most of the estate is efficient to manage — and also a single patch cycle, a single set of known weaknesses, and one exploit chain away from moving laterally through almost everything it touches. Windows 11 isn't unusually weak; the risk is concentration, not the specific vendor. Where check-in kiosks, back-office terminals, and administrative workstations all run identical builds, one working exploit has the run of the estate. Reserving general-purpose Windows deployments for staff who genuinely need them, and moving single-purpose terminals (kiosks, point-of-sale, displays) onto minimal, hardened, purpose-built builds, narrows how far any single weakness can travel.

Where Passenger Data Actually Sits

Where a platform is hosted jurisdictionally matters as much as how it's secured. Data held with a US-headquartered cloud provider can, in some circumstances, be reachable by US legal process regardless of where the servers physically sit — the extraterritorial reach of US law (the CLOUD Act, most notably) is a genuine and frequently cited friction point for EU operators handling EU passenger data. For an airport operating under GDPR, hosting passenger-facing systems and their data with EU-domiciled providers, under EU jurisdiction, closes that particular gap and simplifies accountability. It doesn't make a breach impossible — nothing does — but it removes a layer of legal ambiguity about who else could compel access to the data.

The safest passenger record isn't the one guarded best — it's the one that was never centralised in the first place.

Need to Know, Not a Data Bank

Passenger names, itineraries, and contact details should be scoped to the smallest set of systems and staff who genuinely need them for the task in front of them — not replicated into every convenient dashboard "just in case." A tightly scoped, well-audited system is an inconvenience to an attacker. A sprawling, broadly-accessible one is a data bank sitting there to be pillaged the moment any single credential in the chain gets compromised.

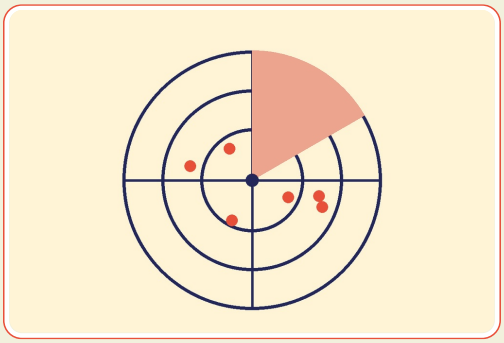
Baseline Moves

- Segment passenger PII into purpose-specific systems (check-in, loyalty, security screening) rather than one shared warehouse everything else queries.
- Apply role-based, need-to-know access — most staff should never be able to pull a full passenger record.
- Prefer EU-domiciled infrastructure and storage for anything containing passenger PII, with contractual data-residency guarantees, not just marketing claims.
- Reduce single-OS dependency on shared-use and single-purpose terminals; reserve general-purpose Windows deployments for roles that need it.
- Log and alert on bulk record access — a single query returning thousands of passenger rows is rarely a normal support task.

Threat Landscape

Who Targets Airports, and Why

- **State-linked actors** — the smallest group by headcount, but the one with the clearest airport-specific track record and the most capability. Motives run from strategic disruption during geopolitical conflict to espionage, not everyday extortion — and, uniquely among the categories below, state-linked activity has been documented both taking down airport administrative systems outright and separately probing infrastructure adjacent to air traffic control (see below).
- **Opportunistic criminal ransomware groups** — typically after administrative or back-office systems, for extortion, not operational sabotage.
- **Financially motivated fraud actors** — targeting payment flows in retail, parking, and car-hire systems.
- **Hacktivist groups** — occasionally target airport-adjacent public sites for visibility, generally via low-sophistication means (defacement, denial-of-service).



Most airport incidents hit administrative systems — but the state-linked cases are the ones that have gone closest to the operational boundary.

Most publicly documented airport cyber incidents land on the administrative or passenger-information side, not the operational side. The state-linked cases below are the exception worth the most attention: they're the category that has actually been caught near the line this report keeps returning to — which is exactly why segmentation (Section 8) is treated as non-negotiable rather than risk-managed.

Documented: State-Linked Activity

Aug 2024 — A United Airlines flight from Delhi toward New York had its GPS spoofed for hours after passing near the Black Sea (Wall Street Journal). EASA logged roughly 6,000 spoofing events over Europe in 2024 alone. Analysts attribute the bulk of this activity to military electronic-warfare operations in conflict zones — state capability, not criminal opportunism.

Jun 2017 — The NotPetya wiper attack, attributed by the CIA and UK Ministry of Defence to Russia's military intelligence service, forced Kyiv's Boryspil International Airport onto manual, whiteboard-and-marker operations alongside Ukrainian banks, the Kyiv metro, and government ministries.

FOR CONTRAST: CRIMINAL, NOT STATE

In September 2018, Bristol Airport's flight-information screens went dark for roughly two days after a ransomware-style attack hit administrative systems. Staff fell back to whiteboards. No ransom was paid, and the airport described the intrusion as opportunistic rather than targeted — the more common, lower-stakes end of the spectrum above.

Physical & Social Engineering Risk

Badge & Access Control

Large airport campuses involve thousands of staff across dozens of employers — airlines, ground handlers, caterers, retail, cleaning, security contractors. That scale makes badge hygiene a genuinely hard problem, not a lazy one:

- Tailgating through secure doors, especially at shift-change peaks when many people move through at once.
- Contractor badges that outlive the contract.
- Badge cloning risk where legacy low-frequency RFID credentials remain in use alongside newer systems.

Phishing & Pretexting

Airports' multi-employer structure is also what makes them attractive pretexting targets — "I'm from the ground-handling contractor, I need access to…" is a plausible sentence in a way it wouldn't be at a single-employer office. Common vectors:

- Invoice-fraud emails impersonating known vendors or ground-handling partners.
- Credential-harvesting pages mimicking internal single sign-on portals.
- Phone-based pretexting toward help-desk staff for password resets.

Baseline Controls

Photo-ID badge checks reinforced by staff culture (not just technology), automated badge expiry tied to contract end dates, phased retirement of legacy RFID credentials, and help-desk verification procedures that can't be talked around under time pressure.

Incident Response & Business Continuity

Preparedness

The Bristol Airport case (Section 10) is a useful template precisely because the response worked: administrative systems were isolated quickly, operational systems were never at risk, and the airport had a manual fallback ready to go. That's the pattern worth planning for.

- A pre-agreed decision tree for isolating affected systems without waiting for full diagnosis.
- Clear, rehearsed communication templates for passenger-facing updates during an outage.
- A defined internal/external reporting chain, including regulatory notification timelines.

Manual Fallback

The unglamorous end of incident response — whiteboards for flight information, paper boarding manifests, radio in place of a down messaging app — deserves the same rehearsal as the technical response. Every documented airport incident that stayed contained had one of these in place before it was needed.

Phase	Typical Actions
Detect	Automated alerting, anomaly monitoring, a clear vendor/contractor reporting channel
Contain	Isolate affected systems, preserve logs, activate manual fallback procedures
Eradicate	Remove attacker access, rotate credentials, patch the exploited weakness
Recover	Restore from clean backups, phased return of systems to service
Review	Post-incident report, regulatory notification, lessons-learned update

Test, Don't Just Document

A continuity plan that's never been rehearsed tends to fail in the same predictable ways: nobody knows who has authority to pull a system offline, and the "backup" process turns out to depend on the very system that just went down.

Recommendations Summary

Network & Access

- Enforce client isolation on all public and transit Wi-Fi.
- Disable VPN split-tunnelling; require MFA on every remote session.
- Quarterly NAT/firewall rule audits with a named owner per rule.
- Independently verify — don't just document — that landside networks have no path to operational systems.

Identity & People

- Per-contractor accounts tied to a joiner/mover/leaver process.
- Badge expiry automatically tied to contract end dates.
- Help-desk verification steps that resist time-pressure social engineering.

Governance

- Named accountability for segmentation decisions, reviewed on a fixed schedule.
- Incident-reporting timelines rehearsed, not just written down.
- Third-party/vendor security requirements written into contracts, not assumed.

Resilience

- A tested manual fallback for passenger information and boarding processes.
- Pre-agreed isolation authority so no one waits for permission mid-incident.
- Rehearsed, plain-language passenger communication templates.

None of this is exotic. Almost every recommendation on this page is "do the boring maintenance task on a schedule" — which is exactly why it keeps not happening.

Risk Overview Matrix

Illustrative ratings, consistent with the risk profiles shown on each section page — not a scored assessment of any real facility.

Category	Likelihood	Impact	Primary Mitigation
Public & transit Wi-Fi	Medium–High	Low–Medium	Client isolation, transit-Wi-Fi parity with terminal Wi-Fi
VPN & remote access	Medium	Medium–High	MFA, no split-tunnelling, scoped routing
NAT & perimeter config	Medium	Medium	Scheduled rule-set audits
Passenger/operational segmentation	Low	Severe	Verified air-gap, periodic testing
Badge & physical access	Medium	Medium	Automated expiry, tailgating culture
Phishing & pretexting	Medium–High	Medium	Verification procedures, staff training
Administrative-system ransomware	Medium	Medium	Backups, tested manual fallback

How to Read This Table

Likelihood and impact are qualitative and directional, meant to prioritise attention — not a substitute for a real, site-specific risk assessment carried out by the airport's own security function.

Glossary

NAT

Network Address Translation — lets many internal devices share a small number of public IP addresses. An addressing tool, not a firewall.

VPN

Virtual Private Network — an encrypted tunnel between a remote device and a network, commonly used for contractor and vendor remote access.

Split Tunnelling

A VPN configuration where only some of a device's traffic goes through the tunnel; the rest reaches the open internet directly.

Client Isolation

A Wi-Fi setting that stops devices on the same access point from communicating with each other directly.

Segmentation

Dividing a network into zones so that a compromise in one zone can't directly reach another — most critically here, passenger-facing vs operational systems.

OT

Operational Technology — systems that monitor or control physical processes (runway lighting, baggage systems, navigation aids), as distinct from general-purpose IT.

MFA

Multi-Factor Authentication — requiring more than one form of proof of identity to log in, e.g. a password plus a hardware token.

Pretexting

A social-engineering technique where an attacker invents a plausible cover story to obtain access or information.

Document Control

Title	Cyber Security Report — Athens International Airport (ATH), Illustrative Edition
Classification	Illustrative / Training Reference — Not an Operational Audit
Version	ATH Ground Systems — 10 September 2026
Prepared for	Internal briefing / portfolio use
Pages	17

Disclaimer

This document does not describe, and is not based on, the actual network architecture, configuration, or security posture of Athens International Airport or any other named facility. All technical content is generic aviation-sector risk framing intended for briefing, training, or design purposes. It should not be relied on as a security assessment of any real system.



Prepared as a design/reference exercise using a house report template.