AirSnitch Defensive Build Specification

Vendor-neutral implementation standard for immediate defensive remediation of AirSnitch-class client-isolation bypasses by MiniMe-Labs Ltd

Version 2.0 | April 15, 2026 | Public build blueprint

Prepared as a neutral technical reference for access-point vendors, controller vendors, hospitality operators, managed Wi-Fi providers, systems integrators, and internal engineering teams.

Intent. This document is written to help builders implement a real defensive answer quickly. It defines what a Phase 1 defensive implementation must do now, what it may defer, and how to validate that the resulting system actually reduces exposure.

Research grounding. This specification is based on the NDSS 2026 AirSnitch paper, the public AirSnitch testing tooling, and current vendor and practitioner guidance. It intentionally focuses on controls that can be implemented with today's firmware, controller, switch, gateway, or edge-enforcement capabilities rather than waiting for new standards or new client software.

1. Scope, audience, and normative language

This specification defines the minimum technical requirements for a Phase 1 defensive implementation against AirSnitch-class attacks. It is intended for any team that needs to build a working mitigation quickly, whether inside an access point, wireless controller, inline edge node, transparent bridge, or an observe-and-orchestrate platform.

The document assumes the reader already understands Ethernet, IPv4/IPv6, DHCP, ARP, NDP, WLAN basics, and enterprise or hospitality network operations. It does not assume access to any single vendor's proprietary internals.

Who should use this document

- Access-point and firmware teams implementing defensive logic inside the AP data path or bridge path.
- Controller and cloud-managed WLAN teams implementing centralized policy, telemetry, or tunnel-path controls.
- Hotel chains, venues, and managed Wi-Fi operators implementing a private edge-based protection layer without waiting for AP firmware updates.
- Systems integrators and internal security/network engineering teams building rapid-response mitigations on existing estates.

What this document is and is not

Type	Description
IS	A build specification for immediate defensive remediation and validation.
IS	A vendor-neutral reference that can be implemented through more than one enforcement point.
IS NOT	A proof that every deployment mode offers identical security coverage.
IS NOT	A replacement for later standards work or deeper per-station cryptographic redesign.
IS NOT	A claim that WPA2/WPA3 cryptography is broken.

2. Threat model and problem statement

AirSnitch-class attacks exploit the gap between how operators think client isolation works and how many products actually implement it. The research shows three root causes that recur across vendors and topologies:

- Group-key and broadcast handling can be abused to inject traffic past policies that only block unicast client-to-client forwarding.
- Isolation is often enforced at one layer only. Traffic blocked at Layer 2 can still be reintroduced through Layer 3 gateway behavior or bridge behavior.
- Client identity is weakly synchronized across the stack. MAC address, IP address, BSSID, VLAN, and forwarding state can drift or be manipulated without the platform treating that as a security event.

A Phase 1 defensive implementation therefore MUST assume that 'AP isolation enabled' is not a sufficient trust boundary. The defensive system MUST validate and enforce separation at the actual forwarding point it controls.

Threat model. The attacker is already associated to the network as an authenticated or otherwise admitted client. This is an insider or admitted-user threat model, not an unauthenticated internet-side

exploit.

Non-goals for Phase 1

- Do not require changes to end-user clients, operating systems, or applications.
- Do not require a new IEEE standard or a redesign of WPA2/WPA3.
- Do not assume universal access to proprietary AP internals.
- Do not promise complete equivalence between AP-native and edge-only deployments.
- Do not rely on long-term architecture changes before delivering immediate protection.

3. Phase 1 design goals

Goal	Meaning
Immediate risk reduction	Block or materially constrain the cleanest exploit paths now.
Cross-environment applicability	Support AP-native, controller, and edge-based implementations.
Operational safety	Avoid breaking guest service through uncontrolled validation or overblocking.
Deterministic validation	Prove whether a deployment is protected rather than assuming it.
Clear upgrade path	Allow stronger identity binding and deeper AP-side hooks later without redesigning Phase 1.

4. Deployment modes and security posture hierarchy

Builders MAY implement Phase 1 in multiple deployment modes, but they MUST describe which mode they support and MUST NOT claim equivalent coverage across modes when the telemetry or enforcement point is weaker.

Mode	Placement	Posture
Mode A	AP-native or firmware/data-path implementation	Strongest for post-decrypt frame awareness and per-BSSID behavior.
Mode B	Controller or tunnel-path implementation	Strong where traffic is centrally owned or tunneled through the controller.
Mode C	Inline edge gateway or property edge appliance	Best immediate cross-vendor operator option; strongest for forwarding and segmentation control.
Mode D	Transparent bridge enforcement	Useful for insertion into legacy estates with minimal addressing changes.
Mode E	Observe-and-orchestrate	Transitional mode only; detection plus pushes to switches/controllers/firewalls.

Important limitation. Mode C, D, and E do not inherently provide the same post-decrypt radio-layer awareness as Mode A. A builder using edge-only enforcement MUST state clearly which GTK- or BSSID-specific behaviors can be directly inspected and which can only be compensated for through stronger forwarding controls, broadcast/discovery containment, and active validation.

5. Mandatory control set

Any implementation claiming compliance with this specification MUST implement the following controls. Optional enhancements are listed later and do not replace the mandatory set.

5.1 Live client identity map

The system MUST maintain a live client identity map representing the best current view of each admitted wireless client and the infrastructure devices whose identities must be protected.

At minimum, each client record MUST include the following fields when available in the chosen deployment mode:

Field	Requirement	Purpose
station_mac	MUST	The client MAC address as observed at the enforcement or telemetry point.
ipv4_set	MUST when IPv4 exists	One or more active IPv4 addresses bound to the station.
ipv6_set	MUST when IPv6 exists	One or more active IPv6 addresses bound to the station.
segment_id	MUST	The VLAN, VRF, tunnel, SSID/ESS, or equivalent trust segment.
ingress_anchor	MUST	The AP, BSSID, controller tunnel, port, or other ingress anchor if visible.
first_seen_ts / last_seen_ts	MUST	Timestamps for state freshness and conflict handling.
state_confidence	SHOULD	High/medium/low confidence derived from source quality and recency.
mobility_window	SHOULD	Roam-grace timing for legitimate movement across APs/BSSIDs/ports.
quarantine_state	MAY	Whether the client is subject to drop, rate-limit, or isolate actions.

Infrastructure identity records MUST also exist for at least the gateway, DNS resolver, DHCP infrastructure, controller uplink peers if relevant, and any other in-scope local service whose impersonation would create a man-in-the-middle condition.

Identity map freshness and conflict handling

- A builder MUST define record expiry timeouts separately for DHCP-learned, ARP/NDP-learned, controller-learned, and observation-only state.
- A builder MUST prefer controller/AP association truth over passive observation when both are available.
- A builder MUST not silently merge conflicting identities; conflicts MUST be logged and resolved by explicit policy.
- A builder MUST preserve historical conflict context long enough to support post-incident review and active validation.

5.2 Default east-west deny for untrusted clients

The enforcement point MUST deny direct guest-to-guest or untrusted-client-to-untrusted-client forwarding by default, regardless of whether the traffic arrives as same-subnet traffic, routed traffic, or hairpinned gateway traffic.

- If the system controls a gateway, bridge, or controller path, traffic between two members of the protected client set MUST be denied unless an explicit exception applies.
- If the system controls an AP data path, direct station-to-station unicast forwarding within the protected client set MUST be denied by default.

- The implementation MUST treat same-subnet reachability through the gateway as east-west traffic, not as ordinary north-south traffic.

5.3 Gateway-bounce suppression

A compliant implementation MUST detect and suppress traffic patterns where the Layer-2 destination or next hop is the gateway or local routing point while the Layer-3 destination resolves to another protected local client.

Reference detection logic

- The system MUST know the local protected client set and the local gateway identity for the segment under inspection.
- If a packet is accepted toward the local gateway and its ultimate Layer-3 destination is another protected local client, the system MUST treat this as candidate bounce traffic.
- The system MUST exempt legitimate gateway-originated traffic, explicitly allowed captive-portal or local service flows, and inter-segment routing where the destination is not another protected local client.
- Candidate bounce traffic SHOULD be dropped inline where the enforcement point permits; otherwise it MUST be surfaced as high-severity telemetry and SHOULD trigger orchestrated blocking.

5.4 Infrastructure identity protection

The implementation MUST protect the identities of local infrastructure devices whose impersonation would permit uplink interception, service spoofing, or control-plane confusion.

- Traffic claiming to originate from the gateway, DNS, DHCP, or other protected infrastructure identity MUST be validated against the expected anchor and segment.
- A client-side appearance of a protected infrastructure MAC or IP on the wireless client side MUST be treated as a security event unless explicitly allowed for a documented reason.
- The system SHOULD pin gateway and resolver identities per segment and SHOULD detect flapping, duplicate appearance, or unexpected movement.

5.5 IPv4 and IPv6 parity

A compliant build MUST not solve the problem for IPv4 while leaving equivalent reachability or spoofing paths open in IPv6.

Area	Requirement
DHCP / DHCPv6	Learn leases and binding state when those protocols are used.
ARP / NDP	Observe and validate L2-to-L3 bindings for both families.
Router Advertisement handling	Define whether RAs are allowed, proxied, filtered, or trusted from specific anchors only.
SLAAC and privacy addresses	Allow multiple active IPv6 addresses per station and handle churn without losing attribution.
ICMPv6 essentials	Preserve required neighbor and path functions while still denying peer abuse.
Multicast control	Constrain MLD and multicast behavior consistently with the discovery policy.

5.6 Broadcast, multicast, and discovery containment

Builders MUST assume that broadcast and multicast are both operationally necessary and dangerous. Phase 1 therefore requires containment and explicit handling, not blanket trust and not blind blanket blocking.

- The system MUST preserve essential network control functions such as DHCP, ARP, NDP, and other protocol-critical control traffic required for basic operation.

- The system MUST define a local policy for service discovery protocols such as mDNS, SSDP, and equivalent mechanisms: proxy, filter, convert, scope, or deny.
- The system MUST NOT leave peer discovery unrestricted by default on guest or otherwise untrusted networks.
- Where the platform can directly inspect post-decrypt frame semantics, unicast IP payloads delivered inside L2 broadcast or multicast delivery SHOULD be treated as suspicious and MAY be dropped unless explicitly allowed by protocol policy.

Decision framework for discovery handling

Action	When to choose it
Proxy	When the service is needed but direct peer discovery should not be exposed broadly.
Convert	When multicast-to-unicast or equivalent optimization reduces blast radius without breaking service.
Scope	When the service should be limited to a room, event, device class, or service VLAN.
Deny	When the service is unnecessary on the protected segment or creates disproportionate risk.
Allow	Only when there is a documented operational reason and the resulting exposure is accepted.

5.7 Active validation engine

A compliant implementation MUST include or support an active validation capability. The system MUST verify that the enforced policy actually works in a live or staged environment instead of assuming that configuration implies protection.

- Validation MUST test direct peer reachability, gateway-bounce exposure, infrastructure impersonation exposure, and identity-conflict behavior.
- Validation MUST support a safe production mode with rate limits and bounded probe intensity.
- Validation MUST maintain a clear separation between production-safe validation and aggressive lab testing.
- Validation results MUST be retained as evidence with timestamps, target segment, result, and confidence.

5.8 Logging, evidence, and operator visibility

The system MUST produce enough evidence for an operator or builder to understand what was blocked, what was merely observed, and why.

- Every security event MUST include timestamp, segment, source identity, destination identity if known, enforcement point, action taken, and the rule or reason.
- The system SHOULD classify events into at least observe-only, blocked, quarantined, degraded-confidence, and validation-failure categories.
- The system MUST expose whether it is operating with full telemetry, partial telemetry, or degraded confidence.

6. Mobility and roaming model

Roaming is one of the easiest places for a defensive implementation to create false positives or, worse, overbroad exceptions that reopen the attack surface. A compliant implementation therefore MUST define explicit mobility handling.

Condition	Required handling
Association or handoff begins	Mark client as moving; reduce duplicate-identity severity for a bounded grace

	period only.
Old and new anchors overlap briefly	Allow overlap only within a configured mobility window and only when other evidence supports legitimate movement.
Identity appears on a non-related segment or impossible anchor	Treat as conflict or spoofing, not mobility.
Mobility window expires with continued conflict	Escalate to security event and apply policy.
Controller/AP truth available	Prefer authoritative roam events over passive inference.

The builder MUST publish the default mobility window and the signals used to prove legitimate movement. The implementation MUST NOT rely on an unbounded 'roam exemption' to avoid false positives.

7. Exception model

A compliant implementation MUST define how exceptions are introduced, scoped, reviewed, and removed. Exceptions are often necessary in real estates, but unmanaged exceptions are one of the fastest ways to recreate the original problem.

- Exceptions MUST be explicit, not implicit side effects of broad allows.
- Exceptions MUST be scope-limited by segment, device class, service, location, or time window whenever possible.
- Exceptions SHOULD be attached to a reason code and review date.
- Exceptions MUST NOT permit unrestricted guest-to-guest reachability.
- If a service requires peer discovery or local reachability, the preferred order is scoped proxying first, then scoped conversion, then narrowly scoped allow.

8. Operational state model

State	Meaning
Observing	Telemetry active, no inline blocking beyond baseline network policy.
Enforcing	Mandatory controls active and inline where available.
Quarantining	Specific clients or segments subject to stricter isolation or drop actions.
Degraded telemetry	Protection continues with reduced confidence because one or more inputs are missing or stale.
Fail-open	Security function has reduced itself to preserve guest service; operator visibility MUST reflect this clearly.
Fail-closed	Security function blocks traffic more aggressively to preserve isolation; operator visibility MUST reflect this clearly.
Validation-running	Safe active checks in progress; rate-limited and bounded.
Unable to validate	Protection state unknown; this MUST be surfaced as a risk condition.

The implementation MUST define entry and exit conditions for each state and MUST expose the current state in a way that operators can see immediately. Security claims MUST be tied to state; a system in degraded telemetry or unable-to-validate state MUST NOT present itself as fully assured.

9. Performance and safety requirements

Phase 1 is a security control for live networks. A build that is technically elegant but operationally disruptive is not compliant in practice.

- The builder MUST define acceptable added latency, throughput impact, and control-plane convergence targets for the intended deployment mode.
- The validation engine MUST be rate-limited and MUST support safe scheduling.
- The implementation MUST bound alert storms through aggregation or backoff without hiding genuine first-occurrence events.
- Where inline blocking is used, the builder MUST define rollback and bypass behavior that preserves basic service during failure.

10. Minimum validation matrix

No builder may claim success until the implementation has passed the minimum matrix below on at least one topology representative of the intended deployment. If the product supports multiple deployment modes, each supported mode SHOULD be tested separately.

Test ID	Test	Pass condition
V-01	Direct peer reachability	Two protected clients on same segment cannot directly exchange traffic absent an explicit exception.
V-02	Gateway bounce	Traffic aimed at the gateway but destined to another protected client is blocked or reliably surfaced.
V-03	Infrastructure impersonation	Client-side appearance of gateway/DNS/DHCP identity is blocked or surfaced as high severity.
V-04	Duplicate identity / mobility	Legitimate roaming is tolerated within the mobility window; spoofing conflicts escalate after the window.
V-05	Dual-stack parity	Equivalent protections hold for IPv4 and IPv6.
V-06	Discovery policy	Approved discovery use cases still function; unapproved peer discovery does not.
V-07	Production-safe validation	Validation runs do not materially disrupt guest service.
V-08	Failure behavior	Fail-open/fail-closed/degraded states are visible and behave as documented.

Recommended topology classes for testing

- Open or captive-portal guest SSID with local breakout.
- WPA2/WPA3 protected guest SSID with local breakout.
- Controller-tunneled deployment where guest traffic is centralized.
- Inline edge deployment in front of guest breakout.
- Dual-stack deployment with IPv6 enabled.
- Discovery-heavy environment such as hotel rooms, casting, or conference systems.

11. Technical build order

1. Establish visibility. Identify what telemetry and enforcement point are truly available in the chosen deployment mode.
2. Build the live client identity map and infrastructure identity registry with explicit freshness and conflict logic.
3. Implement default east-west deny at the strongest available forwarding point.
4. Implement gateway-bounce suppression and infrastructure identity protection.
5. Implement IPv4/IPv6 parity and discovery containment rules.
6. Add active validation in safe production mode.
7. Add operator visibility, evidence retention, and state reporting.
8. Add optional enhancements such as AP-native post-decrypt inspection, tighter controller correlation, or automated quarantine where platform support exists.

12. Optional enhancements

The following capabilities MAY be added, but they do not replace the mandatory control set and MUST NOT be used to postpone it:

- AP-native post-decrypt inspection of suspicious GTK or broadcast-delivery semantics.
- Controller-issued session identifiers or stronger association-state export to downstream enforcement points.
- Automated quarantine or deauthentication after repeated high-confidence events.
- Discovery proxy services integrated with room, event, or device-class policy.
- Temporal or cryptographic attestation layers that bind session identity more strongly over time.

13. Explicit implementation warnings

Warning	Why it matters
Do not assume AP isolation equals segmentation.	The whole point of Phase 1 is to validate and enforce the boundary that products often only imply.
Do not solve IPv4 only.	IPv6 can silently reintroduce local reachability and spoofing paths.
Do not use unlimited roam exemptions.	They suppress false positives by quietly suppressing security.
Do not let edge-only products overclaim AP-native visibility.	Different modes have different strengths.
Do not run aggressive validation in production by default.	The defense must not become a service outage generator.
Do not leave exceptions unmanaged.	Operational exceptions are one of the easiest ways to recreate the original risk.

14. Compliance checklist

- A documented deployment mode and enforcement point exist.
- A live client identity map exists with freshness and conflict rules.
- Infrastructure identities are registered and protected.
- Default east-west deny exists for the protected client set.

- Gateway-bounce suppression exists.
- IPv4 and IPv6 are handled consistently.
- Discovery behavior is explicitly governed rather than implicitly trusted.
- An active validation engine exists in safe production mode.
- Operational states are defined and visible.
- The minimum validation matrix has been executed and recorded.

Appendix A. Required telemetry by deployment mode

Mode	Minimum telemetry
AP-native	Association state, BSSID, client MAC, segment, post-decrypt frame context if available, local bridge/gateway knowledge.
Controller	Tunnel/session ownership, client association state, segment mapping, gateway and service identities, routed/bridged path visibility.
Inline edge	DHCP/ARP/NDP visibility, segment and gateway knowledge, protected client set, local service identities, traffic path control.
Transparent bridge	Same as inline edge plus reliable ingress/egress distinction through the bridge.
Observe-and-orchestrate	SPAN/mirror visibility, switch/controller/firewall APIs for enforcement, and a trusted source for client/segment membership.

Appendix B. External sources consulted

- NDSS 2026 paper: 'AirSnitch: Demystifying and Breaking Client Isolation in Wi-Fi Networks.'
- NDSS 2026 presentation slides for AirSnitch.
- Public AirSnitch testing repository maintained by the researchers.
- Cisco review and recommendations for AirSnitch.
- SANS architectural analysis of AirSnitch and client isolation.

This specification is intentionally conservative where platform internals vary. When a capability depends on proprietary AP or controller behavior, the implementation MUST state that dependency plainly instead of implying universal support.

Prepared by Dino Demetriou

MiniMe-Labs

WWW.MINIME-LABS.COM