

Article

Not peer-reviewed version

---

# The Drift Theorem - Proof of Policy Drift in Cisco ASA Firewalls

---

[Michael Cody](#)\*

Posted Date: 13 August 2025

doi: 10.20944/preprints202508.0977.v1

Keywords: firewall policy drift; Michael Aaron Cody; cross-layer constraint drift; Cisco ASA; network address translation (NAT); access control list (ACL); application layer gateway (ALG); firewall verification; drift theorem; SAT solver; policy verification; firewall evasion modeling; multi-layer security drift



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# The Drift Theorem - Proof of Policy Drift in Cisco ASA Firewalls

Michael Aaron Cody

Independent Theorist; mac92contact@gmail.com

## Abstract

This paper presents a constraint-based modeling framework for analyzing packet evaluation order in Cisco ASA 9.x firewalls, with emphasis on cross-layer interactions between Network Address Translation (NAT), Access Control Lists (ACL), and Application Layer Gateways (ALG). Unlike existing single-layer policy verification tools, this framework encodes multi-layer processing order directly from vendor documentation to identify cross-layer drift conditions. Packets  $p$  are defined as tuples  $(src\_ip, dst\_ip, src\_port, dst\_port, proto, flags)$ ; layer transformations  $\tau$  are deterministic mappings  $(p \rightarrow p')$ ; and policy predicates  $\Phi$  are boolean functions  $(\Phi(p) \in \{ALLOW, DENY\})$ . Based on official ASA 9.x documentation, the inbound processing sequence is modeled as  $p \rightarrow \tau_{NAT}(p) \rightarrow \Phi_{ACL}(p') \rightarrow \tau_{ALG}(p') \rightarrow p''$ , and the outbound sequence as  $p \rightarrow \Phi_{ACL}(p) \rightarrow \tau_{NAT}(p) \rightarrow \tau_{ALG}(p') \rightarrow p''$ . The analysis demonstrates that, under specific configurations (e.g., inbound static NAT to private addresses combined with SIP ALG inspection), a packet denied at the ACL stage can satisfy the final path predicate  $\Phi_{PATH}$  after subsequent transformations, without violating connection-state rules. Scope is restricted to synthetic configurations derived from public vendor documentation; no live systems were accessed or tested. Results indicate that such drift conditions are a deterministic outcome of the documented layer ordering, and can be formally identified through satisfiability analysis.

**Keywords:** firewall policy drift; Michael Aaron Cody; cross-layer constraint drift; Cisco ASA; network address translation (NAT); access control list (ACL); application layer gateway (ALG); firewall verification; drift theorem; SAT solver; policy verification; firewall evasion modeling; multi-layer security drift

## 1. Mathematical Foundation of Cross-Layer Constraint Drift

This section establishes a unified formal framework for modeling packet transformations and policy evaluations in Cisco ASA version 9.x. The aim is to create a mathematically precise structure that captures the processing order, transformation semantics, and rule evaluation logic for:

- Network Address Translation (NAT)
- Access Control List (ACL) evaluation
- Application Layer Gateway (ALG) / inspection logic

The documented order of these layers differs for inbound and outbound traffic [5,6], creating an *order-asymmetry* that can lead to cross-layer constraint drift.

### Packet and State Representation:

Let  $P$  be the set of all packets, each represented as:

$$p = (src\_ip, dst\_ip, src\_port, dst\_port, proto, flags, payload)$$

where each component is drawn from finite domains:

$$\begin{aligned} \text{src\_ip, dst\_ip} &\in \{0, 1\}^{32} && \text{(IPv4 address space)} \\ \text{src\_port, dst\_port} &\in \{0, 1, \dots, 65535\} && \text{(16-bit port space)} \\ \text{proto} &\in \{1, 6, 17, 47, \dots\} && \text{(IANA protocol numbers)} \\ \text{flags} &\in \{0, 1\}^8 && \text{(TCP flags, ICMP type/code)} \\ \text{payload} &\in \{0, 1\}^* && \text{(variable-length data)} \end{aligned}$$

Let  $S$  be the set of possible protocol/session states maintained by the firewall:

$$S = \{(conn\_id, state\_flags, timers, seq\_nums, \dots) \mid \text{active connections}\}$$

### Formal Set Definitions:

$$ACL \subseteq P \quad (\text{packets permitted by ACL rules after priority resolution})$$

$$NAT \subseteq P \times P \quad (\text{packet pairs in original} \rightarrow \text{translated form})$$

$$ALG \subseteq P \times P \quad (\text{packet pairs in pre-inspection} \rightarrow \text{post-inspection form})$$

Address space partitions:

$$\mathcal{A}_{internal} = \{10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16\} \cup \{\text{custom private ranges}\}$$

$$\mathcal{A}_{external} = \{0, 1\}^{32} \setminus \mathcal{A}_{internal}$$

### Layer Transformation Functions:

In the firewall model, each processing stage is abstracted as a transformation function that maps packet structures or packet-state tuples into new representations. This formalism allows different inspection layers, such as NAT and ALG, to be precisely defined in mathematical terms, enabling the model to reason about how address translation and application-layer processing affect packet flow and security policy enforcement.

$$\tau_{NAT} : P \rightarrow P \quad (\text{NAT address/port mapping})$$

$$\tau_{ALG} : P \times S \rightarrow P \times \mathcal{P}(P) \quad (\text{ALG inspection/modification + auxiliary flows})$$

#### 1.1. Detailed NAT Transformation Semantics

This section formalizes how NAT operations are represented in the model, distinguishing between static mappings and port address translation. By expressing these transformations as mathematical functions, the model can precisely capture how internal and external address spaces are related, allowing systematic reasoning about packet rewriting effects on firewall rule evaluation.

Static NAT:

$$\tau_{NAT}^{static}(p) = p' \text{ where } p' = (src\_ip, dst\_ip', src\_port, dst\_port', proto, flags, payload)$$

with mapping function  $\mu : \mathcal{A}_{internal} \rightarrow \mathcal{A}_{external}$ :

$$\mu(a) = \begin{cases} NAT\_table[a] & \text{if } a \in domain(NAT\_table) \\ a & \text{otherwise} \end{cases}$$

Port Address Translation (PAT):

$$\tau_{NAT}^{PAT}(p) = p' \text{ where } p' = (src\_ip, \mu(dst\_ip), src\_port, \pi(dst\_port, dst\_ip), proto, flags, payload)$$

### 1.2. Mapping Table for NAT Transformations

The mapping table enumerates specific NAT translation rules, pairing original source/destination tuples with their corresponding translated forms and conditions. This explicit representation enables a solver to match and apply the correct transformation during packet modeling, ensuring that address and port rewrites are faithfully reflected in constraint evaluations.

| Original (src, dst, port)        | Translated (NAT)                    | Condition         |
|----------------------------------|-------------------------------------|-------------------|
| (10.0.0.5, 203.0.113.10, 22)     | (198.51.100.5, 203.0.113.10, 2222)  | Static Port Remap |
| (10.0.0.6, 198.51.100.8, 80)     | (203.0.113.200, 198.51.100.8, 8080) | PAT w/ Port Remap |
| (192.168.1.0/24, any, any)       | (203.0.113.100, any, high_ports)    | Dynamic PAT Pool  |
| (172.16.0.0/16, dmz_server, 443) | (interface_ip, dmz_server, 8443)    | Interface PAT     |
| ⋮                                | ⋮                                   | ⋮                 |

### Truth Table for ACL/NAT/ALG Interactions

The truth table systematically enumerates the combined effects of ACL permissions, NAT translations, and ALG-driven auxiliary flows. By explicitly defining outcomes, flow types, and enforcement mechanisms for each combination, it enables formal modeling of both intended and emergent behaviors across processing layers. This structure supports precise solver evaluation of rule interactions and drift conditions.

Let:

A : ACL allows packet

N : NAT match/translation occurs

G : ALG creates auxiliary flows

| ACL       | NAT       | ALG       | Outcome          | Flow Type     | Mechanism               |
|-----------|-----------|-----------|------------------|---------------|-------------------------|
| A         | N         | G         | Permit           | Primary + Aux | ACLallow, NAT, ALGflows |
| A         | N         | $\bar{G}$ | Permit           | Primary       | ACLallow, NATtransform  |
| A         | $\bar{N}$ | G         | Permit           | Primary + Aux | NoNAT, ALGflows         |
| A         | $\bar{N}$ | $\bar{G}$ | Permit           | Primary       | StandardACLpermit       |
| $\bar{A}$ | N         | G         | Permit(Drift)    | Post – NAT    | NATchangesevaldomain    |
| $\bar{A}$ | N         | $\bar{G}$ | Deny             | Blocked       | ACLblockspost – NAT     |
| $\bar{A}$ | $\bar{N}$ | G         | Permit(ALGDrift) | Control&Data  | ALGbypassACL            |
| $\bar{A}$ | $\bar{N}$ | $\bar{G}$ | Deny             | Blocked       | StandardACLdeny         |

### 1.3. Theorem: Cross-Layer Constraint Drift

This theorem formalizes the condition under which a mismatch between inbound and outbound evaluation domains can occur due to cross-layer interactions between NAT, ACL, and ALG configurations. By expressing the requirement for differing address translations and rule contexts, it captures the precise scenario in which a packet may be denied under one evaluation path yet permitted under another, providing a verifiable basis for detecting drift in multi-layer firewall policies.

**Theorem 1.** For any ASA 9.x configuration  $\mathcal{C} = (NAT\_rules, ACL\_rules, ALG\_config)$  meeting:

- $|NAT\_rules| > 0$
- $|ACL\_rules| > 0$
- $\exists(p_1, p_2) \in NAT : p_1 \neq p_2$
- $evaluation\_domain_{inbound} \neq evaluation\_domain_{outbound}$

there exists  $p \in P, s \in S$  such that:

$$\Phi_{ACL}(p) = \text{DENY} \quad \wedge \quad \Phi_{PATH}(p, s) = \text{ALLOW}$$

where:

$$\Phi_{PATH_{in}}(p, s) = \Phi_{ACL}(\tau_{NAT}(p)) \wedge (\tau_{ALG}(\tau_{NAT}(p), s) \text{ succeeds})$$

$$\Phi_{PATH_{out}}(p, s) = \Phi_{ACL}(p) \wedge (\tau_{ALG}(\tau_{NAT}(p), s) \text{ succeeds})$$

### Lemma 1: Address Space Overlap Inevitability

Any ASA configuration satisfying the conditions of Theorem 1 necessarily creates overlapping ACL evaluation domains between  $\mathcal{D}_{pre}$  and  $\mathcal{D}_{post}$  in the presence of at least one non-identity NAT mapping.

**Proof.** Let  $P$  be the finite IPv4 packet space, with  $\mathcal{D}_{pre} \subset P$  denoting the set of packets evaluated *before* NAT (outbound domain) and  $\mathcal{D}_{post} \subset P$  denoting the set of packets evaluated *after* NAT (inbound domain).

From ASA processing order [6]:

$$\mathcal{D}_{pre} \cap \mathcal{D}_{post} = \emptyset$$

and NAT is a partial bijection  $\tau_{NAT} : \mathcal{D}_{pre} \rightarrow \mathcal{D}_{post}$  satisfying:

$$\exists (p_1, p_2) \in NAT \text{ such that } p_1 \neq p_2$$

(i.e., at least one non-trivial translation exists).

Let  $\mathcal{A}_{ACL}^{pre} \subseteq \mathcal{D}_{pre}$  and  $\mathcal{A}_{ACL}^{post} \subseteq \mathcal{D}_{post}$  denote the sets of packets *permitted* by ACL rules in each domain. By ASA architecture, ACLs are defined independently per domain:

$$\mathcal{A}_{ACL}^{pre} \text{ and } \mathcal{A}_{ACL}^{post} \text{ are unconstrained by any enforced equivalence.}$$

Since IPv4 address space is finite and ACL rules typically operate on CIDR ranges  $R \subset \{0, 1\}^{32}$ , define:

$$R^{pre} = \bigcup_{r \in ACL_{pre}} r, \quad R^{post} = \bigcup_{r \in ACL_{post}} r$$

where  $r$  is the set of addresses permitted by a single ACL rule.

By the pigeonhole principle: If  $\tau_{NAT}$  maps any  $p \in \mathcal{D}_{pre} \setminus R^{pre}$  into  $p' \in R^{post}$ , then there exists at least one *drift packet*  $p$  such that:

$$\Phi_{ACL}^{pre}(p) = \text{DENY} \quad \text{and} \quad \Phi_{ACL}^{post}(\tau_{NAT}(p)) = \text{ALLOW}.$$

To avoid overlap, it must hold that:

$$\tau_{NAT}(\mathcal{D}_{pre} \setminus R^{pre}) \cap R^{post} = \emptyset.$$

However, since: 1.  $\tau_{NAT}$  is surjective onto  $\mathcal{D}_{post}$  by ASA NAT processing, and 2.  $R^{post}$  is non-empty for any functional network, this disjointness condition can only hold if  $R^{post} = \emptyset$  or  $R^{pre} = P$  (both trivial and unrealistic in operational networks).

Therefore, in any non-trivial NAT + ACL configuration, there must exist at least one drift packet satisfying:

$$p \in \mathcal{D}_{pre}, \quad \Phi_{ACL}^{pre}(p) = \text{DENY}, \quad \Phi_{ACL}^{post}(\tau_{NAT}(p)) = \text{ALLOW}.$$

This proves inevitable address space overlap between ACL evaluation domains.  $\square$

**Scope and Limitations:** This applies to ASA configurations with deterministic NAT/ACL/ALG order [3,6]. Does not extend to dynamic ML-based inspection.

## 2. Processing Order in Cisco ASA 9.x

The Cisco ASA firewall applies deterministic, vendor-documented packet transformations and rule evaluations in a fixed order for each traffic direction [1,2,5,6]. The processing order is not symmetric between inbound and outbound flows, producing a measurable logical divergence in evaluation domains. This section encodes those operations mathematically, quantifying the order-asymmetry and identifying the resulting predicate-space shift that forms the basis for multi-layer constraint drift.

### Packet and Function Model:

Let the packet space be:

$$P = \{p \mid p = (\text{src\_ip}, \text{dst\_ip}, \text{src\_port}, \text{dst\_port}, \text{proto}, \text{flags}, \text{payload})\}$$

with finite domains:

$$\text{src\_ip}, \text{dst\_ip} \in \{0,1\}^{32}, \quad \text{src\_port}, \text{dst\_port} \in \{0,1,\dots,65535\}$$

$$\text{proto} \in \Pi_{\text{IANA}}, \quad \text{flags} \in \{0,1\}^8, \quad \text{payload} \in \{0,1\}^*$$

NAT translation:

$$\tau_{\text{NAT}} : D_{\text{pre}} \rightarrow D_{\text{post}}$$

where  $D_{\text{pre}}, D_{\text{post}} \subseteq P$  are the evaluation domains before and after translation. Static and dynamic NAT rules are modeled as partial bijections  $\mu : \Sigma \rightarrow \Sigma$  over the address-port subspace  $\Sigma \subset P$ .

ACL evaluation:

$$\Phi_{\text{ACL}} : P \rightarrow \{\text{ALLOW}, \text{DENY}\}$$

defined over ordered rule sets  $(r_1, r_2, \dots, r_n)$ , each  $r_i$  being a Boolean predicate on packet fields. ASA enforces the first-match rule; unmatched packets are implicitly denied.

ALG transformation:

$$\tau_{\text{ALG}} : P \times S \rightarrow P \times P(P)$$

where  $S$  is the firewall state table and  $P(P)$  denotes the set of auxiliary flows generated (e.g., SIP control/data split).

State mapping:

$$\sigma : \text{flow\_id} \mapsto \{\text{NEW}, \text{ESTABLISHED}, \text{RELATED}, \text{INVALID}\}$$

with  $\text{flow\_id}$  derived from a normalized subset of packet fields.

This model isolates each transformation and evaluation step into discrete, composable operations. The separation of  $D_{\text{pre}}$  and  $D_{\text{post}}$  is critical: it exposes the translation boundary where logical constraints may diverge, creating attack surface in multi-layer enforcement.

### 2.1. Inbound Processing Order

For inbound traffic:

$$p \xrightarrow{\tau_{\text{NAT}}} p' \xrightarrow{\Phi_{\text{ACL}}} \{\text{ALLOW}, \text{DENY}\} \xrightarrow{\tau_{\text{ALG}}} p''$$

Sequence:

- (i)  $\tau_{\text{NAT}}(p)$  rewrites addresses/ports according to NAT rules [5].



- (ii)  $\Phi_{\text{ACL}}(p')$  applies ACL predicates in the post-NAT domain  $D_{\text{post}}$ .
- (iii)  $\tau_{\text{ALG}}(p')$  may alter payload and headers or create auxiliary sessions.

Formally:

$$\text{Order}_{\text{in}} = \tau_{\text{ALG}} \circ \Phi_{\text{ACL}} \circ \tau_{\text{NAT}}$$

The inbound order applies access control after translation, enabling any NAT mapping from a denied pre-NAT tuple to a permitted post-NAT tuple to bypass policy intent. This outcome is architectural, not misconfigurational, and is the first condition for asymmetric predicate satisfaction.

### Outbound Processing Order

For outbound traffic:

$$p \xrightarrow{\Phi_{\text{ACL}}} \{\text{ALLOW}, \text{DENY}\} \xrightarrow{\tau_{\text{NAT}}} p' \xrightarrow{\tau_{\text{ALG}}} p''$$

Sequence:

- (i)  $\Phi_{\text{ACL}}(p)$  applies predicates in the pre-NAT domain  $D_{\text{pre}}$ .
- (ii)  $\tau_{\text{NAT}}(p)$  translates allowed packets.
- (iii)  $\tau_{\text{ALG}}(p')$  may apply application-layer rewrites.

Formally:

$$\text{Order}_{\text{out}} = \tau_{\text{ALG}} \circ \tau_{\text{NAT}} \circ \Phi_{\text{ACL}}$$

The outbound order keeps ACL predicates independent of NAT rules, which removes the translation boundary from the evaluation space. This creates a non-isomorphic mapping between inbound and outbound constraint spaces.

### 2.2. Asymmetry Exploitation Surface

Since  $\tau_{\text{NAT}}$  and  $\Phi_{\text{ACL}}$  do not commute:

$$\tau_{\text{NAT}} \circ \Phi_{\text{ACL}} \neq \Phi_{\text{ACL}} \circ \tau_{\text{NAT}}$$

it follows that there exists  $p \in P$  such that:

$$\Phi_{\text{ACL}}(p) \neq \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p))$$

Define the drift vector:

$$\delta_{\text{drift}} = \{p \in D_{\text{pre}} \mid \Phi_{\text{ACL}}(p) = \text{DENY} \wedge \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p)) = \text{ALLOW}\}$$

In practical terms,  $\delta_{\text{drift}}$  defines the packet set capable of crossing the translation boundary undetected by ACL intent. Enumeration of  $\delta_{\text{drift}}$  via constraint solvers yields deterministic candidates for traversal through multi-layer policy without violating state tracking.

### Enhanced NAT/ACL Domain Surjectivity Analysis:

**Lemma 2.1:** If  $\tau_{\text{NAT}}$  is surjective over  $D_{\text{post}}$  and  $\Phi_{\text{ACL}}$  admits at least one ALLOW predicate in  $D_{\text{post}}$  not satisfied by any element of  $D_{\text{pre}}$ , then  $\delta_{\text{drift}} \neq \emptyset$ .

**Proof:** Surjectivity ensures that  $\forall y \in D_{\text{post}}, \exists x \in D_{\text{pre}} : \tau_{\text{NAT}}(x) = y$ . By assumption, there exists  $y^* \in D_{\text{post}}$  such that  $\Phi_{\text{ACL}}(y^*) = \text{ALLOW}$  and  $\forall x \in D_{\text{pre}}, x \neq y^*, \Phi_{\text{ACL}}(x) = \text{DENY}$ . Surjectivity guarantees  $\exists x^* \in D_{\text{pre}} : \tau_{\text{NAT}}(x^*) = y^*$ . Therefore,  $\Phi_{\text{ACL}}(x^*) = \text{DENY}$  while  $\Phi_{\text{ACL}}(\tau_{\text{NAT}}(x^*)) = \text{ALLOW}$ , implying  $x^* \in \delta_{\text{drift}}$ .  $\square$

**Statistical Prevalence Analysis:** To address the practical applicability of Lemma 2.1, this section analyze the frequency of its conditions in operational ASA deployments. Enterprise NAT configurations commonly exhibit near-surjectivity through:

- Dynamic PAT pools covering significant address ranges
- Interface-based NAT with broad port allocation
- Multiple static NAT rules creating comprehensive mappings

The condition requiring ACL asymmetry ( $\exists y^* \in D_{\text{post}} : \Phi_{\text{ACL}}(y^*) = \text{ALLOW} \wedge \forall x \in D_{\text{pre}}, x \neq y^*, \Phi_{\text{ACL}}(x) = \text{DENY}$ ) occurs frequently in practice due to:

- Internal-to-DMZ access rules (post-NAT allows  $192.168.x.x \rightarrow \text{DMZ}$ , pre-NAT denies external  $\rightarrow \text{DMZ}$ )
- Service-specific NAT mappings (SSH access allowed to translated  $10.0.0.x$  range, denied to original external IPs)
- Network segmentation policies that differ between address spaces

### 2.3. Bidirectional Drift Vector Definition

To address the complete drift space, I extend the definition:

$$\delta_{\text{drift}}^{\text{full}} = \delta_{\text{drift}}^+ \cup \delta_{\text{drift}}^-$$

where:

$$\delta_{\text{drift}}^+ = \{p \in D_{\text{pre}} \mid \Phi_{\text{ACL}}(p) = \text{DENY} \wedge \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p)) = \text{ALLOW}\}$$

$$\delta_{\text{drift}}^- = \{p \in D_{\text{pre}} \mid \Phi_{\text{ACL}}(p) = \text{ALLOW} \wedge \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p)) = \text{DENY}\}$$

For security analysis,  $\delta_{\text{drift}}^+$  represents the policy bypass surface (the primary concern), while  $\delta_{\text{drift}}^-$  represents legitimate traffic blocking (availability concern).

#### Stateful Tracking Integration:

While the mathematical framework focuses on stateless rule evaluation, practical exploitation must account for connection state tracking. Define the stateful constraint:

$$\Phi_{\text{STATE}} : P \times S \rightarrow \{\text{VALID}, \text{INVALID}\}$$

The effective drift vector under stateful operation becomes:

$$\delta_{\text{drift}}^{\text{effective}} = \{p \in \delta_{\text{drift}}^+ \mid \exists s \in S : \Phi_{\text{STATE}}(p, s) = \text{VALID}\}$$

This refinement maintains mathematical rigor while acknowledging that practical exploitation requires satisfying both rule logic and state constraints.

### 2.4. Drift Discovery Methodology

To bridge the gap between mathematical inevitability and practical identification, this section outline a systematic approach for discovering exploitable drift instances:

#### Step 1: Configuration Analysis

1. Extract NAT rule set  $N = \{n_1, n_2, \dots, n_k\}$
2. Extract ACL rule set  $A = \{a_1, a_2, \dots, a_m\}$
3. Identify evaluation domains  $D_{\text{pre}}$  and  $D_{\text{post}}$

#### Step 2: Constraint Encoding

1. Model  $\tau_{\text{NAT}}$  as constraint relations over packet fields
2. Model  $\Phi_{\text{ACL}}$  as Boolean predicates in CNF form
3. Construct drift detection formula:  $\phi_{\text{drift}} = \neg\Phi_{\text{ACL}}(p) \wedge \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p))$

#### Step 3: Satisfiability Resolution



1. Submit  $\phi_{\text{drift}}$  to SMT solver (Z3, CVC4)
2. Enumerate satisfying assignments representing drift candidates
3. Validate candidates against state tracking constraints

This methodology transforms the theoretical drift vector into a computable enumeration process, enabling systematic identification of policy bypass opportunities.

This enhanced analysis demonstrates that multi-layer constraint drift is both mathematically inevitable under common configuration patterns and systematically discoverable through formal methods. The combination of theoretical foundations and practical methodology establishes a comprehensive framework for analyzing ASA policy vulnerabilities.

### 3. Drift Theorem in Multi-Layer ASA Processing

This section states and introduces the Drift Theorem, here refined as the *State-Enabled Domain Escape Theorem*, which demonstrates that Cisco ASA 9.x firewall architecture mathematically guarantees the existence of packets that would be denied by Access Control List (ACL) evaluation in a new-connection context, but are permitted by the complete multi-layer processing chain when evaluated under certain reachable connection states. This result is a direct consequence of the vendor-documented processing order asymmetry and represents a fundamental architectural characteristic rather than a configuration error [1,2,5,6]. The theorem establishes formal conditions under which ACL-defined denial domains can be traversed without matching their new-connection criteria.

As established in Section 2.2, the NAT-ACL composition in ASA 9.x processing is non-commutative:

$$\tau_{\text{NAT}} \circ \Phi_{\text{ACL}} \neq \Phi_{\text{ACL}} \circ \tau_{\text{NAT}}$$

This property defines the  $\delta_{\text{drift}}^+$  set, which contains packets whose pre- and post-NAT ACL evaluations differ. The existence of  $\delta_{\text{drift}}^+ \neq \emptyset$  does not, by itself, prove a security bypass. However, it creates the necessary condition for state-enabled domain escape, which will be examined in the subsequent subsections.

The security-relevant drift set is:

$$\delta_{\text{drift}}^+ = \{p \in D_{\text{pre}} \mid \Phi_{\text{ACL}}(p) = \text{DENY} \wedge \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p)) = \text{ALLOW}\}$$

For configurations where  $\delta_{\text{drift}}^+ \neq \emptyset$ , state-enabled domain escape becomes mathematically inevitable.

**Formal Multi-Layer Processing Model:** Building on the formal policy representations discussed in Section 2, I now model the Cisco ASA packet-processing pipeline, including NAT translation, ACL filtering, and stateful inspection, as a unified constraint system. While earlier research has applied SAT/SMT and related methods to verify and optimize such configurations. This section is to express the combined behavior of these layers in a form that permits systematic reasoning about all admissible and inadmissible traffic states. This formulation enables us to analyze potential state transitions that, although compliant with individual layer rules, could collectively result in outcomes contrary to the intended security policy.

#### 3.1. Layer Composition

I define the processing chain for inbound and outbound flows as follows:

$$\Phi_{\text{PATH}}^{\text{in}}(p, s) = \begin{cases} \text{ALLOW,} & \text{if } \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p)) = \text{ALLOW} \\ & \wedge \Phi_{\text{STATE}}(\tau_{\text{NAT}}(p), s) = \text{VALID} \\ & \wedge \tau_{\text{ALG}}(\tau_{\text{NAT}}(p), s) = \text{ACCEPT} \\ \text{DENY,} & \text{otherwise} \end{cases}$$

$$\Phi_{\text{PATH}}^{\text{out}}(p, s) = \begin{cases} \text{ALLOW}, & \text{if } \Phi_{\text{ACL}}(p) = \text{ALLOW} \\ & \wedge \Phi_{\text{STATE}}(\tau_{\text{NAT}}(p), s) = \text{VALID} \\ & \wedge \tau_{\text{ALG}}(\tau_{\text{NAT}}(p), s) = \text{ACCEPT} \\ \text{DENY}, & \text{otherwise} \end{cases}$$

State validation is given by:

$$\Phi_{\text{STATE}}(p, s) = \text{VALID} \iff \text{flow\_tuple}(p) \in s \vee \text{can\_establish\_state}(p) = \text{TRUE}$$

ALG transformation with acceptance criteria:

$$\tau_{\text{ALG}}(p, s) = \begin{cases} \text{ACCEPT} & \text{if application-layer constraints satisfied} \\ \text{REJECT} & \text{otherwise} \end{cases}$$

The multi-layer constraint system can be written as:

$$\mathcal{C}(p, s) = \Phi_{\text{PATH}}(p, s) \wedge \text{PolicyCompliance}(p, s)$$

where policy compliance encodes administrative intent, which may not be equivalent to the conjunction of individual layer decisions.

NAT, ACL, and state predicates are represented as Boolean constraints. Inbound and outbound sequences are encoded separately:

$$\text{Inbound: } \exists p \in P : \text{ACL}_{\text{in}}(\tau_{\text{N}}(p)) \neq \text{ACL}_{\text{out}}(p)$$

$$\text{Outbound: } \exists p \in P : \text{ACL}_{\text{out}}(p) \neq \text{ACL}_{\text{in}}(\tau_{\text{N}}(p))$$

Such asymmetries form the basis for detecting potential policy drift conditions.

**Complexity Considerations:** Let  $|P|$  be the space of possible packets,  $|N|$  the number of NAT rules, and  $|A|$  the number of ACL rules. Naive enumeration is  $O(|P| \cdot |N| \cdot |A|)$ . A constraint solving reduces search space significantly but retains exponential worst-case complexity in the bit-width of packet header fields. For realistic enterprise configurations, constraint pruning (e.g., prefix aggregation, port range merging) is essential for tractability.

**Security Interpretation:** This formalism allows identification of traffic patterns that appear permissible in the combined multi-layer pipeline yet would be denied under a single-layer interpretation of policy intent. Such cases represent architectural inconsistencies that can be leveraged for both attack simulation and defensive policy validation.

### 3.2. Theorem Structure

Having established the mathematical foundation for multi-layer policy drift in Section 2, I now provide the formal proof structure for Theorem. This constructive proof demonstrates the existence of packet-state pairs that satisfy ACL denial in policy intent while achieving end-to-end traversal through the complete ASA processing pipeline. The proof leverages the documented processing order asymmetry to show how state-enabled domain escape occurs under realistic configuration conditions.

**Given:**

- $\Phi_{\text{ACL}}$ : Access Control List evaluation function.
- $\tau_{\text{NAT}}$ : Network Address Translation mapping.
- $\tau_{\text{ALG}}$ : Application Layer Gateway transformation.
- $\Phi_{\text{STATE}}$ : Stateful inspection predicate.
- $\Phi_{\text{PATH}}$ : Complete ASA multi-layer processing function.
- ASA 9.x processing order:  $\text{NAT} \rightarrow \text{ACL} \rightarrow \text{ALG}$  (inbound),  $\text{ACL} \rightarrow \text{NAT} \rightarrow \text{ALG}$  (outbound).

From Section 2:

$$\delta_{\text{drift}}^+ = \{p \mid \Phi_{\text{ACL}}(p) = \text{DENY} \wedge \Phi_{\text{ACL}}(\tau_{\text{NAT}}(p)) = \text{ALLOW}\}$$

This set captures packets whose ACL disposition changes after NAT transformation, forming the core candidate set for drift analysis.

Policy Intent Evaluation Function: This function models what the ACL would decide for a given packet  $p$  if there were no pre-existing connection state:

$$\Phi_{\text{ACL}}^{\text{new}}(p) \triangleq \Phi_{\text{ACL}}(p) \quad \text{with } s = \emptyset$$

This is a comparator for policy enforcement analysis; it is *not* a literal stage in ASA processing.

**To Prove:** For configurations with  $\delta_{\text{drift}}^+ \neq \emptyset$ , there exists a packet-state pair  $(p, s)$  such that:

$$\Phi_{\text{ACL}}^{\text{new}}(p) = \text{DENY} \quad \wedge \quad \Phi_{\text{PATH}}(p, s) = \text{ALLOW}$$

where  $\Phi_{\text{PATH}}$  represents actual ASA multi-layer processing.

Proof by Construction:

### Step 1: State Seed Generation

Select a configuration with:

- NAT rule:  $\tau_{\text{NAT}}(192.168.1.100) = 203.0.113.100$  (static, bidirectional mapping).
- ACL rule: permit tcp 192.168.1.0/24 any eq 22.
- Processing asymmetry: inbound NAT  $\rightarrow$  ACL; outbound ACL  $\rightarrow$  NAT.

Define the outbound seed packet:

$$p_{\text{seed}} = (192.168.1.100, 203.0.113.50, 22, 12345, \text{TCP}, \text{SYN})$$

Outbound sequence:

$$\Phi_{\text{ACL}}(p_{\text{seed}}) = \text{ALLOW} \quad (\text{source in } 192.168.1.0/24) \quad (1)$$

$$\tau_{\text{NAT}}(p_{\text{seed}}) = (203.0.113.100, 203.0.113.50, 22, 12345, \text{TCP}, \text{SYN}) \quad (2)$$

$$s^* = \{(203.0.113.100, 203.0.113.50, 22, 12345, \text{ESTABLISHED})\} \quad (3)$$

### Step 2: Drift Candidate Packet

Return packet exploiting state:

$$p^* = (203.0.113.50, 203.0.113.100, 12345, 22, \text{TCP}, \text{ACK})$$

Inbound NAT translation:

$$\tau_{\text{NAT}}(p^*) = (203.0.113.50, 192.168.1.100, 12345, 22, \text{TCP}, \text{ACK})$$

### Step 3: Policy Intent vs Actual Path

Policy intent evaluation:

$$\Phi_{\text{ACL}}^{\text{new}}(\tau_{\text{NAT}}(p^*)) = \text{DENY} \quad (\text{source not in permitted range})$$

Actual path evaluation:

$$\Phi_{\text{STATE}}(\tau_{\text{NAT}}(p^*), s^*) = \text{VALID} \quad (4)$$

$$\tau_{\text{ALG}}(\tau_{\text{NAT}}(p^*), s^*) = \text{ACCEPT} \quad (5)$$

$$\Phi_{\text{PATH}}^{\text{in}}(p^*, s^*) = \text{ALLOW} \quad (6)$$

#### Step 4: Satisfaction of Theorem Condition

$$\Phi_{\text{ACL}}^{\text{new}}(p^*) = \text{DENY} \wedge \Phi_{\text{PATH}}(p^*, s^*) = \text{ALLOW}$$

Thus the theorem holds for  $(p^*, s^*)$ .

For any ASA configuration where:

1.  $\delta_{\text{drift}}^+ \neq \emptyset$ ,
2. Bidirectional NAT permits reverse flow mapping,
3. Stateful inspection allows return packets without ACL re-evaluation,
4. ALG accepts the protocol on established connections,

the existence of  $(p, s)$  satisfying the theorem's condition is *mathematically inevitable* under the above constraints.

**Note:** This result does not imply a software defect; it is an architectural property arising from:

$$\tau_{\text{NAT}} \circ \Phi_{\text{ACL}} \neq \Phi_{\text{ACL}} \circ \tau_{\text{NAT}}$$

in conjunction with asymmetric processing order between directions and stateful overrides of ACL logic.

#### 3.3. Multi-Layer Policy Drift Theorem

This theorem defines the empirically observable condition under which Access Control List (ACL) evaluation and subsequent state-enabled processing in ASA 9.x produce divergent outcomes. This divergence creates a repeatable, verifiable path for packets to traverse ACL-defined denial zones without meeting new-connection criteria.

**ASA 9.x Policy Drift Theorem:** For any ASA 9.x configuration with  $\delta_{\text{drift}}^+ \neq \emptyset$ , there exist packet-state pairs  $(p, s)$  such that:

$$\Phi_{\text{ACL}}^{\text{new}}(p) = \text{DENY} \wedge \Phi_{\text{PATH}}(p, s) = \text{ALLOW}$$

where:

- $\Phi_{\text{ACL}}^{\text{new}}(p)$  is the ACL decision on a new connection attempt (no prior state)
- $\Phi_{\text{PATH}}(p, s)$  is the complete ASA processing result with connection state  $s$
- $s$  is a valid connection state created by a prior allowed packet flow  $p_{\text{seed}}$

This condition can be validated by observing ASA packet traces and ACL logs, confirming that packets matching  $\delta_{\text{drift}}^+$  bypass new-connection ACL enforcement solely due to pre-existing, reachable connection state.

#### 3.4. Constructive Proof

The constructive proof above demonstrates that ASA 9.x firewalls can permit traffic flows that explicitly violate ACL new-connection intent when stateful inspection is active. This drift condition is not the result of a misconfiguration or rule ordering mistake, but an inherent outcome of the platform's processing sequence (NAT  $\rightarrow$  ACL inbound, ACL  $\rightarrow$  NAT outbound). An attacker capable of initiating or influencing an outbound session can leverage this behavior to establish a valid connection state and subsequently deliver inbound traffic from otherwise denied sources, bypassing intended ACL

boundaries without generating new-connection log entries. This represents a measurable reduction in ACL enforcement integrity and should be treated as a security-relevant architectural property.

- Static NAT:  $203.0.113.100 \leftrightarrow 192.168.1.100$
- ACL: permit tcp 192.168.1.0/24 any eq ssh
- ALG: SSH inspection (standard behavior)

**Phase 1: Outbound Connection Establishment:**

$$p_{\text{out}} = (192.168.1.100, 203.0.113.50, 22, 12345, \text{TCP}, \text{SYN}, \emptyset)$$

1. ACL evaluation:  $\Phi_{\text{ACL}}(192.168.1.100, 203.0.113.50, 22) = \text{ALLOW}$
2. NAT translation
3. State creation:  $s^* = \{(203.0.113.100, 203.0.113.50, 22, 12345, \text{ESTABLISHED})\}$
4. ALG acceptance

**Phase 2: Return Traffic Domain Escape:**

$$p_{\text{ret}} = (203.0.113.50, 203.0.113.100, 12345, 22, \text{TCP}, \text{ACK}, \text{data})$$

1. NAT translation to  $(203.0.113.50, 192.168.1.100, 12345, 22)$
2. ACL new-connection evaluation:  $\Phi_{\text{ACL}}^{\text{new}}(\dots) = \text{DENY}$
3. State validation:  $\Phi_{\text{STATE}}(\dots, s^*) = \text{VALID}$
4. Stateful override: ACL denial superseded by state
5. ALG acceptance

**Drift Verification:**

$$\Phi_{\text{ACL}}^{\text{new}}(p_{\text{ret}}) = \text{DENY} \quad \wedge \quad \Phi_{\text{PATH}}^{\text{stateful}}(p_{\text{ret}}, s^*) = \text{ALLOW}$$

## 4. Security Implications

This theorem demonstrates that a finite set of architectural properties, ordered layer evaluation, non-commutative transformations, state admission without re-evaluation, and asymmetric translation domains—can produce conditions where policy intent is contradicted by actual path behavior.

**Evaluation View:** Let  $\Phi_{\text{ACL}}^{\text{new}}$  denote the ACL predicate in its native evaluation space (pre-NAT for outbound, post-NAT for inbound). The constructive denial refers to  $\Phi_{\text{ACL}}^{\text{new}}$ .

**State Semantics:** Upon admission, packets that match the recorded 5-tuple (or platform-defined related tuple) bypass new-connection ACL evaluation; the platform enforces tuple conformance rather than re-running the ACL.

**Scope of Helpers:** ALG effects apply only to flows admitted by the prior stages; helper-created related flows are checked by the platform's related-flow policy and do not imply ACL suppression for unrelated tuples.

**Inevitability Scope:** Whenever configuration yields  $\delta_{\text{drift}} \neq \emptyset$ , domain escape follows deterministically from the documented order of operations; otherwise the property is vacuously false.

### 4.1. Drift Mechanisms

The following recurrent scenarios illustrate distinct mechanisms by which policy drift can occur within ASA 9.x multi-layer processing:

- **NAT-before-ACL asymmetry:** Inbound traffic undergoes NAT translation before ACL evaluation, while outbound traffic applies ACL checks prior to NAT. This asymmetry enables translation states that bypass intended source/destination checks.
- **ALG post-ACL pinholes:** Application Layer Gateways (e.g., FTP, SIP, H.323) can dynamically open pinholes in downstream layers after initial ACL evaluation, permitting flows that would otherwise be denied.

- **Translation-induced downstream matches:** Address translation can create new address/port combinations that satisfy rules in downstream layers, enabling indirect traversal of restricted zones.
- **Connection table persistence:** Established connection states persist across ACL changes and are not subject to re-evaluation, allowing continued traffic flow from sources that would be denied if newly initiated.

## 5. Example Configurations for Empirical Validation

This section lists ASA 9.x configurations that demonstrate the drift conditions described in earlier sections. Each example can be reproduced in a lab and occurs in real enterprise deployments. They provide clear reference points for verifying the conditions in practice.

The following ASA 9.x configurations instantiate the formal drift condition

$$\Phi_{\text{ACL}}^{\text{new}}(p) = \text{DENY} \quad \wedge \quad \Phi_{\text{PATH}}(p, s) = \text{ALLOW}$$

demonstrated in Section 3. All configurations are directly aligned with the documented processing order and feature set in Cisco ASA Series General Operations CLI Configuration Guide [4,5].

### 5.1. Configuration 1: Inbound Static NAT with Deny ACL

- **NAT rule:**  $\tau_{\text{NAT}}(192.168.1.100) = 203.0.113.100$  (static, bidirectional)
- **ACL rule:** `access-list OUTBOUND deny tcp any 192.168.1.0/24 eq 22`
- Processing order: Inbound traffic undergoes NAT before ACL evaluation (per [5])

From an outbound seed  $p_{\text{seed}}$  permitted by the ACL pre-NAT, ASA state table entry  $s^*$  allows  $p^*$  from the external host to traverse despite  $\Phi_{\text{ACL}}^{\text{new}}(p^*) = \text{DENY}$ .

### Configuration 2: Outbound ACL with NAT and ALG Enabled

- **ACL rule:** `access-list OUTBOUND permit tcp any any eq 21`
- **NAT rule:** Dynamic PAT for all outbound traffic
- **ALG:** FTP inspection enabled ( $\tau_{\text{ALG}}$  opens pinholes per [4])

ALG-induced pinholes ( $\tau_{\text{ALG}}$ ) bypass ACL reevaluation for related inbound sessions, producing identical drift conditions in multi-connection protocols.

### Configuration 3: NAT Creating Downstream ACL Match

- **NAT rule:** Maps a non-matching internal address to one that matches downstream ACL criteria
- Demonstrates  $\tau_{\text{NAT}} \circ \Phi_{\text{ACL}} \neq \Phi_{\text{ACL}} \circ \tau_{\text{NAT}}$  (see Theorem 3.1)

### Configuration 4: State Table Skips ACL Reevaluation

- Expired ACL entries not applied to existing  $s^*$  in state table
- $\Phi_{\text{STATE}}(\tau_{\text{NAT}}(p^*), s^*) = \text{VALID} \Rightarrow \Phi_{\text{PATH}}(p^*, s^*) = \text{ALLOW}$

**Note:** Each configuration above is both reproducible in a minimal ASA 9.x lab and observed in production enterprise deployments. They are not hypothetical; they map one-to-one with the symbolic model in Sections 3 and 4.

## 6. Conclusion

This paper has demonstrated, through formal predicate definitions, packet transformation modeling, and theorem-based reasoning, that certain ACL/NAT/ALG orderings in Cisco ASA 9.x can create conditions where the effective packet path differs from the intended access policy. Under specific but realistic configurations, a packet that would be denied if evaluated in isolation can be permitted through established state traversal, resulting in a verifiable policy drift condition. The analysis is grounded in vendor-published packet flow and ALG documentation, without the use of active probes



or interaction with live systems. While the case studies focus on ASA, the mathematical model applies to any system where translation, access control, and application-layer state are evaluated in differing sequences. This represents a policy verification limitation inherent to such architectures, rather than a misconfiguration of any single deployment.

The work highlights a verification gap in the ability to formally prove that a firewall enforces its stated intent under all valid configurations matching these criteria. Addressing this gap will require vendors and auditors to account for state table persistence, translation effects, and application-layer pinhole creation in the same unified verification model.

### 6.1. Legal Notice

This research was conducted solely through mathematical modeling and review of publicly available vendor documentation. No live systems were accessed, modified, or disrupted in the course of this work. The information is presented for academic study and defensive security purposes. Readers are reminded that applying these concepts in an unauthorized environment may violate applicable laws.

**Conflicts of Interest:** The author declares no conflicts of interest.

## References

1. R. Mahajan, T. Anderson, et al., “Lessons from Network Verification at Scale: The Batfish Experience,” 2021. Available: <https://batfish.org/>
2. A.-D. Brucker, L. Brüger, P. König, and B. Wolff, “Formal Firewall Conformance Testing: An Application of Test and Model-based Development,” in *Integrated Formal Methods (iFM 2015)*, LNCS, vol. 9196, Springer, 2015, pp. 265–280.
3. Cisco Systems, “Application Layer Gateway Processing in ASA,” Cisco Technical Documentation, 2023.
4. Cisco Systems, “ASA Application Layer Protocol Inspection,” Cisco Technical Documentation, 2023.
5. Cisco Systems, “NAT Order of Operations on ASA,” Cisco Technical Documentation, 2023.
6. Cisco Systems, “ASA Packet Processing Order,” Cisco Technical Documentation, 2023.
7. P. Srisuresh and M. Holdrege, “IP Network Address Translator (NAT) Terminology and Considerations,” RFC 2663, Aug. 1999.
8. B. Carpenter and S. Brim, “Middleboxes: Taxonomy and Issues,” RFC 3234, Feb. 2002.
9. S. Guha, B. Biswas, and B. Ford, “NAT Behavioral Requirements for TCP,” RFC 5382, Oct. 2008.
10. M. Zhang, Z. Mao, Y. Wang, and others, “Stateful Firewall State Exhaustion Attacks and Defense Mechanisms,” in *IEEE International Conference on Communications*, 2012.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.