From Implicit Execution to Deterministic Control

A Unified Control-Plane Architecture for Network and Execution Security

Skip Middleton
President, Sinteag Ventures, Inc.
skip@sinteag.com · sinteag.com · ivdcontrol.com
Patent applications pending, USPTO · April 2026

Abstract

Modern security failures share a common structural flaw: systems grant execution authority implicitly. Network infrastructures forward traffic based on reachability rather than intent. Application and AI-driven systems execute code and actions based on origin or context rather than admissibility.

This paper defines **implicit execution** as the dominant failure mode across both network-layer and execution-layer systems. It introduces a unified control-plane architecture that enforces deterministic admissibility prior to action. Two implementations are presented: **IVD-N** (Invariant Vector Defense, Network) addresses network-layer attack formation by enforcing upstream control over traffic behavior, and **IVD-ACP** (IVD, Admissibility Control Plane) addresses execution-layer risk by enforcing pre-execution admissibility of actions, code, and AI-generated intent. Together, these systems demonstrate that security can be reframed from detection and mitigation to bounded, pre-authorized execution.

1. The Structural Failure: Implicit Execution

Across domains, modern systems exhibit the same assumption: if an input reaches the system, it is allowed to act.

- In networks: Packets that satisfy routing conditions are forwarded. At scale, this enables distributed attacks to form before mitigation occurs.
- In software systems: Code delivered through trusted channels is executed during installation or runtime. Supply chain attacks exploit this assumption.
- In AI-driven systems: Model outputs are translated into tool calls or system actions. The distinction between suggestion and execution is often absent.

These are not isolated issues. They are manifestations of the same architectural flaw: execution authority is granted implicitly.

2. Comparative Failure Analysis

The table below maps the structural failure mode, current mitigation approaches, and IVD response across both implementation domains.

Feature	Network Perspective (IVD-N)	Execution Perspective (IVD-ACP)
Primary Risk	Formation before control. Attack traffic aggregates before mitigation engages.	Action without admissibility. Untrusted input executes on delivery.
Failure Mode	Traffic aggregates before control mechanisms engage.	Trusted distribution does not guarantee safe execution.
Current Mitigation	Downstream scrubbing and filtering. Reactive.	Manual review or post-hoc detection. Probabilistic.
Core Flaw	Reachability is assumed to equal permission.	Delivery is assumed to imply execution authority.
IVD Response	Upstream Psi-vector pattern extraction; macro-object formation; BGP FlowSpec enforcement prior to aggregation.	Pre-execution admissibility evaluation; deterministic authority assignment; sandboxed or blocked execution.

3. Formal Definitions

Admissibility: A deterministic evaluation of whether a candidate input or action is permitted to influence system state. Admissibility is independent of source trust and is based on structural and behavioral properties of the input.

Authority: The bounded level of execution capability granted to an admissible input. Authority is explicitly assigned and enforced prior to execution. Examples include READ_ONLY, SANDBOXED, and QUARANTINED.

Invariant Vector (Psi-vector): A fixed-dimension representation of behavioral characteristics extracted from observed activity. In networking, this includes header-derived features such as timing, distribution, and protocol attributes. Psi-vectors are constant-cost to compute and enable comparison independent of source cardinality.

Macro-Object: A synthesized representation of a distributed phenomenon derived from clustering related Psi-vectors across observation points. A macro-object captures the structural signature of an attack or coordinated behavior and serves as the unit of control.

Deterministic Control: A property of a system in which identical inputs under identical policy and state conditions produce identical classification and enforcement outcomes. Determinism excludes probabilistic scoring and ensures repeatability, bounded state, and auditability.

4. The Control-Plane Doctrine

Execution must be decoupled from arrival. All inputs must be treated as untrusted until evaluated. Authority must be assigned explicitly before any action is performed.

> ***No input is allowed to act without passing through a deterministic admissibility layer.***

This applies equally to packets in a network and commands in an execution environment. In networking, control planes determine how traffic should be handled. In execution systems, an analogous control plane must determine whether actions should be allowed to execute at all.

5. IVD-N: Upstream Control of Network Behavior

IVD-N implements deterministic control at the network layer. Edge Telemetry Sensors (ETS) extract Psi-vectors from packet headers without maintaining per-flow state. Regional Synthesis Nodes (RSN) correlate anomalies across observation points, and a Global Coordination Engine clusters related Psi-vectors into macro-objects representing attack structures.

Control is expressed as bounded policy derived from macro-object invariants and propagated upstream using Border Gateway Protocol (BGP) Flow Specification (FlowSpec).

Key properties:

- Behavior-based classification, not source-based filtering.
- Control applied prior to large-scale aggregation.
- Bounded rule count independent of attack size or source cardinality.
- No route instability in the base routing table.
- At TRL-6 (Technology Readiness Level 6), FlowSpec actions are logged for validation rather than injected into live forwarding paths.

IVD-N FlowSpec Control Expression (Logged, TRL-6)

A representative macro-object derived from a high-rate TCP SYN pattern produces the following policy expression:

```
match {
  destination 203.0.113.0/24;
  protocol tcp;
  tcp-flags syn;
}
then {
  traffic-rate 0;   // logged; not injected at TRL-6
}
```

In TRL-6 validation, this rule is generated deterministically from clustered Psi-vectors and logged as a policy recommendation without being injected into live forwarding paths.

6. IVD-ACP: Pre-Execution Admissibility Control

IVD-ACP applies the same doctrine to execution. All candidate actions are intercepted prior to execution, normalized into a structured representation, and evaluated against policy. Authority states are assigned deterministically.

Core Authority States:

- READ_ONLY: No system state can be modified.
- SANDBOXED: Execution occurs in a controlled, isolated environment.
- QUARANTINED: Execution is denied entirely and the input is retained for inspection.

This model applies uniformly to AI-generated actions, software packages, scripts, and API-driven commands. Execution is no longer implicit; it is conditional and bounded.

IVD-ACP Decision Trace (Validated Run)

Representative audit log entry from a validated run against a live agent environment:

```
{
  "ts": "2026-04-02T15:47:22Z",
  "principal": "lab-operator-1",
  "action_type": "config_change",
  "target": "test-system",
  "requested_operation": "change",
  "parameters": { "setting": "demo_flag", "value": true },
  "decision": { "authority": "SANDBOXED", "reason":
"state_modifying_operation" },
  "execution": { "mode": "sandbox", "status": "executed" }
}
```

This demonstrates input normalization, deterministic classification, and full auditability. Repeated identical inputs under identical policy conditions produce identical authority assignments.

7. Validation and Application

Validation in a Live Agent Environment

A live agent was tasked with generating a configuration change request. Under IVD-ACP, the request was intercepted and evaluated rather than resulting in direct system modification. The system assigned SANDBOXED authority and executed the action in an isolated environment, with a full audit record of the request, the classification decision, and the execution outcome.

Application to Software Supply Chains

Under a control-plane model, package contents are treated as untrusted input regardless of distribution channel. Execution attempts are intercepted and evaluated prior to execution. In the Axios npm compromise (March 2026), malicious code entered through a trusted distribution path and executed automatically. Under IVD-ACP, that behavior would have been intercepted at the

execution boundary and either blocked or redirected to sandbox execution. The compromise would not have translated into impact.

The critical distinction: the package did not need to evade detection. It only needed to be trusted long enough to execute. IVD-ACP removes that assumption.

8. Unified Architecture

IVD-N and IVD-ACP implement the same control principle at different layers. IVD-N prevents uncontrolled aggregation of malicious traffic before it reaches a target. IVD-ACP prevents uncontrolled execution of actions and code once an input has been delivered.

Shared properties across both implementations:

- Removal of implicit trust from the execution or forwarding path.
- Deterministic pre-action evaluation enforced at the control plane.
- Bounded control state under scale, independent of input cardinality.
- Full auditability of every classification decision.

9. Determinism as a Security Primitive

Determinism is not a performance characteristic; it is a security property. A deterministic system ensures:

- Identical inputs produce identical decisions under identical conditions.
- No reliance on probabilistic scoring or shifting heuristic thresholds.
- Bounded resource consumption independent of input scale.
- Reproducible outcomes for validation, audit, and policy enforcement.

Deterministic control enables reliable policy enforcement, stable behavior under stress, and verifiable security guarantees. This is the property that makes the architecture suitable for federal evaluation and critical infrastructure deployment.

10. Implications and Conclusion

This model requires a shift in system design assumptions:

- Execution authority is no longer inherited from source or context.
- Systems must tolerate delayed or conditional execution.
- Audit logs become decision records rather than event traces.

Conclusion

The dominant failure mode in modern systems is implicit execution. Networks allow traffic to aggregate before control engages; software systems execute code upon delivery; AI systems translate output directly into action. A deterministic control-plane architecture eliminates this assumption. IVD-N and IVD-ACP demonstrate that control can be enforced prior to aggregation

and execution across both domains. Security, in this model, is not reactive; it is enforced at the boundary where action would otherwise occur.

Only admissible actions are allowed to execute.

Sinteag Ventures, Inc. · Clyde, NC / Coral Gables, FL · sinteag.com · ivdcontrol.com