Proof of selective defense

Three distinct test tracks show what HoneyMesh did in a disposable federation lab, an isolated Node integration, and a bounded live Exchange canary. Each claim below is limited to the traffic and environment tested.

Track	Observed result	Boundary
Federation	One manual IPv4 ban reached three hosts; 15 host/traffic cases stopped replying while controls replied.	Disposable AMD64 lab, five traffic cases.
Node integration	A signed app event caused a selective IPv4 packet drop; an isolated matrix also passed IPv6 and release/expiry checks.	Disposable veth, separate reference app; no production multi-app pilot.
Live Exchange	Exact test-source IPv4 and IPv6 HTTPS paths stopped replying after the agent recorded kernel application; controls stayed reachable.	Short public canary, not sustained load or general attack detection.

The local app-to-agent ban in the live canary was **not** forwarded to other networks. The federation test used a separate, operator-issued manual ban.

A manual ban crossed a trusted mesh

An operator banned a test source on host A. Signed sharing reached trusted hosts B and C, which recorded the ban as a remote mesh action. Before the ban, the source reached all three hosts. Afterward, it received no reply on the five tested cases. A separate unbanned source continued to receive replies.

Traffic	Before: test source	After: same source	Control after ban
TCP 80	3/3 replied	3/3 no reply	3/3 replied
TCP 8448	3/3 replied	3/3 no reply	3/3 replied
UDP 53	3/3 replied	3/3 no reply	3/3 replied
UDP 9999	3/3 replied	3/3 no reply	3/3 replied
ICMP echo	3/3 replied	3/3 no reply	3/3 replied

The client probes used a bounded timeout; 'no reply' is not a packet-capture count. The final result followed a corrected probe and fresh ban after an initial script labeling error.

What this establishes

Selective manual-ban propagation and enforcement across the five tested IPv4 cases on three AMD64 hosts. It does not measure automatic detection, propagation speed, other hosts, or production traffic.

Application evidence reached the kernel

A separate plain Node service imported the HoneyMesh Node Guard package. In an isolated test, repeated invalid signatures from a designated source caused one signed event. The agent validated the event, applied a generic-XDP source rule, and the source lost ICMP replies while a control source kept replying. Switching to TEST removed the rule and restored reachability.

A broader isolated matrix

The exact agent binary used on the Exchange also passed selective IPv4 and IPv6 veth drops, unlisted-source exclusion, loopback rejection, TEST-mode release, and a short rule's natural expiry. The public interface was not used in these isolated runs.

Trust boundary

The Node app sends bounded evidence over a protected local socket. The agent decides whether to apply a local packet rule according to mode, source scope, and policy. The current production bridge binds one app identity to one socket credential; a second app has not been attached to the production agent.

The live site stayed reachable to controls

A disposable independent client supplied one exact IPv4 and one exact IPv6 test address. The Exchange machine API returned 40 unauthenticated 401 responses for each family. The agent then recorded a kernel-applied local rule; subsequent HTTPS probes from that test source received no HTTP response. Independent controls continued to receive HTTP 200 over both families.

Recovery check	Observed result
Operator release	Switching the agent to TEST released each rule; the affected client reached the site again.
Natural expiry	A separate 120-second IPv4 rule expired; the test client returned to normal site and API responses.
Final state	Agent TEST, app packet auto-blocking off, public XDP detached, zero agent bans, both services healthy; disposable client deleted.

The application did not return 429 during the public bursts: the packet rule was applied after the 40th 401, so the next test-source requests received no HTTP response. Dedicated XDP drop counters and packet capture were not collected.

How the results were checked

Before: confirm each test and control path reaches the same target. **Action:** issue one manual federation ban or trigger a bounded app event. **After:** compare the affected source with an independent control and read agent decision state. **Recovery where tested:** release or wait for expiry, verify reachability, and return the agent to TEST. Disposable lab resources were removed after the tests.

Evidence kept privately

The detailed operator record retains test timestamps, command outputs, agent audit rows, checksums, rollback actions, and cleanup results. This public edition uses anonymized host labels and summarized probe results. It does not publish source addresses, cloud identifiers, management paths, credentials, signing material, full request headers, or database contents.

Public verification

The proof page links to nine sanitized summaries. Each was generated only after its corresponding private test record passed status and five-case result checks. Those summaries show outcomes, not the underlying command transcript.

Claims we have not made

These tests do not establish line-rate throughput, native or offloaded XDP, detection accuracy, false-positive rates, shared-NAT safety, CDN-fronted attribution, high availability, multi-process Guard state, or sustained-load behavior. The public canary tested HTTPS no-reply for exact test sources; it did not repeat the full TCP/UDP/ICMP lab matrix on the live site.

What comes next

A separately credentialed second-app production pilot, per-app rate isolation, fresh-host install and rollback, exact-artifact vulnerability review, dedicated packet counters, performance measurements, and broader network coverage remain release gates. Optional exchange or direct-peer sharing of app-originated bans requires separate operator opt-in and end-to-end tests.

Current operating state

After the canary, the Exchange agent was returned to TEST, automatic packet blocking was disabled, the public XDP program was detached, and the site was verified over IPv4 and IPv6. Live status can change; consult the current activity view for present operation.

Method note

All results are dated 29-31 October 2026. This publication is a bounded proof of concept, not a general security guarantee or certification.