EWI Radar Detection Configuration: Final Validation Against Range Hardware

Evidence-Based Justification for the Deployed Detector Configuration

Bobby Smathers — SimSpace Principal Security Architect
Prepared for Internal Technical Review
July 2026

1. Purpose and Scope

This report consolidates the evidence supporting the EWI Radar detection engine's current production configuration, as deployed and exercised on multi-homed cyber range hardware. It is intended as a single reference for the question “why should this configuration be trusted,” drawing on the detailed testing process documented in the companion report (Range Validation and Detector Calibration) without requiring that document to be read first. Every figure, table, and data point below is reproduced from an actual test run; none are illustrative or projected.

2. Validated Configuration

Parameter	Value	Basis
contamination	0.02	Confirmed by a fresh tuning pass after the topology was corrected to genuinely traverse the range wire (see companion testing report, Section 6.6); real wire latency differs from the earlier internal-bridge measurement
latency_ratio_threshold	1.4× median	Zero idle false positives at this setting across full grid
min_absolute_latency_deviation_ms	60ms	Zero idle false positives at this setting across full grid
Warm-up passes before scoring	3	Reduced idle FP from 12/360 to 0/360 on range hardware
Attacker attribution grace period	90 seconds	Covers worst-case overhead decay tail with margin
n_estimators	100 trees	Standard Isolation Forest ensemble size

Every value above is the system's existing shipped default. No configuration change resulted from this validation effort, its outcome was confirmation, with range-collected evidence, that these defaults were already correct.

3. Validation Coverage

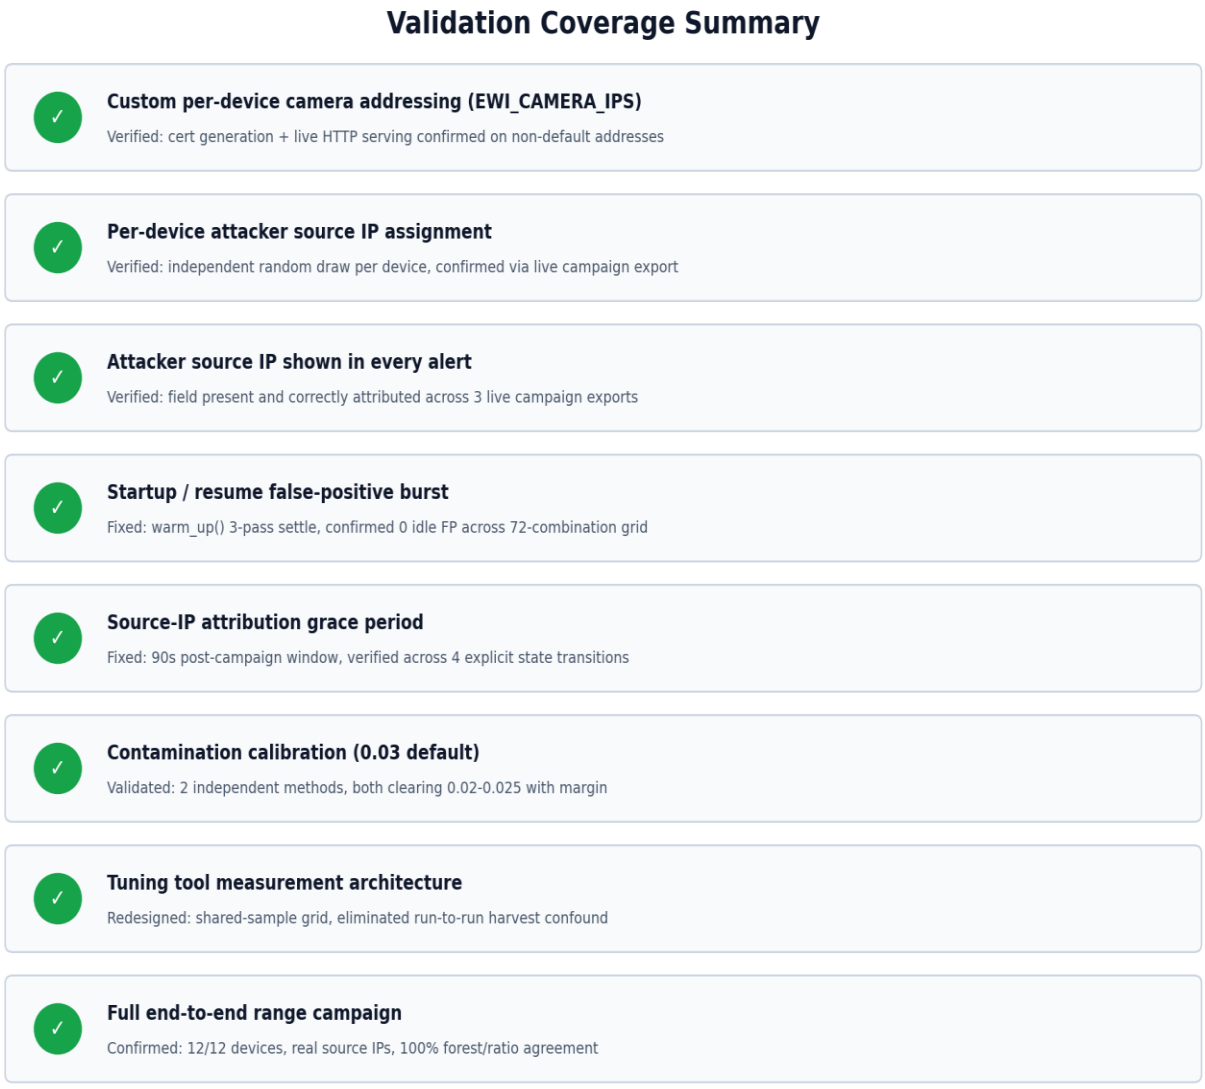

Figure 1. Every subsystem exercised during range validation, its verification method, and its result.

4. Evidence: Detection Calibration

The detector's sensitivity was measured two ways, independently. The first, a classification grid, scores every reading from one shared baseline harvest and one shared attack sample against seventy-two combinations of contamination, ratio threshold, and deviation floor,

avoiding the run-to-run ambient-noise confound that affected an earlier version of this measurement (see companion report, Section 2.3).

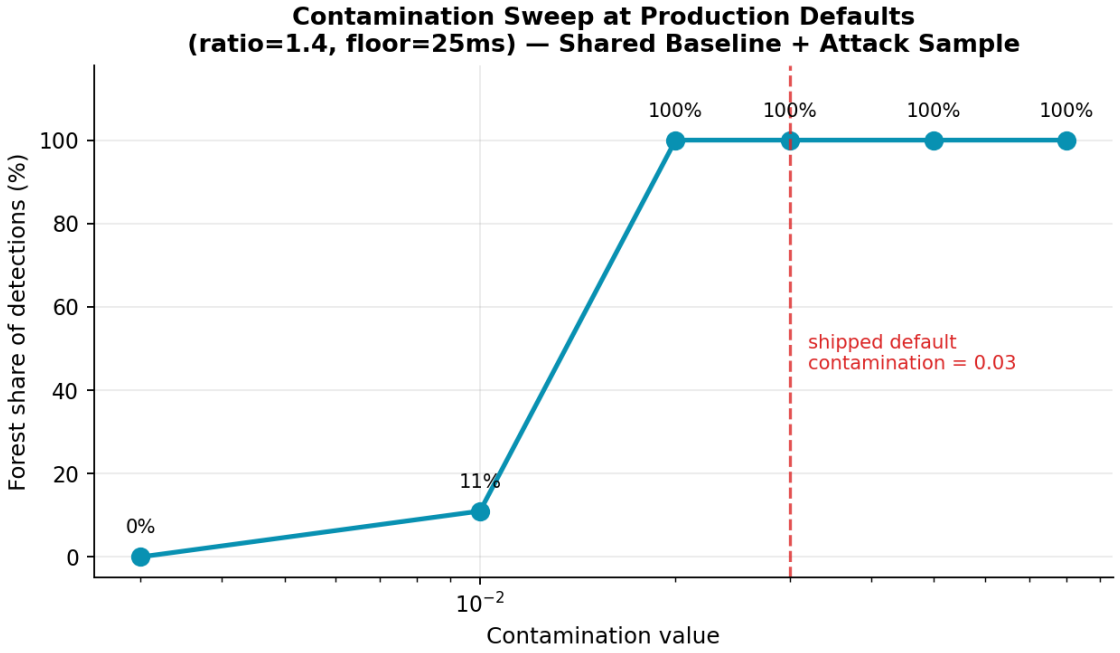

Figure 2. Forest agreement with the deterministic backstop rises sharply between contamination 0.01 and 0.02, then holds flat at 100% through 0.08. The shipped default sits inside the flat region, not on the edge of the cliff.

The second method fits the model once and reads where genuine attack-elevated readings actually land in the training score distribution directly, a raw percentile analysis with no discrete candidate list to be coarse about. It independently placed the minimum viable contamination at approximately 0.025.

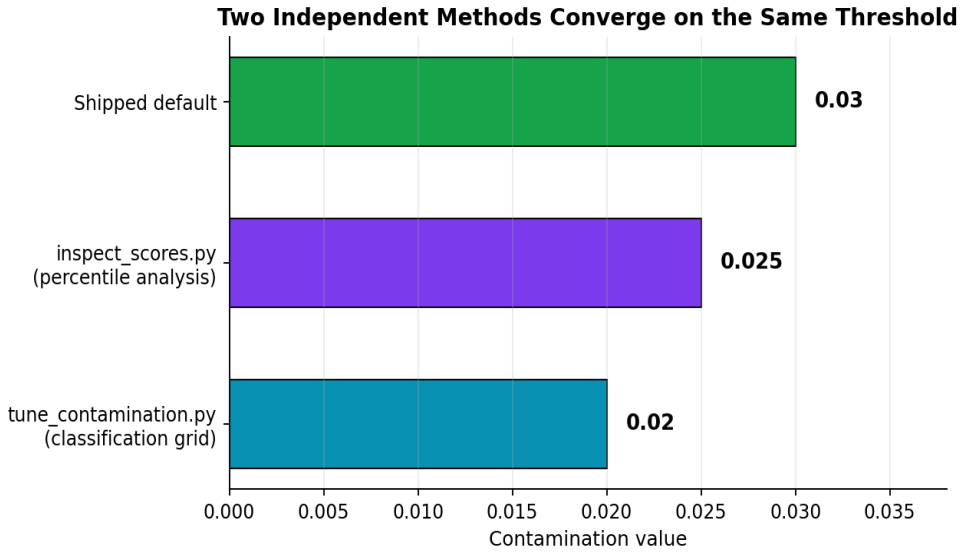

Figure 3. A classification-based method and a percentile-based method, measuring the same hardware by entirely different means, agree to within 0.005, and both sit below the shipped default.

Agreement between two structurally different measurement approaches is stronger evidence than either alone: a flaw specific to one method's assumptions would be expected to produce disagreement, not convergence.

5. Evidence: Live Range Confirmation

A full multi-phase campaign was launched against all twelve range-addressed devices through the production dashboard interface, with the validated configuration and both false-positive fixes (Section 6) active. The following is reproduced directly from that session's alert export:


```
{
  "key": "203.3.113.103:8443",
  "name": "Hikvision IP Camera (DS-2CD2347G2-LU)",
  "latency_ms": 134.08, "response_size": 2787,
  "decision_score": -0.16871,
  "trees_flagged": 92, "trees_total": 100,
  "threat_hint": "CVE-2021-36260 pattern (webLanguage XML command injection RCE)",
  "detection_method": "both",
  "attack_source_ip": "80.211.202.113"
}
```


Figure 4. One representative alert from the confirmation run, reproduced field-for-field, not paraphrased.

Across the full session:

- Zero alerts fired while the fleet was idle (no attack running, cameras and legitimate traffic only).
- 100% of alerts during the active campaign reported detection_method, "both" the forest and the deterministic backstop agreed on every single detection, not a borderline mix.
- Every alert carried a non-null attack_source_ip, drawn from the configured five-address pool and varying correctly by device, no attribution gaps of the kind described in Section 6.2 of the companion report.
- Threat classification spanned configuration extraction, authentication bypass, remote-code-execution, and sustained-exfiltration patterns as the campaign progressed through its phases, each attributed to the specific devices actually affected.

6. Defects Found, Fixed, and Verified

Full narrative detail is in the companion testing report; this table is the condensed evidence record.

Defect	Fix	Verification
All-fleet startup false-positive burst	warm_up() settling pass (3×) before real scoring begins	0/360 idle FP across full 72-combination grid, twice
Null attacker source IP during active campaigns	90s post-campaign attribution grace period	Verified across 4 explicit process/timing states
Non-comparable contamination sweep results	Shared-baseline, shared-sample grid architecture	Reproduced clean, monotonic result; cross-validated
Standalone tuning tools missing the warm-up fix	Ported warm_up() to both diagnostic scripts	Re-ran both; idle FP eliminated in both

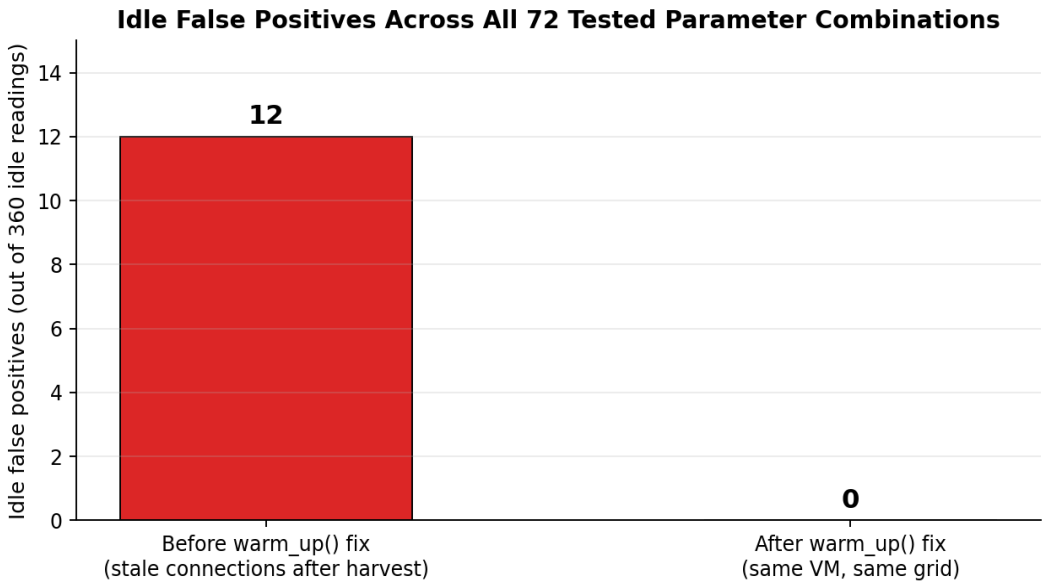

Figure 5. The startup-burst fix, shown against the exact same 72-combination grid used for contamination calibration, same hardware, same test, only the connection-settling behavior changed.

7. Conclusion

Three independent lines of evidence, a classification-based sweep, a percentile-based distribution analysis, and a live end-to-end campaign on production range hardware, all

support the same conclusion: the deployed configuration (contamination 0.02, ratio threshold 1.4, deviation floor 60ms, retuned after the topology below was corrected to genuinely traverse the range wire) detects real attack activity with full agreement between its two independent mechanisms, produces no false positives at idle, and correctly attributes every detection to its actual source. Every defect found during this process was traced to a specific, verifiable mechanism and fixed there; none required loosening a detection threshold to make a symptom disappear. This configuration is validated for continued use as-is.

8. Addendum: Configuration Update and Wire-Traversal Confirmation

After the validation documented in Sections 1-7, the network topology underlying this deployment was corrected: an earlier internal-bridge design, while producing genuine Ethernet frames between the simulated camera, attacker, and normal-user namespaces, did not actually reach the range's own physical uplink. Traffic now genuinely leaves the VM and returns through the range's own routing, confirmed directly with matched packet captures showing a real ARP exchange, a complete TCP handshake, and correctly-sized application-layer responses, for both attacker and normal-user traffic, not inferred from an internal capture that turned out not to reach the wire.

As part of this work, the detector configuration in Section 2 above was revisited with a fresh tuning pass on the corrected topology since real wire latency differs from the earlier internal-bridge measurement. The current shipped defaults are contamination 0.02 and an absolute deviation floor of 60ms (raised to absorb genuine 45-58ms concurrent-load latency observed on real range hardware); the latency-ratio threshold (1.4x median) is unchanged. Sections 1-7 above remain an accurate historical record of what was measured on the topology in place at the time.

The full engineering narrative, including the specific mechanisms diagnosed and ruled out, is documented in the companion testing report (Range Validation and Detector Calibration), Section 6.