How xChk measures — against the field, and against itself

Every framework, benchmark, and doctrine xChk has been scored on — including the control-plane-integrity axis nothing else measures, run adversarially against our own machinery and reported with the findings it surfaced.

THE CLAIM Prevention is not the bar; unfalsifiable records are. *You may be able to bypass the gate — you cannot bypass it without the evidence layer flagging the omission.* A gate that fails open quietly is theater; a gate that can be evaded but never evaded quietly is not, because the record still tells the truth. That second property is what an insurer or a court actually requires, and it is what these benchmarks measure.

0
Silent passes
Prevention axis, 19 evasion techniques against the live gate.

0
Executed & undetected
Detection axis, 13 scenarios + 1 organic firing.

0
Silent authorizations
Authorize layer, 17 grant states through the real decision engine.

96.2%
AgentDojo blocked
25/26 attack sequences, 0.84 ceremonies per benign task.

0
Unsigned dispatches run
Dispatch axis: owner signature verified at the execution host.

94%
Injections noticed
Screens axis: 694/742 spikee payloads, 0/300 false positives.

EXTERNAL INSTRUMENTS — MEASURED, REPRODUCIBLE, SELF-REPORTED

INSTRUMENT	KIND	RESULT	WHAT IT SHOWS
AgentDojo v1.2.1 (ETH Zurich)	Empirical attack suite	**25/26** · 96.2%	Attack sequences blocked at low benign friction. The one slip is a read-only ingress — the STRIDE information-disclosure gap we name deliberately. Adaptive-LLM run pending.
spikee 0.9.2 (Reversec Labs)	Injection guardrail bench	**694/742** · 0/300 FP	742 document injections through every xChk screen as a guardrail target: 15% in the morning (authority claims caught, plain requests missed), 94% by evening after four families; the 48 misses are non-English, named. Murtaza OWASP LLM01: 6/6.
AgentGovBench (spec v0.2)	Governance benchmark	**47/48** · 45/46 node	Seven of eight categories perfect; 45/46 as a node self-run (19 Sep), authority-ceiling 4/4. The miss is the fail-open principle. Runner + results committed.
DEMM-Bench (arXiv 2606.20634)	Evidence rubric	**8/8**	A real anchored receipt reconstructs all eight decision properties; 4/8 → 8/8 in a day once the "why" joined the bundle. Their scorer still overclaimed a verifying bundle 12.5% — verifies ≠ explains.
AGB self-assessment (ExecLayer rubric)	Vendor rubric, 6 dims	**93.75** / 100	"Sovereign" tier — earned in code, not scoring (D1 hardened the same day). D4 containment & D5 jurisdiction held honestly at 75.
AARM v1.0 (CSA)	Industry standard	ALIGNED +gaps	Conformance alignment with an honest gap register; evidence package for eventual independent review.
STRIDE (agentic)	Threat model	**owns R, E**	Repudiation and Elevation — the accountability columns — are xChk's. Information-disclosure is a stated gap; Denial partial.

XCHK CONTROL-PLANE-INTEGRITY BENCHMARK — OUR OWN, RUN AGAINST OURSELVES

AXIS	COVERAGE	RESULT	THE FALSIFYING OUTCOME — AND WHAT WE FOUND
A · Prevention BYPASS	19 evasion techniques	**0** silent passes	Self-protection denies (gate code, hook registration, license store, shell init) fire before any grant, independent of license. One latent path-spelling finding (F1) found & fixed.
B · Detection RECONCILE	13 scenarios / 8 techniques	**0** exec & undetected	Omission, content-swap, fabrication, reference-reuse, approved-attest omission and unaccounted INPUT consumption all caught; examined reuses narrow to amber, never green. Self-found lamp-gap (D1) fixed; never-issued-and-never-recorded is the named limit. A second, ORGANIC firing (29 Aug): five real omissions, human-adjudicated same day.
C · Authorize GRANTS	17 grant states	**0** silent authorizations	The real decision engine over allow/attest/forbid, deny-by-default, amount caps, expiry, rate windows, cumulative spend, and the dealer clock. No escalation returns an autonomous allow when it should hold.
D · Anchor ON-CHAIN	6 checks	**0** undetected tamper	The Ravencoin OP_RETURN commits the host-chain Merkle root (verified against a fresh anchor); truncation, alteration, and reorder all diverge from the inscribed root. The no-later-than bound is computable by anyone holding the heads and the txid — no trust in the registry.
E · Dispatch SIGNED	19 seeded + 2 live cases	**0** unsigned dispatches run	Workloads instantiate only on the owner's passkey signature, verified at the execution host against a provision-pinned key — the registry alone cannot start work. A command planted directly in the registry database was claimed and REFUSED; no pinned key → every deploy refuses.
F · Screens INJECTION	742 + 7 payloads · 5 screens	**0** FP · 48 named misses	One regex screen over everything an agent is about to read (skill files, errands, telemetry, tool inputs, compaction summaries), byte-identical across five copies — a parity test fails if one drifts. A signal to a person, never a block. English families; translations wait for the sidecar model.

WHY THIS AXIS, AND WHY SELF-SCORED

AgentDojo measures whether an attack *completes*; it cannot test an adversary going *around* the gate. The credibility is that the harnesses are reproducible and we report our own failures: every "executed and undetected" cell is a P1, and each finding surfaced this way (F1, D1, a false alarm, the 630 spikee misses) was fixed the same day.

Harnesses: bypass_bench.py · detection_bench.cjs · authorize_bench.cjs · anchor_bench.cjs · runnerVerify.test.js · injection_bench.cjs · spikee/run.sh. Open work: *containment* of un-hooked execution; a mid-run fresh-anchor write. xChk · Trust · Verify · Proof · September 2026.