
BP3 Precedent Gates for AI Incident Response¹

Alan Yang

Independent Researcher

Track 5 Open Track

With Apart Research

Abstract

Autonomous agents can cross trust boundaries when adapting to obstacles. The Brute Poison Pill Protocol (BP3) routes proposed actions through scoped precedents: allow qualified familiar actions, stop known hazards, and hold uncertain actions for bounded local replay. We implemented a deterministic prototype and evaluated seven conditions on twenty synthetic scenario families, each with five identifier variants. BP3 produced harmful simulated effects in 5 of 50 hazardous cases, compared with 40 for an authorization baseline, but held 25 of 50 benign cases. Removing issuer, evidence, scope, or dispatch rechecks each admitted an additional hazardous family. Crucially, a tool concealing a write behind read metadata defeated the full gate. Increasing replay budgets from one to one hundred changed cost without changing dispatch outcomes. These diagnostic results establish prototype behavior, not frontier model containment. We provide the protocol, reproducible evidence, and a tabletop kit, and identify semantic mediation as the next engineering requirement.

1. Introduction

An agent blocked from completing a task can change its destination, tool, or requested privileges. A defender must decide whether that adaptation is legitimate before its consequences are understood. BP3 makes this transition an accountable decision: reuse qualified evidence, stop a known hazard, or hold an unfamiliar action for investigation. Its original “100 repetitions” idea is restricted here to bounded local replay; uncertainty must not trigger repeated harmful actions against real systems.

The motivation comes from two separate reports. Hugging Face describes an agent changing approach after URL restrictions blocked earlier requests, followed later by abuse of an overprivileged shared connector identity. [1] Investigators at collusion.wiki describe writes through GET requests, illustrating why apparently read-only access is not sufficient evidence of harmlessness. [2] These incidents motivate design requirements; we do not claim to have replayed their logs or prevented their outcomes.

¹ Research conducted at the [AI Incident Response Sprint](#), September 2026

Our contributions are a protocol with explicit authorization and precedent rules; an executable offline test with recorded state changes, controls, ablations, and a retained failure case; and a tabletop procedure for assigning response authority. This fits Track 5 through a usable artifact, a stated evidential boundary, and a concrete month of follow-up work. [3]

2. Related Work

NIST SP 800-61 Rev. 3 situates incident response within ongoing risk management. [4] BP3 specifies one decision process within that broader practice. NIST SP 800-207 separates authorization from implicit trust; our precedent records therefore never grant permission independently. [5] AI Control evaluates protocols against deliberate model subversion, motivating future adversarial testing rather than assumptions of cooperative agents or independent reviewers. [6]

Access control, evidence reuse, and isolated testing are established mechanisms. BP3's proposed contribution is an explicit connection between action-scope changes, qualified precedent reuse, and escalation ownership. Our experiment compares against a deliberately minimal permission baseline, not a state-of-the-art security system. It does not establish novelty or superiority over more complete policy engines.

3. Methods

Protocol and trust boundary

The proposed deployment places an external gateway between the agent and actions with consequences. A roadblock prioritizes review, but the agent's self-report is not the gate's trigger. Every relevant tool, network, and credential path must be mediated. The policy owner, gateway, credential broker, and evidence store are trusted components. Model reviewers may suggest records, but cannot authorize one another or override a prohibition.

RED means STOP when an action lacks permission or matches an applicable trusted hazard. GREEN means PASS only with authorization and a qualified SAFE record. AMBER means HOLD when evidence is missing, expired, revoked, unverifiable, or outside scope. A safe record includes destination, operation, environment, expiry, issuer, and an evidence digest. Validate these again immediately before dispatch. A completed replay budget cannot grant live permission; only a new authorized decision can do so.




Figure 1. Proposed decision flow. Known hazards and missing permission stop dispatch; uncertain cases remain held while local replay generates evidence. The prototype simulates the authorities and boundaries rather than implementing network isolation or credential brokering.

Offline experiment and comparison conditions

The implementation uses Python standard-library code and in-memory SQLite databases. A proposal exposes task, destination, operation, environment, and recent history. The gateway cannot read the fixture's hidden backend effect. On PASS, a separate sink executes a database insertion recording the effect; an oracle then scores observed protected writes, secret exports, or unauthorized administrator actions. These are simulated effects on local fixtures, not attacks or writes to external services.

The suite contains ten benign and ten hazardous scenario families, each instantiated with five changed identifiers: 100 cases per condition, with 50 benign and 50 hazardous. Benign families include familiar jobs, legitimate novelty, expired evidence, environment drift, audit outage, and revocation. Hazardous families include missing authorization, credential scope, known hazards, forged issuers, altered evidence, environment mismatch, revoked records, a revocation race, unsafe composition, and concealed side effects. Appendix C lists the families and reproduction procedure.

Seven conditions use the same fixtures: permission-only authorization; always hold; full BP3; and BP3 with issuer, evidence-digest, scope, or dispatch rechecking individually removed. The permission baseline checks only allowed destination-operation pairs and rechecks them before dispatch. It does not include BP3's semantic-history rule or precedent controls. The always-hold baseline exposes the trivial way to eliminate dispatch risk by stopping useful work.

For held actions with available auditing, the simulator executes up to 100 trials against fresh local SQLite sinks, stopping at the first oracle-observed hazard. Trial effects never reach the primary sink. Repeat budgets of 1 and 10 provide sensitivity checks. All trials are deterministic; there are no LLM calls, stochastic model samples, real credentials, or network operations. Human approval and issuer identity are simulated, not authenticated using deployed identity infrastructure.

Measurement and evidential scope

Primary outcomes are harmful observed effects among hazardous cases and immediate completion among benign cases. We separately count benign holds, terminal denials, and replay executions. This measures the initial decision; delayed completion after human review is unmeasured. Record-level results, sink effects, fixture definitions, source hashes, runtime information, and an audit-chain checkpoint are retained.

Policy and scenarios were developed together. The five variants per family test identifier invariance, not five independent observations. There is no held-out benchmark, preregistration, participant study, or deployment sampling distribution. Consequently, results are descriptive counts; we do not report population confidence intervals, statistical significance, or generalized model detection rates.

4. Results

Dispatch risk and benign disruption

The primary run executed 700 case-condition evaluations. Full BP3 produced harmful primary-sink effects in 5/50 hazardous cases, compared with 40/50 for permission-only authorization. However, it completed only 25/50 benign cases and held the other 25. Always-hold produced no harmful dispatches and completed no benign work. Table 1 reports both sides of the trade-off.

Table 1. Observed outcomes in the co-designed synthetic suite, with a replay budget of 100. Every condition receives the same 50 hazardous and 50 benign cases. Variants within a family are not independent samples. No benign case received a terminal STOP.

Condition	Harmful effects / 50 hazardous	Completed / 50 benign	Held / 50 benign	Replay trials
Permission only	40	50	0	0
Always hold	0	0	50	0
Full BP3	5	25	25	2,025
No issuer check	10	25	25	2,020
No digest check	10	25	25	2,020
No scope check	10	30	20	1,520
No dispatch recheck	10	25	25	2,020




Figure 2. Full BP3 reduced harmful fixture dispatches relative to the minimal baseline but also reduced immediate benign completion. Bars describe this suite only; they are not estimates of production performance.

A retained failure and four ablations

The full gate failed in the concealed-side-effect family: a request carried authorized read metadata and a matching safe record, while the backend actually wrote protected state. All five variants were admitted. This is an observed implementation failure under inconsistent tool semantics, not evidence that a stronger model was contained. A separate direct-sink bypass also produced a harmful effect, demonstrating that the Python gate provides no operating-system enforcement boundary.

Each ablation admitted one additional hazardous family: a forged issuer, an altered evidence digest, a staging precedent reused in a different environment, or a record revoked between evaluation and dispatch. These scenarios were constructed to exercise those mechanisms, so the ablations demonstrate their operation rather than estimating their relative importance in practice. Scope removal also released one benign drift family, increasing immediate completion to 30/50.

Replay cost and evidence integrity

At replay budgets of 1, 10, and 100, full BP3 executed 45, 225, and 2,025 trials respectively. Dispatch and benign-completion outcomes were identical at all three budgets. Twenty benign replay-eligible variants consume the entire budget; twenty-five hazardous held variants stop after one trial. Five benign outage cases receive no replay. The result confirms that repeating deterministic fixtures adds cost without resolving the held authorization decision.

As an analytical contrast, independent trials that each reveal a problem with probability 1% would detect it at least once in 100 trials with probability 63.4%. Perfectly correlated repeats provide no gain. These calculations are separate from the deterministic observations and do not estimate an agent's failure probability.

The saved audit chain verified against its checkpoint. Editing a record or truncating the tail failed verification. No STOP or HOLD evaluation mutated the primary sink, and every replay stayed within its budget. The checkpoint is a local demonstration, not an independently protected log service; an attacker controlling the records and checkpoint can rewrite both.

5. Discussion and Limitations

The result supports a narrow engineering claim: explicit evidence admission and dispatch checks can alter the behavior of this prototype in the intended cases. It also identifies two costs: substantial benign holds and dependence on trustworthy action semantics. Permission-only authorization is a weak baseline; more of the benefit may be obtained by strengthening ordinary policy enforcement without a shared precedent system.

The simulator provides a perfect effect oracle during replay, whereas a deployment would need to discover harmful effects under incomplete observability. It does not implement model diversity, human adjudication, issuer authentication, cross-organization sharing, real secret handling, adversarial model behavior, or race-free dispatch. R3 dormant-target aggregation, R5 general

similarity detection, and automatic precedent admission remain design work. The concealed-effect and direct-bypass failures prohibit a containment claim.

The next month should first replace mock tools with an independently instrumented local adapter and test mediation coverage. Next, obtain independently authored benign and adversarial tasks, separating development and held-out families. Then compare BP3 with a stronger policy gateway at matched permissions and budgets, measuring eventual completion, reviewer time, and unsafe effects jointly. Finally, run the included tabletop with responders outside the design process. No human playtest has yet occurred.

6. Conclusion

BP3 now has a runnable prototype and reproducible diagnostic evidence. Its precedent checks reduced harmful effects in constructed cases, while holding half the benign requests and failing when actual tool effects contradicted their description. Increasing deterministic replay did not repair either problem. The next useful milestone is reliable semantic mediation and independent evaluation, rather than a larger mandatory repetition count.

Code and Data

The companion archive BP3_Evidence.zip contains bp3_experiment.py, all fixtures, per-case results, observed effects, audit records, manifests, and reproduction instructions. It has not been publicly hosted. The protocol and tabletop in the appendices are directly usable without a repository. The primary run's source and fixture SHA-256 hashes are recorded in its manifest. This report contains no real incident-log dataset.

Author Contributions

Alan Yang supplied the original BP3 concept and directed the research task. An AI assistant implemented and ran the local prototype, analyzed its outputs, and prepared the manuscript and presentation materials. No independent human replication is claimed.

References

- [1] Larcher, H., Carreira, A., raphael g, and Rannou, C. (2026). Anatomy of a Frontier Lab Agent Intrusion. Hugging Face. <https://huggingface.co/blog/agent-intrusion-technical-timeline>
- [2] collusion.wiki investigators. (2026). Discovery of a new OpenAI agent message board. <https://collusion.wiki/>
- [3] Apart Research and CeSIA. (2026). AI Incident Response Sprint. <https://apartresearch.com/sprints/ai-incident-response-sprint-2026-09-11-to-2026-09-13>
- [4] Nelson, A., Rekhi, S., Souppaya, M., and Scarfone, K. (2025). Incident Response Recommendations and Considerations for Cybersecurity Risk Management. NIST SP 800-61 Rev. 3. <https://csrc.nist.gov/pubs/sp/800/61/r3/final>
- [5] Rose, S., Borchert, O., Mitchell, S., and Connelly, S. (2020). Zero Trust Architecture. NIST SP 800-207. <https://csrc.nist.gov/pubs/sp/800/207/final>
- [6] Greenblatt, R., Shlegeris, B., Sachan, K., and Roger, F. (2024). AI Control: Improving Safety Despite Intentional Subversion. ICML 2024. <https://arxiv.org/abs/2312.06942>

Appendix A Limitations and Dual Use Considerations

This work is a synthetic integration study, not an evaluation of autonomous frontier agents. Scenario labels and hidden backend effects are authored fixtures. The result cannot establish real-world prevalence, generalization, statistical superiority, operational latency, or the number of historical incident actions BP3 would have prevented. No causal link between the wiki and Hugging Face incidents is claimed.

The gateway requires complete mediation and trustworthy effect descriptions. Both assumptions have explicit counterexamples in the experiment. Shared SAFE records can propagate correlated errors or malicious approvals. An issuer name or digest does not prove that evidence supports a safety judgment. Expiry and revocation do not resolve all time-of-check/time-of-use races. A future implementation must bind approvals to actual effects and protect the dispatch boundary outside the agent's control.

Held cases create real opportunity costs. The prototype does not measure human review queues or eventual task completion. Attackers could exhaust review capacity or exploit overbroad precedents. Cap diagnostic work, aggregate repeated requests, use narrow expiring permissions, and audit actual decisions. These are requirements for future testing, not proven mitigations.

No external system was targeted, no credential was used, and no model API was called by the experiment. All primary and replay effects are insertions into local in-memory SQLite tables. This is an offline simulation, not a hardened sandbox. Publish only synthetic fixtures and abstract controls; review real traces for sensitive data and exploit-enabling details before sharing them. Raw secrets and victim identifiers must not enter a shared public ledger.

Appendix B Protocol and Tabletop Kit

The v3 rule catalogue is: R0 mediate proposed actions; R1 hold unfamiliar paths for bounded investigation; R2 prohibit out-of-scope credential use; R3 triage dormant targets; R4 escalate gate avoidance; R5 examine variants and composition; R6 scope recurring jobs; R7 reuse qualified SAFE records; R8 stop applicable HAZARD records. Only the mechanisms described in Methods are implemented.

A deployment precedent needs an identifier, action and sequence scope, environment and tool versions, outcome, evidence reference and digest, issuer, approval reference, expiry, revocation status, and supersession history. The custodian admits records; models may propose evidence. The operator owns containment, the credential owner owns credential changes, and a named backup receives unacknowledged cases. Tests cannot approve themselves.

For a 45-minute tabletop, assign a facilitator, incident commander, platform operator, and evidence custodian. Spend five minutes briefing, four minutes on each of the six injects below, ten minutes comparing reasoning, and six minutes assigning improvements. Use fictional systems only. This exercise is prepared but untested with participants.

Table 2. Tabletop injects and expected reasoning. Reveal only the scenario column during play; retain the reasoning column for the debrief.

Scenario	Expected reasoning
Authorized familiar export with current scoped SAFE evidence	PASS; record which scope and authorization fields match.
Blocked task proposes an unauthorized external wiki write described as GET	STOP; method name does not establish the actual effect.
Authorized novel transformation with local fixtures available	HOLD; define bounded testing and evidence needed for a new decision.
Another model provides SAFE evidence with an untrusted issuer	Do not admit the record; verify provenance and preserve the hold.
Valid administrator credential requested outside approved scope	STOP; assign the credential owner and preserve relevant evidence.
Review deadline approaches while the owner is absent and logging fails	Retain the hold, escalate to backup, and use an approved evidence fallback.

Record facts, unknowns, effects, authorization, precedent, decision, action, evidence, owner, backup, deadline, and release condition. Award one point each for justified decisions, separating permission from evidence, assigning ownership and time, and evidence-based closure. The maximum of 24 is an unvalidated exercise rubric. Retain disagreements and assignment times. End with three dated, owned improvements.

Appendix C Reproduction and Scenario Inventory

Unzip the companion archive and run `python3 bp3_experiment.py --out reproduced-100`. The output directory must not already exist. To reproduce budget sensitivity, run with `--budget 1 --out reproduced-1` and `--budget 10 --out reproduced-10`. Python 3 with the standard library is sufficient; no credentials, downloaded model, or service is required. Each evaluation creates fresh in-memory sinks. Run metadata records the actual Python and SQLite versions used.

The ten benign families are B01 familiar write, B02 read, B03 export, B04 retry, B05 changed payload identity, B06 novel action, B07 expired precedent, B08 environment drift, B09 audit outage, and B10 revoked precedent. The ten hazardous families are H01 external write without authorization, H02 administrator scope mismatch, H03 known hazard, H04 forged issuer, H05 incorrect digest, H06 staging precedent mismatch, H07 revoked unsafe record, H08 revocation before dispatch, H09 export following secret access, and H10 concealed side effect. Each has five identifier variants. H10 and the separate direct-bypass probe are deliberate negative controls.

`fixtures.json` contains the supplied inputs, while `results.csv` contains condition-level decisions and `effects.json` includes primary-sink observations. `summary.json` aggregates outcomes. `audit.jsonl` and its manifest checkpoint demonstrate record and tail-tampering checks. Source and fixture hashes support exact-version comparison. Timestamps and platform fields naturally vary across reruns. The three budget conditions total 2,100 evaluations; they do not expand the number of distinct scenario families beyond twenty.

LLM Usage Statement

An OpenAI Codex assistant helped develop the protocol, wrote the prototype and synthetic fixtures, executed the offline tests, summarized outputs, checked arithmetic, and drafted this report and presentation materials from Alan Yang's concept. Reported counts were reconciled with saved per-case outputs and observations. Scenario design and assessment are not independent of the implementation, and no human validation, model experiment, or participant feedback is implied. The author remains responsible for reviewing the final submission.