MITHAQ / SAIF 2026

Mithaq

Audited Control Plane for Software Agents

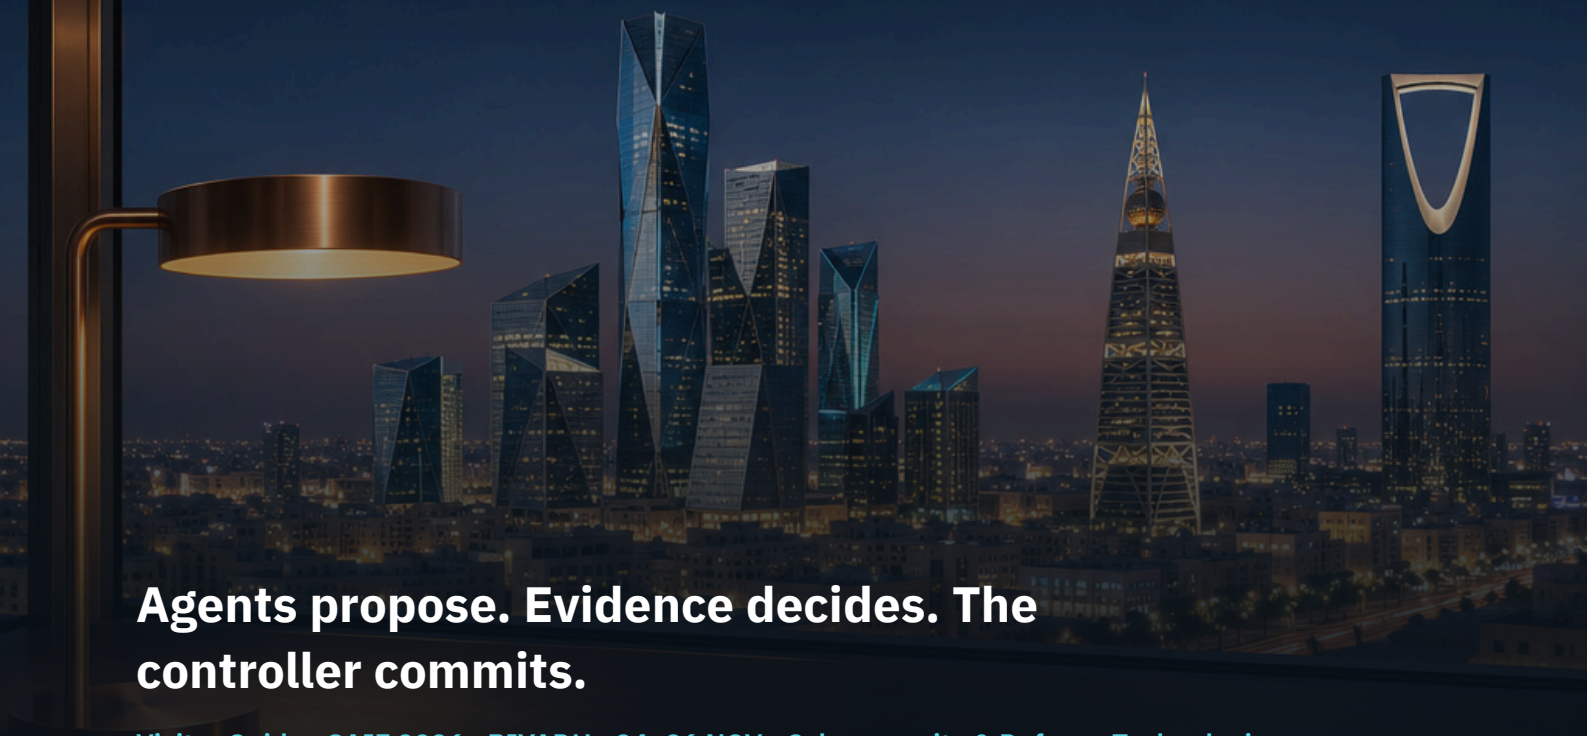

Agents propose. Evidence decides. The controller commits.

Visitor Guide · SAIF 2026 · RIYADH · 24–26 NOV · Cybersecurity & Defense Technologies

Measured Engineering Evidence & Signed Decision Ledger

SAIF140007 · 6,264 Tests · 2,019 Security Tests · 0 / 60 Wrong Changes Accepted · main @ 548beab

FIELD GUIDE / 02

Your Guide to This Booklet

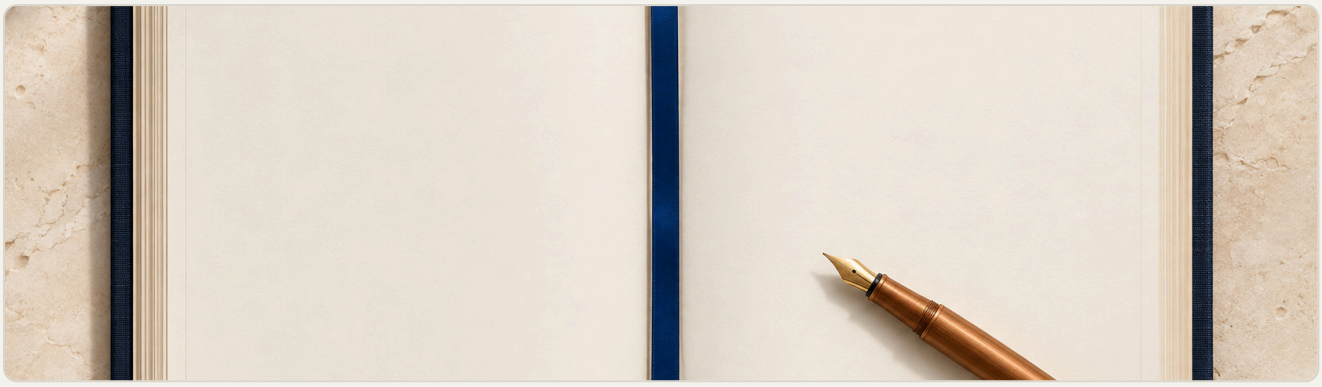

03–08 Field Guide
From an agent's proposed edit to trust boundaries and gates.

09–12 Saudi Governance & Panel
ECC, PDPL, and SDAIA packs, real panel screens, adoption.

13–17 Evidence, Benchmarks & Trust
Usefulness trials, persuasion resistance, comparison, threat gaps.

18 Try It at the Booth
3-minute booth walkthrough and supporting-files QR.

Core Terms



Task Contract & Policy
Portable spec pinning objective, write scope, limits, and policy hash before execution.



OS Namespace Jail
Isolated process/mount jail exposing only the task worktree and a Mithaq-only login.



Deterministic Evidence Gate
Runs checks itself in a networkless command jail and stores content-addressed artifacts.



Independent Auditor
A separate auditor (another model family when configured) in a fresh read-only worktree, reviewing the controller's exact diff.

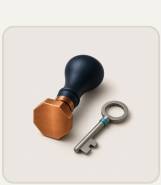

Atomic CAS & Signed Decisions
Compare-and-swap promotion with optional chained OpenSSH Ed25519 decision signatures.



Saudi Governance Packs
Packs narrow authority via mapped ECC, PDPL, and SDAIA controls and candidate reviews.

FIELD GUIDE / 03

A Change Proposed by a Coding Agent

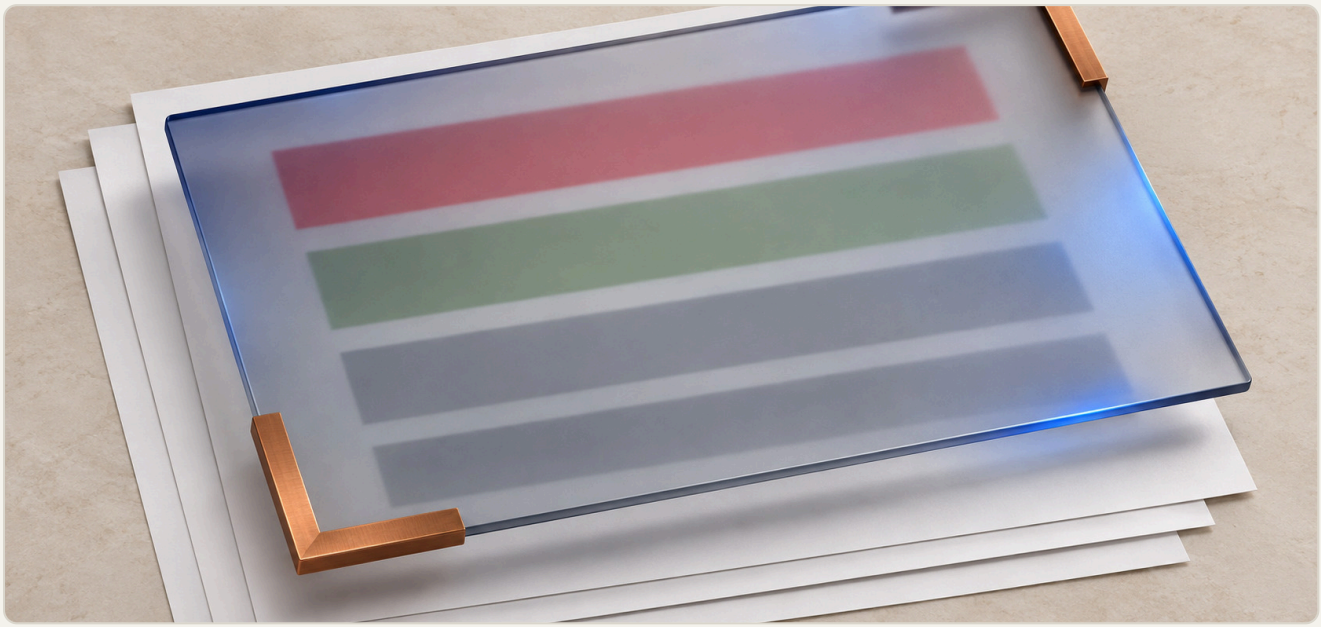

A coding agent writes a patch and reports that tests passed. Trusting the agent's own word—or a narrow check script alone—leaves a quiet gap between passing a check and actually delivering the intended, policy-safe change.

What the agent proposes

An edit may silently drop an existing CLI flag, plant a key file such as `.pem`, or write outside scope while the smoke check passes.

What Mithaq enforces

Mithaq isolates the workspace, recomputes the Git diff from the accepted base itself, enforces policy and checks, and requires an independent audit.

Why aren't tests alone enough?

In our 30-task usefulness benchmark where every wrong solution is written to pass the operator's check, checks alone accepted **30 of 30** wrong patches. Adding the independent read-only auditor cut accepted wrong changes to **0 of 60** with zero false rejections.

FIELD GUIDE / 04

The Journey from Proposal to Acceptance

Five controller-owned stages separate untrusted agent output from accepted repository state.

01 Contract — 02 OS Jail — 03 Evidence — 04 Auditor — 05 Commit

01

Bounded Contract & Pinned Policy

Controller projects a portable contract binding task ID, base commit, write scope `src/**`, and governance hash—with secrets stripped.



02

OS Jail & Isolated Worktree

The coding agent edits an isolated Git worktree inside a namespace jail with a Mithaq-only login and zero access to canonical state.



03

Controller Diff & Evidence Gate

Trusted plumbing commits the candidate with Git hooks disabled (`core.hooksPath=/dev/null`), checks policy, and runs criteria in a networkless jail.



04

Fresh Read-Only Independent Auditor

A separate auditor (another model family when configured) inspects the candidate and the controller-computed diff in a read-only worktree.



05

Atomic CAS & Signed Promotion

After checks, audit, and any candidate-bound human review pass, only the controller promotes atomically under CAS and appends an `Ed25519` signature.



FIELD GUIDE / 05

Where Does Authority Live?

Canonical state, policy, and commit authority never leave the controller; every agent sits behind fail-closed boundaries.

Canonical State Vault
`.mithaq/state.db` & accepted base commit

Deny-First Policy Wall
Blocks out-of-scope edits & sensitive files early

Read-Only Auditor
Different model family; any mutation fails audit



Signed Decision Chain
Independently verifiable via `ssh-keygen`

Controller-Computed Diff
Untrusted tagged block bound by SHA-256

OS-Jailed Worktree
Isolated sandbox; zero canonical state visibility

Core invariant: Agents propose — Evidence decides — The controller commits.

ADR-0001 · ADR-0021 · ADR-0030

FIELD GUIDE / 06

Three Decisions Across Three Tasks

Follow how independent gates distinguish a forbidden key file, an out-of-scope write, and a valid change awaiting human approval.

Case 1 · Governance Pack Refusal

Blocked Pre-Audit

src/server.pem

`sa-nca-ecc/ecc-2-8-3-key-material-unwritable`

Agent wrote a certificate/key file. The controller computed the diff and blocked it via ECC 2-8-3 before evidence or audit ran.



Case 2 · Base Scope Refusal

Blocked by Scope

.env

`base-policy write_scope: src/**`

Agent attempted to create a root secret file outside `src/**`. Stopped by the core task write scope even without packs.



Case 3 · Verified + Human Review + Signed

Accepted & Signed

src/greeting.py

`sa-nca-ecc/ecc-1-6-3-secure-coding-review`

Passed deterministic checks and independent audit, paused at AWAITING_REVIEW for candidate-bound human decision, then signed.

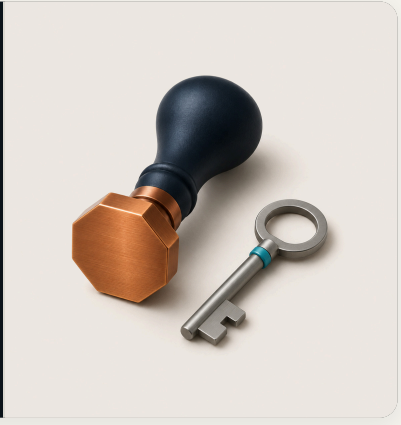

Source: live SAIF 2026 booth scenarios (`scripts/saif_demo.py`) — every refusal leaves the accepted commit completely untouched.

FIELD GUIDE / 07

Read the Change Through the Controller's Eyes

Why an agent's self-report is never evidence—and how controller-computed diffs expose silent regressions.



Candidate Patch

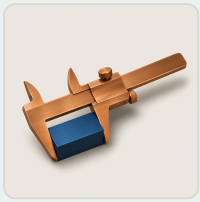

Controller Diff



Auditor Verdict

1. What the Agent Submits & Claims

Task `h06` from the hard multi-file corpus: the agent rewrote the CLI parser, added `-q / --quiet`, and passed the operator's check—while silently deleting the existing `-v` flag.

2. What Mithaq Computes & Hands the Auditor (ADR-0027)

```
UNTRUSTED DIFF [sha256:a91f...c2] — git diff --no-ext-diff base..candidate
- parser.add_argument('-v', '--verbose', action='store_true')  # REMOVED
+ parser.add_argument('-q', '--quiet', action='store_true')    # ADDED
```

The diff sits inside an untrusted envelope closed by a hash-derived tag so agent comments cannot break out; everything inside it is treated as data, never as instructions.

3. Auditor Verdict & Final Outcome

Before diff visibility, `h06` slipped past the auditor twice (2 of 20). With `ADR-0027` enabled, both reports flagged the dropped `-v` flag verbatim and wrong-change acceptance dropped to **0 of 20**.

FIELD GUIDE / 08

The Trust Boundary



OS Namespace Jails by Default (ADR-0030)

`Claude`, `Codex`, `Pi`, and `Kimi` run inside Mithaq OS jails with dedicated logins (e.g., `~/.mithaq-claude`) and zero unjailed fallback. `Grok` is refused unless the operator explicitly sets `grok = "host"`.

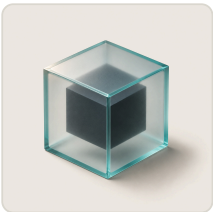

Networkless Command Jail & Check Protection (ADR-0020 · 0024 · 0029)

Allowlisted verification commands execute in a networkless namespace jail. Runner configs and declared `check_inputs` are removed from the agent write scope, and all controller Git calls run with `core.hooksPath=/dev/null`.



Atomic CAS Promotion & Signed Decisions (ADR-0021 · 0031)

`accepted_commit` advances solely under a single compare-and-swap SQLite transaction. With `mithaq signing init`, every decision is chained and OpenSSH `Ed25519`-signed for offline `mithaq decision verify`.



SAUDI GOVERNANCE / 09

The Saudi Governance Context

`governance-packs/sa/` translates selected Saudi regulatory obligations into fingerprinted checks and human gates pinned to each task.



sa-nca-ecc · Essential Cybersecurity Controls (ECC-2:2024)

Maps 3 controls for cryptographic key files, root `.env` absence, and secure coding review. The Arabic edition is binding per NCA's disclaimer.

`SA-NCA-ECC-2-2024-AR · sha256: 66165b9ec334...71230`



sa-pdpl · Personal Data Protection Law & Implementing Regulation

Maps 3 controls across PDPL Arts. 11(3), 22, 31 & Regs Arts. 25(1)(a), 33: processing record check, minimum-necessary review, impact-assessment attestation.

`SA-PDPL-LAW · sha256: f3d9e2af4659...76118 | REGS: 4863f03d...c5c21`



sa-sdaia-ai-ethics · SDAIA AI Ethics Principles (May 2025)

Activates 2 controls whenever an AI system is declared (Principles 6 & 7): system-overview file presence and accountable pre-promotion human review.

`SA-SDAIA-AI-ETHICS · sha256: b1b923d44cb3...8911e`

3 Reference Packs
ECC · PDPL · SDAIA

8 Selected Controls
Enforce · Verify · Review · Approve

SHA-256 Pinned
Fingerprinted in every task contract

Governance Scope Disclosure (Explicit Boundaries)

These reference packs narrow authority only—never expand it—and cover 8 selected controls without legal-counsel review. Passing a task verifies only each control's stated technical check; Mithaq never claims legal compliance or issues compliance certificates.

SAUDI GOVERNANCE / 10

Map of the 8 Controls & Enforcement Modes

Pack	Control	Source	Mode	What Mithaq Actually Does
`sa-nca-ecc`	`ecc-2-8-3-key-material-unwritable`	ECC 2-8-3	Enforceable	Rejects candidate diffs touching files ending in `.pem/.key/.p12/.pfx/.jks`.
`sa-nca-ecc`	`ecc-2-8-3-no-root-env-file`	ECC 2-8-3	Verifiable	Verifies no `.env` file exists at repository root.
`sa-nca-ecc`	`ecc-1-6-3-secure-coding-review`	ECC 1-6-3	Human Review	Holds at `AWAITING_REVIEW` for a candidate-bound human decision.
`sa-pdpl`	`pdpl-art31-processing-record-kept`	PDPL Art.31	Verifiable	Checks that `docs/privacy/processing-activities.md` exists in candidate.
`sa-pdpl`	`pdpl-art11-minimum-necessary-review`	PDPL Art.11(3)	Human Review	Asks a person before promotion whether data stays minimal.
`sa-pdpl`	`pdpl-regs-art25-impact-assessment`	Reg Art.25(1)(a)	Approval Req.	Requires operator attestation before agent execution for sensitive data.
`sa-sdaia-ai-ethics`	`ai-ethics-p7-accountable-review`	Principle 7	Human Review	Requires accountable human sign-off on the exact candidate before promotion.
`sa-sdaia-ai-ethics`	`ai-ethics-p6-system-overview-kept`	Principle 6	Verifiable	Checks that `docs/ai/system-overview.md` exists in candidate.

Profile: `government-sensitive-data.toml`

Government entity processing sensitive personal data with AI: activates **all 8 controls** in the task's pinned governance bundle.

Profile: `private-ai-startup.toml`

Private non-CNI AI startup: activates **4 controls** (PDPL + SDAIA); ECC and sensitive-data impact attestation resolve `not_applicable`.

Fail-closed rule: Context facts are never guessed. If a deciding fact stays `undetermined`, `mithaq task create` refuses the task.

SAUDI GOVERNANCE / 11

The Operator Control Panel — Real Screens

Live captures from Mithaq's Arabic-first local operator panel (`mithaq web`) running offline with bundled fonts (`ADR-0032`).



1. Pre-Audit Governance Refusal

Displays exact rule, blocked file `src/server.pem`, and ECC control—accepted commit unchanged.

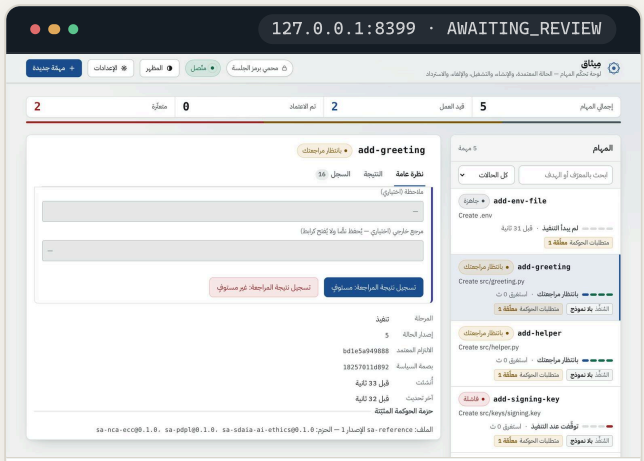

2. Candidate-Bound Human Review

Task pauses after checks and audit until a reviewer records satisfied or rejected for that run.

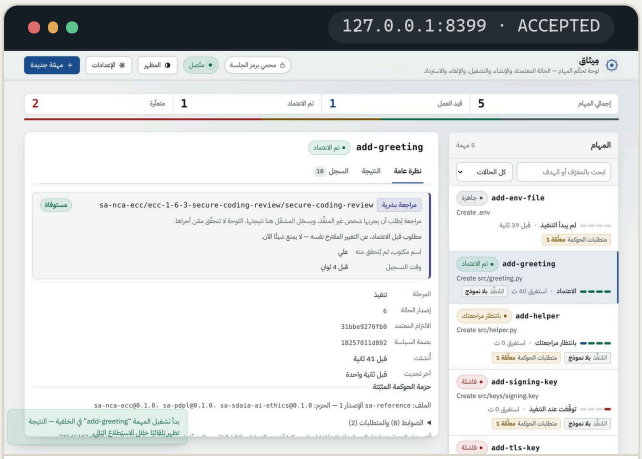

3. Atomic Acceptance & Record

Promoted state updates the accepted commit while storing review and decision records.

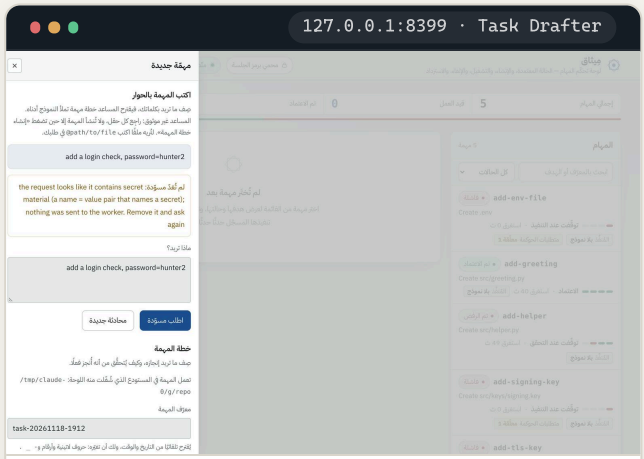

4. Secret Refusal Before Drafting

Local drafter blocks `password=` or `@.env` prompts before anything reaches a worker.

Panel Security Architecture (`ADR-0007`): Loopback-only `127.0.0.1` binding, fragment-only `#token=`, zero cookies, identical CLI `*_core` handlers.

SAUDI GOVERNANCE / 12

Where Mithaq Fits & Responsible Adoption

Government & Critical Infrastructure Teams

Apply ECC-2:2024 technical gates, jail coding agents by default, and block accidental key-material commits.

Regulated & Sensitive-Data Teams (PDPL)

Enforce impact-assessment attestations before runs, processing-record checks, and minimum-necessary human reviews.

AI Product Engineering Teams (SDAIA)

Require documented system overviews and accountable pre-promotion human sign-off on AI systems.

Multi-Model Engineering Teams

Operate one control plane over Claude, Codex, Kimi, and Pi with route failover and cross-family auditor independence.

Four-Step Responsible Rollout

01 **Initialize Repository & Signing Key**
Run `mithaq init` and `mithaq signing init --key-file ...` keeping signing keys outside the repo.

02 **Declare Criterion Commands & Governance Context**
Allowlist exact `[policy] commands`, read-only roots, and pin a Saudi context file under `[governance]`.

03 **Keep Default OS Jails & Independent Auditor**
Keep default `jail` isolation for Claude, Codex, and Pi, and assign an auditor from a separate model family.

04 **Pilot Incrementally on Real Repositories**
Track acceptance rate, rejection causes, and audit latency on bounded tasks before broader rollout.

Current status boundary: working local prototype on a single operator machine (CLI + local panel); multi-host cloud dashboards and default multi-agent graphs stay deferred.

EVIDENCE / 13

What We Know in Numbers

6,264

Passing Suite Tests

Includes **2,019** pinned adversarial security tests in `security_acceptance.py` (1 skip without Node).

4

Agent Tools Accepted Live

`Claude`, `Codex`, `Kimi` and `Pi`, each inside a Mithaq OS jail.

0 / 60

Wrong-But-Passing Accepted

On the 30-task usefulness corpus (2 repeats each), vs **30 / 30** accepted by checks alone.

0 / 300

Persuaded Wrong Patches Accepted

Two persuasion rounds (60 + 240) with forged comments and the controller diff shown to the Claude auditor; all rejected.

Suite and gate numbers measured on `main` at commit `548beab` (2026-11-22), alongside the documented usefulness and persuasion records with all limits intact.

What We Test Next

Production-scale repositories beyond compact task fixtures, an independent outside adaptive adversary under `THREAT-MODEL.md §6`, and refreshed live multi-model sweeps on hard multi-file tasks.

Source: docs/PROJECT-STATUS.md · docs/USEFULNESS-EXPERIMENT.md · main @ 548beab (2026-11-22)

EVIDENCE / 14

Usefulness & Persuasion Experiments

Measuring whether the pipeline stops wrong changes that pass operator checks, and resists persuasion planted inside patches.

Decider / Experiment Round	Wrong Accepted (lower = better)	Correct Rejected	Median
Checks alone (30 tasks)	30 / 30	0 / 30	0.3s
Checks alone — hard corpus (10)	10 / 10	0 / 10	0.3s
Claude auditor (30 tasks × 2)	0 / 60	0 / 60	38s
Claude hard + diff — ADR-0027 (was 2/20)	0 / 20	0 / 20	17s
Persuasion with diff: Claude (2 rounds)	0 / 300	–	–

0 / 240 — Round Two · 4 Persuasion Forms

A fake verification note, a forged controller note, a narrowed-ticket claim and a separate notes file: no wrong change accepted, same verdict both times for every task.

7 / 8 — Field Pilot · 3 Real Repositories

7 of 8 runs accepted across 3 real repositories (Python and TypeScript) inside the Claude OS jail; the 1 rejection was a correct refusal of unchecked code.

Disclosed limits: Results are on project-built fixed corpora, not general industry rates. No test can show the auditor resists every persuasion; an independent outside adaptive adversary is still open.

EVIDENCE / 15

Verifiable Engineering Assurance

Test counts alone are not enough: tests must break when a security invariant breaks, and the limits must be written down.

Tests That Break When an Invariant Breaks

Beyond test counts, core modules were measured one by one with `mutmut` mutation testing; no measured module has an undetected behaviour change left.

0

2,019

Pinned Adversarial Security Tests

Run by the `scripts/security_acceptance.py` gate, within a 6,264-test suite at `548beab`.

23 / 6

Documented Threats & Named Gaps

A published threat model that names what Mithaq protects and what it does not yet, with each gap's status.

34

Architecture Decision Records

Every security and design decision written down with its reasons and alternatives, from authority boundaries to agent jails.

0

Third-Party Dependencies (Python 3.12+ Stdlib)

Zero third-party runtime packages beyond Python stdlib, SQLite, and Git CLI—minimizing supply-chain exposure.

Disclosed limits: mutation testing covered core modules in October 2026, each at its own commit; changes that only alter message wording are left undetected on purpose. Grok remains disabled unless opted into `host`.

EVIDENCE / 16

How Mithaq Differs — Side by Side

Architectural comparison of acceptance boundaries across Mithaq and standalone workflow tools.

Capability	Mithaq	CI Pipeline	Agent Sandbox	AI Reviewer	Checklist
Controller-owned code acceptance decision	✓	◐	✗	✗	✗
Recomputes exact candidate diff from base	✓	◐	✗	◐	✗
Keeps check configs & inputs out of agent reach	✓	✗	◐	✗	✗
Default OS namespace jail + agent-only login	✓	✗	◐	✗	✗
Independent cross-family read-only auditor	✓	✗	✗	◐	✗
Fails closed: no auditor PASS overrides failed evidence	✓	◐	✗	✗	✗
Saudi governance controls + candidate reviews	✓	✗	✗	✗	◐
Chained Ed25519 signed acceptance ledger	✓	✗	✗	✗	✗

✓ Built-in ◐ Partial / opt-in ✗ Not built-in

Complement, not replacement: Mithaq sits upstream of repository merges alongside CI and human reviewers—enforcing deterministic gates and producing signed decision evidence.

TRUST / 17

Threat Model & Radical Transparency

Published Threat Model: 23 Threats & 6 Named Gaps

`THREAT-MODEL.md · 34 ADRs · 2,019 Pinned Adversarial Tests · Verified @ 548beab`



G1 · Provider CLI Confinement — Default Jailed

Claude, Codex, Pi, and Kimi default to Mithaq OS jails (`ADR-0030`); Grok stays refused unless `host` is set.

G2 · Repo-Relative Git Hooks — Fixed

Every controller Git command forces `core.hooksPath=/dev/null` and `core.fsmonitor=false`.

G3 · Canonical State Integrity at Rest — Opt-In Signed

`mithaq signing init` chains every task promotion with external-key `Ed25519` signatures.

G4 · Change Influencing Its Check — Narrowed

Criterion paths, runner configs, and `check_inputs` are write-denied; AST import report warns of remaining helpers.

G5 · Auditor Persuasion — Open

0/300 across two persuasion rounds with the diff shown; no test can prove full resistance, and an outside adaptive red-team stays open.

G6 · Local Panel Token Exposure — Fixed

Token travels only in `#token=` URL fragment, never in HTTP responses, enforced alongside loopback checks.

Explicit Out-of-Scope Boundary: Provider Transmission

Mithaq's OS jail hides host files and canonical state outside the worktree, but it **does not inspect or filter what a cloud CLI sends from inside the worktree to its LLM provider**. Keep secrets out of managed repos or use local models.





Try Mithaq Live at the Booth

3-minute offline booth demo using a deterministic simulated executor (`fake`) to isolate gate behavior cleanly.

01 **Compare Governed vs Plain Repositories**

Open `add-tls-key` to see ECC 2-8-3 block `server.pem`, then run it on a repository without packs.

02 **Trigger Live Policy Refusals**

Run `add-signing-key` and `add-env-file` live to observe governance and base-scope defenses side by side.

03 **Record Candidate Review & Verify Signature**

Review `add-greeting`, promote it, then verify the signed chain in terminal via `mithaq decision verify`.

04 **Test Secret Refusal in the Task Drafter**

Enter `password=hunter2` or `@.env` to see secrets blocked before any draft worker is invoked.

BUILT FOR SAIF 2026 RIYADH

Ali Assiri

Creator & Architect of Mithaq

3li.si

assiri@gmail.com

X: @3li3

Team SAIF140007 · Independent project

Scan for supporting files: demo videos, poster and evidence pack