PLATFORM PROPOSAL

AI and MCP Access Governance Platform

Governed intake, approval and evidence for AI tool, agent workflow and MCP server integrations.

STRUCTURED INTAKE · RISK CLASSIFICATION · HUMAN REVIEW · AUDIT EVIDENCE

This document presents a company-neutral platform concept for controlling how AI-enabled tooling is requested, reviewed and evidenced before it can interact with sensitive data, credentials or production pathways.

Prepared by Markus Walker
Security, AI and Cloud Engineer
Brisbane, Queensland, Australia

markus@markuswalker.com
linkedin.com/in/markus-walker-au
github.com/markus-doc

CONTROL LOOP

Request
Capture owner, purpose, data class, environment and tool details.

v

Classify
Evaluate privilege, data exposure, tool behaviour and vendor maturity.

v

Review
Route material risk to an accountable reviewer with conditions where needed.

v

Evidence
Record the terminal decision in a durable, structured format.

Core idea: treat AI integrations as governed access requests, not unmanaged developer tooling.

01 / EXECUTIVE SUMMARY

The decision problem

AI agents and MCP-connected tools can receive data, invoke APIs and trigger actions across SaaS, cloud and internal platforms. The organisation needs a repeatable way to decide which integrations are allowed, who approved them and what evidence proves the control operated.

1	8	42	4
RUNNABLE WORKFLOW	CONTROL OBJECTIVES	PASSING TESTS	RISK LEVELS

The risk

AI tooling can become an access path before ownership, data exposure, permissions and review evidence are understood.

The response

A lightweight governance layer validates requests, classifies risk, routes review and writes evidence.

The operating value

Teams can enable AI adoption while preserving accountability, least privilege and traceable decisions.

The boundary

The platform is designed for approval, review and evidence. It is not silent auto-provisioning.

Outcome: a practical control loop that turns AI integration requests into visible, reviewable and evidence-backed decisions.

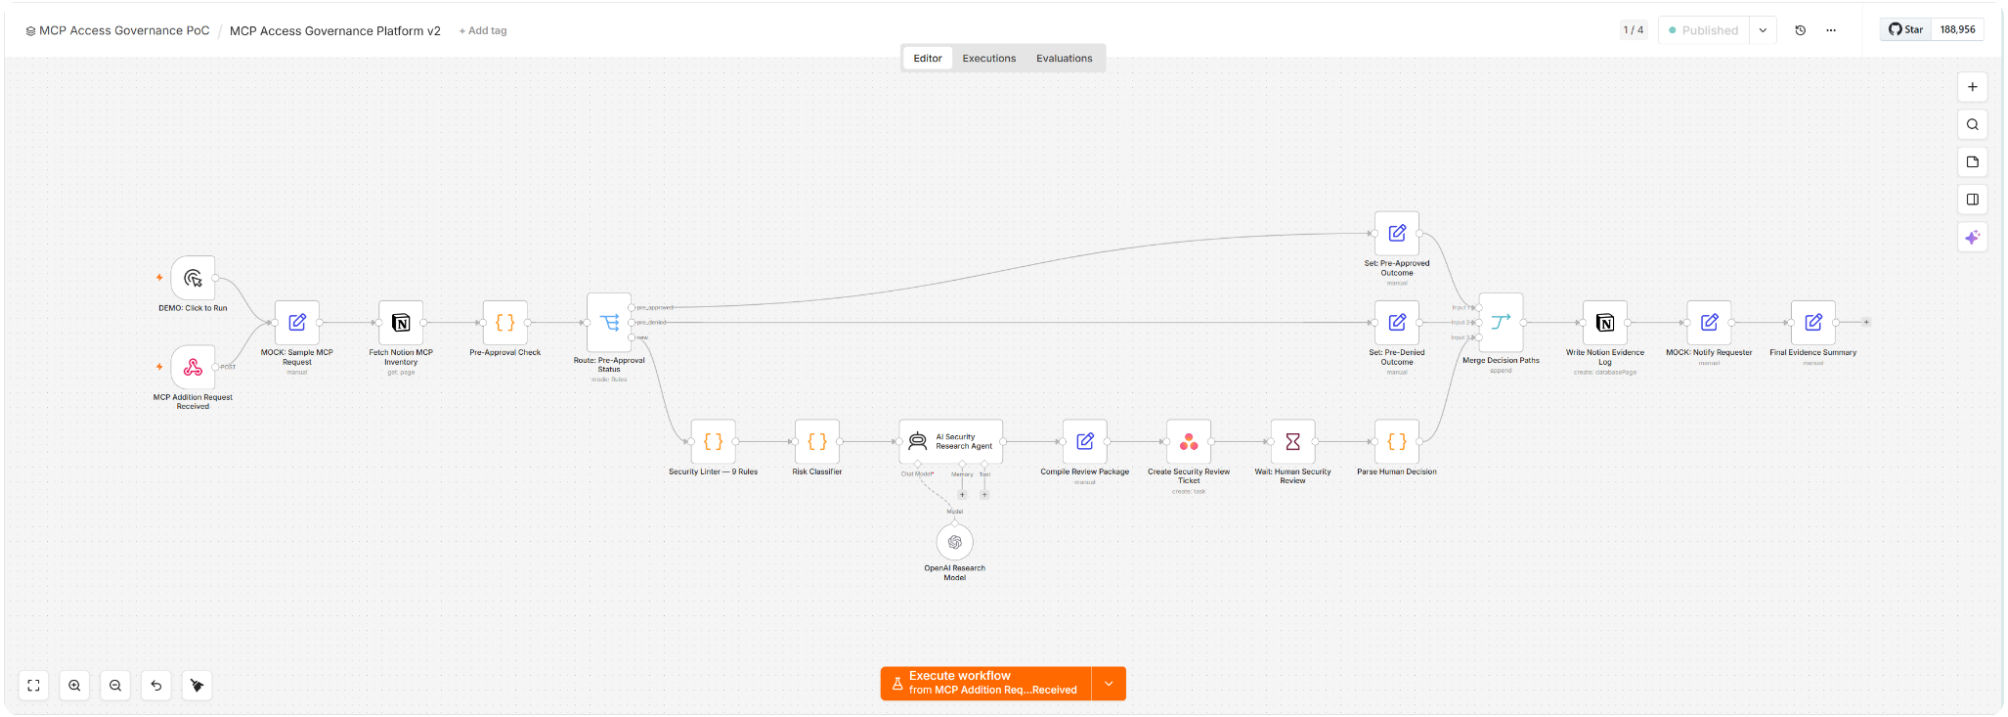

Full n8n workflow canvas: intake, inventory lookup, pre-approval check, route by status, security linter, risk classifier, review package, human wait gate, evidence write and final summary.

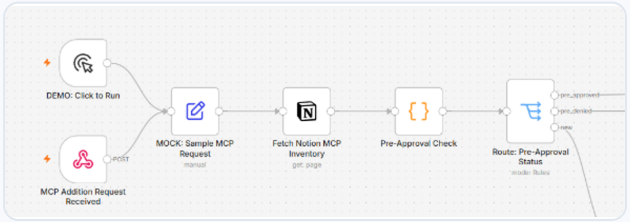

Intake and validation

Request trigger, inventory lookup and pre-approval status check.

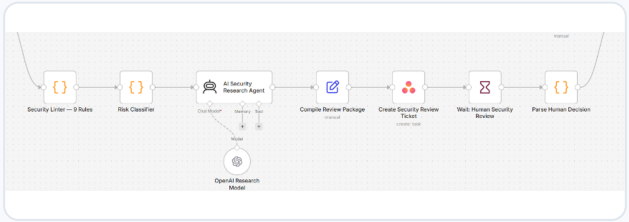

Risk and review

Security linter, risk classifier, AI research support and reviewer package.

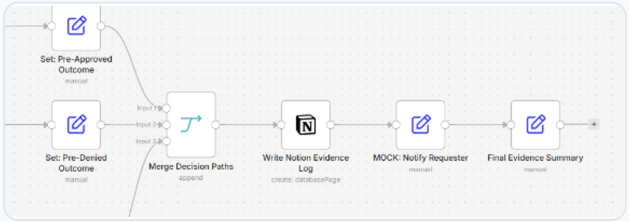

Evidence path

Decision merge, evidence write, requester notification and final summary.

02 / WHY THIS MATTERS

The governance gap created by AI tooling

Traditional access governance often assumes human users, known applications and stable permission models. AI and MCP integrations break that pattern because a workflow may include an agent, a model provider, an MCP server, a SaaS platform, credentials, data movement and a human operator.

Uncontrolled adoption can create

- Sensitive data passing through tools that were never reviewed.
- Unclear ownership for AI workflows and MCP integrations.
- Inconsistent approval criteria between teams.
- Weak evidence when a decision is challenged later.
- Access scope that grows without periodic review.

A governed process should provide

- Required intake fields before assessment can start.
- Consistent risk classification based on documented signals.
- Named accountability for material approvals.
- Conditions for time-bound and purpose-bound access.
- Durable records that show what decision was made and why.

The platform changes AI access from an informal request into a controlled workflow with ownership, routing, review and evidence.

QUESTION	PLATFORM RESPONSE
Who owns the workflow?	Ownership metadata is required at intake and carried into the evidence record.
What can it access?	Classification records data class, environment, privilege scope and MCP tool detail.
Was a human accountable?	Medium and high risk requests require named review before approval.
Can the control be proven later?	Each terminal path produces a structured decision record.

03 / OPERATING MODEL

Request to evidence

The operating model is intentionally simple. It sits between a requester and any approved use of an AI tool, agent workflow or MCP server that may interact with organisational systems.

1. **SUBMIT** — A requester provides purpose, owner, data class, environment and tool detail.
2. **VALIDATE** — Required fields and minimum governance information are checked.
3. **CLASSIFY** — Risk signals are assessed across data, privilege, environment and vendor maturity.
4. **ROUTE** — Low risk can proceed with evidence. Material risk routes to review.
5. **REVIEW** — Reviewer approves, approves with conditions or denies.
6. **EVIDENCE** — The terminal decision is written to a structured record.

Risk classification

LEVEL	ROUTING	OUTCOME
Low	Automated	Allowed with evidence.
Medium	Human review	Named reviewer and decision record.
High	Human review	Escalation with conditions or denial.
Prohibited	Auto-denied	Immediate denial with no exception path.

Decision principles

- Require enough information to classify risk before approval.
- Keep material approvals accountable to a named reviewer.
- Use conditions where access is allowed but constrained.
- Record denial paths as deliberately as approval paths.
- Preserve evidence in a structure that can be reviewed later.

Design boundary: the workflow supports governed decisions first. Any future provisioning integration should sit behind explicit approval and change control.

04 / REFERENCE ARCHITECTURE

Governed orchestration layer

The architecture keeps the governance layer visible. n8n orchestrates intake, classification and routing. Asana provides review tasks. Notion provides a control register and inventory layer. GitHub stores workflow exports, schemas, source code and evidence examples.

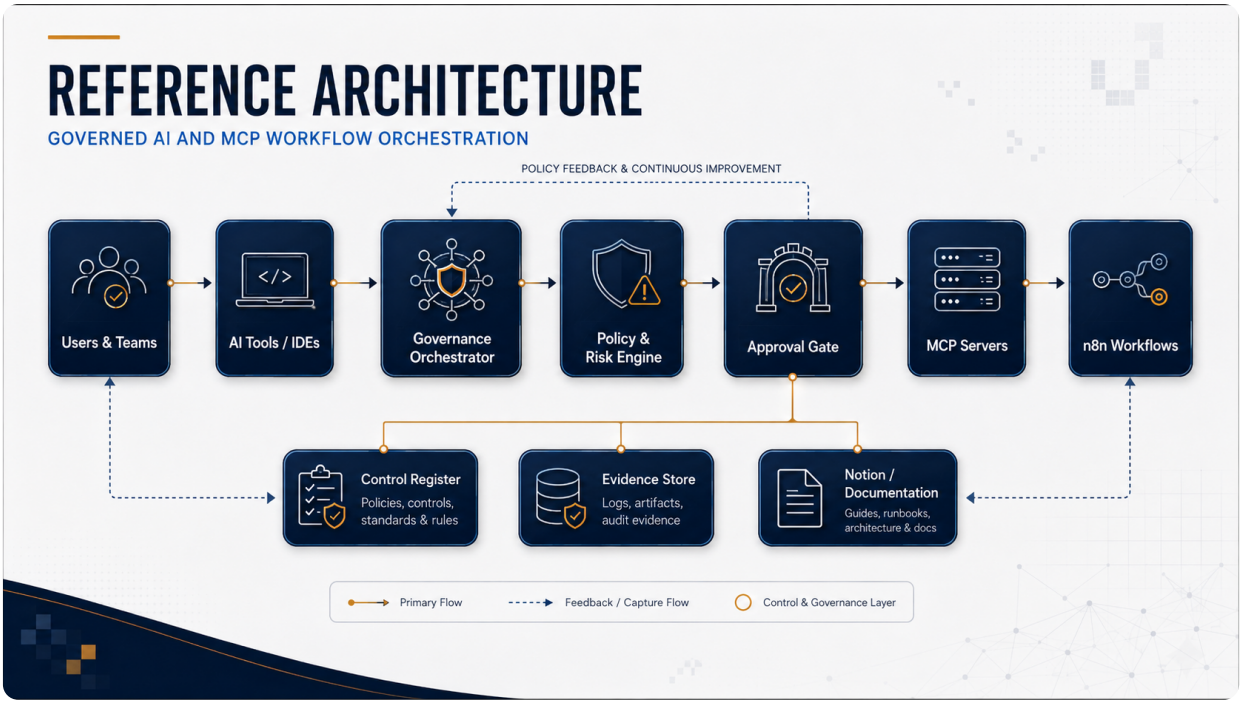

Reference architecture: users and teams, AI tools, governance orchestrator, policy and risk engine, approval gate, MCP servers, n8n workflows, registers and evidence stores.

Related workflow view: the dedicated n8n canvas on page 3 shows this architecture implemented as an operational control flow.

PLATFORM	ROLE IN THE MODEL
n8n Cloud	Workflow orchestration, validation, classification, decision routing and evidence summary.
Asana	Human review task queue, remediation tracking and operational evidence capture.
Notion	Governance knowledge base, MCP server inventory, control notes and decision records.
GitHub	Source control, documentation, workflow exports, JSON schemas and test evidence.

05 / CONTROL OBJECTIVES

Controls for safe AI and MCP adoption

The control set is written for practical operation. It describes what must be captured, how decisions are made, what boundaries apply and how evidence is retained.

REF	CONTROL OBJECTIVE
GOV-01	Structured intake: every request is captured with required fields.
GOV-02	Risk-based decisioning: classification uses documented and consistent criteria.
GOV-03	Human accountability: material risk routes to a named security reviewer.
GOV-04	Audit evidence: every terminal decision produces a structured record.
GOV-05	Least privilege: approved access is time-bound and purpose-bound.
GOV-06	Safe automation boundary: no real provisioning occurs without explicit approval.
GOV-07	Scope guardrails: workflow changes follow documented architecture.
GOV-08	Periodic review: approved access is re-reviewed on material scope change.

Identity and ownership

Every workflow needs a business owner, technical owner and named decision path.

Least privilege

Approval should be constrained by purpose, data class, environment and duration.

Approval gates

Material risk is held for explicit review instead of being allowed by automation alone.

Evidence quality

The record must show the request, risk level, reviewer, decision, conditions and date.

06 / TESTABLE GOVERNANCE

Security linter for n8n exports

A deterministic Python linter inspects n8n workflow JSON exports for AI and MCP governance risk signals. It converts policy intent into checks that can be tested, reviewed and extended.

6	42	0	1
LINTER RULES	PYTEST TESTS	RUNTIME SIDE EFFECTS	STRUCTURED OUTPUT PATH

RULE	SEVERITY	SIGNAL
LINT-001	High	MCP client tool node detected.
LINT-002	High	MCP node with credential references.
LINT-003	Medium	HTTP node calling MCP-like endpoint directly.
LINT-004	High	AI agent node with no human approval gate upstream.
LINT-005	Medium	Workflow missing owner metadata.
LINT-006	Low	Workflow missing risk classification metadata.

Engineering value: governance requirements are not only described in documents. They are turned into repeatable checks with test coverage.

The workflow view on page 3 includes the security linter and risk classifier stages as part of the decision path.

07 / FRAMEWORK ALIGNMENT

Standards-aware without overclaiming

The design uses recognised security and AI governance language to shape control objectives, evidence design and review pathways. It does not claim certification, attestation or formal control operation against any framework.

FRAMEWORK	ALIGNMENT AREA
ISO/IEC 27001	Access control, supplier risk, change management and auditability.
ISO/IEC 42001	AI management, human oversight, accountability and lifecycle governance.
NIST SP 800-37	Risk management structure and control lifecycle thinking.
NIST SP 800-30	Risk assessment language, likelihood, impact and treatment.
CSA AI Controls Matrix	Cloud AI governance, access control and data governance.
MAESTRO	Agentic AI threat modelling, orchestration and tool execution.
OWASP LLM Top 10	Prompt injection, data leakage, tool misuse and excessive agency.
MCP Security Guidance	Consent, authorisation and confused deputy protections.

Control language

Risk and review terms stay consistent from intake to decision evidence.

Audit trail

Terminal decisions create records that can be reviewed without reconstructing chat or ticket history.

Lifecycle view

Approvals can be linked to scope change, expiry and periodic review.

Scope honesty

The design is presented as a proposed governance model and working prototype, not an attested control environment.

08 / ADOPTION PATHWAY

How the model can mature

A production programme can keep the same core control loop while increasing assurance around identity, evidence retention, change control, monitoring and periodic review.

Phase 1
Controlled pilot

- Define intake fields.
- Register AI tools and MCP servers.
- Run approvals in review mode.
- Validate evidence shape.

Phase 2
Operational hardening

- Add identity-aware boundaries.
- Integrate change records.
- Version schemas and exports.
- Review evidence quality.

Phase 3
Governed scale

- Add periodic review.
- Introduce reporting dashboards.
- Map controls to evidence.
- Expand to more workflows.

Decisions supported by the platform

DECISION	WHY IT MATTERS
Which integrations are allowed?	Aligns AI adoption with risk appetite and operational priorities.
Who can approve material risk?	Creates accountable decision ownership.
What conditions must be met?	Supports purpose-bound, time-bound and data-bound approval.
What evidence is retained?	Shows how a request moved from intake to decision.

Recommended first use: pilot against non-production AI integration requests and review the quality of decision evidence before enforcement is expanded.

09 / SCOPE, ASSUMPTIONS AND SOURCE

Clear boundaries for review

The material is intentionally sanitised and company-neutral. It shows an approach to governing AI and MCP access without exposing private environments, proprietary infrastructure, live credentials or customer data.

Included

- Runnable n8n workflow prototype.
- Structured intake and decision flow.
- Governance documentation and schemas.
- Workflow export and evidence examples.
- Security linter with 42 passing tests.

Not claimed

- No claim of production deployment.
- No formal compliance certification claim.
- No complete enterprise control environment claim.
- No live credentials or real customer data.
- No automatic provisioning without approval.

Name	Markus Walker
Role	Security, AI and Cloud Engineer
Location	Brisbane, Queensland, Australia
Email	markus@markuswalker.com
LinkedIn	linkedin.com/in/markus-walker-au
GitHub	github.com/markus-doc
Website	markuswalker.com
Project repository	github.com/markus-doc/MCP_Governance_Platform