PRODUCT BRIEF

MCP Security Governance

Real-time compliance reporting and deterministic security enforcement for the MCP servers your AI agents depend on.

THE COMPLIANCE BLIND SPOT

MCP is where AI compliance obligations are won or lost

MCP has become the de facto standard for AI agent-to-tool integration. AI agents now autonomously invoke tools, read resources and act on enterprise data through MCP servers — shifting integration from deterministic API calls to dynamic, agent-driven interactions. Unmanaged adoption creates shadow integration paths that bypass enterprise controls: a single mis-scoped tool or logging gap produces untraceable agent actions, with no evidence trail to prove which controls were in place, or that they held.

Existing WAFs, AI gateways, DAST tools and GRC platforms cannot evidence MCP compliance.
Without control at the MCP layer, an enterprise cannot prove its AI policy is enforced.

THE 42CRUNCH APPROACH

From discovery to enforcement

42Crunch enables enterprises to govern the MCP servers they expose to agents, partners and customers. It goes beyond discovering servers and identifying risks: it evidences compliance against the frameworks you are accountable to, remediates vulnerabilities and blocks non-compliant servers — ensuring only secure, approved MCP services reach production.

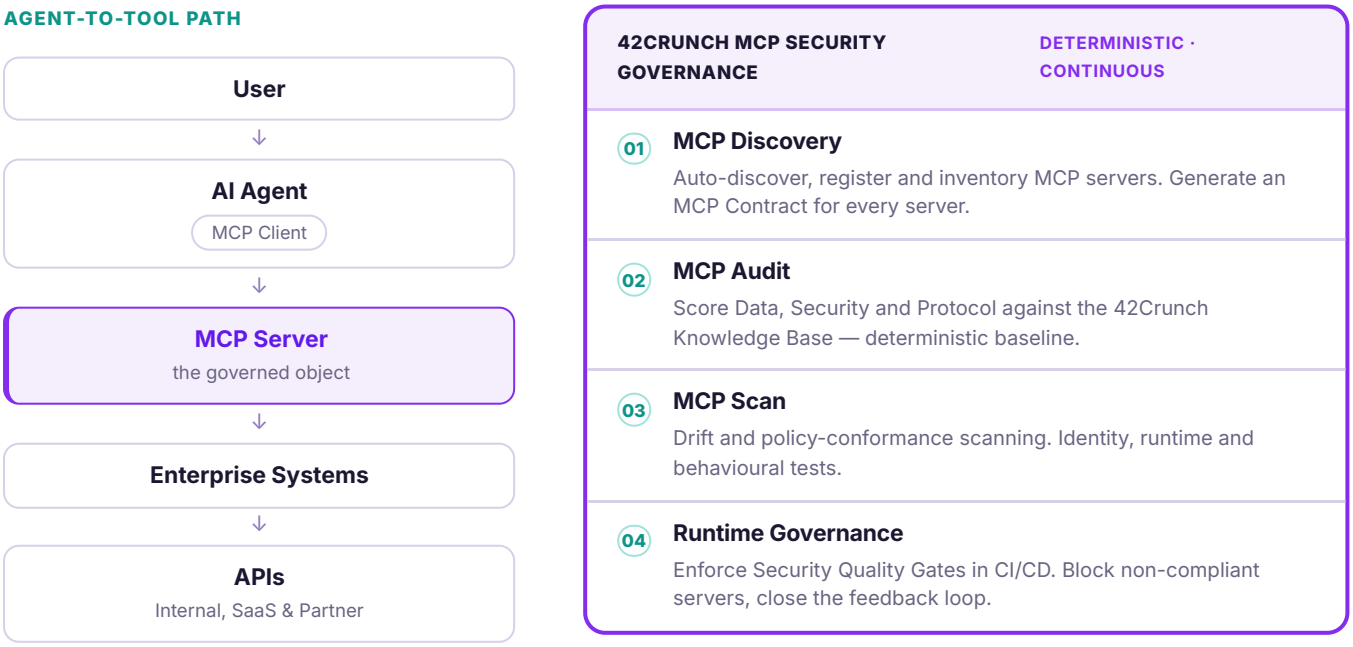

Fig 1. Discovery, audit, scan and runtime governance across the agent-to-tool path.

HEADLINE CAPABILITY

Real-time compliance reporting against regulatory frameworks

42Crunch continuously evaluates every MCP server against the regulatory and industry frameworks your organization is accountable to, and maps each finding to the specific control it breaches. Compliance posture is computed by the control on every audit and every scan — not assembled from documents at audit time.

FRAMEWORK	WHAT 42CRUNCH EVIDENCES	REPORTING
NIST AI RMF	Govern, Map, Measure and Manage functions evidenced for every agent-accessible tool surface.	**Per-finding control mapping**
OWASP MCP Top 10	Deterministic coverage of all ten MCP vulnerability classes, scored per server.	**Scored 0–100 / A–F**
EU AI Act	Technical documentation, logging, human oversight and risk-management obligations for high-risk AI systems.	**Versioned audit record**
ISO/IEC 42001	AI management-system controls operationalized as machine-checkable requirements.	**Continuous conformity**
CSA AICM	AI Controls Matrix domains — including CCC and GRC control families — mapped to concrete findings.	**Per-domain findings**

What this gives the compliance organization

Continuous evidence — every audit, scan and Security Quality Gate decision produces a versioned, scored, framework-mapped record, not a point-in-time attestation.

Deterministic, not LLM-judged — the same server evaluated twice returns the same score and the same findings.

Finding-level traceability — each finding carries a control reference, severity and remediation.

Enforced, not advisory — Security Quality Gates block non-compliant MCP servers in CI/CD before production.

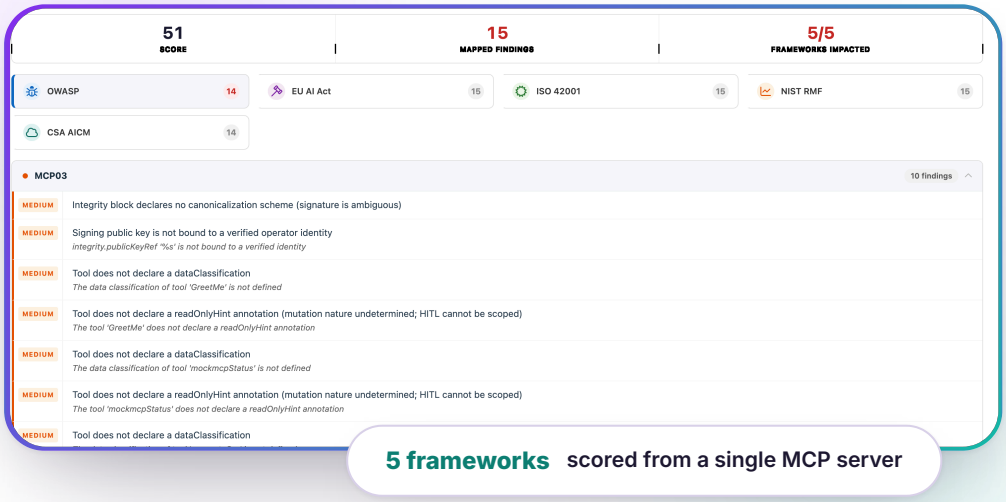

Fig 2. Compliance reporting — one MCP server scored against five frameworks.

HOW IT WORKS

One policy artifact drives every compliance decision

42Crunch generates an MCP Contract for each server: a machine-readable specification of how it must operate. Written against a common enterprise policy schema, it is the single source of truth across discovery, audit, scanning and compliance reporting.

01

Continuous Discovery

42Crunch continuously discovers MCP servers across the enterprise and generates an initial contract for each — an immediate inventory without manual documentation.

02

Drift Detection

Running servers are tested against their approved contracts, identifying OWASP MCP Top 10 vulnerabilities and unauthorized changes before they expose enterprise systems.

03

Data Classification

Each tool is evaluated on the sensitivity of the data it handles and the actions it performs. High-impact tools get stricter controls, including human-in-the-loop approval.

04

Audit Traceability

Every assessment and Security Quality Gate decision creates a versioned, scored governance record, evidencing that required controls were evaluated and enforced.

SECURITY COVERAGE

MCP server vulnerability prevention

By validating tool definitions, inputs, credentials, permissions, isolation and delegated identity against the approved MCP Contract, 42Crunch blocks unsafe MCP servers before connection and detects drift after deployment.

Tool Poisoning
Malicious instructions embedded in tool descriptions that hijack LLM behavior.

Rug Pulls
Dynamic replacement of a trusted tool definition with a malicious one at runtime.

Code & Command Injection
Model-provided inputs passed to shell, SQL or system APIs without validation.

Credential Leakage
API keys and OAuth tokens improperly stored, logged or cached.

Excessive Permissions
Over-privileged tools violating least-privilege — one breach compromises everything.

Insufficient Isolation
Cross-tenant data leakage, shared service accounts, unsandboxed execution.

Confused Deputy
Token passthrough lets the server be tricked into misusing user privileges.

Source: OWASP, A Practical Guide for Secure MCP Server Development.

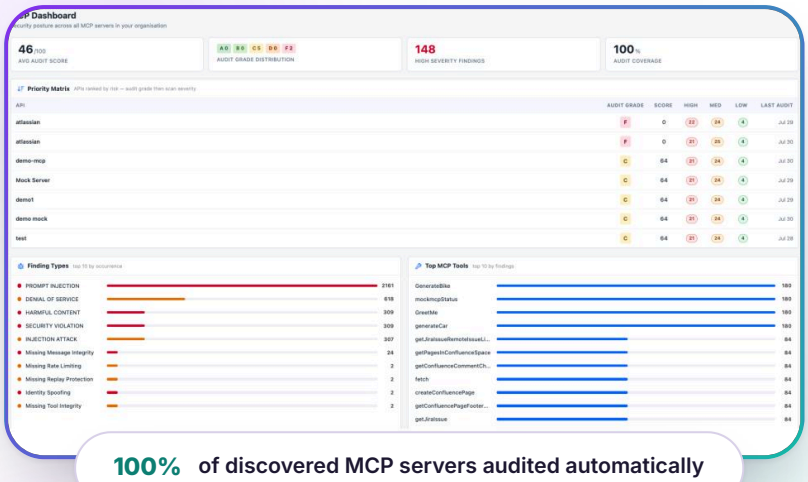

100% of discovered MCP servers audited automatically

Fig 3. MCP Governance Dashboard — posture across every server.

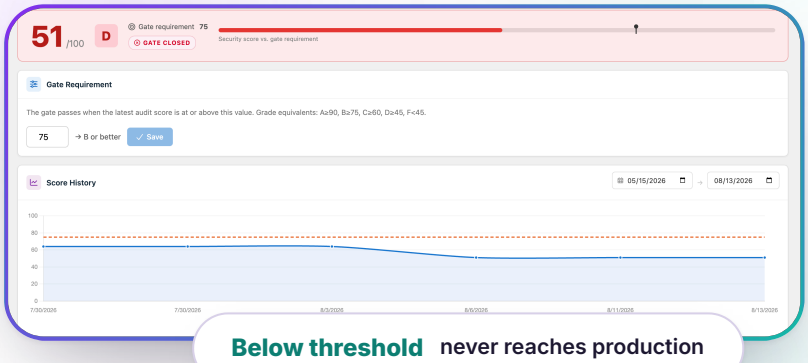

Below threshold never reaches production

Fig 4. Security Quality Gate — enforced pass/fail scoring.

● IMMEDIATE IMPACT

What the agentic enterprise gets

Full AI regulatory framework compliance

NIST AI RMF · OWASP MCP Top 10 · EU AI Act · ISO 42001 · CSA AICM — enforced through 42Crunch Security Quality Gates.

Compliance baseline in 60 seconds

Score any MCP server rapidly, prioritize remediation sooner and shorten the security review to avoid release delays.

Zero-touch MCP discovery

Auto-inventory of MCP servers across cloud infrastructure, with evaluation starting immediately to accelerate time to market.

Full OWASP MCP vulnerability coverage

Deterministic coverage across all OWASP MCP Top 10 vulnerability classes to reduce gaps in defense.

Determine your MCP server compliance baseline in 60 seconds

Point the 42Crunch baseline scan at any MCP server and receive a scored report against the OWASP MCP Top 10 and your regulatory frameworks in under 60 seconds. No agent, no deployment, no commitment required.

Register for your free evaluation → **42crunch.com/mcp-baseline**