# MCP Launch Readiness Audit

Target: **Sample Vulnerable Fixtures**
Auditor: Payload QA
Date: 2026-10-02

This audit reduces risk; it is not a certification, not a guarantee of security, and not penetration testing. A clean scan means no known patterns were found — it cannot prove the absence of vulnerabilities. Re-scan after every material change.

## Executive summary

Static scan found **23** findings: 0 high, 0 medium, 0 low.
Reliability probing found **0** issues: .
Readiness: **8** pass, **0** fail, **5** manual attestations.

This audit reduces risk; it is not a certification, not a guarantee of security, and not penetration testing. A clean scan means no known patterns were found — it cannot prove the absence of vulnerabilities. Re-scan after every material change.

## Findings & remediation

| Sev | Rule | Title | Location | Remedy |
|---|---|---|---|---|
| **HIGH** | `CR-006` | Stripe secret key in source | `fixtures/exposed_server.ts:13` | Roll the Stripe secret key; publishable keys only in source. |
| **HIGH** | `EX-001` | Server binds all network interfaces (0.0.0.0) | `fixtures/exposed_server.ts:26` | Bind localhost; put auth in front for remote access. |
| **HIGH** | `PV-001` | Dangerously powerful tool name | `fixtures/overprivileged_tools.json:1` | Rename to a narrow capability or remove the tool. |
| **HIGH** | `PV-001` | Dangerously powerful tool name | `fixtures/overprivileged_tools.json:1` | Rename to a narrow capability or remove the tool. |
| **HIGH** | `PV-002` | Raw SQL execution tool | `fixtures/overprivileged_tools.json:1` | Replace raw SQL with parameterized, allowlisted queries. |
| **HIGH** | `SS-001` | URL parameter with no SSRF protection | `fixtures/overprivileged_tools.json:1` | Validate URL parameters against an allowlist; block private ranges. |
| **HIGH** | `TP-001` | Instruction override phrase in tool description | `fixtures/poisoned_server.py:15` | Strip behavioral instructions from the tool description; keep a pure capability statement. |
| **HIGH** | `TP-004` | Data exfiltration instruction in tool description | `fixtures/poisoned_server.py:19` | Remove the exfiltration instruction; tools return data to the agent, never transmit it. |
| **HIGH** | `TP-005` | Conceal-from-user directive in tool description | `fixtures/poisoned_server.py:19` | Delete secrecy directives; log the tool's behavior instead. |
| **MEDIUM** | `AU-003` | No rate limiting | `(project):0` | Add per-tenant rate limiting. |
| **MEDIUM** | `EX-002` | Debug mode enabled in server source | `fixtures/exposed_server.ts:14` | Gate debug mode behind an env var defaulting to off. |
| **MEDIUM** | `EX-004` | Stack traces returned to clients | `fixtures/unstable_server.py:35` | Log tracebacks server-side; return a generic error id. |
| **MEDIUM** | `PV-004` | Network fetch tool with no URL allowlist | `fixtures/overprivileged_tools.json:1` | Add a domain allowlist and block private/metadata IPs. |
| **MEDIUM** | `PV-006` | Tool without input schema | `fixtures/overprivileged_tools.json:1` | Define an inputSchema for every tool. |
| **MEDIUM** | `PV-006` | Tool without input schema | `fixtures/overprivileged_tools.json:1` | Define an inputSchema for every tool. |
| **MEDIUM** | `PV-008` | Sensitive path referenced in tool definition | `fixtures/exposed_server.ts:23` | Stop referencing credential paths in tool definitions. |

| Sev | Rule | Title | Location | Remedy |
|---|---|---|---|---|
| **MEDIUM** | `PV-008` | Sensitive path referenced in tool definition | `fixtures/overprivileged_tools.json:1` | Stop referencing credential paths in tool definitions. |
| **MEDIUM** | `PV-008` | Sensitive path referenced in tool definition | `fixtures/poisoned_server.py:16` | Stop referencing credential paths in tool definitions. |
| **MEDIUM** | `PV-008` | Sensitive path referenced in tool definition | `fixtures/poisoned_server.py:19` | Stop referencing credential paths in tool definitions. |
| **MEDIUM** | `TP-006` | Safety override instruction in tool description | `fixtures/poisoned_server.py:17` | Remove safety-weakening language from tool metadata. |
| **MEDIUM** | `TP-012` | External URL embedded in tool description | `fixtures/poisoned_server.py:16` | Audit or remove outbound URLs in tool metadata. |
| **LOW** | `EX-007` | Sensitive route path exposed | `fixtures/exposed_server.ts:22` | Remove or protect sensitive routes. |
| **LOW** | `PV-007` | Overly broad input schema (additionalProperties: true) | `fixtures/overprivileged_tools.json:1` | Enumerate allowed properties explicitly. |

## Reliability test results

No reliability issues found across malformed-input, unexpected-argument, rapid-load, burst, and isolation probes.

- sequential_50: {'p95_seconds': 0.0, 'over_timeout': 0}
- burst_5x20: {'responses': 100}
- isolation: checked

## Deployment readiness checklist

| Status | Check | Detail |
|---|---|---|
| **PASS** | **Authentication configured**<br>Every tools/call must authenticate the caller. | pattern matched (oauth...); live: unauthenticated call rejected (static+live) |
| **PASS** | **Per-tool authorization scopes**<br>Tokens carry scopes; each tool enforces the scopes it needs. | pattern matched (scope...) (static) |
| **PASS** | **Audit logging enabled**<br>Every tool call is logged with timestamp, tenant, tool, outcome. | pattern matched (audit...) (static) |
| **PASS** | **Rate limits configured**<br>Per-tenant/per-tool rate limiting is present. | pattern matched (rate_limit...) (static) |
| **PASS** | **Server-side input validation**<br>Arguments are validated for type, presence, and length. | pattern matched (validat...) (static) |
| **PASS** | **Tenant isolation**<br>Tenant comes from the auth token; client tenant claims are verified. | pattern matched (tenant...) (static) |
| **PASS** | **Error responses are safe**<br>No stack traces or internals leak to clients. | no traceback leakage pattern in source (static) |
| **PASS** | **Versioned with documentation**<br>Server reports a version; a README documents setup and security. | README.md: True; version string: True (static) |

| Status | Check | Detail |
|---|---|---|
| MANUAL | **Secrets rotated and out of source**<br>No credentials in the repo or its history; all secrets in env/secret manager. | requires human attestation (manual) |
| MANUAL | **Dependencies audited**<br>Third-party packages pinned and checked for known vulnerabilities. | requires human attestation (manual) |
| MANUAL | **TLS in front of any remote transport**<br>Remote MCP endpoints are TLS-terminated with valid certificates. | requires human attestation (manual) |
| MANUAL | **Rollback plan exists**<br>You can revert to the previous server version quickly if a deploy breaks. | requires human attestation (manual) |
| MANUAL | **Incident runbook exists**<br>Someone other than you can follow the audit log and respond to abuse. | requires human attestation (manual) |

## Remediation status

For each finding above, record the disposition. Suggested values: **fixed** (code changed + re-scan clean), **accepted risk** (documented business reason), **false positive** (evidence noted). Re-run `mcp-sec-audit` after every fix.

## Sign-off

I have reviewed the findings, reliability results, and readiness checklist. I understand this audit reduces risk and is not a certification, guarantee of security, or penetration test.

Auditor signature / date: _____

Engineering lead signature / date: _____