SANCTUMSHIELD · AGENTIC GOVERNANCE FOR THE OTHER 99%

MSP AI Disclosure — The Standing Answer Sheet

Last updated: October 27, 2026

managed service providers (MSP / MSSP) who want to answer the AI-governance questions before a client, auditor, or insurer asks them.

Why this exists. Your clients are starting to ask what AI is running in their environment, who it reports to, and what happens when it misfires — and the questions arrive first through cyber-insurance renewals, then audits, then regulation. Fixing this client-by-client means re-papering a whole book of agreements. This is the way around that: **one standing AI Disclosure**, filled in once and refreshed quarterly, that your service agreement references by incorporation. New contracts carry the reference from day one; existing ones pick it up at renewal. One artifact instead of two hundred amendments — and handing it over unprompted during procurement makes you look categorically different from a provider who has to go find out.

How this pairs. This is the *answer* artifact. The MSP AI Trust Questionnaire is the *assessment* instrument your clients use — the ten questions this document answers in advance. Each section below is tagged with the client question it resolves.

How to use this — for providers

1. with your real, current facts. Where a control isn't in place yet, say so plainly and note the plan — an honest "not yet, here's the compensating control and the date" reads better than a dodge.
2. A claim resolves to a control, an artifact, a named owner, and a date. "We have a kill switch" is not an answer; the procedure, its scope, the authorized owner, the time-to-effect, and the last-tested date is.
3. ("Provider's current AI Disclosure, published at <URL>, is incorporated into this agreement"). Refresh it quarterly and whenever your platform vendor ships a material AI change.

4. Unprompted. It's a trust asset, not a compliance chore.

Honest limits: completing this Disclosure is a transparency and accountability step; it does not by itself make you or your client compliant. It is not legal advice — have counsel review the incorporation language.

The disclosure

Provider: *[legal entity name]* Version / date:

[vX.Y · YYYY-MM-DD]

1. AI features active in client environments

`ANSWERS CLIENT QUESTION 1`

[List every AI-driven feature currently active in client environments — across cloud, on-premises, and hybrid deployments — and, for each, what it can do and whether it can execute code or take action.]

Strong disclosure: distinguishes AI features from ordinary automation, and states how you learn when your platform vendor turns a new one on.

2. Change control over platform-vendor AI releases

`ANSWERS CLIENT QUESTION 2 · NON-NEGOTIABLE`

[Describe how much notice you get from your platform vendor before a new AI feature ships, and what happens between that notice and it running on client endpoints. Attach or reference your change-control policy.]

Strong disclosure: for cloud-delivered platforms, states plainly whether releases arrive automatically, and what control you do or don't have over them.

3. Approval process — and where the client enters it

`ANSWERS CLIENT QUESTION 3`

[Describe the approval step for AI features that can execute scripts or take action, and where (if anywhere) the client enters it. Reference the written procedure.]

4. Per-client / per-endpoint vs. global controls

ANSWERS CLIENT QUESTION 4

[State whether AI controls can be scoped or disabled for a single client organization (or a single endpoint), or whether enablement is global across your client base.]

Strong disclosure: if per-client scoping isn't available yet, say so and name the compensating control.

5. Evidence artifacts you produce

ANSWERS CLIENT QUESTION 5

[Describe the record of AI-initiated actions you can produce — what acted, on which asset, when, and under whose authority — the cadence (on event and monthly), and attach the example artifact a client can expect.]

6. Human review before execution

ANSWERS CLIENT QUESTION 6

[State whether a technician reviews AI-generated scripts before they execute in a client environment, or whether generated code runs unreviewed. Reference the review policy.]

7. "Kill switch" — scope, time, owner, last tested

ANSWERS CLIENT QUESTION 7 · NON-NEGOTIABLE

[Document the procedure to disable AI for a single client organization or endpoint without degrading the rest of the service: its scope, the time-to-effect, the named person authorized to invoke it, and the date it was last tested.]

Strong disclosure: a switch that only stops future invocations — or one that can only be thrown platform-wide — is disclosed as such, not called a per-client kill switch.

8. Data egress — recipients, terms, retention, training

ANSWERS CLIENT QUESTION 8 · NON-NEGOTIABLE

[State where client data goes when an AI feature processes it — ticket text, hostnames, file paths, session data — who receives it, under what terms, retained how long, and whether any of it is used to train a model. Attach the subprocessor list and data-processing terms.]

9. Breach notification — in hours

`ANSWERS CLIENT QUESTION 9`

[State your notification commitment, in hours, if your platform vendor or one of its AI subprocessors is breached — and reference the clause that commits to it.]

Strong disclosure: gives a number, not "promptly." The client's own regulatory clock runs regardless of the vendor→provider→client relay.

10. Named accountable owner

`ANSWERS CLIENT QUESTION 10`

[Name the specific person accountable for AI behavior in client environments — role, name, contact.]

Additional disclosures

MCP gateways and connectors into client systems

`ANSWERS THE QUESTIONNAIRE'S "ALSO WORTH ASKING"`

[List each MCP gateway or connector your platform uses to reach client systems, the systems behind it, and its owner.]

Agents acting in client environments: who can stop each, and when last proven

`ANSWERS THE QUESTIONNAIRE'S "ALSO WORTH ASKING"`

[For each AI agent that can act on client systems, name the person who can stop or pause it mid-task and the date that stop was last exercised, in a test or for real.]

Staff personal AI accounts and devices

`ANSWERS THE QUESTIONNAIRE'S "ALSO WORTH ASKING"`

[State your rule on technicians using personal AI accounts or personal devices with client data.]

Terms: MCP Gateway, Proven Stop, Inbound third-party agent. The Academy covers the tool and MCP supply chain in module M11 (CTO / CIO / IT track).

Right to decline

[State whether a client can opt out of AI features and remain a customer in good standing, including any price or service-level difference.]

If you are a security provider (MSSP) — autonomy inventory

[For each AI-assisted response action — host isolation, account disablement, alert suppression, detection tuning — state whether the AI acts autonomously or recommends to an analyst, and confirm that AI decisions to suppress or downgrade an alert are retained, reviewable, and present in the evidence you produce for client audits.]

How to incorporate this by reference

Add a clause to your service agreement rather than restating this document inside it, for example:

```
Provider maintains a current AI Disclosure describing the AI operating in
Client environments, published at <URL> and refreshed no less than quarterly.
That AI Disclosure is incorporated into this Agreement by reference. Provider
will notify Client of material changes to it.
```

Have your counsel review and adapt the wording. New agreements carry the reference from day one; existing agreements pick it up at renewal.

Guardrails honored

- Publishing this Disclosure attests transparency, not compliance or safety.
- No guarantee of outcome. It records what you disclose; it does not warrant a result.
- The subprocessor list, example artifact, and by-reference addendum are good practice to *adopt* — describe your real, current state, and note where something isn't in place yet.

Sources

- OWASP GenAI Security Project, *OWASP Top 10 for Agentic Applications (2026)* (ASI01–ASI10) — https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/
- Cloud Security Alliance, *The Agentic Trust Framework* (Feb 2, 2026) — https://cloudsecurityalliance.org/blog/2026/02/02/the-agentic-trust-framework-zero-trust-governance-for-ai-agents
- Cunningham, C. *Agentic Zero Trust* v3.0 (May 2026), Cequence — https://www.cequence.ai/agentic-zero-trust/
- NIST, *AI Risk Management Framework (AI RMF 1.0)* — https://www.nist.gov/itl/ai-risk-management-framework

ACKNOWLEDGMENT · Practitioner red-team review by Nate Moore, CISSP (linkedin.com/in/natemoore).

© PIGENAI LLC · SanctumShield · sanctumshield.com · Last updated October 27, 2026 · This is an educational template, not legal advice.