PORT1

Liminal For MSPs

Secure AI workspace checklist for MSPs

A checklist for MSPs choosing a secure AI workspace for clients: data protection, models, separation, identity, logs and compliance, with how Liminal does it.

November 2026 | Jason Makevich, CISSP

A secure AI workspace for clients should protect sensitive data before prompts reach a model, give people the models they need under one set of controls, keep each client separate, work with the client's identity provider and log every interaction. Use this checklist to size up any platform.

Each item notes how Liminal handles it.

How should MSPs use this checklist?

Use it to compare platforms before you commit, and again in a client's security review to show what's covered.

Specific answers help more than yes or no ones here, and so do notes on anything that depends on configuration, like which models a team gets or which policy applies to each data type. PORT1 runs Liminal demos if you want to see how it handles these items.

Does it protect sensitive data before a model sees it?

Look for a workspace that finds sensitive data in prompts and files and protects it before anything reaches a model.

- **Detects the data your clients handle.** Liminal: as of November 2026, PII, PHI, PCI and intellectual property across more than two dozen built-in data types, plus custom terms like client names.
- **Covers files as well as typed prompts.** Liminal: detection covers prompts and files.
- **Lets admins choose the action for each data type.** Liminal: mask, redact, warn or block, set per model for the organization, teams or individuals.
- **Keeps answers useful.** Liminal: restores the protected details in the response.

Does it give people the models they need?

People who get the models they want in the approved workspace have less reason to go looking elsewhere.

- **Several providers in one place.** Liminal: unlimited use of the latest models from OpenAI, Anthropic, Google, Perplexity and xAI.
- **Clear terms on use.** Liminal: no token limits, metering or per-user caps on included models in normal use, with fair-use terms for automated, high-volume traffic.
- **Room for other models.** Liminal: clients can bring their own models on their own credentials, billed by that model's provider.
- **One set of controls across models.** Liminal: the same protection and logging apply whichever model someone picks.

Does it keep each client separate?

Each client's data should stay in its own environment, in a region that fits the client.

- **A dedicated environment for each client.** Liminal: each client gets its own dedicated tenant, with its services, data, search indexes, logs and encryption keys kept separate.
- **Data residency options.** Liminal: the US (the default), the UK and the EU.
- **No training on customer data.** Liminal never trains on customer data.
- **Encryption.** Liminal: in transit and at rest, with AES-256 for prompt logs.

Does it fit the client's identity and access setup?

The workspace should use the identity provider and roles the client already runs.

- **Single sign-on.** Liminal: SAML 2.0 with any compliant identity provider, including Microsoft Entra ID, Okta and Ping Identity.
- **Existing teams and groups.** Liminal brings in the teams, groups and permissions the client already manages.
- **Role-based access.** Liminal: custom roles decide who uses which models and who manages what.

Can you see and prove what happened?

When a client asks who sent what, to which model and under which policy, the logs should answer it.

- **Detailed logs.** Liminal logs every interaction with the original and protected prompt, the user, the model and the policies applied.
- **Exports your SIEM can use.** Liminal: CSV, JSON or XML downloads, or scheduled exports to AWS S3, Azure Blob Storage or Google Cloud Storage.
- **Usage and value reporting.** Liminal's Insights Engine shows usage, topics and return on investment, with sensitive data only as counts by type.

Does it support the client's compliance needs?

The vendor's own audits and agreements should fit the client's industry.

- **Independent audits.** Liminal: SOC 2 Type 2 certified, with three consecutive audits completed as of November 2026.
- **BAAs for healthcare.** HHS says covered entities and business associates can use a cloud service provider to store or process electronic protected health information, provided they have a HIPAA-compliant business associate agreement with that provider (HHS guidance on HIPAA and cloud computing). Liminal: HIPAA compliant, with BAAs available for healthcare customers.
- **Coverage for the models too.** Liminal holds BAAs with Anthropic, OpenAI and Google for the models it offers. Perplexity and xAI models aren't covered, so teams that handle PHI can be kept on covered models.
- **Other frameworks.** Liminal: FERPA, COPPA and CIPA compliant, and built to support GLBA, CJIS, GDPR and CCPA compliance and GxP requirements.

Can it connect to the systems clients use?

Connections should run inside the same governed workspace.

- **Governed connections.** Liminal: MCP connections to the systems a client chooses, included at no extra cost, with answers that cite their sources.

Can you deliver it as a managed service?

The platform should fit how an MSP works across many clients.

- **A tenant for each client, and one for you.** Liminal: each client gets its own tenant, and your own use runs in a separate one.
- **A way to try it.** Liminal: partners get a shared sandbox with no client data, and PORT1 runs Liminal demos.
- **Help when you need it.** Partners get PORTER, PORT1's partner platform, and PORTER Chat answers Liminal product and technical questions 24x7.
- **Security review support.** Liminal's security documents and BAA requests go through PORT1.

For more detail, see Liminal security and compliance, how Liminal keeps sensitive data out of AI models and how MSPs offer secure AI as a service.

Sources

- HHS, "Guidance on HIPAA & Cloud Computing" (hhs.gov), accessed November 14, 2026.
- Liminal product details come from Liminal's public website, checked November 2026.
- PORT1 program details come from PORT1.

Download: Secure AI workspace checklist for MSPs ›

Book a Liminal demo