ATLAS TECHNICA

Solution Brief: Secure AI Foundation (SAIF)

Institutionalizing Intelligence: From Shadow AI to a Governed AI Operating Boundary

As the industry's first Managed Intelligence Provider (MIP), Atlas Technica uses SAIF to establish a secure, enterprise-grade AI operating boundary that moves firms from unmanaged, shadow AI to a governed model built on approved access, data controls, and ongoing oversight.

Governed Access
Control who can use AI and under what conditions.

Data Boundaries
Keep sensitive information inside managed control.

Auditability
Maintain usable visibility for compliance and regulatory review.

Governed AI control layer typically deployed & tested in 3-4 weeks*.

*subject to client access, readiness & approvals

IMPLEMENTATION MODULES

MODULE	WHAT WE DO	BUSINESS IMPACT
Shadow AI Containment	Identify unsanctioned AI usage patterns and steer users toward approved tools through supported browsers, and endpoint controls.	**Reduces shadow AI risk.** Creates visibility into real usage and helps limit data exposure through unsanctioned AI tools.
Identity & Device Controls	Configure SSO, MFA, conditional access, approved-user groups, and managed-device requirements for sanctioned AI platforms.	**Strengthens access posture.** Ensures approved users access AI through the organization's identity and device posture.
Permission Audit	Review SharePoint for oversharing and inherited access risk.	**Reduces exposure risk.** Helps reduce accidental exposure of sensitive content through AI.
Sensitivity Labeling	Publish and apply a practical label, including restricted content where appropriate.	**Clearer data boundaries.** Creates practical boundaries for AI usage and governance.

ONGOING SERVICE & GOVERNANCE

Atlas provides a structured operating layer after go-live to keep AI access governed, auditable, and supportable as usage expands.

Governance Review
Regular review of approved AI access, block lists, shadow AI categories, and monitor/block settings as usage evolves. Includes review of M365 and related permissions affecting AI access.

Reporting & Oversight
Governance reporting on usage trends, notable activity, key exceptions, and recommended control adjustments. Includes change notices for control or vendor changes.

Operational Support
Ongoing handling of AI-related tickets, access issues, control exceptions, and remediation as adoption expands. Includes incident handling for data-exposure events, with updates until closed.

Guardrail Tuning
Periodic tuning of AI policies and configs as needs, vendor capabilities, and regulations change.

Your SAIF Tier 1 Service

What is covered, and how we support you

1. What you have

YOUR SCOPE	YOUR SERVICE	USER GROUPS COVERED
The AI platforms and controls agreed in your SOW	Ongoing governance and support (see section 4)	All billable users

2. How to reach us

a. Raise the ticket through your normal Atlas Support route. No special AI channel needed.
b. Support picks it up, with our AI specialists engaged behind them as needed.
c. Standard response times from your agreement apply.

Sensitive data? If sensitive data has reached an AI tool it shouldn't have, flag this in the ticket. These cases are handled as a priority.

3. How we deliver SAIF Tier 1

Step	What happens
Kick-off	We confirm your scope, contacts, and how to reach us. You receive this document.
Implementation and validation	Your controls are implemented and tested before anything is switched to enforcement.
Handover to ongoing service	Your day-to-day contact is confirmed; your first review date is set.
Post-implementation check-in	Review what has changed, answer user questions, and confirm cadences.

4. What we do for you on an ongoing basis

Area	What we do for you	How often
Governance Review	Review approved/blocked AI usage, unapproved attempts, policy violations, and consumption trends	Quarterly
	Review M365 and related permissions affecting AI access/visibility	Quarterly
	Governance check-in covered in your quarterly business review (QBR)	Quarterly
Reporting & Oversight	Governance report: usage trends, notable activity, exceptions, recommended control adjustments	Quarterly
	Change notices for control or vendor changes affecting you	As it happens
Operational Support	Support/escalation for AI access, security, config, logging, and audit issues	Whenever you raise a ticket
	Incident handling for data-exposure/governance events, with updates until closed	As it happens, as a priority
Guardrail Tuning	Tune AI policies/configs as your needs, vendor capabilities, and regulations evolve	As needed, and at each review
	Recommendations on controls, access, config, and risk exposure	At each review, and when something changes

Important: how SAIF scope works

a. Ongoing SAIF governs and supports what has already been implemented.
b. New platforms, models, connectors, integrations, or significant expansions of AI use fall outside ongoing support. They require written scope and a separate implementation SOW or an approved written change order before work begins. Examples: adding a platform not in your SOW (e.g. Grok), moving into agentic tooling like Claude Cowork, or a major change to your controls.
c. To ask for more, raise it through your normal ticket route or account manager. We will confirm whether it is covered, a small change, or new work.

Need help with everyday AI-usage questions? Add fractional AI enablement (block of hours).

Atlas Technica is the premier technology partner for the alternative investment industry and the industry's first Managed Intelligence Provider (MIP), purpose-built for capital markets. Founded in 2016, Atlas Technica delivers enterprise-grade IT, cybersecurity, cloud, and AI so firms can place IT on our shoulders.