

Secure AI Toolkit for MSPs

Control AI before it becomes shadow IT

Version	Buyer-facing v0.3
Audience	MSP owners, service managers, vCIO/vCISO advisors, and technical leads
Position	\$97-\$197 self-serve security-first AI control kit
Support	Documentation + async FAQ only

Practical operational guidance only. Not legal advice, not a compliance certification, and not a substitute for client-specific risk assessment.

Version: Buyer-facing v0.3 Audience: MSP owners, service managers, vCIO/vCISO advisors, and technical leads
Product position: \$97-\$197 self-serve security-first AI control kit Support model: documentation + async FAQ only

Control AI before it becomes shadow IT

AI adoption is already happening inside MSPs and their clients. Staff are testing ChatGPT, Copilot, Gemini, Claude, AI meeting tools, browser extensions, and AI features built into everyday SaaS platforms.

The risk is not that AI exists.

The risk is uncontrolled AI use:

- sensitive data pasted into tools no one approved;
- browser extensions quietly reading business content;
- AI-generated advice reaching clients without review;
- staff using personal accounts for work data;
- automations being connected to email, files, tickets, billing, or client systems without proper controls.

The Secure AI Toolkit gives MSPs a practical starting system to bring AI use under control, reduce data leakage risk, and turn AI governance into a client-facing service opportunity.

This is not hype. This is operational control.

What this toolkit helps you do

Use this toolkit to:

- 1. identify where AI is already being used;
- 2. define what staff can and cannot put into AI tools;
- 3. publish a plain-English AI acceptable-use policy;
- 4. approve, restrict, or reject AI tools;
- 5. classify AI use cases by risk;
- 6. give staff safe prompt rules;
- 7. run an AI Security Scorecard;
- 8. package a simple AI readiness service for clients.

By the end, you should have a basic AI control system your MSP can use internally and a repeatable conversation you can take to clients.

Who this is for

This toolkit is for MSPs that:

- already have staff or clients experimenting with AI;
- need a written AI use policy quickly;
- want to reduce sensitive-data leakage risk;
- want a security-led way to discuss AI with clients;
- want to package AI governance or AI readiness as a paid service;
- prefer self-serve templates over live coaching, custom consulting, or complex software builds.

Who this is not for

This toolkit is not for teams that want:

- legal advice;
- a compliance certification;

- a done-for-you implementation;
- custom AI software development;
- unrestricted autonomous agents connected to production systems;
- a prompt pack that ignores governance, data handling, and vendor control.

Important disclaimer

This toolkit is practical operational guidance. It is not legal advice, a compliance certification, or a substitute for client-specific risk assessment.

You should adapt it to your own contracts, privacy obligations, client requirements, regulated-industry needs, vendor terms, and risk appetite.

Quick Start

Day 1 — Stop uncontrolled risk

- 1. Assign an AI owner.
- 2. List the AI tools staff are already using.
- 3. Pause any use involving secrets, credentials, regulated data, security incident details, or sensitive client data.
- 4. Share the Prompt Safety Rules with staff.
- 5. Run the AI Security Scorecard.

Week 1 — Put controls in writing

- 1. Publish the AI Acceptable Use Policy.
- 2. Define the four data classes: Public, Internal, Confidential, Restricted / Sensitive.
- 3. Create the AI Tool Approval Register.
- 4. Classify common AI use cases as Approved, Restricted, or Prohibited.
- 5. Train staff on what must never be pasted into AI tools.

Month 1 — Turn governance into a client service

- 1. Run a client-facing AI readiness discussion.
- 2. Use the discovery questions to identify workflow opportunities.
- 3. Package a simple AI Risk and Workflow Audit.
- 4. Create a 30-day roadmap for one low-risk AI use case.
- 5. Decide whether your next offer is AI governance, fixed workflow automation, hosted assistants, or a reseller-style AI service package.

1. Secure AI Operating Model

Core principle

AI is a business system. It needs ownership, data boundaries, approved use cases, vendor control, human validation, and monitoring.

Control layer 1 — Ownership

Every AI tool and workflow needs an accountable owner.

Minimum standard:

- assign one business owner who approves use cases;
- assign one technical or security owner who reviews access, data, and vendor risk;
- record one review date for each approved tool or workflow.

Questions to answer:

- Who approves new AI tools?
- Who decides what data can be used?
- Who reviews risky outputs before they reach clients?
- Who can disable an unsafe AI tool or automation?

Control layer 2 — Data boundaries

Users need to know which data can and cannot be used with AI.

Minimum standard:

- Public data may be used in approved tools.
- Internal data may be used in approved tools with suitable privacy settings.
- Confidential data requires a business need, approved tool, and minimisation.
- Restricted / Sensitive data is prohibited unless an explicitly approved secure workflow exists.

Questions to answer:

- What data must never leave controlled systems?
- Which client data can be used only with explicit approval?
- Which tools are approved for each data class?
- What redaction rules should staff follow?

Control layer 3 — Tool approval

AI tools should not be adopted just because they are convenient.

Minimum standard:

- review data retention and training settings;
- require MFA or SSO where available;
- check admin controls;
- check audit logs;
- check export and deletion processes;
- separately review browser extensions, plugins, file connectors, email connectors, ticketing integrations, RMM/PSA integrations, and identity integrations.

Questions to answer:

- Does the vendor train on customer data?
- Can admins control users and settings?
- Are logs available?
- What happens when a user leaves?
- What connected systems can the AI tool access?

Control layer 4 — Human validation

AI output is draft work, not authority.

Minimum standard:

- client-facing messages require human review;
- security advice requires qualified review;
- technical changes require testing;
- factual claims require source checking;
- legal, HR, financial, medical, and compliance-sensitive outputs require specialist review.

Questions to answer:

- What outputs can be used internally as drafts?
- What outputs need review before client use?
- What outputs are too risky for AI assistance?
- How should staff label assumptions and uncertainty?

Control layer 5 — Automation control

AI automation carries higher risk than manual prompting.

Minimum standard:

- least-privilege access;
- human approval for high-impact actions;
- logs for AI actions;
- rollback or disable procedure;
- testing before production.

Questions to answer:

- Can the AI send messages?
- Can it update tickets, files, billing, identity, or client systems?
- Can it trigger irreversible actions?
- Is there a clear stop button?
- Are logs reviewed?

Five operating rules

- 1. Classify the data before using AI.
- 2. Use approved tools only.
- 3. Never paste secrets, credentials, private keys, session tokens, MFA recovery codes, regulated data, or confidential client data into unapproved AI systems.
- 4. Treat AI output as draft until reviewed by a human.
- 5. Do not connect AI to production systems, email, ticketing, billing, RMM/PSA, identity, code repositories, or client environments without explicit approval and least-privilege controls.

2. AI Acceptable Use Policy Template

Use this as a starting policy. Adapt it to your own organisation, clients, contracts, and legal obligations.

Purpose

This policy defines how staff may use AI tools safely while protecting client data, company data, credentials, intellectual property, regulated information, and business operations.

Scope

This policy applies to:

- employees;
- contractors;
- temporary staff;
- company-owned devices;
- personal devices used for work;
- public AI tools;
- embedded AI features in SaaS platforms;
- browser extensions;
- copilots;
- chatbots;
- automation tools;
- custom AI assistants and agents.

Approved uses

Staff may use approved AI tools for:

- drafting internal notes, outlines, summaries, and first-pass documentation;
- brainstorming non-sensitive ideas;
- rewriting non-confidential content;
- creating generic checklists and templates;
- summarising content that has been cleared for AI use;
- generating draft scripts or code where reviewed and tested before use;
- producing internal draft communications where a human reviews before sending.

Restricted uses

The following require manager, security, or nominated AI owner approval:

- client-facing responses;
- security recommendations;
- legal, HR, financial, compliance, or contract-related content;
- uploading internal documents;
- uploading client documents;
- using AI to assess vulnerabilities, incidents, or security architecture;
- automation that sends messages or updates systems;
- connecting AI to email, ticketing, CRM, file storage, code repositories, RMM/PSA, SIEM, EDR, billing, identity, backup, monitoring, or client systems.

Prohibited inputs

Do not input or upload:

- passwords;
- API keys;
- private keys;
- tokens;
- session cookies;
- MFA recovery codes;
- client confidential information unless explicitly approved and protected;
- personal information not cleared for the tool;
- regulated data unless approved under a documented secure process;
- security incident details that could expose client weaknesses;
- vulnerability details before remediation or disclosure approval;
- proprietary pricing, contracts, or commercial strategy unless approved;
- private source code unless tool/vendor approval covers it.

Prohibited uses

Do not use AI to:

- make final hiring, firing, disciplinary, legal, financial, medical, or compliance decisions;
- send client communications without human review;
- execute production changes without approval;
- bypass internal approval processes;
- generate phishing, malware, credential theft, evasion, or harmful content;
- impersonate clients, staff, vendors, or regulators;
- conceal AI use where disclosure is required by policy, contract, or law.

Human review rule

AI output must be reviewed before it is used for:

- client advice;
- security decisions;
- technical changes;
- public content;
- legal, financial, HR, medical, or compliance-sensitive material;
- automation logic;
- code, scripts, firewall rules, identity changes, or production runbooks.

Incident reporting

Report immediately if:

- sensitive data was pasted into an unapproved tool;
- AI generated harmful, false, or unsafe advice that reached a client;
- an AI tool gained unintended access to systems or data;
- an automation behaved unexpectedly;
- a staff member installed an unapproved AI browser extension or plugin;

- AI output caused or almost caused a client-impacting incident.

3. AI Data Classification and Handling Rules

Data class	Examples	AI handling rule	Approval requirement
Public	Public website copy, public vendor docs, published marketing material, public job ads	Allowed in approved AI tools	User discretion within policy
Internal	Internal SOPs, process notes, non-sensitive meeting notes, generic internal training material	Allowed in approved AI tools if privacy/retention settings are acceptable and no client-confidential data is included	Team lead or AI owner if unsure
Confidential	Client names/environments, private client communications, internal financials, sales pipeline, commercial strategy, non-public security findings	Restricted. Requires approved tool, business need, and data minimisation	AI owner plus business owner
Restricted / Sensitive	Passwords, API keys, private keys, tokens, session cookies, MFA recovery codes, regulated personal information, incident response details, unreleased vulnerability details, legal matters, client security architecture	Prohibited unless explicitly approved under a documented secure workflow	Security owner plus executive/business approval

Data minimisation checklist

Before using AI, ask:

- Can I remove names, domains, IPs, usernames, ticket numbers, and identifiers?
- Can I generalise the scenario?
- Can I use synthetic data?
- Can I ask for a framework instead of uploading source data?
- Can I summarise the problem without including client-specific details?
- Is this tool approved for this data class?
- Is there a contractual, privacy, or regulatory restriction on this data?

Redaction examples

Risky input:

Client ACME Corp at 10.20.30.0/24 has exposed RDP on host srv-finance-01.

Safer prompt:

A client has an exposed remote access service on an internal server. Draft a general remediation checklist and communication template. Do not assume technical details not provided.

Risky input:

Here is a client incident report with names, IPs, EDR alerts, and suspected attacker behaviour.

Safer prompt:

Create an incident report structure for a suspected endpoint compromise. Include sections for timeline, scope, evidence, containment, eradication, recovery, and client communication.

4. Approved / Restricted / Prohibited AI Use Cases

Use this register to classify common AI activities before staff use them.

Use case	Default status	Notes / minimum control
Brainstorming non-sensitive ideas	Approved	Use approved tools only
Rewriting public or non-confidential content	Approved	Human review before publishing
Drafting internal SOPs from non-sensitive notes	Approved	Human review; label assumptions
Summarising public vendor documentation	Approved	Check source accuracy
Drafting client-facing emails	Restricted	Human review required before sending
Drafting security recommendations	Restricted	Qualified technical/security review required
Summarising internal documents	Restricted	Approved tool; confirm document classification first
Generating scripts or code	Restricted	Review, test, and approval before use
Connecting AI to email, ticketing, CRM, RMM/PSA, SIEM, EDR, billing, identity, or file storage	Restricted	Separate approval; least privilege; logs; rollback
Uploading client data	Restricted / normally prohibited	Requires explicit client/data approval and approved secure workflow
Uploading secrets, tokens, private keys, passwords, session cookies, MFA recovery codes	Prohibited	Never paste into AI
Uploading regulated data without approved process	Prohibited	Requires documented secure workflow
Autonomous production changes	Prohibited by default	Only allowed under separately approved controlled automation
Generating malware, phishing, credential theft, or evasion content	Prohibited	Security testing exceptions require formal authorization and controlled process

5. AI Tool Approval Register

Use one row per AI tool. Review browser extensions, plugins, connectors, and integrations separately from the base AI product.

Field	Entry
Tool name	
Vendor	
Business owner	
Technical/security owner	
Intended users	
Approved use cases	
Prohibited use cases	

Field	Entry
Data classes allowed	Public / Internal / Confidential / Restricted
Retention/training setting reviewed	Yes / No
Training on customer data disabled where available	Yes / No / Not available
SSO/MFA available	Yes / No
Admin controls available	Yes / No
Audit logs available	Yes / No
Export/delete process known	Yes / No
Contract/DPA reviewed	Yes / No / Not required
Connected systems/integrations	None / Email / Files / Ticketing / CRM / Code / RMM-PSA / Identity / Other
Approval status	Approved / Restricted / Pilot / Rejected
Review date	
Review owner	

Approval decision guide

Approve when:

- the business use case is clear;
- the allowed data classes are understood;
- retention and training settings are acceptable;
- MFA, SSO, and admin controls are sufficient for the use case;
- users know the handling rules.

Restrict when:

- the tool is useful but data risk is unclear;
- only some teams should use it;
- only Public or Internal data is allowed;
- integrations need separate approval;
- client-facing use needs additional review.

Reject when:

- the tool lacks required controls;
- vendor or data terms are unacceptable;
- users want to paste sensitive or regulated data without a secure workflow;
- the tool introduces uncontrolled plugins, browser access, or system actions.

6. AI Risk Checklist for MSPs

Governance

- ☐ AI owner assigned.
- ☐ Approved AI tools documented.
- ☐ AI acceptable-use policy published.

- ☐ Staff know what data cannot be entered into AI.
- ☐ Client-facing AI use is explicitly approved before use.
- ☐ Tool reviews have owners and review dates.

Data security

- ☐ Data classifications are defined.
- ☐ Restricted data is prohibited from public AI tools.
- ☐ Client data use has explicit approval.
- ☐ Knowledge-base sources are reviewed before upload.
- ☐ Data deletion/export process is known for approved tools.
- ☐ Redaction guidance is available to staff.

Tool and vendor control

- ☐ AI tools are reviewed before adoption.
- ☐ MFA or SSO is used where available.
- ☐ Admin ownership is not tied to one employee.
- ☐ Browser extensions and plugins are controlled.
- ☐ AI tools connected to email, files, ticketing, or identity are separately approved.
- ☐ Vendor retention and training settings are reviewed.

Output validation

- ☐ AI output is treated as draft.
- ☐ Technical and security recommendations require human review.
- ☐ Client communications require review before sending.
- ☐ AI-generated code and scripts are reviewed and tested before use.
- ☐ Sources are checked for factual claims.
- ☐ Assumptions and missing information are labelled.

Automation and agents

- ☐ AI automations have least-privilege access.
- ☐ High-impact actions require human approval.
- ☐ Logs are available for AI actions.
- ☐ Rollback or disable process is documented.
- ☐ Automations are tested before production use.
- ☐ Production access is separated from experimentation.

Client service readiness

- ☐ Discovery questions are ready.
- ☐ A fixed assessment scope is defined.
- ☐ A 30-day roadmap template exists.
- ☐ Pricing and deliverables are documented.
- ☐ Out-of-scope custom work is explicitly excluded.

7. Prompt Safety Rules

- 1. Never paste secrets.
- 2. Never paste unredacted client data into unapproved tools.
- 3. Ask for frameworks, checklists, or generic examples when possible.
- 4. Use placeholders for names, domains, IPs, usernames, and identifiers.
- 5. Tell the AI what not to include.
- 6. Require the AI to label assumptions.
- 7. Require human review for technical and security advice.
- 8. Do not rely on system prompts as security controls.
- 9. Do not let prompts override policy.
- 10. Treat AI output as untrusted until reviewed.
- 11. Do not connect prompts directly to production actions without separate approval.
- 12. Do not use AI output as evidence unless the source is verified.
- 13. Do not use AI to hide uncertainty from clients.
- 14. Do not allow a prompt to bypass data classification rules.
- 15. Keep client-identifying details out unless the workflow is explicitly approved.

Safe prompt pattern

Use this structure:

- 1. Role — who the AI should act as.
- 2. Task — what output is needed.
- 3. Context — only the minimum safe context.
- 4. Constraints — data, security, style, and scope limits.
- 5. Verification — ask the AI to label assumptions and checks.
- 6. Output format — specify the table, checklist, email, SOP, or report structure.

Example:

You are assisting an MSP service manager. Create a general client-facing checklist for safe AI adoption. Use plain English. Do not include legal claims, compliance guarantees, or technical assumptions. Include assumptions and a section titled "What a human must verify".

8. Secure Prompt Vault v1

Policy drafting prompt

You are helping draft an internal AI acceptable-use policy for an MSP. Use plain English. Include approved uses, restricted uses, prohibited uses, data-handling rules, human review requirements, and incident reporting steps. Do not include legal claims or compliance certifications. Ask clarifying questions if required. Label assumptions.

Data classification prompt

Create a data classification table for an MSP using Public, Internal, Confidential, and Restricted categories. For each category, list examples, AI handling rules, approval requirements, and redaction guidance. Include warnings for credentials, client data, incident details, regulated personal information, private source code, and security architecture.

Workflow audit prompt

Create an AI workflow audit for an MSP. Identify repetitive tasks across sales, service desk, reporting, documentation, onboarding, admin, finance, and client communication. Score each task by time saved, data sensitivity, client impact, implementation difficulty, and suitability for AI assistance. Recommend the first three low-risk use cases.

Client discovery prompt

Create a client discovery questionnaire for a secure AI readiness assessment. It should identify current AI tool use, sensitive-data risks, workflow bottlenecks, approved tools, staff behaviour, client-facing use cases, business systems connected to AI, and automation opportunities. Group questions by governance, data, tools, workflows, and client service opportunity.

SOP conversion prompt

Convert the following process notes into a structured SOP with purpose, scope, prerequisites, roles, step-by-step procedure, exception handling, quality checks, rollback/stop conditions, and review date. Do not invent missing technical steps. Label assumptions and missing information.

Output validation prompt

Review the following AI-generated output as a security-conscious MSP. Identify unsupported claims, risky recommendations, missing assumptions, sensitive-data exposure, client-impacting issues, and what a human should verify before use. Return a table with issue, severity, reason, and recommended fix.

Client email drafting prompt

Draft a client email explaining a proposed AI readiness assessment. Use plain English. Avoid hype. Emphasise safe adoption, data protection, staff behaviour, workflow value, and practical next steps. Do not promise compliance certification or guaranteed business outcomes.

AI tool review prompt

Create a vendor review checklist for an AI tool. Include data retention, training on customer data, SSO/MFA, admin controls, audit logs, integrations, export/delete process, contractual review, approved data classes, and review cadence.

Low-risk workflow selection prompt

Given a list of workflows, rank them for first AI implementation. Score each by time saved, sensitivity of data, client impact, repeatability, reversibility, and human-review ease. Recommend 1-3 low-risk workflows and explain why higher-risk workflows should wait.

9. Client-Facing AI Discovery Questions

Use these questions for a lightweight AI readiness discussion. The goal is to find risk, value, and safe first use cases without turning the engagement into custom consulting.

Current AI use

- 1. What AI tools are staff currently using?
- 2. Are staff using personal accounts, business accounts, or both?
- 3. Are browser extensions, meeting bots, copilots, or embedded AI features in use?
- 4. Are staff allowed to paste client data into AI tools?
- 5. Has any AI use already reached client-facing work?

Policy and ownership

- 6. Do you have a written AI use policy?
- 7. Who approves new AI tools?
- 8. Who owns AI risk in the business?
- 9. Are staff trained on what data cannot go into AI?
- 10. Is there a process to report unsafe AI use?

Data and vendor risk

- 11. What data must never leave your controlled environment?
- 12. Do you know whether AI vendors use your data for training?
- 13. Are AI tools reviewed for retention, audit logs, admin controls, and deletion/export options?
- 14. Do any AI tools connect to email, files, CRM, ticketing, billing, identity, or source code?
- 15. Are client contracts or privacy obligations relevant to AI use?

Workflow opportunity

- 16. Which workflows consume the most repetitive admin time?
- 17. Which client-facing tasks are most delayed by manual work?
- 18. Which internal documentation or SOP gaps slow the team down?
- 19. Which reports, summaries, or communications are repeatedly drafted by humans?
- 20. Which AI use case would create the fastest measurable business value?

Safe first implementation

- 21. Which workflows use only Public or Internal data?
- 22. Which workflows can keep a human approval step?
- 23. Which workflows are easy to reverse if output is poor?
- 24. Which workflows have clear examples of good output?
- 25. Which workflows should be explicitly avoided for now?

10. MSP AI Service Packaging Starter Kit

Positioning

Position this as a practical security and workflow service, not generic AI coaching.

Core message:

We help you understand where AI is already being used, reduce sensitive-data risk, and identify safe workflow opportunities before AI becomes uncontrolled shadow IT.

Service name ideas

- Secure AI Readiness Assessment
- AI Risk and Workflow Audit
- Secure AI Operations Starter Kit
- Managed AI Governance Pack
- AI Business Systems Setup
- AI Use Policy and Risk Review

- AI Shadow IT Discovery

Starter service package

Suggested offer:

Secure AI Readiness Assessment

Suggested price:

\$750-\$2,500 depending on client size and depth.

Suggested scope:

- one AI tool inventory;
- one AI acceptable-use policy draft;
- one data-handling rules table;
- one risk-ranked use-case list;
- one workflow opportunity map;
- one secure prompt starter pack;
- one 30-day implementation roadmap.

Explicit exclusions

Exclude these from the starter package unless sold separately:

- custom software development;
- complex workflow automation;
- production system integrations;
- legal or compliance certification;
- 24/7 AI monitoring;
- custom model training;
- regulated-data processing design;
- incident response.

Upsell paths

- Ongoing AI governance review.
- Approved AI tool rollout.
- SOP and knowledge-base cleanup.
- Fixed workflow automations.
- Hosted assistant from a fixed catalog.
- MSP reseller license for repeat delivery.

Client-facing deliverable outline

- 1. Executive summary.
- 2. Current AI tool inventory.
- 3. Key data-risk observations.
- 4. Policy gaps.
- 5. Approved / restricted / prohibited use cases.
- 6. Low-risk workflow opportunities.

- 7. High-risk workflows to defer.
- 8. 30-day roadmap.
- 9. Recommended next service.

11. AI Security Scorecard Operating Guide

The AI Security Scorecard is the self-serve diagnostic that replaces a sales call.

Scorecard shape

- 30 questions.
- 10 domains.
- 0-90 maximum score.
- Four risk bands.
- Product-routing recommendation.

Risk bands

Score	Band	Meaning	Route
0-29	Critical / uncontrolled AI adoption	AI use likely lacks ownership, policy, data boundaries, and vendor control	Secure AI Toolkit immediately
30-54	Exposed / inconsistent controls	Some practices exist but are inconsistent across tools, data, users, and workflows	Toolkit + self-paced course waitlist
55-74	Developing / ready to systemise	Foundations exist; the MSP is ready to package governance and workflow services	Course + reseller license waitlist
75-90	Advanced / reseller-ready	Strong baseline; candidate for reseller system and fixed hosted assistant catalog	Reseller license waitlist

Follow-up rule

Use the scorecard to recommend the next self-serve product path. Do not turn it into a custom diagnostic call.

Customer-facing buttons should remain limited to:

- 1. Download PDF report.
- 2. Print / Save PDF.
- 3. Reset.

Any owner copy or lead notification should happen in the background through the hosted endpoint when available.

12. 30-Day Implementation Checklist

Day 1

- ☐ Assign AI owner.
- ☐ List current AI tools.
- ☐ Pause unapproved sensitive-data use.

- ☐ Publish prompt safety rules.
- ☐ Run AI Security Scorecard.

Week 1

- ☐ Publish AI acceptable-use policy.
- ☐ Define data classifications.
- ☐ Create AI Tool Approval Register.
- ☐ Approve/restrict/reject common AI tools.
- ☐ Train staff on data-handling rules.

Week 2

- ☐ Run workflow audit.
- ☐ Select 1-3 low-risk AI workflows.
- ☐ Create prompt library.
- ☐ Define human-review requirements.
- ☐ Document excluded high-risk workflows.

Week 3

- ☐ Package Secure AI Readiness Assessment.
- ☐ Draft client-facing discovery script.
- ☐ Create 30-day roadmap template.
- ☐ Define pricing and out-of-scope boundaries.
- ☐ Identify first 10 clients/prospects to approach.

Month 1

- ☐ Review internal AI usage.
- ☐ Close policy gaps.
- ☐ Document first client-safe AI service offer.
- ☐ Identify next automation candidate.
- ☐ Decide whether the next offer is governance review, fixed automation, hosted assistant, or reseller licensing.

Source alignment note

This toolkit is aligned to public guidance themes from NIST AI RMF, NIST Generative AI Profile / NIST AI 600-1, OWASP Top 10 for LLMs and Gen AI Apps, and ASD ACSC AI security/data-security guidance.

It is not legal advice, a compliance certification, or a substitute for client-specific risk assessment. MSPs should adapt it to client contracts, local privacy obligations, regulated-industry requirements, vendor terms, and their own risk appetite.

Version notes

v0.3 changes:

- converted the prior internal draft into cleaner buyer-facing product copy;
- removed internal draft framing from the opening section;

- sharpened the core promise around shadow AI control;
- improved buyer fit and exclusion language;
- made the quick-start sequence more direct;
- preserved the security-first guardrails, self-serve delivery model, and no-custom-work positioning;
- retained all major templates and operational sections needed for the product bundle.