AI Assurance & Control Validation Assessment

October 20, 2026

This work was performed under contract to:

Larkspur Financial Partners — fictional sample
A2 Advanced engagement — illustrative sample report

ioSENTRIX
Trusted by Security Conscious Companies
https://iosentrix.com
contact@iosentrix.com

CREST PEN TEST

Revision History

Version	Modification	Date	Author
0.9	Draft issued for client review	14/10/2026	ioSENTRIX
1.0	Final report delivered	20/10/2026	ioSENTRIX

Document Confidentiality Statement

The information in this document is confidential to the person to whom it is addressed and should not be disclosed to any other person. It may not be reproduced in whole, or in part, nor may any of the information contained therein be disclosed without the prior consent of **ioSENTRIX LLC** ('the Company'). A recipient may not solicit, directly or indirectly (whether through an agent or otherwise) the participation of another institution or person without the prior approval of the directors of the Company.

Any form of reproduction, dissemination, copying, disclosure, modification, distribution and or publication of this material is strictly prohibited.

The information in this documentation is subject to change without notice and should not be construed as a commitment by the Company. The Company makes no representations or warranties, express or implied, with respect to the documentation and shall not be liable for any damages, including any indirect, incidental, consequential damages (such as loss of profit, loss of use of assets, loss of business opportunity, loss of data or claims for or on behalf of user's customers), that may be suffered by the use.

Copyright © 2026, ioSENTRIX LLC.

Table of Contents

Revision History **2**
Document Confidentiality Statement **2**
Table of Contents **3**
Executive Summary **4**
Dashboard **6**
1 Introduction **8**
2 Methodology **8**
2.1 The AI Assurance Method 8
2.2 Assurance states and evidence model 8
2.3 Severity rankings 9
2.4 Scope 10
3 Assessment Limitations **11**
4 AI Discovery: What Is Actually in Use **11**
5 Policy versus Reality **13**
6 Control Baseline by Domain **14**
7 Technical Validation **15**
7.1 Data-leakage validation matrix 15
7.2 Browser extensions and desktop assistants 16
7.3 Detection and alerting chain 16
7.4 Microsoft 365 Copilot configuration 17
7.5 ChatGPT Enterprise workspace configuration 17
8 Technical Risks **18**
8.1 Summary of findings 18
9 Findings **19**
9.1 Critical Severity Risks 19
9.1.1 AI-DLP-01 · Personal AI accounts indistinguishable from enterprise 19
9.2 High Severity Risks 20
9.2.1 AI-DLP-02 · File uploads to AI services are not inspected 20
9.2.2 AI-END-01 · AI browser extension bypasses all data controls 21
9.2.3 AI-COP-01 · Copilot surfaces over-shared SharePoint content 23
9.3 Medium Severity Risks 24
9.3.1 AI-DLP-03 · Unapproved AI services outside the blocked category 24
9.3.2 AI-MON-01 · AI activity not retained, forwarded or alerted 25
9.3.3 AI-GOV-01 · AI policy lacks enforcement map and attestation 26
9.3.4 AI-TEN-01 · ChatGPT Enterprise lifecycle and compliance logging gaps 27
9.4 Low Severity Risks 28
9.4.1 AI-END-02 · Desktop AI assistant outside software allowlist 28
9.5 Informational Severity Risks 29
9.5.1 AI-VEN-01 · Embedded SaaS AI features not reviewed 29
10 Framework Traceability **30**
11 Residual Risk and Remediation Roadmap **31**
About ioSENTRIX **33**

Copyright © 2026, ioSENTRIX LLC.

Executive Summary

ioSENTRIX was engaged by Larkspur Financial Partners, a regional wealth-management and insurance brokerage with approximately 850 staff, to answer a question its AI Acceptable Use Policy asserts but had never been tested: can customer, financial or employee data reach an AI service it should not, and would anyone know? The engagement was delivered as an A2 Advanced assessment over four weeks (14 September – 9 October 2026) and covered AI discovery across the whole organization, a control baseline against ioSENTRIX's 137-control AI assurance catalog, technical validation with synthetic data, and configuration review of the two sanctioned enterprise AI tenants (Microsoft 365 Copilot and ChatGPT Enterprise).

Larkspur has a policy; it does not yet have the gates to enforce it. Discovery across proxy, identity, endpoint and procurement evidence found 27 distinct AI services in use against a sanctioned list of three. 412 of 850 staff (48%) touched an unapproved AI service in the 30-day window and 118 submitted content to one. Of the nine operative clauses in the Acceptable Use Policy, two are enforced by a technical control that held under test, two are partially enforced, four have a control that failed under test, and one has no technical control behind it at all.

Technical validation with synthetic customer records confirmed the gap. Pasted text containing customer-data markers was correctly blocked at consumer AI services, but the same records uploaded as files were not inspected and reached three unapproved services. A personal ChatGPT account used on a managed laptop was indistinguishable from the sanctioned ChatGPT Enterprise tenant to the web proxy and DLP, so personal accounts can receive any data the enterprise tenant can. An AI browser extension installed on 140 endpoints sent page content to a third-party API that no control inspected. None of the blocked or allowed events generated an alert.

ioSENTRIX identified ten (10) findings: one (1) Critical, three (3) High, four (4) Medium, one (1) Low and one (1) Informational, consolidated from seventeen (17) controls that failed validation. Of 118 applicable controls, 31 reached the Assured state.

The encouraging news is that most fixes are configuration, not new tooling: Larkspur already owns the SSE proxy, DLP, Purview and endpoint management needed to close nine of the ten findings. ioSENTRIX recommends a 90-day program that enforces tenant and account restrictions at the proxy, extends DLP to file uploads, brings Copilot's data exposure under control, and wires AI events into monitoring, followed by Continuous AI Assurance to keep the inventory and the validation current as tools and policy change.

Confidential
Copyright © 2026, ioSENTRIX LLC.

Key findings

- **Personal AI accounts are invisible (Critical).** A personal ChatGPT login on a managed device passed through the same proxy allow-rule as the enterprise tenant; synthetic customer records were uploaded without a block, a log distinction or an alert.
- **File uploads are not inspected (High).** DLP inspects pasted text but not file uploads; synthetic customer files reached Gemini, Claude and Perplexity.
- **A browser extension bypasses every control (High).** An AI summarizer extension on 140 endpoints sends page content to an API domain outside the AI web category; nothing blocks, inspects or logs it.
- **Copilot surfaces over-shared content (High).** A test user retrieved synthetic HR and client files from SharePoint sites shared with "Everyone except external users"; Restricted SharePoint Search is off and sensitivity labels cover 12% of sites.
- **Nobody is watching (Medium).** AI activity is logged for 30 days at the proxy only, never forwarded to the SIEM, and no alert rule exists for blocked data-leakage attempts.

Strategic recommendations

- Configure tenant and instance restrictions at the SSE proxy so that only the enterprise ChatGPT workspace and the corporate Microsoft tenant are reachable; block or redirect personal logins.
- Extend DLP inspection to file uploads and to the full AI web category, including services not yet categorized; close the extension gap with a managed-browser extension allowlist.
- Enable Restricted SharePoint Search until oversharing is remediated; apply sensitivity labels to client and HR sites; restrict Copilot Studio agent creation.
- Forward AI activity logs to the SIEM, retain for 12 months, and alert on blocked data-leakage attempts and on first use of any new AI service.
- Add an enforcement map and annual attestation to the AI Acceptable Use Policy so every clause names the control that enforces it.

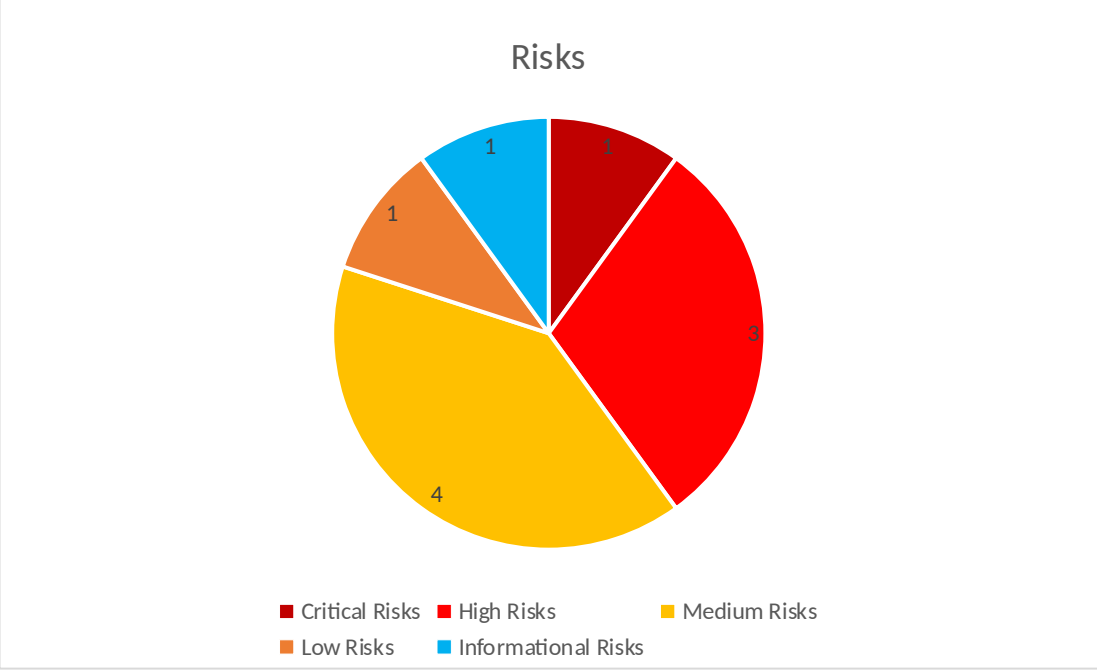
Confidential
Copyright © 2026, ioSENTRIX LLC.

Dashboard

Name	AI Assurance & Control Validation Assessment (A2 Advanced)
Method	ioSENTRIX AI Assurance Method (Discover → Map → Baseline → Validate → Assure); NIST risk severity matrix
Dates	14 September – 9 October 2026 (4 weeks)
Targets	Enterprise AI usage across all departments (850 staff); AI Acceptable Use Policy v1.2; Microsoft 365 Copilot and ChatGPT Enterprise tenants; SSE proxy, DLP, endpoint and browser controls
Depth	137-control catalog, 118 applicable · 36 controls technically validated with synthetic data · 27 AI services discovered

Findings Breakdown (Chart)

Critical Priority Risks	**1**
High Priority Risks	**3**
Medium Priority Risks	**4**
Low Priority Risks	**1**
Informational	**1**
Total Findings	**10**



Copyright © 2026, ioSENTRIX LLC.

Assurance State Distribution (118 applicable controls)

Assurance state	Count	Executive reading
Assured	31	Works and proven to work (technical validation passed)
Implemented — not proven	34	Present and configured; looks protected, not tested
Defined — on paper only	19	Documented, not shown operating
Asserted — unverified	9	Claimed in interview only
Failing — proven exposed	**17**	**Tested and does not work; remediate now**
Absent — missing	8	Control not present

Discovery and Validation Indicators (read together, never as a lone number)

AI services observed / sanctioned	Staff using unapproved AI (30 days)	Policy clauses enforced / partial / failing / none	Technical validations / failed	Alerts raised during testing
27 / 3	**412 of 850**	**2 / 2 / 4 / 1**	**36 / 17**	**0**

Confidential
Copyright © 2026, ioSENTRIX LLC.

1 Introduction

Larkspur Financial Partners engaged ioSENTRIX to perform an AI Assurance & Control Validation assessment: to discover what AI is actually in use across the organization, to compare that reality against the AI Acceptable Use Policy and the controls that are supposed to enforce it, and to test, with safe synthetic data, whether those controls hold. Fieldwork ran from 14 September to 9 October 2026.

Larkspur does not build AI systems. It consumes them: Microsoft 365 Copilot and ChatGPT Enterprise are sanctioned for all staff, GitHub Copilot for the technology team, and a growing number of AI features inside existing SaaS tools. The risk this assessment addresses is therefore not model security but data handling: customer financial records, employee data and client documents moving into AI services under terms, controls and visibility that Larkspur may or may not have.

Unlike a governance questionnaire, this assessment does not stop at whether a policy exists. Every material control is graded on the strength of the evidence behind it, and the controls that matter most were tested by attempting the very thing the policy prohibits.

2 Methodology

2.1 The AI Assurance Method

ioSENTRIX assesses AI risk in five phases. At A2 Advanced depth, validation is control validation: ioSENTRIX attempts what the policy prohibits, using synthetic data, and records whether each control blocks, logs or alerts. Adversarial red-teaming of AI systems is a separate service and was not in scope.

Phase	Question answered	Primary output
1 Discover	What AI is sanctioned, observed, unknown or unapproved, and by whom?	AI inventory with confidence rating
2 Map	How is AI used; what data moves into it; which policy clauses apply?	Use-case register; policy-to-control map
3 Baseline	Are the governance and security controls designed and present?	Evidence-graded control assessment (137-control catalog)
4 Validate	Do the controls actually work when someone tries to break the policy?	Validation matrix with Pass / Partial / Fail per tool and path
5 Assure	What is the residual risk, and what gets fixed first?	Findings, residual-risk view and 30/60/90-day roadmap

2.2 Assurance states and evidence model

ioSENTRIX separates maturity from proof. Each control receives a maturity level (0–5) and an evidence grade from E0 (asserted) to E3 (technically validated); the effective score is the lower of the

Copyright © 2026, ioSENTRIX LLC.

two, so a confident interview cannot produce a high assurance rating. Each applicable control resolves to one of six states:

State	Meaning	Evidence required
Assured	Works and proven to work	Technical validation passed (E3)
Implemented — not proven	Present and configured, not tested	Configuration or demonstration (E2)
Defined	Documented, not shown operating	Policy or procedure (E1)
Asserted	Claimed in interview only	Assertion (E0)
Failing	Tested and does not work	Validation failed
Absent	Not present	No control

2.3 Severity rankings

ioSENTRIX ranks the severity of each finding using a method developed by the United States National Institute of Standards and Technology (NIST), calculating risk severity as a derived value of likelihood and impact. Likelihood is the probability of a given threat exploiting the weakness, incorporating ease of exploitation and the determination and capability of the threat agent; impact is the possible consequence of occurrence. The risk value is calculated using the table below.

		Impact				
		Critical	**High**	**Medium**	**Low**	**Informational**
Likelihood	**Critical**	Critical	High	Medium	Low	Informational
	High	Critical	High	Medium	Low	Informational
	Medium	High	Medium	Medium	Low	Informational
	Low	Medium	Low	Low	Low	Informational
	Informational	Low	Low	Informational	Informational	Informational

Table 1: NIST Risk Severity Matrix

The levels of risk severity have the following meanings:

- **Critical:** A critical risk is characterized by a very high likelihood and a high or very high impact. Remediation must be prioritized immediately and scheduled as the highest priority. Critical issues are considered blockers for new releases and must be resolved before deployment.
- **High:** High-severity issues present a substantial risk, though they may not block new releases. Remediation should be scheduled promptly and treated as a high priority.

Copyright © 2026, ioSENTRIX LLC.

- **Medium:** Medium severity issues require remediation. The organization should create and implement a remediation plan to address these risks effectively.
- **Low:** For low-severity issues, the organization may choose to either mitigate the risk or formally accept it, depending on its risk tolerance and strategic priorities.
- **Informational:** Informational findings have minimal or no direct impact on the target scope. However, they are reported as they may contribute to risk under specific circumstances, such as when combined with other vulnerabilities or technologies. Remediation is optional.

2.4 Scope

The scope of the engagement was limited to the organization, policies, tenants and controls listed below.

In scope	Coverage
Organization	All Larkspur departments and managed endpoints (850 staff, 4 offices, remote workforce)
Policy	AI Acceptable Use Policy v1.2 (June 2026) and the controls that enforce it
AI discovery	SSE/proxy logs, identity-provider sign-ins, endpoint software inventory, managed-browser extension inventory, procurement and expense records, SaaS admin consoles
Sanctioned AI tenants	Microsoft 365 Copilot (M365 E3 + Copilot add-on) and ChatGPT Enterprise workspace; GitHub Copilot (technology team)
Data protection controls	SSE web and DLP policies, endpoint DLP, Microsoft Purview sensitivity labels and DLP, email DLP
Endpoint and browser	Endpoint management policies, software allowlisting, Chrome and Edge extension management
Monitoring	Proxy and DLP logging, SIEM forwarding, alert rules, retention
Validation	36 controls validated with synthetic customer, employee and client-document records carrying ioSENTRIX markers

Out of scope

The following were considered out of scope for this assessment:

- Adversarial testing (red teaming) of AI systems; no AI system was attacked, only data-handling controls were tested.
- Security testing of AI systems built in-house (none exist at Larkspur).
- Personal and BYOD devices, and mobile applications of AI services.
- Implementation of remediation (available under a separately scoped statement of work).

Copyright © 2026, ioSENTRIX LLC.

- ISO/IEC 42001 certification: ioSENTRIX is not a certification body and this report does not constitute certification against any framework.
- Social engineering.

3 Assessment Limitations

The ever-changing technology landscape and the increasing sophistication of attacks against networked systems are reasons for which no entity can truthfully claim to identify all security issues, nor guarantee the lifetime security of an organization's AI estate. This point-in-time assessment was performed on a best-effort basis on the environment provided by Larkspur Financial Partners; changes to the environment may affect the applicability of these results. ioSENTRIX cannot guarantee 100% coverage for any security assessment.

> **Discovery is risk-based, not exhaustive**
> AI discovery is triangulated across independent evidence planes and reported with a confidence rating; it is a floor, not a ceiling. Validation used synthetic data generated by ioSENTRIX from Larkspur's record schemas: no Larkspur production customer, employee or client data was submitted to any AI service, and no Larkspur data was ingested into ioSENTRIX systems or any external AI service to perform this assessment. Known blind spots: personal and BYOD devices, mobile applications, and AI use on the two offices' guest wireless networks, which bypass the SSE proxy.

4 AI Discovery: What Is Actually in Use

Larkspur's sanctioned list contains three AI services. Discovery across five evidence planes over the 30-day window ending 30 September 2026 identified 27 distinct AI services in use from managed endpoints and corporate accounts. The table below groups them by category; a per-service register with user counts was delivered separately.

Sanctioned versus observed

Category	Services observed	Status	Notes
Enterprise AI (sanctioned)	3	Approved	ChatGPT Enterprise (612 active users), Microsoft 365 Copilot (734 licensed), GitHub Copilot (41)
Consumer chat assistants	5	Unapproved	Consumer ChatGPT (personal accounts), Gemini, Claude, Perplexity, Copilot (consumer). 263 users; paste blocked at two of five, file upload blocked at none
Writing and grammar tools	4	Unapproved	Free tiers of writing assistants; 171 users, concentrated in Client Services and Marketing
Meeting transcription	3	Unapproved	Bots joining client calls from personal accounts; 29 users in Advisory

Copyright © 2026, ioSENTRIX LLC.

Category	Services observed	Status	Notes
Browser extensions with AI	3	Unapproved	AI summarizer (140 endpoints), writing assistant (88), translation (22); installed without IT approval
Desktop AI assistants	1	Unapproved	Installed on 37 endpoints outside the software allowlist
AI features inside sanctioned SaaS	8	Status unknown	AI companions and summarizers in the video-conferencing, CRM, document-signing, design and note-taking platforms; terms not reviewed; several enabled by default

Who is using what

Department	Staff	Using unapproved AI	Dominant unapproved use
Advisory and Wealth Management	214	131 (61%)	Consumer chat for client-letter drafting; meeting bots on client calls
Client Services	188	104 (55%)	Writing assistants; consumer chat for correspondence
Insurance Brokerage	142	68 (48%)	Consumer chat for policy summaries; file uploads of policy documents
Finance and Operations	96	41 (43%)	Spreadsheet and reconciliation questions in consumer chat
Marketing	38	29 (76%)	Writing, image and design AI features
Technology	71	22 (31%)	Code assistants beyond GitHub Copilot
HR, Legal, Compliance	61	17 (28%)	Consumer chat for policy drafting; one HR file-upload event
Executive and other	40	—	No material unapproved use observed

AI inventory confidence: Moderate

Discovery reconciled the sanctioned list against SSE/proxy logs, identity-provider sign-ins, endpoint software and browser-extension inventories, and procurement records, with every discrepancy dispositioned. Confidence is Moderate rather than High because the SSE proxy does not cover the guest wireless networks in two offices, and because AI features embedded in sanctioned SaaS platforms are not visible in network telemetry; those were identified from admin consoles and vendor documentation only. Population reach: 94% of managed endpoints.

Confidential
Copyright © 2026, ioSENTRIX LLC.

5 Policy versus Reality

The AI Acceptable Use Policy v1.2 contains nine operative clauses. For each, ioSENTRIX identified the technical control that should enforce it and tested whether it does. A clause with no technical control depends entirely on staff reading and following the policy; the discovery results above show how well that works.

Policy clause (v1.2)	Enforcing control	Result	Evidence
§1 Sanctioned AI tools are accessed only through company SSO with MFA	Identity provider: SSO enforcement on both enterprise tenants	**Pass**	SSO enforced; password login rejected on both tenants (VAL-IAM-02, VAL-TEN-01)
§2 Only approved AI tools may be used (ChatGPT Enterprise, M365 Copilot, GitHub Copilot)	SSE web category block for "AI and ML"	**Partial**	Paste blocked at Gemini and Claude; Perplexity and the consumer Copilot not in category; 16 unapproved standalone services observed, plus 8 embedded SaaS AI features of unknown status
§3 No customer, financial or employee data in any AI tool other than the approved enterprise tenants	SSE DLP (text); endpoint DLP (file)	**Fail**	Pasted markers blocked; file uploads not inspected on any path (VAL-DLP-02)
§4 Personal accounts are prohibited on company devices	SSE tenant / instance restriction	**Fail**	Not configured; personal ChatGPT login indistinguishable from enterprise tenant (VAL-DLP-03)
§5 AI browser extensions require IT approval	Managed-browser extension allowlist	**Fail**	No extension policy in Chrome or Edge management; 3 AI extensions, 250 installations
§6 AI output used in client deliverables must be reviewed by a human	None (procedural)	**Asserted**	No workflow control, attestation or sampling; not testable
§7 Prompts and outputs are logged and may be reviewed	Proxy logging; ChatGPT Enterprise compliance logging	**Partial**	Proxy logs 30 days, not forwarded; compliance logging not enabled in the ChatGPT workspace
§8 Staff complete AI awareness training annually	LMS assignment and tracking	**Fail**	Course assigned June 2026; 38% completion at assessment date
§9 Company data in sanctioned tools must not be used to train vendor models	Enterprise tenant settings: training exclusion	**Pass**	Excluded by default on both tenants; settings documented (VAL-TEN-03)

Copyright © 2026, ioSENTRIX LLC.

Governance firms ask. We test.

Every clause above would have been marked "in place" by a questionnaire-based review, because the policy exists and management believes it is followed. Two of nine clauses are enforced by a control that held under test.

6 Control Baseline by Domain

Each domain is reported as a state mix rather than a single percentage, so the reader sees how much is proven, how much is assumed, and how much is failing. The eight domains below account for all 118 applicable controls (31 Assured, 34 Implemented, 19 Defined, 9 Asserted, 17 Failing, 8 Absent). Nineteen catalog controls that apply only to organizations building AI systems were marked not applicable.

Domain	State mix (Assured · Impl · Def · Asserted · Failing · Absent)	Assessment	Evidence conf.
Governance & Accountability (14)	4 · 3 · 4 · 2 · **1** · 0	Policy exists and is owned; no enforcement map, attestation or training completion	High
AI Discovery & Inventory (10)	2 · 2 · 2 · 1 · **2** · 1	No inventory owner or review cadence; new services not detected	Moderate
Data Protection & Leakage (22)	6 · 6 · 2 · 1 · **6** · 1	Text DLP strong; file upload, personal accounts and uncategorized services uncontrolled	High
Identity & Access (16)	5 · 5 · 3 · 1 · **1** · 1	SSO on both enterprise tenants; SCIM missing on ChatGPT Enterprise	High
Endpoint & Browser (12)	2 · 4 · 2 · 1 · **2** · 1	Software allowlist exists; no browser-extension management; desktop assistant installs	High
Enterprise AI Configuration (16)	4 · 5 · 2 · 1 · **3** · 1	Copilot oversharing exposure; Copilot Studio open; ChatGPT connectors broadly enabled	High
Monitoring, Detection & Response (14)	4 · 5 · 2 · 1 · **1** · 1	Logging exists at the proxy; no forwarding, retention or alerting for AI events	High
Vendor & Third-Party AI (14)	4 · 4 · 2 · 1 · **1** · 2	Enterprise vendor terms reviewed; embedded AI features in SaaS not reviewed	Moderate

Copyright © 2026, ioSENTRIX LLC.

7 Technical Validation

The validation stories below are the evidence behind the assurance ratings and the findings. Every test used synthetic records generated by ioSENTRIX from Larkspur's data schemas and carrying ioSENTRIX-specific markers recognizable to DLP; no real customer, employee or client data was used. Tests were run from a standard managed laptop with a standard user account, from the corporate network and from a home network via the SSE client. Thirty-six controls were validated; the stories below cover the ones that failed or partially failed.

7.1 Data-leakage validation matrix

Objective: verify that synthetic customer records cannot reach an AI service other than the sanctioned enterprise tenants, by each path a user could realistically take, and that attempts are logged and alerted.

Path → Service	ChatGPT Enterprise (SSO)	ChatGPT (personal account)	Gemini / Claude (consumer)	Perplexity
Paste text with markers	Allowed (sanctioned); logged	**Allowed; logged as chatgpt.com, not distinguished from enterprise**	Blocked by SSE DLP; logged; **no alert**	**Allowed; domain not in AI category; not inspected**
Upload file with markers	Allowed (sanctioned); logged	**Allowed; file not inspected**	**Allowed; file not inspected**	**Allowed; file not inspected**
Paste from home network (SSE client)	Allowed; logged	**Allowed; not distinguished**	Blocked; logged; no alert	**Allowed**
Direct API / CLI	N/A	Blocked (egress policy)	Blocked (egress policy)	Blocked (egress policy)

Result: Fail on file upload, personal accounts and uncategorized services

Text DLP at the proxy works where the destination is categorized as AI. Everything else passed: file uploads on every path, personal ChatGPT accounts on every path, and any AI service the category list does not yet know about. Zero alerts were raised by any blocked or allowed event. Remediation direction: enable file inspection in the SSE DLP policy; add instance awareness so personal and enterprise ChatGPT sessions are distinguished and personal sessions blocked; move from a category block-list to an allow-list for AI destinations.

Personal versus enterprise accounts

Objective: verify that a personal consumer account on a sanctioned vendor's service is blocked or distinguished from the enterprise tenant on a managed device.

Test step	Safe test design	Result
VAL-DLP-03a Login path	Sign in to chatgpt.com with an ioSENTRIX-owned personal test account from a managed laptop	**Fail — login allowed; same proxy rule as enterprise tenant**

Copyright © 2026, ioSENTRIX LLC.

Test step	Safe test design	Result
VAL-DLP-03b Data path	Paste and upload synthetic records in the personal session	**Fail — allowed on both paths; logs show only the domain**
VAL-DLP-03c Tenant header	Check for tenant or workspace restriction headers injected by the SSE proxy	**Fail — none configured**
VAL-IAM-02 Enterprise SSO	Confirm enterprise workspace enforces SSO and rejects password login	Pass — SSO enforced

7.2 Browser extensions and desktop assistants

Objective: verify that AI functionality installed outside the sanctioned tools cannot exfiltrate page or document content without inspection.

Test step	Safe test design	Result
VAL-END-01 Extension install	Install the AI summarizer extension observed on 140 endpoints, as a standard user	**Fail — installation allowed; no extension policy in Chrome or Edge management**
VAL-END-02 Extension data path	Open an intranet page containing synthetic client records; invoke "summarize"	**Fail — page content sent to the vendor API domain; not categorized as AI; not inspected; not logged as AI activity**
VAL-END-03 Desktop assistant	Install the desktop AI assistant observed on 37 endpoints	**Fail — allowed; not on the software allowlist, but allowlisting is in audit mode only**
VAL-END-04 Desktop assistant data path	Submit a synthetic document through the desktop assistant	Partial — paste blocked by DLP; file attachment allowed

7.3 Detection and alerting chain

Objective: verify that AI-related security events (blocked data-leakage attempts, new AI service first use, policy violations) are logged, retained, forwarded and alerted.

Test step	Safe test design	Result
VAL-MON-01 Logging	Locate the blocked VAL-DLP events in the SSE console	Pass — present with user, destination and policy name
VAL-MON-02 Retention	Confirm retention period for proxy and DLP events	**Partial — 30 days at the proxy; not forwarded to the SIEM**
VAL-MON-03 Alerting	Trigger six blocked events across two days; check for any alert to security staff	**Fail — no alert rule exists; zero alerts**

Copyright © 2026, ioSENTRIX LLC.

Test step	Safe test design	Result
VAL-MON-04 New-service detection	Introduce a benign, previously unseen AI service from a test endpoint; check whether it is detected within 7 days	**Fail — not detected; no new-service report or review**

7.4 Microsoft 365 Copilot configuration

Objective: verify that Copilot cannot surface content to a user who should not see it, and that its agent and web-grounding features are governed.

Test step	Safe test design	Result
VAL-COP-01 Oversharing	Seed synthetic HR and client files into a SharePoint site shared with "Everyone except external users"; ask Copilot about them as a test user in another department	**Fail — content returned; 61 such sites identified in tenant review**
VAL-COP-02 Sensitivity labels	Check label coverage and whether a labelled synthetic file is excluded or watermarked in Copilot responses	**Partial — labels defined; applied to 12% of sites; labelled file correctly handled**
VAL-COP-03 Restricted SharePoint Search	Check whether Restricted SharePoint Search is enabled as a temporary oversharing control	**Fail — not enabled**
VAL-COP-04 Copilot Studio agents	Attempt to create and share an agent as a standard user	**Fail — any user can create and share agents; no approval**
VAL-COP-05 Audit	Confirm Copilot interactions are captured in the unified audit log	Pass — captured; 180-day retention

7.5 ChatGPT Enterprise workspace configuration

Objective: verify that the sanctioned ChatGPT workspace is configured so that only current staff can use it, that usage is reviewable, and that data-sharing features are governed.

Test step	Safe test design	Result
VAL-TEN-01 SSO	Confirm SSO is enforced for all members	Pass
VAL-TEN-02 Lifecycle	Reconcile workspace members against the identity provider	**Fail — SCIM not enabled; 23 stale accounts, including 6 leavers with active access**
VAL-TEN-03 Data use	Confirm workspace data is excluded from model training by default and retention settings are documented	Pass — excluded by default; retention at vendor default, documented
VAL-TEN-04 Compliance logging	Confirm compliance logging or export is enabled and reviewed	**Fail — not enabled; policy §7 cannot be met for this tool**

Copyright © 2026, ioSENTRIX LLC.

Test step	Safe test design	Result
VAL-TEN-05 Connectors and sharing	Check which data connectors and custom-GPT sharing settings are enabled, and for whom	**Partial — custom GPTs workspace-only (good); cloud-storage connectors enabled for all members without review**

8 Technical Risks

8.1 Summary of findings

Below is the summary of the findings identified during the assessment. Seventeen failing controls were consolidated into ten findings; complete details are provided in the following section.

No	Security Risk	Severity	Status
9.1.1	AI-DLP-01: Personal AI accounts indistinguishable from enterprise	Critical	Open
9.2.1	AI-DLP-02: File uploads to AI services are not inspected	High	Open
9.2.2	AI-END-01: AI browser extension bypasses all data controls	High	Open
9.2.3	AI-COP-01: Copilot surfaces over-shared SharePoint content	High	Open
9.3.1	AI-DLP-03: Unapproved AI services outside the blocked category	Medium	Open
9.3.2	AI-MON-01: AI activity not retained, forwarded or alerted	Medium	Open
9.3.3	AI-GOV-01: AI policy lacks enforcement map and attestation	Medium	Open
9.3.4	AI-TEN-01: ChatGPT Enterprise lifecycle and compliance logging gaps	Medium	Open
9.4.1	AI-END-02: Desktop AI assistant outside software allowlist	Low	Open
9.5.1	AI-VEN-01: Embedded SaaS AI features not reviewed	Informational	Open

Copyright © 2026, ioSENTRIX LLC.

9 Findings

9.1 Critical Severity Risks

9.1.1 AI-DLP-01 · Personal AI accounts indistinguishable from enterprise

Risk Rating	Impact	Likelihood
Critical	Critical	High

Description:

Larkspur sanctions ChatGPT Enterprise, so the SSE proxy allows chatgpt.com. The proxy has no instance awareness or tenant restriction configured, so a personal consumer ChatGPT account used from a managed laptop receives exactly the same treatment as the enterprise workspace: the login is allowed, pasted and uploaded content is allowed, and the proxy log records only the domain. DLP text inspection does not apply because the destination is on the allow-list.

Discovery shows this is not hypothetical: identity-provider and proxy evidence indicate consumer ChatGPT logins from 97 managed endpoints in the 30-day window, 41 of them by staff who also hold an enterprise seat. Data submitted to a personal account is outside Larkspur's enterprise terms (no training exclusion, no retention control, no compliance logging) and outside any ability to review it.

Affected systems:

1. SSE proxy: web access and DLP policies for the AI category and for chatgpt.com
2. Managed Windows and macOS laptops (all users)

Validation performed:

1. VAL-DLP-03a: sign in to chatgpt.com with an ioSENTRIX-owned personal test account from a managed laptop — allowed (fail).
2. VAL-DLP-03b: paste and upload synthetic customer records in the personal session — allowed on both paths; proxy log shows domain only (fail).
3. VAL-DLP-03c: inspect proxy configuration for tenant or workspace restriction headers — none configured (fail).

Evidence:

- Proxy log entries for the test session showing chatgpt.com allowed under the enterprise rule, with no account or workspace attribute.
- SSE policy export showing no instance-awareness or tenant-restriction configuration for the sanctioned AI vendors.

Copyright © 2026, ioSENTRIX LLC.

- Discovery register: 97 endpoints with consumer ChatGPT sign-ins in the window (identity and proxy evidence).

Business impact:

Customer financial records, client documents and employee data can be submitted to a consumer AI account that Larkspur does not control, cannot log and cannot review, under terms that permit retention and may permit training. For a regulated brokerage this is a direct confidentiality and regulatory-reporting exposure, and it defeats the purpose of having bought the enterprise tenant.

Policy clauses and controls:

Policy §3, §4 · Controls DLP-04 (instance and tenant restriction), DLP-05 (account-aware logging), IAM-06 (sanctioned-tenant enforcement)

Remediation:

- Configure instance awareness / tenant restriction at the SSE proxy for ChatGPT and for the Microsoft tenant, so that only the Larkspur enterprise workspace is reachable and personal logins are blocked or redirected to the enterprise SSO.
- Add the account or workspace attribute to AI activity logging so personal and enterprise sessions are distinguishable in review.
- Communicate the change to staff with the enterprise tenant as the sanctioned path; retest VAL-DLP-03 after configuration.

9.2 High Severity Risks

9.2.1 AI-DLP-02 · File uploads to AI services are not inspected

Risk Rating	Impact	Likelihood
High	High	High

Description:

The SSE DLP policy inspects text submitted in web forms and correctly blocks pasted content carrying customer-data markers at categorized AI destinations. File uploads (multipart form data and the upload APIs used by AI chat services) are not inspected on any path: synthetic customer files reached Gemini, Claude and Perplexity without a block or a log entry beyond the domain. Endpoint DLP is installed but has no rule for uploads to AI destinations.

Affected systems:

1. SSE proxy: DLP policy for the AI category (file inspection disabled)
2. Endpoint DLP policy (no AI-destination upload rule)

Copyright © 2026, ioSENTRIX LLC.

Validation performed:

1. VAL-DLP-02: upload synthetic customer and client-document files to Gemini, Claude and Perplexity from a managed laptop on the corporate network — allowed on all three (fail).
2. VAL-DLP-02b: repeat from a home network via the SSE client — allowed (fail).
3. VAL-DLP-01: paste the same records as text — blocked at Gemini and Claude (pass), allowed at Perplexity (see AI-DLP-03).

Evidence:

- Proxy log entries showing the upload requests allowed, with no DLP policy match.
- SSE DLP policy export showing file inspection disabled for the AI category.
- Endpoint DLP policy export with no rule covering uploads to AI destinations.

Business impact:

The most common way a staff member shares a client document with an AI tool is to upload it, and that path is uncontrolled. Discovery recorded 118 staff submitting content to unapproved services in 30 days; uploads from Insurance Brokerage included policy documents.

Policy clauses and controls:

Policy §3 · Controls DLP-02 (upload inspection), DLP-03 (endpoint upload control)

Remediation:

- Enable file inspection in the SSE DLP policy for the AI category and for all uncategorized destinations.
- Add an endpoint DLP rule blocking file uploads to AI destinations outside the sanctioned tenants.
- Retest VAL-DLP-02 on both networks after configuration.

9.2.2 AI-END-01 · AI browser extension bypasses all data controls

Risk Rating	Impact	Likelihood
High	High	High

Description:

An AI summarizer browser extension is installed on 140 managed endpoints. It reads the content of the current page and sends it to the vendor's API domain, which is not categorized as AI by the SSE proxy and is therefore neither blocked nor inspected; the traffic is logged only as a generic web request. Neither Chrome nor Edge management has an extension allowlist or blocklist, so any user can install

Copyright © 2026, ioSENTRIX LLC.

any extension. Two further AI extensions (a writing assistant on 88 endpoints and a translation tool on 22) have the same characteristics.

Affected systems:

1. Managed browsers (Chrome and Edge): extension management policy absent
2. SSE proxy: AI category coverage for extension API domains

Validation performed:

1. VAL-END-01: install the summarizer extension as a standard user — allowed (fail).
2. VAL-END-02: open an intranet page containing synthetic client records and invoke summarize — page content sent to the vendor API domain; not inspected; logged as generic web traffic (fail).

Evidence:

- Browser-extension inventory from endpoint management showing the three AI extensions and install counts.
- Proxy log entries for the extension API domain showing allowed, uncategorized requests during the test.
- Chrome and Edge management policy exports showing no extension controls.

Business impact:

Any page a staff member views, including CRM records, portfolio screens and HR pages, can be sent to a third party with a single click, with no block, inspection or AI-specific log. Discovery shows the extension is most prevalent in Advisory and Client Services, the departments handling the most sensitive client data.

Policy clauses and controls:

Policy §3, §5 · Controls END-03 (browser extension management), DLP-06 (destination coverage)

Remediation:

- Deploy an extension allowlist through Chrome and Edge management; remove the three AI extensions and require IT approval for any extension with page-content access.
- Add the identified extension API domains to the AI category at the proxy, and move AI destinations to an allow-list model so new services are blocked by default.
- Retest VAL-END-01 and VAL-END-02.

Confidential
Copyright © 2026, ioSENTRIX LLC.

9.2.3 AI-COP-01 · Copilot surfaces over-shared SharePoint content

Risk Rating	Impact	Likelihood
High	High	High

Description:

Microsoft 365 Copilot respects existing permissions, which means it also inherits existing oversharing. Tenant review found 61 SharePoint sites shared with "Everyone except external users", including HR, client-onboarding and brokerage sites. A test user in another department asked Copilot about synthetic HR and client files seeded into one such site and received their contents. Sensitivity labels are defined but applied to 12% of sites, and Restricted SharePoint Search, Microsoft's temporary control for exactly this situation, is not enabled. Separately, any user can create and share Copilot Studio agents without approval.

Affected systems:

1. Microsoft 365 tenant: SharePoint Online site permissions (61 broadly shared sites)
2. Microsoft Purview: sensitivity label coverage
3. Copilot Studio: agent creation and sharing settings

Validation performed:

1. VAL-COP-01: seed synthetic HR and client files into a broadly shared site; query Copilot as a test user from another department — content returned (fail).
2. VAL-COP-02: query a labelled synthetic file — correctly handled (pass); label coverage 12% (partial).
3. VAL-COP-03: check Restricted SharePoint Search — not enabled (fail).
4. VAL-COP-04: create and share an agent as a standard user — allowed (fail).

Evidence:

- Copilot responses to the test queries containing the synthetic file contents.
- SharePoint admin report listing the 61 broadly shared sites.
- Purview label coverage report; Copilot Studio tenant settings export.

Business impact:

Copilot turns latent oversharing into a one-question search: any staff member can ask for HR, compensation or client information that a permissions mistake made visible years ago. The exposure exists today; Copilot makes it discoverable.

Copyright © 2026, ioSENTRIX LLC.

Policy clauses and controls:

Policy §3 · Controls COP-01 (oversharing remediation), COP-02 (label coverage), COP-03 (Restricted SharePoint Search), COP-05 (agent governance)

Remediation:

- Enable Restricted SharePoint Search immediately, limiting Copilot to an allow-list of reviewed sites until oversharing is remediated.
- Remediate the 61 broadly shared sites; apply sensitivity labels to HR, client and brokerage sites first; use the SharePoint oversharing reports to track progress.
- Restrict Copilot Studio agent creation and sharing to an approved group; retest VAL-COP-01 after remediation.

9.3 Medium Severity Risks

9.3.1 AI-DLP-03 · Unapproved AI services outside the blocked category

Risk Rating	Impact	Likelihood
Medium	High	Medium

Description:

The SSE "AI and ML" category block is the only control implementing policy §2, and it depends on the vendor's category list. Perplexity, the consumer Microsoft Copilot and several writing and meeting tools are not in the category and are reachable with no inspection. Sixteen unapproved standalone services were observed in 30 days; the category blocked two of them.

Affected systems:

1. SSE proxy: AI category block-list

Validation performed:

1. VAL-DLP-01: paste synthetic records at Perplexity — allowed; destination uncategorized (fail).
2. VAL-MON-04: introduce a benign unseen AI service — reachable, not detected (fail).

Evidence:

- Proxy category lookup results for the 16 unapproved standalone services (14 uncategorized or categorized outside AI).
- Discovery register with user counts per service.

Copyright © 2026, ioSENTRIX LLC.

Business impact:

A block-list that depends on a third party's categorization will always lag the market; staff found 14 services it does not cover. Any of them can receive pasted or uploaded data without inspection.

Policy clauses and controls:

Policy §2 · Controls DLP-06 (destination coverage), DIS-03 (new-service detection)

Remediation:

- Move AI destinations to an allow-list: block uncategorized and newly seen AI services by default and approve by exception.
- Add the 16 observed services to the AI category or block them explicitly; review new-service reports monthly (see AI-MON-01).

9.3.2 AI-MON-01 · AI activity not retained, forwarded or alerted

Risk Rating	Impact	Likelihood
Medium	Medium	High

Description:

AI-related events are logged at the SSE proxy with user, destination and policy name, which is a sound foundation. They are retained for 30 days, are not forwarded to the SIEM, and no alert rule exists for blocked data-leakage attempts or for first use of a new AI service. Six blocked events triggered deliberately over two days produced no alert to anyone.

Affected systems:

1. SSE proxy logging and retention
2. SIEM: log sources and alert rules

Validation performed:

1. VAL-MON-01: locate blocked test events in the SSE console — present (pass).
2. VAL-MON-02: confirm retention and forwarding — 30 days, not forwarded (partial).
3. VAL-MON-03: trigger six blocked events; check for alerts — none (fail).
4. VAL-MON-04: introduce an unseen AI service — not detected in 7 days (fail).

Evidence:

- SSE retention settings; SIEM log-source inventory with no SSE/DLP source.
- Absence of alerts for the six test events, confirmed with the IT security lead.

Copyright © 2026, ioSENTRIX LLC.

Business impact:

Policy §7 promises review of AI activity that cannot happen beyond 30 days, and a staff member repeatedly attempting to exfiltrate customer data to an AI service would never be noticed.

Policy clauses and controls:

Policy §7 · Controls MON-02 (retention and forwarding), MON-04 (alerting), DIS-03 (new-service detection)

Remediation:

- Forward SSE and DLP events to the SIEM and retain for 12 months.
- Create alert rules for blocked data-leakage attempts (threshold per user per day) and for first use of any AI destination not on the allow-list.
- Produce a monthly new-AI-service report for the policy owner.

9.3.3 AI-GOV-01 · AI policy lacks enforcement map and attestation

Risk Rating	Impact	Likelihood
Medium	Medium	High

Description:

The AI Acceptable Use Policy v1.2 is owned, approved and published. It does not identify which control enforces each clause, four of its nine clauses rely on a control that failed under test and one has no technical control at all, staff have not attested to it, and the associated awareness course assigned in June 2026 shows 38% completion. The policy is therefore a statement of intent that most staff have not read and that systems do not enforce.

Affected systems:

1. AI Acceptable Use Policy v1.2 and its governance process
2. Learning management system: AI awareness course

Validation performed:

1. VAL-GOV-01: map each policy clause to an enforcing control and test it — 2 enforced, 2 partial, 4 failing, 1 none (see Policy versus Reality).
2. VAL-GOV-02: confirm attestation and training completion — no attestation; 38% completion (fail).

Copyright © 2026, ioSENTRIX LLC.

Evidence:

- Policy-to-control map (section 5).
- LMS completion report for the AI awareness course.

Business impact:

Without an enforcement map, policy reviews will continue to conclude that controls are "in place" because the policy says so. Without attestation and training, the clauses that depend on staff behaviour rest on people who have mostly not completed the course.

Policy clauses and controls:

Policy §6, §8 · Controls GOV-03 (enforcement mapping), GOV-05 (attestation), GOV-06 (training)

Remediation:

- Add an enforcement-map appendix to the policy naming the control, owner and evidence for each clause; review it with each policy revision.
- Require annual attestation; drive training completion to a target with manager follow-up; for §6, add a sampling review of AI-assisted client deliverables.

9.3.4 AI-TEN-01 · ChatGPT Enterprise lifecycle and compliance logging gaps

Risk Rating	Impact	Likelihood
Medium	Medium	Medium

Description:

The sanctioned ChatGPT Enterprise workspace enforces SSO and excludes workspace data from model training, both confirmed. It is not connected to the identity provider for provisioning, so membership drifts: reconciliation found 23 stale accounts, including 6 leavers whose enterprise seats remain active. Compliance logging is not enabled, so policy §7 cannot be met for this tool, and cloud-storage connectors are enabled for all members without review of which data sources they expose.

Affected systems:

1. ChatGPT Enterprise workspace: membership, compliance settings, connectors

Validation performed:

1. VAL-TEN-02: reconcile workspace members against the identity provider — 23 stale, 6 leavers (fail).
2. VAL-TEN-04: confirm compliance logging — not enabled (fail).

Copyright © 2026, ioSENTRIX LLC.

3. VAL-TEN-05: review connector and sharing settings — connectors enabled for all; custom GPTs workspace-only (partial).

Evidence:

- Workspace member export reconciled against HR leaver list.
- Workspace admin settings export.

Business impact:

Former staff retain access to a tool that holds client correspondence drafts; usage cannot be reviewed; and connectors can expose cloud-stored documents to the workspace without a decision having been made.

Policy clauses and controls:

Policy §7 · Controls IAM-04 (lifecycle), TEN-02 (compliance logging), TEN-04 (connector governance)

Remediation:

- Enable SCIM provisioning from the identity provider; remove the 23 stale accounts now.
- Enable compliance logging and define who reviews it and how often.
- Restrict connectors to an approved group and reviewed data sources.

9.4 Low Severity Risks

9.4.1 AI-END-02 · Desktop AI assistant outside software allowlist

Risk Rating	Impact	Likelihood
Low	Low	High

Description:

A desktop AI assistant is installed on 37 managed endpoints. Software allowlisting exists but runs in audit mode, so the installation is recorded and allowed. Pasted content from the assistant is blocked by DLP; file attachments are not (see AI-DLP-02).

Affected systems:

1. Endpoint management: software allowlisting (audit mode)

Copyright © 2026, ioSENTRIX LLC.

Validation performed:

1. VAL-END-03: install the assistant as a standard user — allowed; allowlisting in audit mode (fail).
2. VAL-END-04: submit synthetic content — paste blocked; attachment allowed (partial).

Evidence:

- Endpoint software inventory showing 37 installs; allowlisting policy export in audit mode.

Business impact:

Limited on its own because the data path is largely covered by the DLP findings, but it demonstrates that the software allowlist does not currently prevent anything.

Policy clauses and controls:

Policy §2 · Controls END-01 (software allowlisting)

Remediation:

- Move software allowlisting to enforce mode with a short exception process; remove the assistant where no business need exists.

9.5 Informational Severity Risks

9.5.1 AI-VEN-01 · Embedded SaaS AI features not reviewed

Risk Rating	Impact	Likelihood
Informational	Informational	Informational

Description:

Eight AI features inside platforms Larkspur already sanctions (video-conferencing companion, CRM assistant, document-signing summaries, design and note-taking tools) are enabled, several by default, without a review of their data-handling terms. Because they ride on sanctioned platforms they are invisible to network controls and to the AI category. Enterprise vendor terms for the three sanctioned AI tools, by contrast, have been reviewed.

Affected systems:

1. Vendor management: AI-feature terms for sanctioned SaaS platforms

Copyright © 2026, ioSENTRIX LLC.

Validation performed:

1. VAL-VEN-01: check whether each embedded AI feature has a documented terms review and an enable/disable decision — none documented.

Evidence:

- SaaS admin console settings for the eight features; vendor-management records.

Business impact:

Several of these features process meeting audio and client documents. The exposure is probably acceptable under existing enterprise agreements, but nobody has checked.

Policy clauses and controls:

Controls VEN-03 (embedded AI feature review), DIS-02 (inventory coverage)

Remediation:

- Add embedded AI features to the AI inventory; review each vendor's AI data-handling terms and record an enable, restrict or disable decision.

10 Framework Traceability

The methodology uses recognized AI governance and security references as traceability anchors, then extends them with evidence requirements and technical validation. Versions are current as of the engagement date.

Reference	How it is used
NIST AI RMF 1.0 (Jan 2023) and NIST AI 600-1 (Jul 2024)	Govern / Map / Measure / Manage backbone; generative-AI risk profile for an AI-consuming organization.
ISO/IEC 42001:2023	AI management-system concepts for policy, roles and continual improvement. ioSENTRIX is not a certification body.
CSA AI Controls Matrix v1.1 (Jul 2026)	247-objective, 18-domain control crosswalk; the consumer-side domains (governance, data, identity, logging, third party) map to this engagement.
OWASP Top 10 for LLM Applications 2026	Sensitive information disclosure and supply-chain categories, applied to AI consumption rather than AI development.
EU AI Act, as amended by Regulation (EU) 2026/1744	Deployer transparency obligations apply from 2 August 2026; relevant to any EU client-facing use of AI-generated content.
GLBA Safeguards Rule; state insurance data-security regulations	Customer financial-data protection obligations that frame the business impact of the data-leakage findings.

Copyright © 2026, ioSENTRIX LLC.

Traceability, not certification

Framework references indicate alignment with recognized guidance. They do not constitute certification against, or compliance with, any framework or regulation, and alignment with one framework does not imply compliance with another.

11 Residual Risk and Remediation Roadmap

Residual risk by data path

Data path	Inherent risk	Control effectiveness	Residual risk
Personal AI accounts on managed devices	Critical	Failing — no instance or tenant restriction	Critical
File upload to unapproved AI services	High	Failing — no file inspection	High
Browser extensions and uncategorized services	High	Failing — no extension policy; block-list model	High
Copilot over sanctioned M365 data	High	Partial — permissions respected; oversharing unremediated	High
Paste to categorized consumer AI	High	Implemented — text DLP blocks; no alert	Medium
Sanctioned enterprise tenants	Medium	Implemented — SSO, training exclusion; lifecycle gaps	Low

30 / 60 / 90-day roadmap

Window	Priority actions	Outcome
0–30 days	Configure SSE instance and tenant restrictions for ChatGPT and Microsoft (AI-DLP-01); enable file inspection and endpoint upload rules (AI-DLP-02); enable Restricted SharePoint Search and restrict Copilot Studio (AI-COP-01); deploy browser-extension allowlist (AI-END-01).	The Critical and three High paths are closed using tools Larkspur already owns.
31–60 days	Move AI destinations to an allow-list (AI-DLP-03); forward AI events to the SIEM with alerting (AI-MON-01); enable SCIM and compliance logging on ChatGPT Enterprise (AI-TEN-01); enforce-mode software allowlisting (AI-END-02).	New AI services are blocked by default; attempts are visible and alerted; sanctioned tenants are governed.

Copyright © 2026, ioSENTRIX LLC.

Window	Priority actions	Outcome
61–90 days	Remediate the 61 over-shared sites and extend label coverage (AI-COP-01); publish the policy enforcement map, attestation and training push (AI-GOV-01); review embedded SaaS AI features (AI-VEN-01); ioSENTRIX retest of all Failing controls.	Policy clauses map to proven controls; Failing controls move to Assured with evidence.
Ongoing	Continuous AI Assurance: quarterly re-discovery and re-validation, new-service and vendor-feature review, policy updates, trended board dashboard.	Sustained, board-ready assurance as AI tools and policy change.

ioSENTRIX can support remediation independently

Using Larkspur's existing stack (SSE proxy and DLP policy, Purview labels and Restricted SharePoint Search, browser and endpoint management, SIEM content) under a separately scoped statement of work. Continuous AI Assurance keeps the inventory current and re-tests the controls that matter between annual assessments.

Confidential
Copyright © 2026, ioSENTRIX LLC.

About ioSENTRIX

ioSENTRIX LLC is a security consulting firm accredited by CREST, an internationally recognized hallmark for professional cybersecurity services, certified to ISO/IEC 27001 and attested to SOC 2 Type 2. As a CREST-accredited penetration testing company, we adhere to the highest standards of service, ensuring that our clients receive the most comprehensive and reliable security testing available.

We provide a wide range of security consulting services to clients worldwide. Our clientele includes Fortune 500 companies, large enterprises, small start-ups, financial institutions, and several high-tech companies. At ioSENTRIX, we are committed to delivering innovative cybersecurity solutions.

As an innovative consulting company, we offer a full range of cyber security services tailored to businesses of all sizes and budget requirements. Our focus is on helping clients identify, mitigate, and prevent vulnerabilities in their software, infrastructure, and cloud environments.

Our services include design review, threat modeling, penetration testing, code review, red teaming, vCISO advisory and AI Assurance & Control Validation. Our offensive-security heritage is what lets us do more than review AI governance on paper: we test whether the controls actually work.

Learn more about how our CREST accreditation and our services can benefit your business at **https://www.iosentrix.com.**

ioSENTRIX LLC.

13800 Coppermine Rd, Suite 190
Herndon Virginia 20171 (USA)

Sales: 1 (888) 958-0554
Email: sales@iosentrix.com

Confidential
Copyright © 2026, ioSENTRIX LLC.