SAMPLE • ANONYMIZED

AI Readiness & Security Assessment

Executive findings, prioritized risk, opportunity map, and 90-day roadmap

ASSESS → SECURE → IMPLEMENT → TRAIN → MANAGE

Aegis IT Consulting

Prepared for: Confidential Sample Organization
Assessment status: Sample demonstration
Report version: 1.0

This document is a fictionalized and anonymized demonstration based on a completed small-business assessment. Names, identities, and identifying details have been removed. The findings illustrate the form and depth of a professional Aegis deliverable.

LEADERSHIP VIEW

Executive Summary

Overall classification: DEVELOPING

The organization is already using generative AI for practical business work and demonstrates several positive controls: managed endpoints, human review of AI-generated work, and deliberate avoidance of autonomous integrations.

However, foundational identity, data-management, and governance controls require remediation before AI is broadly connected to organizational repositories or production systems.

Current recommendation: Controlled AI use may continue. Advanced integration: Remediate before broad deployment.

Executive Decision

PROCEED WITH GUARDRAILS	Permit approved, human-supervised AI use. Pause broad data connections, autonomous agents, and write-capable integrations until priority controls are completed.

What is working

- AI access is limited to managed and secured endpoints.
- AI-generated work receives human review before use or delivery.
- No autonomous agents or direct AI-to-production integrations are currently deployed.
- Leadership has identified concrete use cases and is approaching adoption deliberately.

72 CONTROLS • 9 DOMAINS

Readiness by Domain

Domain scores identify where adoption can move forward and where remediation should come first. High-risk findings override the average score so urgent exposure is never hidden inside an otherwise acceptable rating.

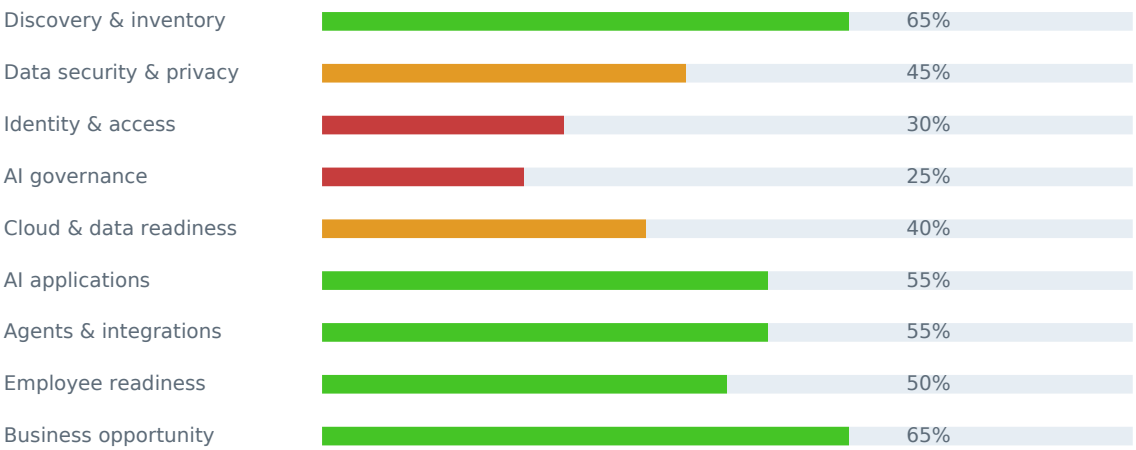

0-34 High risk	35-49 Needs improvement	50-74 Developing strength	75-100 Mature

Interpretation

Strongest current areas	Most urgent areas
AI discovery, practical business opportunity, and controlled initial adoption.	Identity and access security, governance, organization-controlled data repositories, and formal policy.

RISK BEFORE ROLLOUT

Prioritized Findings

Priority	Finding	Business impact and recommended action
HIGH	Multi-factor authentication gaps	Business AI and cloud accounts used for organizational work are not consistently protected with MFA. Action: Enable MFA immediately and confirm unique, password-manager-generated credentials.
HIGH	Privileged account separation	A daily-use cloud account also performs administrative functions, increasing the impact of credential compromise. Action: Create a dedicated privileged identity and apply least-privilege administrative roles.
HIGH	Individually controlled business repository	Authoritative company files are synchronized to an individually controlled cloud-storage account. Action: Move records into an organization-controlled repository with managed ownership and recovery.
MODERATE	No AI acceptable-use policy	Employees lack documented rules for approved tools, sensitive data, human review, and prohibited activity. Action: Adopt an AI Acceptable Use & Security Policy and communicate it through role-based training.
MODERATE	No data classification standard	The organization has no consistent Public, Internal, Confidential, and Restricted handling model for AI use. Action: Define classifications and state which categories may enter each approved AI platform.
MODERATE	Cloud permissions not verified	External sharing, administrative roles, and repository permissions have not been reviewed for AI exposure. Action: Complete a permissions and external-sharing audit before connecting AI to cloud data.

Critical-control override

The readiness score does not replace security findings. A High or Critical issue may trigger an implementation hold for sensitive-data connections or autonomous capabilities until the underlying control is corrected.

VALUE WITH CONTROL

AI Opportunity Map

Potential use cases were ranked by business value, feasibility, risk, and implementation effort. Initial pilots should assist people rather than act autonomously.

#	Opportunity	Value	Effort	Starting model	Outcome
1	Sales & prospect research	High	Med.	Assist only	Research organizations, identify decision-makers, enrich account information, and draft personalized outreach for human approval.
2	Service-management assistant	High	Med.	Read / analyze / draft	Summarize tickets, identify recurring issues, prepare client research, and draft technician notes or reports.
3	Proposal intelligence	High	Med.	Controlled knowledge	Use approved service descriptions, templates, pricing rules, and prior materials in an organization-controlled knowledge system.
4	Cybersecurity workflows	Med.-High	Med.	Human verified	Assist with security research, policy drafting, alert context, assessment preparation, and client-facing recommendations.

Required pilot guardrails

• Approved accounts and organization-controlled ownership	• MFA, least privilege, and documented access	• No restricted data unless explicitly approved
• Human review before communication or consequential action	• Logging, revocation, and defined incident reporting	• Measured outcomes: time saved, quality, risk, and adoption

FROM FINDINGS TO ACTION

30 / 60 / 90-Day Roadmap

FIRST 30 DAYS

Secure the foundation

- Enable MFA on business AI, cloud, and administrative accounts.
- Separate privileged administration from daily-use identities.
- Verify password-manager usage and administrative roles.
- Inventory AI tools, cloud repositories, sharing, and current data ownership.

DAYS 31-60

Establish governance

- Adopt AI acceptable-use and security policy.
- Define approved tools and a vendor-review procedure.
- Create data-classification and AI-retention standards.
- Extend incident response and offboarding to AI accounts, OAuth grants, tokens, agents, and workspaces.

DAYS 61-90

Organize and pilot

- Select the authoritative organization-controlled data repository.
- Correct external-sharing and permission issues.
- Launch one human-supervised, low-autonomy pilot.
- Measure value, security exceptions, quality, time saved, and employee adoption.

Success at day 90

SECURE	GOVERNED	MEASURABLE
Priority identity and ownership gaps closed.	Policy, approvals, and data rules established.	One controlled pilot producing documented value.

EVIDENCE, NOT GUESSWORK

Assessment Methodology

The full Aegis assessment evaluates 72 controls across nine risk and readiness domains. Questionnaire responses are supplemented by technical verification where authorized and applicable.

- AI application and browser-extension inventory
- MFA status, administrative roles, SSO, and account ownership
- Cloud permissions, external sharing, and repository ownership
- OAuth grants, connected applications, APIs, agents, and MCP servers
- AI workspace administration, privacy settings, and retention
- Endpoint security, backup/recovery, and incident-response readiness

What the client receives

Executive score Clear readiness and risk posture.	Prioritized findings Evidence, impact, priority, and action.	Opportunity map Use cases ranked by value, feasibility, and risk.
30/60/90 roadmap Owners, milestones, and practical next steps.	Executive review Plain-English walkthrough and decisions.	Implementation path Policy, remediation, rollout, training, and management.

MOVE FASTER WITH AI. NOT RECKLESSLY.

Start with an Aegis AI Readiness & Security Assessment.

412-695-9900
aegisitconsulting.net

Important: This sample is not a guarantee of security, legal compliance, regulatory compliance, or suitability of any AI system. A complete assessment depends on the information and access authorized by the client at the time of review.