FUSION TECHNOLOGY GROUP
WHERE AI MEETS IT.

FUSION WORKING PAPER

Using AI safely in your business

A practical framework for South African companies that want the benefits of AI without losing control of their information.

Fusion Technology Group (Pty) Ltd

November 2026

In short

AI has arrived in most businesses before anyone decided it should. Staff are drafting emails, summarising documents and analysing spreadsheets with whatever tool is closest, often on personal accounts, and leadership usually can't see what information has gone where.

Banning it rarely works. People simply carry on out of sight, and the business loses the real productivity gains on offer. The better answer is a clear, proportionate set of rules, and a controlled way of connecting AI to company information.

This paper sets out the framework we use. It is written for owners and executives, and it is deliberately scaled for small and medium businesses rather than large corporates.

> **The rule everything rests on**
>
> AI only sees what the person asking is already allowed to see. Enforce this where the data lives, through sign-in, permissions and controlled connections. Never rely on telling the AI to behave.

Eight principles

PRINCIPLE	WHAT IT MEANS
Company AI only	Company information goes only into approved tools, signed in with a company account. No personal or free accounts.
Same access as today	Nobody gets more data through AI than they already have.
Read-only by default	AI connections may look information up. They may not create, change, delete, post or send without a specific, logged exception.
People stay accountable	AI drafts and suggests. A named person checks the output and owns the result.
Risk-based	The more sensitive the data and the bigger the consequence of a mistake, the more approval and oversight it needs.
Least data necessary	Share only what the task needs, and personal information only where there is a business reason.
Visible and auditable	The business can see who used what, when, and what data was returned or refused.
Sensitive data last	Low-risk uses go first. Financial, HR and payroll data are connected last, each with its own approval.

What makes AI risk different

Some AI risks are familiar, such as data leakage and access control. Others are new. These are the ones to plan for:

RISK	WHAT IT LOOKS LIKE	MAIN CONTROLS
Shadow AI	Staff use personal accounts, browser extensions or meeting bots nobody has assessed.	An approved tools list, an easy way to request new tools, training.
Data leakage	Customer, pricing or staff information goes into a personal account or a public chat link.	Company platform only, data classification, sharing settings.
Oversharing through search	AI search surfaces a file shared too widely years ago, such as a payroll sheet shared with everyone.	Audit and clean up sharing before connecting AI to your files.
Confident but wrong	AI invents a figure, specification or legal requirement, and it ends up in a quote or customer email.	Human review, checking against the source.
Prompt injection	An email, PDF or web page hides text telling the AI to reveal information or do something else.	Read-only connections, agent features off, user awareness.
Unintended actions	AI agents with browser control or write access change or send something nobody approved.	Off by default, specific approval, full logging.
Bias and unfairness	AI-assisted screening or disciplinary wording treats people unfairly.	No AI decisions about people, HR review.
Over-reliance	People stop checking, or lose skills.	Accountability, manager review, training.
Supplier and cost	A provider changes terms or prices, has an outage, or usage costs climb.	Vendor review, usage limits, a manual fallback.

A simple way to rate each risk is likelihood multiplied by impact, each from 1 to 3. Treat 6 to 9 before going live, 3 to 4 within 90 days, and accept and monitor 1 to 2. Give every risk an owner and a date.

Step 1: Decide your risk appetite

Agree at executive level how much risk the business will accept, by area. A typical position:

AREA	APPETITE	WHAT THAT MEANS
Everyday productivity on public and internal information	High	Encouraged, with normal checking.
Confidential business information	Moderate	On the company platform, within each person's role.
Personal information about staff, customers or suppliers	Low	Only where the purpose needs it, for authorised roles.
Payroll, HR records, ID and bank details	Very low	Only through an approved connection, after an impact assessment.
AI acting on its own in your systems	None for now	Revisit case by case.
AI making decisions about people	None	AI may help prepare. A person decides.
Passwords and credentials given to AI tools	None	Never.

Step 2: Classify your information

Four classes are enough for most businesses. The class decides where information may go.

CLASS	EXAMPLES	AI RULE
Public	Website content, brochures, published prices	Any approved tool.
Internal	Procedures, internal emails, meeting notes	Company AI platform only.
Confidential	Customer details and orders, cost prices, margins, supplier terms, designs, contracts, financial reports	Company platform, within the person's role. Use approved connections rather than pasting large extracts.
Restricted	Salaries, employee and medical records, ID numbers, bank details, passwords	Never pasted or uploaded into a chat. Approved connections only, for authorised roles. Passwords never.

Step 3: Tier your use cases

Not every use of AI needs the same approval. If it is unclear which tier applies, use the higher one.

TIER	WHEN IT APPLIES	EXAMPLES	APPROVAL
1. Low	Personal productivity on public or internal information.	Drafting internal email, summarising a procedure.	Covered by the acceptable-use policy.
2. Medium	Confidential information, anything going outside the business, or read-only connections.	Customer replies, quotes, sales analysis, order look-ups.	AI owner and data owner. A person reviews before it leaves.
3. High	Restricted information, HR or payroll, anything automated or able to act.	Payroll summaries, automated reports, write access.	Executive and data owner sign-off, impact assessment, testing, 30-day review.

> **Never allowed**
> AI making or carrying out decisions on hiring, pay, performance or dismissal. Monitoring or profiling individual staff. Company information in personal tools. Giving AI passwords. Sending, paying, approving or changing records without a person approving each action. Fake reviews or misleading claims. Getting around access controls.

Step 4: Set the level of human oversight

LEVEL	WHAT HAPPENS	WHERE IT IS USED
1. Assist	AI suggests, explains or summarises. The user decides.	The default for most work.
2. Draft and approve	AI prepares something. A named person checks it against the source before it leaves or is relied on.	Customer and supplier communications, quotes, reports to executives.
3. Act	AI carries out an action or sends something on the company's behalf.	Not at first. Needs a specific decision, assessment and full logging.

AI tools increasingly do more than chat: they control browsers, run scheduled tasks and write to systems. These carry the highest risk of unintended actions and prompt injection. Keep them off by default. Anything approved needs a named owner, the narrowest access possible, logging, and a way to stop it at once. The business process must still work without it.

Step 5: Control access

Group people by role and decide what each group can reach. Each department head signs off their own column before that data is connected.

CAPABILITY	SALES	FINANCE	HR	OPS	EXEC	IT
Chat and own documents	Yes	Yes	Yes	Yes	Yes	Yes
Own email and files	Yes	Yes	Yes	Yes	Yes	Yes
Customers and orders	Yes	Yes	No	View	Yes	No
Stock and production	View	View	No	Yes	Yes	No
Debtors and creditors	Status	Yes	No	No	Summary	No
Financial reports	No	Yes	No	No	Yes	No
Payroll and employee records	No	Totals	Yes	No	Summary	No

An illustration only. Note that IT does not get business data through AI just because it is IT.

Handle joiners, movers and leavers the same way as any other system. Remove access on the last working day, and compare AI seats with active staff every month.

Step 6: Connect AI to your systems properly

The biggest gains come when AI can answer questions from your own ERP, files and reports. This is also where most of the risk sits, so the design matters.

- **Never connect AI directly to a database.** Put a controlled gateway between them that exposes specific, approved questions.
- **Use prepared read-only views** and a restricted database login for each role.
- **Log every request,** allowed or blocked, with the person, tool, time and result.
- **Test the refusals,** not just the answers. Confirm that the wrong role is turned away.
- **Clean up before connecting files.** Audit SharePoint, OneDrive and Teams, and remove anything shared with everyone.
- **Prefer open standards,** so you can change AI provider later.

Step 7: Look after POPIA

- **Use business terms.** Confirm in the current terms that your data isn't used to train models. Personal plans aren't suitable for company work.
- **Treat the AI provider as an operator,** and keep its terms and data processing addendum on file.
- **Check cross-border transfers.** AI services often process data outside South Africa, so confirm that section 72 is satisfied before connecting sensitive personal information.
- **Minimise** personal information, and add AI to your processing records.
- **Complete an impact assessment** before high-risk uses involving personal information.
- **Include AI in incident response.** Your Information Officer decides whether an AI incident is a security compromise that must be notified under section 22.

Who does what

ROLE	RESPONSIBILITY
Executive	Approves the framework, risk appetite, budget and high-risk uses.
AI owner	One named person who owns the framework, the registers and first response to incidents.
Data owners	Department heads who sign off access to their data and recheck it quarterly.
Information Officer	POPIA compliance, impact assessments and notifications.
IT partner	Sign-in, connections, settings, audit logs, and monthly reporting on usage, cost and issues.
Line managers	Justify access, make sure staff are trained, and check AI-assisted work.
Every user	Follows the policy, checks the output, and reports mistakes the same day.

Reporting a mistake quickly should always be treated better than hiding it, or nobody will report anything.

A sensible rollout

PHASE	WHAT HAPPENS
1. Foundation	Find out what is already in use. Agree the rules and risk appetite. Set up company accounts and clean up file sharing.
2. Pilot	Eight to ten people across departments, on low-risk work. Train them, and measure what is actually useful.
3. Company data	Connect read-only sources for lower-risk areas, usually sales and operations first, with sign-off from each data owner.
4. Expand	More users and more use cases. Finance, HR and payroll last, each with its own approval.

Review usage, cost and logs monthly, access quarterly, and the framework every six months. AI changes quickly, so anything fixed once will soon be out of date.

Are you ready? A ten-point check

- ☐ We know which AI tools our staff are already using.
- ☐ We have an approved list of tools, and an easy way to request new ones.
- ☐ Everyone uses a company account with work sign-in.
- ☐ Our staff know what information must never go into AI.
- ☐ We have reviewed who can see what in SharePoint, OneDrive and Teams.
- ☐ A named person owns AI in our business.
- ☐ Anything leaving the business is checked by a person first.
- ☐ AI connections to our systems are read-only and logged.
- ☐ We have checked our provider's terms and our POPIA position.
- ☐ Staff know how to report an AI mistake, and that they won't be punished for it.

> **Fewer than seven ticks?**
> You're in good company, and the fixes are mostly straightforward. Fusion offers a free 45-minute AI conversation to work out what's worth doing first.
>
> 087 265 0624 | hello@fusiongroup.co.za | fusiongroup.co.za/ai

---

This paper is general guidance and is not legal advice. Have your legal or compliance advisor confirm your position under POPIA and other applicable law. Its structure draws on the NIST AI Risk Management Framework, ISO/IEC 42001, the Protection of Personal Information Act 4 of 2013 and King Code principles on governing technology and information, without claiming certification against any of them. © 2026 Fusion Technology Group (Pty) Ltd.