

CLIENT GUIDE

Connecting AI to your business systems

Switching AI on takes minutes. Making it safe, lawful, affordable and still working in two years takes planning. This guide explains why, and how we do it with you.

Connecting AI to your email, files or business systems takes minutes. Doing it so that it is safe, lawful, affordable and still working in two years takes careful planning, and that is the part we want to get right with you.

01 Why your request has come to a senior engineer

You asked our help desk to connect an AI tool to your company systems. **We have not said no.** We have routed it to a senior engineer, because this is a change to how your business data is accessed, not a routine support task.

Our help desk is not permitted to grant AI tools access to mailboxes, files, accounts or line-of-business systems. That rule protects you. A single approval screen can give an AI service standing access to everything the approving user can see, and sometimes everything the whole organisation can see.

We treat every AI integration as a **high-risk change request**. Before we make one, we need three things that are usually missing when the ticket arrives.

1

A risk assessment

Of what the AI will be able to read, change and send, and who it acts for.

2

A view of your environment

Your Microsoft Secure Score and how well access, sharing and devices are controlled today.

3

Your written acceptance

Of the risks and responsibilities that remain once the connection is live.

This guide explains why, so the conversation with your engineer is quick and well-informed.

02 The honest picture

Cloud and AI providers make connecting AI look like the easiest thing in the world, and technically it often is: sign in, click Allow, done. The click is the easy 20%. The other 80% decides whether it is safe.

An AI tool connected to Microsoft 365, a SaaS platform or a server does not bring new security with it. **It inherits yours.** If a folder is shared too widely today, the AI can find it in seconds. If an old account still has access, so does anything acting on its behalf. If multi-factor sign-in is patchy, an attacker who gets in now has an assistant that can search, summarise and send.

Assistants answer. Agents act.

A chat assistant answers questions. An **agent takes actions**: it sends email, updates records, moves files and calls other systems, often without a person checking each step.

	AI assistant	AI agent
What it does	Reads, searches and drafts when a person asks	Carries out tasks across systems on its own
Who is in control	The person, who reviews every answer	The agent, within whatever access it was given
Worst case	Someone sees information they should not	Data sent out, records changed or deleted, costs run up
Planning needed	Weeks: permissions, sharing, policy	Months: scope, approvals, testing, monitoring, exit plan

Most requesters we speak to do not know their current Microsoft Secure Score, and their environment has not been reviewed for AI. That is normal. **It is exactly where we start.**

03 What an AI connection can reach

The same eight risks come up in almost every AI integration request we assess. Each is manageable, but managing it before the connection goes live costs a fraction of cleaning up afterwards.

Area	What usually happens	What it costs you
Oversharing	AI search surfaces files, mailboxes and chats shared too widely years ago: salaries, HR cases, board papers, client data.	People see information they were never meant to, and nobody notices until it is quoted back
Permissions granted	One user clicks Allow and gives an app broad, standing access to mail, files or the whole tenant. Nobody reviews it again.	A third party with more access to your data than most of your staff
Unknown security baseline	Secure Score, MFA coverage, device compliance and conditional access have not been checked. Old accounts and guests still exist.	AI amplifies every weakness: a compromised account can now search and act at speed
Agents taking actions	Agents send email, change records or delete data with no human approval step and no spending or rate limits.	Wrong emails to customers, corrupted records, runaway costs, actions you cannot undo
Prompt injection	A document, email or web page contains hidden instructions that the AI follows.	Data quietly sent outside the business by a tool that was only meant to read
Data location and GDPR	Data is processed by a US or global service, logs keep personal data, and no DPA or DPIA is in place.	GDPR exposure, reportable breaches (72 hours to the ICO), failed security questionnaires
Logging and audit	No record of what the AI read, what it did or who asked it.	You cannot investigate an incident, answer a data request or prove compliance
Cost	Usage-based AI, API and licence charges with no caps or monitoring.	Bills that grow faster than the benefit, or a surprise invoice after one busy week

04 Why AI projects stall before production

We regularly meet businesses that have invested weeks or months in vibe-coded apps or home-built agents, only to find they can never put them into reliable production use. The demo worked; the business case did not survive real users and real data.

Built for a demo

It cannot scale

It worked for one person on test data. With fifty users, real volumes and real permissions it slows down, fails or costs too much to run.

It is not reliable enough

AI output varies. Without testing, guardrails and a human in the loop for important actions, occasional wrong answers become business errors.

Built on sand

The model goes obsolete

Providers retire models and change products on their own timetable, often within a year or two. An agent built tightly around one model may need rebuilding.

It is not portable

Prompts, workflows, data and logic are locked inside one platform. Moving to a better or cheaper option means starting again.

Governance

Compliance was never considered

No DPIA, no record of processing, no data location decision, no answer for a customer who asks where their data goes.

Ownership

Nobody can support it

Only the person or AI session that built it understands it. No documentation, no testing and no one on call when it breaks.

Our view is simple: **don't rent intelligence, own it.** Design for portability from day one, so the value you build stays yours whichever model or provider wins next year.

05 Who carries the risk

Connecting AI does not move responsibility for your data to us or to the AI provider. Under UK GDPR your business is the **controller**; the AI provider is a **processor** acting for you. A provider's certifications secure its platform, not the way you connect it or what you let it see.

What we ask you to sign off


When a senior engineer carries out an AI integration, we ask a director or authorised decision-maker to confirm in writing that:

1. The business wants the connection, understands its purpose and has named an owner for it.
2. The scope is agreed: which systems, which data, which users, and what the AI or agent is allowed to do.
3. The risks found in the assessment have been explained, and any not yet fixed are accepted by the business.
4. A DPIA has been completed where personal or sensitive data is involved, and the provider's DPA is in place.
5. Actions that send information outside the business, spend money or delete data keep a human approval step.
6. The business accepts responsibility for how staff use the tool, supported by an acceptable use and responsible AI policy.

This is not us stepping back. **It is making sure the right person has decided, with the facts in front of them, before business data is opened up.**

06 The safe route to production

Every AI integration we deliver follows the same route. Each gate is a decision point: if it is not met, the work goes back a step rather than forward.



↪ The route loops: when a model is retired, your data changes or the scope grows, it starts again at Assess.

07 Before anything is connected

If you cannot tick every line, the environment is not ready for AI to be connected to it.

- ☐ We know our current Microsoft Secure Score (or equivalent) and have a plan for the gaps
- ☐ Every user and admin signs in with multi-factor authentication
- ☐ Old accounts, guest users and unused app permissions have been reviewed and removed
- ☐ Sensitive files, sites and mailboxes are shared only with the people who need them
- ☐ Sensitive data is labelled, so AI tools can be told what they may not use
- ☐ We have agreed exactly which systems and data the AI or agent may reach, and what it may do
- ☐ Actions that send data out, spend money or delete data need a person to approve them
- ☐ We know where the AI service processes and stores our data, and the DPA is signed
- ☐ A DPIA has been completed where personal or sensitive data is involved
- ☐ AI activity is logged, and someone reviews it
- ☐ Spending caps and alerts are set on every usage-billed AI service
- ☐ Prompts, workflows and data can be exported, so we are not locked to one provider
- ☐ Staff have an acceptable use and responsible AI policy, and know what not to paste into AI tools
- ☐ A named owner is responsible for the integration, and a director has signed off the risks

08 How Total Group can help

The right first step is an **AI Readiness Assessment** of the platform you want to connect, whether that is Microsoft 365, a SaaS application or your own servers. It tells you, in plain English, whether it is safe to proceed, what needs fixing first, and what the connection should and should not be allowed to do.

1

Check where you are

AI Readiness Core Assessment

Our engineers scan your Microsoft 365 tenant and security posture: Secure Score, identity, sharing and devices. You get an executive summary and the priority actions.

AI Readiness

Everything in Core, plus a CREST-accredited external and internal penetration test, a full vulnerability scan and reporting on how AI is already being used.

2

Mind the gap

That's where cost, risk and viability reside.

We measure the gap between where you are and where you need to be before AI is connected, then help you close it: access, sharing, labelling, policy, DPIA and DPA. You get an honest view on whether to proceed, pilot or wait.

3

Engineer it properly

A senior engineer designs and delivers the connection or agent: least-privilege access, human approvals, logging, UK or EU data location, portability, and private AI on your own infrastructure where it fits. An AI capability you own, can move, and can show to customers, insurers and auditors. If you already have a prototype, it often becomes the specification.

4

Stay AI-ready, in perpetuity

AI Readiness as a Service runs the same checks continuously, so your posture stays safe as people, data, settings and models change, with a team accountable when something does.

Talk to a local team

01727 881 224 · hello@totalgroup.co.uk

totalgroup.co.uk · Hertfordshire, the Home Counties and London

