GECKO | AI Platform

THE AI OPERATING SYSTEM FOR HIGHER EDUCATION

The AI Governance Briefing for University IT.

Moving AI from scattered pilots to a governed institutional capability — the problem, the change needed, and the questions to put to any AI vendor. Including us.

Written for

CIOs, IT Directors and Digital Transformation leaders in higher education.

≈1M
conversations handled each month

200+
AI agents launched across institutions

<1 wk
to a production agent with managed launch

PART ONE · THE PROBLEM

Every university is already using AI. The only question is whether IT can see it.

Staff paste student queries into consumer chatbots. Departments trial point solutions on corporate cards. Academics build their own automations.

Each instance is individually reasonable — and collectively it is an ungoverned flow of institutional and personal data into systems nobody procured, with no audit trail, no accessibility review, and no answer when a regulator, an auditor, or a journalist asks: **"What is your AI doing, and how do you know?"**

59–63%

of agent conversations arrive outside office hours in published university deployments — demand no staffing model covers.

Career risk

is the real currency of this decision: an ungoverned AI incident lands on IT's desk, whoever bought the tool.

Pilot fatigue

sets in when every department's experiment is a separate contract, a separate risk review, and a separate dashboard.

Why banning fails

A ban doesn't stop AI use; it keeps it invisible. The queries keep flowing — from personal accounts, on personal devices, into free tiers whose terms nobody read. Prohibition converts a governance problem into a discovery problem.

Why pilots stall

Most university AI pilots die between demo and production — not for lack of capability, but because nobody can answer the second-meeting questions: who approved that answer, where does the data live, what did it cost, and who is accountable when it's wrong?

The institutions getting this right are doing something different: giving their teams a sanctioned path that is genuinely faster than going rogue — with governance built into the platform, rather than bolted onto a policy document.

PART TWO · THE DIRECTION OF TRAVEL

The sector has already decided. The open question is governance.

Independent research across the UK and US tells one consistent story: adoption is bottom-up and near-universal, institutional strategy is catching up fast, and agents are the next wave. The gap is governed delivery.

95%

of UK undergraduates now use generative AI — up from 92% a year earlier — yet only **36%** feel encouraged by their institution to do so, and only **38%** are provided with AI tools. Demand has outrun sanctioned provision.

HEPI / Kortext, Student Generative AI Survey 2026 (1,054 UK undergraduates) [1]

57%

of higher education institutions now describe AI as a strategic priority, up from 49% the year before — while EDUCAUSE warns of a widening **"digital AI divide"** between institutions that can resource governed adoption and those that cannot.

EDUCAUSE, 2025 AI Landscape Study (≈800 respondents) [2]

17% → 60%+

of organisations have deployed AI agents today versus the share expecting to within two years — the steepest adoption curve of any emerging technology Gartner measures. Gartner predicts **40% of enterprise applications** will carry task-specific AI agents by the end of 2026, up from under 5% in 2025.

2026 Gartner CIO & Technology Executive Survey; Gartner press release, Aug 2025 [3]

2 of 10

of EDUCAUSE's Top 10 IT Issues for 2026 are explicitly AI: **#2 "The Human Edge of AI"** and **#9 "AI-Enabled Efficiencies and Growth."** AI has moved from experiment to standing item on the IT leadership agenda.

EDUCAUSE, Top 10 IT Issues 2026 [4]

Jisc's national research adds the texture behind the numbers: students use AI daily — for study, job preparation, wellbeing and everyday admin — and their principal ask of institutions is not more tools but **clear, consistent, course-specific policy**. Even with 86% of universities publishing AI guidance, students still report being unclear about what is allowed.[5]

Read across the studies and the conclusion is hard to avoid: adoption is no longer the differentiator — everyone's students and staff are already using AI. The institutions that pull ahead will be the ones that can adopt it **governed**: visibly, safely, and at departmental scale.

SOURCES

1. HEPI & Kortext, *Student Generative Artificial Intelligence Survey 2026* (Policy Report 199; Savanta polling of 1,054 full-time UK undergraduates, December 2025). hepi.ac.uk
2. EDUCAUSE, *2025 AI Landscape Study: Into the Digital AI Divide* (≈800 respondents, surveyed November 2024). educause.edu
3. Gartner, *2026 CIO & Technology Executive Survey*; Gartner press release, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," August 2025. gartner.com
4. EDUCAUSE, *Top 10 IT Issues, 2026*. educause.edu
5. Jisc, *Student Perceptions of AI 2025* (170+ student discussions; 1,274 survey responses across UK FE and HE). jisc.ac.uk

PART THREE · THE CHANGE NEEDED

From banning tools to governing a capability.

The shift is organisational before it is technical: IT stops being the department that says no, and becomes the owner of the platform every department builds on.

Departments will keep coming — student services, admissions, the IT desk, registry, advancement. Some will build their own agents. Some will want IT alongside them. Some will want it run for them. The governance question is not *whether* they build, but whether it all lands on one governed platform with one set of controls — or on eleven separate contracts.

A reference governance model — five controls that form a workable institutional AI operating model, whatever platform you choose. They are deliberately boring; boring is what auditable looks like.

01 **A curated model allowlist**

IT connects AI providers and explicitly enables the models departments may use. Builders choose from the approved list, never the open internet. New models arrive by decision, not by default.

02 **Central prompts, local agents**

Institutional voice, policy constraints and legal language live in centrally managed, reusable prompt components. Departments build agents on top of them, not around them.

03 **A guardrail baseline that ships on**

Prompt-injection detection and credential/PII redaction enabled institution-wide from day one. Departments may tighten, never loosen. Public-facing API keys carry stricter rules than staff access.

04 **Approval gates on outward actions**

Any workflow that messages a student, changes a record or spends money pauses for a named human. Review latency is monitored like any other service metric.

05 **A weekly review habit**

Someone owns the flagged-conversation queue and the failed-run list. Governance nobody reads is theatre.

PART FOUR · THE VENDOR TEST

Seven questions to ask any AI vendor. Including us.

These separate a demo from an operating capability. Each comes with the answer we give — so you can hold us to it.

1 Where does our data live — and is it used to train models?

Demand a region and a jurisdiction, then ask the training question separately: the platform vendor and the model provider are different parties. The strongest position is bring-your-own-keys, so inference runs under the agreement your institution already negotiated.

✓ **Gecko:** independent regional deployments in Ireland, the US and Canada, each with its own database and infrastructure. Gecko does not train models on institutional data. BYO keys across seven provider families — or Gecko-managed access to start faster.

2 What happens when the best model changes?

It will — roughly every quarter. If switching models means re-implementing agents, you have signed a lock-in disguised as a platform. Ask to see a model swapped live, and ask what you can export.

✓ **Gecko:** the model behind an agent is a dropdown — OpenAI, Anthropic, Amazon Bedrock, Azure OpenAI, Google, OpenRouter, or any approved OpenAI-compatible endpoint. Tools use the open MCP standard; workflows export as portable JSON; the public API is the same one our own console runs on.

3 Who can build, who can run, and who can only look?

"Everyone is an admin" is how pilots work — and how incidents happen. Run and edit must be separate permissions: the person operating an agent on results day should not be able to change its prompt.

✓ **Gecko:** role-based access across 26 resource types with distinct read / write / run permissions; built-in admin, builder, operator and viewer roles plus custom roles; API tokens that are scoped, expiring, stored only as hashes, and capped at their creator's permissions.

4 What stands between a user's message and the model?

Prompt injection is not theoretical — instructions hidden in pasted text, invisible Unicode and encoded payloads are documented attack patterns. And users will paste things they shouldn't. Ask what screens input *before* it reaches a provider, and whether your security team can test the rules.

✓ **Gecko:** 30+ built-in input detectors for prompt injection (including invisible-Unicode smuggling and encoded instructions) and sensitive data; per-rule flag / redact / block; custom rules with an in-console test bench. Output risk is managed by grounding, approval gates and full tracing.

5 Where does a human sign?

Some actions should never be autonomous. The governance question is not "can a human review?" but *where* the review happens, how it is authorised, and whether it can be forged. Approval by Slack emoji is not a control.

✓ **Gecko:** workflows pause for approval as interactive cards in Slack or Teams (or email/SMS), backed by single-use hashed action tokens that cannot be replayed. Email code verification confirms identity mid-conversation. Approver, timestamp and review latency are recorded — average human-review time is a dashboard metric.

6 When leadership asks "is it working?", what do you show them?

If the answer is a vendor slide, you don't have observability. You need, per run: what executed, with which model, using which tools, seeing which sources, costing how many tokens, checked by whom — across every agent in one view.

✓ **Gecko:** every agent and workflow run keeps a complete step-level trace — model inputs and outputs, every tool call, retrieval sources with relevance scores, token usage, guardrail hits, approval outcomes. One unified explorer filters it all; content telemetry shows which pages actually power answers.

7 Will it pass your accessibility review — and who enforces that?

Universities are public bodies with legal accessibility duties. A student-facing AI interface that fails contrast or keyboard requirements is a compliance incident waiting for a complaint. The strongest answer is not "we have a VPAT" but "the platform will not let anyone publish an inaccessible interface."

✓ **Gecko:** web embeds target WCAG 2.2 AA with 7:1 body-text contrast — and the platform rejects themes that fail contrast validation. Enforcement in code, not guidance in a PDF. Keyboard-complete, screen-reader support, reduced motion, 200%/400% zoom, and a dyslexia-friendly font option built in.

The pattern behind all seven: governance that lives in the platform survives staff turnover, budget cycles and vendor promises. Governance that lives in a policy document survives until the first busy week.

PART FIVE · HOW THE PLATFORM HELPS

One governed platform, however each department builds.

Gecko AI Platform is the AI Operating System for higher education: some departments build their own agents, some build with IT, some have Gecko run it fully managed — and it all lands on one platform, under central IT's controls and observability.

01 Build any agent

Agents assembled from a central prompt library in minutes — or described in plain English and drafted by the AI builder for your team to review before anything is created. A visual workflow builder adds branching, parallel steps, sandboxed code and human approvals.

02 Connect it to your data

Crawl university websites (JavaScript pages and PDFs included), ingest SharePoint content, and make the SIS agent-ready through Ellucian Ethos — plus Salesforce, Gecko Engage and custom MCP servers for proprietary systems. Hybrid retrieval, with per-source relevance scores in every trace.

Ellucian Ethos · Salesforce · SharePoint · MCP

03 No vendor lock-in

Seven AI provider families with live model discovery. Bring your own keys or use ours. Swap the model behind a live agent from a dropdown. Workflows export as portable JSON. The model is a setting, not a marriage.

OpenAI · Anthropic · Bedrock · Azure · Google · OpenRouter

04 Deploy where people already are

One agent serves the university website (chat or AI search, one script tag), Microsoft Teams, Slack, and an actual phone number — an AI receptionist that routes callers or warm-transfers to a human team — plus SMS, email and API. Same knowledge, same guardrails, one trace view.

Web · Teams · Slack · Voice · SMS · Email · API

05 Govern and measure everything

Role-based access, scoped expiring API tokens, input guardrails with flag/redact/block, human approval gates with single-use cryptographic tokens, AES-256-GCM encrypted secrets that are write-only after save, and complete step-level tracing of every run. Hosted on AWS with SOC 2 and ISO 27001 assurance; regional deployments in Ireland, the US and Canada; VPAT available; console in English, Welsh and Spanish. Security pack on request.

The operating model: start fully managed — Gecko defines the outcome, connects systems, configures guardrails and goes live in under a week on average — then build together, then let trained university teams build their own. Same platform, same controls, throughout.

PART SIX · MAKING IT REAL

The first 90 days.

WEEKS 1–2

Pick one bounded process

High demand, trusted knowledge, measurable outcome. Out-of-hours enquiry handling is the classic first case for a reason: published deployments report 59–63% of conversations arriving when nobody is staffed.

WEEKS 2–4

Baseline and build

Record current volumes, handling time and response times — this becomes your evidence base. Ground the agent in approved sources only; configure guardrails, roles and a named escalation path. Managed implementation compresses this to under a week on average.

WEEKS 4–8

Run in production, review weekly

Watch the run traces, the flagged queue and the sources actually used. Fix content gaps — most "AI errors" turn out to be content errors with better logging.

WEEKS 8–12

Report and templatise

Compare against the baseline. Turn what worked — roles, guardrail set, approval policy, review cadence — into the pattern the second and third departments inherit. That pattern, not the first chatbot, is the pilot's real deliverable.

The one-page checklist

DATA & RESIDENCY

- ✓ Named hosting region per your jurisdiction
- ✓ Written "no training on our data" statement
- ✓ BYO provider keys supported
- ✓ Credentials encrypted, write-only after save

VENDOR INDEPENDENCE

- ✓ Multiple model providers, switchable without rebuild
- ✓ Open tool standard (MCP)
- ✓ Exportable workflows and configuration
- ✓ Public API equivalent to the vendor's own UI

ACCESS & RUNTIME CONTROLS

- ✓ SSO (SAML) and run-vs-edit role separation
- ✓ Expiring, scoped machine tokens
- ✓ Input screening with a redact option
- ✓ Authenticated, non-replayable approval gates

EVIDENCE & COMPLIANCE

- ✓ Step-level run traces incl. tool calls and sources
- ✓ One view across all agents and workflows
- ✓ WCAG conformance — enforced, not promised
- ✓ Assurance documentation available for review

MAKE YOUR NEXT PILOT A PATH TO PRODUCTION

Bring us one process or service bottleneck.

We'll help you define the outcome, the controls and the evidence needed to turn it into a governed production agent — and every answer in this briefing can be demonstrated live.

Book an AI strategy conversation

build@gecko.ai · gecko.ai

Security pack available on request for IT review.

GECKO | AI Platform