

■ ELEVED.AI / ATLAS

Atlas.

The AI operating map for higher education

EXECUTIVE BRIEF

This brief is the leadership view of the Atlas framework. The full operating map — 24 domains, each with its key questions and the observable signals of good governance — is the companion workbook your domain owners will use. A president reads this; a cabinet works from the map.

[Institution Name] — [Date]

THE MAP, NOT THE GPS

Why a map, not another maturity model

Most AI frameworks in higher education are maturity models: they tell you how far along you are. Atlas is an operating map: it tells you whether you are looking at the whole territory in the first place.

The distinction matters because **you cannot assess maturity in domains you didn't know existed**. Institutions rarely fail at AI because they moved too slowly in the areas they were watching. They fail in the areas nobody was watching — the embedded AI feature no one evaluated, the department that wrote its own policy, the vendor contract no one read for data rights, the assessment model that quietly stopped measuring learning.

And the default outcome — the thing that happens when no one decides otherwise — is fragmentation: every college and unit developing its own AI posture independently, producing inconsistent student experiences, ungoverned data exposure, duplicated spending, and an institution that cannot answer basic questions about its own AI use. Nearly everything in Atlas is, at bottom, a defense against fragmentation. If you take one thing from this brief: **someone must own the whole map**.

The map: six verbs

The full territory is 24 domains, but leadership needs to hold only six verbs — each answering a question a president actually asks.

NO. 01 Decide How do we govern this? FOUNDATION	NO. 02 Protect How do we avoid getting hurt? DEFENSE
NO. 03 Educate How does the academic core change? MISSION	NO. 04 Serve How do we win, keep, and launch students? OFFENSE
NO. 05 Run How do we operate and sustain the enterprise? CAPACITY	NO. 06 Equip How do we bring our people with us? CULTURE

The order encodes a soft sequence without imposing a staircase: an institution that tries to **Serve** with AI before it can **Decide** and **Protect** is taking risks it cannot see; an institution that invests in **Run** without **Equip** buys tools nobody adopts. Mature institutions work all six simultaneously.

Note the postures. **Protect** is defensive spending — avoiding the incident that ends careers and dominates a news cycle. **Serve** is offensive spending — enrollment, retention, revenue. Boards and CFOs sort investments this way instinctively; the map lets you present AI work in the fiduciary language they already speak.

Beneath the verbs, the 24 domains cover the full institutional surface: governance bodies, policy, board oversight, and monitoring (Decide); data governance, security, compliance, and vendors (Protect); teaching, assessment, AI literacy, and research (Educate); the student journey from recruitment through alumni relationship (Serve); operations, infrastructure, finance, advancement, and community (Run); and professional development, change management, and the workforce and labor dimensions of AI (Equip).

Three considerations are deliberately **not** domains, because confining them to a single chapter is how institutions get them wrong. They are lenses applied to every group: **Equity & Ethics** (every domain contains an equity question; quarantining ethics in one committee's chapter is how biased

systems ship), **Agentic AI** (autonomous AI that acts rather than responds changes the risk math everywhere at once), and **Horizon** (a standing scanning practice owned by governance, not a static list of predictions).

■ EXECUTIVE BRIEF · 03

The three instruments

Atlas turns the map into a management tool with three instruments.

The maturity scale

○ 1 UNAWARE — ○ 2 REACTIVE — ● 3 GOVERNED — ○ 4 STRATEGIC

Every domain is scored 1–4: **Unaware** (the institution hasn't recognized the domain), **Reactive** (responding to incidents as they arise, no owner or process), **Governed** (a named owner, a written position, a working process, a review cadence), **Strategic** (the domain is leveraged for differentiation and measurable outcomes). The goal is not Level 4 everywhere — part of strategy is choosing your target per domain, and for some, Level 3 is the right ceiling. The unacceptable level anywhere is Level 1: an unmade decision being made by default.

Risk tiers

T0 PROHIBITED

T1 CONSEQUENTIAL

T2 OPERATIONAL

T3 EXPLORATORY

Every AI use is classified by consequence, and oversight scales with the tier: **T0 Prohibited** (uses the institution rules out entirely — emotion inference, fully automated adverse decisions, covert monitoring), **T1 Consequential** (AI influencing decisions about individual people — admissions, aid, grading, hiring — requiring human accountability, bias testing, and central approval), **T2 Operational** (AI touching institutional data without influencing decisions about individuals — approved tools, local approval within guardrails), and **T3 Exploratory** (public data, sandboxes — light-touch, experimentation encouraged). Classification attaches to the use, not the tool, and ambiguity resolves upward. The U.S. regulatory picture is fragmented and contested; tiered governance is the posture that holds up whichever regime arrives, and it crosswalks cleanly to the standards auditors already cite (NIST AI RMF, ISO/IEC 42001).

Sequencing

NOW

NEXT

LATER

ALWAYS

Every domain carries a default tag — **NOW** (the clock is running), **NEXT** (within the first year), **LATER** (as capacity matures, deferred *intentionally*), and **ALWAYS** (the lenses) — which your institution adjusts to its context: high online enrollment raises assessment urgency, high transfer volume raises credit-mobility urgency.

■ EXECUTIVE BRIEF · 04

Where to start: the first three moves

Twelve domains carry a NOW tag, and no institution can stand up twelve workstreams in ninety days. NOW means *the clock is running*, not *do everything at once*. The entry point is three moves, in order:

1 Name the owner.

Designate the executive owner of AI strategy and charter a governance body with real authority. Brief the board that this is happening — a board that learns about AI governance a year in reads it as being managed around.

2

Inventory what is already running.

Embedded features, shadow tools, personal accounts, vendor AI switched on by default. You cannot govern what you have not seen, and every institution that looks finds more than it expected.

3

Set the floor.

Publish an interim acceptable-use position and one bright-line data rule people can remember. Interim and imperfect beats polished and absent — the ungoverned use is happening today.

Everything else grows from these three. The remaining NOW domains define the first year's core agenda, not the first quarter's.

AN HONEST WORD ON RESOURCING

Operating the map is not free: named owners, a living inventory, scheduled bias validation, tiered training. At a large institution that is realistically two to four FTE of coordination capacity plus tooling; at a small one, a fraction of several people's jobs, honestly assigned. A resource-constrained institution should run the three moves, cover the non-negotiable compliance items (FERPA, GLBA, accessibility), and put a named owner on assessment integrity — then sequence the rest deliberately. That is a defensible posture. Pretending to run all 24 domains with no capacity is not.

■ EXECUTIVE BRIEF · 05

Five realities the map is built around

These are the positions that distinguish Atlas from generic AI governance guidance — each one is developed fully in the workbook.

01

Detection is a losing arms race; assessment design is the durable response.

AI agents have publicly demonstrated the ability to complete entire online courses autonomously. For institutions with significant online enrollment this is a present capability, not a projection — and the strategic answer is redesigning how learning is verified, not buying detection tools whose false positives fall disproportionately on international students and students with disabilities.

02

Shadow AI is a governance failure, not a user failure.

Most faculty and staff already use AI at work; only about half can point to the policy that governs it. If approval takes six months, people work around the process. The only strategy that works is constructive: governed alternatives that are genuinely better than the ungoverned ones.

03

The academic core is governed by faculty.

Curriculum, assessment standards, and graduation requirements run through the faculty senate and curriculum committee — the AI governance body coordinates those decisions; it does not make them. And on unionized campuses, AI in evaluation, hiring, monitoring, and job design is a mandatory subject of bargaining. Institutions that treat this as policy rollout rather than shared governance convert an adoption challenge into a grievance pipeline.

04

Adoption failures wear technology costumes.

Most AI failures in higher education are people failures: tools without training, mandates without resourcing, "augmentation, not replacement" messaging contradicted by budget behavior. The Equip group — fluency, change management, and workforce evolution — is the difference between an institution that owns AI tools and one that uses them.

05

Governance must be evidence, not aspiration.

Every domain in the workbook defines "What Governed Looks Like" as observable signals — a policy findable in under a minute, a bias audit with documented results, at least one AI initiative sunset on evidence, an approval gate that has actually stopped a deployment. If a signal can't be checked by a third party, it isn't governance yet.

What leadership does with this

The president & cabinet	Hold the six verbs, make the three moves, and put the offense/defense framing in front of the board — including the competitive question: what happens if peer institutions adopt AI meaningfully faster, and student expectations outpace delivery, while a shrinking enrollment pool raises the stakes of both.
The board	Gets a structured education session in the first ninety days, then AI updates on a defined cadence that distinguish defensive investment (Protect) from offensive investment (Serve), with a documented risk appetite that guides administrative decisions.
The working group	Starts with the scorecard: every domain gets a named owner, a current maturity score, and a target. The gaps set the agenda; the biggest gap is the investment conversation. Re-score quarterly and chart the trend.
Everyone	Can answer the accreditor's, auditor's, or insurer's question — "what framework are you using?" — with an instrument whose controls crosswalk to the standards they cite, and whose completed scorecard is the evidence.

The full operating map is the companion document: 24 domains, six groups, three lenses, with the key questions your institution needs a position on and the observable signals that each domain is under control. Not every question requires an immediate answer. All of them require an intentional decision about when, how, and by whom they will be addressed.



Maturity models tell you how far you've come. The map tells you whether you're looking at the whole territory. You need both — start with the map.

Developed by [elevated.ai](#).

elevated.ai

THE MAP, NOT THE GPS