



FORWARD UNIVERSITY · FOR APPLICANTS AND PARTNERS

Who keeps the promises.

Every promise a university makes needs a person whose whole job is to keep it. This is Forward University's disaggregated faculty model: named mentors, single-course instructors, independent evaluators, curriculum architects, practice tutors and supervised AI.

“A university's real prospectus is not its brochure. It is its faculty model.”

THE PROBLEM, HONESTLY

One person, six jobs.

Think of the best professor you ever had. Now count what the job asks of them at once.

Design the curriculum	Deliver the teaching	Mentor the students
Mark the assessments	Research and publish	Administer the programme

The professor marks their own homework

The person who taught the unit decides whether the unit was learned. Even with total integrity, that is a conflict of interest built into the timetable.

Mentoring is nobody's actual job

It is what a good academic does after teaching, marking, designing, researching and administering. It is the first thing a busy term deletes.

Quality depends on one person's bandwidth

A brilliant lecturer with a sick child, a grant deadline or four extra tutorials becomes, for that month, a worse teacher. The student pays for it.

It cannot scale without diluting

Doubling students means doubling generalists, each stretched the same way. Quality is diluted by arithmetic, not by intent.

It cannot deliver mastery-based progression

Progression on evidence rather than on the calendar needs marking capacity that is independent of teaching capacity. One person cannot be both clocks.

None of these criticise academics. Every one of them criticises a job description.

THE PRECEDENT

Western Governors University asked the question first.

Founded in 1997 by nineteen United States governors, WGU was built to answer one question: what would a university look like if it were designed around competency rather than the credit hour? Its answer reshaped its faculty — programme mentors, single-course instructors, independent evaluators and central curriculum teams — and the university scaled to hundreds of thousands of students without the model breaking.

THE MODEL

Five governing principles.

Principle	What it means in practice
Disaggregation	The professorship is divided into specialised full-time roles. No role wears two hats that conflict.
Separation of teaching and evaluation	The person who teaches a student never marks that student's assessment. Amazezi enforces the separation in software, not in memory.
Named accountability	Every student can name their mentor. Every course names its instructor and its architect. Every promise in the prospectus maps to a role.

Principle	What it means in practice
Mastery needs service levels	Because students progress on demonstrated mastery, evaluation, feedback and guidance carry published turnaround times, tracked like any other promise.
Humans decide, AI assists	AI teaches tirelessly, drafts and flags; a named human confirms everything that counts, exactly as in the AI Oral Examination Room.

THE SPECIALISTS

Six roles where there used to be one.

Programme Mentor — **One named person, from enrolment to graduation.**

Not a personal tutor you meet once. A named professional whose entire job is your progress: your pace, your pathway, your blockers, your entry dates, your co-op timing. They call you on a fixed cadence, and they call you when the data says something has changed.

Course Instructor (subject specialist) — **Subject answers within one working day.**

Specialists who do one thing: teach their subject to students who are actively stuck in it. They hold clinics, run live sessions and answer directly — because answering is the job, not the overflow from five other jobs.

Evaluator (independent assessor) — **They never taught you. That is the point.**

Trained assessors who mark against the published rubric and nothing else. They do not know your seat, your effort or your excuses. They know your evidence. Marking is returned inside 72 hours.

Curriculum Architect — **They design the unit; they never mark it.**

Academics and industry practitioners who build competence units, rubrics and assessment instruments — then hand them to instructors to teach and evaluators to judge. Design quality stops being a lottery of who was assigned the module.

Practice Tutor (cooperative education) — **The paid-work promise, delivered by somebody.**

Your co-op is not a form you file. A practice tutor prepares you for placement, sits between you and the employer while you are there, and assesses the work-integrated learning when it is done.

AI Learning Systems Lead — **Available at 11 p.m. Never the final authority.**

Musomesa explains, drills, rehearses and scores first-pass attempts at any hour. Every one of its behaviours sits under a named academic supervisor, and every decision that counts is confirmed by a human. AI expands contact time; it never replaces academic authority.

THE WALL

The wall between coaching and judging

No one marks a student they taught. The rule is enforced in software, not in memory.

Coach side	Judge side
Wants you to pass, and helps you get there.	Has no stake in the outcome, only in the truth.
• Programme Mentor • Course Instructors • Practice Tutor • Musomesa, in practice mode	• Evaluators who never taught you • Assessment integrity review • Independent re-mark panel • Human confirmation of every AI score

A FORTNIGHT AT FORWARD

Amina's two weeks, and whose job each moment was.

When	Who	What happened
Monday, 09:40	Programme Mentor	The call that is already in the diary Amina's mentor opens with her pace, not with small talk: two units cleared, one drifting, co-op window moving to the July block. They agree the week before it happens to her.
Wednesday, 23:04	Musomesa	Stuck on gradient descent No human is awake for her. Musomesa is. It works the derivative through with her line by line, three times, then sets two practice items and logs where she struggled.
Thursday, 11:15	Course Instructor	A specialist answers personally The flagged sticking point reaches her instructor, who replies the same day with the one intuition she was missing — well inside the one-working-day promise.
Monday, 14:00	Musomesa, then an Evaluator	Oral exam: AI-scored, human-confirmed Her oral defence is scored first by AI against the published rubric, then confirmed by an evaluator who has never taught her, never met her, and cannot be charmed. Result returned in under 72 hours.
Friday, 16:30	Practice Tutor and Programme Mentor	Placement interview, then five human words Her practice tutor runs the employer interview prep and hands her over. Afterwards her mentor messages: proud of you — prepare well.

At no point did Amina wonder whose job her problem was.

ONE COURSE, FOUR PAIRS OF HANDS

How a course travels through the model.

Stage	Who owns it	What they do
Built once, centrally	Curriculum Architects	Machine Learning Foundations, 4 competence units, is built from the programme's competency register: outcomes, content, activities, the assessment, its rubric, an oral question bank with three approved variants per slot, and the grounding pack for the AI assistant.
Approved as version 1.0	Programme Chair	No course reaches a student until the academic owner of the programme signs the version.
Delivered in both pathways	Course Instructor	Weekly sessions in the Term Model, an intensive fortnight in the Block Model, plus the course room, the clinics, and supervision of the AI assistant. Three hundred students take it this cycle.
Mastery demonstrated when ready	The student	A project submission or the oral examination — the student's choice, when the evidence is ready rather than when the calendar says so.
Marked blind, inside 72 hours	Evaluators	Evaluators mark against the rubric and confirm every AI provisional score. The instructor sees results only in aggregate; the evaluator sees the student only as evidence.
Version 1.1	Architects, then the Chair	Criterion 3 confused a fifth of the cohort, so the architects add a worked example and the Chair signs it. The course improved without any student's result depending on who happened to teach them.

Four pairs of hands touch the course; one standard governs it; the marker never taught the student. That is what an institutional asset means.

EVERY PROMISE, ITS OWNER

The promises table.

A promise without a named owner and a place it is measured is only marketing.

The promise	Whose whole job it is	Where it is measured
You will never be anonymous here.	Programme Mentor	Contact cadence on the mentor board
Someone who knows the subject will answer you.	Course Instructor	First-response time, one working day
Your result means something outside Forward.	Evaluator	Independent marking, blind to teaching
You will be marked against a rubric you were shown first.	Curriculum Architect	Rubric publication date on every unit
You will do real, paid work before you graduate.	Practice Tutor	Placement and co-op assessment records
Help exists at the hour you are actually stuck.	Musomesa, supervised	Round-the-clock availability logs
A human answers for everything that counts.	Every role	The human authority rule, audited

SERVICE LEVELS

Published targets.

Commitment	Target	Measured from
Subject answer turnaround	1 working day	From a student question reaching a course instructor to a substantive reply.
Marking turnaround	72 hours	From submission to a returned evaluator decision against the published rubric.
Programme mentor calls	Every 2 weeks	Scheduled human contact with a named mentor, whatever a student's engagement level.

CASELOADS

What each role is designed to carry.

Enrolment growth translates into published hiring ratios, not silent overload. A breached caseload is a Dean's action item, visible on the same dashboards as everything else.

Role	Caseload	Service promise	Scales by
Programme Mentor	1 to 80 in year one; 1 to 120 at steady state	Contact every 2 weeks in first year; risk response in 2 working days	Hiring mentors per 120 enrolments

Role	Caseload	Service promise	Scales by
Course Instructor	1 per 250 active students per course	Subject answers in 1 working day; clinics before every window	Adding instructors per course as cohorts grow
Evaluator	About 40 assessments per week each	Marking returned in 72 hours; calibration weekly	A marking pool sized by Amagezi's forecast
Curriculum Architect	Team of 3 per about 40 live courses	Course versions each cycle; fixes in 10 working days	One team per School, growing with the catalogue
Practice Tutor	1 to 60 concurrent placements	Two visits per placement; issues acted on in 3 days	Tutors per placement volume, per region
AI Teaching Assistant	Unlimited students, one course	Instant, always on, always cited	Compute, under instructor supervision
Programme Chair	1 per programme, all 91	Boards on schedule; monthly dashboard review	Fixed by catalogue
Dean of School	1 per School, all 10	Staffing plan reviewed each term	Fixed by structure

ENFORCED IN SOFTWARE

A faculty model that lives in a chart dies in the first busy term.

So this one does not live in a chart. The roles, the separations, the caseloads and the promises are objects Amagezi knows, enforces and reports — term in, term out, whether anyone is watching or not.

Engine	What it enforces
Assignment engine	Routes every assessment to an eligible evaluator, enforcing the separation rule and balancing load against the 72-hour promise.
Scheduling engine	Mentor call cadences, clinics and boards, created as calendar objects both sides can see.
Promise engine	Every service level is a timer. Breaches notify the role holder, then the Chair, then the Dean.
Staffing forecast	Enrolment projections times the published ratios equal the hiring plan, per School, per term, automatically.
Audit trail	Every assignment, decision, override and escalation writes to one append-only trail.

What each role sees

- The Mentor's caseload board, with a nightly risk light per student.
- The Instructor's course room, with unanswered questions aged against the one-day promise.
- The Evaluator's blind marking queue, with the 72-hour clock visible.
- The Architect's studio, with per-criterion analytics and pathway equivalence checks.
- The Chair's and Dean's boards, with every promise as a tracked line: green or breached.

WHAT MAKES THE FORWARD VERSION UNIQUE

Beyond the precedent.

The AI teaching assistant, governed

A tireless assistant in every course — and then the harder thing: its human supervisor named on screen, grounded on the course, every answer logged, and banned from assessing.

Practice inside the faculty

Cooperative education is not an office down the corridor. The Practice Tutor is faculty, marking workplace learning into the Capability Transcript with the same rubric discipline as any course.

Two pathways, one standard

Architects certify Term and Block equivalence course by course, so a choice of pace never becomes a difference of standard.

Assessment anywhere, authority always

With the AI Oral Examination Room, the evaluator role reaches any student on earth, on video, with the same human confirmation of every result.

Promises with names on them

Every sentence the University says to a student maps to a role, a caseload and a timer — live in Amagezi, not filed in a policy folder.

THE NINE ROLES WE RECRUIT AGAINST

Hiring structure.

Role	Brief
Programme Mentor	Progress, pace and pathway for a named caseload of students, from enrolment to graduation.
Course Instructor	Deep subject expertise on one course, taught in both pathways, answering inside one working day.
Evaluator	Independent, rubric-bound marking, blind wherever possible, never of a student they taught.
Curriculum Architect	Courses as institutional assets: outcomes, rubrics, question banks and AI grounding packs.
Practice and Placement Tutor	Sourcing placements, brokering matches, and assessing workplace learning with the employer.
AI Teaching Assistant (Musomesa)	Software, supervised: instant grounded help, unlimited practice, and never a mark to give.
Programme Chair	One academic owner per programme: competency register, course approvals, progression and award boards.
Dean of School	Holds the chairs to account and owns the School's staffing plan across all nine roles.
External Examiner and QA Officer	Outside the delivery chain: independent sampling, and audit of separations, turnarounds and calibration.

QUESTIONS APPLICANTS ASK

Mentors, evaluators, measurement and the wall.

Mentors

Who is my Programme Mentor, and what are they actually for?

One named academic, assigned before your first week and kept for the whole programme. Their entire job is your progress: pacing, entry dates, module sequence, and the honest conversation when you are drifting. They are not marking you, so there is nothing you need to hide from them.

How often will I hear from my mentor?

A scheduled call on a published cadence, plus proactive contact whenever the risk light on your record turns. You do not have to be the one who notices you are falling behind — the caseload board notices first, and reaching out is their job, not a favour.

How many other students does my mentor carry?

A published caseload ratio, forecast from enrolment before each term and staffed against it automatically. If enrolment grows, the hiring plan grows with it. Caseloads are a number we publish, not a number we hope you never ask about.

Evaluators**Who marks my work?**

An independent Evaluator who has never taught, mentored or coached you. They mark against the published rubric and nothing else, and they work a blind queue with your identity masked wherever the artefact allows.

Does AI decide my grade?

No. AI may produce a provisional score to speed the queue, but a human Evaluator confirms or overturns it case by case. That is the human authority rule, and it is enforced in software, not by good intentions.

What if I think a decision is wrong?

Feedback is tied to a specific criterion, so a disagreement is about evidence, not taste. Escalation runs to the Programme Chair and, where needed, the Dean, and an External Examiner samples decisions independently of everyone inside the School.

Measurement**How do I know these promises are kept?**

Every service level is a timer inside Amagezi, not a line in a brochure. Marking returned inside 72 hours, course questions answered inside one working day: when a timer breaches it notifies the role holder, then the Chair, then the Dean.

Are the numbers public, or only internal?

The commitments and the reporting are published, including where we miss. Our accountability pages carry founding-intake numbers as they actually are — including zeros — because a promise you cannot check is only marketing.

Does splitting the roles change how I am graded?

It is what makes mastery grading credible. Because your judge is not your coach, we can hold a fixed standard and let you attempt again without penalty. Nothing is averaged down for having needed a second attempt.

Separation**Why separate teaching from judging at all?**

Because one person cannot both champion you and dispassionately rule on your evidence. When those jobs sit in one lap, the grade quietly becomes a report on the relationship. Separating them means your result is evidence about you.

What stops the wall from being crossed quietly?

The assignment engine will not route your work to anyone who has taught, mentored or coached you, and every assignment, override and escalation writes to an append-only audit trail. Evaluators also log any conflict and refuse the work.

How does this protect me as a student?

Three ways: you can be completely honest with your mentor without it affecting a grade; your transcript carries weight outside Forward because an independent examiner stands behind it; and no single person's mood, memory or workload can decide your degree.

THE WORDS, DEFINED**Glossary.**

Term	Meaning
Disaggregation	Dividing the traditional professorship into specialised full-time roles.
Separation of teaching and evaluation	The rule that no one marks a student they taught, enforced by Amagezi's assignment engine.
Caseload	The number of students, courses or assessments a role is designed to carry — published, not private.
Service promise	A published turnaround, such as marking in 72 hours, tracked by a timer in Amagezi.
Blind marking	Evaluation with the student's identity masked wherever the artefact allows.
Calibration	Evaluators marking a shared sample until their scores converge, held weekly.
Grounding pack	The architect-published course materials an AI assistant is permitted to answer from.
Human authority rule	No AI output becomes a student's result without a named human's confirmation.
Capability Transcript	The verified record of what a graduate can do, fed by evaluators and practice tutors.
Competence unit	The University's measure of learning volume; one unit is fifteen notional hours.

Read the live page, or talk to a person.

forwarduniversity.live/who-keeps-the-promises — every role above opens a full detail panel with its responsibilities, measurement rules and hiring process. Recruiting now: forwarduniversity.live/careers