



TALENTBANK TECH HACKATHON 2026

Continuum

Evidence-to-Outcome Career Operating Layer

for Talentbank Career OS

From evidence to outcome — one shared language that carries a candidate's proof all the way to an accountable career result, for candidates, employers, universities and Talentbank.

Mohd Hafiz Ali Bin Mohd Anuar
Team WeLead, Universiti Sains Malaysia

Mentoring & Brainstorming
Coach LENAg haz

The problem

Applications should not disappear into *silence*.

Most platforms help people create a profile, apply, and wait. Continuum turns evidence, opportunities and feedback into clear next actions — an explainable record of evidence, action and career growth over time.



One gap, felt by four groups



Candidates

Apply and wait in silence — no clarity on where they stand or what to improve.



Employers

Review many CVs without enough context — keywords, not evidence.



Universities

Receive outcome signals too late to act on employability.



Talentbank

Opportunities can go stale with no visible follow-up.

A shared evidence language across the ecosystem

Continuum is a clickable Evidence-to-Outcome Career Operating Layer that complements the Talentbank Career Passport and opportunity marketplace. It does not replace either — it makes the journey through them explainable, accountable and useful over time.

Connects

Candidates, employers, universities and Talentbank through one evidence language.

Organises

Evidence across Atlas (204 skills) and EDGE (five evidence signals).

Explains

Every application becomes a record of evidence, action and growth over time.




Continuum

Evidence-to-Outcome
Career Operating Layer

Complements, never replaces
Talentbank Passport + marketplace

Human decisions, always
AI drafts; people decide and verify

Evidence → Context → Outcome → Growth



The outcome becomes the growth action that produces stronger evidence — the loop closes and repeats.

One governed skills graph

Atlas is a governed taxonomy that connects 204 skills across recognised standards, so evidence and opportunities stay comparable across every workspace. It is computed, not a decorative taxonomy.

204

skills mapped



Three tiers


Tier 1
Broad field (ISCED-F)

Tier 2
Narrow fields inside it

Tier 3
Skills + cross-field skills

Five evidence signals


EDGE organises what a person can show into five signal areas. Evidence is mapped across them — never merged into a single score.



E·D·G·E — Reasoning, Delivery, Collaboration, Learning, Initiative


Two honest numbers, never one score

Continuum shows Resume Readiness and Evidence Depth as two separate readings. They are never added together — a low number is framed as expected, not failure.



Resume Readiness

Does your evidence speak to what a role asks for?



Evidence Depth

How much of it can be checked by a source?

Why two numbers?

A resume alone caps Evidence Depth near 33%. Naming the document behind a claim is what raises it. The two answer different questions, so they stay apart.

Meridian Map — claim vs. shown

The Meridian Map plots two layers across the five EDGE areas: what your evidence claims, and how much of it is source-backed. The gap between them is what is left to prove.

What you claim

Dashed outline — evidence you approved

What you can show

Filled shape — source-backed or verified

A per-domain gap panel shows each area's backed count — e.g. "2 of 3 backed".

5

EDGE areas, each labelled

Reasoning · Delivery · Collaboration · Learning · Initiative

The candidate decides what is shared



Granular

Consent is per-item — the candidate chooses exactly what an employer sees.



Revocable

Access can be withdrawn; a 30-day window keeps sharing time-boxed.



No verdict

Evidence to an employer never carries a percentage or score for the person.



Audited

Every share, access and change is written to an audit log.

Trust is the product: write-quality signals are hidden from employers, and nothing is shared without the candidate's action.

Build evidence, control the story

Candidates turn a CV and lived experience into approved, source-backed evidence — and decide what each employer sees.

P

Passport Intake Upload a CV or use a guided evidence survey.

R

Review & approve Every AI suggestion is tied to a source; nothing joins the profile without approval.

M

Meridian Map See claim vs. shown across the five EDGE areas.

G

Growth plan 30/60/90-day actions built from real gaps to a target role.

Boundary Continuum never ranks the candidate or issues an employability verdict.

Context and evidence, not keywords

Employers compose roles from Atlas and review candidates against real requirements — with the human always making the decision.

A

AI Role Composer Draft a role from Atlas; the field narrows to relevant skills.

E

Evidence per requirement See the evidence behind each requirement, not just CV keywords.

C

Clarifications Ask a candidate a focused question; the answer goes on the record.

R

Respectful feedback Feedback becomes a Growth Action the candidate chooses to act on.

Boundary The AI drafts and the employer decides — Continuum never auto-shortlists, auto-rejects or hires.

Aggregate insight, never the individual

Universities read cohort-level demand and alignment signals — with permission, for planning — and never an individual student's record.

C

Cohort signal Aggregate demand across fields and skills.

A

Alignment gap Where curriculum meets — or misses — market demand.

P

Permission-based Insight is shared with consent, for planning purposes.

T

Timely Outcome signals arrive early enough to act on.

Boundary Universities see aggregates only — never an individual student, ever.

Steward the ecosystem, never hire

Talentbank's Control Tower surfaces stale opportunities, operational exceptions and ecosystem signals for follow-up — and governs Atlas.

C **Control Tower** Surface stale opportunities and operational exceptions.

A **Atlas governance** Steward the one governed graph that keeps the ecosystem aligned.

E **Ecosystem signals** Follow up where the loop has stalled.

O **Operational, not decisional** Red is used only for genuine operational urgency.

Boundary Talentbank stewards Atlas and never decides who is hired.

What the AI does — and never does

The AI drafts

- **CV mapper** — Structures a CV into suggested, source-tied evidence
- **Role composer** — Drafts a role from Atlas for employer review
- **Feedback drafter** — Turns chips into respectful, requirement-focused feedback
- **Career Coach** — Maps experience to evidence; a shape to speak in, not a script

Humans decide

- ✗ **Never invents facts or examples**
- ✗ **Never auto-verifies evidence**
- ✗ **Never ranks or scores a person**
- ✗ **Never decides who is hired**

What Continuum refuses to be

These are product decisions, not missing features. Each protects a person from being reduced to a number a machine assigned.



No personality typing



No culture-fit or human ranking



No black-box hiring score



No auto-shortlist or auto-reject



**No evidence to employers
without consent**



**No single merged score for a
person**

"Not a verdict on a person. You decide. Never merged into one score."

From Lumina Graph to Continuum

The product began as Lumina Graph — an AI talent-intelligence layer. Through mentoring and brainstorming with Coach LENAgaz (three online sessions and continued guidance), assumptions were challenged and the direction was sharpened into Continuum.

Lumina Graph

- AI talent-intelligence layer
- Scoring-led framing
- Talent graph as the headline



Continuum

- Evidence-to-outcome operating layer
- Evidence & consent led — no ranking
- The accountable loop as the headline

Six ways into one system

Guided Journey

A 34-step story from evidence to the next career action

Explore Continuum

Pick a question, see how the system responds

View a Workspace

Candidate, Employer, University, Talentbank

Try it Yourself

Step into any role and move through the system

How it Works

See the whole system connect at a glance

Documents

Summary, acknowledgement and materials

A public, clickable prototype — no login, real state management, consent and AI/role boundaries enforced throughout.

Built to adopt in stages

1

Problem clarity

A sharp, specific gap: applications lost to silence, felt by four groups.

2

Trust by design

Granular, revocable consent; no ranking; human oversight everywhere.

3

Modular adoption

Six components adopted independently — Start, Connect, Scale.

4

Real prototype

A gated, clickable system with real state — not a slideshow.



The Continuum promise

No potential should be lost between a CV and a career.

Continuum makes evidence visible, opportunity accountable and every outcome useful — it does not decide who deserves an opportunity; it ensures potential is not lost because evidence is unclear, decisions are invisible, or an outcome becomes a dead end.