

AI-NATIVE EDUCATION, LEARNING AND  
DEVELOPMENT

# BuildMyIQ

One platform for authoring,  
teaching, assessing and  
credentialing — with an AI  
tutor that answers only from  
your own lessons.



---

Pre-launch, waitlist open. Prepared for educators, institution buyers and the  
technical reviewers who must approve them.

[buildmyiq.com](https://buildmyiq.com)

---

## 01 The short version

BuildMyIQ is one system for the whole business of teaching: build the course, enrol the learners, tutor them, mark their work, and issue the certificate — without a single record being re-typed between tools.

Most organisations that teach run on five to fourteen products that do not agree with one another. The rosters drift apart by lunchtime, and nobody can answer the only question that matters — is the teaching working.

BuildMyIQ replaces that stack with one set of connected records. Every one of those is the same record the administrator sees, not a nightly copy of it.

### THE PROMISE

**Teach, assess and credential in one place — and let the AI do the drafting, never the deciding.**

# In numbers

---

**1**

## **Set of records**

Courses, enrolments, progress, assessments and certificates read and write one source of truth.

---

**6**

## **Question types**

The objective ones are marked against your stored answer key, never by a model.

---

**0**

## **AI-written grades**

Drafts wait for a person to accept or reject them.

---

**Free**

## **A real starting tier**

Tiers differ by capacity, not by which features are switched on.

---

# Who it is for

---

## Universities, colleges and online programmes

Paying twice for a learning system and a records system that still have to be reconciled. Here they are the same row, so the reconciliation stops being a job.

---

## Training companies and certification providers

Your product is the credential: a workspace-unique serial, a full issue, expire and revoke lifecycle, and a stored PDF.

---

## Schools and districts

One governed platform, one identity per person, and learner records held to a higher access bar than everything else.

---

## Corporate learning and development teams

Onboarding, upskilling and mandatory training are three shapes of the same problem, and all three run on one engine.

---

# Who it is for, continued

---

## Independent educators, creators and tutoring services

A genuine free tier, published prices with no sales call, and the same authoring and assessment engine the institutions get.

---

## IT and identity teams

Single sign-on with the identity provider you already run, and authorisation decided before a request ever reaches the product.

---

## Data protection officers and compliance leads

Learner education records are sensitive by default, and every action lands on an activity timeline with an actor and a result.

---

## 03 What breaks today

None of these problems are exotic. They are the ordinary, expensive friction of running education on tools that were never designed to know about one another.

### **Five to fourteen systems that disagree**

Rosters, sections and grades live in two products, drift apart, and someone spends their week reconciling spreadsheets.

---

One connected set of records, with nothing copied between systems on the way.

### **A learner gets stuck and waits**

Most learners fill the gap by giving up on that lesson, and the educator never finds out it happened.

---

A tutor inside the lesson player that has the lesson text in hand and leads with a hint.

### **Marking eats the teaching week**

Grading is the part of teaching that scales worst, and where feedback arrives too late to change anything.

---

Objective questions marked instantly; written answers drafted for an educator to accept or reject.

## What breaks today, continued

### Nobody can put a budget on AI

The feature ships either crippled or unpriced, and finance blocks it either way.

---

Two separate monthly allowances, checked before any provider is contacted. Reaching the cap is a clear refusal.

### Completion does not prove competence

A finished course tells an employer that someone sat through something.

---

A skill is a first-class object mapped to courses and lessons, and a certificate attests the skills a learner proved.

### Buying it is a negotiation, not a decision

You cannot budget for next year without a sales cycle, and the number depends on how well you negotiate.

---

Published tiers with published limits, permanent regional bands, and a free tier you can evaluate on.

---

## 04 What you get

Nine areas that share one set of records. The value compounds because none of them need to be integrated with each other.

### DELIVER

#### Learning management

Courses with a draft, published and archived lifecycle; ordered modules; lessons as video, article, interactive, live session, quiz or assignment.

### RECORD

#### Learner records

Enrolments, cohorts, sections and completion state. A records layer rather than a full student-information system.

### AUTHOR

#### Authoring

Structure modules and lessons, attach media through signed-URL storage, target the skill, and publish through validated transitions.

### TUTOR

#### AI tutoring

A panel inside the lesson player that retrieves the lesson being studied plus the closest related ones, and answers from that alone.

### ASSESS

#### Assessment and grading

Passing marks and attempt limits, six question types, answer keys held server-side, and drafts for educator review.

### CREDENTIAL

#### Certificates and skills

A workspace-unique serial, a status, a final score, an optional expiry and a stored PDF, with revocation visible on lookup.

# What you get, continued

## COHORT

### Cohorts and sections

Take a self-paced course and run it as a scheduled, time-boxed group with a facilitator. Same course, same enrolments, no second roster to keep in step.

## INSIGHT

### Learning analytics

Completion and pass-rate views with per-learner and per-cohort progress. Cohort dashboards, mastery views and early-warning signals are maturing rather than finished.

## GOVERN

### Identity and governance

One identity holding many scoped roles, six seeded product roles, authorisation enforced before the request reaches the product, and an activity timeline behind it all.

# Put to work: the learner



## INDEPENDENT LEARNER

### Find the right course and start in minutes

Progress is written per lesson, so tomorrow the platform opens on the next thing rather than the front page.



## INDEPENDENT LEARNER

### A tutor that can only answer from your own lessons

Help at eleven at night that agrees with the course, and a record of what it was built from.




## CAREER CHANGER

### Turn finished courses into a credential you can show

A portable record of proven skills, and an issuer that can stand behind it.


# Put to work: the educator



## INSTRUCTOR

### Get the marking week back

Feedback reaches learners in days rather than weeks, without anyone lowering their standard for a grade.



## INSTRUCTOR

### AI drafts the grade, you still commit it

Hours of marking become minutes of review. Available through the API today; the grading queue screen is next.



## SCHOOL ADMINISTRATOR

### Read exactly what the AI told your learners

Oversight backed by transcripts rather than assurances. Built behind the API today; the review console is a named gap.

---

## 06 Built AI-native

**Retrieval first, model second —  
and nothing an AI produces lands  
on a learner's record by itself.**

---

### PRINCIPLE ONE

The platform goes to its own records first: the lesson being studied, plus the closest related lessons in that course. Only then is a model called, instructed to answer from that content alone.

---

### PRINCIPLE TWO

A model is used only where a person would genuinely have to exercise judgement. Choosing the next lesson needs none, so no model runs — no latency, and no invented lesson title.

---

### PRINCIPLE THREE


AI proposes and a person disposes. There is no code path by which an AI writes a score, a pass mark, feedback or a completion; the path does not exist rather than being switched off by policy.

# Four pillars of the AI design



## **Grounded in your own courses**

Your lesson text is the tutor's only permitted source, and retrieval never leaves the workspace that owns the course.



## **Drafts, never silent actions**

It drafts; an educator accepts or rejects, and the committed decision carries their name and a timestamp.



## **AI only where judgement is needed**

Answer keys are compared, not judged. Course order is followed, not guessed.



## **Metered, scoped and on the record**

Each request is checked against your monthly allowance first and stored with the provider, model and token counts.

---

# What the AI does today

---

## Answers grounded in your own lessons

If the material does not cover the question, the tutor says so instead of guessing.

---

## Objective questions never reach a model

The objective half of a grade is exact, instant and costs nothing.

---

## Hints first, answers second

Off-topic requests are declined, so it stays a tutor rather than becoming a general chatbot.

---

## Monthly allowances checked before the call

Reaching the cap is a plain refusal with an explanation, never a silent overage.

---

## Explanation depth that follows recent scores

A struggling learner gets one idea at a time; a learner well ahead gets the deeper connections.

---

## Three providers behind one router

The learner's question is saved before any provider is contacted, so a failed call means retry rather than lost work.

---

## 06 What is in progress

---

### Rubric-anchored grading drafts

A proposed score and one to three sentences of feedback per question, using the private rubric notes the instructor wrote. It works through the API today; the educator screen is the next piece of work.

---

### What to do next, with the reason written out

The platform walks your published course order and returns the first unfinished lesson with a rationale and a pacing hint. No model runs, so it cannot invent a lesson.

---

### The lessons behind each answer, recorded

The sources are captured on every answer today; the citations drawn under the answer are still being built.

---

### Role-gated transcript review

The permission split is enforced today; the review console and the ability to annotate or redact an answer are not yet built.

# What keeps it in bounds

---

## Grounding never leaves your organisation

One institution's material can never ground another institution's answer.

---

## Conversations are private to the learner

A thread is scoped by authenticated identity rather than by an identifier in the request, so it cannot be read by guessing.

---

## Nothing runs unasked

There is no scheduled or autonomous AI in the product. Every call is triggered by a person in the moment.

---

## AI runs on the same permissions as people

Drafting a grade requires the right to award one, so AI cannot be a back door around authorisation.

---

## Rubric notes stay instructor-only

Learners never see them, and a question returned to a client carries no correct-answer field.

---

## Every answer carries its provenance

A decision taken without a model records no provider, which makes the difference visible.

---

## A person commits every consequential action

Committing a grade runs through the ordinary grading action with the grader's identity and timestamp on it.

---

## No training on your content

A request sends only what it needs: the lesson text and the question, or the question, the rubric note and the answer.

---

## Accept and reject outcomes are stored

The accept rate becomes a measurable quality signal rather than an impression.

STATED PLAINLY

# What the AI does not do yet

- Retrieval ranks lessons by wording overlap rather than by meaning, so a question phrased very differently from the lesson can miss it.
- There is no way for an educator to annotate, correct or redact a specific AI answer.
- AI-assisted authoring does not exist at all, and the adaptive engine in its full sense is a roadmap item.
- Grading drafts, next-lesson suggestions, draft history and transcript review all work through the API and have no user-facing screen.
- There is no usage dashboard, and no warning before an allowance is reached.
- There is no benchmark yet for answer quality, and none for agreement between a drafted grade and the grade finally awarded.

# Who does what

Six product roles are seeded automatically into every workspace. A role grants one level per resource type, and leaving a resource out of a role grants nothing on it. The omission is the access control, not a hidden setting.

|                               |                                                                                                                                                                                                       |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Learning Administrator</b> | Owns the catalogue, skills taxonomy, cohorts, learner records and issued credentials. The only role with destructive rights, because accountability for learner data should sit with one named owner. |
| <b>Program Lead</b>           | Takes courses and assessments through their full lifecycle, schedules cohorts and oversees grading — stopping short of destructive rights on enrolments, cohorts and certificates.                    |
| <b>Course Author</b>          | Structures courses, defines skills, and writes assessments and questions. Enrolments, submissions, cohorts and certificates are omitted from the role entirely.                                       |
| <b>Instructor</b>             | Runs the live cohort: roster, progress, assessments and grading. The catalogue stays read-only because grants are scoped per workspace rather than per course.                                        |
| <b>Registrar</b>              | Enrols learners, keeps cohorts accurate, and issues, expires or revokes certificates. Submissions are deliberately omitted — the records desk never needs a learner's raw answers.                    |
| <b>Learner</b>                | Their own enrolments, results and certificates come back from queries that take no learner identifier, so they keep a full personal record with no workspace-wide list right.                         |

## Under the hood

Delivered as software as a service and built API-first. State is served over a single federated GraphQL API that also carries every other capability in the suite, so a client fetches a course, its lessons, its cover file, its tags and its threads in one request rather than five.

Every query and mutation carries an access rule the API gateway evaluates before the request reaches the product. The product itself never validates a token and never re-implements authorisation.

Three owned domains: catalogue, learning and assessment. Everything cross-cutting is inherited from the shared platform rather than rebuilt, which is why each release arrives with governance already attached.

### THE SPLIT THAT MATTERS

Authorisation is decided before the request reaches the product. The rule lives with the operation, the decision is made once at the edge, and the product's own workspace filter is defence in depth rather than the only line.

# Capability by area

---

## Catalogue

Courses with a lifecycle, a category, a level and a target skill; ordered modules; six lesson types; media in signed-URL storage.

---

## Credentials

A workspace-unique serial, a status, a final score, an optional expiry and a stored PDF. Revoking flips the status.

---

## Learning

Enrolments carrying progress and completion, per-lesson progress as the fine-grained signal, and scheduled cohorts with a facilitator.

---

## Identity and access

Ten workspace resource types back six seeded roles; sensitive learner records are excluded from read-only grants by default.

---

## Assessment

Passing marks and attempt limits, six question types, and submissions auto-scored or graded manually with written feedback.

---

## AI runtime

Three providers behind one router with a per-attempt time limit, retries across the chain and a ceiling on answer length.

---

## Assessment integrity

Answer keys are never returned to a client, so an editor-level grant on questions cannot leak a grading key.

---

## Metering and quotas

Courses, enrolments, assessments, submissions and both AI allowances are checked before the operation rather than after the invoice.

## Open standards it speaks

- OpenID Connect for single sign-on with your existing identity provider
- OAuth 2.0, including device-code flow and client credentials
- GraphQL, served as one federated API across the whole suite
- Model Context Protocol, so AI clients act through the same governed operations
- TLS in transit and AES-256 encryption for stored data
- WCAG 2.1 AA as the accessibility target, with a VPAT recording what is verified
- The ISO 27001 framework as the model for security practice, not a held certification

## Plugs into

- Enterprise identity providers over single sign-on
- Roster and human-resources systems for learner and staff records
- Webhook delivery into your own automation and data pipelines
- Email and in-app notification channels
- Calendar for live sessions, cohort windows and due dates
- Analytics destinations for learning and product events
- AI providers: Anthropic, OpenAI and Google Gemini behind one router

# How it deploys

## **Multi-tenant cloud**

The delivery model today. Isolation is enforced at authorisation and again in the query.

## **Composed into your workspace**

One sign-in, one set of roles and one activity trail across the suite.

## **Dedicated tenancy**

Planned and sales-led: isolated database and compute, from \$120,000 a year. Not yet available to buy.

## **Self-hosted**

Planned and sales-led: from \$90,000 a year for up to 2,000 employees. Not yet available to buy.

---

## 09 Security and trust

Security here is mostly a set of decisions about who cannot see what, made early enough that they shaped the data model.

---

### Authorisation decided before the product sees the request

One place to audit rather than one per feature.

---

### Workspace filtering as defence in depth

A permission cannot reach another organisation's records even if the first decision were wrong.

---

### Learner education records held above read-only

A plain read-only grant does not inherit read or list on enrolments, submissions or certificates.

---

### Self-scoped access for learners

Queries take no learner identifier and filter on the calling identity.

---

### Answer keys never leave the server

A question returned to a client carries no correct-answer field.

---

### Least-privilege role catalogue, seeded not configured

One level per resource type, and omitting a resource grants nothing on it.

---

### Escalation ceiling on the organisation preset

A preset holder cannot assign a role of higher priority than the one they hold.

---

### Quota enforcement before the work runs

Reaching a cap is a clear refusal rather than an overage discovered later.

---

### Activity timeline and attributable change

Actor, target and result on a searchable, exportable timeline. One identity with many roles never makes a change anonymous.

STATED PLAINLY

## **BuildMyIQ is pre-launch and holds no third-party security certification.**

Practices are designed against the ISO 27001 framework, SOC 2 criteria and GDPR principles — but designing against a framework is not being certified to it.

Accessibility is audited against WCAG 2.1 AA with automated scanning; the VPAT records which criteria are verified and which are not yet evaluated.

Any certification will be claimed only once independently awarded.

The privacy-by-design data handling work is in progress toward launch rather than finished. Compliance evidence export and SIEM streaming are roadmap.

Two access-control limits: permissions are scoped per workspace rather than per course, and certificate revocation is not separable from issuance.

# One product, many front doors

## AVAILABLE

### Web application

The full product: catalogue, lesson player with the AI tutor, authoring, assessment, cohorts, certificates and analytics.

### Composed workspace experience

The same product inside a shared application shell, under one sign-in.

### Marketing site

Product story, use cases, the honest AI status page and the launch waitlist.

## PUBLISHED OR IN PROGRESS

### MCP server

31 product tools over the same authenticated API a person uses, acting with the operator's own identity. Destructive tools have no confirmation step.

### Documentation

Product guides, API reference and integration tutorials, published at the documentation subdomain.

## PLANNED

### Public API, CLI, Node and Python SDKs

Application-scoped access and automation against the same governed operations.

### Desktop and mobile applications

Native workspaces for authoring and operations, and iOS and Android for lessons, tutoring and progress.

### Browser and editor extensions

Capture source material, run quick actions, and support technical-curriculum workflows.

## 11 Plans

Pricing is published rather than quoted, and tiers differ by capacity rather than by which features are switched on. Only Free and Starter are live in billing today.

### FREE

#### \$0 forever

15 active learners, 1 live course, 50 enrolments a month

30 tutor questions and 20 grading drafts a month

Live today. No card required.

### STARTER

#### \$99/mo

50 active learners, 3 live courses, 1,000 enrolments a month

500 tutor questions and 300 grading drafts a month

Live today. \$79/mo billed annually.

### GROWTH

#### \$399/mo

250 active learners, 15 live courses, 3 workspaces

2,500 tutor questions and 2,500 grading drafts a month

Priced, not yet purchasable.

### BUSINESS

#### \$999/mo

1,000 active learners, 50 seats, 10 workspaces

25,000 tutor questions and 25,000 grading drafts a month

Priced, not yet purchasable.

### ENTERPRISE

#### From \$12 per learner

Custom quotas, unlimited seats and workspaces, dedicated tenancy or self-hosting

Dedicated service-level agreement from 99.95%

Sales-led, annual contract.

# What you can actually buy today

## Free and Starter are live

Quotas are enforced on courses, enrolments, assessments, submissions, tutor questions and grading drafts. Growth, Business and Enterprise are documented and priced but not yet purchasable.

## Regional pricing, not a discount

Four permanent market bands. Starter is \$99 in Band A and \$35 in Band D. Band follows a verified billing address rather than an IP address; automatic detection is still roadmap.

## Add-ons and enterprise shapes

Planned rather than live: learner packs from \$39, extra courses from \$29, an AI Pro pack at \$79 a month. Published enterprise starting points: dedicated tenancy from \$120,000 a year.

## Discount programmes await billing support

Annual billing is already about 20% below monthly. Education and non-profit pricing at 50%, a startup programme at 40%, and volume tiers are documented policy pending coupon support.

---

# Why teams choose it

---

## One set of records, not an integration

A course becomes an enrolment becomes tracked progress becomes a graded submission becomes a certificate, with nothing copied on the way.

---

## A skill is a first-class object

Skills are mapped to courses and lessons and attested by certificates, so a credential says what someone can do rather than what they attended.

---

## Retrieval before the model, by construction

The lessons that grounded each answer are stored on the message, so grounding is checkable rather than a marketing claim.

---

## Credentials with a real revoke

A lookup proves standing at the moment of checking rather than at the moment of issue.

---

## Determinism where determinism belongs

Objective questions are compared with your answer key; the next lesson is the first unfinished one in your published order. Neither calls a model.

---

## Governance inherited, not rebuilt

Tenancy, roles, audit, billing, quotas, files and accessibility come from the shared platform, so the engineering here goes into learning instead.

# Honest answers to hard questions

---

## **"You are pre-launch. Why would I commit to you?"**

You should not commit blind, and this document deliberately makes that easy to avoid: every capability is labelled available, in progress or roadmap, and the free tier is real enough to evaluate the authoring, assessment, tutoring and credential paths without a contract.

## **"AI marking on a real transcript is a liability we will not accept."**

So do not accept it — the product does not offer it. It produces a proposal in a separate record with a written rationale, an educator accepts or rejects it, and committing runs through ordinary grading with the grader's name on it.

# Honest answers, continued

---

## **"Adaptive learning is your headline and it is not built."**

Correct. What exists is the mechanism underneath it: per-lesson progress, submissions and grades on every enrolment, a mastery reading from recent scores, and a next-lesson decision with a rationale. What does not exist is adaptation across a whole skill journey, and the learner-facing card.

## **"Is it a real student information system?"**

Not yet. What exists is a records layer sharing one source of truth with learning and assessment, which removes the reconciliation problem. What does not exist is attendance, transcripts, degree progress, timetabling and financial integration.

# Where it sits against the alternatives

## VS ACADEMIC LMS

They leave you buying a records system, a plagiarism tool and a media host separately. The concession: decades of hardened workflow and formal certifications we do not have.

## VS CORPORATE PLATFORMS

They ignore academic records and treat skills intelligence as a paid add-on. The concession: their content marketplaces and enterprise track record are real, and ours are pre-launch.

## VS OPEN-SOURCE PLATFORMS

The cost moves rather than disappearing — into hosting, plugins and upgrades. The concession: nothing here is open source, and self-hosting is planned rather than available.

## VS STUDENT INFORMATION SYSTEMS

We remove the sync by making records and learning the same rows. The concession: if you need a registrar-grade system on day one, keep yours.

## VS CONTENT MARKETPLACES

Here the tutor is grounded in your own material and the credential is yours to issue and revoke. The concession: no catalogue of ready-made content, and no marketplace yet.

## VS GENERAL-PURPOSE TOOLS

It stops working when you need a defensible completion record or an access boundary around learner data. The concession: for a dozen learners, the spreadsheet genuinely is cheaper.

START TODAY

# Look at the honest version before you look at a demo.

BuildMyIQ is pre-launch and the waitlist is open. The status page for the AI work marks every capability shipped, in progress or roadmap.

---

WEBSITE

[buildmyiq.com](https://buildmyiq.com)

---

HOW THE AI ACTUALLY WORKS

[buildmyiq.com/ai](https://buildmyiq.com/ai)

---

SECURITY AND ACCESSIBILITY

[buildmyiq.com/security](https://buildmyiq.com/security)

---

TALK TO US

[contact@buildmyiq.com](mailto:contact@buildmyiq.com)



A product of Burdenoff Consultancy Services Pvt. Ltd. Pre-launch: capabilities marked in progress or roadmap are not available today.