

About ReadyEngine™ and the readiness resources.

A set of readiness resources built to increase the number of tradespeople in Canada who earn certification — a Certificate of Qualification, carrying a Red Seal endorsement in the trades that have one — by improving pass rates on first and subsequent attempts.

ReadyEngine™ is the measurement engine they run on.

IN SUMMARY

- 01 A complete set of readiness resources per trade — skill map, teaching content and question bank — built on ReadyEngine™, which turns every answer into evidence about one specific skill.
- 02 Most regulated training measures what was delivered. This measures what was learned, continuously.
- 03 The gap it is best placed to close is structural: apprentices learn a common standard through whatever work their employer happened to have.
- 04 It is built on retrieval practice, the most robust finding in learning science, which makes every measurement a learning event rather than a cost.
- 05 It never claims certainty. Scores cap at 95% in the code, and they can fall.
- 06 It applies before a first sitting, at the challenge route, and after a failed attempt — where one provincial certifying body's exam support unit reached 5,852 candidates in 2024, of whom 26% passed on retake.
- 07 No trades result exists yet. The next useful step is a working session and a demonstration, so a standards and certification team can judge the mechanism for themselves.



Resources per trade, not a platform to fill.

The aim is to increase the number of tradespeople in Canada who earn certification — a Certificate of Qualification, carrying a Red Seal endorsement in the trades that have one — by improving pass rates on first and subsequent attempts.

There are two parts. ReadyEngine™ is the measurement engine. On top of it sit the readiness resources: a complete set for each trade, authored one trade at a time against that trade's published standard, so an institution already delivering the training has something to strengthen it with rather than a tool waiting for someone to supply the content.

The skill map	The teaching content	The question bank
The trade's published standard turned into an addressable structure — every skill named, so every answer can attach to one.	Lessons and explanations authored against that structure, covering the standard rather than whatever the job happened to include.	Items written to depth for each skill, because a skill with three questions behind it cannot be measured, however good the engine is.

On top of that sit the parts a learner recognises: a daily study plan, practice exams, spaced review, a tutor and the readiness score.

Who this is built for

Not a replacement for in-school training, and not a competing qualification. These are resources for the organisations already carrying that work — union training centres, colleges and polytechnics, and other training delivery agents — to put in front of their own apprentices, alongside their own curriculum. Certification stays entirely with the certifying body.

Built trade by trade, and portable across jurisdictions

Each trade is a content project in its own right: map the standard, author the material, write items to depth, validate. That work is substantial and it is ours to do. Once a trade is built, the same resources serve an apprentice in any jurisdiction that adopts them — what changes between provinces is the standard the engine maps to, not the engine or the teaching content.

What the engine underneath actually does.

ReadyEngine™ is the measurement layer that makes the programme above it more than a course. It can run under material an institution already owns and teaches; in the trade programmes it runs under material we author. Either way, what it does is the same.

Every answer is about one skill


Each question tells us something about one specific skill in the standard — not just whether the learner got it right.

Scores build upward

Skill scores combine into topic scores, and topic scores combine into one readiness number.

One number, and the reasons behind it

The readiness score says whether someone is ready. The skill scores underneath say what for — and what they are not. An exam result gives only the first, and only at the end.



Every question a learner answers updates all of it — so the picture is current, not a snapshot taken at the end.

A DIAGRAM OF HOW IT WORKS, NOT A SCREEN CAPTURE

What the learner sees is an adaptive study experience: lessons, quizzes, flashcards, practice exams, a daily plan and a tutor. What sits underneath is a continuously updated picture of which parts of the standard they hold and which they do not.

What an institution sees is the same picture at cohort level — which skills a class is arriving weak on, while there is still term left to do something about it. That is the part no exam result can give them, because it arrives after the cohort has dispersed.

What an employer sees is their own apprentices, in their own reporting: where each one stands against the standard, who is ready to write, and who is not yet. A sponsor who has invested years in an apprentice currently learns how it went at the same moment everyone else does. And because uneven on-the-job exposure is the gap this instrument is best at finding, the same report shows a sponsor which parts of the standard their own work is not covering.

Four regulated domains, one engine.

ReadyEngine™ was built for one regulated licensing exam and has since been carried into three more. Each new credential comes online by mapping its standard into the structure, not by rebuilding the platform.

DOMAIN	PROGRAMME	STATUS
Accounting	Professional accounting exam preparation. The original build.	LIVE
Real estate	Regulator-accredited licensing. The deployment our published results come from.	LIVE
Securities	Banking and securities licensing.	LIVE
Insurance	Life licence qualification — approved by the regulator for mandatory pre-licensing training in Alberta.	LIVE

Two of the four sit inside programmes a regulator has accredited or approved, where the training itself has to satisfy the regulator before it can be offered at all. The trades would be the fifth domain, and the first where the standard is set by a public certifying body rather than a professional regulator.

What a learner actually gets

Daily study plan

A fresh plan each day, built from where concept mastery sits and which lessons are still unwatched.

Quizzes

The primary interaction. Every answer is evidence about one named skill, with feedback attached.

Practice exams

Timed mock exams following the published blueprint mix, so the format stops being unfamiliar.

Concept mastery

The per-skill picture, updated continuously — visible to the learner, not held back for a report.

Lessons

Short instruction written against the standard, so every lesson attaches to named skills rather than to a chapter.

Flashcards

Spaced review on a fixed schedule, for terms and definitions between study sessions.

An AI tutor


Available at any hour and grounded in the exam content. It can explain a concept, and hold the conversation, in the learner's own language.

ReadyRating™


One readiness number on a 1–5 scale, recalculated after every graded interaction and capped at 95%.

The same instrument, on screen.


Live screens from one of those deployments — a securities licensing exam, not a trade. The engine is the same one; only the standard behind it changes.



- **The learner's dashboard.** Today's plan, the stage they are in, concept mastery, the readiness rating, and days to exam.



- **Every skill, scored separately.** Each square is one concept in the standard, coloured by mastery level. This is the picture an exam result cannot give you, and it exists from week one.



- **Readiness, stated plainly.** At level 4 of 5 it says well-prepared, aim for excellence — not that a pass is assured. The prediction underneath is capped at 95%.

Most training measures delivery. This measures knowledge.

No instrument measures competence directly. Every one infers it from a sample of behaviour, and what separates a strong instrument from a weak one is how close that sample sits to the real task — and whether anyone has checked that the inference holds. Placed on that ladder, the category ReadyEngine™ competes in sits near the bottom.

- | | |
|---|---|
| 7 | Sustained job performance over years — what everyone means by competence. Not directly measurable by anyone. |
| 6 | Independent practical assessment — shop test, road test. Closest to the task, expensive, samples once. |
| 5 | Certifying examination — broad and defensible, and once. |
| 4 | Continuous knowledge mastery — a separate score for every skill, updated with every question answered. This is where ReadyEngine sits. |
| 3 | A quiz passed once — a single sample, rarely revisited. |
| 2 | Module or course completion — a record that content was delivered. |
| 1 | Hours logged — attendance, standing in for everything above it. |

Most regulated training operates on the bottom two rungs and lets everyone in the chain infer competence from the record. Moving the measurement to the fourth is a real step of two rungs — and it is not the top of the ladder. Any claim implying the sixth or seventh rung is an overclaim, and we treat it as one internally.

Three things that are genuinely uncommon

The score can go down. A question bank cannot say that. It is the plainest proof that something is being measured rather than counted.

We write both the skill map and the questions ourselves, and we are the accredited course publisher as well as the software company — so when one area is measured thinly, we fix it by writing more questions rather than by adjusting a model.

The ceiling is built in. Every readiness score caps at 95%, in the code, on every deployment. Nobody else in this category has a number that stops short of certainty.

These are hypotheses. Nobody has published the answer.

Remarkably little is published on why candidates fail trade certifying exams. Pass rates are reported; the analysis behind them is not, in any Canadian jurisdiction we can find. What follows are working hypotheses drawn from adjacent research and our own experience in four regulated domains. The list is certainly not complete.

POSSIBLE CAUSE	WHAT IT WOULD TAKE TO ADDRESS	IN SCOPE
Uneven on-the-job exposure	The employer decides what work an apprentice actually sees. Measuring against the whole standard finds what the job never covered.	YES
Regional and sector specialisation	Per-requirement diagnosis that is blind to which industry the hours were logged in.	YES
Decay across a multi-year apprenticeship	Continuous measurement and spaced review across the whole arc, rather than a cram before the sitting.	YES
No diagnosis before the sitting	A readiness signal visible weeks out, concept by concept, so the first feedback is not the result.	YES
Unfamiliarity with the exam format	Repeated exposure to the question form reduces novelty. Test anxiety itself is not ours to fix.	PARTLY
Language and reading load	The tutor can explain a concept, and translate it, in the learner's first language while they study. Not available in the exam room.	PARTLY
Financial and scheduling pressure	Time off work, travel, fee relief, sitting availability.	NO

The first hypothesis is the one with direct Canadian evidence behind it. In Statistics Canada's National Apprenticeship Survey, about 80% of apprentices said their work experience had prepared them well for the exam — which is to say roughly one in five said it had not, rising to more than one in four among carpenters. The last row we do not claim at all: a candidate who cannot get the day off work is not helped by a better study plan.

So the claim is narrower than a pass rate. The instrument says whether the knowledge is there before a seat is booked — and refuses to say someone is ready when it cannot see that they are.

There are almost certainly causes we have not thought of. A certifying body's own research is a better source for this table than ours, and if it names different factors we would rather rebuild it than defend this.

The standard is common. The work is local.

Of the causes above, two are worth taking on their own, because they are structural rather than individual — and because they are the reason a certifying body sees pass-rate variation it cannot attribute to candidate quality.



What an apprentice is examined on, and what they were exposed to

A trade standard describes the whole trade. The work available in a given shop, sector or province covers a slice of it. Two apprentices can complete the same hours, in the same trade, against the same standard, and have met entirely different parts of it — because their employers had different work. Both are then examined on all of it.

That gap is not a training failure and not a candidate failure. It is the predictable consequence of learning a common standard through whatever work happened to be available. It is also invisible: nobody involved — the apprentice, the sponsor, the training provider — holds a map of which parts of the standard the job never touched.

By employer

One shop runs service work, another runs installation. The hours are identical and the exposure is not.

By sector

The same trade sits in oil and gas, in manufacturing, in food processing, in institutional maintenance — each covering a different part of the standard.

By jurisdiction

Provincial economies differ, so the typical exposure profile of an apprentice differs with them, while the standard does not.

This is the clearest case for measuring skill by skill. A gap created by uneven exposure is a knowledge gap against a published standard, which is exactly what this is built to find. It measures against the standard rather than against the work history — so it does not care which sector the hours came from, and it surfaces the untouched areas while there is still time to close them.

An exam finds this gap once, at the end, and reports it as a fail. Measuring skill by skill finds the same gap months earlier, and names which parts of the standard it sits in.

Three points in the journey where a readiness signal changes something.

The same instrument, doing the same thing, at three moments where a certifying body currently has no visibility and the candidate currently has no diagnosis.

01 Before a first sitting

A candidate, their sponsor and their training provider can all see whether the knowledge is there before a seat is booked. A cohort's weak areas are visible while there is still term left to act on them, rather than after everyone has dispersed.

02 At the challenge and equivalency route

An experienced or internationally trained worker gets an honest baseline against the published standard before committing to an assessment or a sitting. Assessment can begin while a worker is still abroad, and the result is a better-evidenced application rather than a coin flip.

03 After a failed attempt

This is the group with the sharpest need and the thinnest resources. They have finished — on-the-job hours complete, in-school training done, the training provider relationship usually ended. They are no longer anybody's student. What they have is a result with no diagnosis attached, a thirty-day wait before they can sit again, and a \$150 fee when they do. Those thirty days are currently the least used period in the system, because a result does not say which parts of the standard were weak — so the strategy available is to revise everything, or whatever felt hardest.

Certifying bodies see this too: one provincial exam support unit reached **5,852 candidates in 2024**, with a **26% pass rate** on the next attempt. A skill-by-skill diagnosis in the first few sessions is what would let those thirty days be spent on what was actually missing.

Both figures: a provincial certifying body's annual report, 2024–25. The prior year's report gives 32% for the trade qualifier subset — a different year and population. Source available on request.

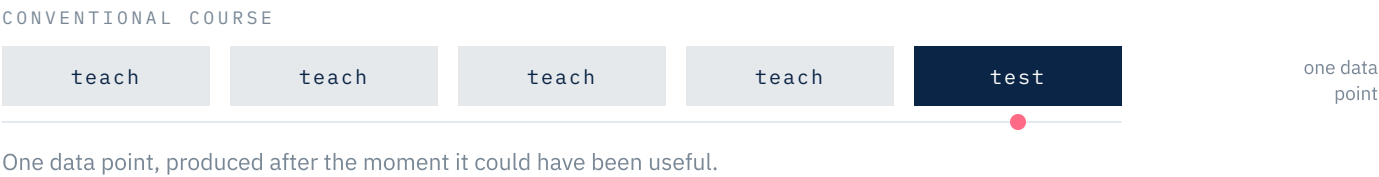
A result tells a candidate they are not ready. It does not tell them what for. That is a diagnosis problem before it is a study problem — and it is the same problem at all three moments above.

Measuring usually costs learning time. Here it does not.

In conventional training, assessment is a cost — time spent finding out whether learning happened rather than causing it. Because it is a cost it happens rarely, so the picture is coarse and late. One finding removes that trade-off.

Actively retrieving something strengthens it far more than re-reading it. Roediger and Karpicke's 2006 demonstration is the canonical one: re-reading feels more productive and produces markedly worse long-term retention than being tested. It replicates across ages, materials and delays, and is one of the least contested findings in the field. So questions here are the primary interaction rather than an assessment bolted onto content — every answer is simultaneously a measurement and a learning event.

Every question is both practice and evidence. There is no trade-off to manage: more measurement is more learning.



What that rests on, named precisely

Learning science is easy to misuse, so we hold to a rule: say what we implement, not what we admire.

EFFECT	HOW IT SHOWS UP HERE	EVIDENCE
Retrieval practice	Questions, flashcards and practice exams are the primary interaction, not an add-on.	STRONG
Mastery learning	Per-skill evidence, with a Proficient bar earned rather than reached by finishing — mastery measurement more than mastery learning.	DIRECTIONAL
Dual coding	Lesson video alongside the questions. The same learner is 5.8 points more accurate on concepts where they watched at least half the video first.	MEASURED
Spacing	Review returns on a fixed schedule rather than one reacting to performance — a simplification, and we label it as one.	SIMPLIFIED

How we would prove this in a single trade.

Nothing in this document is a proven trades outcome, and the only way that changes is a designed trial. This is the shape we would propose. If the design is wrong, we would rather run the one a certifying body specifies.

One trade, defined first

A single trade and cohort agreed before anything begins, with the standard, population and period fixed in writing.

The success measure set by the certifying body

Defined by the certifying body, not by us. If it is first-attempt pass rate, it is first-attempt pass rate.

A comparison group

Without one, any change is unattributable. This is the condition we would insist on whether or not it were asked of us.

Published either way

Methodology and result, whichever direction it goes. A result that can only be reported when it is flattering is not evidence.

What a trial would need

- A published training or curriculum standard for the trade, broken down to individual skills.
- Exam weighting if it exists; if not, we proceed without it and say so.
- A route to confirmed outcomes — the one thing that turns a calibration claim from an argument into a measurement.

The step before any of that

None of that is a first step. The useful next step is smaller: a working session with the standards and certification teams, and a live demonstration. An hour is enough to see how a published standard is mapped, what the measurement looks like at skill level, and where the readiness score comes from — and a team that has seen it working can judge it better than one that has read about it.