The Living Skills Framework

How to keep your people's skills current when the work will not hold still. A working guide for L&D, talent, and change leaders in financial services, with an audit you can run this week.

PARTS	THE AUDIT	READ TIME
Three, plus the audit	Twelve questions	Fifteen minutes

INTUITION KNOWLEDGE NETWORK

Intuition Knowledge Network
Free from IKN. No registration required.

WHY WE WROTE THIS

A method, not more awareness.

Most skills frameworks in financial services are reviewed once a year. The jobs they describe now change every quarter. A new regulatory expectation lands, a system gets replaced, an AI tool is rolled out to the front office, and the framework that was accurate in January is describing a job that no longer quite exists by June.

The people who own that framework know it. What they are short on is not awareness. It is a method: a way to decide what to change and what to leave alone, a way to prove the framework is producing capability rather than completions, and a way to keep it honest across two very different populations, the people who have been in the building for fifteen years and the people who arrived last month with a different set of skills and a different set of gaps.

This guide is that method. It comes out of three questions we keep hearing from learning leaders at banks, insurers, asset managers, and capital markets firms, and it takes a position on each.

We have spent four decades building capability inside financial institutions, and the one thing that has not changed in that time is the shape of this problem. The tools change. The question of whether people can actually do the work does not.

WHAT IS INSIDE

PART ONE

The annual review is the problem.

A skills framework is a description of what a role requires. The trouble is that it is usually written as if the role were stable, and then reviewed on a calendar that has nothing to do with when the role actually changes. Annual review cycles made sense when the main driver of change was a new product line every few years. They do not survive a world where a compliance team’s daily tools change twice in one budget cycle.

The result is a familiar gap. Your framework says one thing, your managers know another, and your learning programs are built to the framework. People complete the programs and then get told on the desk that the real job is different.

> **That is not a content problem.** It is a trigger problem. Nobody decided when the framework should change, so it changes when someone gets around to it.

Three triggers that should force a review

Replace the calendar with triggers. When any of these happens, the affected part of the framework gets reviewed inside the quarter, not at year end.

01

A tool changes how the work is done

Any system rollout, AI tool, or platform migration that touches a role rewrites part of that role. Review the affected skills before the rollout, not after adoption stalls.

02

A regulator changes what the work has to prove

New examination priorities, a model bulletin, a fair lending focus. Each one turns a background skill into a front-line one for somebody. Compliance already tracks these.

03

The reviewers start correcting the same thing

When managers, second-line reviewers, or quality teams keep sending back the same category of error, that is a skill gap announcing itself.

Ask reviewers quarterly what they corrected most. It is the cheapest skills data in the building, and the trigger most firms miss because it lives in people’s heads.

THE TWO-LAYER SPLIT

Adapt, do not rebuild.

The instinct when a framework feels stale is to redo it. Resist that. A full rebuild takes a year, and by the time it ships it is stale again. The firms that keep frameworks current treat them as two layers.

ANCHOR SKILLS

Judgment under uncertainty, client trust, regulatory reasoning, the ability to explain a decision to someone who will challenge it.

These have defined the role for a decade and will define it for the next one. They change slowly and should be written to last.

EDGE SKILLS

This platform, this reporting pack, this year's examination focus.

Tied to a specific tool, process, or moment. Rewritten on the cycle that changes them, not on the annual calendar. Give every edge skill a review date.

> **Separate them once** and the framework becomes something you maintain, not something you replace. A platform change stops invalidating the whole document.

What this changes in practice

- **The review gets smaller.** You are reviewing the edge layer against a trigger, not rewriting a document.
- **The anchors stay quotable.** Managers and reviewers can name the four or five skills that actually define the role.
- **The expiry does the work.** An edge skill with a review date cannot quietly go stale, because the date arrives whether anyone is watching or not.

PART TWO

Completions are not capability.

Every skill in the framework should be measurable at four levels, and each level should have a named owner outside L&D. The rung matters. The owner matters more.

RUNG	WHAT IT MEASURES	EXAMPLE	WHO OWNS IT
1. Activity	**Did the learning happen**	Completion, attendance, time in program	L&D
2. Capability	**Can the person now do the thing, tested**	Assessment score, practice task passed, baseline to progress delta	L&D with the line manager
3. Behavior	**Are they doing it on the job**	Reviewer corrections down, escalations down, tool used inside the sanctioned workflow	The line manager
4. Outcome	**Did the business change**	Time to desk-ready, error rate, examination findings, client work returned less often	The business owner

Rung one is where most reporting stops. Rung two is where any serious assessment-based program lives: a baseline before, a measured result after, per role. Rung three is where the manager comes in, and rung four is where the executive question finally gets an answer.

> **You will not have all four rungs for every skill.** You do not need to. Pick the five skills the business cares about most this year, get all four rungs for those, and report on them relentlessly. Five skills measured to the outcome are worth more in a budget conversation than fifty measured to attendance.

CONNECTING LEARNING DATA TO BUSINESS DATA

Borrow the number, gain an ally.

The reason rung four is hard is that the numbers live in different systems owned by different people. Learning data is in the LMS. Error rates are in operations. Examination findings are in compliance. Time to desk-ready is in whoever runs the graduate program. Nobody is lying to anyone. The data simply never meets.

The practical move is to agree the outcome metric before the program starts, in writing, with the person who already owns that number. Not a new metric L&D invents. Their existing one. If the head of operations already tracks reconciliation errors, that is the rung-four measure for the reconciliation skills, and the head of operations is now a co-owner of the learning result. You have not built a data pipeline. You have borrowed one, and you have made an ally.

Shared accountability, made specific

Everyone agrees skill development should not sit only with HR and L&D. Almost nobody makes that specific. Here is the version that works: accountability follows the rung.

01

L&D owns activity, and co-owns capability

Whether the learning happened, and whether the skill was acquired, tested.

02

Managers own behavior

Because behavior is what they see, and nobody else can.

03

Business owners own outcomes

Because outcomes are what they are already measured on.

Write those three names next to every skill that matters and the accountability conversation is over before it starts.

> **The manager rung fails silently.** A manager who does not know what good looks like cannot own behavior change. Training managers to recognize the skill is often the highest-return program in the whole framework, and the one least likely to be in it.

PART THREE

The gap nobody plans for.

TENURED

Knows the products, the clients, the regulators, and the unwritten rules.

Widening gaps in the tools and methods that arrived after they learned the job.

NEW

Fluent in the tools.

Short on everything else: the products, the judgment, the sense of what a regulator will ask.

One framework, written to one standard, serves neither. It either asks tenured staff to sit through foundations they mastered years ago, which they will quietly refuse, or it assumes new hires have context they do not have, which shows up later as a client error. Most firms respond by building two parallel programs, which doubles the maintenance and splits the culture.

Same framework, two on-ramps

Keep one framework. Change how people enter it. The anchor and edge layers do most of the work here, because the two populations have opposite gaps: tenured staff are strong on anchors and weak on edges, new hires are the reverse.

- **Assess before assigning.** A short baseline per skill, before any program is assigned, routes each person to what they actually lack. It also produces the rung-two data as a by-product.
- **Let tenured staff test out.** Know it, show it. If someone can demonstrate an anchor skill, they skip the program and the framework records the evidence. Respect for expertise is what keeps experienced people engaged with the framework at all.
- **Pair the populations on purpose.** The tenured employee who needs the new tool and the new hire who needs the product context are each other’s fastest teacher. Structured pairing on real work is cheaper than any program.

SO YOU STOP REACTING

Most frameworks are behind, and the trigger was a complaint.

The last question learning leaders ask about their framework is whether it is ahead of the curve or behind it. Leading indicators are the fix. They are not exotic. They are already in the building.

EXTERNAL

- Regulator examination priorities for the coming year
- Vendor roadmaps for the platforms your people use
- What your competitors are hiring for

Each is a preview of a skill demand twelve months out.

INTERNAL

- Help desk and support tickets by topic
- The correction patterns from part one
- Escalation volumes by team
- Which tutorials or resources people search for on their own

Self-directed searching is the purest demand signal you have, because nobody assigned it.

Review these quarterly against the edge layer. When an external signal and an internal signal point at the same skill, that is the review trigger firing early, which is the entire point.

THE AUDIT

Twelve questions, one afternoon.

Score each question 0, 1, or 2. Zero means no, or not really. One means partly, or informally. Two means yes, and it is written down. Total each part. The lowest-scoring part is where to start, and the specific questions scoring zero are your first three actions.

Part one: is the framework current?

Is there a defined trigger, other than the calendar, that forces a review of the framework?	0 1 2
Are anchor skills and edge skills separated, so a tool change does not invalidate the whole document?	0 1 2
Does every edge skill have a review date or an expiry?	0 1 2
Do you ask reviewers and managers, at least quarterly, what they corrected most?	0 1 2

PART TOTAL

Part two: can you prove it worked?

For the five skills the business cares about most, do you measure beyond completion?	0 1 2
Is there a baseline before and a measured result after, per role?	0 1 2
Is the business outcome metric agreed in writing with the person who already owns it?	0 1 2
Does each priority skill have a named owner at the behavior rung and at the outcome rung, outside L&D?	0 1 2

PART TOTAL

Part three: does it serve both populations?

Are people assessed before programs are assigned, rather than assigned by role alone?	0 1 2
Can experienced staff demonstrate a skill and skip the program?	0 1 2
Is there deliberate pairing between tenured staff and new hires on real work?	0 1 2
Do you review at least two leading indicators, one external and one internal, against the framework each quarter?	0 1 2

PART TOTAL

Reading your score

0 to 8

The framework is a document, not a system.

Start with part one. Separate anchors from edges this month; it costs nothing and everything else depends on it.

9 to 16

The framework is maintained but not defended.

Start with part two. Pick five skills, agree one outcome metric each with its existing owner, and report on those before the next budget round.

17 to 24

You are ahead of most of the industry.

Start with whichever part scored lowest, and consider whether your leading indicators are actually leading or just early lagging ones.

UNDERNEATH ALL THREE PARTS

Two things a framework cannot supply on its own.

MANAGEABLE

None of this works if it lands on a two-person team as a year of extra work.

The anchor and edge split, the five-skill rule, and the borrowed outcome metric are all there to make the framework something a small team can keep current. If a step in this guide would take you a quarter, cut it to the version that takes a week.

DEFENSIBLE

You will be asked what changed.

The ladder exists so that you can answer with a number the person asking already trusts, because it is their number. A framework you can defend in a budget conversation is a framework that survives one.

Skills frameworks fail quietly, by drifting out of date while everyone is busy. They succeed the same way, by being maintained in small, regular, evidenced steps. This guide is the smallest set of those steps we know works.

WHERE THIS GOES NEXT

FREE FROM IKN

The first guide in IKN's series on modern skills building.

If you want the assess, path, and practice loop from part three running inside your firm rather than on paper, that is what Intuition Know-How was built for: a baseline per role, a path to the gap, practice on real financial services work, and the progress data the ladder needs. We'd welcome the conversation.

ASSESS	PATH	PRACTICE
A baseline per role	**Straight to the gap**	**On real financial services work**

INTUITION KNOWLEDGE NETWORK

intuitionknowledgenetwork.com

Join the community for the rest of the series.