



L&D in Healthcare

A Playbook for Better Outcomes

Ravinder Tulsiani, DBA

© 2026 Dr. Ravinder Tulsiani. All rights reserved.

Frameworks referenced (CLEAR, Domino Map, Performance Blueprint, IMPACT, AI Value Spectrum) are the author's.

First edition: October 2026.

Published by Catalyst Learning & Strategy. rtulsiani.com

Contact: ravindertulsiani.com

Worked examples are illustrative unless marked as from the author's experience.

Contents

Executive Summary	4
Chapter 1: Diagnose the Function Before You Fix It	6
Chapter 2: Start From the Outcome, Not the Request	12
Chapter 3: Rebuild the Function Around Outcomes	18
Chapter 4: Prove It in the Board's Language	24
Chapter 5: The First Twelve Months	30
Chapter 6: Where AI Earns Its Place	35
How the Frameworks Connect	39
About the Author	40

Executive Summary

Your organization spends serious money on learning. Mandatory modules, onboarding, competency checks, leadership programs, a learning management system, and the hours clinicians spend away from patients to complete it all. What you cannot get, in most healthcare organizations, is a straight answer to a simple question: what did any of it produce?

That gap is the subject of this report. It describes what a learning function built around outcomes looks like, how a chief learning officer would build it in the first twelve months, and what the executive team should expect to see along the way.

The method has four movements. First, diagnose the function honestly: find out what the learning investment actually buys today, where the waste sits, and which problems are not training problems at all. Second, start every significant investment from the outcome it is supposed to move, with a falsifiable claim about the performance it will change, so that money follows evidence rather than requests. Third, rebuild the function around those outcomes, so the org chart serves the aims instead of the other way around. Fourth, prove the contribution in the language the board already uses: regulatory assurance, workforce performance, business contribution with its limitations stated, and investment decisions with validated benefits.

Three convictions run underneath. Solve the right problem: much of what passes for a training need is a management, process, or tooling problem wearing a training costume, and the most valuable thing a learning leader does is refuse to train the wrong thing. Treat compliance as an obligation and capability as the investment: mandatory education is legitimate, and meeting it well can build real capability, but it is not the same as building a workforce. And make evidence the price of trust: a function that states its aims, measures its contribution honestly, and reports its limitations earns the authority to make harder decisions.

Two later chapters complete the picture. One sequences the first twelve months, quarter by quarter, so the board can see what progress looks like and when to expect it. The other places artificial intelligence on a value spectrum: adopt it where it earns its place, choose simpler solutions where they win, and refuse it where risk outweighs benefit.

What this means for you

For the chief executive. You get a learning function accountable for workforce performance rather than learning activity, a budget attached to named aims with evidence behind each, and a first year that produces visible momentum without disrupting required operations. The question this report answers for you is not whether learning matters. It is how you will know whether yours is working.

For the chief nursing executive and clinical leaders. You get a partner who understands that your people's time is the scarcest resource in the building. Clinical leaders own patient outcomes and approve clinical content through their own professional authorities. The function brings design and evaluation expertise to your problems. It does not bring you more mandatory modules.

For the board. You get a function you can govern. A quarterly account in four parts: regulatory assurance that obligations are met, workforce performance that shows competence changing, business contribution with evidence and limitations stated, and investment decisions with validated benefits. The board's attention belongs on the material risks, the investment decisions, and the assurance questions. Operational detail stays with management.

Diagnose the Function Before You Fix It

The CLEAR Diagnosis Model (from Beyond Courses)

EXECUTIVE VIEW

The issue. Healthcare organizations invest substantial staff time in learning activities, but completing training does not establish that workplace performance has improved.

The executive decision. Before approving new programs, new technology, or a restructuring of the function, determine which learning activities are required, which ones improve performance, and which ones need further evidence.

What the diagnosis must establish. What the organization needs the function to accomplish. Where resources are currently going. Whether learning is being applied in the work. Which activities warrant protection, improvement, or discontinuation.

What management should bring forward. A verified inventory of learning obligations, a baseline of workforce performance, and prioritized recommendations with named decision owners.

A healthcare organization can report excellent training completion rates and still be unable to answer a more important question: did the learning change how people do their work? Before changing the learning function, the executive team needs a diagnosis. What is working? What is consuming staff time without sufficient evidence of benefit? What capabilities are missing?

That is where CLEAR begins. Clarify success. Locate current performance. Explore causes. Assess transfer. Recommend interventions. The purpose is to make better decisions about workforce performance, using evidence from this organization rather than assumptions about what training should accomplish.

A note on how to read this chapter. What follows describes patterns I have seen repeatedly across healthcare organizations. They are diagnostic hypotheses worth testing in your organization, not universal verdicts. The CLEAR process exists precisely because the diagnosis must be built from local evidence. The patterns tell you where to look; the evidence tells you what is true here.

Clarify success: what would "working" actually look like?

Forget the course catalog for a moment. If a healthcare organization's learning function were working properly, what would you see?

You would see new nurses reaching full competence on a predictable timeline, not whenever the unit happens to have a quiet week. You would see mandatory education completed without pulling people off the floor during a staffing crisis. You would see training that shows up in behavior on the unit, not just in a learning management system report. And you would see the function answering to the same outcomes the board cares about: safer care, retained staff, audit-ready compliance, and a workforce

that can absorb change instead of drowning in it.

That is the standard. Hold it in mind, because everything that follows is measured against it.

Locate current performance: what healthcare L&D actually does today

Start with what the function is really doing, not what its strategy deck says it does.

In my experience, L&D activity in Canadian healthcare organizations falls into three buckets. The first and largest is mandatory education: the modules required by regulation and internal policy. In Ontario, the College of Nurses of Ontario does not set a fixed number of continuing-education hours the way many US states do. Instead, every nurse in the General or Extended Class is accountable under legislation to reflect on their practice each year, identify learning needs, and complete a learning plan with at least two learning goals. That is the professional obligation. On top of it sit employer-mandated modules: privacy, health and safety, infection control, workplace violence prevention, and whatever the latest incident or audit added. Accreditation Canada's Required Operational Practices have long included patient-safety education and training as a required practice, alongside hand-hygiene education and training, and the current Qmentum program carries "patient safety education and training" forward as a required safety practice. Each mandate is reasonable on its own. Stacked together, they become a second job.

The second bucket is clinical and role-based education: onboarding, competency checks, skills refreshers, leadership development. This is the work that actually builds capability. In most organizations it is smaller than the mandatory bucket, funded last, and measured least.

The third bucket is everything else that lands on L&D's desk because nobody else wants it: system rollouts, policy change communications, the occasional morale initiative dressed up as training.

Now add the structural reality. In many organizations, learning is split between clinical education (nursing professional practice, credentialing) and corporate training (HR, IT, compliance), running on separate systems with separate teams. A nurse can complete the same infection-control content in two different portals in the same month and nobody sees the duplication, because nobody has a single view of the learner's journey. Shift work makes classroom training impractical for large parts of the workforce, so the default answer is more e-learning modules, which deepens the mandatory bucket without touching the capability bucket.

And this is not only a hospital story. In home care, personal support workers and visiting nurses learn alone, supervised intermittently, with training that must reach people who rarely gather in one place. In long-term care, provincial legislation sets specific staff training and retraining obligations that the home must meet regardless of what else is on the calendar. The dispersed workforce is where the gap between assigned training and applied practice is widest, and it is where the diagnosis matters most.

Then there is the staffing context, which is not background noise but the central fact. Ontario needs an estimated 26,000 additional registered nurses just to match the RN-to-population ratio of the rest of Canada, a comparison the Registered Nurses' Association of Ontario reported in July 2024 using CIHI data. Ontario has held the country's lowest RN-to-population ratio for nine consecutive years.

Nationally, Statistics Canada reports the RN job vacancy rate at 7.0 percent in 2024, up from 2.6 percent in 2016. CIHI counted roughly 28,000 vacant positions for registered nurses and registered psychiatric nurses across 2022-23, nearly a quarter of all healthcare vacancies in the country. Personal support workers accounted for the largest share of all, at 30,800 vacant positions.

Note what these numbers do and do not say. The nursing workforce is not shrinking in absolute terms; CIHI reports the regulated nursing supply has grown. The problem is that demand is outpacing supply: in 2022 there were 52 nurses in direct care roles per 1,000 older adults in Canada, down from 59 in 2013. Read those numbers as an L&D leader and the picture is stark: the function is spending most of its energy on mandatory modules while the workforce it is supposed to develop is stretched past what the old training models were built for.

Explore causes: how it got this way

The causes are structural, not personal. Blaming the L&D team misses the point; the team is responding rationally to the incentives around it.

First, mandate layering. Every new rule, incident, or audit finding produces a new course, and old courses almost never get retired. Nobody owns the decommissioning decision, so the catalog grows in one direction. A US health system, Mercy Health, offers a useful illustration of the scale of the waste: working with a vendor's adaptive platform over a two-month rollout to more than 50,000 learners in late 2024, it cut compliance training time from roughly 75,000 hours to roughly 26,000 hours, reporting more than 49,000 hours saved and over \$2 million in cost savings. Those are the vendor's own reported figures, not an independent audit, but the shape of the finding is credible and familiar: a large share of mandatory training time was adding nothing.

Second, the reporting line. L&D often sits inside HR, accountable for completions and costs rather than for clinical outcomes. When your performance is measured in completion rates, you produce completions. The function has no structural reason to ask whether the training changed anything on the unit, because nobody above it is asking either.

Third, the training-request reflex. When something goes wrong on a unit, the organizational reflex is to assign training. A medication error becomes a mandatory module. A patient complaint becomes a communication course. L&D, lacking the standing to push back, converts every operational problem into a course. This keeps the function busy and keeps it irrelevant to the problems that actually matter.

Fourth, cost-center positioning. L&D is typically funded as overhead, which in my experience makes it first in line for cuts and last in line for investment. The irony is visible in the staffing numbers above: organizations cannot hire their way out of the shortage, and the function that could shorten time-to-competency, improve retention through development, and reduce the burden of mandatory training is being starved.

Assess transfer: is any of this reaching the floor?

This is the step most diagnoses skip, and it is the most important one. Transfer asks a simple question: after the training, did behavior change?

For mandatory modules, the honest answer in most organizations is that nobody knows, because nobody measures it. Completion is recorded; behavior is not. And an important discipline applies here: if a unit's fall rate stays unchanged after mandatory education, that does not by itself establish that the education failed. The patient population may have changed, acuity may have risen, or other factors may be masking its contribution. CLEAR does not infer failure from flat trend lines. It investigates whether staff actually applied the required practice, what got in the way, and what else was affecting the outcome. The question is never "did the numbers move." It is "did the behavior change, and if not, why."

For clinical education, transfer is better where it is structured: preceptorships, competency checklists, simulation with feedback. It is worse where it is not: the one-off in-service, the emailed policy update, the annual skills fair that everyone attends and nobody remembers. The pattern is consistent. Learning has a better chance of changing performance when employees can practise it, receive feedback, and apply it under the conditions in which they work.

The accreditation environment is beginning to force this question. Accreditation Canada's Qmentum program has been moving in the same direction, streamlining 38 Required Operational Practices into 24 Required Safety Practices with an explicit emphasis on measurable safety improvement. "Patient safety education and training" remains a required practice. The message from accreditors is consistent: training that cannot show a result is losing its standing.

Recommend interventions: the verdict, not a course list

A proper diagnosis ends with recommendations that fit the evidence. Different findings justify different decisions, and the CLO's job is to match the intervention to what the evidence actually shows.

The organizing idea is this: compliance is a tax, and capability is an investment. The organization must satisfy every legitimate obligation and then reduce unnecessary administrative burden to the minimum. Capability investment should be directed where the return is highest and judged on the improvement it can reasonably contribute to the outcomes the organization exists to produce.

One qualification matters, because a CNO will spot the overstatement immediately. The two are not always separate businesses. A well-designed mandatory infection-prevention program is both a compliance obligation and a genuine capability investment. Some essential education addresses rare but severe events: nobody retires code-blue retraining because the event is infrequent. The target was never legitimate mandatory education. It is the unnecessary burden layered on top of it: the duplicated modules, the content nobody can tie to a requirement or an outcome, the annual re-certifications that assess nothing.

With that distinction in place, the interventions follow from the diagnosis:

- **Inventory and review.** Catalog every mandatory module: who requires it, what evidence of learning need supports it, and whether the current delivery is still appropriate. Retire what cannot be justified. In long-term care, provincial legislation sets specific training and retraining obligations, so the review documents the requirement rather than assuming it can be cut.

- **Unify the learner view.** Clinical and corporate education stop running parallel systems that cannot see each other. One record, one learner, no duplicate completions.
- **Move the accountability.** L&D's performance measures get rewritten around outcomes the clinical side cares about: time to full competence for new hires, completion of high-risk competencies, retention of staff in their first two years. Completion rates stay as an operational metric, not the headline.
- **Install the "do not train" verdict.** Some findings of the diagnosis will not be training problems. If errors persist after training, the cause may be the process, the staffing model, or the tool, and the recommendation is to fix that instead. An L&D function that cannot say "do not train" will keep being handed problems training cannot solve.

A worked example: how the verdict changes with the evidence

WORKED DECISION EXAMPLE

The following is a worked decision example, not a case study.

An Ontario home-care organization faces a recurring problem: inconsistent care documentation by personal support workers, and delayed escalation when a client's condition deteriorates. The standing request from operations is to retrain the PSWs. The CLO runs CLEAR instead of scheduling the course.

The transfer check is harder here than in a hospital, because the workplace is dispersed. Supervisors cannot observe every home visit. So the diagnosis uses what is available: documentation audits, a sample of ride-along observations, and interviews with PSWs and their supervisors. The questions are direct. Do staff understand the documentation requirement? Can they access the correct information during a home visit, on the tools they actually carry? And when they do escalate, what happens next?

Two different verdicts are possible, and the evidence decides between them. If the investigation finds that PSWs know the requirement but cannot reach the documentation system from the field, or that escalations go into a queue nobody monitors, the recommendation is "do not train": fix the tool and the escalation pathway first. If it finds that staff genuinely do not recognize the early signs of deterioration, the intervention is targeted coaching on recognition and escalation, delivered by supervisors in the field, not another e-learning module assigned to everyone.

That is the discipline. The same framework, applied honestly, produces opposite recommendations depending on what the evidence shows. A function that can do this earns something no course catalog can buy: the trust of the operational leaders whose problems it is supposed to solve.

Decisions for the executive committee

The chapter ends where the executive reader lives: with decisions. But note what these decisions are. The chapter has argued that CLEAR must precede conclusions, so the executive committee is not being asked to approve a verdict. It is being asked to authorize the work that makes a defensible verdict possible. A CLO bringing this diagnosis to the executive committee would put forward the

following, each with a named owner:

- **Approve a review of training obligations and a controlled retirement process.** L&D is authorized to examine every mandatory module against its requirement and its evidence, and to retire what cannot be justified, with clinical sign-off from the Chief Nursing Officer or designate on any item touching clinical practice. Decision owner: Chief Nursing Officer, with L&D executing.
- **Require an assessment of the learning systems landscape.** Before any consolidation, an assessment of duplicate records, interoperability options, and the cost of available solutions. Decision owner: Chief Information Officer, with L&D defining requirements.
- **Approve the development of an outcome-linked scorecard.** Not a rewrite of measures on day one, but a mandate to build one: time-to-competence, high-risk competency coverage, and early-tenure retention, with clinical and operational co-owners, reported quarterly. Decision owner: Chief Executive Officer, on the CLO's recommendation.
- **Authorize targeted verification of high-risk competencies.** A small, recurring investment in checking whether required practices are actually performed in the workplace, starting where the risk is highest. Decision owner: Finance, validating the business case L&D prepares.

None of these requires a restructuring or a new platform on day one. Each is a decision the executive committee can make in a single meeting, and together they authorize the diagnosis the rest of this report depends on.

Endnotes

- College of Nurses of Ontario, Quality Assurance program: nurses in the General and Extended Classes complete an annual self-assessment and learning plan with at least two learning goals under the Regulated Health Professions Act, 1991. No fixed CE-hour requirement. cno.org/quality-assurance.
- Accreditation Canada, Qmentum Global program: 38 Required Operational Practices streamlined into 24 Required Safety Practices, with emphasis on measurable safety improvement; "Patient Safety Education and Training" is a required safety practice (Leadership standard). accreditation.ca. International comparison: the US Joint Commission's "Accreditation 360" (announced June 30, 2025) removed 714 additional prescriptive requirements and merged the remainder into 14 National Performance Goals effective January 1, 2026. jointcommission.org.
- RNAO media release, July 25, 2024 (CIHI data): Ontario needs an estimated 26,000 additional RNs to match the RN-to-population ratio of the rest of Canada; Ontario has held the lowest ratio for nine consecutive years.
- Statistics Canada, "Regional dynamics of vacancies in health-related occupations in Canada, 2024" (national figures): RN job vacancy rate 7.0% in 2024, up from 2.6% in 2016; licensed practical nurses 12.8%; personal support workers 6.7%.
- CIHI, "The state of the health workforce in Canada, 2022" (Feb 2024): average 120,140 healthcare vacancies per quarter across 2022-23; registered nurses and registered psychiatric nurses 28,000 (23.3%); personal support workers 30,800 (25.7%). In 2022 there were 52 nurses in direct care roles per 1,000 older adults, down from 59 in 2013.
- Amplifire case study, Mercy Health (vendor-reported, not independently audited): system-wide rollout to 50,000+ learners, November 2024 to January 2025; training hours cut from 75,065 to 25,984; "over 49,000 hours of training time saved, translating to more than \$2 million in cost savings." amplifire.com.
- On longer-range projections: a 2018 forecasting analysis projected a Canadian nursing shortfall of 117,600 by 2030; it was cited by the House of Commons Standing Committee on Health in its March 2023 report. Treat as dated context, not a current forecast.

Start From the Outcome, Not the Request

The Domino Map (from the Domino Map framework)

EXECUTIVE VIEW

The issue. Most learning investments begin with a training request. Requests describe solutions, not problems, so the function builds courses for outcomes nobody defined.

The executive decision. Require every significant learning investment to trace backward from a named patient, resident, or client outcome to the performance that outcome requires, and from that performance to the capability gap worth closing.

What the method must establish. Which outcome is being aimed at. What frontline performance that outcome depends on. Where the capability gap actually sits. The smallest intervention likely to close it.

What management should bring forward. A short list of priority outcomes, each with its performance chain mapped and its proposed intervention sized to the gap.

Chapter 1 ended with a diagnosis. This chapter decides what matters.

The most expensive habit in healthcare L&D is starting with the request. A unit asks for a course on communication. A director asks for mandatory training after an incident. A new policy needs "education." Each request arrives pre-solved: the answer is training, and the only remaining question is what the training should cover. The Domino Map exists to break that habit. It forces the conversation to start at the other end: from the outcome the organization needs, backward through the performance that outcome requires, to the capability gap that is actually worth closing. Only then does it ask what the intervention should be.

The map has seven steps:

- **Aim:** choose the outcome.
- **Define:** describe the performance the outcome requires.
- **Map:** trace backward to the capability gap.
- **Performance Blueprint:** design the rebuilt function.
- **Solution Design / Minimum Effective Dose:** size the intervention to the gap.
- **Measurement:** check whether the prediction held.
- **Scaling:** extend what works.

This chapter walks through the first three steps and Solution Design. The Performance Blueprint, the rebuilt function itself, is the subject of Chapter 3. Measurement gets its full treatment in Chapter 4, under the IMPACT framework.

Aim: choose the outcome first

Everything begins with a named outcome, chosen with the people who own it. Not "improve communication." Not "more training on falls." A specific, measurable outcome that a clinical leader would recognize as theirs: fewer pressure injuries in long-term care. Faster sepsis recognition in the emergency department. Fewer avoidable readmissions in home care.

The aim must be small enough to act on and important enough to matter. An aim like "improve patient safety" is a value, not a target; nobody can map a performance chain to it. An aim like "reduce the rate of stage 2 and greater pressure injuries acquired in our homes" can be mapped and measured, which means someone can own it. The discipline of the aim is the discipline of the whole method. Vague aims produce vague training.

This is also where the CLO earns a different seat. Clinical leaders own patient outcomes; that accountability does not transfer to L&D. What L&D brings is accountability for its contribution: the capability the outcome depends on, built deliberately and measured honestly. When L&D walks into a quality meeting able to show exactly which part of the outcome its work moves, it is no longer the department that delivers courses. It is a partner in the clinical agenda.

Define: describe the performance the outcome requires

Once the outcome is named, the next question is purely operational: what would frontline staff have to do, consistently, for this outcome to improve?

Take pressure injuries in long-term care. The performance that prevents them is not mysterious: personal support workers observing and promptly reporting skin changes, nurses assessing and acting on those reports the same day, repositioning on schedule, and nutrition and hydration monitoring.

This step is deliberately boring. It resists the urge to talk about "awareness" or "culture" and insists on observable behavior. If you cannot describe what a PSW does differently on a Tuesday night when the outcome is improving, you do not yet understand the performance. And if you do not understand the performance, any training you design is a guess.

The define step also surfaces something the training-request reflex hides: most outcomes depend on performance by more than one role, coordinated across handoffs. Pressure injury prevention fails at the handoff between the PSW who notices reddened skin and the nurse who must assess and act. A surprising amount of what looks like a knowledge gap turns out to be a coordination gap, and coordination gaps are rarely fixed by courses.

Map: trace backward to the capability gap

Now the map gets drawn. Starting from the outcome, trace backward: outcome, to required performance, to the capabilities that performance depends on, to the gaps in those capabilities today.

For the pressure injury aim, the chain might run: fewer pressure injuries requires consistent repositioning and early identification of skin changes; that requires PSWs who can recognize and promptly report those changes, and nurses who assess and respond the same day; the gap might be that PSWs were never shown what early-stage skin breakdown looks like on darker skin tones, or that

the escalation pathway exists on paper but nobody follows it because the response never comes.

Notice what the map does to the original request. "We need pressure-injury training for PSWs" becomes a set of precise, testable questions: Is the gap knowledge, or is it the escalation pathway? Is it skill, or is it that night-shift staffing makes the repositioning schedule physically impossible? The map does not assume the answer is training. It locates where training could help and, just as importantly, where it could not. A gap caused by an unstaffed night shift is a resourcing decision, not a curriculum.

This is also where the method connects to evidence the clinical side already trusts. CIHI's analysis of hospital harm through the pandemic found that harms linked to staffing levels rose: urinary tract infections and pneumonia by about 20 percent, aspiration pneumonia by 25 percent, and pressure ulcers by over 50 percent compared with the pre-pandemic period. The finding is about staffing, not training, and the map respects that distinction. Its job is to separate the factors a learning function can move from the ones it cannot, so the function stops being blamed for outcomes it was never equipped to change.

Solution design: the minimum effective dose

Only now, with the outcome aimed at, the performance defined, and the gap mapped, does the method ask what the intervention should be. And the rule is the minimum effective dose: the smallest intervention likely to close the gap reliably.

The minimum effective dose is a discipline against the two most common errors. The first is the enterprise rollout for a local problem: a full curriculum, assigned to everyone, when the gap sits with one role on one shift. The second is the module for a non-training problem: a course assigned when the map shows the gap is the escalation pathway, the staffing model, or the tool.

For the pressure injury example, suppose the map shows PSWs missing early skin changes and a slow escalation response. The minimum effective dose might be a short, hands-on coaching session on skin assessment, delivered by the home's own clinical educators during shift overlap, paired with a one-page escalation pathway posted where PSWs actually work and a commitment from nursing leadership on response times. Not a four-hour e-learning module. Not a system-wide campaign. The smallest thing that closes the gap, delivered where the work happens.

The minimum effective dose also protects the budget conversation. When every intervention is sized to its gap, the function can show exactly what each dollar is buying and why a smaller spend would not have worked. That is a very different conversation from defending a training budget in the abstract.

Measurement and scaling: the handoff

The map's final steps are measurement and scaling, and they belong to the next two chapters, with one distinction worth making now because the two are easy to confuse. The Domino Map's Measurement step asks a narrow question: did this intervention move this aim? IMPACT, in Chapter 4, is the function-level system: the ongoing evidence operation that proves the function's contribution across aims, in language the board uses. The map measures the bet; IMPACT measures the

portfolio.

Chapter 3 builds the function that can execute at scale. The point to carry forward is this: the Domino Map makes a falsifiable claim. It says that closing this gap, with this intervention, will move this outcome. Measurement then checks whether it did. A method that cannot be proven wrong cannot be trusted when it claims to be right.

A worked example: the map on one page

WORKED DECISION EXAMPLE

The following is a worked decision example, not a case study.

An Ontario long-term care organization sets its aim: reduce the rate of pressure injuries acquired in its homes. The quality team and the incoming CLO define the required performance: nursing skin assessments on admission and weekly thereafter, PSW observation and prompt reporting of changes between assessments, repositioning on the schedule the care plan specifies, and same-day nursing assessment when a PSW flags a concern.

The map traces backward. Ongoing observation requires PSWs who can recognize early changes and report them where nurses will see them; the formal assessments remain nursing accountability. Repositioning requires enough staff on nights to follow the schedule. Escalation requires a pathway that produces a same-day response.

The investigation finds three distinct gaps. PSWs report they were shown what to watch for once, at orientation, years ago, and are unsure what early breakdown looks like. The documentation system requires the PSW to leave the resident's room to log the observation, so observations get batched and delayed. And night-shift staffing on two units cannot cover the repositioning schedule as written.

Three gaps, three different answers. The knowledge gap gets the minimum effective dose: brief refresher coaching on recognizing and reporting early skin changes, delivered in the home, with practice on real cases. The documentation gap gets "do not train": move the logging to the point of care. The staffing gap gets escalated to operations as a resourcing decision, with the map as evidence that no training intervention can close it.

One aim. One map. Three verdicts.

WORKED DECISION EXAMPLE · PRESSURE INJURIES IN LONG-TERM CARE (HYPOTHETICAL)

1 · AIM

- Reduce the rate of pressure injuries acquired in our homes.



2 · DEFINE: THE REQUIRED PERFORMANCE

- Skin assessments on admission, then weekly.
- PSW observation and prompt reporting between assessments.
- Repositioning on the schedule the care plan specifies.
- Same-day nursing assessment when a concern is flagged.



3 · MAP: THREE GAPS FOUND

- Knowledge: PSWs unsure what early breakdown looks like.
- System: logging means leaving the room, so reports are batched and delayed.
- Resourcing: night staffing on two units cannot cover repositioning.



4 · MINIMUM EFFECTIVE DOSE: THREE VERDICTS

- Coach: brief refresher in the home, practice on real cases.
- Do not train: move logging to the point of care.
- Escalate: the staffing gap goes to operations as a resourcing decision.



5 · MEASURE

- Pressure injury rate: did the intervention move the aim?

The core insight: the map is a budget document

Here is the idea that organizes this chapter. The Domino Map looks like a learning design tool. Its real value is narrower and more useful: every map ends with a short list, these outcomes, these gaps, these interventions, sized to the evidence. That list lets executives judge learning investments the way they judge any other investment: on the return.

This is also why the aim matters more than the intervention. A modest intervention aimed at a pressure injury rate the board already watches will always beat a beautifully designed program aimed at a vague goal. Start where the accountability already lives.

What this means in year one

For the executive reader: in the first year, the CLO maps a small number of aims, not the whole organization. Two or three outcomes, chosen with clinical leadership, mapped carefully, intervened on minimally, measured honestly. The point of year one is not coverage. It is proof that the method works on real outcomes, in this organization, with evidence the board can see. Coverage comes in year two, built on that proof.

Endnotes

- Canadian Institute for Health Information, "Hospital staffing and hospital harm trends throughout the pandemic": harms linked to staffing levels rose compared with the pre-pandemic period, including urinary tract infections and pneumonia by about 20 percent, aspiration pneumonia by 25 percent, and pressure ulcers by over 50 percent (the last partly attributed to improved data capture). [cihi.ca](https://www.cihi.ca). Supporting US evidence: Aiken LH et al., JAMA 2002;288(16):1987-1993 (each additional patient over four per nurse associated with a 7 percent increase in 30-day mortality following common surgeries).

Rebuild the Function Around Outcomes

The Performance Blueprint (the Domino Map's blueprint step)

EXECUTIVE VIEW

The issue. A learning function cannot reliably improve workforce performance when its authority and operating practices are designed primarily around training delivery.

The executive decision. Determine what organizational structure, decision rights, and professional capabilities are required to deliver the priorities established through diagnosis.

What the blueprint must establish. Who owns each significant decision, how clinical authority is protected, what work belongs within L&D, and how performance evidence informs continued investment.

What management should bring forward. An operating design supported by the diagnosis, with clear accountabilities, implementation dependencies, and an approved approach to managing the transition.

Chapters 1 and 2 diagnosed the function and decided what matters. This chapter rebuilds the function itself. In the Domino Map, this is the Performance Blueprint step: the design of the organization that will execute everything the map prescribes.

A warning before the design. Structures do not fix functions; accountability does. An L&D department can be reorganized and re-staffed, and still produce the same course catalog if nothing changes about what it answers for. The blueprint below is drawn around a single principle: every element of the function, its reporting line, its roles, its governance, and its daily operating model, must answer to outcomes, not to activity.

The reporting line: decision authority matters more than the box on the chart

The most consequential design decision is also the most easily oversimplified: who the head of learning reports to.

It is tempting to prescribe a single answer. In practice, the reporting line matters less than the decision authority it confers. A strong chief human resources officer can be an effective executive sponsor of enterprise capability. A weak arrangement fails wherever it sits. What the CLO actually needs is three things: direct access to enterprise priorities, a defined relationship with clinical leadership, and accountability for the capability contribution the function can reasonably control. Whether the position reports to the CHRO, the COO, or the CEO should depend on the organization's structure and strategic priorities, not on a universal rule.

That said, the current default deserves scrutiny. When L&D reports through human resources and is measured the way HR is measured, completions, costs, compliance, it behaves accordingly. If the

function is being rebuilt around workforce performance, its accountability has to change first, and the reporting relationship has to support that accountability. The test is simple: can the head of learning tell the chief executive, in the executive's language, what the function moved last quarter? If not, the arrangement is wrong, wherever the box sits.

For the CNO reading this: what protects your interests is not where L&D reports. It is the defined clinical relationship, the dotted line with real authority behind it, that gives nursing leadership sign-off over nursing education and a shared say in how the function serves the clinical agenda.

Structure: designed for every setting, arranged to fit the organization

Healthcare organizations do not operate in one setting. A regional system runs hospitals, long-term care homes, and home-care services, each with different workforces, different rhythms, and different regulatory obligations. The learning function must serve all of them without becoming separate functions that cannot see each other.

The blueprint calls for dedicated expertise for each major setting the organization serves: acute care, long-term care, and home and community care. In a large system that may mean dedicated leads. In a smaller or more focused organization it may mean one senior practitioner covering two settings, or contracted expertise for the setting the organization touches lightly. The form follows the organization's shape. What does not change is the requirement: whoever designs learning for a setting must understand its work, its workforce, and its constraints. The home-care lead knows what it means to train people who work alone in clients' homes. The long-term care lead knows the provincial training obligations that cannot be negotiated away.

Under the setting expertise sits a shared core: learning design, learning technology, and measurement. These are centralized where it makes sense because they are scarce skills, and because shared capability is what makes reliable enterprise-wide learning information possible.

The executive decision here is the outcome, not the technology. The organization needs a trustworthy enterprise-wide view of learning and competence, and consistent accountability for it. Whether that is achieved through one consolidated platform or through connected systems depends on the organization's legal entities, existing technology commitments, and clinical governance arrangements. Prescribe the outcome; let the diagnosis determine the arrangement.

Roles: fewer course builders, more performance partners

The roles follow the work. The blueprint replaces the traditional training-department staffing, coordinators who schedule courses and administrators who track completions, with roles designed around the Domino Map.

Performance consultants are the pivotal role. Each is embedded with an operational area, a clinical program or a home-care region, for example, and their job is to run the map's opening steps with anyone who requests training, before anything gets designed. When a director requests training, the performance consultant does not take the order. They run the diagnosis. They are the people who say "do not train" when the evidence says it, and they need the standing, and the CLO's backing, to survive saying it.

Clinical education specialists carry the clinical credibility the function needs. They are regulated professionals who design and deliver the clinical curriculum and who hold the relationship with professional practice leadership. Clinical approval sits with the appropriate professional authority: nursing leadership for nursing matters, and each regulated profession's own authority for its scope. The learning function contributes design and evaluation expertise. It does not hold a veto over clinical judgment, and clinical leaders retain their professional accountability.

Learning technologists own the systems: the learning platform, its integration with HR and clinical systems, and the data plumbing that makes measurement possible. They are not IT support. They are the people who ensure the technology serves the operating model instead of dictating it.

A measurement analyst owns the evidence. This role builds the dashboards, runs the transfer checks from Chapter 1, and prepares the quarterly account of what the function moved. Without this role, measurement is everyone's third priority, which means it is nobody's.

For the CNO, this role design is the answer to the most common complaint about L&D: that it sends generic training designed by people who do not understand the work. Every clinical initiative in this blueprint is designed with clinical educators, delivered with clinical oversight, and measured against clinical outcomes.

Governance: a council with real decision rights

Structure without governance drifts. The blueprint establishes a learning council chaired by the CLO, with members drawn from the organization's actual accountabilities: nursing and professional practice leadership, quality and patient safety, human resources, information technology, and finance, with meaningful representation from regional operations in a dispersed organization like home care. The council does four things: it approves the annual learning agenda against the organization's priority outcomes; it reviews the quarterly evidence of what the function moved; it adjudicates conflicts between mandatory obligations and capacity; and it exercises authority to continue, redesign, or discontinue programs.

That last power matters most, and it has a boundary worth stating plainly. A council cannot simply retire what the law requires. Ontario's Fixing Long-Term Care Act, 2021 and its regulation set specific staff training and retraining obligations, including annual retraining intervals, with limited flexibility for individual training-needs assessment. Discontinuation authority applies to discretionary programming and to the delivery method of required programming, not to the legal obligation itself. Every retirement decision is subject to regulatory obligations and to sign-off from the relevant clinical authority.

The featured case: governance that could say no

I have seen what disciplined governance looks like from the inside. As chair of an education technology committee with a proactive mandate to identify and evaluate new learning tools, I ran a structured evaluation of a third-party engagement platform requested by two departments. The committee put it through its established filters, ran a risk assessment with security and technology colleagues involved, and approved only a limited trial on two computers for planned events. After the trial, the departments found the tool clunkier than expected: it added no value and created frustration. It was never rolled out broadly. It was canceled.

That is a governance success story, and it is worth stating plainly why. The request was received, evaluation criteria were applied, risk was assessed, a limited trial was authorized, and the evidence led to cancellation. The process said no, contained the exposure to two computers, and produced a clean decision with evidence behind it. Most governance bodies cannot do this. They approve pilots that never end, tools that nobody evaluates, and spending that nobody revisits. The authority to discontinue, exercised on evidence and within regulatory bounds, is what separates governance from ceremony, and it belongs in the council's charter from day one.

For the board: this is your assurance mechanism. A council with discontinuation authority, reporting quarterly on what was redesigned or retired as well as what was launched, is how you know the function is being managed rather than merely funded.

The operating model: how work actually flows

With the structure and governance in place, the operating model describes the daily flow of work.

Every significant request for learning enters through the Domino Map. The performance consultant runs the map's opening steps with the requestor and the relevant clinical lead. The output is not a course design. It is a short document: the outcome, the performance chain, the gap, the proposed intervention at its minimum effective dose, and the measure that will tell whether it worked. The learning council reviews anything above a defined threshold of cost or scale. Below the threshold, the CLO approves.

Delivery follows the gap, not the catalog. Some interventions are courses. Many are not: coaching, job aids, workflow changes, supervisor toolkits, peer mentoring. The function maintains a small catalog of proven interventions and a design practice for new ones, and it treats the catalog as inventory to be managed, not assets to be protected.

Technology serves this model. The learning platform is chosen and configured to support the single learner record, point-of-care access for dispersed workers, and the data feeds the measurement analyst needs. Platform decisions go through the same governance as program decisions, with trials scoped the way the engagement-platform trial was scoped: limited, evaluated against pre-set criteria, and killed cleanly when they fail.

FROM THE AUTHOR'S EXPERIENCE

A real example, from my own record. At a healthcare organization where I led the learning function, we were paying over \$500,000 annually for a learning platform that was not meeting our needs. We moved to a new vendor, and after the first year the annual cost was \$101,000. That is an actual, contractual, recurring annual saving from switching vendors, not a one-time benefit and not a projection. The lesson for the blueprint: platform economics belong in the technology decision alongside functionality. These are real, recurring savings. Whether an organization reinvests them in capability or returns them to the bottom line is a separate decision, but it should be made with the full number visible.

A worked example: the blueprint in a home-care organization

WORKED DECISION EXAMPLE

The following is a worked decision example, not a case study.

An Ontario home-care organization with several thousand personal support workers and visiting nurses decides to rebuild its learning function. The current state is familiar: a small corporate training team administering mandatory modules, clinical education scattered across nursing managers, no single view of who has completed what, and supervisors who describe training as something that happens to their staff rather than something that helps them.

The blueprint starts with decision authority, not boxes. The CLO's reporting relationship is settled against the criteria above: direct access to enterprise priorities, a defined dotted line to the chief nursing executive with real sign-off authority, and accountability for the capability contribution. The first capability added is setting expertise: a home-care learning lead, a former visiting nurse who knows what supervision looks like when the workplace is a client's living room. Performance consultants are aligned to service regions. A measurement analyst begins the work of building reliable enterprise-wide learning information out of disconnected systems.

The learning council's early decisions follow the sequence in Chapter 1: the mandate inventory, the first Domino-mapped aims, and a controlled review of modules that cannot be tied to a requirement or an outcome, with the recovered effort redirected to the coaching the map prescribed. Nothing is retired by fiat; every discontinuation goes through the council with regulatory and clinical checks.

In a smaller organization, the same blueprint would look different: one learning lead covering two settings, performance consulting contracted for the diagnostic work, the council meeting less often with fewer seats. The accountabilities do not change. The shape does. That is the point of designing from criteria rather than from a fixed chart.

For the CEO, what changes is accountability: one executive, one reliable record, one quarterly account of what moved. For the chief nursing executive, what changes is that clinical education now has design discipline and nursing sign-off behind it. For the board, what changes is that the learning budget is traceable to outcomes for the first time.

What this means in year one

For the executive reader: the rebuild does not happen all at once, and it must not disrupt what is legally required while it is being built. In year one, the CLO establishes the reporting relationship and the council, fills the pivotal roles the diagnosis justifies, and runs the new operating model on the first two or three mapped aims. Required education and clinical services continue uninterrupted. Restructuring and system consolidation follow confirmed needs and approved business cases, not the other way around.

Headcount follows the work. A better-designed function may need fewer course administrators and more specialized capability, particularly in a dispersed organization where performance consulting and measurement are new disciplines. The staffing plan is built from the aims, not from a predetermined reduction.

Year one ends with the new governance tested: the council has met, the first aims are mapped and measured, and it has demonstrated the ability to challenge a proposed intervention and reach a defensible decision, on evidence and within regulatory bounds. A council that cannot do that is a decoration.

Endnotes

- Role titles vary by organization; the functions they describe do not. "Performance consultant" is used here for the embedded diagnostic role regardless of what a given organization calls it.
- Ontario's Fixing Long-Term Care Act, 2021, s. 82, and O. Reg. 246/22, ss. 259-260: licensees must ensure staff training in prescribed areas, with retraining at annual intervals; individual training-needs assessment may narrow what a given staff member must receive. ontario.ca/laws.

CHAPTER 4

Prove It in the Board's Language

The IMPACT Framework (from Proving Impact)

EXECUTIVE VIEW

The issue. Learning functions typically report activity, completions, hours, satisfaction scores, while the board makes decisions on outcomes. The two conversations never meet, so learning budgets are defended with anecdotes and cut with spreadsheets.

The executive decision. Require the learning function to produce evidence of its contribution in the measures the organization already uses to judge itself.

What you get. A CEO gets a learning budget that can be judged on returns. A CFO gets the contribution stated in financial terms, with the assumptions visible. A board gets a quarterly account it can actually read.

What management should bring forward. An outcome-linked scorecard for the function's major investments, each tied to a named business problem, with baseline measures and a reporting cadence.

Chapter 2 made a promise: every Domino Map ends with a falsifiable claim, that closing this gap with this intervention will move this outcome. This chapter keeps that promise. It uses the IMPACT framework:

- **Identify** the Business Problem
- **Map** the Capability Pathway
- **Pick** the Measures That Matter
- **Assemble** Credible Evidence
- **Communicate** Contribution Clearly
- **Track, Test, and Refine**

A distinction from Chapter 2, stated once so it holds for the rest of the report. The Domino Map's Measurement step tests whether the map's specific prediction held. IMPACT is the ongoing function-level evidence system: it can establish credible evidence for an individual initiative as well as for the portfolio, in the language the board uses.

Identify the business problem: start where the money already goes

Every significant learning investment gets tied to a business problem the executive team already recognizes. Not "nurses need more training on documentation." The business problem: new nurses take too long to reach full productivity, and too many leave in the first year, in a province that needs 26,000 more registered nurses to match the rest of the country. That is a problem the CEO, the CNO,

and the CFO all feel, and it is stated in their terms before learning is mentioned at all.

The discipline here is a defensible purpose. Every significant learning investment needs one. For mandatory education, the organization must establish the requirement and determine how to meet it effectively; the evidence question is whether the requirement is being met, not whether the program earns a financial return. For discretionary investment, management should state the performance problem being addressed and the evidence that would justify the expenditure. That distinction would withstand scrutiny from any Canadian healthcare board, because it respects what the law requires while refusing to let everything else escape accountability.

For the CFO: this is where the function's credibility begins. A learning leader who opens with the business problem, stated in operational and financial terms, is someone you can have a budget conversation with. A learning leader who opens with the program is someone you manage as overhead.

Map the capability pathway: show the causal chain

For each business problem, the function maps the pathway from its work to the outcome. Training does not directly reduce turnover. The pathway runs: structured onboarding and preceptorship shorten time to full competence; competence reached sooner means less moral distress in the early months; supported early months mean fewer first-year departures; fewer departures mean lower agency spend and safer staffing.

The pathway must be explicit enough to argue with. That is its value. If the chief nursing executive believes the real driver of first-year departures is scheduling rather than onboarding, that disagreement happens now, on paper, before money is spent, instead of later, when the results disappoint and everyone has a different theory. A pathway that cannot survive a skeptical clinical leader is not ready for funding.

Note what the pathway does not claim. It does not claim that learning alone moves the outcome. Each link in the chain is a hypothesis to be tested, not an established fact: the intervention is expected to change something directly, subsequent outcomes may be influenced, and other operating factors, staffing levels, compensation, leadership quality, workload, need to be named alongside it. The pathway states what evidence would support or challenge the proposed contribution. And no financial benefit is presented to the board until finance has validated the calculation. Lower turnover does not automatically produce a cash saving; vacancies may persist, agency contracts may continue, retained capacity may be used differently. The ROI figure reaches the CFO only with its assumptions visible and finance's sign-off behind it. That discipline is exactly what makes IMPACT valuable.

Pick the measures that matter: the board's metrics, not the LMS metrics

The measures are chosen from the organization's existing scorecard, not invented for learning. Retention and turnover by tenure cohort. Time to full competence for new hires in critical roles. Incident and harm rates for the aims in Chapter 2. Compliance audit results. Cost per proficient employee, which translates the function's efficiency into a number finance can compare year over

year.

What gets left out of the impact story matters as much as what gets in. Hours delivered and satisfaction scores are operational metrics. They tell the function whether its machinery is running. They tell the board nothing about whether the machinery matters. They stay in internal reporting.

Completion rates deserve a finer distinction. Completion is inadequate as proof of impact, but a board may legitimately need assurance that required education and high-risk competency obligations are being met. That assurance belongs on the board's page, in concise form, as regulatory assurance, not as evidence of contribution. The board's reporting exhibit therefore has four areas: regulatory assurance (mandatory obligations met, outstanding risks, corrective action); workforce performance (demonstrated competence and operational performance change); business contribution (outcome trends, evidence of contribution, and its limitations); and investment decisions (expenditure, validated benefits, recommended next action). That is what the CLO produces, quarterly, on one page.

Each measure needs three properties: a baseline taken before the intervention, an owner who is not the learning function alone, and a reporting cadence the executive team already follows. A measure without a baseline is a story. A measure owned only by L&D is a claim nobody else will defend. A measure reported on a separate cycle from everything else is a ritual.

Assemble credible evidence: pilots before proclamations

Evidence is assembled the way the clinical side assembles it: comparatively and with a baseline. The strongest design available to a learning function is the pilot with a comparison group. Redesign onboarding in two units, hold two comparable units on the current program, and compare time-to-competence and early retention at six and twelve months. It is not a randomized trial, and it does not need to be. It needs to be honest about what it can and cannot show.

Two cautions keep the honesty real. First, the design must be proportionate to the decision. A two-unit pilot can justify continuing the pilot; it rarely justifies enterprise-wide investment on its own, particularly for a noisy outcome like retention, where staffing changes, acuity shifts, or a single manager departure can swamp the signal. The chapter's standard is staged: pilot evidence earns continuation, stronger evidence earns scale. Second, where pilots are not practical, the function uses before-and-after measurement with careful attention to confounders, the same discipline Chapter 1 applied to the falls-prevention example. If the population changed, if acuity rose, if staffing shifted, the evidence says so. Credibility compounds: one honest report with caveats buys more trust than three glossy ones.

FROM THE AUTHOR'S EXPERIENCE

A real example, from my own record. At a healthcare organization where I led the learning function, we rebuilt onboarding for clinical staff, nurses, behavioral supports workers, and rehab staff, around a simple change: strip out the frontloading, deliver the essentials first so new hires reach productivity sooner, and move the remaining training into ongoing professional development over time using a drip-feed approach. Onboarding got shorter, new hires became productive faster, and turnover roughly halved over the following two to three years. The honest version includes the caveat: compensation structures and contracts changed in the same period, so the redesign was a contributor, not the sole cause. New hires also reported feeling less overwhelmed, which mattered for early retention. That is what credible evidence looks like. It claims a contribution, names the confounders, and lets the reader judge.

FROM THE AUTHOR'S EXPERIENCE

Another example from my own experience shows a different measurement point. A VR-based dementia program was built for an emotional gap rather than a knowledge gap: long-term care staff understood dementia intellectually but could not relate to it. Thirty staff went through an immersive experience replicating the hallucinations and sensory disruption of dementia. Every participant self-reported increased empathy immediately afterward. At sixty days, with managers involved in the follow-up, 90% reported they were applying what the training taught them in real situations. Self-reported, and labeled as such. But it was measured at application rather than reaction, which is where most evaluation stops.

Skepticism about learning's claims is healthy, and this section is built for it. The evidence comes with its limitations stated, its comparison groups named, and its assumptions visible. That is what makes it usable in a budget decision.

Communicate contribution clearly: translation, not decoration

The evidence then gets translated. Not into a learning report, but into the organization's language. The format below is illustrative: it shows how a CLO should structure the communication, with real numbers filled in from the organization's own evidence, never invented.

ILLUSTRATIVE EXHIBIT

Illustrative executive reporting format. Business problem: delayed readiness for independent nursing practice. Intervention: revised onboarding and supervised practice. Evidence: demonstrated proficiency and retention compared against agreed baselines, with comparison groups where feasible, limitations stated. Financial assessment: finance validates costs and any savings attributable to the intervention before the figures reach the board. Decision: continue, adjust, or expand according to the strength of the evidence.

Every number in the real version has a source, a baseline, and a named owner. The CFO can check the arithmetic. The CNO can check the competence claim. The board can see the return.

What this is not: a dashboard of completions, a testimonial slide, or a claim that the program "drove engagement." Executives fund outcomes. The communication states the outcome, the contribution, and the cost, and stops.

ILLUSTRATIVE BOARD REPORT

Hypothetical: quarterly, one page

REGULATORY ASSURANCE

- All mandatory education current.
- Exception: fire-safety recertification in two LTC homes, now scheduled; interim risk noted.

WORKFORCE PERFORMANCE

- Time to independent practice for new RNs fell against baseline in the two pilot units.
- Comparison units: unchanged.

BUSINESS CONTRIBUTION

- First-year departures lower in pilot units than in comparison units.
- Confounder noted: manager change in one comparison unit.
- Finance validated the avoided agency-cost arithmetic.

INVESTMENT DECISION

- Continue the redesign.
- Extend the evaluation to two further units before any wider rollout.

Operational detail stays with management and the council.

The board's attention belongs on the material risks, the investment decisions, and the assurance questions. The operational detail behind each line stays with management and the council.

Track, Test, and Refine: decisions after the first results

The final step turns measurement from an event into an operation. Each aim's evidence is reviewed on the cadence the council set in Chapter 3. Interventions that moved their aim get scaled, following the Domino Map's Scaling step. Interventions that did not get redesigned or retired, through the council's discontinuation authority. New aims enter the map as old ones graduate.

Over time, this produces the portfolio view: a short list of active investments, each with its business problem, its pathway, its measures, its evidence to date, and its status. That list is the function's budget submission, and it gives the board a direct line from dollars to outcomes.

A worked example: measuring an onboarding redesign

WORKED DECISION EXAMPLE

The following is a worked decision example, not a case study. It shows how decisions change with the evidence; no outcomes are assumed.

An Ontario hospital network redesigns nurse onboarding in response to the business problem: too many first-year departures in a market where every vacancy takes months to fill. The capability pathway is mapped with nursing leadership as a hypothesis: structured preceptorship and competency milestones shorten time to full competence; earlier competence reduces early-career distress; reduced distress lowers first-year departures. The pathway names its rival explanations from the start, including scheduling and unit leadership.

The measures are picked from the existing scorecard: time to independent practice for new RNs, first-year retention by cohort, and agency hours on the pilot units. Baselines are taken from prior cohorts. Two units pilot the redesign; two comparable units continue the current program.

At twelve months, the council faces the evidence and three possible verdicts. If proficiency improved but retention did not, the question is whether onboarding was ever the retention problem, and the next step is to re-examine the diagnosis, not to expand the program. If neither improved, the question is whether the redesign was actually implemented as designed, and whether the original map identified the right gap. If both improved against baselines and comparison units, with confounders stated, the redesign earns continuation and a larger pilot, not immediate enterprise rollout: the evidence supports the next test, not the final one.

One element, a mentoring component, shows no measurable difference in the pilot. It is not dropped on that evidence alone; a small pilot cannot carry that conclusion. It continues into the next cycle with a sharper measure, and the decision to keep or remove it waits for sufficient evidence. That restraint is the point. A function that kills programs on thin evidence will soon be afraid to try anything, and a function that scales on thin evidence will soon be defending failures.

What this means in year one

For the executive reader: in year one, IMPACT runs on the two or three aims mapped in Chapter 2, not on the whole function. The scorecard is built, the baselines are taken, the first pilots report. A basic portfolio report can start in year one, even though credible outcome evidence takes longer to mature. What the board sees at the end of the first operating cycle is straightforward: which problems the function is addressing, what resources have been committed, what the available evidence shows, and which decisions management recommends next. The purpose is not to claim credit for every improvement. It is to make the function's contribution visible and credible in resource decisions.

Endnotes

- Baselines drawn from at least two prior cohorts or periods, where available, help distinguish the intervention's effect from normal variation. Treat this as a working standard rather than a universal requirement; the rigor of the design should be proportionate to the decision it supports.

The First Twelve Months

Plain practitioner guidance: no framework brand

EXECUTIVE VIEW

The issue. New executives often spend their first year producing strategy documents while the function keeps running on old habits. By the time the strategy launches, credibility is spent.

The executive decision. Sequence the first year so that diagnosis comes before design, proof comes before scale, and every quarter produces something the executive team can see.

What you get. A CEO gets visible momentum without disruption to required operations. A CNO gets early evidence that the new leader understands the clinical reality. A board gets a year that ends with a measured account, not a slide deck.

What management should bring forward. A quarter-by-quarter plan with named deliverables, decision points for the executive committee, and the evidence each quarter is expected to produce.

The previous four chapters established the approach: diagnose with CLEAR, use the Domino Map to connect outcomes to performance, design the function through its Performance Blueprint step, and evaluate contribution through IMPACT. This chapter puts that approach in calendar order. It describes what a chief learning officer actually does in the first twelve months, in what sequence, and what the organization should expect to see.

One rule governs the whole year: nothing in this sequence disrupts legally required education or clinical services while it is being built. The new model proves itself on a small number of aims before it touches anything else. That constraint is not caution for its own sake. It is how a new executive earns the right to change things.

Months 1 to 3: diagnose and build trust

The first quarter belongs to the diagnosis in Chapter 1 and to relationships. The CLO meets the people who will determine whether any of this works: the chief nursing executive, the quality and patient safety leaders, the HR and IT executives, the finance partner who will eventually validate the business cases, and a sample of frontline managers and educators who understand where the operational problems persist.

In parallel, the mandate inventory begins. Every mandatory module is catalogued against its requirement and its evidence, on a risk-prioritized basis; a comprehensive inventory in ninety days may not be realistic in a large organization, and the review starts where the risk and the likely waste are highest. This produces the year's first visible progress: an initial performance baseline, opportunities for improvement supported by evidence, and early changes where the necessary approvals are already in place. Early movement matters less for its financial value than for what it signals. An executive who returns hours to the floor in the first ninety days, on evidence and with the

proper sign-offs, has earned attention for everything that follows.

The quarter ends with two deliverables for the executive committee. First, the CLEAR diagnosis: what the function does today, where the waste sits, where the gaps are, and which findings are "do not train." Second, the proposed learning council: its charter, its membership, and its first meeting date. The committee is not asked to approve a restructuring. It is asked to authorize the diagnosis to continue into decisions.

For the CNO: this is the quarter where you learn whether the new leader understands your world. A CLO who spends it listening to your educators and your frontline managers, rather than presenting a vision, is one worth backing.

Months 4 to 6: map the first aims and pilot the new model

With the diagnosis in hand, the second quarter runs the first Domino Maps. Two or three aims, chosen with clinical leadership, mapped carefully through the map's opening steps. The performance consultants, if hired, run this work. If the roles are not yet filled, the CLO runs it directly with the relevant clinical leads, which is slower but has the advantage of putting the new leader shoulder to shoulder with the operation from the start.

The first minimum-effective-dose interventions launch as pilots. One aim gets its intervention; the measurement for that aim begins, with baselines taken before launch. The learning council meets for the first time and takes its first decisions: the mandate review findings, the pilot designs, and, if the evidence supports it, the first controlled retirements.

This is also the quarter where the governance gets its first real test. Sooner or later, usually sooner, a request arrives that the map says should not be training: a new mandatory module after an incident, a system-wide rollout for a local problem, a vendor pilot with no evaluation criteria. What matters is not that the function says no on schedule, but that it demonstrates the ability to challenge a proposed intervention and reach a defensible decision, with evidence, with clinical partnership, and without making an enemy. A council that approves a necessary investment on strong evidence is functioning just as well as one that redirects an unnecessary one. The council's backing is what makes a well-reasoned challenge survivable, and how the CLO handles the first of these sets the tone for the function's authority for years.

For the CEO: this is the quarter where the method becomes visible. You should expect to see the first maps, the first pilots, and the first documented decision to challenge or redirect a proposed intervention. That decision is the leading indicator that the rebuild is real.

Months 7 to 9: judge the pilots, build the scorecard

The third quarter produces the first real evidence cycle. By now, management should have enough implementation evidence to determine which interventions can continue, which need adjustment, and which may be ready for controlled expansion. Decisions that depend on longer-term clinical or workforce outcomes stay open until the evidence is sufficient; an apparent improvement in a short pilot does not by itself justify scaling. This is the Track, Test, and Refine discipline from Chapter 4 operating on its first cycle, and the organization watches closely, because this is where most

initiatives quietly stop being evaluated.

The IMPACT scorecard goes live this quarter. Baselines from the first pilots, the four reporting areas from Chapter 4, regulatory assurance, workforce performance, business contribution, investment decisions, reported on the cadence the council set. It will be thin in places; a scorecard with honest gaps is more credible than a complete one built on soft numbers.

The operating model from Chapter 3 extends beyond the pilot aims. New requests now enter through the map as a matter of routine. The council exercises its authority, to continue, redesign, or discontinue, on evidence and within regulatory bounds. Each exercise of that authority makes the next one easier.

A word on resources, because the plan so far assumes the CLO can staff it. Canadian healthcare executives often face recruitment restrictions, collective agreements, and competing demands on clinical staff time. The first year must be achievable without assuming new headcount or technology expenditure. Where performance consultants cannot be hired, the CLO redefines part of an existing role, shares analytical capacity with the quality team, or leads the initial diagnosis personally. The second quarter already allows for the CLO doing the mapping directly; extend that pragmatism across the year. The method does not require the ideal organization. It requires the discipline, applied with whatever capacity exists.

For the board: the third quarter is when the learning budget starts to look different. The same dollars, but now attached to named aims, with early evidence and a visible method for deciding what continues.

Months 10 to 12: consolidate and report

The final quarter consolidates. The practices that proved themselves get written into how the organization works: the map becomes the standard intake for significant learning requests, the scorecard becomes part of the quarterly rhythm, the council's calendar is set for the coming year. Practices that did not prove themselves are either retired or scheduled for further evaluation, with the decisions recorded and the reasoning preserved, so the organization does not repeat the experiment out of forgetfulness.

The year ends with the account Chapter 4 designed: which problems the function addressed, what resources were committed, what the evidence shows, and which decisions management recommends next. It is presented to the executive committee in detail and to the board, or the appropriate board committee, in summary, focused on material risks and investment decisions requiring its attention. It does not claim transformation. It claims a foundation: a diagnosed function, a working method, early evidence, and a credible plan for year two.

Year two is sketched, not promised in detail: more aims mapped, the scorecard deepened, the technology consolidation assessed through a proper business case. The sketch matters because it shows the first year was a beginning, not a project with an end date.

A worked example: one year in a home-care organization

WORKED DECISION EXAMPLE

The following is a worked decision example, not a case study. It follows the organization from Chapters 1 and 3. No outcomes are assumed; the example shows how the executive response changes with the evidence.

The home-care CLO's first quarter produces the risk-prioritized mandate review and an early finding to test: personal support workers appear to be completing similar privacy and safety content in two different systems. The finding goes to the council with its evidence and its required approvals, not as a promised saving. The council authorizes the investigation, and the CLO begins the documentation aim from Chapter 1 by mapping it with the clinical leads.

In the second quarter, the documentation map is built. The investigation tests the Chapter 1 hypothesis: part knowledge gap, part tooling gap. A coaching pilot launches in one region; the documentation system's field access is referred to IT as a non-training intervention. When operations requests a mandatory module after a documentation audit, the council demonstrates the governance working: the map shows the gap is supervisory follow-up rather than knowledge, and the council reaches a defensible decision to redirect the request to the regional managers with a follow-up protocol, recording its reasoning.

The third quarter brings the first evidence cycle, and the response depends on what it shows. If documentation competence improves with credible supporting evidence, the intervention continues and controlled expansion is considered. If workers understand the procedure but cannot use the system effectively in the field, the system problem goes back to the accountable operational and technology leaders. If documentation improves but client-care outcomes remain unchanged, the question is whether the expected relationship between the two was valid. If the evidence is inconclusive, the evaluation is extended or revised before any wider commitment. This is the chapter's central discipline: the calendar sets when evidence is reviewed, but the evidence determines what is decided.

The year ends with the account: what was learned, what was changed, what evidence supports those changes, and what management is prepared to do next. Nothing in it is dramatic. Everything in it is defensible.

What the board should expect to see

For the executive reader, the year's value can be stated as expectations, one per quarter, with the calendar treated as illustrative. The evidence required for a decision determines when that decision can be made.

ILLUSTRATIVE EXHIBIT

Illustrative executive exhibit: what management should demonstrate during year one. Months 1 to 3: initial diagnosis, regulatory obligations, priority performance gaps; the executive decision is to approve priorities and scope of change. Months 4 to 6: mapped outcomes, intervention proposals, baseline measures; the executive decision is to authorize appropriate trials and resources. Months 7 to 9: early implementation findings and expenditure, with evidence limitations stated; the executive decision is to continue, adjust, or expand where justified. Months 10 to 12: validated results to date, unresolved questions, and next-year priorities; the executive decision is to confirm ongoing investment and accountabilities.

In plain terms: after the first quarter, a diagnosis you can read and a council with a charter. After the second, the first mapped aims, the first pilots, and the first documented decision to challenge or redirect a proposed intervention. After the third, a live scorecard and the first evidence-based continuation decisions. After the fourth, a measured account and a year-two plan.

If any quarter produces none of these, the rebuild has stalled, and the executive committee should ask why. That checkability is deliberate. A plan the board cannot monitor is not a plan; it is a hope.

By the end of year one, management should be able to show what it learned, what it changed, what evidence supports those changes, and what it is prepared to do next. That is a defensible standard. It demonstrates operational leadership without promising outcomes that depend on circumstances the CLO does not fully control.

Endnotes

- Timelines compress or extend with organizational size and complexity. A single-site organization may run this sequence faster; a multi-province system will need longer for the technology and governance work. The order of operations matters more than the calendar.

Where AI Earns Its Place

The AI Value Spectrum

EXECUTIVE VIEW

The issue. Healthcare executives face two opposite pressures on AI: vendors and boards pushing rapid adoption, clinicians and regulators demanding caution. Learning functions are being asked to "AI-enable" their programs without a clear account of where the technology actually helps.

The executive decision. Require every AI proposal in the learning function to pass through the same value test: where AI adds real value, where a simpler solution wins, and where risk outweighs benefit.

What you get. A CEO gets an AI posture that is neither hype nor paralysis. A board gets a governance story it can defend. Clinical leaders get assurance that each proposal faces the authorities responsible for the risks it creates before anything is piloted.

What management should bring forward. A short register of AI proposals and pilots in the learning function, each placed on the spectrum with its risk assessment, its clinical approvals, and its kill criteria.

This chapter is deliberately short. AI is one instrument in the function's toolkit, not its purpose. The report has held it back until now for a reason: an organization that has not yet learned to diagnose its gaps and govern its interventions has no business deploying AI at scale. It will automate its confusion. The chapters before this one are the precondition. This chapter is the discipline applied to the newest tool.

The discipline is the AI Value Spectrum: every proposed use of AI is placed on a spectrum with three zones. In the first zone, AI adds real value. In the second, a simpler solution wins. In the third, risk outweighs benefit. The CLO's job is to keep proposals honest about which zone they are in, and to move them between zones as evidence accumulates.

Where AI may add real value

AI earns its place where the work is high-volume, pattern-rich, and currently bottlenecked by human capacity. In a learning function, the clearest candidates sit behind the scenes. Drafting scenario-based practice cases from clinical protocols, so educators spend their time refining judgment calls instead of writing boilerplate. Generating practice variations for competency assessment, so a nurse preparing for a skills validation faces fresh scenarios rather than the same three. Summarizing evaluation data across cohorts, so the function sees patterns in transfer failures that no one person could read across hundreds of responses. Translating mandatory content into the languages a diverse workforce actually speaks, with clinical terminology checked by a qualified reviewer.

These are legitimate uses to investigate. They are not yet demonstrated value. AI earns its place when the organization can show it improves the work after accounting for verification, correction, operating costs, and the consequences of errors. Human review may consume much of the time initially saved. Clinical terminology errors may create additional work. Faster content production alone is not sufficient evidence of value. That is the central argument of *AI for Learning & Development: Faster Content Isn't Better Learning*: speed is not the measure. The comparison that matters is practical: quality, verified hours saved, total cost, and risk, measured against the conventional alternative.

Where a simpler solution wins

The second zone is the one vendors least want to discuss. Many problems presented as AI opportunities are solved better by simpler means. A team struggling to find the right policy document does not need a generative assistant; it needs a searchable repository and five minutes of training on how to use it. Inconsistent onboarding does not need a personalized AI learning path; it needs a standard checklist and a preceptor who follows it. Low completion rates do not need predictive analytics; they need someone to ask the staff why they are not finishing.

The test is straightforward: what would the non-AI solution cost, and what would it achieve? If a checklist, a job aid, a process fix, or a conversation closes most of the gap at a fraction of the cost and risk, the simpler solution wins. This is the minimum-effective-dose thinking from Chapter 2 applied to technology choice. A CLO who cannot say "we do not need AI for this" will spend the AI budget on problems that were never AI problems.

Where risk outweighs benefit

The third zone is where the CLO's governance role matters most, and where the decision rights must be precise. AI proposals must receive approval from the authorities responsible for the risks they create. Clinical leaders retain responsibility for clinical content and professional judgments. Privacy, information security, and legal reviewers determine whether the proposed use of information is permissible. HR and operational leadership address employee-related impacts, including collective agreement considerations. One clarification matters: privacy law applies to personal health information within its legal scope, so employee performance data does not automatically fall under Ontario's health privacy law merely because the employer is a healthcare organization. Other legal obligations under privacy and employment law may also apply, and the reviewers named above are the ones who determine which.

Where risk outweighs likely benefit, the proposal should not proceed. The right distinction is between risks that can be controlled and risks that prevent deployment. Some concerns can be addressed through changes to the design or a controlled evaluation. Others, including an unacceptable privacy risk or insufficient evidence of safety, require rejection. A pilot is appropriate only after the organization has established that the proposed use is permissible and that the remaining risks can be managed. A small trial does not make an unlawful or unacceptable use safe.

The question behind the risk zone is straightforward: what would cause this trial to stop? The organization does not need to fear AI, and it does not need to chase it. It needs a register of what is

being tried, where each item sits on the spectrum, what the risks are, and what would cause the trial to stop.

What the workforce needs to work with AI

There is a second AI question the learning function cannot dodge, and it is larger than the function's own tools. As AI enters clinical documentation, scheduling, triage support, and decision support, the workforce needs the capability to work with it safely. That means knowing what the tool was designed to do and what it was not, checking its output before acting on it, recognizing when it is uncertain or wrong, and understanding the privacy obligations that travel with patient information into any system. These are not technical skills for a specialist team. They are basic professional competencies for anyone whose work touches an AI-assisted process, and building them is squarely the learning function's job. The same spectrum applies: teach the judgment first, add the tool second, and never let facility with the tool substitute for accountability for the outcome.

A worked example: triaging three proposals

WORKED DECISION EXAMPLE

The following is a worked decision example, not a case study. Each proposal receives a conditional decision; no outcomes are assumed.

A learning function receives three AI proposals in one quarter. The first: a generative tool that drafts personalized practice scenarios for wound-care competency validation, with every scenario reviewed by a nurse educator before use. The second: an AI chatbot to answer staff questions about policies and procedures. The third: an AI tool that scores recorded patient interactions to flag staff who may need coaching.

The spectrum triages them conditionally. The first is a candidate for a limited evaluation using approved, non-identifiable material. Compare the educator effort and total cost against existing development methods, and check the accuracy of what it produces, before any wider use. One caution travels with it: AI-generated practice questions are useful preparation material, but they are not automatically valid measures of clinical competence. Formal competency decisions require professionally defensible assessment methods, and the evaluation must test that, not assume it.

The second goes to the simpler-solution test: test ordinary search and document improvements first, and reconsider AI only if a meaningful unmet need remains. The third is withheld until necessity and legal authority are established. Fairness, clinical implications, and the appropriate decision rights must be assessed before any further evaluation is considered. The privacy implications alone are serious; add bias risk and labor-relations exposure, and the concerns dwarf the proposed efficiency.

Three proposals, three conditional decisions, one method. That is the spectrum working.

What this means in year one

For the executive reader: in year one, the CLO should establish which AI applications are worth evaluating, confirm the governance arrangements that apply, and authorize only the uses supported

by sufficient evidence and appropriate safeguards. Wider deployment depends on demonstrated value and the organization's ability to manage the associated risks. A modest start is not caution for its own sake. It is sequencing: the organization proves, on small and well-governed evaluations, that it can tell value from hype before it commits at scale.

The report closes with a brief note on who wrote it and why it can be trusted.

Endnotes

- AI proposals involving patient data in Ontario must satisfy the Personal Health Information Protection Act and the organization's own information governance. The CLO's role is to ensure these reviews happen before any pilot, not to conduct them.
- Information and Privacy Commissioner of Ontario, AI Scribes: Key Considerations for the Health Sector (January 2026), <https://www.ipc.on.ca/en/resources/ai-scribes-key-considerations-health-sector>. Its subject is AI scribes rather than learning technology, but its treatment of vendor assessment and contractual safeguards, with monitoring and accountability over time, applies to any AI handling personal health information.
- Canada's privacy commissioners, Principles for responsible, trustworthy and privacy-protective generative AI technologies (December 2023), https://www.priv.gc.ca/en/privacy-topics/technology/artificial-intelligence/gd_principles_ai/. Relevant to legal authority and to the necessity and proportionality of the use.
- College of Physicians and Surgeons of Ontario, Using Artificial Intelligence in Clinical Practice (August 2025), <https://www.cpso.on.ca/physicians/policies-guidance/advice-to-the-profession/using-artificial-intelligence-in-clinical-practice>. Relevant to physician accountability for AI tools, including bias and privacy.
- The Pan-Canadian AI for Health Guiding Principles, approved by the country's health ministers, set out shared values for responsible AI adoption across Canada's health systems.

How the Frameworks Connect

One page a board member can grasp in seconds: from diagnosis through execution, with AI as a decision filter on technology choices.



About the Author

Ravinder Tulsiani is Principal of Catalyst Learning & Strategy. His twenty-plus years in learning and organizational performance span the healthcare and financial sectors, along with public-sector and post-secondary work.

He spent eight and a half years leading learning and development at one of Canada's largest healthcare organizations, where as Director of Learning & Development he led a team of 14 with an annual budget of approximately \$2M, serving a network of more than 60,000 learners. During that tenure he was a member of the organization's enterprise-wide AI Governance Committee and chaired its Education Technology Committee. His work has centered on a single question: how organizations know whether their investment in people is producing the performance they need.

He is the author of the Capability Architecture series, including *Proving Impact* and *AI for Learning & Development: Faster Content Isn't Better Learning*. He holds a DBA in Organizational Capability and Human Resource Management. His designations include the CTPD and PMP, and he is a Certified Six Sigma Black Belt (CSSBB).

© 2026 Dr. Ravinder Tulsiani. All rights reserved.