ESOURCE — CHARLESLWAND

The High-Impact eLearning Design Guide

How to Design Training That Changes Behaviour, Not Just Ticks a Compliance Box

A comprehensive guide for L&D; professionals, HR managers, and anyone responsible for training that is supposed to produce results. Built on learning science, refined through 100+ real-world projects, and written for people who are tired of building courses nobody remembers.

By Charles (Audi) Lwanda

Learning Experience Architect · Nairobi, Kenya

charleslwanda.com

6+	100+	7
Years in L&D;	Projects Completed	Sections in this Guide

CONTENTS

SECTION 01

Why Most Training Fails Before It Starts

The single most common mistake in corporate learning is building the right solution to the wrong problem. Before you write a single learning objective, before you open Storyline, before you book the SME interview — you need to answer one question: is this actually a training problem?

Thomas Gilbert's Behaviour Engineering Model (1978) identified six causes of performance problems. Only two of them — knowledge deficiency and skill deficiency — are addressable through training. The other four (unclear expectations, inadequate feedback, wrong tools, misaligned incentives) require process changes, management interventions, or system redesign. Training applied to these causes produces courses that are technically sound and completely useless.

The Six Causes of Underperformance

Information / Expectations	Do people know what they are expected to do? Is the standard clear and communicated regularly?	**Not training**
Feedback	Do people know when they are performing correctly or incorrectly? Is feedback timely and specific?	**Not training**
Resources / Tools	Do people have what they need to perform correctly? Time, tools, authority, support?	**Not training**
Incentives	Is correct performance rewarded? Is incorrect performance allowed or even incentivised?	**Not training**
Knowledge	Do people genuinely lack the knowledge to perform? Training addresses this.	**Training ✓**
Skill / Practice	Do people lack the practised ability to execute, even if they know the theory? Training addresses this.	**Training ✓**

The Diagnostic Question to Ask Before Every Brief

"If we gave everyone perfect knowledge of this topic tomorrow morning, would the performance problem disappear?" If the answer is no — or even maybe — you have not yet identified a training problem. Keep asking until you find one, or be honest that training is not what this situation needs.

SECTION 02

The Performance Gap Framework

Every training project I take on starts with the same five questions. They are not a needs analysis template — they are a diagnostic protocol that separates a genuine training problem from everything that looks like one.

01

What are people doing vs. what should they be doing?

Describe the observable behaviour gap in concrete terms. Not "they need to understand data protection" but "they are sharing client data in unencrypted email when they should be using the secure portal." If you cannot describe the gap in observable behaviour, the problem is not yet defined well enough to design a solution.

02

Why aren't they doing it?

Work through Gilbert's six causes before assuming training is the answer. Talk to people who perform the task well and people who do not. The gap between their answers tells you more than any survey.

03

What does "good" look like, and how will we measure it?

Define success before designing anything. A specific, measurable outcome agreed with the business stakeholder before development begins. If you cannot agree on a measure, you cannot agree on a solution.

04

What is the minimum effective intervention?

What is the simplest thing that would produce the required behaviour change? Build the minimum, measure it, and iterate. Over-engineering is expensive and often counterproductive.

05

What will reinforce the learning after it ends?

Training without reinforcement decays. What happens in the 30 days after the course? Manager conversations, job aids, spaced retrieval, practice opportunities — these are not optional extras. They are the mechanism by which transfer actually occurs.

SECTION 03

Designing for How the Brain Actually Works

Most eLearning is built on how designers think learning should work, not on how the brain actually processes and retains information. The gap between those two things is where most of your learning budget disappears.

Cognitive Load Theory

John Sweller, 1988

Working memory has a limited capacity. When instructional content exceeds it, learning degrades. Design implication: chunk content into 5–7 concept units maximum per screen. Remove all content that does not directly serve the learning objective. Use worked examples for complex tasks before asking learners to solve independently.

The Forgetting Curve

Hermann Ebbinghaus, 1885

Without reinforcement, 70% of new information is forgotten within 24 hours. Design implication: single-exposure training is a waste of money. Build spaced retrieval into every programme — even if it's just three follow-up questions sent via email at 7, 14, and 30 days post-training.

Dual Coding Theory

Allan Paivio, 1971

Information encoded both verbally and visually is retained more strongly than information encoded in one mode alone. Design implication: pair narration with relevant visuals, not with on-screen text. Narration plus on-screen text creates split-attention effect and degrades retention.

The Testing Effect

Roediger & Karpicke, 2006

Retrieving information from memory strengthens the memory trace more than re-reading the material. Students who took three practice tests retained 80% after one week. Those who re-read the material retained 40%. Same time. Double the retention.

Recognition-Primed Decision Making

Gary Klein, 1998

Experts do not make decisions by analysing all options. They pattern-match to familiar situations and act on the recognised pattern. Design implication: build scenario-based training that gives learners practice recognising the patterns that trigger the right response — not training that teaches them checklists.

SECTION 04

Scenario Design That Actually Works

A scenario where the wrong answer is obviously wrong teaches nothing. Real decisions are hard because all the options are defensible. If your learner can guess the right answer without engaging, your scenario is decoration.

The Five Markers of a Broken Scenario

✗ **One option is clearly wrong**

If learners can eliminate it without reading, you have not created a decision. You have created the appearance of one.

✗ **The scenario is generic**

Vague characters and situations produce vague learning. The more specific and real the scenario, the more transfer to real work.

✗ **Wrong answers lead to "Try Again"**

In real life, mistakes have consequences that unfold over time. Your scenario should replicate this.

✗ **The correct answer always aligns with what learners want to hear**

Real ethical and professional dilemmas are uncomfortable. A scenario without discomfort is a fable.

✗ **The debrief gives the answer without exploring the reasoning**

Reflection is where learning consolidates. "Here is the right answer" closes thinking. "Here is what different choices set in motion" opens it.

The Five Elements of a Scenario That Works

✓ **Authentic tension**

All choices are defensible. The learner has to think, not recognise.

✓ **Specific, named characters**

Real names, real roles, real pressures. Generic characters produce generic decisions.

✓ **Consequence fidelity**

Outcomes match what would actually happen — including delayed, unexpected consequences.

✓ **Emotional stakes**

The decision should matter to someone. Stakes produce engagement. Engagement produces encoding.

✓ **Reflective debrief**

Not "here is the right answer" but "here is what different choices set in motion and why."

SECTION 05

Assessment That Measures What Matters

Assessment is not the end of learning — it is one of the most powerful learning tools available to you. The testing effect shows that retrieval strengthens memory more than re-exposure. Design your assessments as learning events, not gatekeeping mechanisms.

The Assessment Design Principles

- ☐ **Test application, not recall**
  Multiple choice tests recognition. Scenarios test judgment. Reflection tests metacognition. Decide which you need — then design for it, not for what is easiest to build.
- ☐ **Make distractors genuinely plausible**
  Every wrong answer should be something a person who does not know the material would seriously consider. If the wrong answers are obviously wrong, you are measuring the ability to read, not the ability to perform.
- ☐ **Justify the pass threshold with evidence**
  80% is not a meaningful standard unless you can explain why 79% is insufficient and 80% is adequate. Agree the threshold with the business based on the risk of the performance gap.
- ☐ **Use feedback as a teaching tool**
  Every answer — correct and incorrect — should generate feedback that explains what the choice means in practice. This is where the majority of learning from assessment occurs.
- ☐ **Separate practice from summative assessment**
  Practice should be safe to fail. Summative assessment should reflect actual performance conditions. Mixing them produces data that is impossible to interpret.
- ☐ **Space assessment over time**
  A single end-of-course quiz measures short-term retention, not durable learning. Build in retrieval moments at 7, 14, and 30 days. The data you get back will be more honest and more useful.

★ *Red-outlined boxes = critical items. If these fail, your assessment is measuring the wrong thing.*

SECTION 06

The 50-Question Pre-Launch Checklist

Use this before any course goes live. Not as a formality — as a genuine quality gate. Every unchecked box is a known risk. Decide consciously whether to accept it.

DIAGNOSIS & STRATEGY

- [ ] Is this genuinely a knowledge or skill problem?
  Run through Gilbert's six causes before proceeding.
- [ ] Can you describe the behaviour gap in observable terms?
  "Do A instead of B" — not "understand X better."
- [ ] Have you spoken directly to people who perform this task?
  Not just the SME or manager — the actual performers.
- [ ] Have you spoken to learners directly?
  Their reality differs from what stakeholders describe.
- [ ] Have you defined success with a business metric?
  A number, a timeframe, and an agreed measurement method.
- [ ] Has the stakeholder agreed that training is the right intervention?
  Written or verbal — you need alignment before building.
- [ ] Do you know what happens if you do nothing?
  The counterfactual should inform the design brief.
- [ ] Is there a plan for what happens after the training ends?
  Manager reinforcement, spaced retrieval, job aids.

LEARNING OBJECTIVES

- [ ] Are objectives written in observable, measurable behaviour?
  Bloom's action verbs: identify, apply, evaluate, construct. Not "understand" or "appreciate."
- [ ] Is each objective traceable to a performance gap?
  If you cannot trace an objective to a consequence, cut it.
- [ ] Have performers (not just SMEs) reviewed the objectives?
  SMEs describe what they know. Performers describe what they do.
- [ ] Are there 5 or fewer objectives per module?
  More than 5 suggests scope creep rather than instructional necessity.
- [ ] Does each objective have a corresponding assessment item?
  If it cannot be tested, it should not be taught.

CONTENT & STORYBOARD

- [ ] Has all content been reviewed for accuracy by an SME?
  And have you pushed back on content that does not serve an objective?

- [ ] Have you cut at least 20% from the first content draft?
  First drafts are always too long. Ruthless editing improves learning.
- [ ] Is on-screen text minimal — under 30 words per slide?
  Narration plus on-screen text creates split-attention. Use one or the other.
- [ ] Are all visuals functional, not decorative?
  Stock photos of people smiling contribute nothing. Diagrams and decision trees do.
- [ ] Has the storyboard been approved before development began?
  Changes in development cost 3–5x more than changes in storyboard.
- [ ] Are branching logic and variables specified in the storyboard?
  Not improvised during build.
- [ ] Can every screen answer the question "why does this screen exist"?
  If not, cut it.

DEVELOPMENT & QA

- [ ] Has the course been tested in the actual delivery LMS?
  SCORM on one platform ≠ SCORM on all platforms.
- [ ] Have all branching paths been tested end-to-end?
  Including every wrong answer path and edge case.
- [ ] Is completion tracking verified to trigger correctly?
  Check the xAPI or SCORM statements in the actual environment.
- [ ] Have accessibility requirements been checked?
  Alt text, keyboard navigation, colour contrast, captions.
- [ ] Has a learner (not the SME) piloted the course?
  SMEs are the worst beta testers. They know too much.
- [ ] Has load time been tested on a slow connection?
  Corporate learners often access training on restricted networks.
- [ ] Is the file size within the LMS upload limit?
  Compress all media before packaging.
- [ ] Have all placeholder links and dummy text been replaced?
  Ship-ready means zero placeholders.

POST-LAUNCH & MEASUREMENT

- [ ] Are you tracking more than completion rate?
  Completion is the lowest signal. Start with time-on-task and assessment performance.
- [ ] Is there a 30-day check-in scheduled with the business owner?
  Is the performance gap closing? What are managers observing?
- [ ] Is qualitative feedback being collected beyond star ratings?
  "What will you do differently?" beats "Rate your experience 1–5."
- [ ] Is there a plan for updating the course when content changes?
  Outdated courses are worse than no course.
- [ ] Has a post-project review been scheduled?
  What worked, what you would change, and what the data showed.

SECTION 07

Measuring Real Impact: Kirkpatrick Without the Excuses

The Kirkpatrick Model has existed since 1959. Most L&D; teams are still primarily measuring at Level 1. This is not a resource problem. It is a priority problem — and it is costing the learning profession its seat at the strategic table.

Level 1 — **Reaction** — Did learners enjoy it?

Useful for identifying design problems. High satisfaction does not predict learning or transfer. Never report Level 1 data as evidence of training effectiveness.

Level 2 — **Learning** — Did they acquire knowledge or skill?

Assessment scores, pre/post comparisons, observed demonstrations. Meaningful but limited — measures performance in a training environment, not on the job.

Level 3 — **Behaviour** — Are they applying it?

Manager observation, peer feedback, performance data. This is where transfer becomes visible. Measure at 30 and 60 days post-training. Most L&D; teams never get here.

Level 4 — **Results** — Did it move the business?

The metric the stakeholder actually cares about. Customer scores, incident rates, conversion, turnover. Establish a baseline before training. Measure at 60 and 90 days. This is the only level that justifies training budgets.

The One Metric That Changes Everything

Before your next training programme launches, agree on one Level 4 metric with your business stakeholder. One number. A baseline. A target. A measurement date. That conversation — and the discipline of going back to check — will do more for your credibility as an L&D; professional than any course design improvement.

A closing note: This guide is built on one belief — that learning can be genuinely excellent. Not just adequate. Not just compliant. Excellent in the sense that it changes what people do, produces outcomes the business can measure, and treats learners as the capable adults they are. If even one project you build after reading this is better because of it, the guide has done its job.

charleslwanda.com · info@charleslwanda.com · Available for freelance, retainer & full-time remote L&D; roles