

# Harland & Finch: The Diane Vignette

*A scenario-based cybersecurity awareness learning experience, built from brief to deployed web application using AI-assisted instructional design.*

Scenario-Based Learning	Cybersecurity Awareness	HTML / CSS / JS	AI-Assisted Design	ADDIE
3 Learner Personas	9 Learning Outcomes	40 QA Test Cases		1 HTML File

## 01 — THE BRIEF

### A fictional brief, grounded in real risk

Harland & Finch is a fictional mid-size UK law firm facing a documented and growing risk of Business Email Compromise (BEC) attacks. The firm had one near-miss in the previous twelve months. A Training Needs Analysis identified that staff could identify phishing in the abstract but failed to apply that knowledge under realistic working conditions — time pressure, competing demands, and the social complexity of a professional services environment.

Existing training was click-through slides with no consequence mechanics — described internally as "a tick-box exercise." The brief called for scenario-based learning that tested decision-making under realistic conditions, not knowledge recall in a calm environment.

*"The problem was not that staff didn't know what phishing was. The problem was that knowing and doing are different skills — and the existing training only tested the first one."*

## 02 — ANALYSIS

### Understanding the learners before designing anything

The analysis phase produced three learner personas, nine learning outcomes, and a set of design constraints. An anthology structure was chosen because a single protagonist risks alienating learners whose role doesn't match. Each persona carries different responsibilities, different risk exposure, and a different relationship to the firm's data and communication systems.

<b>DIANE</b>	Legal secretary covering for a colleague. High email volume, time pressure, limited authority to delay decisions.
<b>TOM</b>	Senior fee-earner. High trust, high-value decisions. Primary target for impersonation and payment-diversion attacks.
<b>PRIYA</b>	Administrative staff. Broad access to systems, lower perceived risk. Entry point for credential and data harvesting.

### *Learning Outcomes*

Nine LOs were developed from the TNA, behaviour-specific and verb-led, mapped to real decision points in the firm's operations. Coverage across the two completed interactions:

LO	Outcome	Coverage
LO1	Analyse messages for phishing indicators and determine the appropriate verification or reporting action before responding.	Task 1
LO2	Recognise the need to report suspicious messages and near-misses promptly, following the firm's reporting procedure.	Task 1
LO3	Apply a deliberate recipient-checking process to verify the correct person before sending or releasing documents.	Task 2
LO4	Check attachments and documents to confirm they are correct, necessary, and relevant before release.	Task 2
LO5	Evaluate the risk of alternatives to secure sending and choose the approved method for sensitive documents.	Deferred
LO6	Assess the sensitivity of outgoing information and apply the appropriate level of protection.	Deferred
LO7	Verify the identity of callers or requesters before acting on verbal or written instructions.	Task 1
LO8	Follow the firm's clean desk and device policy when leaving a workstation.	Deferred
LO9	Follow the correct escalation procedure when unable to complete a task safely or without appropriate authority.	Task 2

## Where the thinking happened

The decisions below represent points where professional judgement shaped the design — where a default option was rejected, where domain knowledge overrode a suggestion, or where a constraint was acknowledged honestly rather than designed around.

### 01 No countdown timers — pressure mechanics rooted in accessibility

**SITUATION** The TNA finding was clear: staff fail under real-world pressure, not calm conditions. Time-based mechanics were the obvious response. They were rejected on accessibility grounds before the argument had been made in full.

**DECISION** Pressure mechanics ruled out. Pressure delivered instead through information overload, interruption/context-switching, and social pressure built into the scenario structure.

**PRINCIPLE** Under the Equality Act 2010, equivalent experience is required, not just equivalent content. A module where a screen reader user is penalised for their assistive technology's reading speed doesn't meet that bar.

### 02 Realistic independent interaction — no prompted option lists

**SITUATION** Standard branching scenarios present a labelled decision point with options to choose from. This eliminates the most important cognitive skill: recognising that a decision is needed.

**DECISION** Learners must read the email, hover or click links, examine content, and act without being told where the decision is. The interface simulates a real desktop environment.

**PRINCIPLE** Option lists tell the learner where to look. In the real world, nobody presents a phishing email with three labelled response options. The learning value is in the identification, not just the choice.

### 03

#### Enterprise email banner — domain knowledge over design convenience

**SITUATION** The original design used an interactive popup triggered when Diane clicked the sender address — a discovery mechanism for the domain mismatch.

**DECISION** Replaced with a static amber "external sender" banner already present when the email is opened. No click required.

**PRINCIPLE** Harland & Finch was established as having enterprise-level email protection. External sender banners are standard. Requiring Diane to discover information her email client would already surface is both unrealistic and undermines the scenario logic.

### 04

#### Acknowledged LO3 limitation — constraint named, not hidden

**SITUATION** LO3 (verify correct recipient before releasing documents) is harder to make genuinely interactive in a physical document context on a desktop screen than in an email-based task.

**DECISION** Proceed, but log the limitation explicitly. The design thinking is sound; the execution is constrained by the portfolio format.

**PRINCIPLE** A portfolio piece that pretends constraints don't exist is a less honest demonstration of professional judgement than one that owns them. In a real commission, richer interaction types would be available.

## Working with AI as a design partner

This project was built using Claude as an active thinking partner throughout — not as a content generator, but as a pressure-tester, a researcher, and a first-draft challenger. In each exchange below, the direction came from the designer; the AI's role was to examine whether that direction held up.

### *The incident report frame — an idea Claude didn't suggest*

Claude offered standard narrative frame options for the anthology: a covering email, a case file wrapper, a static introduction. None were compelling. The frame that was adopted came from a different direction entirely.

**Amy** What about a frame along the lines of 'the incident report', but where the learner's decisions during interaction ultimately decide whether there ends up being a problem or not?

**Claude** That's a strong idea — and it's yours, so log it. What you're describing is a shift from passive narrator ("here's what happened") to active architect ("you decide what happens"). Done well, the incident report at the end either reads as a near-miss summary or a breach notification — depending entirely on how the learner played it.

*The idea came first; the AI's job was to confirm it was sound and flag what would make it fail. Claude was prompted to log the decision with the source correctly attributed.*

### *Pressure mechanics — pushing back before Claude finished exploring*

Claude began exploring pressure mechanisms including time-based scoring. The response came before the exploration had run its course.

**Claude** Countdown timers and response-time scoring are the most direct mechanisms for replicating real-world urgency in a digital environment...

**Amy** I don't want any countdowns or metrics based on time spent responding to a challenge. Information overload, interruption and context switching, and social pressure I think can all be built into a learning scenario, however.

**Claude** Log this? "Pressure mechanics: time-based constraints ruled out on accessibility grounds. Information overload, interruption/context-switching, and social pressure identified as the three legitimate mechanisms."

*The redirect was immediate and grounded in the Equality Act before the argument had been made in full.*

### *Stress-testing scenario realism — conveyancing access permissions*

Claude proposed reframing Diane's cover role to a family law matter. The suggestion was corrected within the established scenario parameters.

**Claude** Reframing the vignette so that Diane is covering a family matter would heighten the access-restriction conflict. Family law has additional ethical walls, creating a clearer instructional tension...

**Amy** Reframing the entire vignette so that Diane is covering a family matter would require significant rework. She is covering conveyancing and the whole interaction is based around transferring monies for a completion. Stress-test this in that context.

**Claude** You're right — I overstepped there. Diane's context is already established: she's covering for Sarah in conveyancing. Let me stress-test both options within conveyancing.

*The pattern — AI overreaches, designer redirects, AI recalibrates — is characteristic of effective AI collaboration. The question is not whether the AI makes incorrect suggestions, but whether the designer has the domain knowledge to catch them.*

## 05 — THE BUILD

# A single HTML file, deployed from GitHub to Cloudflare

The Diane vignette is a single self-contained HTML file — one document, one CSS design system, one JavaScript state object managing both interactions. The choice was deliberate: portfolio clarity and portability over modularity. A single file is easier to deploy, inspect, and share, and it ensures that Diane's progress persists naturally throughout the session.

The interface simulates a realistic workplace desktop environment: a persistent sidebar showing Diane's task list, an email inbox, a document management system view, and a directory panel. Learners interact directly with these elements rather than being presented with abstracted branching choices.

The project is hosted on Cloudflare Pages, deployed via GitHub push. The tech stack — HTML, CSS, vanilla JavaScript — was chosen because it produces artefacts that work anywhere, require no toolchain, and are fully inspectable by any reviewer who looks at the source.

## 06 — QUALITY ASSURANCE

# 40 test cases before a line of production JavaScript

The QA matrix for Task 1 was written before the interaction was built. If the learner can interact with many elements in many orders, the possible paths must be mapped before they can be validated. Each test case specifies: the trigger action, the expected system response, the feedback content and tone, the resulting state change, and any branching conditions.

<b>40</b> Total test cases	<b>5</b> Interaction branches	<b>3</b> Outcome types	<b>1</b> QA matrix, pre-build
-------------------------------	----------------------------------	---------------------------	----------------------------------

## What this demonstrates, and what a real commission would look like

Portfolio context vs. real commission This project was built as a portfolio case study, which means some design decisions were constrained by format. In a real commission, I would use a proper authoring tool (Articulate Storyline or Rise, Adobe Captivate, or a custom build depending on LMS requirements), which would give richer interaction types, SCORM compliance, and proper accessibility audit tooling. The LO3 limitation in Task 2 is the clearest example. The design thinking is sound; the execution is constrained. I have named that rather than papering over it, because a portfolio piece that hides its constraints is a less honest demonstration of professional judgement than one that owns them. What this project shows A full ADDIE process, from a structured TNA through to a deployed artefact. Instructional design decisions made with domain knowledge, not just design instinct. Quality assurance as a design discipline, not a final check. And how AI can be used as a thinking partner without replacing the designer's judgement — which is, at this point in the profession's development, probably the most useful thing a senior ID can demonstrate.