FIBER INTERNET PROVIDER | SALES AND CUSTOMER CARE TRAINING

Course Design Document

Version 2.0 | November 25, 2026 | Designer: Martine Brown, M.Ed. | Instructional design work sample | Supersedes v1.0

1. Project overview

Field	Detail
Working title	Selling Fiber with Confidence, Module 1: **You're Live**, The Right Speed Conversation
Format	Immersive, voiced HTML5 module built as a softphone "agent desktop." Packaged as SCORM 1.2 for an LMS and published as a web demo on Netlify.
Length	About 14 minutes. Eight scenes, five animated films (about 2 minutes in total), 65 voiced lines, one live branching call, one four-call knowledge check.
Audience	Sales and customer care team members who sell or support residential internet. Mixed product knowledge. Most can describe the plans; fewer can diagnose why a household is unhappy with one.
Delivery	Plays in any modern browser, desktop or phone, no plugin. Reports completion, pass status and a percentage score to the LMS. The web demo runs the same build without an LMS.
Portfolio role	Instructional design work sample: curriculum design for sales and customer-facing teams, scenario-based learning, AI-assisted media production (voice, animation, illustration), and a full design document and storyboard set.
What changed from v1.0	v1.0 was a silent, seven-screen reading module with one branching scenario. v2.0 keeps the same objectives, method and scoring, and turns the practice into a live call the learner hears and feels. Section 8 records each change and its reason.

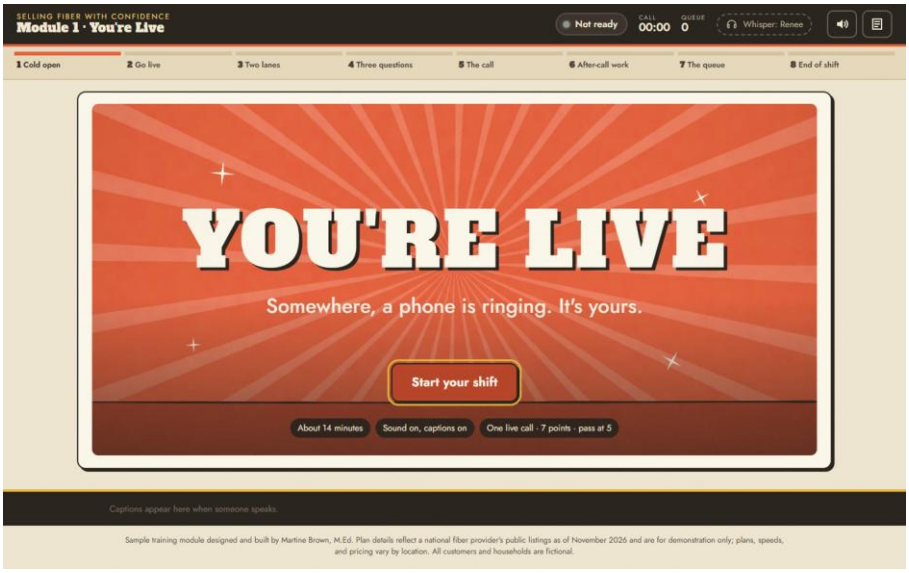

Figure 1. The start screen. The learner begins their "shift" on a softphone agent desktop.

2. Business need

Frontline team members field the same complaint every day: "my internet is slow." The stated need is speed. The unstated need, which only surfaces when someone asks, is a household whose daily life has outgrown its plan, very often on the upload side. Cable uploads run at a fraction of download speed, so a home where two adults send video all day will feel broken on a plan that looks generous on paper. When a team member diagnoses download for an upload problem, the customer gets a bigger plan that does not fix anything, and the next call is angrier.

This module teaches the diagnosis before the pitch: ask about the household, listen for the stated need and the one underneath it, then match the home to the right fiber tier and say the reason in the customer's own words.

3. Audience

Attribute	Profile
Role	Residential sales and customer care team members, phone and chat
Product knowledge	Can name the plans and prices. Shaky on the cable versus fiber difference, especially symmetrical upload.
Selling experience	Mixed. Comfortable offering a promotion; less practiced at discovery questions.
AI experience	Has access to AI assistants such as Microsoft Copilot. Uses them informally, if at all. No shared guardrails yet.
Motivation	Practical: fewer repeat calls, recommendations that stick, confidence on the hard calls.
Time available	Low. Training happens between calls. Anything over 15 minutes will be rushed or skipped.

What that implies for the design

Learner reality	Design response
Short seat time	One module, about 14 minutes, one customer story carried all the way through
Knows the catalogue, not the diagnosis	Product content is one 55-second film and three tap-to-hear plan cards. The method and the live call get the rest.
Learns by doing	A live branching call with three scored decisions, a listening task and an optional spoken rehearsal
Learns by feeling the stakes	The module opens on the customer's side of the problem, and every wrong recommendation plays out as a callback weeks later. The learner hears what a bad match costs before being told.
Needs it on the next call	A printable pocket guide, Match the Household, that maps what a customer says to the tier that fits
AI is new and ungoverned	The AI section is practiced, not read: the learner builds a prompt, is stopped from pasting private data, and catches an error in the AI's draft

4. Learning objectives

By the end of the module, the learner will be able to:

#	Objective	Level	Practiced and assessed in
1	Explain the difference between cable and fiber internet, including symmetrical speeds, in plain customer language	Understand	Scene 2 (Two Lanes film, plan cards); Queue call 1
2	Use three discovery questions to surface a household's stated needs and the needs underneath them	Apply	Scene 3; Scene 4 Decision 1, the change question; Queue call 3
3	Diagnose an upload-driven problem from a customer's story	Analyze	Scene 4 "Tap the clues" and Decision 2; Queue call 2
4	Recommend the right speed tier for a household profile and justify it in the customer's own words	Evaluate	Scene 4 Decision 3; optional "Say it yourself" rehearsal
5	Use an AI assistant prompt to prepare for and summarize a customer conversation, within data guardrails	Apply	Scene 5 prompt builder and draft check; Queue call 4

Out of scope: provisioning and installation, hardware troubleshooting, billing and retention offers, mobile service, and multi-gig tiers beyond a mention. Each belongs in its own module if this becomes a series.

5. Instructional strategy

Story first. The learner feels the problem before learning the fix.

The module opens with **The Freeze**, a 17-second film from the customer's side: Dana is presenting to her team at 9:02 AM when her video freezes mid-sentence, a coworker says "Dana? Dana, you froze again," and she whispers "Every single morning." A phone rings, and the title card says it is yours. The learner takes the call already knowing what it feels like to be Dana, which is the empathy the job asks for.

Method before tool. The human skill is taught before the tool skill.

A short huddle teaches the two-lane model of upload and download and the three discovery questions. The AI prompts arrive only after the call, once the learner has something worth prompting about. A prompt is only as good as the discovery behind it.

Practice over presentation. The call is the curriculum.

- **Whisper-line coaching.** The learner's team lead, Renee, coaches on a whisper line the customer cannot hear, the way real call-center supervisors do. Every piece of feedback lives inside the story, in the voice of someone who has done the job. There is no bare "try again" anywhere in the module.
- **The house reveal.** A cutaway of Dana's house sits beside the call with every room dark. As she describes her morning, the rooms light up in sync with her words: her office on video, Theo uploading in the back room, the teens' room "after 3 PM." The unstated need becomes something the learner can see appear.
- **Consequence callbacks.** A wrong recommendation is not marked wrong. It plays out: Dana calls back three weeks later, angrier, or two months later, doubting her bill and ready to shop around. Then Renee explains why, a rewind plays, and the learner chooses again.

- **"Say it yourself."** At the recommendation, the learner may record themselves saying it in their own words and then hear a senior rep say it. Saying it out loud is the behavior the job needs.
- **A payoff.** The right recommendation earns **The next morning**, a mirror of the cold open with no freeze: "Same morning. Different call."

Transfer at exit.

The module ends with a performance support tool, not a summary slide: the Match the Household pocket guide, printable and designed to be open during a live shift.

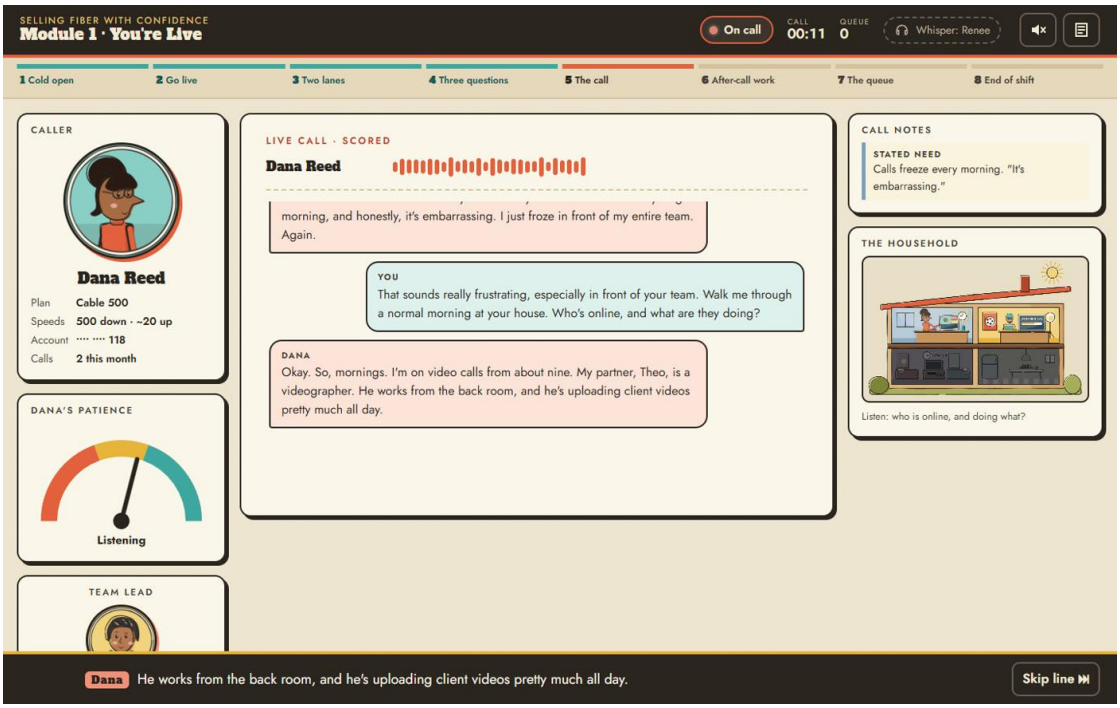

Figure 2. The live call. Dana's caller card and patience meter (left), the call and decisions (center), the notes panel and the house cutaway lighting up room by room (right).

6. Course structure

#	Scene	Purpose	Primary interaction	Obj.
0	Cold open: The Freeze	Feel the problem from the customer's side; set the stakes	Film F1 (17.5 s), then Take the call	-
1	Go live	Meet Renee on the whisper line; state the objectives	Voiced intro; Join the huddle	-
2	Huddle: Two lanes	Teach cable versus fiber and symmetrical speeds; the three tiers	Film F2 (55 s); three tap-to-hear plan cards	1
3	Huddle: Three sticky notes	Teach the three discovery questions and what each reveals	Tap each sticky note to hear Renee; Go available	2
4	The call	Practice the full method on one household	Live branching call: 3 scored decisions, Tap the clues, the change question, optional recording; callback films F4a and F4c; payoff film F5	2, 3, 4

#	Scene	Purpose	Primary interaction	Obj.
5	After-call work	Use AI for the follow-up, within guardrails	Drag notes into a prompt; private data blocked; catch the draft's error; two prompts to copy	5
6	The queue	Confirm the core ideas under check conditions	Four short voiced calls, three choices each, scored	1, 2, 3, 5
7	End of shift	Score, completion, and the pocket guide	Scorecard; Download pocket guide; Replay Dana's call; Restart	-

7. Assessment and completion

Where	Type	Scored	Purpose
Scene 4	Live branching call, three decisions	Yes, 3 points (first attempt only)	Shows whether the learner can run the method, not just recall it. Retries are allowed so every consequence can be heard, but only the first choice at each decision counts.
Scene 4	Tap the clues; the change question; Say it yourself	No	Formative listening and rehearsal. Tap the clues is required (all four clues found) before the call continues.
Scene 5	Prompt builder and draft check	No	Practiced guardrails: private cards are blocked with the reason; the draft's planted error must be caught to continue.
Scene 6	Knowledge check, four calls	Yes, 4 points	Symmetrical speed, upload diagnosis, the change question, data guardrails
Scene 7	Completion	Combined 7 points	Pass at 5 of 7. SCORM 1.2 reports cmi.core.score.raw as a percentage (min 0, max 100) and lesson_status "passed" at 5 or more, otherwise "completed." Mastery score 71 in the manifest.

Feedback design

Every incorrect choice gets specific, voiced coaching that names why it fails on a real call, in the customer's terms. Correct choices get the reasoning, not praise. Wrong recommendations are shown, not told, through the callback films. Correct and incorrect states are always paired with words, never color alone.

8. Design decisions

Each of these was settled during the build and is recorded with its rationale.

Revised: from "No narration" to "Voiced, with captions"

v1.0 was designed to work silently, like a rep reading between calls. In review, that choice protected seat time but gave up the one thing a phone skill needs: hearing the customer. Tone, frustration and relief are what a rep actually responds to, and a silent page cannot carry them. v2.0 is voiced throughout, and the

reversal keeps v1.0's reason intact: captions are always on, a full transcript is one click away, every film has burned-in captions and a caption file, and the module is fully usable on mute.

The rep is never voiced

The learner is the rep. Their choices appear as text in the call transcript and the customer answers out loud. Giving the rep a voice that is not the learner's own would break the illusion. The one exception is "Say it yourself," where the learner records their own voice; nothing is saved or sent.

Production voices: Articulate Storyline 360 AI voices

All 65 lines were generated in Storyline 360's Generate AI Audio using the ElevenLabs v3 model, which takes emotion tags such as [frustrated], [quietly] or [exhales, softening] that shape delivery without being spoken. Seven voices were cast and approved by ear: Natasha (Dana), Matilda (Renee), Chris (Marcus), Eric (the model rep), and Bill, Jessica and Roger (the three queue callers). Every line was speech-checked against the script, and two takes were regenerated when a word was misheard.

Art direction: Mid-century Cartoon, chosen from three style tests

The same 10-second cold open was produced in three styles with an identical soundtrack, so only the look changed: an isometric cutaway, a 1950s mid-century cartoon, and an urban pen-and-wash sketch. Mid-century Cartoon was chosen as the most professional fit for a corporate sales audience: it gives the customer a face for the emotional beats while staying calm and polished, and its classroom-film heritage suits instruction.



Figure 3. The three style tests at the same moment: isometric, mid-century cartoon (chosen), urban sketch.

Films built in code

The five films were produced with the open Lemo-Opuscar film pipeline: every frame is drawn by code from a single timeline, voices are lip-synced from word timings, and a sampled cool-jazz score with a four-note sonic logo ties the films together (stated on the title card, minor in the callbacks, resolved in the payoff). Any film can be rebuilt with one command when a line or a plan detail changes.

Content gleans from the provider. The look does not.

The teaching content reflects a national fiber provider's public offer: the published fiber tiers, the symmetrical speed difference, and the way sales and care teams talk to households. The visual identity borrows nothing. No provider branding, colors, or wordmarks appear anywhere; the palette, type, characters and chrome are the designer's own.

One module, fully built, instead of a series outline

A work sample proves more with one module a reviewer can finish than with three thin ones. Depth in the call, the feedback and the media carries the evidence.

HTML5 build, with a Storyline production path

The module plays anywhere with no plugin, including on a phone, and sits next to a Storyline piece in the portfolio as a second medium. The voices were produced in Storyline, and the same design document and storyboard would drive a Storyline or Rise build unchanged.

Speeds and household fit, no prices

Pricing varies by location and changes with promotions; speeds and the symmetrical upload difference are the durable teaching content. The module names the tiers and what each fits, and date-stamps its source note instead of teaching numbers that will date.

Scenario score and check score combined; first attempt only

One combined score out of 7 keeps the stakes honest: the call (3) shows application, the queue (4) shows recall and judgment. Retries exist so the learner hears every consequence, but scoring only the first attempt keeps the score meaningful.

The AI section teaches guardrails as practice, not as a footnote

The guardrail is enforced in the interaction itself: private cards bounce with a NEVER PASTE stamp and the reason, and the learner must catch the AI draft's wrong plan before continuing. The guardrail card sits at the same visual weight as the two prompts.

9. Visual design system

Direction: 1950s mid-century cartoon. Flat color, printed slightly off its ink line, warm paper, geometric characters, colour-block rooms. It should read as a well-made training product, not as a telecom marketing page and not as any provider's branding. The course shell, films and illustrations share one palette and one set of characters.

Token	Value	Use
--paper	#F4EAD5	Page background; warm paper ground
--white	#FBF6EA	Cards, stage, caption cards
--ink	#2B2420	Body text, outlines, top bar and caption bar
--coral	#E4613F (dark #B8452B)	Primary buttons, Dana's accent, scene progress, NEVER PASTE stamp
--mustard	#E8B43A (light #F3D37F)	Whisper: Renee indicator, sticky notes, highlights
--teal	#3FA7A0 (dark #257A74)	Secondary actions (Go available, Run the prompt), fiber, correct states
--blue	#7F9CC0	Walls and cool plates in films and illustrations
--good	#1B5753	Correct states, always paired with a written verdict
--warn	#A8351C	Incorrect and blocked states, always paired with a written verdict

Type: Alfa Slab One for headlines and numerals; Jost for interface and body text; Oleo Script for a single kicker. All fonts are open-licensed (SIL OFL) and self-hosted in the package, so the module renders identically offline and inside an LMS.

Components: softphone top bar (agent status, call timer, queue counter, Whisper: Renee indicator), caller card with talking portrait, patience dial, call transcript with speech bubbles, reply cards, notes panel, house cutaway, sticky notes, plan cards, prompt builder, AI draft, scorecard, pocket guide. Motion is purposeful: the patience dial eases, rooms fade in on cue, portraits lip-sync to the audio, and everything respects reduced-motion settings.

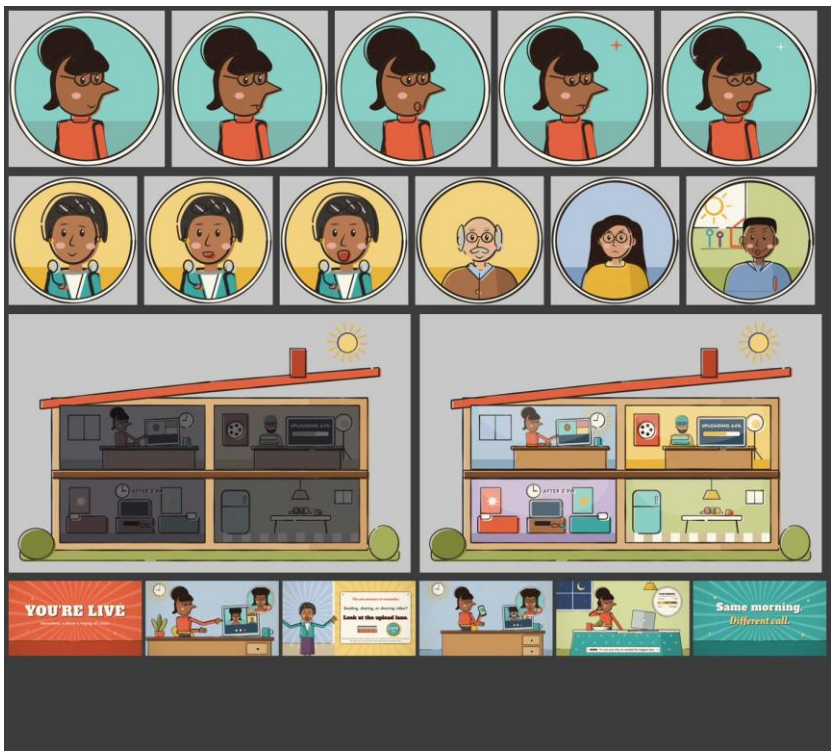

Figure 4. The illustration set: Dana's five portrait states, Renee and the three callers, the house cutaway dim and lit, and the film posters.

10. Accessibility

Target: WCAG 2.1 AA.

- **Captions everywhere.** Every voiced line is captioned in the caption bar with the speaker's name; every film has burned-in captions and a WebVTT file. A full transcript panel is available at any time.
- **Works on mute.** With sound off, captions run on timed display and every decision, cue and piece of feedback is in text. Nothing essential exists only in audio.
- **Keyboard.** Every interaction is a real button and fully keyboard operable, including the prompt builder (select a note and press Enter instead of dragging). Visible focus ring; focus moves to the next action after audio ends.
- **Screen readers.** Semantic headings, aria-live regions for captions and feedback, labelled controls.
- **Color and contrast.** Body text contrast exceeds 4.5 to 1. Correct and incorrect are always paired with words.
- **Motion.** prefers-reduced-motion shortens or removes animation.

- **Optional microphone.** "Say it yourself" is optional, handles permission denial gracefully, and never stores or sends the recording.
- **Phone width.** Columns stack; a compact patience meter stays pinned above the choices; notes and the house open as tabs.

11. Technical approach

Field	Detail
Build	Vanilla HTML5, CSS and JavaScript. No frameworks, no backend, no external network requests.
Media	5 films (H.264, about 19 MB total), 65 voice lines (MP3, about 6 MB; call and caller lines carry a light phone-line filter), 24 illustrations (about 3 MB). SCORM package about 28 MB.
Performance	Each scene preloads its own lines and the next scene's audio and film so lines start on cue. Portrait mouth frames are held in memory to avoid re-requests.
Lip sync	A per-line loudness envelope (40 ms steps) is precomputed for every voice file; portraits open their mouths only while that line is audibly playing.
Tracking	SCORM 1.2 wrapper: LMSInitialize on load; score.raw as a percentage of 7; lesson_status passed (5 or more) or completed; suspend_data and lesson_location store the scene so the module resumes; LMSCommit at key events; LMSFinish on unload. Without an LMS, progress falls back to the browser's local storage, which is how the Netlify demo runs.
Devices	Responsive from phone to desktop. Tap targets at least 44 px.
Source data	Plan tiers and speed figures reflect a national fiber provider's public listings as of November 2026, noted on the module itself.

12. Production pipeline

Step	Tool	Output and check
1. Script and voice sheet	Design document, storyboard, script	Every line with an ID, speaker, delivery note and text; numbers spelled for speech
2. Voices	Storyline 360 Generate AI Audio (ElevenLabs v3)	Lines generated in batches per character with emotion tags; exported from the Media Library
3. Speech check	faster-whisper	Batches split into one file per line ID from word timings; every line transcribed and compared with the script; misheard takes regenerated
4. Style test	Lemo-Opuscar pipeline	Same 10-second scene in three styles with one shared soundtrack; Mid-century Cartoon selected
5. Films	Lemo-Opuscar pipeline (code-drawn frames, sampled score, mux to -14 LUFS)	Five films with captions; contact sheets reviewed for continuity and factual accuracy
6. Illustrations	Same character code as the films	Portraits, house layers, callers, posters, so everything matches on screen

Step	Tool	Output and check
7. Course build	HTML, CSS, JavaScript	Scenes, interactions, captions, lip sync, SCORM
8. Test	Browser runs and a mock LMS	All-correct (7/7, passed), all-wrong (every callback and rewind plays), muted and keyboard-only, phone width, SCORM calls logged

13. Assumptions and risks

- Learners have handled live customer calls and know the current plan catalogue by name.
- Learners have access to an approved AI assistant at work; the module teaches prompts, not tool access.
- A reviewer will spend about 14 minutes, the full seat time, with the module, with sound on where possible.

Risk	Mitigation
Plan details change after the build	Speeds and household fit only, no prices; the source note is date-stamped; films rebuild from code with one command
Voice licensing	The AI voices were generated with Articulate 360's AI audio under Articulate's terms. Their use in an exported HTML build should be confirmed against Articulate's current terms before any production deployment; a Storyline build of the same storyboard avoids the question.
Audio cannot play (muted office, blocked autoplay)	Captions always on, timed display on mute, transcript panel; audio starts only after the learner's first click
Microphone unavailable or declined	"Say it yourself" is optional, with a Skip button and a plain-language message
Reviewer expects a Storyline file	Design document and storyboard are tool-agnostic and specify a Storyline production path; voices were produced in Storyline
Fictional customer read as a real account	The module footer states it is a work sample with a fictional customer and household; account numbers on screen are labelled fictional

14. Evaluation plan

Measured the way I measure in practice: against the goals set in this design document. In the module, the call and queue scores show whether the diagnosis skill landed, and the first-attempt scoring keeps those numbers honest. After launch, I would use end-of-module feedback and the questions learners ask afterward to build short follow-up microlearning, such as a second call with a household whose problem really is download or WiFi.

The behavior that matters is discovery questions being used on live calls. I would check for it with team leads and call listening a few weeks after launch, watch repeat-call rates for "slow internet" complaints, and refine the module from what shows up. If the module were deployed, completion and score data from the LMS would be reviewed against those same goals, not reported as an end in themselves.

Sample content created by Martine Brown as a work sample. The customer and household are fictional. Plan details reflect a national fiber provider's public listings as of November 2026.