

Northwind Freight

SAMPLE · FICTIONAL DATA

A record of the AI literacy measure Northwind Freight deployed this period: what was deployed and when, what it is designed to develop, who completed it, and where people's judgement was weakest. For your internal records and audits.

ASSIGNED	COMPLETED	COMPLETION	MEDIAN TIME	SOURCES CHECKED
1,180	1,034	87.6%	24 min	1 December 2026

1 What was deployed

VERSION	LIVE FROM	UNTIL	CHANGES
2026.4	15 Oct 2026	1 Dec 2026	Version first deployed.
2026.4.1	2 Dec 2026	current	Primer 2 wording on approved tools clarified. All sources re-checked 1 Dec 2026. No situation changed, so results from both versions are combined.

Delivered through Northwind Freight's LMS as one SCORM package. Contents: primer, 1 The Thursday shortcut, 2 The confident clause, 3 The easy answer, 4 the current situation (rewritten each quarter). Engaged time target 25 minutes.

2 What the module is designed to develop

LEARNING OUTCOME	ARTICLE 4 THEME (COMMISSION Q&A)	ISO/IEC 42001 7.3	NIST AI RMF
Explain why generative AI produces confident errors, and what that means for checking its output.	General understanding of AI: how it works and what it can and cannot do	7.3 b) Contribution to the effectiveness of the AI management system	GOVERN 2.2: AI risk management training
Decide what information may go into which AI tool, using the organisation's approved-tools list and policy.	The organisation's AI systems and the risks of using them	7.3 a) The AI policy	GOVERN 2.2: AI risk management training
Recognise when to check an AI recommendation before acting on it, and when overriding it is the poorer choice.	Role and context of use, including human oversight	7.3 b) Contribution to the effectiveness of the AI management system	GOVERN 3.2: Roles for human-AI configurations and oversight
Respond to AI-enabled deception by checking the source and channel, not by judging how real something looks.	Risks of the AI in use, including misuse by others	7.3 c) Implications of not conforming	GOVERN 4.1: A critical-thinking, safety-first culture
Know who to tell, and what is worth telling, when something goes wrong with AI at work.	The organisation's role, policies and procedures	7.3 c) Implications of not conforming	GOVERN 4.3: Incident identification and information sharing

3 Completion by group

GROUP (FROM THE LMS)	ASSIGNED	COMPLETED	RATE	MEDIAN TIME
Operations	420	371	88.3%	23 min
Customer service	310	289	93.2%	24 min
Finance	140	118	84.3%	26 min
Sales	190	151	79.5%	21 min
Head office	120	105	87.5%	25 min
All	1,180	1,034	87.6%	24 min

4 Judgement report: where choices were weakest

Aggregate only. No individual's choices are reported; any group under 10 learners is suppressed. These are the 5 decisions where a costly or over-cautious choice was most common, with what each suggests about policy, tools or culture.

- 1 SITUATION 3 · TUESDAY, 10:10

41% released a held payment on a team lead's say-so without opening the assistant's reasons.

What it suggests: Pressure to override the claims assistant is common and goes unchallenged.

Consider: Show the assistant's reasons on the payment screen itself, and ask team leads to record a reason for every override.
- 2 SITUATION 2 · WEDNESDAY, 14:10

36% sent an AI-drafted claim reply without checking the clause it quoted against the agreement, and a further 15% asked the tool to confirm its own draft.

What it suggests: Staff treat a draft that quotes a contract as already checked.

Consider: Add a "checked against the agreement" step to reply templates that quote terms or figures, and sample sent replies each month.
- 3 SITUATION 4 · SATURDAY, 07:42

27% paid a purchase-card request from a voice that sounded like the managing director without calling back on a known number; another 17% relied on a personal question.

What it suggests: Site staff treat a familiar voice as proof of identity.

Consider: Make a call back on a directory number, and a second approver for card payments, the written rule for every urgent phone request.
- 4 SITUATION 1 · FRIDAY, 11:05

34% would put customers' bank details into an approved AI tool because "it's approved".

What it suggests: Staff read "approved" as "approved for everything".

Consider: Print the "never for" list where the tool is launched, not only in the policy.

29% agreed to switch off duplicate checks for top accounts without looking at how often the flags were right.

What it suggests: Appetite to switch controls off rather than tune them.

Consider: Route requests to change AI checks through a short review that looks at the flags' accuracy first.

Chosen actions

At the end, each learner picked one thing to do that week (a chosen action, not evidence of behaviour).

Look up which AI tools Northwind Freight approves, and for which kinds of data.	38%
Save how to reach the Privacy Desk somewhere I can find it in a hurry.	27%
Move one task I do in a personal AI account over to Copilot (work account).	22%
None of these fits my role	13%

5 Every decision, across the organisation

Share of learners choosing each option. n is the number who reached that decision; later decisions depend on earlier choices.

● Strong call ● Partly there ▲ Costly call ◆ Safe, but costly

Situation 1: The Thursday shortcut Data and unapproved AI tools

16:40 · Thursday, 16:40 · $n = 1,034$

How do you get the summary done?

A	Remove names and account numbers, then use Parley. ● PARTLY THERE		21%
B	Do it by hand. Customer data and AI do not mix. ◆ SAFE, BUT COSTLY		9%
C	Use Copilot (work account), after checking it is approved for complaints. ● STRONG CALL		52%
D	Paste the export into the free Parley app, as Sam says. ▲ COSTLY CALL		18%

08:20 · Friday, 08:20 · $n = 403$

What now?

A	Tell the Privacy Desk this morning: what went in, where and when. ● STRONG CALL		58%
B	Ask Sam whether it is a big deal. ● PARTLY THERE		20%
C	Delete the chat. It is done, and nothing bad has happened. ▲ COSTLY CALL		22%

11:05 · Friday, 11:05 · $n = 631$

The sheet holds customers' bank details. What do you do?

A	Refuse, and tell Sam to report themselves for asking. ◆ SAFE, BUT COSTLY		11%
B	Take the bank details out first, then run it. ● STRONG CALL		55%
C	Run it. Copilot (work account) is approved, so it is covered. ▲ COSTLY CALL		34%

09:15 · Monday, 09:15 · $n = 1,034$

What do you do?

A	Report Sam to the Privacy Desk by name, without telling them. ● PARTLY THERE		19%
B	Post a reminder about the AI rules in the team channel. ● PARTLY THERE		12%
C	Nothing. It is Sam's call, not mine. ▲ COSTLY CALL		8%
D	Tell Sam it needs reporting, and show them Copilot (work account). ● STRONG CALL		61%

Situation 2: The confident clause Fluent, confident errors

14:10 · Wednesday, 14:10 · $n = 1,034$

The draft is ready. What do you do with it?

A	Ask Copilot (work account) to confirm that clause 9.3 says this. ▲ COSTLY CALL		15%
B	Send it. It reads well, quotes the agreement and is due today. ▲ COSTLY CALL		36%
C	Write the reply yourself. AI should not touch contracts. ◆ SAFE, BUT COSTLY		8%
D	Open the agreement and the customer record, and check it. ● STRONG CALL		41%

09:30 · Thursday, 09:30 · n = 527

Marlow have found the errors. How do you reply?

A	Forward it to your manager to deal with, and stay out of it. 🟡 PARTLY THERE		15%
B	Say clause 9.3 is in a newer version of the terms. ▲ COSTLY CALL		6%
C	Tell your manager, then reply: the letter was wrong; the limit is 5,000.00. ● STRONG CALL		52%
D	Send a corrected letter without saying the first was wrong. 🟡 PARTLY THERE		27%

14:25 · Wednesday, 14:25 · n = 507

You have checked the draft. What do you do now?

A	Drop the clause number but keep the 1,500.00 offer. ▲ COSTLY CALL		7%
B	Correct it, send it, and tell your manager what Copilot (work account) did. ● STRONG CALL		44%
C	Stop using Copilot (work account) for letters and warn the team off it. ◆ SAFE, BUT COSTLY		11%
D	Correct it from the agreement and send it. Done. 🟡 PARTLY THERE		38%

11:00 · Friday, 11:00 · n = 1,034

How do you handle the summary?

A	Draft it with Copilot (work account), read it once, and post it. ● STRONG CALL		51%
B	Let Copilot (work account) post it to the channel unread. 🟡 PARTLY THERE		14%
C	Draft it, check every line against the notes, and get Jo to approve it. ◆ SAFE, BUT COSTLY		27%
D	Write it by hand. Copilot (work account) cannot be trusted now. ◆ SAFE, BUT COSTLY		8%

Situation 3: The easy answer Oversight, and when not to override

10:10 · Tuesday, 10:10 · n = 1,034

What do you do with the hold?

A	Open the two claims side by side before deciding. ● STRONG CALL		38%
B	Accept the hold and send the standard rejection. 🟡 PARTLY THERE		12%
C	Override it and pay. Priya knows these customers. ▲ COSTLY CALL		41%
D	Keep the hold and refer Brightwell to the fraud team. ◆ SAFE, BUT COSTLY		9%

10:52 · Tuesday, 10:52 · n = 486

The assistant was right about the duplicate. What do you do now?

A	Report Priya to the Privacy Desk for pressuring you to override a control. ◆ SAFE, BUT COSTLY		5%
B	Pay it anyway. Brightwell matter more than 1,240.00. ▲ COSTLY CALL		17%
C	Override it, and record "customer relationship" as the reason. 🟡 PARTLY THERE		12%
D	Keep the hold, and show Brightwell and Priya the earlier payment. ● STRONG CALL		66%

09:05 · Wednesday, 09:05 · n = 124

Brightwell are offended. What do you do?

A	Apologise, and show them CX-5512 was paid on the 3rd. ● STRONG CALL		64%
B	Pay the claim to keep a twelve-year customer happy. ▲ COSTLY CALL		15%
C	Resend the standard wording. The decision was right, so it stands. ▲ COSTLY CALL		21%

09:40 · Thursday, 09:40 · n = 424

Finance wants to know what happened.

A	Quietly ask Brightwell to send it back, and tell finance it's handled.		23%
	● PARTLY THERE		
B	Say the assistant's flag was not clear enough. ▲ COSTLY CALL		14%
C	Explain, ask finance to recover it, and fix your note. ● STRONG CALL		63%

14:00 · Thursday, 14:00 · n = 1,034

What do you say?

A	Say checks should never change. Controls are controls. ◆ SAFE, BUT COSTLY		7%
B	Suggest checking how often last quarter's flags were right. ● STRONG CALL		48%
C	Agree. The flags waste everyone's time on good customers. ▲ COSTLY CALL		29%
D	Stay quiet. It is not your call. ● PARTLY THERE		16%

Situation 4: This quarter: the voice on the line Cloned voices and payment fraud

07:42 · Saturday, 07:42 · n = 1,034

What do you do with the request?

A	Block the number and tell all staff Martin was hacked. ◆ SAFE, BUT COSTLY		10%
B	Pay it. You know his voice, and collections are on Monday. ▲ COSTLY CALL		27%
C	Call Martin back on the number in the company directory. ● STRONG CALL		46%
D	Ask something only Martin would know, and pay if it fits. ● PARTLY THERE		17%

08:30 · Saturday, 08:30 · n = 579

The call was not Martin. What do you do now?

A	Call the number back to tell them you know. ▲ COSTLY CALL		6%
B	Block the number and delete the texts. Nothing was paid. ● PARTLY THERE		24%
C	Get the site card cancelled and ask for requests by email only. ◆ SAFE, BUT COSTLY		12%
D	Report it to the Privacy Desk with the number, times, link and request. ● STRONG CALL		58%

09:10 · Monday, 09:10 · n = 455

Finance want to know what happened. What do you tell them?

A	Offer to repay the 4,850.00 yourself so nobody looks into it. ▲ COSTLY CALL		12%
B	Tell finance and the Privacy Desk everything: the call, number, texts and payment. ● STRONG CALL		61%
C	Explain, but leave out the instruction to keep finance out of it. ● PARTLY THERE		18%
D	Say it looked like a genuine invoice, and leave out the call. ▲ COSTLY CALL		9%

10:00 · Wednesday, 10:00 · n = 1,034

What do you do with the request?

A	Hang up and call him back on the number on file before agreeing. ◆ SAFE, BUT COSTLY		24%
B	Agree, and read him the site card number for a late fee. ▲ COSTLY CALL		15%
C	Agree. It is routine and asks for nothing that could be misused. ● STRONG CALL		46%
D	Refuse, and report the call to the Privacy Desk as a possible scam. ◆ SAFE, BUT COSTLY		15%

6 Deployer addendum

To be completed by Northwind Freight. It places this module within your wider AI literacy measures, suited to your systems and roles.

AI systems in use and who uses them

Roles with higher exposure, and the extra measures for them

Other AI literacy measures this period

Approved tools and policy reference	AI Acceptable Use Policy; Copilot (work account), DeepL Pro
--	---

Reporting contact	The Privacy Desk
--------------------------	------------------

Reviewed by, role and date

7 What this record does and doesn't say

It says

- Which version of the module was deployed, when, and what changed.
- What the module is designed to develop, and how that maps to common frameworks.
- Who completed it, by group, and how long it took.
- Where choices across the organisation were weakest.

It doesn't say

- That Northwind Freight is "compliant" or "certified".
No certificate exists or is required for Article 4.
- That any individual reached a particular level.
Individual choices are not reported.
- That behaviour changed. Chosen actions are intentions, not outcomes.

Sources: Regulation (EU) 2026/1744, Article 1(5) replacing Article 4 of Regulation (EU) 2024/1689; European Commission, AI Literacy Questions & Answers, updated 27 July 2026; ISO/IEC 42001:2023 clause 7.3; NIST AI RMF 1.0. All people, companies and figures in this sample are fictional.