PROJECT 3 OF 5 | START HERE LEARNER GUIDE

Third-Party Risk Decision

Engineer a risk decision

Harborline Health wants to sign a near-final MSA and DPA and go live on NovaBridge Relay, a support-operations SaaS product, for internal IT and People employee support. Implementation has already started. You will assess the evidence, separate findings from risks, judge evidence strength, and recommend a decision a leader can accept.

> **Companion video**
> Watch the walkthrough before you begin: https://youtu.be/354-jZ-MDuI

Your project folder

Item	What it is for
START_HERE_Learner_Guide.pdf	This guide.
Master_Workbook.xlsx	Your main workspace: registers, recommendation, Break, Improve, and Defend tabs.
1_Assess_Evidence	The NovaBridge evidence packet for Stage 1 (eleven files).
3_Build_Source_Data	Structured copies of the packet (questionnaire, evidence index, subprocessors, metadata).
4_Break_Incident_Evidence_Open_After_Build	New facts for Stage 4. Do not open until your initial recommendation exists.
Deliverable_Templates	Word templates: the initial Decision Memo (Stage 3) and the Decision Brief (Stage 6).
Optional_Python_Evidence_Hygiene	Optional Deep Dive: evidence hygiene flags in Python. Never outputs a decision.

Contents

Welcome to Project 3 . . . 3
Stage 1: Assess . . . 6
Stage 2: Design . . . 8
Stage 3: Build (Analyze) . . . 10
Stage 4: Break . . . 12
Stage 5: Improve . . . 14
Stage 6: Defend . . . 15
Stage 7: Career Translation . . . 17
Before you call this finished . . . 18
Optional: Python evidence hygiene (Deep Dive) . . . 19

Welcome to Project 3

Read this once before Stage 1.

The engagement

Organization	Harborline Health (fictional)
Vendor / product	NovaBridge, Inc. / NovaBridge Relay
Review freeze	2026-10-20 12:00 UTC (`2026-10-20T12:00:00Z`). Interpret every date against it.
Stated Day 1 use	Internal IT and People (HR operations) employee support ticketing and knowledge base
Not in Day 1 use	EHR, member or patient chat, production cloud administration
Secondary feature	Optional NovaBrief AI summarization. Not the project topic.
Target go-live	2026-11-17 (sponsor)
Your role	GRC practitioner. You recommend. A named authority accepts residual risk.

Business urgency is real: implementation has started and an alternative vendor would slip a quarter. GRC is not the villain. Urgency does not erase risk.

How to think in this project

- Questionnaires provide claims, not certainty.
- SOC 2 provides scoped assurance, not a universal security certificate.
- Not every finding is a material risk, and not every control gap requires rejection.
- Compensating controls may materially reduce risk. Contract protections are part of treatment.
- Residual risk remains after controls. Assessment confidence may be lower than a tidy risk label suggests.
- The strongest learner is not the one who finds the most issues or rates harshest. It is the one who makes the most defensible decision.

Words you need first

Term	Meaning
Finding / evidence	A fact in the packet: a date, a scope line, a clause, a mismatch.
Risk	The exposure that fact creates for Harborline. A finding is not automatically a risk.
Treatment / condition	What Harborline plans to do: control, compensating control, contract language, time-bound exception, deferral, or rejection.
Inherent risk	Risk from the business use and data **before** taking much comfort from the vendor's controls.
Residual risk	What remains after the treatments you rely on.
Assessment confidence	How much you trust your own risk picture, given incomplete, stale, scoped, or contradictory evidence. Separate from residual risk.
Evidence strength	Strong, Moderate, Weak, or Unknown, judged on independence, recency, relevance, scope, specificity, completeness, and fit to Harborline's use. Never a hidden numeric score.
CUEC	Complementary user entity control: something the customer must operate for the vendor's SOC 2 story to hold.
DPA	Data processing agreement: contractual rules for how the vendor processes Harborline's data.
Fourth party	A vendor's vendor (subprocessor) that may process Harborline's data.

> **Only four recommendation states**
>
> Approve, Approve with conditions, Defer, or Reject. There is no official correct answer. No spreadsheet formula or script decides it for you.

Choose your depth and path

Option	Time	What it includes
Standard	About 4 to 6 hours	Full loop, registers, initial recommendation, Break, Decision Brief, Career Translation.
Deep Dive	About 8 to 12 hours	Standard plus a fuller register, deeper contract and CUEC work, both Build paths, optional Python hygiene, and interview defense.

Build path (Stage 3)	Where	Status
Decision Memo (recommended)	**Deliverable_Templates/3_Build_Decision_Memo_Initial.docx**	Complete. Can reach the top of the rubric.
Spreadsheet	Tabs **3 Decision Matrix** and **Conditions**	Complete. Equal, not lesser.
Python evidence hygiene	**Optional_Python_Evidence_Hygiene**	Optional Deep Dive only. Never replaces the memo or spreadsheet.

Everyone completes the Design registers in the workbook and the **3 Initial Recommendation** checkpoint.

What you will produce

#	Artifact	Where it lives
1	Current-state assessment	**1 Assess**
2	Inherent risk profile	**2A Inherent Risk**
3	Evidence Evaluation Framework	**2B Evidence Strength**
4	Findings and Questions Register	**2C Findings Register**
5	Third-Party Risk Register	**2D Risk Register**
6	Initial recommendation	**3 Initial Recommendation** plus **Deliverable_Templates/3_Build_Decision_Memo_Initial.docx** or **3 Decision Matrix**
7	Break investigation	**4 Break**
8	Improve: what changed and what did not	**5 Improve**
9	Exactly three pre-launch conditions plus deferred residual risk	**6 Pre-Launch Conditions**
10	Third-Party Risk Decision Brief	**Deliverable_Templates/6_Defend_Decision_Brief.docx**
11	Career Translation (Standard set)	**7 Career Translation**

Stage 1: Assess

Read the NovaBridge packet and classify what you see.

What you are trying to accomplish

Label packet facts without collapsing them into a single "vendor risk" blob, and identify what the packet cannot tell you.

Why this matters

Assess is where you practise keeping finding language (what the packet shows) apart from risk language (what it might mean for Harborline). It is also where you record at least one thing that looks loud but is not a Harborline risk.

Open these

What	Where to find it
The eleven evidence files	**1_Assess_Evidence**
Structured copies	Tabs **Business_Context**, **Questionnaire**, **Evidence_Index**, **Subprocessors**
Your assessment table	**Master_Workbook.xlsx**, tab **1 Assess**

File	What it is
1_Assess_Evidence/Sponsor_Request.pdf	Why Support wants a decision now
1_Assess_Evidence/Vendor_Profile.pdf	Who NovaBridge is (sales snapshot)
1_Assess_Evidence/Architecture_Data_Flow.pdf	How Relay is described to work
1_Assess_Evidence/Questionnaire_Response.pdf	Vendor claims
1_Assess_Evidence/SOC2_Extract.pdf	Independent assurance extract
1_Assess_Evidence/Pentest_Summary.pdf	Independent testing summary
1_Assess_Evidence/Privacy_One_Pager.pdf	Vendor privacy language
1_Assess_Evidence/Contract_DPA_Excerpts.pdf	Draft MSA and DPA clauses
1_Assess_Evidence/Subprocessor_List.pdf	Vendor-supplied subprocessors
1_Assess_Evidence/Engineering_Notes.pdf	Harborline launch identity model
1_Assess_Evidence/Evidence_Index.pdf	Dates and sources

The seven labels (pick one per row)

Label	Use it when
Observation	A fact. Not yet a judgment.
Claim vs evidence	A questionnaire or sales statement that does not match another artifact, or that you cannot verify.
Control gap	A control is missing, delayed, or not in the launch design.
Assurance gap	Independent evidence is missing, stale, out of scope, or carved out.
Contract / legal issue	A clause, liability, audit-right, or deletion term that may not match the claim or the use.
True negative / not relevant	Looks loud but does not apply to the stated use, or a control the packet actually supports.
Unknown	You cannot tell yet. Say what artifact would let you tell.

Do this

1. Read every file in **1_Assess_Evidence**. Start with **1_Assess_Evidence/Sponsor_Request.pdf** and **1_Assess_Evidence/Evidence_Index.pdf**.
2. Use the structured tabs to compare claims and dates quickly: **Questionnaire** (vendor claims), **Evidence_Index** (artifact dates), **Subprocessors** (subprocessors as of 2026-06-12).
3. On **1 Assess**, add at least ten rows: the packet fact, one label, the possible exposure for Harborline (or "none / not yet"), and what you would still want to know. Leave the exposure blank for true negatives and pure observations.
4. Use more than two labels. Mark at least one true negative or not-relevant item and name at least one unknown.
5. Write the two sentences: what worries you most for the **stated** Day 1 use, and what you still cannot know.
6. Tick the four self-checks at the bottom of the tab.

What you produce	You are done when
• At least ten labeled rows, two sentences, and self-checks on **1 Assess**.	☐ At least ten rows and more than two labels. ☐ At least one true negative and one unknown. ☐ You did not treat every questionnaire Yes as proven or every gap as a launch blocker.

Next: Move to Stage 2: Design.

Stage 2: Design

Build the structures you will use to decide.

What you are trying to accomplish

Create an inherent-risk profile, an evidence-strength evaluation, a findings and questions register, and a risk register that keeps finding, risk, treatment, residual risk, and assessment confidence in separate fields.

Why this matters

You are not scoring the vendor yet. You are making sure Finding, Risk, and Treatment cannot collapse into one concept, and that evidence strength is judged on purpose.

Open these

What	Where to find it
Inherent risk profile and evidence strength	Tabs **2A Inherent Risk** and **2B Evidence Strength**
Findings register and risk register	Tabs **2C Findings Register** and **2D Risk Register**

Rating rules

Risk: High, Medium, or Low (a short range is fine). Evidence strength: Strong, Moderate, Weak, or Unknown. Assessment confidence: free text such as "Low to Moderate". No hidden score.

Do this

1. On **2A Inherent Risk**, describe the stated use in your own words, rate each factor High, Medium, or Low with a reason, and write a 3 to 6 sentence inherent-risk statement for the **stated Day 1 use**. There is no required rating.
2. On **2B Evidence Strength**, rate each listed artifact Strong, Moderate, Weak, or Unknown, and say what claim it can actually support. Record disagreements between artifacts as findings. Do not average them away.
3. On **2C Findings Register**, log findings (cite artifact and date), open questions (what evidence would answer each, and who could produce it), and true negatives.

4 On **2D Risk Register**, add material risks only. Fill every field. At least one material control gap must evaluate a compensating control and say why it is or is not sufficient for a time-bound launch. The launch identity and SSO timing in **1_Assess_Evidence/Engineering_Notes.pdf** is the obvious candidate.

5 Stay on the Assess packet. Do not invent Break facts.

What you produce	You are done when
• Completed **2A Inherent Risk**, **2B Evidence Strength**, **2C Findings Register**, and **2D Risk Register**.	☐ Finding, risk, and treatment sit in separate fields. ☐ Residual risk and assessment confidence are separate values. ☐ At least one compensating-control judgment is written. ☐ Open questions are still questions, not decisions.

Next: Move to Stage 3: Build (Analyze).

Stage 3: Build (Analyze)

Write an initial recommendation before anything changes.

What you are trying to accomplish

Produce a defensible recommendation in one of four states, tied to packet facts, with inherent risk, residual risk, assessment confidence, evidence-strength notes, and named owners. Then stop.

Why this matters

In this project, Analyze is the Build stage: you build a decision, not a TPRM product. Making a decision **before** Break lets you show later what changed and why.

Open these

What	Where to find it
Initial recommendation checkpoint (everyone)	Tab **3 Initial Recommendation**
Decision Memo path (recommended)	**Deliverable_Templates/3_Build_Decision_Memo_Initial.docx**
Spreadsheet path (equal)	Tabs **3 Decision Matrix** and **Conditions**
Your Design work	**2A Inherent Risk**, **2B Evidence Strength**, **2D Risk Register**
Optional Python hygiene	**Optional_Python_Evidence_Hygiene/README.md**

The quality bar

There is no official correct state. Unconditional Approve is hard to defend on this packet. Reject is allowed if you argue it. Approve with conditions and Defer are both available. Write the one you can defend.

Do this

1. Choose at least one complete path: the Decision Memo or the spreadsheet. You may do both.
2. **Memo path:** open **Deliverable_Templates/3_Build_Decision_Memo_Initial.docx**, save a copy, and complete every section: what you were asked to decide, stated use and data, evidence and strength, findings then risks then treatments, the compensating-control judgment, contractual treatment, the recommendation and why it is defensible, conditions sketched, residual risk and who would accept it, and what would change it.
3. **Spreadsheet path:** on **3 Decision Matrix**, complete all four option rows, mark exactly one as selected, and explain why it is defensible and why not the others. On **Conditions**, list candidate conditions as a working list (more than three is fine here).
4. **Everyone:** complete **3 Initial Recommendation** even if your memo already covers it. It is the freeze-point decision you will compare after Break.

5 **Deep Dive (optional):** run the Python evidence hygiene script (end of this guide). It flags stale dates and claim mismatches. It must never output a decision.

6 Stop. Do not open the Break folder until **3 Initial Recommendation** is complete.

What you produce	You are done when
• A completed **3 Initial Recommendation**. • A completed memo (**Deliverable_Templates/ 3_Build_Decision_Memo_Initial.docx**) or **3 Decision Matrix** and **Conditions**.	☐ Your recommendation is one of the four states. ☐ Inherent risk, residual risk, and assessment confidence are all stated, and confidence is not a restatement of residual risk. ☐ At least three material issues are written as finding, risk, treatment. ☐ You named a risk owner and an acceptance authority (a role, not "GRC generally").

Next: Only now open the Break evidence folder and move to Stage 4.

Stage 4: Break

New facts arrive after your recommendation.

What you are trying to accomplish

Decide what the new facts mean for data classification, inherent risk, privacy, fourth-party exposure, your recommendation, and who should accept the risk.

Why this matters

Relationships change after the first review. The packet gives you facts, not conclusions. The skill is re-asking the right questions without assuming the first recommendation is automatically void.

Open these

What	Where to find it
Break evidence (four files)	**4_Break_Incident_Evidence_Open_After_Build**
Your Break worksheet	Tab **4 Break**
Your initial position	**3 Initial Recommendation**

Open this evidence only after your initial recommendation

Do not open **4_Break_Incident_Evidence_Open_After_Build** until **3 Initial Recommendation** is complete.

File	What it is
4_Break_Incident_Evidence_Open_After_Build/Update_From_Support.pdf	New request from Support leadership
4_Break_Incident_Evidence_Open_After_Build/Sample_Ticket_Fields.pdf	Fields Support wants in Relay
4_Break_Incident_Evidence_Open_After_Build/Novabridge_Trust_Notice.pdf	Vendor notice dated 2026-09-28
4_Break_Incident_Evidence_Open_After_Build/Data_Location_Note.pdf	Processing location in that notice

Break work is structured, so it lives entirely in the workbook. There is no separate Word file for this stage.

Do this

1. Confirm **3 Initial Recommendation** is complete.
2. Read the four Break files, starting with **4_Break_Incident_Evidence_Open_After_Build/Update_From_Support.pdf**.
3. On **4 Break**, summarize what is new, in your own words, for the Support request and the NovaBridge notice.
4. Answer the six determination questions with your determination, the facts you are using, and what is still unknown.
5. If you promote anything new, write it as finding, risk, and treatment.
6. Record the recommendation movement: initial state, state after Break (one of the four, or the same state with different conditions), and why that is defensible.

What you produce	You are done when
• A completed **4 Break** tab.	☐ You can explain what is new and what is still unknown. ☐ You decided how (if at all) the recommendation moves, without quoting a hidden correct outcome.

Next: Move to Stage 5: Improve.

Stage 5: Improve

Revise the analysis without pretending the first answer was sacred.

What you are trying to accomplish

Record what changed and why, what did not change and why, and update the register, recommendation, conditions, evidence requests, contract asks, monitoring, and authority.

Why this matters

Changing everything without a reason is as weak as changing nothing. A reader should see the before and after, and residual risk should still be named rather than implied to be zero.

Open these

What	Where to find it
Your improve record	Tab **5 Improve**
Risk register	**2D Risk Register**
Conditions working list	**Conditions**

Do this

1. On **5 Improve**, answer **What changed, and why?** (evidence, use case, fourth parties, inherent risk, residual risk, confidence, recommendation, authority).
2. Answer **What did not change, and why?** (for example, conclusions that still hold for the original employee-support path).
3. Add or amend rows on **2D Risk Register** and log each change in the register-updates table on **5 Improve**.
4. Update **Conditions**: mark each condition Pre-launch or Deferred and give it an owner. It may still have more than three rows.
5. Write evidence requests, contract and DPA asks, monitoring and reassessment after launch, and the risk owner and acceptance authority now.

What you produce	You are done when
• Completed **5 Improve**, updated **2D Risk Register** and **Conditions**.	☐ A reader can see before and after without watching you work. ☐ Residual risk is still named.

Next: Move to Stage 6: Defend.

Stage 6: Defend

Choose exactly three pre-launch conditions and brief the acceptance authority.

What you are trying to accomplish

Under a 30-day-class launch window, commit to exactly three pre-launch conditions, defer everything else with residual risk named, and write a one-page Decision Brief.

Why this matters

A list of ten blockers is not a decision. Leaders need three things that must be true before go-live, an honest deferred list, and a named person who accepts what remains.

Open these

What	Where to find it
Conditions, deferred items, and authority	Tab **6 Pre-Launch Conditions**
One-page Decision Brief	**Deliverable_Templates/6_Defend_Decision_Brief.docx**

Condition ideas (a menu, not a key; there is no official trio)

SSO and MFA compensating package; DPA addendum; data-classification boundary; retest; HelioInfer assurance packet; privileged-support restriction; incident timelines; exit and deletion test; scheduled reassessment.

Do this

1. On **6 Pre-Launch Conditions**, write exactly three specific, testable pre-launch conditions with the related risk, owner, and due date. Extra ideas go in the Deferred table.
2. For each deferred issue, state the residual risk, why it is not pre-launch, and the follow-up or reassessment trigger.
3. Complete the owner-versus-authority section: risk owner and why, acceptance authority and why it fits after Break, and what GRC would recommend if that authority declined.

4 Open **Deliverable_Templates/6_Defend_Decision_Brief.docx**, save a copy, and complete it on one page. Include both residual risk and assessment confidence. Point to your three conditions; do not list a fourth.

5 Read the brief aloud; it should take about five minutes.

What you produce	You are done when
• Completed **6 Pre-Launch Conditions** and your copy of **Deliverable_Templates/6_Defend_Decision_Brief.docx**.	☐ Exactly three pre-launch conditions. ☐ Deferred residual risk is explicit. ☐ Residual risk and assessment confidence both appear and are not used as substitutes. ☐ An authority is named.

Next: Move to Stage 7: Career Translation.

Stage 7: Career Translation

Turn the decision into an honest professional story.

What you are trying to accomplish

Write the Standard career set: honesty sentence, portfolio description, resume bullets, 30-second explanation, and a STAR story.

Why this matters

Interviewers test whether you can explain finding versus risk, SOC 2 scope, and residual risk versus confidence. Harborline is fictional: never describe it as an employer or client.

Open these

What	Where to find it
Career prompts and writing areas	Tab **7 Career Translation**

Do this

1. Say the honesty sentence: *"This was an independent portfolio project where I simulated a third-party risk decision for a fictional company. I assessed vendor evidence, designed registers that separate findings from risks from treatments, recommended a decision, investigated a change in use and subprocessors, improved the analysis, and wrote a decision brief with residual risk, assessment confidence, and named acceptance authority."*
2. Complete the Standard set on **7 Career Translation**. Ratings and counts must match your own analysis.
3. **Deep Dive (optional):** README paragraph, skills, the ten interview questions, and the defense drill. If asked for a vendor score out of 100, practise explaining why you will not give one.

What you produce	You are done when
• The Standard career set on **7 Career Translation**.	☐ You can say the honesty sentence and STAR without reading. ☐ No claims such as "I built a TPRM platform" or "passed the vendor because they had SOC 2".

Next: Complete the final checklist.

Before you call this finished

A final check across all seven stages.

- ☐ Assessment
- ☐ Inherent risk profile
- ☐ Evidence Evaluation Framework
- ☐ Findings and Questions Register
- ☐ Third-Party Risk Register
- ☐ Initial recommendation (memo or spreadsheet, plus checkpoint tab)
- ☐ Break worksheet
- ☐ Improve record
- ☐ Exactly three pre-launch conditions plus deferred risk
- ☐ Decision Brief
- ☐ Career Translation Standard set

Reasoning questions

- Can I explain why I chose my recommendation state, and which alternative I rejected?
- What evidence did I trust (independence, recency, relevance) and what did I not trust?
- After Break, did I keep finding, risk, and treatment distinct?
- How is my residual risk different from my assessment confidence?
- What would cause me to reassess this vendor decision?
- Can I defend this to a sponsor who wants to launch this week?

Optional: Python evidence hygiene (Deep Dive)

Flags for human analysis. Never a vendor decision.

This optional script loads the structured files and flags evidence-hygiene issues: artifact dates older than a threshold **you document**, and claim-versus-artifact mismatches such as an "annual testing" claim against the pen-test date. It writes **Optional_Python_Evidence_Hygiene/outputs/evidence_hygiene.json**.

> **Hard rule**
> The script must not output Approve, Approve with conditions, Defer, or Reject, a residual-risk rating, or a recommendation. Forbidden keys anywhere in the output: recommendation, decision, verdict, approve, defer, reject, residual_risk. The shape test checks this.

Files: **src/hygiene.py** (TODOs), **expected_output_shape.md** (output shape), and **tests/test_hygiene_shape.py** (shape and guardrail checks, not an answer key).

Python setup (one time)

1. Install Python 3.10 or newer from https://www.python.org/downloads/. On Windows, tick **Add Python to PATH** during setup.
2. Open a terminal. On Windows use PowerShell or Command Prompt. On macOS use Terminal.
3. Check the version: `python3 --version` (on Windows the command is often `python --version`). You want 3.10 or higher.
4. Change into the optional Python folder inside this project. The easiest way: type `cd ` followed by a space, then drag the **Optional_Python_Evidence_Hygiene** folder onto the terminal window and press Enter.
5. Nothing needs to be installed. The exercises use the Python standard library only.

Run commands

```
python3 src/hygiene.py
python3 tests/test_hygiene_shape.py
```

On Windows, type `python` wherever these commands say `python3`. Scripts find the data folder on their own, so you can run them from any location as long as the project folder structure is unchanged. Output is written to the **Optional_Python_Evidence_Hygiene/outputs** folder, which the scripts create if it is missing.

What you see	What it means
`NotImplementedError`	Expected. You still have a TODO in that function. Read the HINT lines next to it.
`python3: command not found`	Try `python`. If neither works, install Python and reopen the terminal.
`FileNotFoundError ... outputs/0N_...json`	Run the earlier numbered step first. Each step reads the previous step's output.
Wrong SLA or freshness results	Use the review freeze timestamp from this guide, never your computer clock.