

AI Agents · Workflows · Data Protection

MODULE 3 — VIRTUAL FINANCE TEAM


3 AI Roles. 1 Finance Agent. Automated.

3-5 hours

Individual Coaching

Live Agent Demo

Note:

 This training is conducted by an AI system as a trainer (e.g., ChatGPT, Gemini, Claude) – load both Markdown files (Systemprompt.md and content.md), then upload them and enter “Start”. No human trainer required.

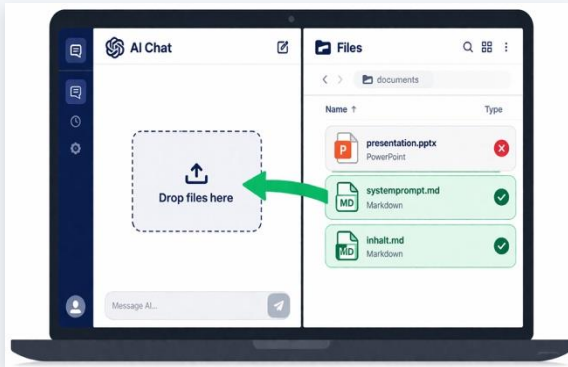
 Download both files: <https://www.foundic.org/category/trainings/>



FOUNDIC.org



1 Upload Files to Chat

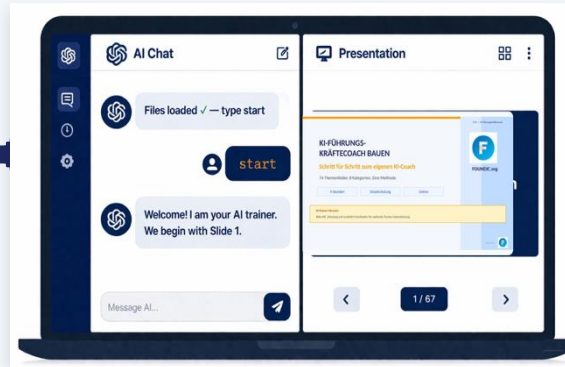


- ① **Open your LLM in the browser**
ChatGPT, Claude, or Gemini — Text Mode (no audio!)
- ② **Upload Both Markdown Files**
systemprompt.md + content.md via Drag & Drop
- ❌ **Do Not Upload the PowerPoint**
Only the two .md files belong in the chat

⚠️ No Audio Yet — Files Only Loadable in Text Mode

📁 Files Missing? Download: foundic.org/trainings

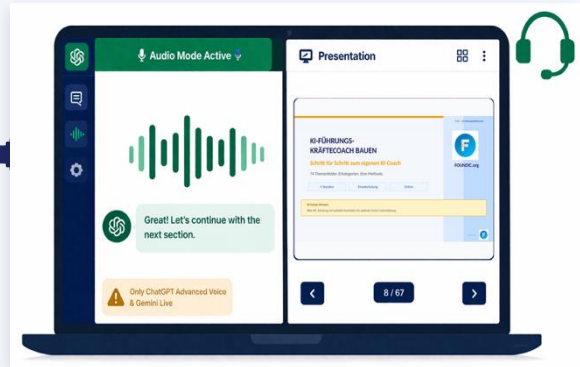
2 Begin Training with "Start"



- ① **Type the Word: Start**
A single word is enough — the trainer starts automatically
- ② **Split Your Screen**
Left: Chat window · Right: Open PowerPoint slides

✓ **The Slides Are Your Visual Anchor — the Trainer Tells You When to Advance.**

3 Optional: Enable Audio Mode



- ① **Complete Steps 1 & 2 First**
Upload Files + Type "Start" — Then Audio
- ② **Enable Audio Mode**
ChatGPT: Headphones icon · Gemini: Microphone icon
- ③ **Put On Headset & Speak**
Review Slides, Speak Freely with the Trainer

Only available in ChatGPT (Advanced Voice) and Gemini Live — not in Claude.

Transfer-Check Module 2 — What Have You Implemented?

Transfer-Check Module 2

Which PowerQuery pipeline did you build?
 How much time per month do you save?
 → Brief Report — 2 Minutes
 → Celebrate Success — Then Next Step

What went well?


What Time Savings Have You Achieved?
 → Specific Number: "X Minutes for Y Task"
 → Note for Quick-Win Plan
 → What Was More Difficult Than Expected?

Kick-Off Check Module 3

Which Manual Finance Process Costs You the Most Time?
 → Today We Build an Agent Workflow for It.
 → Trainer Listens Actively — Note Examples

Your Automation Challenge Today

Brief: Which Process Would You Like to Automate?
 → Will Be Addressed in the Building Session.
 → No Problem Too Large or Too Small.

 Files Ready? automation_readiness_scoring.xlsx + Northern Light Corporation Dataset.xlsx

Agenda — What to Expect Today

15 min	Transfer-Check M2 + Kick-Off Check	Kick-Off
10 min	Input: Chat to Agents — Paradigm Shift	Input
15 min	Agentic AI: Roles + Finance Use Cases	Input
15 min	LIVE DEMO: 3-Agent Cascade	Demo
20 min	EXERCISE 1: Processes + Workflow Design	Exercise
10 min	LEARNING CHECK 1 (6 Multiple-Choice)	Check
10 min	<i>Break</i>	Break
10 min	Automation Stack: 4 Levels + Tool Comparison	Input
10 min	Guardrails, HITL Gates + Audit Trail	Input
15 min	Agent Prompt Design + Cascade Patterns	Input
20 min	EXERCISE 2: Build Multi-Agent Workflow	Exercise
10 min	LEARNING CHECK 2 (6 Multiple-Choice)	Check
10 min	<i>Break</i>	Break
40 min	BUILDING SESSION: Your Finance Agent	Hands-on
15 min	Use Cases + Mini-Case	Analysis
10 min	Final Check + Key Takeaways + Transfer	Conclusion

2–4 hours. 2 learning checks. 1 building session. 1 virtual finance team — ready to take home.



From Answer to Action — What AI Agents Do Differently

AI Chat (Module 1)

You Ask a Question — AI Responds Once.

No memory, no plan, no tools.

Output: Text. Action: none.

AI Agent (Module 3)

You set a GOAL — Agent plans and acts.

Uses tools (Excel, API, email, web search).

Output: Result + Action + Log.

The Difference in Practice

Chat:

"Write me a variance commentary for EBIT -8%."

Agent:

Loads ERP data automatically, calculates variances, writes commentary, has it reviewed, posts it into the reporting template — all without your involvement.

Agentic AI — 3 Characteristics That Change Everything



Goal Orientation

Agent receives a goal, not an instruction. The agent determines how to achieve the goal autonomously.



Tool Usage

Agent can use external systems: Excel API, Email, Web Search, ERP queries.





Iterative Planning

Agent checks its output, corrects errors and repeats steps until the goal is achieved.

Finance example: Agent receives goal "Create the monthly variance report". It loads ERP data → calculates delta → writes commentary → sends for approval. All without your intervention.

Where AI Agents Directly Help in Finance Daily Work

 <p>Monthly Report Automation Today: 6-8 hrs./month manual</p>	<p>With Agent: 45 min. + Agent Review</p>	<p>Savings: ~6 hrs.</p>
 <p>Variance Analysis + Commentary Today: 3-4 hrs./month</p>	<p>With Agent: 20 min. + HITL</p>	<p>Savings: ~3 hrs.</p>
 <p>Anomaly Detection Journal Entries Today: 2-5 hrs./month</p>	<p>With Agent: 10 min. + Approval</p>	<p>Savings: ~4 hrs.</p>
 <p>Stakeholder Reporting via Email Today: 2-3 hrs./month</p>	<p>With Agent: 5 min. + Approval</p>	<p>Savings: ~2.5 hrs.</p>

 Total conservative: 13-26 hrs./month → with Agent Workflows: ~2-3 hrs./month. That is 3-5 hrs./week for strategy instead of routine.

5 Roles — Your AI-Powered Finance Team



Analyst Agent

Loads data, calculates KPIs, identifies variances



Reviewer Agent

Checks Analyst output for plausibility and format



Compliance Agent

Checks all output for GDPR, EU AI Act, and internal policies



Finance Proxy Agent

Formulates management-appropriate commentary for board



Escalation Rule

Human HITL Gate: Controller checks + approves upon threshold breach

A virtual team needs no vacation request. It works in parallel, 24/7, and escalates when necessary.

Analyst Agent + Reviewer Agent — The Core Duo

Analyst Agent — Scope & Prompt Elements

ROLE: Finance Analyst with access to ERP data and KPIs.

TASK: Load [Source], calculate [KPI], identify top 3 variances.

OUTPUT FORMAT: Table with delta, % variance, trend, and 1-sentence comment.

GUARDRAIL: No personal data. Aggregated data only. GDPR compliant.

Reviewer Agent — Scope & Prompt Elements

ROLE: Experienced Controller who checks finance reports for quality.

TASK: Check Analyst output for: Completeness, Plausibility, Format.

OUTPUT FORMAT: APPROVED or REVISE with specific correction guidance.

ESCALATION: If Delta > 15% or unknown position: always trigger HITL gate.

Cascade: Analyst Output → Reviewer Input → Reviewer Output → HITL Gate → Approval → Final Product

Compliance Agent + Finance Proxy — Security + Language

Compliance Agent — What It Checks

GDPR: Does output contain personal data? Name, email, account number? → STOP.

EU AI Act Art. 4: AI literacy proof present? Audit trail complete? → LOG.

Internal policy: Does output violate finance guidelines? → FLAG + HITL escalation.

Amount threshold: Is an action triggered that affects amounts > EUR 10,000? → HITL MANDATORY.

Finance Proxy Agent — Management Language Automatic

ROLE: Experienced CFO advisor, writes for C-level audience.

TASK: Articulate the analyst output in 3 paragraphs: Trend, risk, recommendation.

STYLE: Management-ready, active, concrete, no jargon without explanation.

OUTPUT: Ready to copy directly into board presentation or board report.

IMPORTANT: Compliance Agent can stop a workflow. Finance Proxy runs AFTER compliance clearance. Never simultaneously.

When must humans intervene? — 3 HITL Scenarios



Scenario 1: Amount Threshold

Agent wants to trigger an action affecting amounts > EUR X (payment, booking, approval).
Rule: ALWAYS HITL Gate. No exceptions. Amount = your internal limit.



Scenario 2: Uncertainty > Threshold

Agent has confidence < 80% or deviation > 15% without a clear explainable cause.
Rule: Set flag, hold result, inform controller.



Scenario 3: External Recipient

Agent wants to send output to external recipients (customers, regulators, board).
Rule: ALWAYS human review before sending. No automation without approval step.

LIVE DEMO — I will now show you step by step the 3-agent cascade

1

1. Data Foundation:

Open northern_light_corporation_dataset.xlsx, tab Revenue_Q3. Aggregated export: 6 months, no personal data.

2

2. Analyst-Prompt:

ChatGPT: "You are a finance analyst. Analyze this revenue export: [insert table]. Calculate EBIT delta, top-3 deviations, 1-sentence comment per line item."

3

3. Reviewer-Prompt:

ChatGPT: "You are an experienced controller. Review this analyst output: [insert output]. Is it complete, plausible, management-ready? Give APPROVAL or REVISION."

4

4. HITL Gate:

Controller reads reviewer output, checks numbers (P-Q-R), gives approval. Then proceed to Finance Proxy.

5

5. Finance-Proxy:

ChatGPT: "You are a CFO advisor. Formulate this review in 3 paragraphs: Trend, risk, recommendation. Target audience: Board."



Demo Result — P-Q-R Check of Agent Output

Analyst Output (Example)

EBIT Delta: -8.3% vs. prior year (Plan: -5%). Main drivers: Energy costs +18%, Personnel costs +4%.

Top 3: (1) Energy +156T EUR (2) Material costs +89T EUR (3) Travel costs -12T EUR (savings).

Reviewer Output

APPROVED. Output complete. Delta calculation correct. Management language OK.

Note: Travel costs are positive (savings) — highlight as quick win in CFO report.

P-Q-R Check of the Agent Output (MANDATORY before Approval)

P — Plausibility: Is the figure -8.3% correct? ERP cross-check: Yes. Trend consistent: Yes. No outliers.

Q — Quality: Complete? Yes. Table format? OK. Technically correct? Yes. Management-appropriate? Yes.

R — Review: Human reviewed and signed: B. Mueller, 15.05.2025. Ready for board.

EXERCISE 1 (20 min.) — Which of your processes are suitable for agents?

Task: Check process suitability

1. Open automation_readiness_scoring.xlsx, tab Process-Check.
2. List 5 recurring finance tasks from your daily work.
3. Rate each according to 4 criteria (generic suitability check — Module 4 includes an in-depth 5-criteria AP check):
Repetitive / Rule-based / Data-driven / Verifiable.
4. Calculate suitability score (max. 4 points) — all 4 = ideal.
5. BONUS: Draft an Analyst Agent prompt for your top candidates.

Reference guide: Suitability Criteria (generic, 4 points — applies to all AI use cases)

Repetitive: Does this task occur at least monthly?

Rule-based: Are there clear rules / formulas / standards?

Data-driven: Is the task based on structured data?

Verifiable: Can a reviewer agent verify the output for accuracy?

Example Results

Monthly Commentary: **4/4 — Ideal for Agents!**

Tax Return: **2/4 — Borderline**

Budget Alert Email: **3/4 — Good fit**



Learning progress check 1 — 6 questions | 1 correct answer each | Trainer discusses

1. What characterizes an AI agent vs. chat LLM?

- A) Larger parameter count
- B) Autonomous planning, tools, goals — multi-step
- C) Faster response time
- D) Lower-cost license

2. Which role checks the analyst output for plausibility?

- A) Compliance Agent
- B) Finance-Proxy
- C) Reviewer Agent
- D) HITL Gate

3. What is a HITL gate?

- A) New AI model
- B) Human approval point in the workflow
- C) Automation software
- D) Data protection standard

4. When must a HITL gate be triggered?

- A) Only in annual reports
- B) For irreversible actions: payment, external sending
- C) Never — AI is reliable
- D) Only for >50,000 EUR

5. What is the suitability score 4/4 for agents?

- A) Very expensive
- B) Repetitive + Rule-based + Data-driven + Verifiable
- C) Not automatable
- D) Only for ERP systems

6. Which data must NOT be used in cloud AI?

- A) Aggregated monthly revenue
- B) MAPE values
- C) Account numbers + personal data
- D) Budget totals without personal reference

Short break: Answer all 6 questions — then we will discuss together.



PAUSE



Screen off — take a short break — back in 10 minutes

After the break: Automation Stack — Which tool for which finance process?

Thought exercise: Which of your processes from the exercise would you automate today?

The Automation Stack — 4 Levels at a Glance

4

Level 4: Agentic AI

LLM-driven agent cascades. Multi-step, autonomous planning, feedback loop.

Make AI + n8n AI + LangChain

3

Level 3: Workflow Automation

Visual no-code workflows. If-then rules, API connections, notifications.

Make.com, n8n, Zapier

2

Level 2: API Integration

Direct connection of systems via APIs. Connect ERP + CRM + AI tool.

REST-API, SAP OData, Graph API

1

Level 1: RPA

Rule-based automation via UI. Click macros, screen scraping, RPA bots.

UiPath, Power Automate Desktop



Tool Comparison — Which Tool for Finance Controllers?

Criterion	Make.com	n8n	LangChain
Target Audience	Business users / Controllers	Tech-savvy User	Developer
Code Required?	No — 100% no-code	Minimal (JSON)	Yes (Python)
AI Agents	Native AI Steps (Claude, GPT)	AI nodes available	Complete Framework
Data Protection / On-Premise	Cloud (EU-Server)	Self-host possible	Self-host possible
Cost (Paid)	From 9 EUR/month	Approx. 20 EUR/month (Cloud) or self-host	Open Source (API costs extra)
Free Tier?	✅ 1,000 Ops/Month Sufficient for Getting Started	✅ Open Source Self-Host = free	✅ Free API calls: approx. €0.01–0.03 per 1k tokens
Finance Use Case	Ideal for Controllers	Great for IT-Savvy Teams	For Developer Teams
Recommendation	★★★ RECOMMENDED	★★☆ Good	★☆☆ Developer

Which tool fits which finance process?

Process	IT required?	Volume	Level 3+?	Recommendation
Automate monthly commentary	No	medium	✓	Make.com + ChatGPT
Fetch ERP export via API	Yes	high	—	Make.com + REST-API
Anomaly scanning journal entries	No	high	✓	Make.com AI + Excel
Generate stakeholder emails	No	low	✓	Make.com + GPT-4o
Compliance report to regulatory authority	Yes	low	✓	n8n Self-Host + GPT
Dashboard refresh from ERP	Yes	high	—	Power Automate + API

4 Guardrail Patterns That Secure Every Finance Agent



Output-Validation

Agent validates its own output against defined rules BEFORE passing it on.

Example: "Does the output contain a figure outside the plausible range? STOP."



Data-Masking

Agent input is automatically cleaned: no account numbers, no names, no address data.

Example: Regex filter in the Make.com step before the AI call.



Confidence Threshold

Agent triggers HITL gate when confidence score < 80% or variance > defined threshold.

Example: Evaluate GPT response with confidence tag.



Audit Trail (EU AI Act)

Every agent call is logged: timestamp, input hash, output hash, user ID, approval.

EU AI Act Art. 4+26 requires this for high-risk systems.

Implementing HITL Gates Correctly — 3 Implementation Patterns

Pattern 1: Approval Step in Make.com

Workflow pauses after agent output.

Make sends email to controller with 2 buttons: APPROVE / REJECT.

Controller clicks button — Make.com continues workflow or cancels it.

Setup Time: Approx. 30 Minutes. No Code Required.

Pattern 2: Review Form

Workflow generates output + PDF summary.

Controller receives link to review form (Google Form or Typeform).

Form response triggers next workflow step.

Advantage: Revision comment can be entered directly.

Pattern 3: Escalation via Slack / Teams Message

Workflow sends critical output to dedicated Slack/Teams channel.

Controller writes "OK" or "STOP" as a reply.

Make.com monitors the channel and responds to the reply.

Advantage: Approve on mobile, no separate tool needed.

EU AI Act 2025 — What Finance Controllers Need to Know

Art. 4	AI Competency Requirement 10B981
Art. 13	Transparency Requirement 3B82F6
Art. 26	HITL Requirement (High-Risk) 8B5CF6
Audit Trail	Practical Implementation F59E0B

GDPR Rules for AI Agents in Finance Daily Work

🚫 NOT PERMITTED in Cloud AI (ChatGPT, Claude, etc.)

Personal data: name, email, IBAN, tax ID, account number.

Employee data: salary, performance reviews, attendance.

Customer data: contract content with personal reference, creditworthiness data.

🟡 Gray Area — only with contractual safeguards

Aggregated financial data without personal reference (e.g. revenue by product group) — usually acceptable.

Anonymized benchmarks and KPI reports — usually acceptable, have reviewed.

Emails to business partners — only if a data processing agreement with the AI provider exists.

🟢 Safe Without Restriction

Public market data (stock prices, industry reports, interest rates).

Internal process descriptions and workflow documentation without personal reference.

Anonymized historical data and aggregated forecasts without personal reference.

EXERCISE 2 (20 min.) — Design your own Finance Agent Workflow

Task in 4 steps

Step 1: Take your top process from Exercise 1 (Score 3+/4).

Step 2: Define input (Input), agent steps, output (Output).

Step 3: Mark: Where does the workflow need a HITL gate?

Step 4: Check: Does the input contain GDPR-sensitive data? If yes: Add a data cleaning step.

Template: Workflow Sketch

TRIGGER: When does the workflow start? (Time / File arrival / Manual)


INPUT: What data? From which system? GDPR category?

AGENT 1: What does it do? Prompt outline in 2 sentences.

AGENT 2 (optional): Reviewer / Compliance

HITL GATE: Who approves? How? (Email / Teams / Form)

OUTPUT: What comes out? In what format?

 Tip: Prepare your workflow design — it will be directly implemented in the building session!



Learning progress check 2 — 6 questions | 1 correct answer each | Trainer discusses

1. What does a guardrail pattern do?

- A) Improves the LLM model
- B) Protects output from uncontrolled results
- C) Reduces cloud costs
- D) Automates approvals

2. GDPR: What must NOT be used in cloud AI?

- A) Aggregated monthly revenue
- B) IBAN, name, tax ID
- C) Anonymized benchmarks
- D) Public market data

3. EU AI Act Art. 4 requires:

- A) Complete AI ban
- B) Appropriate AI competency of users
- C) Technical certification
- D) External audits every 2 years

4. What is data masking in an agent workflow?

- A) File encryption
- B) Removal of sensitive data before AI call
- C) UI overlay for data privacy
- D) Cloud backup procedures

5. When is an audit trail required?

- A) Never — voluntary
- B) For AI-generated decisions (EU AI Act)
- C) Only during year-end closing
- D) Only for personal data

6. HITL approval step in Make.com — how does it work?

- A) AI approves itself
- B) Controller clicks email button — workflow continues
- C) Automatic timer approval after 24h
- D) Supervisor must enter code

Short break: Answer all 6 questions — then we will discuss together.



PAUSE



Screen off — take a short break — back in 10 minutes

After the break: Building Session — Your Finance Agent Takes Shape Now!

Keep your workflow design from Exercise 2 ready.

BUILDING SESSION (40 min.) — Today you build your first finance agent. No demo — your real data.

1



Step 1: Define Roles + I/O

2



Step 2: Draw Workflow Diagram

3



Step 3: Formulate Analyst Prompt

4



Step 4: Formulate Reviewer Prompt

5



Step 5: Test Agent Cascade

6



Step 6: Define HITL Gate + QC Checklist



At the end you will have: Analyst Prompt + Reviewer Prompt + HITL gate specification — ready to use immediately.



1 Step 1: Define Roles + I/O

Task: Define Your Agent Workflow

- A) WHAT is the input? (Which file / which database export / API source?)
- B) WHAT is the desired output? (Format, length, target audience)
- C) Which agent role is needed? (Analyst / Reviewer / Finance-Proxy / Compliance)
- D) Who is the HITL Gate decision-maker? (Name or role)

Template: Your Workflow Card

PROCESS NAME: _____

INPUT (GDPR Category Red/Yellow/Green): _____

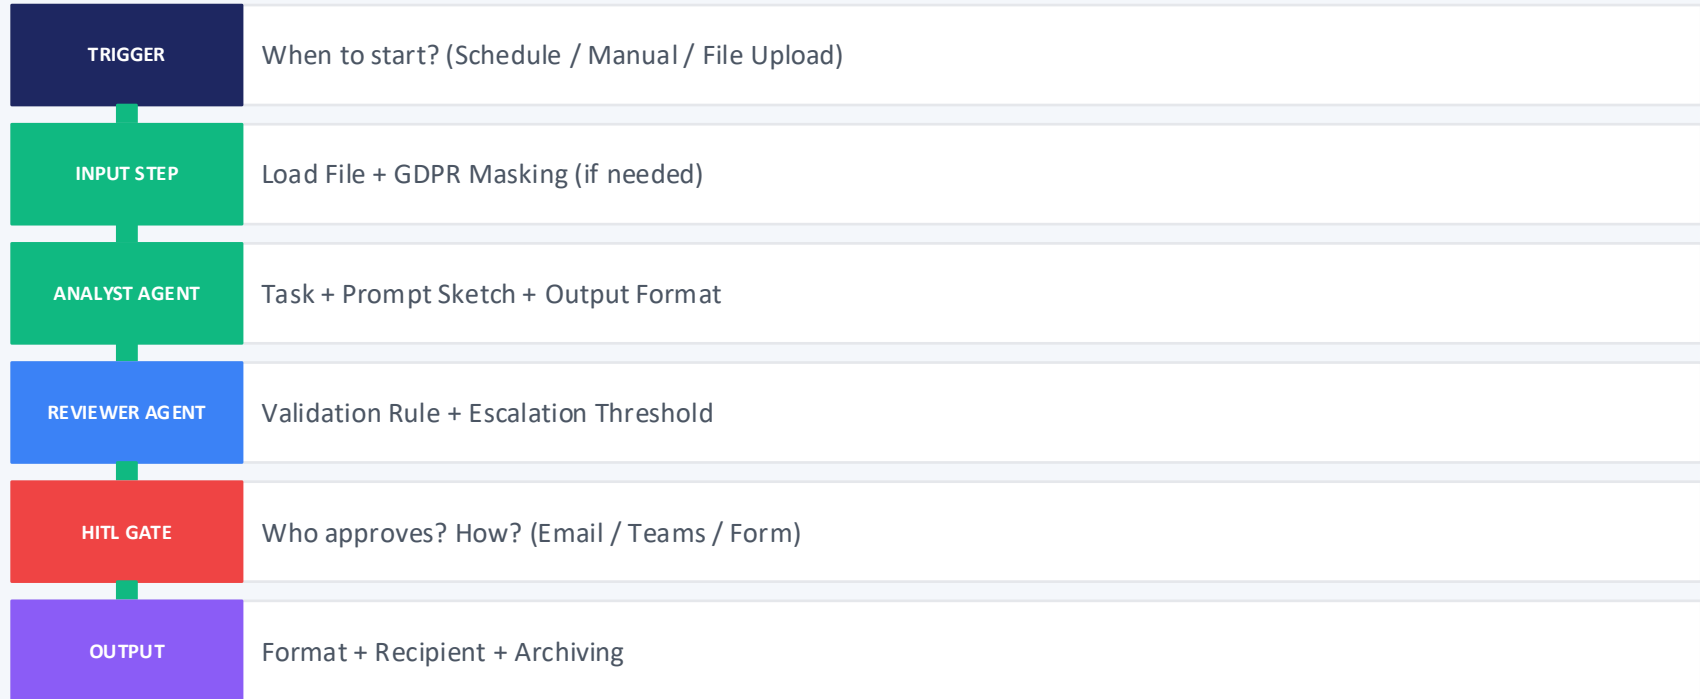
DESIRED OUTPUT: _____

AGENT ROLES: _____

HITL GATE-PERSON: _____

5 minutes. Fill out your workflow card. This is the foundation of your agent.

2 Step 2: Workflow Diagram



Draw the boxes — with arrows. Paper or notepad. 5 minutes.

3

Step 3: Formulate Analyst Prompt

Analyst Prompt: CRAFT Structure (from Module 1)

C — Context: You are a finance analyst for [Company Type] focused on [Area].

R — Role: Analyze precisely, number-based, without assumptions beyond the data.

A — Task: [Specific task — e.g. "Calculate EBIT Delta Q3 vs. Prior Year"].

F — Format: Table with Delta, %, Trend arrow, 1-sentence comment per line item.

Example Prompt (Northern Light Corporation)

You are a finance analyst for Northern Light Corporation. Analyze the following revenue export:

[Insert table here]

Calculate: (1) EBIT Delta Q3 vs. Prior Year in EUR and %. (2) Top 3 variance line items. (3) 1-sentence comment per line item.

Format: Table. No personal data. Only aggregated data.

4 Step 4: Formulate Reviewer Prompt

✓ Reviewer Prompt: Review Checklist + Escalation

- R1 — Completeness: Are all required fields/KPIs included in the Analyst output?
- R2 — Plausibility: Is any figure outside the expected range? (>20% variance: FLAG)
- R3 — Format: Is the required format correct (table, sentences, length)?
- R4 — Escalation: If a rule violation → respond with: "REWORK: [Feedback]"

Example Prompt (Northern Light Corporation)

You are an experienced controller. Review the following Analyst output for:

1. Completeness (all 3 KPIs present?)
2. Plausibility (no Delta > 30% without explanation?)
3. Format (table correctly structured?)

Respond with either: APPROVED or REWORK: [specific feedback].

Analyst Output: [Insert output here]

5 Step 5: Test Agent Cascade

TEST NOW: Run Your Analyst Prompt + Reviewer Prompt in ChatGPT

Test-Input

Open Northern Light Corporation Dataset.xlsx Tab Revenue_Q3. Copy the Table (without Personal References).

Analyst

Enter Test Input + Your Analyst Prompt in ChatGPT. Document Output.

Reviewer

Enter Analyst Output + Your Reviewer Prompt in ChatGPT. Do you get APPROVAL?

P-Q-R-Check

Check the output: plausibility, quality, accuracy. Note what is still missing.



6 Step 6: Define HITL Gate + QC Checklist

Specify HITL Gate

- A) Who approves? Name + Role: _____
- B) How is the output delivered? (Email / Teams / Form): _____
- C) What happens on REJECTION? (Workflow abort / Revision + Retry): _____
- D) Timeout Rule: If no response within X hours → automatically escalate to: _____

QC Checklist: Is Your Agent Ready?

- Analyst Prompt: CRAFT structure complete?
- Reviewer Prompt: APPROVAL / REVISION clearly defined?
- HITL Gate: Person + Method defined?
- Data Protection: Input Data Protection Compliant (No Personal References)?
- Audit Trail: Where is the agent output stored/logged?
- Test passed: Reviewer gave APPROVAL?

Building Session Complete — What Have You Built?

Please briefly present (2 min. per participant): Process + Agent + HITL Gate + 1 learning point

Bernd's Result (Example)

Process: Monthly EBIT variance commentary for Northern Light Corporation.

Analyst Agent: CRAFT Prompt with ERP export as input. Output: Table + Commentary.

Reviewer Agent: 5 review rules, escalation if >20% Delta.

HITL Gate: Teams message to Bernd Mueller — Approval via reply.

Estimated time savings: 4-5 hours per month.

Tanja's Result (Example)

Process: Budget alert email to stakeholders if EBIT variance > 10%.

Agent: GPT-4o formulates email based on template.

HITL Gate: Email approval before sending. GDPR: no personal data in prompt.

Key learning: The data masking step before the AI call was crucial.

6 Finance Agent Use Cases Already Running in Practice



Monthly Report Automation

Analyst loads ERP, calculates KPIs, Reviewer checks, Finance Proxy formats for board.

Time savings: 5-7 hrs./month



Anomaly Detection — Booking Journal

Agent scans all bookings for deviations $>2\sigma$, reports to controller.

Time savings: 3-4 hrs./month



Intercompany Reconciliation

Agent compares IC positions of both entities, creates difference list, escalates.

Time savings: 4-6 hrs./month



Cash Flow Forecast Update

Agent loads actuals, updates forecast model, explains top-3 drivers.

Time savings: 2-3 hrs./month



Stakeholder Reporting

Agent creates personalized reports for CFO, board, and investors from a shared data foundation.

Time savings: 3-5 hrs./month



Compliance Pre-Check

Agent checks bookings for policy compliance BEFORE month-end closing. Reduces rework.

Time savings: 2-4 hrs./month

These 6 use cases cover 80% of all routine finance tasks. Which will you build next?



MINI-CASE: Northern Light Corporation — An Agent Workflow in Crisis

Situation: Friday evening, 6:00 PM. CFO needs report by Monday 8:00 AM.

Bernd Mueller, Controller at Northern Light Corporation, receives a Teams message Friday at 6:00 PM:

"Bernd — I need the complete Q3 variance report for the board of directors by Monday 8:00 AM.

"5 entities. EBIT, Cashflow, Working Capital. With commentary and trend analysis."

Bernd thinks: Normally 2 days of work. With his new agent workflow?

He has: ERP export (aggregated, GDPR-compliant), his analyst prompt from the building session, his reviewer prompt, HITL gate via Teams, and... 3 hours this evening.


 **Task: What does Bernd do? Step by step.**

1. Does Bernd check the input for GDPR compliance first? 2. Which agent does he start first?
3. What does he do if the reviewer says "REWORK REQUIRED"? 4. When does he send it to the CFO?

You have 2 minutes. What would you do in Bernd's situation? Think out loud.

Bernd's Workflow — Step by Step


18:05

 GDPR Check: ERP export verified. Aggregated, no personal data. Green.

18:10

 Analyst Agent: CRAFT prompt + ERP table in ChatGPT. 5 entities, 3 KPIs. Output in 4 min.

18:20

 Reviewer Agent: Output in reviewer prompt. Result: "REWORK REQUIRED: Working Capital entity 3 missing."


18:28

 Bernd corrects: Entity 3 data reloaded, analyst re-run. Reviewer: "APPROVED".


18:35


 HITL Gate: Bernd reviews numbers personally (P-Q-R check). Signs off.

18:42

 Finance Proxy: Management commentary for board of directors. Board language, 3 paragraphs.

18:50

 Bernd sends Teams message to CFO: "Report is ready — earlier than expected."

 Result: 45 minutes instead of 2 days. Bernd sleeps well. The board gets its report.

MINI-CASE 2: Tanja + Budget-Alert — What's going wrong??

Tanjas Workflow

Tanja builds a Budget-Alert email workflow.

When EBIT deviation >10%: Agent drafts email to investors.

Workflow should send email automatically.

No data masking step. Investor email addresses in prompt. No HITL gate before send.

First email goes out: "Dear Mr. Sample, Your portfolio shows -12%..."

What's wrong?

GDPR violation: Investor names and email addresses in prompt = personal data in cloud AI.

No HITL gate: Email sent directly to external recipients without human review = high risk.

No approval: Automatically sent stakeholder communication = liability risk.

Correct solution: Data masking + HITL gate + approval before send

Investor IDs instead of names in prompt. Approve email draft via HITL gate. Only then send automatically.



Final check: 6 questions — Module 3 complete. 1 correct answer each.

1. What is agentic AI compared to chat LLM?

- A) Lower-cost model
- B) Autonomous, goal-oriented AI system with tool usage
- C) New OpenAI version
- D) Language model with larger dataset

2. Which role gives final approval before the finance proxy?

- A) Analyst Agent
- B) Compliance Agent
- C) Reviewer agent with HITL confirmation
- D) CFO directly

3. What does the Make.com approval step mean?

- A) Automatic approval after 24h
- B) Controller clicks email link: APPROVE or REJECT
- C) AI self-verification
- D) API certification

4. Which data is GDPR-compliant for cloud AI?

- A) IBAN and account number
- B) Aggregated revenue without personal reference
- C) Employee salary
- D) Customer addresses

5. EU AI Act Art. 26 applies to:

- A) All AI uses
- B) High-risk AI that influences decisions
- C) Only health data
- D) Only government systems

6. What are the 3 HITL triggers?

- A) Length + Format + Tone
- B) Amount + Uncertainty + External sending
- C) Only in case of errors
- D) Only during year-end closing



5 Things You Take Away Today

1

Agents vs. Chat:

AI agents pursue goals, use tools, and plan iteratively — chat LLMs respond only once. That is the fundamental difference.

2

Virtual Finance Team:

Analyst + Reviewer + HITL Gate = the minimum. Compliance + Finance Proxy = the complete solution. Building Session: Your agent is ready.

3

Automation Stack:

Make.com (Levels 3+4) is the right tool for finance controllers without IT support. No-code, EU servers, native AI integration.

4

Guardrails + GDPR:

Data masking, output validation, confidence threshold, audit trail — not optional, but mandatory. Sensitive data stays local.

5

HITL is not a weakness:

Human-in-the-Loop is not a bottleneck — it is your protection. Money, uncertainty, external sending = always human gate.



Transfer Task — Choose an Option

A**Option A: Analyst cascade in production**

Put your building session agent into practice.
Next real report: Use agent workflow, document result.

- Date: next monthly report
- Protocol: Agent output + Reviewer result + Approval timestamp
- Objective: Complete P-Q-R Check

B**Option B: Set up Make.com account**

Set up Make.com free tier and build your building session workflow as a real Make workflow.
No programming required.

- Create Make.com account (free)
- Configure Analyst Prompt as "ChatGPT" Module
- Integrate HITL Approval via Email

C**Option C: Audit Trail + Governance Documentation**

Create a simple governance document for your finance agent:
Purpose, data flow, guardrails, HITL rules, audit log schema.

- Template: ai_governance_pack_template.docx from Module 1
- Integrate EU AI Act Article 4 documentation
- Get supervisor signature

All three options produce production-ready results. Choose what fits your next step.



Digital Finance Masterclass — Your Progress

The Masterclass Roadmap

1

Kick-Off & Quick Wins

AI Fundamentals, CRAFT, Data Privacy, Custom GPT

✔ Completed

2

Data Analytics

PowerQuery, KPI Dictionary, Forecasting, IBCS

✔ Completed

3

Virtual Finance Team

Agents, Make.com, Guardrails, HITL

✔ Completed

4

Accounting & Operations

RAG, AP Automation, FraudDetection, SoD

→ Next Session

5

Strategy & Governance

EU AI Act, ROI Calculator, Roadmap, Change

→ Conclusion

My Quick-Win Plan — Starting Tomorrow

1. My Workflow:

Build an Agent Cascade for one Finance Use Case

2. Automation Scan:

Complete Readiness Score for Top-3 Processes

3. My HITL Gate:

Define a critical review step for Aloutput

4. Governance setup:

ai_governance_pack_template.docx — Components B1–B2

Quick-Win Worksheet — Your Action Items Before the Next Session

1

My Building Session Result

Process Name + Agent Configuration:

2

Transfer Option (A/B/C)

Selected: _____. Deadline: _____

3

Make.com Step

Account created by: _____

4

HITL Gate configured

Person: _____. Method: _____

5

First real agent run

Planned for: _____

6

Governance Document

Started: _____. Complete by: _____



Goal: By Module 4, at least 1 real agent run with documented results.





Module 3 Completed

Virtual Finance Team — built, tested, ready.



Module 4: Accounting Ops

Journal Entry Automation · Period-end close with AI · Matching + Reconciliation






Glossary — Key Terms Module 3

Term	Definition
Agentic AI	AI system that pursues goals, uses tools, and plans iteratively. Not just text responses.
Analyst Agent	LLM role that loads data, calculates KPIs, and creates structured outputs.
Reviewer Agent	LLM role that checks outputs: completeness, plausibility, format. Issues APPROVAL or FLAG.
HITL Gate	Human-in-the-Loop Gate: Human approval step before irreversible actions.
Make.com	No-code workflow automation tool. Connects APIs and AI models without programming.
Guardrail	Security rule for agents: output validation, data masking, confidence threshold.
Audit Trail	Log of every AI call: timestamp, user, model, input hash, output hash, approval.
EU AI Act Art. 4	Competence obligation: AI users must demonstrate adequate AI competence.
Data Masking	Removing personal data from AI inputs. Mandatory in GDPR gray areas.
P-Q-R Check	Plausibility + Quality + Review: Human final check before approval of every AI output.

You made it!

Thank you for your time and trust.

-  **I hope you've taken something away** — and will put it to use tomorrow.
-  **Suggestions for improvement?** Leave a comment on foundic.org — we read every suggestion and continuously improve our training.
-  **Our goal:** Free training for everyone — because professional development should not be a matter of budget.

 **If you enjoyed this training** — buy us a coffee. This helps us develop more free training.

 **Treat us to a coffee**

→ [Feedback & Comment: foundic.org/training](https://foundic.org/training)