



AI Implementation Curriculum Guide

A facilitator's path for turning a promising AI idea into a workflow people can use, review, improve and decide whether to scale.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 54 / EMBER & TEAL

EXECUTIVE WORKING EDITION - 08 OCT 2026

54

AT A GLANCE

A facilitator's path for turning a promising AI idea into a workflow people can use, review, improve and decide whether to scale.

This guide explains how to run the Halden & Company AI Implementation path through four stages: Contain, Design, Pilot and Learn. A pilot here is not a product demonstration. It is a managed learning loop around a meaningful piece of work. For each stage the guide gives goals, a timed session plan, working outputs, facilitation notes, prework, sticking points and signs that the stage landed. It also covers role tracks, written exercises, assessment beyond attendance, a six-week cadence and ways to keep the learning loop running after the pilot decision.

WHO IT IS FOR Business owners, process owners, delivery and technology teams and pilot participants responsible for an AI pilot, and the facilitators who guide them.

TIME TO USE 50 minutes to read; six weeks to run

What you will have at the end

- A clear view of who the path is for and what must exist before a pilot starts
- Timed plans for Contain, Design, Pilot and Learn with checkable outputs
- Exercises written so a delivery lead can run them
- Measures of learning and pilot quality beyond attendance
- A routine that carries the learning loop into the next pilot

Contents

01	Who the implementation path is for and what it assumes	03
02	Contain the use case	04
03	Workflow and quality design	06
04	Prepare people, controls and the supported launch	08
05	Review evidence and decide	10
06	Executives, managers and practitioners	12
07	Four exercises a facilitator can run as written	13
08	Measuring beyond attendance	15
09	A six-week schedule	17
10	Carrying the learning loop forward	18
11	A first quarter that produces evidence	19
12	Terms used in this publication	20

PROGRAM OVERVIEW

Who the implementation path is for and what it assumes

01



The path helps a team run one contained pilot well, and leave with a method it can repeat.

The AI Implementation curriculum uses co-design sessions, workflow prototyping, implementation support and retrospectives. Its typical shape is pilot design, workflow build, supported launch and evidence review. It deliberately avoids the product-demonstration pattern, where a tool is shown, enthusiasm builds and no one is accountable for what happens next. Instead, each stage produces a document someone owns: a charter, a workflow design, a readiness checklist, a feedback log and a decision record.



Who it is for

- Business owners who will be accountable for a pilot outcome
- Process owners who know how the work really runs
- Delivery and technology teams who will configure or integrate tools
- Pilot participants who will use the new workflow every day
- Organizations with several pilots that need a repeatable way to compare them



Prerequisites

You need a candidate use case, a business owner with authority over the workflow, a small group of people who will use it, and access to the information the work needs. A basic responsible-use position should exist, even if brief. You do not need to have chosen a tool. In fact, choosing the tool during Design is usually better than before the path begins, because the workflow decides what the tool must do.

STARTING POINT	WHAT TO EMPHASIZE	WATCH FOR
Starting from zero	A small, reviewable use case and safe learning	Aiming for broad automation first
Already experimenting	Roles, inputs, quality checks and learning loops made explicit	Pilots with no owner and no end date



STARTING POINT	WHAT TO EMPHASIZE	WATCH FOR
Preparing to scale	Repeatable patterns, governance and decision points	Scaling on the strength of one enthusiastic team


NOT LEGAL ADVICE

Implementation raises privacy, security, contract and employment questions. Check your pilot design with counsel, security and procurement colleagues, and follow current law and your own policies.


WORKING QUESTION

Who is the one person who will be accountable for the outcome of your pilot, and have they agreed?


STAGE 1: CONTAIN

Contain the use case

02

Choose a user, workflow, boundary and outcome specific enough to test.

Goals

- Name the user, the workflow and the desired change in one page
- Set the information boundary: what the pilot may and may not touch
- Agree the person accountable for the outcome and the people who will use it
- Define a start and end date, so the pilot cannot drift


Session plan (120 minutes)

TIME	ACTIVITY	OUTPUT
0-15	Owner presents the idea and why it matters now	Problem statement
15-45	Describe the current workflow and where it hurts, with a real example	Current-state sketch
45-70	Draw the boundary: users, volume, information, systems, excluded cases	Boundary
70-95	Write the desired change and the first measure of progress	Outcome and measure
95-115	Draft the charter and list assumptions to test	Charter draft
115-120	Assign charter finalization	Owner and date

Working outputs

- A pilot charter: purpose, scope, users, information boundary, owner, dates, measures, stop conditions
- A list of assumptions to test
- A list of what the pilot will not do

Facilitation notes

Containment is a creative act. The owner will want to include more users and more cases because the idea feels valuable. Ask which single type of case, performed by which five people, would teach you the most. Write stop conditions now, while no one is emotionally invested: for example, stop if reviewers find the same serious error type twice in a week. Keep the charter to one page.

Pework

- Owner brings a recent, real example of the work, anonymized where needed
- Technology lead lists the systems and data the work touches
- Facilitator interviews two likely pilot participants

Common sticking points

- Scope is a department rather than a workflow. Narrow to one case type
- The measure is vague, such as better efficiency. Ask what would be countable
- No one owns the outcome. Pause until there is a name
- The boundary is silent on information. Ask what the tool will read and write

How to tell it landed

- The charter fits on one page
- The owner can state the stop conditions from memory
- Participants have been named, not described as a team
- At least one thing was moved out of scope

WORKING QUESTION

If your pilot succeeded beyond expectations, what is the single case type you would want to have proven first?

STAGE 2: DESIGN

Workflow and quality design

03

Define inputs, roles, quality checks, exceptions, information boundaries and a measure of progress.

Goals

- Map the future workflow step by step, marking what AI does and what a person does
- Define the review points, the quality standard and the exception route
- Specify inputs and how information enters and leaves the tool
- Choose the tool or configuration that fits the workflow, not the reverse

Session plan (two sessions, 120 minutes each)

TIME	ACTIVITY	OUTPUT
S1 0-40	Map the future workflow, step by step, with the people who do it	Future-state map
S1 40-70	Mark each step: person, AI-assisted, AI-drafted with review, or not automated	Role allocation

TIME	ACTIVITY	OUTPUT
S1 70-120	Define quality: what good looks like with five examples and five failures	Quality standard
S2 0-40	Write the exception route: what happens when output is poor or inputs are missing	Exception design
S2 40-80	Specify inputs, outputs, information handling and logging	Information design
S2 80-120	Compare tool options against the design and decide what to configure	Tool decision

Working outputs

- A workflow design with roles and review points
- A quality standard with examples of good and poor output
- An exception procedure
- A shortlist or choice of tool with reasons

Facilitation notes

Review design determines the pilot's honesty. If review is optional or rushed, the pilot will appear to work while risk accumulates. Specify who reviews, what they check, how long they have and what they do when output is wrong. Insist on real examples for the quality standard. Teams that define good only in the abstract tend to accept anything that looks fluent. Involve security and procurement colleagues in the tool decision from the start.

Pework

- Participants gather ten real inputs and the outputs they would have produced
- Security, privacy and procurement contacts are told the pilot is coming
- Facilitator prepares a blank workflow template with swim lanes

Common sticking points

- The design assumes perfect inputs. Test with the messiest example
- Review is assigned to a manager who has no time. Assign to the person who knows the work
- Quality is described as accuracy only. Add tone, completeness, timeliness and fit
- The tool is chosen first and the workflow bent around it

How to tell it landed

- Every step has a named role

- The quality standard includes failure examples
- The exception route has been walked through with a real case
- The tool choice can be explained without reference to a demo



WORKING QUESTION

At which step in your workflow would a wrong answer cause the most harm, and who catches it?

STAGE 3: PILOT

Prepare people, controls and the supported launch

04



Introduce the change with the people closest to the work, and give them room to say what is not working.

Goals

- Make sure participants know the purpose, boundary, standard and escalation route before launch
- Confirm controls: access, logging, information handling and rollback
- Launch with support in the room or on call
- Capture friction daily and adjust within agreed limits



Session plan (launch workshop 90 minutes, then weekly check-ins of 30 minutes)

TIME	ACTIVITY	OUTPUT
0-20	Walk through the charter, boundaries and stop conditions	Shared understanding
20-50	Practise the workflow on three sample cases with coaching	Practiced launch



TIME	ACTIVITY	OUTPUT
50-70	Confirm access, logging, escalation contact and rollback steps	Controls checked
70-85	Complete the launch readiness checklist and sign	Go decision
85-90	Explain the feedback channel and daily log	Feedback route
Weekly	Review log, errors, time, participants' views, adjust within limits	Improvement log



Working outputs

- A completed launch readiness checklist
- A pilot feedback and improvement log
- Weekly review notes with changes made
- A running list of issues escalated and how they were resolved



Facilitation notes

Support matters more than training in this stage. Be present in the first days, either physically or in a channel, and answer quickly. Make it easy to report problems: a single form or thread, and a promise that reports lead to action. Review the log weekly and publish what changed, because participants who see their feedback acted on keep reporting. Do not change the design silently. Record every adjustment with a date and a reason.



Pework

- Participants read the charter and the quality standard
- Owner confirms time for participants to work in the new way
- Technology team tests access and rollback before launch day

Common sticking points

- Participants revert to old habits under pressure. Keep the first cases few and manageable
- The feedback log is empty because people are polite. Ask direct questions in check-ins
- Adjustments creep beyond the charter. Review against scope weekly
- Support is thin after the first week. Keep the on-call arrangement until the pilot ends



How to tell it landed

- Participants can state the boundary and the stop conditions

- The log has entries from every participant by the end of week one
- At least one change was made because of a participant's report
- Rollback was tested before launch



WORKING QUESTION

On day three of the pilot, how will a participant tell you something is going wrong, and what will you do within a day?

STAGE 4: LEARN

Review evidence and decide

05

Evaluate what happened, improve the workflow and consciously choose whether to extend, change, pause or stop.

Goals

- Assemble the evidence: measures, review findings, participants' views, incidents
- Compare results with the charter's outcome and stop conditions
- Decide to continue, change, pause or stop, and record the reasons
- Capture reusable patterns for the next pilot

Session plan (150 minutes)

TIME	ACTIVITY	OUTPUT
0-25	Present the evidence pack: measures, log, incidents, views	Shared facts
25-55	Participants give their account in their own words	Participant view
55-85	Compare to charter outcome and stop conditions	Gap analysis



TIME	ACTIVITY	OUTPUT
85-115	Discuss options: scale, change, pause, stop; list conditions for each	Options
115-140	Decide and write the record	Decision record
140-150	Capture lessons and assign follow-up	Lessons and owners

Working outputs

- A scale, pause or redesign decision record
- A lessons-learned note with reusable patterns
- An updated workflow and quality design
- A revised charter if the pilot continues

Facilitation notes

A decision to stop is a valid result if the evidence supports it, and teams need to hear that early. Protect participants who report problems. Separate evidence from advocacy: ask the owner and a skeptical reviewer to present their reading of the same facts. Where evidence is thin, extend the pilot with a clear question rather than scaling on hope. Record what you would need to see to change the decision.

Pework

- Facilitator compiles the evidence pack three days before
- Owner and one reviewer write separate one-page readings
- Risk or security lead summarizes any incidents or near misses

Common sticking points

- Positive anecdotes outweigh weak data. Ask for the measure first
- Costs of review time are omitted. Include the time reviewers spent
- Scaling pressure from above. Ask what additional controls scaling would require
- The pilot ends and no one writes the record. Schedule the writing in the session

How to tell it landed

- The decision matches the evidence and stop conditions
- Dissent is recorded in the decision record
- Participants were told the result and the reason
- The lessons are filed where the next pilot team will find them

WORKING QUESTION

Which result from your pilot would make you stop, and have you written it down before the pilot begins?

ROLE-BASED TRACKS

Executives, managers and practitioners

06

A pilot needs sponsors who unblock, managers who protect time and practitioners who tell the truth.

ROLE	ATTENDS	CONTRIBUTION
Executives	Contain session and the Learn decision	Sponsor, remove blockers, accept the decision
Business owner	All stages	Accountable for outcome, signs charter and decision
Managers	Design, launch workshop, weekly check-ins	Protect time, support participants, handle exceptions
Practitioners	Design, launch, weekly check-ins, Learn session	Use the workflow, review outputs, report friction
Technology and security	Design and launch	Configure, control access, support rollback

Executive track

Executives are asked for three things: attend the Contain session so scope decisions are made once, respond within two days to escalations, and attend the Learn decision ready to hear a no. Brief them on the stop conditions at the start, so a stop does not feel like a surprise or a failure.

Manager track

Managers adjust workload so participants have the time to review properly. They join the launch workshop and weekly check-ins, and they must resist the urge to count AI-assisted output volume as success before quality is shown. Equip them with a brief explaining what good output looks like.



Practitioner track

Practitioners are co-designers, not subjects. Involve them in Design, give them a clear way to challenge the tool and the process, and treat their view as evidence. Rotate a participant into the weekly review so the loop includes first-hand accounts.

WORKING QUESTION

Which practitioner in your pilot is most likely to tell you the uncomfortable truth, and have you asked them?



EXERCISES

Four exercises a facilitator can run as written

07

Each exercise has numbered steps and a defined output.



Exercise 1: Shrink the pilot (Contain, 40 minutes)

01 Start wide

The owner writes the pilot as they first imagine it, in three sentences.

02 Cut users

Ask: which five people would teach us the most? Remove everyone else.

03 Cut cases

Ask: which one case type is frequent and reviewable? Remove the rest.

04 Test

Read the shrunken version aloud. If the owner says it is too small to matter, ask what they would learn if it worked.



Exercise 2: Good and poor (Design, 45 minutes)

- 
- 01 Collect**
Bring ten real past outputs from the workflow, anonymized.
 - 02 Sort**
Participants place each in good, acceptable or not acceptable and write one reason.
 - 03 Extract**
Pull out the criteria that appeared in the reasons.
 - 04 Write**
Turn criteria into a checklist with examples. Test it on two new samples.

Exercise 3: Walk the exception (Design, 30 minutes)

- 
- 01 Pick a bad case**
Choose a real input that is incomplete or unusual.
 - 02 Walk the design**
Trace it through the new workflow, step by step, aloud.
 - 03 Mark gaps**
Note every point where nobody knows what to do.
 - 04 Fix**
Add rules or roles for each gap, and rewalk the same case.

Exercise 4: Evidence reading (Learn, 40 minutes)

- 
- 
- 01 Split**
Two readers prepare one page each from the same evidence, one arguing to scale, one arguing to stop.
 - 02 Present**
Each presents in five minutes without interruption.
 - 03 Test**
The group asks what evidence would change each reader's view.
 - 04 Decide**
Choose an outcome, and write the conditions that would reopen it.

ILLUSTRATIVE EXAMPLE

A shrunk pilot might read: five coordinators will use an assistant to draft first replies to one category of routine supplier inquiry for four weeks, with each reply reviewed before sending.



WORKING QUESTION

Which exercise would most change the pilot you are currently planning?

ASSESSMENT OF LEARNING

Measuring beyond attendance

08

K'

Success is a better decision and a workflow people can trust, not a pilot that merely ran.



Meeting attendance and tool usage statistics do not show learning. Evaluate at three levels: whether the team learned to run pilots, whether the workflow is better, and whether decisions improved. Set the baseline in Contain by recording how long the work takes now and how often it is redone. Treat estimates as estimates and label them.

LAYER	WHAT TO MEASURE	HOW
Team capability	Can the owner and team write a charter, quality standard and stop-use case conditions unaided	Ask them to draft one for a second case
Workflow quality	Review pass rate, rework, time per case compared with baseline	Sample cases weekly with the reviewer's notes
Participant experience	Confidence, workload, and trust in the workflow	Three-question check-in weekly



LAYER	WHAT TO MEASURE	HOW
Decision quality	Decision matched the evidence and the stop conditions	Compare the record with the charter



Signals worth watching

- Reviewers catch problems early and the design changes as a result
- Participants volunteer improvements to the workflow
- Review time stays reasonable instead of ballooning
- Stop conditions are discussed openly
- The second pilot starts faster because the first left reusable pieces



AVOID VOLUME AS SUCCESS

Counting AI-generated items says nothing about their value. Pair every volume figure with a quality figure and the review time it consumed.



WORKING QUESTION

What baseline measure can you take this week, before the pilot, so that you can show a change later?



SAMPLE CADENCE

A six-week schedule

09

Six weeks fits a small pilot only if Contain and Design finish in the first two weeks.

WEEK	SESSION OR TASK	WHO
1	Contain session (120 minutes); charter finalized	Owner, executives, facilitator
2	Design sessions 1 and 2; tool decision; controls checked	Owner, practitioners, technology
3	Launch workshop (90 minutes); supported launch begins	Participants, managers
4	Pilot running; weekly check-in; log review	Participants, owner
5	Pilot running; weekly check-in; evidence pack drafted	Participants, owner
6	Learn session (150 minutes); decision record; lessons	Executives, owner, participants

Running the cadence

- Do not launch until the readiness checklist is signed
- Keep the pilot to a fixed number of participants and cases
- Hold the weekly check-in at the same time, even if there is little to report
- Update the improvement log within a day of any change
- If the pilot needs longer, extend it with a stated question, not by default

Complex integrations may need a longer build before launch. Keep the Contain and Design outputs and the Learn decision, and move the dates. Do not shorten the review step to save time.

WORKING QUESTION

What would force you to delay the launch date, and who has the authority to say so?

SUSTAINING THE WORK

Carrying the learning loop forward

10

The method is the asset: keep it in use after the first pilot.

01 File the pack

Store the charter, design, checklist, log and decision record in a shared place.

02 Run a retrospective

Meet within two weeks of the decision to record what to repeat and what to change.

03 Monitor live workflows

If the workflow continues, keep sampling quality monthly and review incidents.

04 Start the next one

Use the same templates for the next pilot, and compare results using the same measures.

Sustainment checklist

- A named owner for the live workflow
- A monthly quality sample with a threshold for action
- A route for participants to raise concerns
- A register of pilots, owners and decisions
- A review of controls when tools or terms change

Free assessments are available at halden.company/assessments, and you can contact the team at halden.company/contact if you want help with a pilot.

WORKING QUESTION

If the person who led your pilot left tomorrow, what would the next person need to continue?

30 / 60 / 90

A first quarter that produces evidence

11

In 90 days you can run one contained pilot, reach a recorded decision and have templates for the next.

First 30 days: ground it

- Name the business owner and choose one contained use case
- Write the pilot charter with stop conditions
- Gather ten real inputs and take a baseline measure
- Inform security, privacy and procurement contacts

Days 31 to 60: test it

- Complete the workflow design, quality standard and exception route
- Choose and configure the tool, and test rollback
- Hold the launch workshop and sign the readiness checklist
- Start the supported pilot and the weekly log

Days 61 to 90: decide

- Assemble the evidence pack and hold the Learn session
- Record a continue, change, pause or stop decision
- Hold the retrospective and file reusable templates
- Select the next candidate use case

GLOSSARY

Terms used in this publication

12

TERM	MEANING
Pilot charter	A one-page statement of purpose, scope, users, information boundary, owner, dates and stop conditions.
Stop condition	A pre-agreed result that ends or pauses the pilot.
Information boundary	The limits on what information a pilot may read, write or store.
Quality standard	A description with examples of what acceptable output looks like.
Exception route	The agreed steps when output is poor or inputs are missing.
Readiness checklist	A pre-launch list confirming people, controls and support are in place.
Rollback	The steps to return to the previous way of working.
Improvement log	A dated record of friction reported and changes made during the pilot.
Evidence pack	The measures, log, incidents and participant views gathered for the decision.
Decision record	A short document of the decision, reasons, dissent and conditions to reopen it.

A note on use

This guide is an educational resource, not legal, technical or financial advice. Adapt it to your organization and confirm legal and regulatory points with qualified advisers.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.