PUBLIC-FACING PDF | REVIEW BEFORE PRODUCTION | FOLIUM SYSTEMS

AI PROFITABILITY ENGINEERING

AI Profitability Engineering

Many AI programs start from model access and then search for value. Folium starts from the business workflow and engineers backward into the smallest capable system. The profitable path is often not the largest model. It may be a focused model, RAG lane, rules-assisted workflow, local CPU-capable route, semantic cache, tool router, private endpoint, human-gated action system, or hybrid architecture that completes a specific job at a controlled cost.

AUDIENCE

Owners, operators, finance leaders, technical buyers, investors, AI sponsors, and teams trying to make AI economically useful

PURPOSE

Show how Folium designs AI around margin, cost-to-capability fit, workflow value, and measurable business outcomes

UPDATED

May 2026

- AI profitability starts with the work, not the model.
- Right-sized systems can protect margin by avoiding large-model defaulting, repeated token waste, and unfocused chat usage.
- Focused models, local CPU-capable lanes, deterministic tools, retrieval, caching, batching, and model cascades can outperform broad AI usage when the workflow is specific.
- Folium measures AI by useful work completed, cost to run, risk controlled, revenue recovered, and operating burden reduced.

SERVICE ARCHITECTURE

Folium service lines are organized around the work buyers need to control.

Audits, RAG, agents, software, integrations, governance, private AI, commerce AI, modernization, and AI operations become one visible service map.

SERVICE MAP

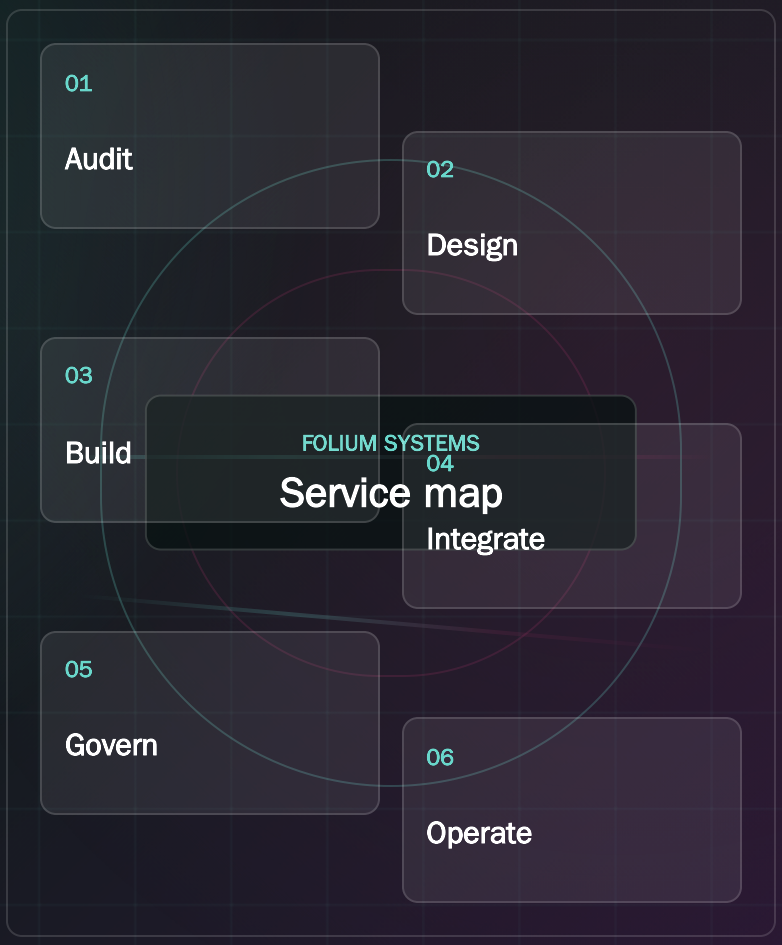

01
Helps first-time buyers understand the offer quickly.

02
Shows that services connect instead of living as scattered pages.

03
Turns broad capability into a controlled next move.

R

NAVIGATION MAP

Choose the review route before reading cover to cover.

This packet is meant to support a real decision meeting. Different reviewers should enter through different routes, then come back together around the same controlled next step.

DECISION ROUTE | OPERATING ROUTE | TRUST ROUTE

EXECUTIVE ROUTE

Decision first

Start with the cover, visual summary, executive read, controls, first ninety days, and handoff. This route helps leaders decide whether the next move is education, audit, first build, pilot, or operations.

- Outcome
- Risk
- Owner
- Next gate

OPERATIONS ROUTE

How the work will run

Read the workflow map, procedures, operating roles, metrics, first sprint, and buyer worksheet. This route shows whether staff can actually use, review, and improve the future process.

- Workflow
- Staff
- Support
- Improve

TECHNICAL AND TRUST ROUTE

Where the boundaries live

Focus on records and work products, controls, risk assumptions, reference work products, source truth, runtime placement, and launch conditions before any private access expands.

- Source
- Access
- Runtime
- Rollback

BUYER SESSION ROUTE

Turn reading into a working session

Use the discovery questions, role review route, buyer worksheet, and engagement fit ladder to prepare one process, one owner, one source map, and one next decision.

- Process
- Examples
- Questions
- Decision

> Best use: bring one workflow, the people who own it, the systems it touches, the data classes involved, and the decision this packet should help leadership make.

01

EXECUTIVE READ

AI profitability engineering in plain language.

Many AI programs start from model access and then search for value. Folium starts from the business workflow and engineers backward into the smallest capable system. The profitable path is often not the largest model. It may be a focused model, RAG lane, rules-assisted workflow, local CPU-capable route, semantic cache, tool router, private endpoint, human-gated action system, or hybrid architecture that completes a specific job at a controlled cost.

RECORD | BOUNDARY | ACTION

NARROW

Start with one profitable job

Name the workflow, owner, current cost, desired output, acceptable risk, and value before choosing the model route.

- Workflow
- Owner
- Value

RIGHT-SIZE

Use the smallest capable system

A focused model, RAG lane, rules-assisted agent, local CPU path, private route, or cloud model should be chosen by task economics.

- Focused
- Local
- Hybrid

CONTROL

Stop silent AI cost leaks

Caching, batching, route reuse, tool contracts, rate limits, and human gates keep recurring cost tied to useful output.

- Cache
- Batch
- Gate

MEASURE

Profitability is a release gate

The system should show cost per useful output, saved time, avoided rework, recovered revenue, and support burden.

- Cost
- Save
- Earn

REUSE

Turn delivery into compounding machinery

Templates, eval cases, prompts, caches, routing patterns, launch gates, and review files reduce future delivery friction.

- Tools
- Tests
- Records

RETIRE

Remove weak AI lanes

AI routes that do not earn their cost should be simplified, cached, localized, rerouted, paused, or shut down.

- Simplify
- Pause
- Reroute

> This packet is public-facing. It is written for serious review without exposing private infrastructure, customer data, credentials, live provider wiring, or internal project labels.

02

WORKFLOW MAP

The operating path should be visible before anyone trusts the outcome.

Folium uses workflow maps to turn broad AI ambition into inspectable work. Each phase names the procedure, the visible output, and the decision gate that prevents excitement from outrunning control.

DECISION GRID | REVIEW LENS | NEXT STEP

PHASE	PROCEDURE	VISIBLE OUTPUT	DECISION GATE
Cost the current work	Measure labor time, rework, delays, support load, error cost, missed revenue, and customer friction.	Current-state cost map.	The workflow has measurable economic pressure.
Define the useful output	Name what AI must retrieve, classify, draft, route, validate, summarize, notify, or prepare for approval.	Useful-output definition.	Output is tied to work, not novelty.
Select the system route	Compare focused model, RAG, local CPU-capable path, private endpoint, rules-assisted agent, cloud model, or hybrid route.	Right-size route matrix.	The smallest capable route is justified.
Exploit existing capacity	Check whether current hardware, local CPU lanes, browser-adjacent tools, databases, or existing software can carry part of the workload.	Existing-capacity worksheet.	The customer does not buy unnecessary infrastructure.
Design a model cascade	Route easy work to rules, retrieval, caches, small models, or focused models before escalating to larger models for hard cases.	Model cascade policy.	Large-model usage is reserved for work that needs it.
Reduce recurring cost	Add caching, batching, prompt reuse, retrieval, tool routing, rate limits, quota alerts, and fallback lanes.	Cost-control design.	Repeated work does not become repeated waste.
Gate authority	Keep expensive, sensitive, customer-facing, or state-changing actions human-reviewed until records justify expansion.	Authority gate map.	Autonomy does not outrun economics or trust.
Tie work to revenue	Identify where AI can recover missed leads, improve follow-up, reduce returns, speed quote response, improve content quality, or prevent customer churn.	Revenue-recovery map.	AI is measured against value created, not only cost avoided.
Measure the result	Track cost per useful output, cycle time, quality, rework, support load, revenue recovery, and next improvement.	AI profitability ledger.	The build can prove whether it should expand.
Improve or retire	Tune, simplify, cache, reroute, replace, localize, scale, or retire AI lanes based on value and cost evidence.	Improvement decision record.	The system earns its place each cycle.

03

RECORDS AND WORK PRODUCTS

The work should leave behind material a buyer can inspect.

A serious engagement should produce more than conversation. Folium packages records, diagrams, checklists, routes, system surfaces, launch gates, and handoff material so the buyer can keep control after the first win.

DECISION GRID | REVIEW LENS | NEXT STEP

WORK PRODUCT	WHAT IT CONTAINS	HOW THE REVIEWER USES IT
Current-state cost map	Labor, delay, error, support, subscription, rework, and missed-revenue pressure.	Shows where AI could create economic value.
Useful-output definition	The specific work product AI must produce, route, or prepare.	Prevents chat volume from replacing business value.
Right-size route matrix	Focused model, RAG, CPU/local lane, private endpoint, cloud API, rules, agent, and hybrid options.	Chooses AI placement by economics and risk.
Existing-capacity worksheet	Current hardware, CPU paths, local storage, databases, browsers, SaaS tools, and staff systems that can support the lane.	Prevents infrastructure spend before it is justified.
Model cascade policy	Escalation order from deterministic logic to retrieval, focused models, small models, local/private routes, and larger models.	Reserves expensive routes for the work that needs them.
Cost-control design	Caching, batching, retrieval, prompt reuse, quota limits, alerts, and fallback.	Reduces recurring spend and runaway usage.
Authority gate map	Human approvals, allowed actions, blocked actions, spend thresholds, and rollback.	Controls financial and operating risk.
Revenue-recovery map	Lead response, quote speed, abandoned workflow recovery, support resolution, retention, content quality, and commerce lift opportunities.	Shows where AI can create new value, not only reduce expense.
Profitability ledger	Cost per useful output, time saved, quality, rework avoided, revenue recovered, and support burden.	Tells leaders whether the AI lane should expand.
Retirement and simplification record	What should be cached, localized, simplified, merged, paused, or retired.	Keeps AI from becoming permanent overhead.

04

PROCEDURES

The procedure is the product as much as the technology.

The goal is not to make AI look impressive for one meeting. The goal is to make the operating path repeatable, explainable, reviewable, and safe enough to improve.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- Do not start with a general model when the business has not named the work.
- Choose the smallest capable model or route that can perform the task safely.
- Use focused models, RAG, rules, tools, and local CPU-capable lanes where they fit the job.
- Reserve larger frontier models for tasks that actually need broad reasoning, language quality, or complex synthesis.
- Route easy work through deterministic logic, retrieval, cached answers, small models, or focused models before escalating to expensive routes.
- Check whether existing hardware or software can carry the first lane before recommending new infrastructure.
- Cache repeated answers, prompts, retrieval results, and workflow decisions where safe.
- Batch and schedule non-urgent AI work instead of paying premium runtime costs for everything.
- Measure cost per useful output, not only token volume or chat count.
- Tie every AI lane to a workflow owner and a business metric.
- Track revenue recovered, churn avoided, faster response, better quality, and customer confidence in addition to labor savings.
- Keep human gates on expensive, sensitive, customer-facing, or state-changing actions.
- Retire, simplify, or reroute AI lanes that do not earn their operating cost.

05

CONTROLS

Governance, quality, and launch gates keep speed honest.

Folium keeps the buyer's next decision tied to observable gates: source truth, authority, access, testing, ownership, support, rollback, and improvement cadence.

DECISION GRID | REVIEW LENS | NEXT STEP

GATE	WHAT MUST BE TRUE	STOP OR REFINE SIGNAL
Work gate	The workflow, owner, current cost, and useful output are named.	AI is being added because it is fashionable.
Route gate	The model/runtime choice is justified by task, risk, cost, latency, privacy, and support.	Everything routes to the largest model by default.
Existing-capacity gate	Current hardware, CPU lanes, databases, or owned tools were checked before new spend was added.	The customer buys new infrastructure before proving need.
Cascade gate	Easy tasks route to lower-cost lanes before expensive models are used.	Large-model calls become the default path.
Cost gate	Expected usage, recurring cost, cache opportunities, alerts, and thresholds are visible.	The system can spend silently.
Value gate	The system measures saved time, avoided rework, revenue recovery, quality, or risk reduction.	No one can tell if the AI paid for itself.
Authority gate	Human approval, rollback, and blocked actions exist for costly or risky behavior.	AI can create expensive mistakes without review.
Retirement gate	There is a path to simplify, localize, cache, merge, pause, or shut down low-value lanes.	AI becomes overhead forever.

06

DISCOVERY QUESTIONS

The right questions expose the real project.

These prompts help a buyer and Folium decide whether the next step should be education, audit, first build, security review, pilot, or an operating support path.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- Which workflow costs the business money every week?
- What is the useful output AI must create, route, validate, or prepare?
- What is the current cost of doing the work manually or with existing tools?
- What would make the AI lane too expensive to justify?
- Can a focused model, RAG lane, rules-assisted workflow, or CPU-capable local route do the job?
- Which tasks truly need a larger model?
- Can the first lane run on existing hardware, current databases, existing SaaS tools, or a modest private route?
- What should be handled by rules, retrieval, cache, or a small model before reaching a larger model?
- Which repeated work can be cached, batched, reused, or routed to a cheaper lane?
- Where could AI recover missed revenue instead of only reducing manual effort?
- What action must stay human-reviewed until profit and safety are proven?
- What metric would prove the AI lane earned expansion?
- What would cause the business to retire or reroute the AI lane?

07

VISUAL DIGESTION

Diagrams, charts, and overlays make the work easier to review.

Dense AI work should not only be explained in paragraphs. The reviewer should be able to inspect maps, scorecards, matrices, lanes, and before-after views that reveal where the value and risk live.

RECORD | BOUNDARY | ACTION

AI profit path

A flow from current cost to useful output to right-sized route to cost gate to margin ledger.

- Cost
- Output
- Route
- Ledger

Model sizing ladder

A ladder from rules and retrieval to focused models, local CPU lanes, private endpoints, and larger cloud models.

- Rules
- RAG
- Focused
- Frontier

Cost leak map

A map of token waste, repeated prompts, overpowered models, unmanaged agents, no caching, and silent provider spend.

- Tokens
- Routes
- Agents
- Spend

Margin gate dashboard

A dashboard for cost per useful output, time saved, revenue recovered, quality, support load, and rollback triggers.

- Cost
- Time
- Revenue
- Risk

Existing-capacity route

A diagram showing where CPU, owned databases, current software, private endpoints, and cloud APIs can share the lane.

- CPU
- Data
- Private
- Cloud

Revenue recovery loop

A loop from missed work to faster response to cleaner follow-up to recovered opportunity.

- Miss
- Respond
- Follow
- Recover

08

OPERATING ROLES

Every serious AI path needs named owners before it becomes dependency.

The same technology can be safe or unsafe depending on who owns the workflow, data, quality, launch authority, support, and improvement loop. Folium makes those responsibilities explicit so no buyer inherits an orphaned system.

DECISION GRID | REVIEW LENS | NEXT STEP

ROLE	OWNS	RECORD TO INSPECT
Executive sponsor	Priority, budget, risk tolerance, stop/continue decision, and expansion timing.	Decision note, value hypothesis, and approval boundary.
Business process owner	The day-to-day work, acceptance criteria, staff impact, and operational usefulness.	Workflow map, user feedback, and adoption notes.
Technical owner	Systems, APIs, databases, runtime placement, deployment, monitoring, and fallback.	Architecture map, integration log, and support route.
Knowledge owner	Source truth, document freshness, policies, retrieval scope, and correction workflow.	Source inventory, freshness cadence, and review exceptions.
Security or risk reviewer	Data classes, credentials, access, logs, retention, blocked actions, and incident path.	Boundary map, permission table, and rollback trigger.
Folium delivery lead	Build coordination, review file, known limits, quality checks, and handoff completeness.	Launch room, eval record, and improvement backlog.

09

QUALITY SCORECARD

A max-detail packet should tell reviewers how to judge the work.

Folium uses scorecards to make a subjective AI conversation more inspectable. The score is not a substitute for judgment; it helps leadership see whether the next step is education, repair, sandbox, pilot, or operations.

DECISION GRID | REVIEW LENS | NEXT STEP

SCORE AREA	STRONG SIGNAL	WEAK SIGNAL
Business fit	The workflow is specific, painful, owned, and tied to measurable operational improvement.	The project is framed as adding AI generally.
Source truth	Approved sources are known, fresh, classified, and connected to the answer path.	The system mixes stale, unknown, or unapproved sources.
Behavior quality	Representative tasks pass, wrong-answer behavior is known, and edge cases are recorded.	The review build only shows a polished happy path.
Authority control	AI actions are separated into draft, retrieve, recommend, route, execute, block, and escalate.	The system can act without visible permission.
Staff readiness	Users can explain the tool, correct it, escalate, and understand their role.	Staff feel replaced, confused, or unsupported.
Operations readiness	Support, monitoring, rollback, release rhythm, and source refresh are owned.	No one knows who maintains the system after launch.

10

THIRTY / SIXTY / NINETY

The work should have a believable first ninety days.

A controlled first ninety days keeps ambition high without turning uncertainty into production risk. Folium uses the period to move from understanding into a narrow working example, then into reviewable operating rhythm.

DECISION GRID | REVIEW LENS | NEXT STEP

WINDOW	FOCUS	EXPECTED OUTPUT
First 30 days	Discovery, source inventory, first-lane selection, staff interviews, data boundary, and build plan.	Process map, owner map, first-build scope, source list, and launch blockers.
Days 31-60	Working surface, RAG or agent behavior, integration stub, evaluation cases, browser checks, and staff review.	Sandbox, evaluation file, screenshots, known limits, and repair list.
Days 61-90	Architecture review, pilot conditions, governance layer, training guide, support path, and improvement cadence.	Launch room, go/no-go record, operations guide, and next-stage recommendation.

11

RISK AND ASSUMPTION REGISTER

The hidden assumptions should be visible before they become expensive.

Every AI engagement contains assumptions about data, people, systems, cost, behavior, and authority. Folium treats those assumptions as review material, not background noise.

DECISION GRID | REVIEW LENS | NEXT STEP

ASSUMPTION	WHY IT MATTERS	HOW FOLIUM REVIEWS IT
The source is authoritative	AI can only be as reliable as the sources and business rules it is allowed to use.	Source inventory, owner confirmation, retrieval tests, freshness cadence.
The process is ready	A broken process can become a faster broken process when AI is added too early.	Workflow mapping, bottleneck review, owner interview, first-lane narrowing.
The runtime fits the data	Cloud, private, local, and hybrid routes carry different privacy, cost, latency, and support tradeoffs.	Runtime matrix, data classification, provider review, fallback plan.
Staff will adopt the tool	Adoption fails when users do not understand, trust, correct, or benefit from the system.	Training notes, staff review, feedback loop, manager visibility.
Authority is clear	The system can create harm if it sends, updates, approves, or routes without permission.	Permission table, blocked actions, human review, audit trail.
The system can be supported	A useful first build becomes fragile if nobody owns incidents, source updates, or cost review.	Support guide, owner map, release rhythm, rollback trigger.

12

FIRST SPRINT PROCEDURE

The first sprint should produce something real and reviewable.

Folium prefers a narrow first sprint that creates a working surface or review file the buyer can challenge. The first sprint is not the final system; it is the safest way to make the future visible.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- Confirm the single process and the decision the sprint must support.
- Collect approved example material, redacted review records, public references, screenshots, workflow notes, and source rules.
- Define what will be built: portal, dashboard, RAG assistant, agent route, integration adapter, audit file, or launch room.
- Create the visual workflow: intake, source, model or agent route, human review, output, record, and next gate.
- Run representative tasks, edge cases, bad input, missing data, and blocked-action tests.
- Prepare browser screenshots, known limits, support questions, and next-stage blockers.
- Review with staff and leadership before expanding data, access, authority, or dependency.
- End with a decision: stop, refine, rebuild, pilot, or prepare an operating plan.

13

REFERENCE WORK PRODUCTS

The packet should make the invisible work tangible.

AI work often fails because the important pieces are invisible until something breaks. Folium turns those pieces into work products the buyer can open, print, challenge, and improve.

RECORD	BOUNDARY	ACTION

Process map

A before-and-after workflow showing people, systems, data, decision points, blockers, and expected output.

- Before
- After
- Owner
- Gate

Data boundary map

A map of source classes, approved use, blocked use, retention, provider exposure, and custody.

- Public
- Internal
- Private
- Blocked

Model and agent route

A path showing which model, tool, retrieval source, or agent lane is used and where humans approve.

- Route
- Tool
- Review
- Escalate

Evaluation file

A record of tasks, expected outcomes, failures, repairs, known limits, and acceptance criteria.

- Cases
- Failures
- Repairs
- Limits

Launch room

A board for owners, support, training, rollback, incidents, go/no-go, and improvement backlog.

- Owner
- Support
- Rollback
- Backlog

Handoff guide

A plain-language guide staff can use to understand what the system does, cannot do, and how to report problems.

- Use
- Limit
- Correct
- Report

14

METRICS AND REVIEW RHYTHM

The business should know how improvement will be measured.

Folium keeps measurement practical. The first goal is not a perfect dashboard; it is a clear set of signals that shows whether the process is saving time, reducing risk, strengthening staff, or improving customer outcomes.

DECISION GRID | REVIEW LENS | NEXT STEP

SIGNAL	WHAT TO WATCH	DECISION IT SUPPORTS
Time recovered	Manual steps removed, average handling time, repeated work reduced, faster routing.	Should this workflow expand to more users or adjacent processes?
Quality improved	Wrong answers, missing sources, correction rate, review exceptions, customer rework.	Is behavior strong enough for pilot or does it need repair?
Risk reduced	Blocked unsafe actions, escalations, data-boundary violations avoided, rollback readiness.	Can authority expand or should controls remain tight?
Staff confidence	Training completion, feedback volume, adoption friction, override rate, manager notes.	Does the workforce need more support before launch?
Cost and runtime	Provider cost, local infrastructure cost, latency, uptime, fallback use, subscription sprawl.	Should runtime placement change?
Customer impact	Response speed, consistency, issue resolution, conversion support, satisfaction signals.	Is the capability improving the business outcome?

15

ROLE REVIEW ROUTE

Each reviewer should know what to inspect first.

A max-detail packet is only useful when different reviewers can find their lane quickly. Folium separates executive, operations, technical, security, finance, and staff questions so the buyer can bring the right people into the right part of the review.

DECISION GRID | REVIEW LENS | NEXT STEP

REVIEWER	START WITH	DECISION THEY SUPPORT
Executive sponsor	Value hypothesis, launch gate, first ninety days, and stop/refine/continue choices.	Whether the process deserves a controlled engagement.
Operations lead	Workflow map, operating roles, support rhythm, and staff feedback loop.	Whether the future process can be run by the team.
Technical lead	Runtime placement, data path, integration surface, monitoring, and fallback.	Whether the architecture can be supported safely.
Security or risk reviewer	Data classes, permissions, blocked actions, logs, retention, and rollback.	Whether access can expand beyond public review.
Finance or owner	Cost signals, subscription overlap, runtime tradeoffs, labor impact, and support burden.	Whether the first build has a practical business case.
Staff user	Plain-language use, limits, escalation, correction path, and training expectations.	Whether the tool strengthens the job instead of confusing it.

16

BUYER WORKSHEET

The packet should turn into a working session, not only reading material.

Before a call, Folium wants the buyer to gather the real operating pieces that make the review useful. The worksheet keeps the conversation grounded in one process, one owner, one source map, and one next decision.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- Bring one workflow that is slow, risky, expensive, repetitive, customer-visible, or staff-heavy.
- Name the systems touched by the workflow: store, CRM, ERP, inbox, spreadsheet, database, portal, document folder, or legacy application.
- Separate approved public material from internal, customer, regulated, confidential, credential, and blocked material.
- Write down who owns the work today, who reviews exceptions, and who will own the AI-assisted version.
- List the decisions AI may draft, retrieve, recommend, route, block, or escalate, and the decisions that stay human-owned.
- Bring examples of good output, bad output, common exceptions, missing data, and customer-facing risk.
- Name the first useful working surface: dashboard, portal, assistant, queue, control room, commerce lane, integration, or review file.
- Decide what record would make leadership comfortable with the next stage.

17

ENGAGEMENT FIT LADDER

The next step should match the maturity of the record.

Folium does not need every buyer to start at the same altitude. The right offer depends on how much process clarity, source truth, owner alignment, and launch readiness already exists.

DECISION GRID | REVIEW LENS | NEXT STEP

IF THE BUYER HAS	BEST NEXT FOLIUM MOVE	OUTPUT TO EXPECT
AI interest but no clear process	AI systems audit or first workflow finder.	Pressure map, source inventory, first-lane recommendation, and risk view.
A clear process but no working surface	Forward engineering first sprint.	Clickable surface, route map, known limits, and next-stage blockers.
A tool that works in parts but not in operations	Architecture and launch readiness review.	Permission map, runtime decision, support model, and go/no-go record.
A failed or frightening rollout	AI recovery and staff enablement path.	Issue register, staff training plan, repair roadmap, and confidence loop.
Sensitive data or cost pressure	Local, private, or hybrid AI placement review.	Runtime matrix, data custody plan, fallback route, and vendor-exit view.
A useful pilot that needs care	AI operations support.	Monitoring rhythm, source refresh, release notes, incident path, and improvement backlog.

18

HANDOFF

The last page of a packet should create the next controlled move.

Folium's handoff view separates what can be done now, what needs customer records, what needs approval, and what should wait until the review file is stronger.

DECISION GRID | REVIEW LENS | NEXT STEP

HANDOFF LANE	OWNER	NEXT RECORD
Executive sponsor	Defines the business outcome, expansion decision, and acceptable cost range.	Profitability thesis and go/no-go gate.
Finance or operations owner	Tracks cost baseline, savings, recovered revenue, and support burden.	AI profitability ledger.
Technical owner	Owns model route, runtime placement, cache strategy, monitoring, and fallback.	Right-size route matrix and cost controls.
Risk or support owner	Defines human gates, blocked actions, escalation, rollback, and customer-impact controls.	Authority gate map and incident path.
Folium delivery lead	Connects workflow value, model route, operating records, and improvement backlog.	Launch record and economic review file.

> The strongest next step is narrow: one process, one owner, one source map, one working surface, one review file, and one decision gate.

19

NEXT STEP

AI should not be a cost center disguised as innovation.

Use this packet to identify the first AI lane that can earn its operating cost through focused work, right-sized models, controlled runtime, and measurable business outcomes.

Bring the process

Name the business process, the systems involved, the people affected, and the decision this PDF should support.

Separate review from production

Keep public examples, sandbox review, pilot access, and production dependency in separate stages with clear owners.

Ask for the record

Request screenshots, browser checks, known limits, launch blockers, support plans, and the next approval path.