



Folium

PUBLIC-FACING PDF

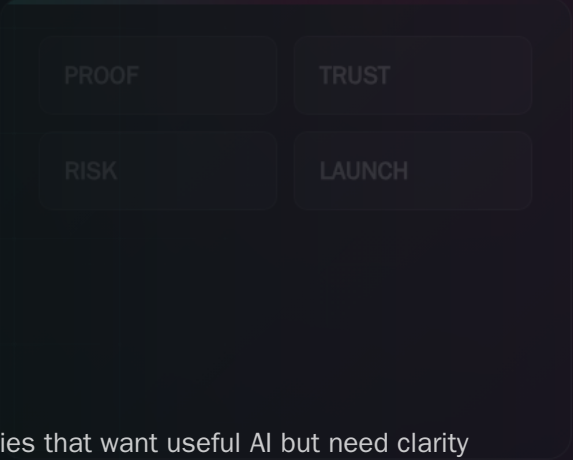
REVIEW BEFORE PRODUCTION

FOLIUM SYSTEMS

AI SYSTEMS AUDIT PACKET

# AI Systems Audit Packet

The AI Systems Audit is the controlled entry point for companies that want useful AI but need clarity first. Folium inspects processes, tools, data, staff impact, cost, risk, integrations, governance, and first-build candidates before recommending a path.



## AUDIENCE

Executives, IT, operations, compliance-aware teams, and first-build sponsors

## PURPOSE

Define what a buyer should expect from a Folium audit and what records it should produce

## UPDATED

May 2026

The audit turns AI pressure into a ranked set of safe, useful, reviewable next steps.

The first output is clarity: process map, source map, risk map, tool inventory, and first-build recommendation.

The audit protects the buyer from overbuying tools, exposing data, or launching without owners.

SCENE

BOUNDARY

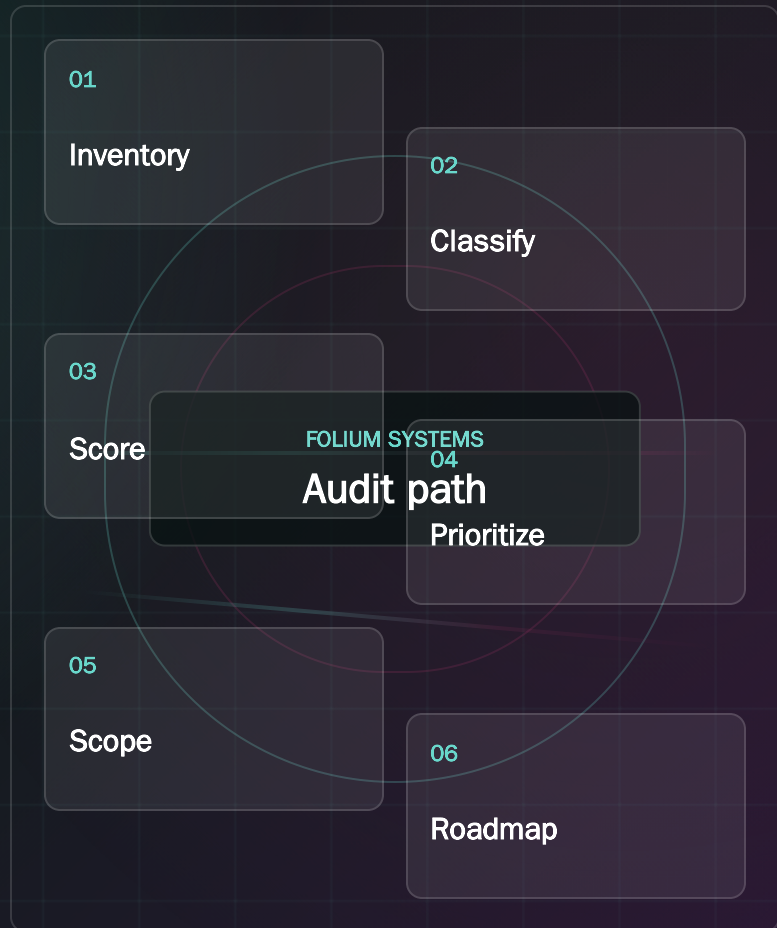
NEXT GATE

AI SYSTEMS AUDIT

# The audit finds the first safe lane before the buyer overbuilds.

The audit packet maps processes, systems, data, tools, subscriptions, staff impact, risk, runtime fit, and first-build candidates.

AUDIT PATH



01

Turns the audit into an operating product.

02

Shows the buyer what records they receive.

03

Connects discovery to a first build instead of a static report.

# 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.

# AI systems audit packet in plain language.

The AI Systems Audit is the controlled entry point for companies that want useful AI but need clarity first. Folium inspects processes, tools, data, staff impact, cost, risk, integrations, governance, and first-build candidates before recommending a path.

RECORD

BOUNDARY

ACTION

## INVENTORY

**Find what already exists**

List AI tools, automations, subscriptions, documents, databases, websites, stores, CRMs, and staff workarounds.

- Tools
- Sources
- Owners

## RISK

**Name what can go wrong**

Map data exposure, wrong-answer risk, vendor sprawl, unsupported workflows, compliance concerns, and staff resistance.

- Data
- Access
- Support

## VALUE

**Rank the first safe lane**

Choose the first workflow by value, feasibility, risk, and reviewability.

- Time
- Cost
- Quality

## PLAN

**Create the next build path**

Produce a scoped first-build, governance, launch, and operations recommendation.

- Scope
- Build
- Operate

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.

# 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
<b>Kickoff</b>	Name sponsor, audit goal, business pressure, constraints, and review audience.	Audit charter.	The audit has a business owner.
<b>Tool inventory</b>	List AI subscriptions, automations, agents, APIs, databases, documents, and shadow workflows.	AI estate inventory.	Unknown tools are surfaced.
<b>Workflow scan</b>	Identify repeated tasks, customer moments, staff bottlenecks, support issues, and document-heavy work.	Workflow candidate list.	At least one useful lane is visible.
<b>Data boundary review</b>	Classify sources by public, internal, private, regulated, confidential, and blocked.	Data boundary map.	Sensitive data is not guessed.
<b>Risk and readiness</b>	Score risk, quality, staff adoption, source readiness, integration difficulty, and support burden.	Readiness matrix.	Launch blockers are clear.
<b>First-build selection</b>	Choose a narrow build with visible value and controlled risk.	First-build recommendation.	The buyer knows what to test first.
<b>Launch posture</b>	Define owner, governance, support, rollback, training, and improvement needs.	Launch readiness outline.	Pilot path is credible.
<b>Audit handoff</b>	Package findings, decisions, diagrams, and recommended next step.	Audit packet.	Leadership can decide from records.

# 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
<b>AI estate inventory</b>	Tools, accounts, automations, prompts, agents, APIs, data stores, and owners.	Shows sprawl, overlap, cost waste, and hidden dependencies.
<b>Workflow candidate matrix</b>	Potential AI use cases ranked by value, feasibility, risk, and reviewability.	Selects the first safe build.
<b>Data boundary map</b>	Source classes, allowed use, blocked data, retention, custody, and provider exposure.	Protects sensitive information.
<b>Risk heatmap</b>	Wrong-answer risk, privacy risk, operational risk, staff risk, customer risk, vendor risk.	Prioritizes controls.
<b>First-build scope</b>	Process, users, systems, data, acceptance criteria, and excluded actions.	Moves from audit to action.
<b>Launch readiness checklist</b>	Owners, support, rollback, training, monitoring, evaluation, and approvals.	Prevents unsupported pilots.

# 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

- Start with one executive sponsor and one operations owner.
- Collect existing tool, system, source, and process information before recommending new software.
- Name duplicated AI subscriptions and unsupported automations.
- Use staff interviews to find hidden workarounds and fear points.
- Classify data before discussing model/provider placement.
- Separate easy productivity gains from high-risk operational dependency.
- Rank workflows by value and readiness, not hype.
- Document what should stop, what should continue, and what should be built next.
- Produce one first-build recommendation that leadership can understand.
- End with a decision meeting, not a vague report.

# 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
<b>Access gate</b>	The audit only uses approved public/internal materials and agreed interviews.	Private or regulated data is being shared casually.
<b>Inventory gate</b>	AI tools and business systems are sufficiently mapped.	Unknown shadow systems remain central to the process.
<b>Readiness gate</b>	A first workflow has owner, sources, value, and risk controls.	Candidate workflows are not ready.
<b>Scope gate</b>	The first build has acceptance criteria and excluded actions.	The build is too broad or unbounded.
<b>Leadership gate</b>	The audit packet supports stop/refine/build/pilot decision.	Decision makers lack records.

# 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 AI tools are already being paid for?
- Which workflows rely on staff memory or spreadsheet patches?
- Where are customers waiting too long?
- Which document or knowledge process is repeated every week?
- Which current AI output is trusted without enough review?
- Which integration would save the most manual handoff?
- Which data should never enter a public AI service?
- Which process would prove value fastest without production risk?

# 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 estate map

A system map of tools, subscriptions, automations, documents, stores, owners, and gaps.

- Tool
- Source
- Owner
- Gap

## Workflow candidate matrix

A scoring grid for value, risk, data readiness, staff adoption, and build difficulty.

- Value
- Risk
- Readiness
- Effort

## Data boundary overlay

A color-coded map for public, internal, private, regulated, confidential, and blocked sources.

- Public
- Internal
- Private
- Blocked

## First-build funnel

A funnel from many ideas to one process with the strongest safe first outcome.

- Ideas
- Rank
- Scope
- Build

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

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

# 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
<b>First 30 days</b>	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.
<b>Days 31-60</b>	Working surface, RAG or agent behavior, integration stub, evaluation cases, browser checks, and staff review.	Sandbox, evaluation file, screenshots, known limits, and repair list.
<b>Days 61-90</b>	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.

# 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
<b>The source is authoritative</b>	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.
<b>The process is ready</b>	A broken process can become a faster broken process when AI is added too early.	Workflow mapping, bottleneck review, owner interview, first-lane narrowing.
<b>The runtime fits the data</b>	Cloud, private, local, and hybrid routes carry different privacy, cost, latency, and support tradeoffs.	Runtime matrix, data classification, provider review, fallback plan.
<b>Staff will adopt the tool</b>	Adoption fails when users do not understand, trust, correct, or benefit from the system.	Training notes, staff review, feedback loop, manager visibility.
<b>Authority is clear</b>	The system can create harm if it sends, updates, approves, or routes without permission.	Permission table, blocked actions, human review, audit trail.
<b>The system can be supported</b>	A useful first build becomes fragile if nobody owns incidents, source updates, or cost review.	Support guide, owner map, release rhythm, rollback trigger.

# 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.

# 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

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

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

# 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.

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

# 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
Audit sponsor	Executive owner	Audit charter and final decision.
Process owner	Operations lead	Workflow map and first-build scope.
Technical reviewer	IT or systems owner	System/source boundary and integration notes.
Folium delivery	Forward engineering lead	Recommendation packet and next-stage plan.

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

# A good audit makes the first build obvious and the bad ideas visible.

Use this packet to prepare for a Folium Systems AI audit and decide what records your team should gather first.

## 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.