



folium


PUBLIC-FACING PDF

REVIEW BEFORE PRODUCTION

FOLIUM SYSTEMS

CUSTOMER-OWNED AI INFRASTRUCTURE

# Customer-Owned AI Infrastructure And Data Residency



Some AI systems should not become an opaque vendor dependency. Folium can plan customer-owned infrastructure, self-hosted services, private databases, local or hybrid inference, customer-controlled audit trails, data-residency boundaries, backups, restore drills, portability, provider-exit paths, monitoring, and support ownership.

## AUDIENCE

Technical buyers, security reviewers, operators, data-sensitive businesses, executives, and teams evaluating private or local AI

## PURPOSE

Show how Folium can design AI infrastructure where the buyer controls data, runtime, audit, recovery, and exit paths

## UPDATED

May 2026

Runtime placement is a business, data, risk, cost, latency, and ownership decision.

Customer-owned infrastructure can include private services, databases, audit custody, local or hybrid inference, backup, restore, and provider-exit paths.

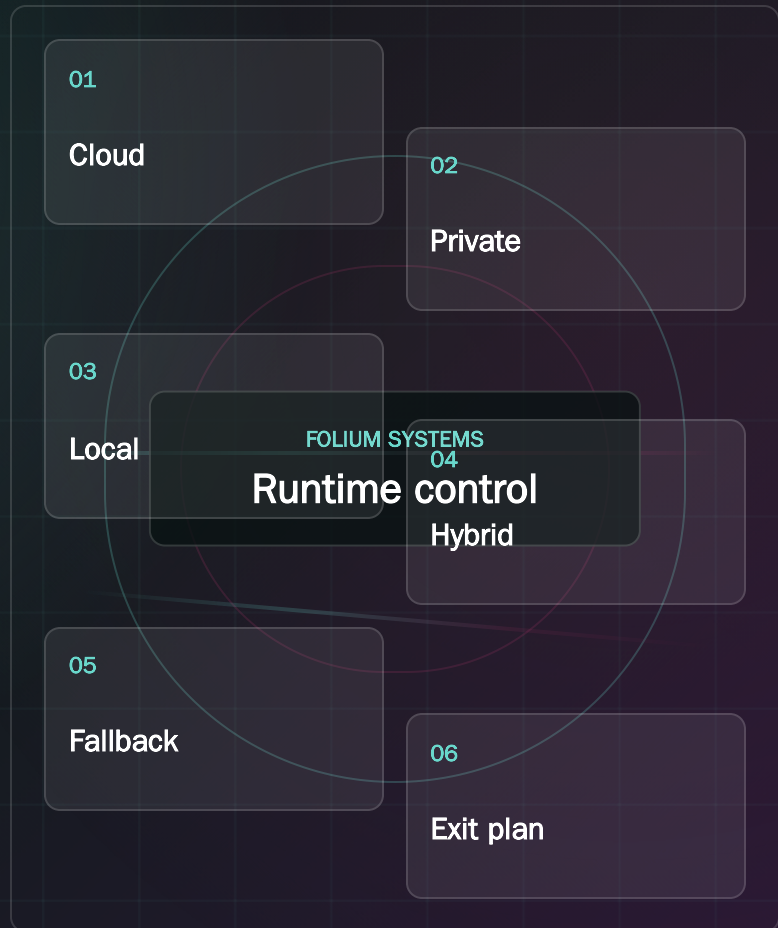
Public proof can describe the method without exposing private topology, credentials, contracts, model names, or customer data.

## RUNTIME PLACEMENT

# Local, private, and hybrid AI choices should be business decisions.

The runtime guide compares cloud APIs, private endpoints, local models, hybrid routes, RAG placement, agent lanes, costs, latency, custody, and fallback.

## RUNTIME CONTROL



01

Shows Folium can reason beyond one provider.

02

Makes privacy, cost, latency, and control tradeoffs visible.

03

Protects data custody before architecture locks in.

# 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

### 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

## OPERATING ROUTE

### 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

## TRUST ROUTE

### 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.

# Customer-owned AI infrastructure in plain language.

Some AI systems should not become an opaque vendor dependency. Folium can plan customer-owned infrastructure, self-hosted services, private databases, local or hybrid inference, customer-controlled audit trails, data-residency boundaries, backups, restore drills, portability, provider-exit paths, monitoring, and support ownership.

## RECORD

## BOUNDARY

## ACTION

## CUSTODY

**Data control is designed**

Data classes, residency needs, logs, audit trails, database custody, retention, and export paths are mapped.

- Data
- Logs
- Export

## RUNTIME

**Placement follows the workflow**

Cloud APIs, private endpoints, local inference, open-source runtimes, customer-owned systems, and hybrid routes are compared.

- Cloud
- Local
- Hybrid

## RESILIENCE

**Restore is planned before crisis**

Backups, restore drills, degraded mode, route contracts, monitoring, and support owners are named.

- Backup
- Restore
- Monitor

## EXIT

**Provider lock-in is not assumed**

Portability, provider-exit notes, replacement routes, source registers, and exportable evidence support control.

- Port
- Exit
- Replace

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>Classify data</b>	Name data classes, residency needs, retention, logs, audit custody, and blocked movement.	Data custody map.	Sensitive material has a route.
<b>Compare runtimes</b>	Evaluate cloud, private, local, open-source, commercial, non-AI, customer-owned, and hybrid routes.	Runtime placement matrix.	Placement is justified.
<b>Design ownership</b>	Assign database custody, audit trail, monitoring, backup, restore, support, and export owners.	Ownership map.	The buyer controls critical records.
<b>Plan continuity</b>	Create degraded-mode, backup, restore drill, model replacement, provider failure, and route fallback paths.	Continuity plan.	Failure has a response.
<b>Package exit</b>	Document source registers, export formats, portability notes, provider-exit paths, and replacement route options.	Provider-exit packet.	Lock-in risk is visible.

# 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>Customer-owned infrastructure map</b>	Services, databases, logs, models, runtime routes, owners, support, and data boundaries.	Shows what the buyer controls.
<b>Data residency and custody plan</b>	Data classes, location, retention, movement, redaction, audit trail, and blocked paths.	Protects sensitive records.
<b>Runtime placement matrix</b>	Cloud, private, local, open-source, commercial, non-AI, and hybrid routes compared.	Prevents one-route assumptions.
<b>Backup and restore drill checklist</b>	Backups, restore owners, test cadence, degraded mode, and recovery evidence.	Makes recovery practical.
<b>Provider-exit packet</b>	Export paths, source registers, route contracts, replacement options, and portability notes.	Reduces vendor lock-in.

# 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

- Classify data before choosing a model route.
- Compare cloud, private, local, open-source, commercial, non-AI, customer-owned, and hybrid options.
- Name database, log, audit, backup, restore, and support owners.
- Plan degraded mode and provider failure behavior.
- Design export and provider-exit before dependency grows.
- Keep private customer data, credentials, contracts, private topology, private model names, and internal project labels out of public review material.
- Use explicit public-safe boundaries when work touches providers, money, regulated-adjacent review, customer-impacting actions, or live operating authority.
- Leave behind records that a buyer can inspect: source, scope, owner, date, evidence class, known limits, blocked states, and next-stage gate.

# 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>Custody gate</b>	Data class, residency, retention, logs, and movement are documented.	Data route is assumed.
<b>Runtime gate</b>	Placement is selected by workflow, risk, cost, latency, quality, and support.	One vendor path is defaulted.
<b>Restore gate</b>	Backup and restore drill path exists.	Recovery is theoretical.
<b>Exit gate</b>	Export, portability, and provider-exit path are known.	The buyer is trapped.
<b>Boundary gate</b>	Private topology, credentials, contracts, and customer data stay out of public proof.	Sensitive implementation detail leaks.

# 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

- What data must remain customer-controlled?
- Which logs and audit trails must stay private?
- Which workloads require local, private, cloud, or hybrid routes?
- What happens if a model provider is down or too expensive?
- How would the buyer exit or replace a provider?
- Which proof can be public and which proof must stay private?
- Which owner can approve the next stage?
- What would make the workflow pause, rollback, or stay blocked?

# 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

## Runtime placement matrix

Cloud, private, local, open-source, commercial, customer-owned, non-AI, hybrid.

- Cloud
- Private
- Local
- Hybrid

## Custody map

Data, logs, audit trail, database, source registers, exports, and retention.

- Data
- Log
- Audit
- Export

## Provider-exit loop

Export, route contract, replacement option, restore drill, degraded mode, support owner.

- Export
- Replace
- Restore
- Support

# 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
<b>Business owner</b>	Customer-owned AI infrastructure priority, acceptance criteria, and stop/continue decision.	Workflow decision note and next-stage gate.
<b>Technical owner</b>	Architecture, integrations, runtime placement, data movement, monitoring, and fallback.	Architecture map, route contract, and support guide.
<b>Risk or review owner</b>	Data boundaries, authority classes, blocked actions, evidence, and release conditions.	Boundary map, gate table, and evidence packet.
<b>Operator or support owner</b>	Daily workflow state, exception handling, escalation, recovery, and improvement rhythm.	Queue map, incident route, and operating handoff.
<b>Folium delivery lead</b>	Public-safe build coordination, evaluation, known limits, launch room, and handoff completeness.	Review file, release notes, and improvement backlog.

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

# Customer-owned AI is an architecture decision, not a slogan.

Use this packet when data residency, private infrastructure, local AI, portability, or vendor exit matters.

## 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.