PUBLIC-SAFE PACKET | PROOF BEFORE PRODUCTION | FOLIUM SYSTEMS

TRUST PACKET

Folium Systems Trust Packet

This packet is built to be worth printing. It explains how Folium creates trust before private data, live systems, regulated actions, or production dependency enter the room.

AUDIENCE
Executives, operators, security reviewers, procurement, staff leaders

PURPOSE
Make proof boundaries, data handling, AI limits, and launch trust visible

UPDATED
May 2026

- Trust begins when proof, sandbox, pilot, and production are separated.
- Folium treats data boundaries, permissions, ownership, and stop signs as part of the build.
- AI adoption succeeds when staff understand the system and reviewers can inspect the evidence.

01

TRUST MODEL

Trust is built by separating proof, access, and authority.

AI work becomes dangerous when a polished demonstration is treated like production readiness. Folium's trust model keeps the buyer's confidence tied to evidence: what is public, what is sandboxed, what has access, what can act, and who owns the next decision.

EVIDENCE	BOUNDARY	ACTION

PUBLIC PROOF

Safe to inspect without private access

Public pages, packets, tools, and proof stories should explain the model without requiring customer data, production credentials, or hidden claims.

SANDBOX REVIEW

Real enough to evaluate, limited enough to control

A sandbox or redacted proof lets staff inspect behavior, handoffs, limits, and value before private systems enter the workflow.

PILOT AUTHORITY

Controlled access with named owners

A pilot needs approved systems, data classes, permissions, owners, support paths, rollback, and user training.

PRODUCTION TRUST

Earned by operating evidence

Production depends on evidence over time: quality gates, monitoring, incidents, updates, cost control, and staff feedback.

02

OPERATING PRINCIPLES

The principles Folium applies before AI enters a workflow.

These principles are meant to be simple enough for executives and concrete enough for technical reviewers.

CHECKLIST | OWNER PATH | RELEASE SIGNAL

- Do not confuse a beautiful demo with a launch-ready system.
- Do not expose private or regulated data before the buyer approves the data boundary.
- Do not let AI take sensitive action without explicit permission, role design, and human review where needed.
- Do not hide known limits; record them so the next stage can repair or accept them knowingly.
- Do not force staff to trust a mystery; train them on what AI can do, cannot do, and where humans remain responsible.
- Do not treat compliance as a slogan; build quality gates, evidence, controls, and customer-side review paths.
- Do not make the buyer dependent on one model, one vendor, one cloud, or one unreviewed automation path.
- Do not move from proof to production without a support, incident, rollback, and improvement plan.

> Folium's practical trust position is simple: prove first, govern before access, and keep humans in the places where judgment, responsibility, or regulated impact matters.

03

DATA HANDLING

Data trust starts with knowing what AI is allowed to see.

Before a customer-specific workflow is built, Folium helps classify data and choose the safest route for each workflow. The goal is to avoid accidental exposure, unnecessary vendor sharing, and unclear retention.

DECISION GRID | REVIEW LENS | NEXT GATE

DATA CLASS	TYPICAL HANDLING	TRUST QUESTION
Public	Can be used in public proofs, examples, and general website content.	Does this reveal anything the business has not approved publicly?
Internal	Can support a scoped sandbox if owners approve access and retention.	Which staff roles should see it and for what purpose?
Customer	Requires stronger access rules, redaction, purpose limits, and review.	Could exposure harm a customer, relationship, or legal obligation?
Financial or regulated	Requires customer-side legal/compliance review before use.	What standards, disclosures, retention rules, and audit needs apply?
Secrets and credentials	Should not enter model prompts, public packets, or general logs.	How are keys, tokens, passwords, and provider credentials stored and rotated?
Blocked	Must stay out of the workflow until explicit approval changes scope.	Who can approve a change and what evidence must exist first?

04

AI OUTPUT LIMITS

Folium separates useful AI assistance from unapproved AI authority.

Most buyers do not need AI to run wild. They need AI to make work faster, clearer, more consistent, and easier to review. The permission model is how Folium keeps that difference visible.

EVIDENCE | BOUNDARY | ACTION

Safe early actions

Draft, summarize, explain, classify, compare, retrieve, route, and recommend when the source and limitations are clear.

- Draft customer responses
- Summarize internal notes
- Find relevant policy or product detail
- Prepare a review packet

Actions needing review

Anything affecting money, customer treatment, legal claims, employee decisions, permissions, records, or provider actions should pass through a human gate.

- Send final notices
- Approve terms
- Change records
- Trigger external provider actions

Blocked until approved

Production credentials, regulated decisions, hidden data sharing, unlogged tool use, and irreversible actions stay blocked unless scope and approval change.

- Unapproved live actions
- Credential exposure
- Regulated final decisions
- Silent data movement

Evidence required

The system needs logs, screenshots, known limits, failed-case review, and clear owner signoff before authority expands.

- Quality gate
- Owner signoff
- Rollback path
- Incident route

05

SECURITY POSTURE

Security is a design conversation, not a final checklist.

Folium helps buyers ask the right questions before access expands. Security review should map systems, users, vendors, logs, credentials, retention, dependencies, and failure paths.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- Map the systems involved and the direction data flows between them.
- Name the runtime for each AI function: public-demo, cloud API, private endpoint, local model, hybrid route, or future production service.
- Define how prompts, retrieved sources, outputs, logs, uploaded files, and transcripts are retained or discarded.
- Separate human-readable configuration from secrets and credentials.
- Name who approves provider use, API access, database access, and model/runtime placement.
- Prepare fallback behavior when an API, model, retrieval store, commerce system, or legacy app fails.
- Document support ownership, incident severity, communication path, and rollback trigger.
- Keep public proof packets free of internal IPs, private tokens, customer data, and secret-like values.

06

ACCESSIBILITY AND ADOPTION

AI is only useful if people can understand and use it.

Trust also depends on the people who will work with the system. Folium treats accessibility, staff clarity, training, and role confidence as part of the trust packet.

EVIDENCE | BOUNDARY | ACTION

Accessible public surfaces

Pages and tools should be readable, keyboard-friendly, responsive, and usable on desktop, tablet, and mobile.

- Readable type
- Touch-friendly controls
- No hidden essential actions
- No horizontal overflow

Plain-language AI explanations

Staff should know what the AI does in everyday language, not only in technical architecture terms.

- What it can do
- What it cannot do
- When to escalate
- How to correct it

Adoption without fear

Folium positions AI as a force multiplier for staff knowledge, not a mystery tool imposed without training.

- Role clarity
- Feedback loops
- Training packets
- Human review points

Operations ownership

A workflow needs people who own quality, updates, incidents, source freshness, and improvement after launch.

- Workflow owner
- Knowledge owner
- Support owner
- Executive sponsor

07

LAUNCH DISCIPLINE

The safest AI launch has stop signs.

A trust packet should tell a buyer when to move forward and when to pause. The stop signs are not weakness; they protect speed by preventing expensive mistakes.

DECISION GRID | REVIEW LENS | NEXT GATE

SIGNAL	WHY IT MATTERS	FOLIUM MOVE
No owner	A workflow without an owner becomes an orphan after the demo.	Assign business, technical, support, and executive ownership before pilot.
Unclear data boundary	The system may expose or misuse information without a shared data plan.	Map data classes, blocked data, retention, redaction, and approval.
Unapproved live action	AI may trigger customer, financial, legal, or operational impact without authority.	Separate draft, recommend, route, execute, and human approval.
No quality gate	A polished answer can hide weak retrieval, reasoning, or tool routing.	Run evaluation, browser proof, known-limits review, and failed-case repair.
No rollback path	A problem becomes harder to contain after staff depend on the workflow.	Define degraded mode, rollback trigger, incident route, and owner communication.
Staff not ready	Adoption fails when users fear the system or do not understand the boundaries.	Prepare training, plain-language guides, feedback loops, and support.

08

NEXT STEP

Trust is the operating system around the AI system.

Use this packet to bring security, leadership, operations, and staff into the same conversation before access expands. The next move should be a scoped review of one workflow, one data boundary, and one evidence path.

Bring the workflow

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

Separate proof from production

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

Ask for the evidence

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