PUBLIC-SAFE PACKET | PROOF BEFORE PRODUCTION | FOLIUM SYSTEMS

SECURITY AND PROCUREMENT PACKET

Folium Security Procurement Review

This packet is designed for the people who have to say yes responsibly: security, procurement, IT, counsel, operations, leadership, and workflow owners. It turns AI review into a staged operating process instead of a trust-me demo.

AUDIENCE	PURPOSE	UPDATED
Security, procurement, IT, counsel, operations, leadership	Prepare buyer-ready review questions, artifacts, gates, and owner responsibilities	May 2026

> Access should expand in stages: public proof, scoped discovery, sandbox, architecture review, pilot, production plan.

> AI permissions must separate drafting, retrieval, recommendation, routing, execution, blocking, and escalation.

> Every serious review should leave behind evidence, known limits, owners, and a next decision gate.

01

REVIEW FRAME

Security review should begin before the workflow gets private access.

AI procurement gets cleaner when the buyer can see scope, data flow, runtime placement, permissions, evidence, support, and customer responsibilities before credentials or sensitive data enter the build.

EVIDENCE | BOUNDARY | ACTION

SCOPE

What is the workflow?

Name the business process, users, systems, data classes, exclusions, reviewers, assumptions, and success criteria.

BOUNDARY

What can the AI see?

Separate public, internal, customer, regulated, confidential, secret, and blocked data before the proof expands.

RUNTIME

Where will the AI run?

Choose public-demo proof, cloud API, private endpoint, local model, hybrid route, or production service based on risk and value.

AUTHORITY

What can the AI do?

Define draft, retrieve, recommend, route, execute, block, escalate, and human-approve actions.

02

PROCUREMENT MATRIX

The questions buyers should not have to chase.

Folium turns common procurement questions into a review packet instead of leaving them scattered across calls, emails, screenshots, and assumptions.

DECISION GRID | REVIEW LENS | NEXT GATE

QUESTION	FOLIUM ANSWER PATTERN	EVIDENCE TO PREPARE
What data is involved?	Data classification, approved sources, blocked data, retention, and redaction.	Data boundary map, source list, redaction notes, retention stance.
Where does it run?	Runtime placement by workflow and sensitivity.	Runtime decision table, cost/privacy rationale, fallback route.
Who can use it?	Role-based access and permissioned workflow steps.	User roles, approval chain, blocked actions, owner signoff.
What can it change?	Separate draft/recommend from execute/update/send.	Tool permission table, human-review gates, audit trail plan.
How is quality measured?	Evaluate workflow behavior, not only answer polish.	Eval cases, browser checks, known limits, failed-case notes.
What happens on failure?	Support path, incident class, rollback, degraded mode, and repair loop.	Runbook, severity ladder, rollback trigger, communication path.

03

ACCESS LADDER

Access should expand only when evidence supports it.

The safest AI implementation does not jump from idea to production. It moves through staged access with explicit gates.

DECISION GRID | REVIEW LENS | NEXT GATE

STAGE	ALLOWED ACCESS	GATE TO ADVANCE
Public proof	Public pages, packets, screenshots, public tools, and public-safe examples.	Buyer understands value and agrees to scoped discovery.
Scoped discovery	Customer describes workflow, systems, risks, users, blockers, and desired outcome.	Scope is narrow enough for a safe sandbox proof.
Redacted sandbox	Approved safe samples, synthetic examples, public policies, or redacted documents.	Behavior is useful enough for technical architecture review.
Architecture review	System map, runtime plan, provider plan, permissions, logging, and support assumptions.	Security, operations, leadership, and owners approve a controlled pilot.
Controlled pilot	Limited users, limited data, limited actions, monitoring, support, and rollback.	Evidence shows readiness or identifies repairs before production.
Production plan	Approved live access only after legal, security, operational, and owner gates.	Final responsibilities, monitoring, training, and change process are accepted.

04

SECURITY CONTROLS

The review should cover the whole operating path.

Procurement should not evaluate a model in isolation. The real risk lives across data, software, people, providers, logs, permissions, and change management.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- Document source systems and destination systems before any integration is approved.
- Separate public packets from customer-specific evidence that belongs in controlled review.
- Define what gets logged, where logs live, who can access them, and how long they remain.
- Keep secrets, API keys, customer credentials, provider tokens, and passwords outside public files and prompts.
- Use least-privilege tokens and service accounts when automation touches external systems.
- Name who can approve runtime choices, provider choices, data access, production credentials, and launch timing.
- Prepare fallback behavior for model outage, retrieval failure, tool failure, stale source, and user rejection.
- Review dependency posture, headers, public file handling, PDF downloads, and privacy-safe analytics before launch.

05

AI PERMISSION MODEL

Not every AI action belongs in the same permission bucket.

Folium helps customers design permissions so the business can gain speed without giving AI unreviewed authority.

DECISION GRID | REVIEW LENS | NEXT GATE

ACTION CLASS	EXAMPLES	REVIEW POSTURE
Draft	Write first-pass emails, FAQs, summaries, work notes, task descriptions.	Generally safe in sandbox when sources and limits are visible.
Retrieve	Find policy, procedure, product, order, ticket, document, or knowledge-base detail.	Requires source boundaries, freshness notes, and access controls.
Recommend	Suggest next workflow step, routing choice, response, escalation, or evidence packet.	Needs human accountability for sensitive outcomes.
Route	Send work to queue, owner, reviewer, or next system without changing final records.	Needs audit trail and blocked-route rules.
Execute	Update record, send official message, place order, issue refund, approve account, change access.	Requires explicit approval, logging, rollback, and human review where impact is material.
Block or escalate	Refuse unsafe request, route to human, stop launch, flag missing data.	Should be encouraged when the system lacks authority or evidence.

06

OPERATING OWNERSHIP

Security review fails when nobody owns the workflow after launch.

Folium pushes ownership into the packet so the buyer can see who maintains quality, sources, support, and change after the exciting demo ends.

EVIDENCE | BOUNDARY | ACTION

Business owner

Owns the workflow outcome, success criteria, staff adoption, and priority decisions.

Technical owner

Owns integration choices, runtime placement, access, deployment, monitoring, and fallback.

Knowledge owner

Owns source accuracy, document freshness, policy updates, and retrieval boundaries.

Security reviewer

Owns data sensitivity, credentials, access control, logs, retention, and external provider posture.

Support owner

Owns incidents, user questions, rollback, degraded mode, and post-incident improvement.

Executive sponsor

Owns the decision to continue, stop, repair, expand, or move into production planning.

07

RED FLAGS

These signs should slow the deal down, not speed it up.

A serious AI partner should welcome these checks. They prevent a buyer from adopting an impressive but ungoverned system.

CHECKLIST	OWNER PATH	RELEASE SIGNAL

- The demo requires private customer data before a data boundary is approved.
- The vendor cannot explain where prompts, files, outputs, and logs live.
- The system can take live action without a named approval path.
- No one can show failed cases, known limits, or repair decisions.
- The buyer has no owner for source freshness, support, incidents, or rollback.
- Staff are expected to trust the AI without role clarity or training.
- Procurement, security, counsel, and operations are introduced after production pressure exists.
- The proposal assumes one AI vendor, one model, or one cloud can solve every workflow.

08

REVIEW PACKET OUTPUT

A good review should leave behind artifacts.

After a serious review, the buyer should not be left with a vague call recap. The next decision should be supported by clear documents and evidence.

EVIDENCE | BOUNDARY | ACTION

Scope packet

Workflow, users, systems, exclusions, assumptions, desired outcome, and decision criteria.

Boundary map

Data classes, approved sources, blocked data, provider handoffs, retention, and secrets handling.

Permission table

Draft, retrieve, recommend, route, execute, block, escalate, and human-approve rules.

Evidence packet

Screenshots, browser checks, eval cases, known limits, failed-case notes, and repair plan.

Operations runbook

Owners, support, incidents, rollback, degraded mode, monitoring, and improvement cadence.

Decision gate

A plain recommendation: stop, repair, continue sandboxing, pilot, production-plan, or move to AI operations.

09

NEXT STEP

A cleaner review creates a faster yes or a safer no.

Use this packet before private access expands. Folium can help turn one workflow into a reviewable AI path with boundaries, evidence, owners, and a launch decision the business can defend.

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.