SurgeONE.ai
AI-Powered. Practitioner-Led

Regulatory Exams are Changing

How AI Helps Firms Stay Ready

A practical guide for RIAs, broker-dealers, compliance leaders, and technology teams

Executive Summary

Regulatory exam readiness can no longer be treated as a last-minute document collection exercise. As firms use more technology, vendors, AI-enabled workflows, and cybersecurity controls, they must be able to demonstrate ongoing supervision, clear ownership, consistent documentation, and evidence that controls operate in practice.

AI can reduce the manual burden of locating records, comparing documents, monitoring workflows, identifying gaps, and preparing first drafts. It cannot replace legal or compliance interpretation, accountable decision-making, or experienced human review. A defensible program uses approved tools, governed data, traceable sources, defined review requirements, testing, and retained evidence.

The goal is continuous readiness: policies match actual practices, controls create evidence as work occurs, exceptions trigger remediation, and responsible people remain accountable for conclusions and actions.

Regulatory Readiness Starts Before Exams Begin

www.surgeone.ai
info@surgeone.ai
+1 650-294-8889

SurgeONE.ai
AI-Powered. Practitioner-Led

1. Why Regulatory Readiness Is Changing

Traditional exam preparation often begins when a request arrives. Teams then collect policies, vendor files, system reports, approvals, exception records, and supporting communications. This approach is difficult to sustain because relevant evidence may be distributed across systems, ownership may be unclear, and decisions may not have been documented when they were made.

A continuous readiness model connects five elements: obligation, control, evidence, review, and remediation. An obligation is translated into an operating control. The control produces time-stamped evidence. A qualified person reviews the result. Any exception is assigned, tracked, and closed with supporting proof.

A firm should be able to answer the following questions without relying on individual memory:

- What requirement or policy does the control address?
- Who performs and reviews the control?
- How often does it operate?
- Where is the execution evidence stored?
- What happened when an exception occurred?
- Who approved remediation and closure?

Build Evidence Before the Exam Begins

SurgeONE.ai
AI-Powered. Practitioner-Led

2. Where AI Can Help

AI is most useful for repeatable, high-volume activities where outputs can be checked against authoritative sources. Practical use cases include:

- Evidence retrieval: searching approved repositories and ranking potentially relevant records.
- Document comparison: identifying differences or possible gaps across policies, procedures, questionnaires, and current practices.
- Workflow monitoring: detecting overdue reviews, missing approvals, incomplete records, or broken handoffs.
- Exception triage: grouping and prioritizing alerts or missing evidence for human review.
- Regulatory change support: suggesting policies and controls that may be affected by a new notice or requirement.
- Drafting support: preparing initial versions of evidence indexes, control narratives, test summaries, and response materials.

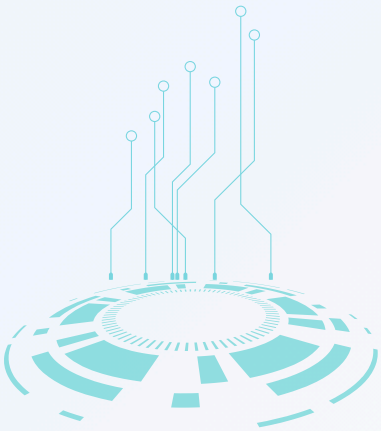

AI Makes Compliance Work Smarter and Faster

Every output should be treated according to its risk. Internal summaries, metadata tagging, duplicate detection, and draft checklists are generally more suitable starting points than client-facing representations, regulatory conclusions, automated account decisions, or autonomous remediation.

SurgeONE.ai
AI-Powered. Practitioner-Led

3. Defensible AI Governance

A defensible AI-enabled compliance program is a layered control system, not only an AI policy. It should address the use case, data, model or vendor, workflow, human decision, and retained evidence.

- **Governance and Inventory** - Maintain an inventory of AI tools and embedded AI features, including the approved purpose, business owner, users, data types, integrations, risk tier, vendor, and current status. Unapproved tools and prohibited uses should be clearly defined and, where practical, technically restricted.
- **Data Controls** - Define what information may enter each tool. Apply data classification, minimization, access control, retention, deletion, and residency requirements. Sensitive client, employee, business, or regulatory data should not be submitted to an unapproved public AI service.
- **Model and Vendor Controls** - Document the model or service used, material limitations, security controls, data-use terms, subprocessors, change-notification practices, business continuity, and the firm's ability to export or delete information. Due diligence should reflect the actual use case and configuration, not only the vendor's general security statement.
- **Workflow Controls** - Use approved source repositories and prompt templates where appropriate. Require source references for material statements. Define confidence or exception thresholds, mandatory review points, escalation paths, stop-use criteria, and the people authorized to approve or override results.
- **Evidence and Auditability** - Retain enough information to explain and reproduce material work. Depending on risk, this may include the use case, input, source documents, output, tool or model context, reviewer, review date, decision, override reason, approval, and final disposition.

SurgeONE.ai
AI-Powered. Practitioner-Led

4. Risk-Based Human Review

Human review is meaningful only when the reviewer has authority, subject-matter competence, context, and access to the supporting sources. A simple approval click is not sufficient if the reviewer cannot evaluate the basis of the output.

A practical risk model may include four levels:

- **Assistive:** internal, reversible, and low-sensitivity use. Require an approved tool, user review, and basic logging.
- **Advisory:** output influences analysis or prioritization. Require grounding, testing, documented review, and override records.
- **Decision-support:** output may have a material client, financial, regulatory, or operational effect. Require independent validation, enhanced approval, monitoring, and formal change control.
- **Restricted or prohibited:** the data, purpose, or level of autonomy is unacceptable. Apply technical restrictions, detection, escalation, and periodic attestation.

Validation should examine source fidelity, completeness, consistency, privacy, security, operational resilience, inappropriate outcomes, reviewer overrides, and changes in performance. Material changes to a model, vendor, data source, prompt template, integration, or workflow should trigger reassessment, as well as require documentation and written approval.

AI Decisions Need Qualified Human Oversight

www.surgeone.ai
info@surgeone.ai
+1 650-294-8889

SurgeONE.ai
AI-Powered. Practitioner-Led

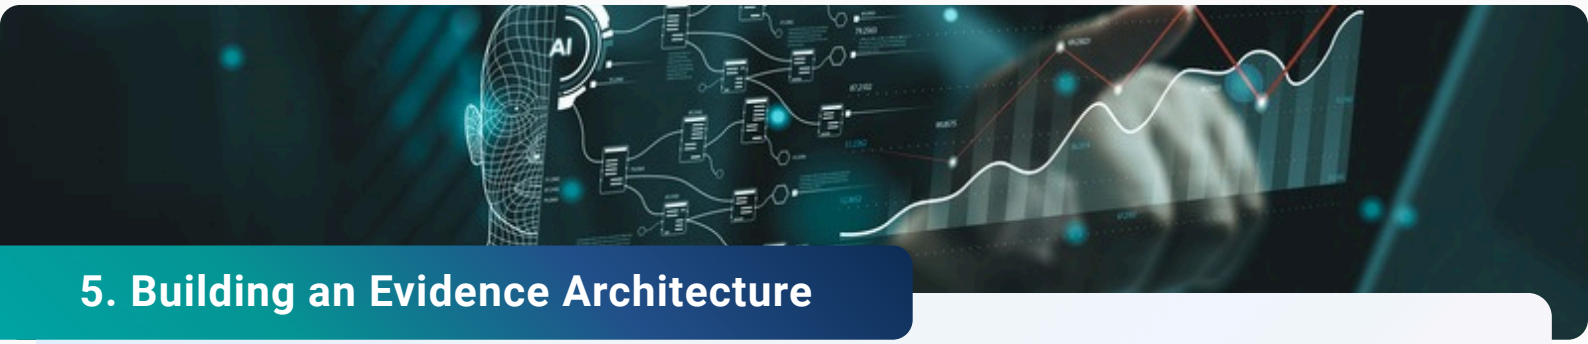

5. Building an Evidence Architecture

Exam readiness depends on the quality and connectivity of records. Evidence should show what happened, when it happened, who or what performed the activity, what was reviewed, what decision was made, how exceptions were resolved and who gave final approval.

Core evidence records include:

- **Policy or procedure:** owner, version, approval date, effective date, and mapped obligations.
- **Control:** purpose, frequency, performer, reviewer, system, and testing method.
- **Execution record:** timestamp, input or population, status, performer or system identity, output, and evidence location.
- **AI interaction:** approved use case, source references, output, reviewer, decision, and disposition.
- **Exception or issue:** severity, root cause, owner, target date, interim action, and closure evidence.
- **Vendor record:** service, data access, risk tier, due diligence, contractual controls, and monitoring date.

High-quality evidence is authentic, attributable, complete, reproducible, protected, and retrievable. The record should include failed results, overrides, exceptions, and remediation, not only successful outcomes.

A model-generated explanation is not evidence by itself; defensibility comes from retained sources, system records, testing, review, and decision documentation.

SurgeONE.ai
AI-Powered. Practitioner-Led

6. Connecting Compliance, Cybersecurity, Vendor, and AI Risk

AI readiness should not operate separately from compliance, cybersecurity, data governance, vendor management, and operational resilience. A compliance use case may depend on a third-party model, connectors to production systems, sensitive data, identity controls, and downstream workflow actions. Weakness in any layer can affect the reliability and defensibility of the final result.

The integration boundary deserves particular attention. Apply least-privilege access, separate read and write permissions, restrict autonomous actions, log connector activity, protect credentials, and require written approval for material configuration changes.

Protect Every Layer Behind AI Decisions

Questions for an AI or technology vendor should cover:

- What data is collected, retained, deleted, or used for model training?
- Is the data that is collected anonymized?
- How is data encrypted, segregated, accessed, and logged?
- Which models and subprocessors are used, and where is data processed?
- How are material service or model changes communicated?
- What independent assurance, testing, or control evidence is available?
- How are incidents, outages, recovery, termination, and data export handled?
- Can the firm retain prompts, sources, outputs, approvals, and administrative events when needed?

SurgeONE.ai
AI-Powered. Practitioner-Led

7. A Practical 90-Day Roadmap

DAYS 1–30 — DISCOVER AND GOVERN

- Inventory AI tools, embedded features, vendors, and current use cases.
- Assign accountable owners and identify sensitive data flows.
- Define approved, restricted, and prohibited uses.
- Map key obligations, controls, evidence sources, and known gaps.

DAYS 31–60 — DESIGN AND PILOT

- Assign risk tiers and minimum control requirements.
- Choose one low-risk, reversible pilot using approved internal data.
- Define the baseline, validation method, reviewer, evidence, and escalation process.
- Complete use-case-specific vendor and security review.

DAYS 61–90 — OPERATE AND TEST

- Monitor workflow status, exceptions, overrides, and overdue remediation.
- Test a sample of outputs for source fidelity, completeness, and review quality.
- Retain evidence and report material risks to management.
- Update controls based on pilot results before expanding the use case.

A Practical 90-Day AI Readiness Roadmap

SurgeONE.ai
AI-Powered. Practitioner-Led

8. Compact Exam-Readiness Checklist

- [ ] AI and technology use cases are inventoried and assigned to accountable qualified owners.
- [ ] Policies and procedures reflect actual workflows and currently approved tools.
- [ ] Key controls identify the performer, reviewer, frequency, system, and required evidence.
- [ ] Sample records show that controls operated, including exceptions and remediation.
- [ ] Material AI outputs are grounded in approved sources, validated, and retained with reviewer disposition.
- [ ] Sensitive data flows, access, retention, deletion, and residency requirements are documented.
- [ ] Vendor due diligence is current and linked to the actual service, configuration, and use case.
- [ ] Cybersecurity access, logging, incident response, continuity, and recovery controls are tested.
- [ ] Material model, vendor, data, prompt, integration, and workflow changes trigger reassessment.
- [ ] Users understand permitted use, prohibited data, verification duties, and escalation paths.
- [ ] The firm can assemble a coherent response package from indexed, controlled sources.
- [ ] Open risks, overdue remediation, readiness measures, and significant overrides are reported to management.

SurgeONE.ai
AI-Powered. Practitioner-Led

Conclusion

Questions to Ask Now

- Where is AI already being used, including embedded features and unapproved tools?
- Which compliance evidence such as supervisory approvals would be hardest to produce today?
- Which controls depend on spreadsheets, email, or individual memory?
- Which AI outputs could materially affect clients, filings, disclosures, or regulatory responses?
- What low-risk pilot could improve evidence without increasing decision risk?

Conclusion

AI can make regulatory readiness more efficient by improving retrieval, comparison, monitoring, drafting, and gap identification. It should not be used to transfer accountability or replace qualified judgment. The stronger operating model connects policies, controls, evidence, human review, supervisory approvals, exceptions, remediation, cybersecurity, vendor oversight, and data governance. The central principle is simple: evidence should be created by the process, not reconstructed after the fact. Firms that build proof into daily work are better positioned to respond clearly, consistently, and defensibly when regulatory scrutiny occurs.

Important Notice

www.surgeone.ai info@surgeone.ai +1 650-294-8889