COGNICOR

ART:
A Framework for Building Useful, Compliant AI

Accuracy. Relevance. Trust. ART offers CIOs, compliance leaders, and technology partners a clear blueprint for building responsible AI systems in wealth management and beyond. By embedding verifiable accuracy, contextual relevance, and transparent trust, ART ensures AI solutions remain auditable, explainable, and aligned with advisor and client outcomes.

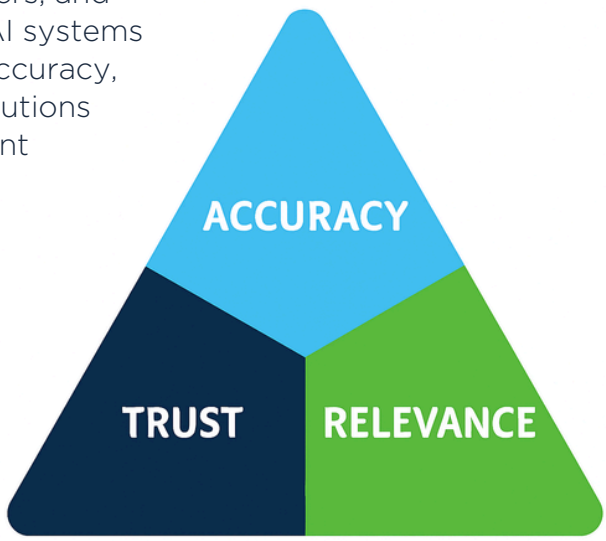

1. ACCURACY — VERIFIABLE TRUTH

Accuracy ensures every AI output is factual, verifiable, and traceable to its data source. For regulated industries like wealth management, accuracy is not optional—it's the foundation of compliance, advisor confidence, and client trust.

Key Design Principles:

- Source Integrity & Provenance – Every statement is traceable to its data origin with retrieval time and confidence metadata.
- Conflict Awareness – Surface disagreements between systems instead of averaging or guessing.
- Temporal Validity – Show when each fact was last verified and which dataset or version it came from.

CIO Evaluation Questions:

- Can the platform display the exact data sources supporting each response?
- How are discrepancies between data systems identified and surfaced?
- Is there a timestamp record for every data source in the response?

KPIs:

- Evidence coverage rate (% of answers with sources)
- Conflict detection rate
- Hallucination error rate (target: < 0.1%)

In Practice:

- Portfolio Reconciliation: The AI flags a mismatch between custodian and CRM-reported AUM, citing exact data fields and time stamps.
- Regulatory Reporting: FINRA query automation uses only verified source documents to draft compliant responses.

2. RELEVANCE — CONTEXTUAL & IMPACT DRIVEN

Relevance ensures AI insights are meaningful, contextual, and prioritized based on business impact and user preference. A useful framework is IURRE — Impact, Urgency, Risk, Relationship, and Effort — to rank and surface the most valuable insights to decision-makers.

Key Design Principles:

- Prioritize using IURRE (Impact, Urgency, Risk, Relationship, Effort)
- Personalize recommendations by user role and firm policy
- Continuously learn from advisor behavior and feedback

CIO Evaluation Questions:

- How does the AI quantify and rank insights by priority?
- Can relevance weights be tuned without developer support?
- Does the system adapt over time based on user actions?

KPIs:

- Advisor adoption rate of AI recommendations
- Time-to-action from surfaced insights
- Reduction in missed client tasks or opportunities

In Practice:

- Client Meeting Prep: AI flags a 529 funding gap with high impact and low effort, surfacing it for immediate discussion.
- Task Prioritization: The Copilot reminds advisors of pending RMD notices or private fund capital calls ranked by urgency.

SINDHU@COGNICOR.COM

COGNICOR

3. TRUST & SECURITY — EXPLAINABLE & AUDITABLE AI

Trust emerges when users can see not only what the AI recommends but how it got there. In a regulated industry, explainability drives both compliance and user adoption.

Key Design Principles:

- Explainability & Auditability – Expose reasoning paths, inputs, and evidence.
- Data Privacy & Security – Honor consent boundaries, encrypt all data, enforce least-privilege access.
- Governance & Human Oversight – Retain human review, approval, and override options.

CIO Evaluation Questions:

- Does each output display its data sources, time stamps and reasoning path?
- Can compliance export a full audit trail on demand?
- How are data and model versions managed for traceability?

KPIs:

- Evidence completeness score
- Compliance acceptance rate
- Advisor trust/confidence index (NPS)

In Practice:

- *Audit-Ready Reporting:* Compliance exports an evidence bundle showing every data source used in client performance summaries.
- *Model Governance:* Each AI model update is version-controlled, with records of rationale, training data, and reviewer approval.

IMPLEMENTATION CHECKLIST

- ○ Embed transparency in every AI interface
- ○ Implement one-click exportable audit bundles
- ○ Track advisor trust via post-interaction surveys

IMPLEMENTATION PLAYBOOK

Embedding ART across governance, technology, and experience design requires a coordinated strategy. The following playbook outlines practical steps for operationalizing ART principles.

Data & Governance

- Map systems of record and establish precedence rules.
- Apply encryption, RBAC, MFA, and activity logging.
- Monitor model drift and maintain explainability documentation.

Product & UX

- Enforce a "no evidence, no answer" policy.
- Visualize data conflicts with one-click resolution options.
- Allow users to adjust relevance weights interactively.

Quality Operations

- Conduct weekly accuracy audits and classify errors.
- Schedule quarterly compliance reviews for evidence and traceability.
- Continuously track user trust, adoption, and engagement metrics.

Stage	Description	Key Milestones
Foundational	Establish basic evidence tracking and auditability	50% evidence coverage, initial relevance tuning
Operational	Integrate ART principles across workflows and systems	Weekly audits, firm-wide relevance scoring
Scaled	Continuous learning, automation, and compliance integration	<1% hallucination rate, automated audit exports

IF IT ISN'T EVIDENCED, PRIORITIZED, AND EXPLAINABLE IN ONE PANE, IT ISN'T PRODUCTION-GRADE AI.

SINDHU@COGNICOR.COM