THE TRUST GAP

Bridging Human Inquiry and AI Workflows in Regulated Industries

White Paper
June 2026

Contents

Executive Summary

Artificial Intelligence is advancing rapidly. Models can summarize conversations, classify information, draft recommendations, automate workflows, and support decisions at a scale that would have seemed unrealistic only a few years ago. Yet adoption in insurance, financial services, and other regulated industries continues to move more slowly than the technology itself.

The reason is not hard to understand. In regulated industries, useful output is not enough. Organizations must be able to trust the workflow that produced it.

The primary obstacle is not capability. It is trust.

That distinction matters because AI is moving closer to the point where sensitive information is first discovered, interpreted, structured, and used. In insurance and financial services, that point is often not a database, form, CRM system, or model. It is a conversation. A client explains a concern. An advisor asks a better question. A planning fact becomes visible for the first time. From there, the information may later influence underwriting, suitability, recommendations, applications, service, planning, supervision, and enterprise workflows.

This is where the promise of AI meets its hardest adoption problem. If AI participates in the movement from conversation to action, regulated organizations must know more than whether the model is capable. They must know what information was captured, what was protected, what the AI system saw, what it produced, who reviewed it, what changed, what was verified, and what the organization later relied upon.

This paper refers to that challenge as the Trust Gap.

The Trust Gap is the distance between what AI may be capable of doing and what regulated organizations can responsibly trust it to do. It appears when information moves from human inquiry into AI processing and then into regulated work. It is not a rejection of AI. It is the reason AI adoption requires a stronger operating model.

The migration to cloud computing offers a useful parallel. Enterprises did not adopt cloud infrastructure simply because the technology worked. Adoption accelerated when organizations became confident that security, governance, compliance, and operational control could be maintained in a new operating model. Cloud adoption was primarily concerned with where information was stored and processed. AI adoption is increasingly concerned with how information is interpreted, transformed, and used to influence decisions.

Existing governance frameworks remain essential. Regulators and standard-setters are not creating a separate legal universe for AI. They are applying familiar obligations to AI-enabled workflows: supervision, privacy, cybersecurity, vendor oversight, recordkeeping, fairness, accountability, disclosure, suitability, best interest, and anti-discrimination. The terminology varies across insurance, securities, cybersecurity, and AI governance frameworks, but the underlying question is consistent: can the organization prove that the workflow remains trustworthy when AI participates?

Those frameworks provide important guidance. But they do not fully address the smaller operational unit that often matters most in advisory work: the individual fact discovered through human inquiry, transformed by AI, reviewed by a person, and later relied upon in regulated action.

This paper proposes Trust Architecture as the missing operational layer between trusted human inquiry and regulated action. Trust Architecture is not a replacement for NIST, NAIC, FINRA, NYDFS, cybersecurity programs, records management, CRM systems, or model governance. It is a practical framework for preserving trust as information crosses boundaries between human relationships, AI systems, and institutional workflows.

The framework has a simple hierarchy. Trust Architecture is the overall framework. Trust Boundaries show where trust must be preserved as information changes state, purpose, system, or control environment. DAPT - Discovery, Action, People, Timestamp - gives material advisory facts a durable

structure. Clarify/Verify governs whether those facts can be corrected, updated, and trusted before reliance. Discovery Agreement is the living record that emerges as verified discoveries accumulate over time.

This matters especially at the distribution edge. In insurance and financial services, the information that later powers underwriting, suitability, service, planning, and enterprise workflows is often first discovered in the advisor-client relationship. If that information is captured poorly, structured inconsistently, or stripped of trust before it reaches the enterprise, downstream AI cannot fix the problem. A distribution-first AI strategy starts where the information starts.

OMQ, the working system discussed in Section 9, illustrates this approach in practice. A fact-finding conversation becomes a source record. AI assists in identifying Points of Discovery. The advisor reviews, clarifies, and verifies those discoveries. A verified Discovery can then support a next action, such as quote readiness, while still showing what information remains missing or unresolved. As verified discoveries accumulate, they begin to form a Discovery Agreement.

The purpose of this paper is not to claim that every component of Trust Architecture is new. Consent, privacy, auditability, retention, cybersecurity, fairness, and governance are established obligations. The stronger claim is more specific: regulated AI adoption needs an operational layer that governs how individual advisory facts move from conversation to AI processing to human review to institutional action.

Trust is not the result of AI adoption.

Trust is the condition that makes AI adoption possible.

Section 1 — The Trust Gap

Bridging Human Inquiry and AI Workflows in Regulated Industries

Artificial Intelligence is advancing at a remarkable pace. Models can summarize conversations, generate recommendations, classify information, automate workflows, and produce insights that would have required teams of people only a few years ago.

Yet adoption in insurance and financial services remains uneven.

The issue is not simply whether the technology works. In many cases, it does. The harder question is whether organizations, advisors, clients, compliance officers, carriers, and regulators can trust the process by which AI is introduced into regulated advisory work.

For decades, technology initiatives in insurance and financial services have focused on moving information more efficiently through systems. Forms became digital. Portals replaced paper. CRMs, connected databases, cloud infrastructure, and back-office automation changed the way work moved. More recently, AI has entered many of those same processes.

Each innovation improved the movement of information. But none fully addressed a more basic question:

Where does the information come from in the first place?

The highest-value data in insurance and financial services does not originate in forms, databases, CRM systems, or AI models. It originates in trusted human inquiry.

That is easy to forget because the industry tends to see information only after it has already been captured. By the time a carrier receives an application, an underwriter reviews a file, a supervisor examines a recommendation, or a CRM displays a client record, the information has already passed through an earlier and more important stage. Someone had to ask the right question. Someone had to hear the answer. Someone had to recognize that the answer mattered.

A client may reveal anxiety about a child’s future, uncertainty about retirement income, concern about a spouse, an unresolved business succession issue, or a planning assumption that no longer fits. Those

statements often do not exist anywhere else before the conversation. They are not merely pieces of data waiting to be processed. They are discoveries.

This process has traditionally been called Fact Finding. Despite decades of technological change, Fact Finding remains one of the highest-value activities in an advisory relationship because it produces the information on which everything else depends.

The best advisors understand something that technology often overlooks: advisors demonstrate their expertise more by the questions they ask than the answers they give. In fact, better questions is a pretty good definition of expertise. The right question can uncover a risk, obligation, opportunity, concern, or objective that would otherwise remain invisible. The answer matters, but the answer rarely appears without the question.

This observation becomes more important as AI enters advisory workflows. AI can help organize information, summarize conversations, identify patterns, and support decision-making. But before any of that can happen, the information must first be discovered. An AI system cannot analyze what was never shared, summarize what was never revealed, or automate a workflow based on information that was never trusted to the system in the first place.

That creates what may be the central challenge of AI adoption in regulated industries: the Trust Gap.

The Trust Gap is the distance between what AI systems can do and what clients, advisors, enterprises, carriers, compliance officers, and regulators are willing to trust. It appears whenever sensitive information moves from a human conversation into a system that may summarize, transform, structure, or use it to support action.

Without trust, information is withheld. When information is withheld, workflows become incomplete. When workflows are incomplete, adoption stalls.

For years, organizations have described data silos as a technology problem. In many cases, they are really trust problems. People do not willingly share high-value information with systems, organizations, advisors, or AI workflows they do not trust.

This leads to a broader observation:

> **FIELD OBSERVATION:** Lack of adoption is the ultimate data silo.

A system may be technically capable, elegantly designed, and economically attractive. But if people do not trust it enough to use it, the system never gains access to the information necessary to create value.

This distinction matters because many current discussions about AI focus almost entirely on capability. Models continue to improve, costs continue to decline, and performance continues to increase. But adoption does not automatically follow capability.

History provides a useful parallel.

When enterprises first began migrating to cloud infrastructure, the primary challenge was not whether the technology worked. The technology already worked well enough to offer real advantages in cost, scalability, flexibility, and speed. The harder question was whether organizations could trust a new operating model in which sensitive information moved into environments they did not own and could not physically see.

Could security be maintained? Could regulators be satisfied? Could governance survive the shift? Could organizations preserve control while relying on infrastructure operated by someone else?

The cloud era required new frameworks, controls, certifications, audit practices, and governance models that helped organizations build trust in a new way of operating. Adoption accelerated when enterprises

became confident not only in the capability of cloud infrastructure, but in the control environment surrounding it.

AI presents a similar challenge, but it extends beyond infrastructure.

Cloud adoption was primarily concerned with where information was stored and processed. AI adoption is increasingly concerned with how information is interpreted, transformed, and used to influence decisions. That distinction matters in regulated advisory work because AI does not merely move information through a system. It may help determine what information matters, what action should follow, and what record the organization later relies upon.

> **MIGRATION LESSON:** AI changes the trust question from where information lives to how information is interpreted, transformed, and used.

The question is no longer whether AI can generate useful outputs. The question is whether AI can be trusted within workflows involving sensitive client information, regulated advice, supervisory obligations, recordkeeping requirements, and fiduciary responsibilities.

That challenge cannot be solved by model capability alone.

It requires an architecture of trust.

The remainder of this paper proposes that trust should not be viewed as an outcome of AI adoption. Trust is the condition that makes AI adoption possible. And if trust is the prerequisite, then it must be designed intentionally into the workflow itself.

That is the purpose of Trust Architecture.

Section 2 — The Next Migration

Cloud required trust in a new infrastructure model. AI requires trust in a new decision-support model.

Every major technology shift requires more than technical capability. It requires enough trust for organizations to change how they operate.

The history of enterprise cloud adoption provides a useful example. When cloud computing first emerged, many of the underlying technologies already worked. Organizations could reduce infrastructure costs, increase flexibility, and accelerate deployment. The technical benefits were apparent.

Yet adoption did not occur overnight. Enterprises were being asked to place sensitive information into environments they did not own and could not physically see. For many organizations, the decisive questions involved security, governance, and operational control. In hindsight, cloud adoption can appear inevitable. At the time, it felt anything but inevitable.

Cloud providers responded by building trust alongside capability. They invested in security frameworks, compliance certifications, audit controls, and documentation that helped enterprises evaluate and manage cloud risk. Adoption accelerated not because the technology alone became useful, but because organizations gained confidence in the control environment surrounding it.

The result was trust-enabled adoption.

The AI transition shares some of that history. Organizations are once again evaluating a technology capable of reducing costs, improving productivity, and transforming workflows. And once again, the technology itself is advancing faster than enterprise adoption.

Many of the concerns are familiar. Organizations want confidence that sensitive information can be protected, governance maintained, vendors trusted, and control preserved. Those concerns are legitimate and, in many cases, necessary.

But AI introduces a challenge that cloud computing largely avoided.

Cloud adoption was primarily concerned with where information was stored and processed. AI adoption is increasingly concerned with how information is interpreted, transformed, and used to influence decisions.

That difference changes the trust problem. Cloud infrastructure moved and protected data. AI systems increasingly participate in workflows that analyze information, generate recommendations, and shape outcomes. The question is no longer simply whether information is secure once it enters the system. The question is whether the processes acting upon that information can be trusted.

This distinction becomes particularly important in insurance and financial services, where the most valuable information often originates in conversations between clients and advisors. That information may pass through AI systems before it becomes part of a recommendation, application, underwriting review, or service workflow.

The challenge is therefore no longer limited to securing enterprise infrastructure. It includes preserving trust as information moves from human inquiry into regulated action.

This is where the comparison between cloud adoption and AI adoption begins to diverge. Cloud adoption required organizations to trust a new infrastructure model. AI adoption requires organizations to trust a new decision-support model.

The difference is significant.

Infrastructure decisions largely concern technology. Decision-support systems influence people, recommendations, workflows, and outcomes. They operate much closer to the point where human judgment, professional responsibility, and regulatory obligations intersect.

That may prove to be one of the defining challenges of the next decade. The question facing enterprises is not whether AI can produce useful outputs. The question is whether those outputs can be trusted inside workflows governed by privacy obligations, supervisory responsibilities, and regulatory oversight.

Answering that question requires more than model capability.

It requires an architecture of trust.

Section 3 — What Regulators Are Actually Saying

Regulators are not asking whether AI can be useful. They are asking whether existing obligations still hold when AI participates in the workflow.

One of the motivations for this paper was a simple question:

What are regulators actually saying about AI?

The volume of discussion surrounding artificial intelligence can make the answer surprisingly difficult to find. New guidance appears regularly, and the pace of commentary can create the impression that an entirely new regulatory regime is emerging.

After reviewing major guidance from insurance, financial services, cybersecurity, and AI governance sources, I came away with a different conclusion.

The most striking observation is not how much disagreement exists. It is how much agreement exists.

Across the major sources reviewed — including NIST, FINRA, the NAIC, NYDFS, and related cybersecurity and financial-services guidance — the message is remarkably consistent. Regulators are not primarily focused on whether organizations use AI. They are focused on whether organizations continue to meet their existing obligations while using it.

Privacy, supervision, recordkeeping, consumer protection, and accountability remain the responsibility of the organization deploying the technology. The introduction of AI does not eliminate those obligations. If anything, it increases the need to demonstrate them.

That distinction matters because much of the public discussion around AI assumes that regulators are creating an entirely new rulebook. That was not my conclusion. Instead, I found regulators asking a more practical question:

Can organizations continue to protect consumers, supervise activity, maintain records, secure information, and demonstrate accountability when AI is participating in business processes?

Viewed through that lens, many frameworks begin to look like different expressions of the same concern. NIST emphasizes trustworthy AI through governance and risk management. Insurance regulators emphasize governance, fairness, and consumer protection. Financial services regulators emphasize supervision, communications, and recordkeeping. Cybersecurity frameworks emphasize sensitive information, third-party risk, and operational resilience.

The terminology differs, but the underlying concern is consistent.

Trust.

> **REGULATORY SIGNAL:** Regulators are asking whether existing obligations survive when AI enters the workflow.

Not trust as marketing language or public perception. Trust in the operational sense: can the organization explain what happened, show who was responsible, demonstrate how information was handled, and prove that appropriate controls were followed?

This observation became increasingly important as the research continued.

Most regulatory frameworks focus on systems, organizations, controls, and governance programs. They describe what responsible oversight should look like and what outcomes organizations should be able to demonstrate. Those frameworks are necessary and valuable.

Yet they largely begin after information already exists.

Those frameworks generally assume that relevant information has already been collected, recorded, processed, and incorporated into organizational workflows. They say less about the earlier path: how a fact was discovered, how conversation became record, how meaning changed along the way, and how the information became part of a recommendation, application, workflow, or institutional decision.

This is not a criticism of the frameworks. It is an observation about their scope.

Regulators are focused on governing organizations and systems. The challenge emerging in AI-enabled advisory workflows is that information itself is becoming more dynamic. Conversations are transcribed, discoveries are extracted, summaries are generated, and workflows may be triggered from information that only moments earlier existed as human conversation.

The point where human conversation becomes operational action is therefore becoming more important, not less.

That observation revealed what appeared to be a missing layer.

Not a replacement for existing governance frameworks.

A layer that connects them.

A layer concerned with how information moves from trusted human inquiry into AI processing and ultimately into regulated action.

That is the challenge explored in the next section.

Section 4 — The Missing Layer

Existing frameworks govern systems. The missing layer governs information as it crosses from conversation into action.

The research reviewed in this paper points to an important conclusion: existing frameworks are not wrong. In fact, they are largely addressing the right concerns.

NIST focuses on governance, accountability, transparency, and risk management. FINRA emphasizes supervision, communications, and recordkeeping. NAIC guidance focuses on insurer governance, consumer protection, and accountability. Cybersecurity frameworks such as NYDFS Part 500 focus on protecting sensitive information and managing third-party risk.

Each framework addresses a critical aspect of AI adoption. Yet something important remains largely unaddressed.

Most existing frameworks govern systems, models, infrastructure, controls, records, vendors, policies, or organizations. They remain essential, but they do not fully address the smaller unit that often matters most in advisory work: the individual fact discovered through human inquiry, transformed by AI, reviewed by a person, and later relied upon in regulated action.

That distinction matters because the highest-value information in insurance and financial services rarely originates inside enterprise systems. It originates in conversations.

A client may tell an advisor something that has never been entered into a form. A business owner may disclose a concern that exists nowhere in a CRM. A family may reveal an obligation, fear, objective, or circumstance unknown to the carrier, the enterprise, and the technology supporting them.

Only after that information is shared does it begin its journey through systems, workflows, and controls.

The industry often talks about enterprise trust boundaries as if they are the primary challenge. Enterprises invest enormous resources protecting information once it enters their systems. They build secure environments, access controls, vendor oversight programs, and increasingly, private AI environments intended to keep sensitive information within a trusted perimeter.

These efforts are rational and necessary. But they focus on only one side of the problem.

There is another trust boundary that receives far less attention: the boundary where client conversation becomes advisor knowledge, where advisor knowledge becomes enterprise information, and where human inquiry becomes institutional action.

That is where the highest-value information originates.

And that is where trust is first established or lost.

The enterprise often asks, “How do we protect information once we have it?” Clients, advisors, compliance officers, and regulators are asking a related but earlier question: “How do we trust the process that turns discovered information into action in the first place?”

The distinction is subtle but important. One question begins with data already inside the system. The other begins with the human relationship that created the data.

AI-enabled advisory workflows operate precisely at the intersection of those two trust problems. Information originates in a human conversation, crosses into AI processing, is transformed into structured information, and may later influence a recommendation, application, underwriting review, or service workflow.

At each stage, trust must be preserved.

This is where many discussions about AI become incomplete. The conversation often begins with the model: a prompt enters, the model processes it, and an output emerges. But regulated advisory workflows contain something more important than prompts and outputs. They contain trust boundaries.

Information does not simply move from input to output. It moves across a series of boundaries where identity, context, purpose, control, and responsibility may change.

A practical example illustrates the pattern:

Source State, where real identities and full human context exist.

De-identification Boundary, where sensitive information is protected before AI processing.

AI Processing State, where substituted or limited context is used by the model.

Restoration Boundary, where AI output returns to the human workflow.

Human Working State, where real identities are restored and a responsible person can review, clarify, verify, and act.

The existence of a boundary alone is not enough. The boundary must be observable, auditable, and capable of producing evidence that the control was actually applied.

Trust cannot depend solely on policy statements.

Trust requires proof.

This observation reveals the missing layer between traditional governance frameworks and practical AI adoption. Organizations have developed sophisticated methods for governing systems. What remains less developed is the governance of information as it crosses trust boundaries between human inquiry, AI processing, and regulated action.

This is not a replacement for existing frameworks. It is the layer that connects them.

Without it, organizations may successfully govern their models while struggling to govern the information those models consume and produce.

The challenge is not merely protecting enterprise data. The challenge is preserving trust as information moves from human inquiry to regulated action.

That challenge is the foundation for Trust Architecture.

Section 5 — Trust Architecture Framework

Trust Architecture begins one layer closer to the source: the material fact discovered through human inquiry.

The missing layer described in the prior section is not theoretical. It appears whenever human conversation becomes operational data. A client says something important, an advisor hears it, an AI system may summarize it, a workflow may be triggered, and a record may eventually be relied upon by the organization.

At that point, the question is no longer simply whether the AI model performed well. The question is whether the process can be trusted.

Trust Architecture begins with that question.

It is not a replacement for existing governance frameworks. It does not replace NIST, FINRA, NAIC, NYDFS, cybersecurity rules, CRM systems, data lineage, or enterprise governance programs. Those frameworks and systems remain necessary.

Trust Architecture addresses a different problem.

It asks how trust is preserved as information moves from human inquiry into AI processing and then into regulated action.

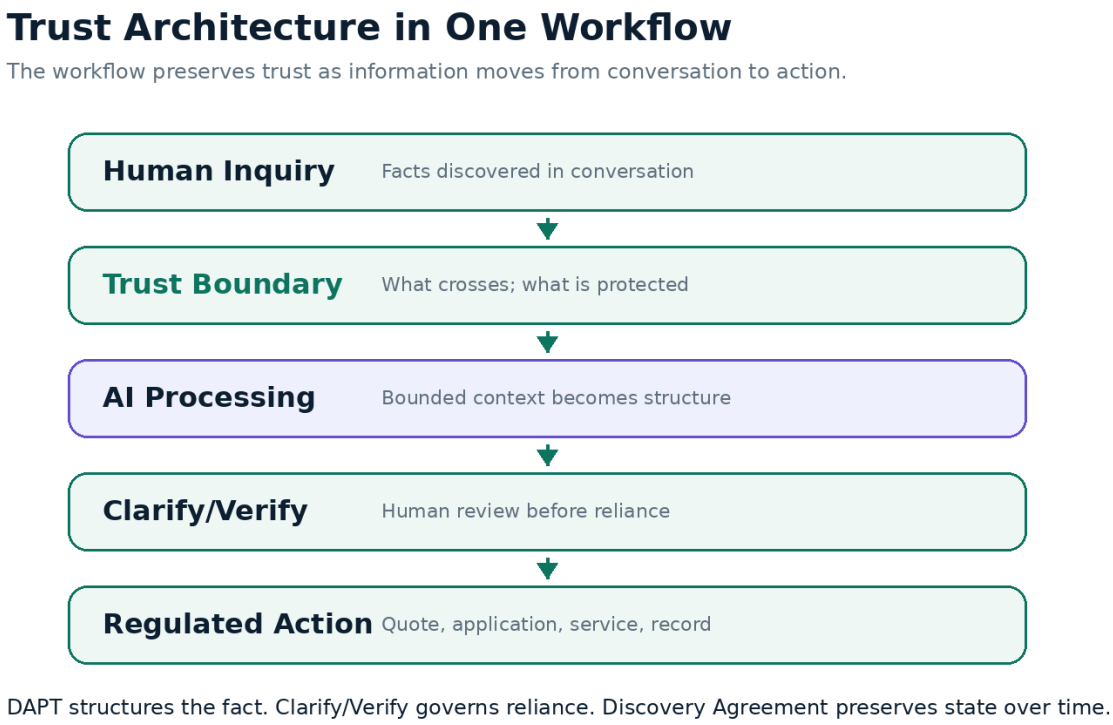

Figure 1. Trust Architecture in one workflow. Human inquiry produces the discovery; Trust Architecture governs how that discovery moves through AI processing, human review, and regulated action.

Existing governance frameworks typically organize around systems and controls. Trust Architecture begins one layer closer to the source. It begins with the discovery itself.

The premise is simple: before information becomes part of a recommendation, application, or institutional record, it first exists as a fact discovered through human inquiry.

If trust cannot be preserved at that level, it becomes increasingly difficult to preserve later.

This is the key distinction. In many enterprise settings, governance begins after information has already entered the system. The record exists, the model has been selected, the vendor has been approved, and the control environment has been established.

All of that matters.

But advisory workflows begin earlier. They begin when a client says something that was not previously known: a family concern, a business succession issue, or an assumption that no longer fits.

These are not merely entries in a database. They are material advisory facts.

That is why Trust Architecture is built around a simple observation: trust must survive each transformation of information.

As information is captured, transformed, verified, shared, retained, and eventually relied upon, each change in state creates both value and risk. The risk is not only that data may be exposed or misused. The deeper risk is that meaning may change, the source may become unclear, consent may be incomplete, human review may be skipped, or the organization may later be unable to prove what happened.

A trusted AI workflow must therefore answer practical questions inside the workflow itself. Was the information captured with appropriate consent? Who was involved? What was discovered? Was the information client-stated, advisor-stated, or AI-inferred? What did the AI system actually see? Was sensitive information protected before AI processing? Was the output reviewed by a human? Was the discovery clarified, corrected, or verified? What action did it support? Can the organization prove the answer?

These questions are not separate from compliance. They are how compliance becomes operational.

Consent matters because trusted inquiry cannot depend on vague permission. A client may agree to a conversation, but that is not the same as agreeing to recording, transcription, AI summarization, storage, sharing, or later use in a recommendation or application.

Identity matters because advisory facts belong to real people and are acted upon by real professionals. The system must know who participated, who captured the information, who verified it, who updated it, and who relied on it.

Provenance matters because the most important unit is often not the document or the record. It is the fact.

A transcript may contain hundreds of statements. A CRM record may contain years of accumulated information. A planning file may contain many documents. But the operational question is often narrower: what specific fact was discovered, when was it discovered, who was involved, and has it been verified?

Without fact-level provenance, AI can produce summaries, but it cannot produce reliable advisory records.

Human oversight matters because AI output should not become institutional truth merely because a model generated it. The advisor must be able to review the discovery, compare it to the source, correct what is wrong, verify what is accurate, and accept responsibility for the next step.

Data protection matters because not every artifact belongs in every part of the workflow. Raw audio, transcripts, AI prompts, structured discoveries, client profile fields, and advisor notes serve different purposes and carry different risks. A trusted workflow should expose only what is necessary, protect sensitive information at boundary crossings, and preserve evidence of what was actually sent to the AI system.

De-identification should not be treated as a complete privacy solution. In regulated advisory workflows, it is one control among several. Technical substitution of PII should be paired with appropriate consent, vendor and model-use restrictions, no-training commitments where applicable, retention rules, and audit evidence showing what the model actually received.

Retention, governance, fairness, and audit evidence complete the structure. Regulated industries cannot keep everything forever or delete everything in the name of privacy. Organizations cannot claim AI is governed unless they can show what model was used, what data was processed, what output was created, who reviewed it, what changed, and what action followed. Nor can they treat all AI use cases as equivalent. Summarizing a client-stated fact is different from recommending a product. Identifying missing information is different from classifying underwriting risk. Drafting a next question is different from generating a client-facing recommendation.

This is the role of the Trust Architecture Framework.

TAF organizes the controls required to preserve trust as advisory information moves from human inquiry to regulated action. Its purpose is not to invent a new compliance regime. Its purpose is to make existing obligations operational at the point where the most valuable information is created.

That point is the trusted conversation.

The next sections examine the parts of Trust Architecture that may be most distinctive. DAPT provides a compact schema for governing the lifecycle of material advisory facts. Clarify/Verify provides a process for challenging, correcting, updating, and verifying those facts before reliance. Trust Boundaries identify the transitions where information moves between human inquiry, AI processing, and operational use.

Together, these concepts help address the Trust Gap by making trust observable at the level where it is most often created, transformed, and lost: the individual fact.

Section 6 — DAPT: Governing the Lifecycle of Facts

The important unit in advisory workflows is often not the document, the record, or the model. It is the fact.

The prior section introduced the central operating premise of Trust Architecture: the important unit in advisory workflows is often not the document, the record, the model, or the system.

It is the fact.

That distinction matters because regulated workflows usually do not fail in the abstract. They fail around specific facts: something was assumed rather than verified, discussed but never recorded, or changed without the record changing with it.

These are not merely data quality problems. In advisory work, they are trust problems.

DAPT is a compact schema for governing those facts.

DAPT stands for Discovery, Action, People, and Timestamp. Each Point of Discovery is structured around those four elements, not to create a more elegant note, but to convert a material advisory fact into a structured, traceable, and actionable unit of trust.

A Discovery is what was learned, stated, observed, or inferred. It is the substance of the fact. In an advisory workflow, the Discovery might be a client's stated concern, a planning need, or a change in circumstances.

Action is what should happen next. It connects the Discovery to responsibility. A fact that does not lead to any possible action may still be interesting, but it may not be material. In advisory work, the value of a fact often depends on whether it changes what someone should do.

People identifies the participants in the Discovery event: the people who were present for, contributed to, or can verify the fact as it was established. This matters because facts in advisory relationships are rarely floating abstractions. They are discovered by people and become trustworthy only when the workflow can show who participated in the Discovery.

Timestamp identifies when the information was discovered, updated, verified, or relied upon. Time matters because advisory facts age. A fact that was true during one conversation may become stale after a life event, health change, or planning update.

Together, these four elements answer a simple set of questions: what was discovered, what should happen next, who participated in the Discovery, and when did this enter the record?

The strength of DAPT is that it operates at the fact level.

Traditional systems usually govern larger containers: client records, files, tasks, or supervisory review. Those tools are useful, but they do not necessarily govern the lifecycle of the individual advisory fact that emerged from a conversation and later influenced action.

That is the gap DAPT is designed to address.

A transcript may preserve what was said, and a summary may identify themes, but neither necessarily governs the fact that should carry forward. CRM notes and workflow tasks can preserve context or assign action, but they often fail to connect the action back to the Discovery that made it necessary.

DAPT sits between those layers.

It does not replace the transcript, CRM, workflow, or compliance record. It creates a structured bridge between human inquiry and institutional action.

This is why DAPT should be understood as an operational schema rather than a governance theory. At the level of theory, DAPT does not replace metadata, records management, or data governance. Those disciplines already exist and remain necessary.

DAPT's contribution is more practical.

It takes the core questions those disciplines care about and applies them to the smallest unit that matters in advisory work: the material fact discovered through human inquiry.

That makes the schema especially useful in AI-enabled workflows because AI changes the volume, speed, and form of discovery. A human advisor may conduct a conversation. An AI system may transcribe it, summarize it, and extract possible discoveries. Without structure, those outputs can become a new kind of clutter: fluent, plausible, useful-looking, and difficult to trust.

DAPT forces structure at the moment of extraction.

The AI system should not merely say, "The client needs more insurance." That is a summary. A governed advisory fact requires more. It should preserve what was discovered, what action follows, who is involved, and when the discovery entered the system.

For example, a DAPT-structured discovery might establish that James stated his current coverage is inadequate and confirmed a need for $1.5 million of 20-year term coverage to protect Linda and Tyler if he dies before the mortgage is paid and Tyler is grown.

That fact can now be clarified, verified, acted upon, and audited.

This is the movement from conversation to record.

It is also the movement from AI output to trust-controlled workflow.

DAPT also helps separate types of information that are often blended together. A client-stated fact is different from an advisor observation. An AI inference is different from either. A recommendation is different from a Discovery. A carrier underwriting decision is different from the information submitted to support it.

Those distinctions matter because each type of information carries a different trust burden. Client-stated facts may need confirmation, AI inferences may need explanation or limitation, and recommendations may need appropriate support.

DAPT does not solve all of those obligations by itself. But it gives the workflow a clearer starting point by preventing material facts from dissolving into unstructured notes or opaque AI summaries.

The People field is especially important in this respect.

In advisory work, facts often depend on who said them and who can verify them. A spouse may reveal a concern the other spouse had not raised. An advisor may observe a planning gap the client did not understand. Another professional may later need to confirm or act on the information.

By identifying the people who participated in the Discovery, DAPT attaches accountability to the fact. The fact is no longer an isolated statement. It becomes part of a human context.

The Timestamp field serves a similar purpose.

Many advisory failures are not caused by information that was always wrong. They are caused by information that became stale. A client moves, a business is sold, or a health condition changes.

DAPT makes time part of the fact's structure, not an afterthought. A Discovery can therefore be understood not only by what it says, but by when it was discovered and whether it has been updated or verified since.

That temporal structure becomes more important as AI systems support ongoing advisory relationships. Static records decay. Trustworthy records must remain alive enough to be challenged, corrected, updated, and relied upon only when appropriate.

This leads directly to the next control.

DAPT can structure a fact, but structure alone does not make the fact true. A Discovery may be incomplete, an Action may be wrong, a Person may be missing, or the fact may become outdated after it enters the system.

For that reason, DAPT should not be treated as the end of the trust process. It is the beginning of a governed lifecycle.

The fact must still be clarified. It must be verified. It may need to be corrected. It may need to be updated as circumstances change.

This is where Clarify/Verify becomes essential.

DAPT gives the advisory fact a durable form. Clarify/Verify governs whether that fact can be trusted over time.

Together, they create the foundation for fact-level governance in AI-enabled advisory workflows. DAPT answers what the fact is. Clarify/Verify determines whether the fact is ready to be relied upon.

Section 7 — Clarify/Verify: Governing the Lifecycle of Trust

Structure alone does not create trust. A fact must remain open to challenge before reliance.

DAPT gives the advisory fact a durable form. Clarify/Verify governs whether that fact can be trusted over time.

That distinction matters because structure alone does not create trust. A fact can be organized and still be wrong, traceable and still incomplete, accurate when captured and stale when later relied upon.

This is why Trust Architecture cannot stop at extraction.

An AI system may identify a material discovery from a conversation, place it into a structured format, connect it to an action, identify the people involved, and preserve the time it entered the system. That is useful, but the output should not become institutional truth simply because the system produced it.

In advisory work, the movement from output to trust requires human judgment.

Clarify/Verify is the process by which that judgment enters the workflow.

A Discovery should remain open to challenge before it is relied upon. If a Discovery is wrong or incomplete, it should be clarified. If it changes over time, it should be updated. If it is accurate enough to support action, it should be verified.

This may sound like ordinary review, but in an AI-enabled advisory workflow it becomes something more important. Clarify/Verify is not merely a user experience feature. It is a trust control.

The reason is that AI-generated records create a new kind of risk. A traditional advisor note is visibly human. It may be incomplete or imprecise, but everyone understands that it reflects one person's record of a conversation. An AI-generated summary carries a different kind of authority. It can appear polished,

complete, and objective even when it has misunderstood the client, overstated a conclusion, missed context, or blended fact with inference.

That fluency can create false confidence.

Clarify/Verify exists to interrupt that false confidence.

> **TRUST CONTROL:** Clarify/Verify turns human oversight into a fact-level checkpoint before reliance.

It creates a required moment of review between AI output and durable reliance. The advisor is not merely receiving a summary. The advisor is deciding whether the Discovery is accurate, whether the Action follows from it, whether the People field is complete, and whether the record is ready to become part of the working file.

That decision matters because a verified fact has a different status than an unreviewed AI output. It has moved through a human checkpoint. Someone has reviewed it, compared it against the conversation or known context, and accepted responsibility for using it.

This is where Clarify/Verify connects directly to the regulatory themes discussed earlier in the paper.

Regulators are not asking whether AI can be useful. They are asking whether organizations can supervise it, explain it, preserve records, protect consumers, and demonstrate accountability. Clarify/Verify helps make those obligations operational at the fact level.

The control depends on separating three states that are often blurred together: AI-generated output, clarified information, and verified information. AI-generated output may be useful, but it remains provisional. Clarified information may be improved, but it may still require review. Verified information is the point at which a responsible human has accepted that the fact is accurate enough to support action.

Without that separation, a workflow can move too quickly from conversation to conclusion. A model extracts and displays a fact, the system treats it as true, and a recommendation, application, or service action follows.

The problem is not that any single step is unreasonable. The problem is that the trust transition was never governed.

Clarify/Verify governs that transition.

Clarify allows the record to improve without pretending the first output was final. It gives the advisor a way to correct a misunderstanding, add missing context, refine vague language, or distinguish what the client actually said from what the system inferred.

Verify marks the point of reliance.

It is the advisor’s attestation that the Discovery is no longer merely a generated artifact. It is a working fact.

This structure becomes more important because advisory facts do not remain static. A client’s health, business, family obligations, and planning assumptions can all change. A fact that was accurate when discovered may become misleading if it is relied upon later without review.

Clarify/Verify supports trust over time by making change part of the lifecycle. The question is not only, “Was this fact accurate when it was first discovered?” The better question is, “Is this fact still accurate enough to support the action being taken now?”

That is a different standard.

It recognizes that advisory records are living records. They do not become trustworthy once and remain trustworthy forever. They require a process by which meaning can be revisited as circumstances change.

This is also why Clarify/Verify should not be reduced to a simple edit function. In a CRM, a user may change a field because new information came in or because the old information was wrong. That may be useful, but the update often says little about the underlying trust process. The system may not know what prompted the change, whether the client confirmed it, whether AI was involved, or whether the change affected a recommendation.

Clarify/Verify should carry more weight than that.

A clarification should preserve the idea that a fact is being refined because accuracy matters. A verification should preserve the idea that a human has reviewed the fact and is willing to rely on it. Over time, the record should show not merely the current answer, but the lifecycle of trust around that answer.

That lifecycle has practical consequences. If a client later challenges a recommendation, the organization should be able to show what fact was discovered, how it was structured, whether it was clarified, who verified it, and what action it supported. If a compliance officer reviews a file, the officer should not have to guess whether a summary was raw AI output or advisor-reviewed information. If an advisor returns to the client months later, the advisor should know which facts are verified, which remain uncertain, and which may need to be updated.

This is the operational value of Clarify/Verify.

It does not make AI perfect. It makes AI output accountable.

It also preserves the human element in a way that is practical rather than symbolic. The advisor is not pushed aside by automation or reduced to a rubber stamp. The advisor remains responsible for interpreting the client's situation, challenging the system when necessary, and deciding when information is ready to support action.

That matters because advisory trust is not created by the system alone. It is created through the relationship between the client, the advisor, and the process they use to turn information into action.

Clarify/Verify strengthens that relationship by giving the advisor, and eventually the client, a way to say whether a Discovery is accurate, incomplete, outdated, or ready to rely upon.

DAPT governs the structure of the fact.

Clarify/Verify governs the trust status of the fact.

Together, they move AI output closer to something regulated industries can actually use: information that is generated, reviewed, accountable, and ready for responsible action.

Section 8 — Trust Boundaries

Trust is not a static property of a system. It must be preserved as information changes state.

Trust Architecture depends on a simple idea: trust must be preserved as information changes state.

That idea becomes especially important when AI enters the advisory workflow because AI does not merely store information. It receives information, transforms it, and returns something new. A conversation may become a transcript; a transcript may become a summary; a summary may become a structured discovery; and that discovery may later support an action or institutional record.

Each transformation can create value, but each also crosses a boundary.

A trust boundary is the point where information moves from one state, system, purpose, or control environment into another. In traditional enterprise discussions, the trust boundary is often understood as the perimeter around the organization or its technology environment. Information is inside the trusted environment or outside it. A vendor is approved or not approved. A system is authorized or not authorized.

Those boundaries still matter.

But AI-enabled advisory workflows introduce a different kind of boundary, one closer to the source of the information itself.

The relevant boundary is not only between the enterprise and the outside world. It is also between the client conversation and the AI system, between AI output and the advisor's working record, and between that working record and the action that follows.

That is where trust is most vulnerable.

In a trusted conversation, information has human context. The advisor knows who is speaking, what prompted the answer, what was uncertain, and what still needs to be confirmed. Once that information enters an AI system, some of that context can be lost or transformed. The model may summarize accurately, but it may also compress nuance, blend fact with inference, or present uncertain information with more confidence than it deserves.

The boundary matters because AI is useful enough to be relied upon.

If the system turns a conversation into structured information, the organization must be able to explain how that transformation occurred. It must know what crossed the boundary, what was protected before crossing, what the AI system actually saw, what came back, and what human process governed the result.

This is where Trust Boundaries become operational.

In practical terms, the workflow moves through three states. Source State contains real identities and full human context. AI Processing State receives protected or substituted context. Human Working State restores the information to a form the advisor can review, clarify, verify, and rely upon when appropriate.

The De-identification Boundary is the point where information moves from the source conversation toward AI processing. In this context, de-identification primarily means protecting PII and other regulated client information before it is sent into AI processing. Identity-linked fields, such as names, dates of birth, account identifiers, or policy numbers, should not cross the boundary unless they are necessary for the task and permitted by the workflow.

The purpose is not to pretend the remaining information is anonymous or harmless. In advisory work, even substituted or de-identified context may remain sensitive if it can be linked back to a client profile or combined with other facts. De-identification reduces unnecessary exposure, but it does not eliminate the need for consent, vendor controls, no-training commitments where applicable, retention rules, and audit evidence. Its value is that it limits what crosses the boundary and helps preserve proof of what the model actually saw.

> **TRUST BOUNDARY NOTE:** De-identification is a boundary control, not an anonymity guarantee. It reduces exposure while consent, vendor controls, retention rules, and audit evidence carry the remaining burden.

The Restoration Boundary is the point where AI output returns to the human advisory environment. The advisor cannot work forever with substituted identities or stripped context. The information must be restored into a form that is usable, reviewable, and connected to the real client relationship.

Restoration is not merely a technical step. It is a trust step. The system is saying, in effect, that the AI processed a protected version of the information and that the result has now been returned to the human workflow, where review and responsibility belong.

The value of the boundary pattern is the evidence it produces.

A boundary that cannot be observed is only a policy statement. A boundary that leaves evidence becomes a control. In an AI-enabled advisory workflow, the organization should be able to show what was protected before AI processing, what the model actually saw, what output was returned, and how the restored information moved into human review.

That evidence becomes especially important when sensitive client information is involved. A carrier, compliance officer, regulator, or enterprise buyer may not be satisfied with a general assurance that “data is protected” or “AI is governed.” They will want to know how the boundary works in practice.

The organization should be able to show what crossed the boundary, what was withheld, whether the model could train on the data, whether client identity was exposed, whether retention followed policy, whether human verification occurred, and whether those answers can be proven.

Trust Boundaries help convert those concerns into workflow controls.

They also help distinguish between different types of records. A raw transcript, a de-identified AI prompt, a model output, a restored advisory discovery, and a verified Point of Discovery are not the same artifact. They should not be governed as if they carry the same meaning, risk, or evidentiary value.

This distinction matters because one of the easiest mistakes in AI governance is to treat the workflow as a single event.

The client conversation happened, the AI summarized it, and the advisor acted.

That description is too simple for regulated environments. The more accurate view is that information changed state several times, and each state carried different obligations.

Trust Architecture does not eliminate that complexity.

It makes the complexity visible.

This is one reason the preserved AI-processing record is so important. If an organization can show what the model actually saw, it can separate the source conversation from the AI input, the AI input from the AI output, and the AI output from the human-verified record. That separation supports accountability because it prevents later confusion about whether a statement came from the client, the advisor, the model, or the verified record.

In OMQ, this idea appears in the distinction between source information, de-identified processing context, restored working information, and verified Points of Discovery. The narrative context preserved for AI processing is intentionally de-identified, which means it can serve as evidence of what the model actually saw without exposing the full source identity. That is not a complete governance program by itself, but it is a meaningful Trust Boundary control.

The larger point is that de-identification alone is not the whole answer.

A workflow also needs consent before information is captured or processed. It needs DAPT to structure the fact that emerges. It needs Clarify/Verify to govern whether that fact can be trusted. It needs retention rules to determine what is preserved, and it needs audit evidence to show what happened.

Trust Boundaries connect those controls.

They show where trust must be protected as information moves between human inquiry, AI processing, and regulated action.

That is why Trust Boundaries are central to the Trust Architecture Framework. They make clear that trust is not a static property of a system. It is something that must be preserved through movement.

The client trusts the advisor with sensitive information. The advisor trusts the workflow to handle that information properly. The enterprise trusts the controls, the carrier or institution trusts the record, and the regulator trusts the evidence.

If any boundary fails, the chain weakens.

The goal of Trust Architecture is not to remove every boundary. That would be impossible. The goal is to make the boundaries visible, controlled, and auditable so that information can move safely from the conversation where it was discovered to the action it is meant to support.

Section 9 — Trust in Practice: From Fact-Finding to Action

The case study shows how a conversation becomes a verified, actionable record without treating AI output as final truth.

The prior sections describe Trust Architecture as a framework. This section shows how the framework appears inside an actual advisory workflow.

The example is a life insurance fact-finding conversation. An advisor meets with James and Linda Kowalski to discuss family protection, existing coverage, mortgage exposure, and the need for additional insurance. During the conversation, James confirms that his current coverage is inadequate and states a need for $1.5 million of 20-year term coverage to protect Linda and Tyler if he dies before the mortgage is paid and Tyler is grown.

The purpose of this case study is not to present OMQ as a complete implementation of Trust Architecture. It is not. A mature enterprise deployment would still require stronger client-facing consent, formal retention rules, audit export, version history, enterprise identity controls, vendor governance, and use-case boundaries.

The purpose is narrower and more practical: to show how trusted human inquiry can become structured, reviewable, and actionable information without treating AI output as final truth.

Exhibit 1 — Source Conversation

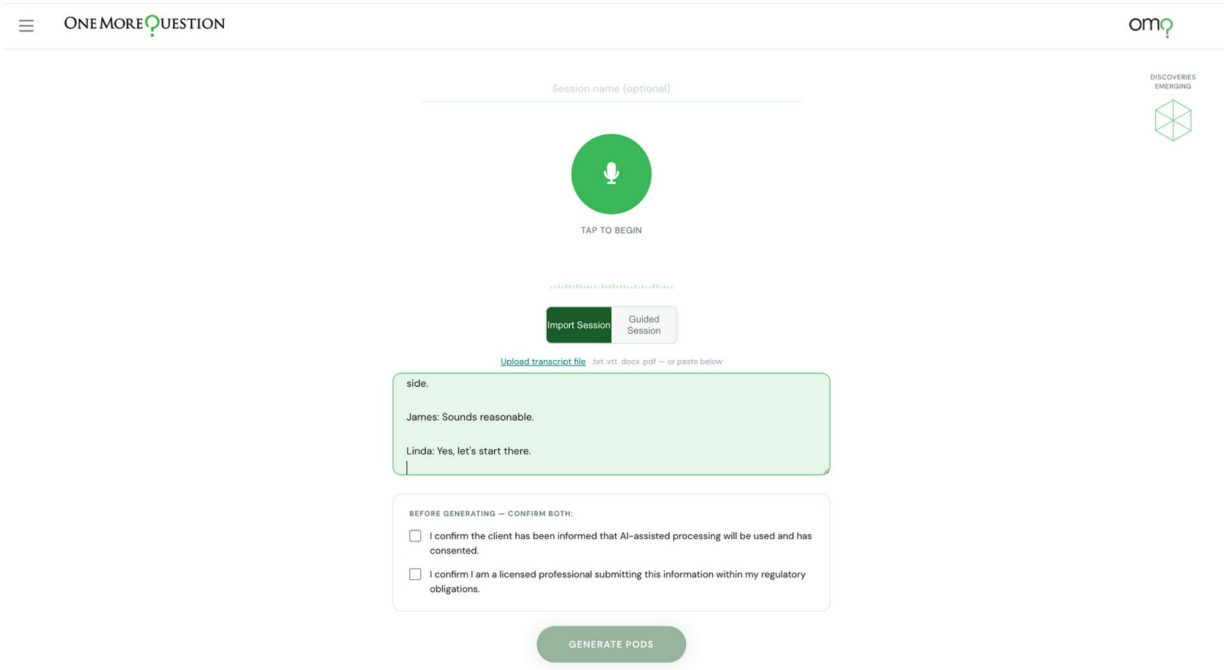

Exhibit 1. Source Conversation. The workflow begins with a consented fact-finding session before AI-generated discoveries are created.

The workflow begins with a consented advisory conversation. The meeting may occur over video, phone, or in person, but the principle is the same: the conversation is source material, not the final record.

A transcript preserves what was said, but it does not determine which facts matter, what actions may follow, or whether any statement is ready to support a regulated workflow. In a traditional advisory process, those judgments may remain inside the advisor's notes, memory, CRM entries, or follow-up tasks. Some information is captured, some is summarized, and some is lost.

OMQ begins from a different premise.

The conversation is the source of value, but the transcript is not enough. The system must identify the material advisory facts inside the conversation and preserve enough context to explain where those facts came from.

In this example, the important discovery is not simply that James and Linda discussed life insurance. The important discovery is that James stated a specific coverage need, for a specific purpose, tied to specific family obligations.

That is the difference between a meeting summary and an advisory fact.

Exhibit 2 — Trust Boundary and De-identification

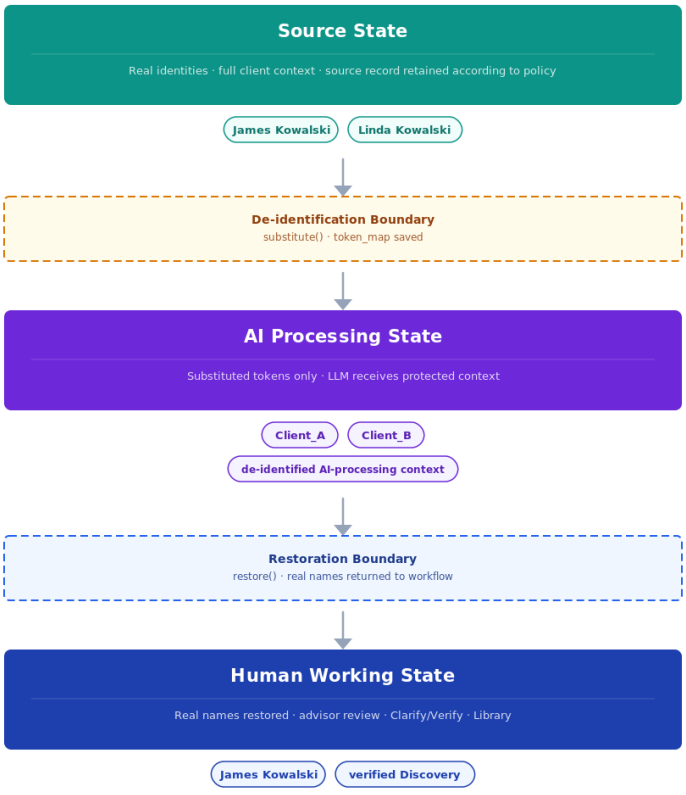

Exhibit 2. Trust Boundary Model. Source information is protected before AI processing, transformed into substituted-token context, and restored when it returns to the human working state for advisor review and Clarify/Verify.

Before information is sent into AI processing, OMQ applies a Trust Boundary.

Personally identifiable information is protected before the model receives the planning context. Identity-linked fields, such as names, dates of birth, account identifiers, or policy numbers, should not cross into AI processing unless they are necessary for the task and permitted by the workflow.

The client remains identifiable to the advisor, but not to the AI processing context.

This distinction matters. The purpose of de-identification is not to pretend the remaining information is anonymous or harmless. Advisory context may remain sensitive even when identity-linked fields have been substituted or removed. The purpose is to reduce unnecessary exposure, preserve evidence of what the model actually saw, and make the boundary auditable.

In practical terms, the workflow moves from Source State, where real identities and full human context exist, through a De-identification Boundary into AI Processing State. After the model produces output, the workflow passes back through a Restoration Boundary into Human Working State, where the advisor can review, clarify, verify, and decide what can be relied upon.

The boundary is valuable because it leaves evidence. The organization should be able to show what crossed into AI processing, what was withheld, whether client identity was exposed, whether the model could train on the data, and whether human verification occurred before reliance. Retention should then follow the applicable policy for each artifact.

Trust Architecture should not mandate one retention answer for every enterprise. It should enforce the organization's regulatory interpretation by artifact type, so one firm may retain source materials while another may minimize or delete them when permitted.

That evidence is what turns a privacy promise into a workflow control.

Exhibit 3 — Point of Discovery

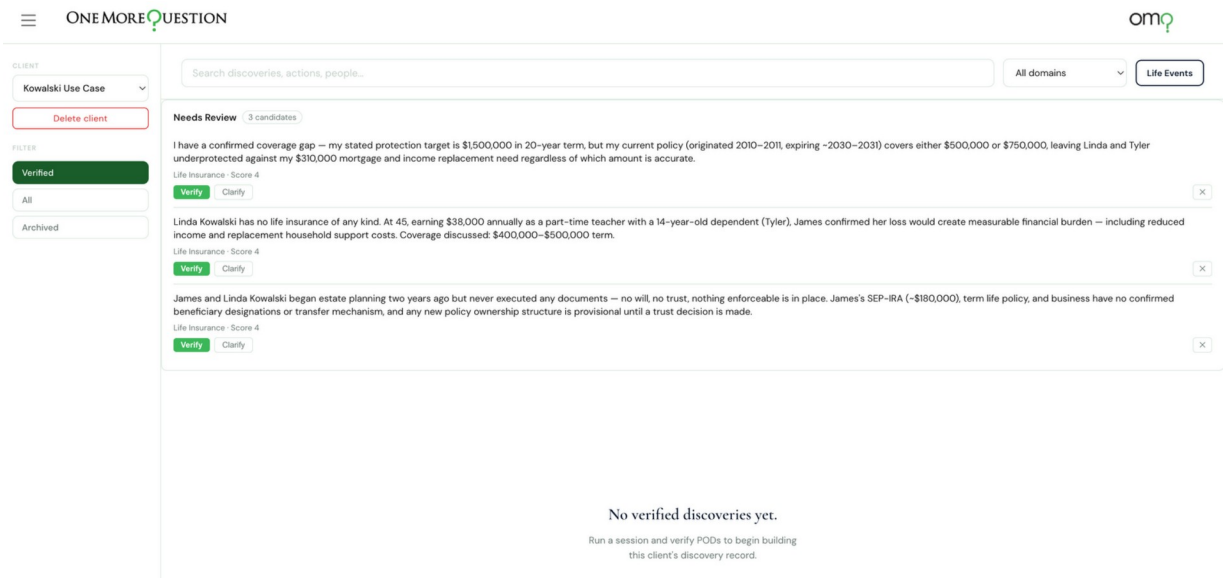

Exhibit 3A. Needs Review State. AI-generated candidate discoveries remain provisional before advisor verification.

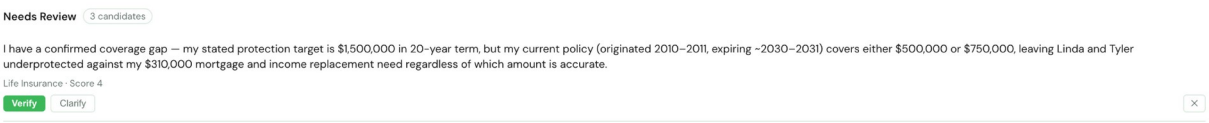

Exhibit 3B. Point of Discovery Detail. The $1.5 million, 20-year term coverage need appears as an AI-generated candidate before it becomes a verified advisory fact.

After AI processing, OMQ identifies Points of Discovery, or PODs.

A POD is not a generic summary. It is a structured advisory fact that may matter to the relationship, support a next step, or become part of the working client record.

In the Kowalski example, the system may identify a discovery like this:

James confirmed that his current life insurance coverage is inadequate and stated a need for $1.5 million of 20-year term coverage to protect Linda and Tyler if he dies before the mortgage is paid and Tyler is grown.

That sentence carries more advisory value than a general summary that says the clients “discussed life insurance needs.” It identifies the person, the concern, the amount, the duration, and the purpose of the coverage.

This is where DAPT appears in practice.

DAPT Element	Case Study Application
Discovery	James confirmed that his current coverage is inadequate and stated a need for $1.5 million of 20-year term coverage to protect Linda and Tyler.
Action	Prepare for a life insurance quote request after missing quote information is confirmed.
People	James, Linda, and the advisor, as participants in the fact-finding conversation.
Timestamp	Discovered during the fact-finding session.

The structure is simple, but the shift is significant. The system is not merely preserving content. It is creating a fact-level record that can be reviewed, clarified, verified, acted upon, and audited.

As verified PODs accumulate, they can form a Discovery Agreement: a living record of what has been learned, who participated in the Discovery, what action may follow, and when the facts were established or updated.

Exhibit 4 — Clarify/Verify Review

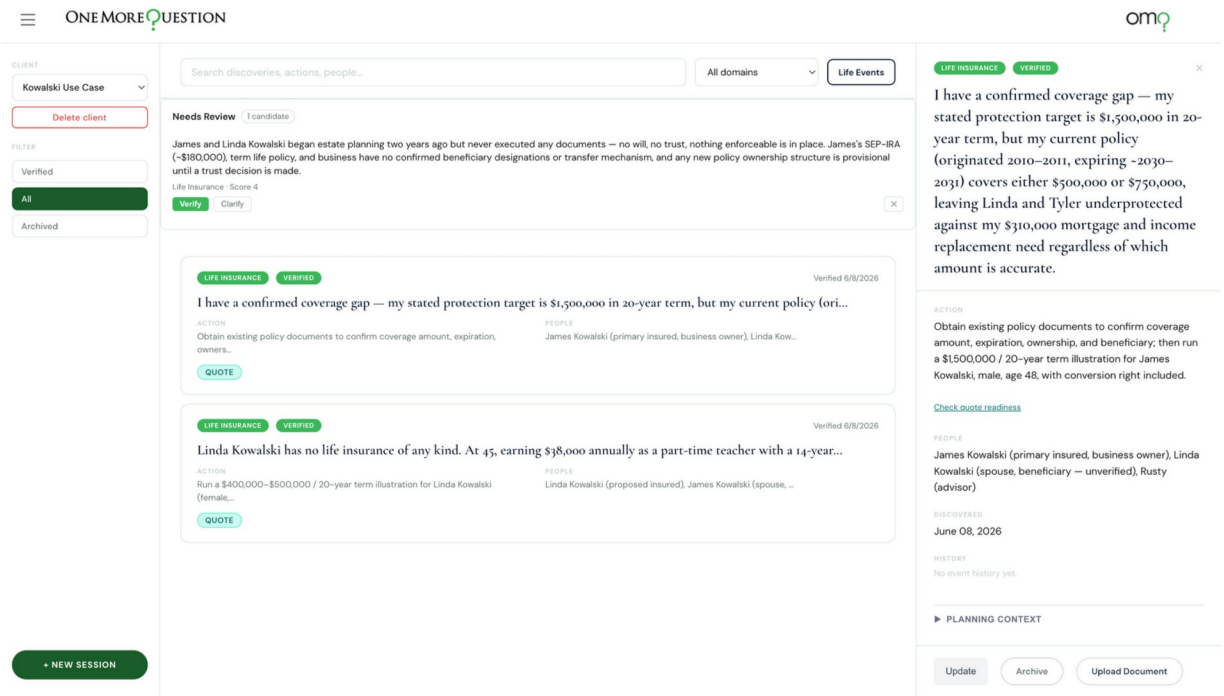

Exhibit 4. Human Review and Verification. The advisor remains responsible for reviewing discoveries, verifying what is ready, and leaving unresolved items in review.

The Point of Discovery does not automatically become institutional truth.

It enters review.

This is where Clarify/Verify becomes a trust control rather than a user interface feature. The advisor can compare the Discovery against the conversation, correct what is incomplete, refine what is ambiguous, and verify only what is accurate enough to support action.

In this example, the advisor may review the stated amount, duration, purpose, insured person, beneficiary concern, and any changed circumstances before treating the Discovery as part of the working record. If the amount should be $1.25 million rather than $1.5 million, the Discovery can be clarified. If the duration should be 30 years rather than 20, the record can be corrected. If James confirms the amount and duration, the Discovery can be verified.

At that point, something important changes.

The information is no longer merely AI-generated output. It has passed through human review and become a verified planning fact.

That distinction matters for compliance and supervision. A reviewer should not have to guess whether a statement came directly from the client, was inferred by the model, or was reviewed by the advisor. Clarify/Verify creates a visible transition between generated output and information that is ready to support action.

Exhibit 5 — From Discovery to Action

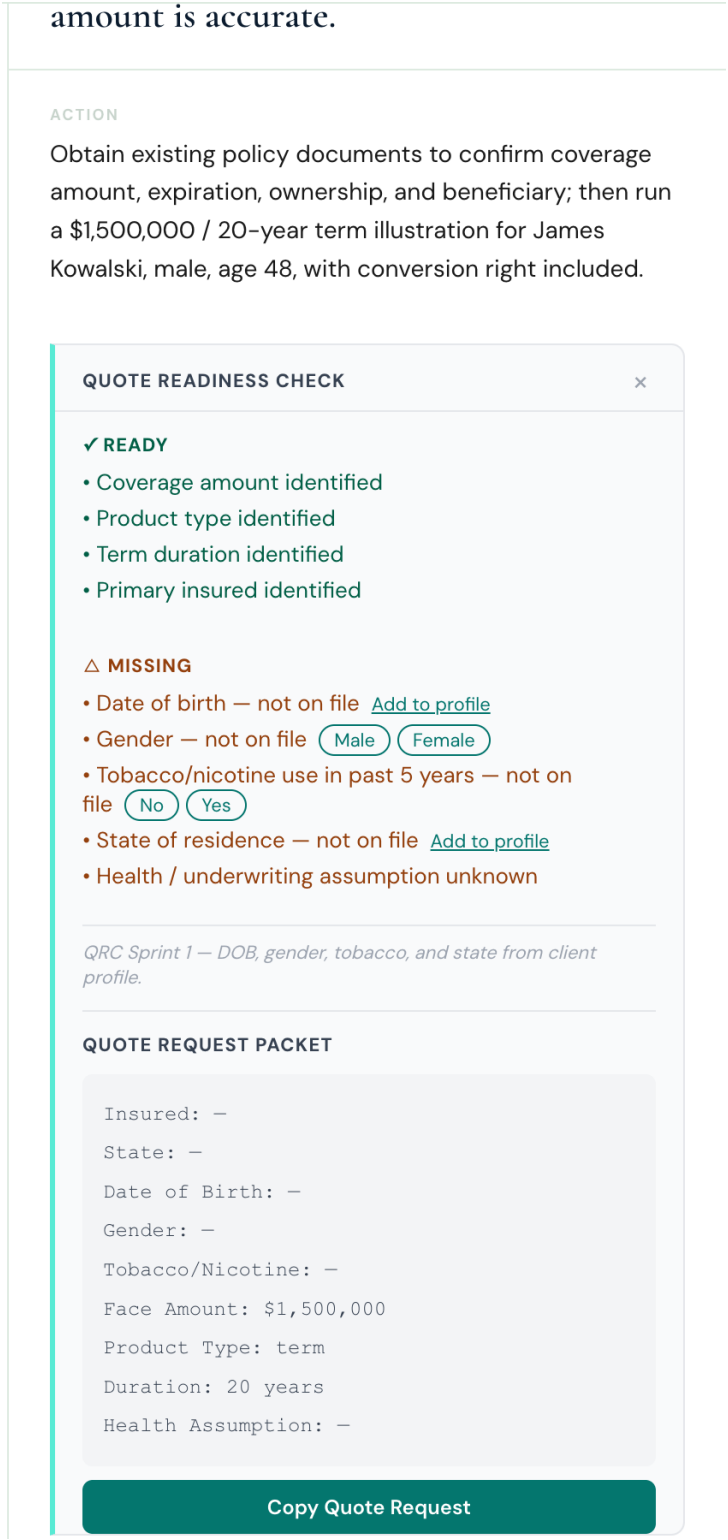

Exhibit 5. Quote Readiness Check. The verified Discovery populates core quote fields while showing missing information that still requires human confirmation.

Once a Discovery is verified, it can support action.

In the life insurance workflow, the next step may be a quote request. The verified Discovery establishes a need for coverage, but a quote request still requires operational details. Some information may be ready, some may exist in the client profile, and some may still need confirmation.

The quote-readiness view shows the practical value of DAPT's Action component. The system can connect the verified Discovery to the next workflow while showing what remains unresolved before the advisor proceeds.

For example, the face amount, product type, and duration may be available from the conversation. Other fields, such as state of residence, date of birth, tobacco status, or underwriting assumption, may need to be confirmed before the advisor sends a request to a brokerage general agency or carrier.

The value is not that AI completes the regulated workflow on its own. The value is that AI helps identify the material fact, connect it to the next responsible action, and reveal what still needs human confirmation.

That is a more realistic model for AI adoption in regulated industries. The system accelerates discovery and workflow preparation, but the advisor remains responsible for judgment, verification, and action.

Exhibit 6 — Verified Record / Discovery Agreement

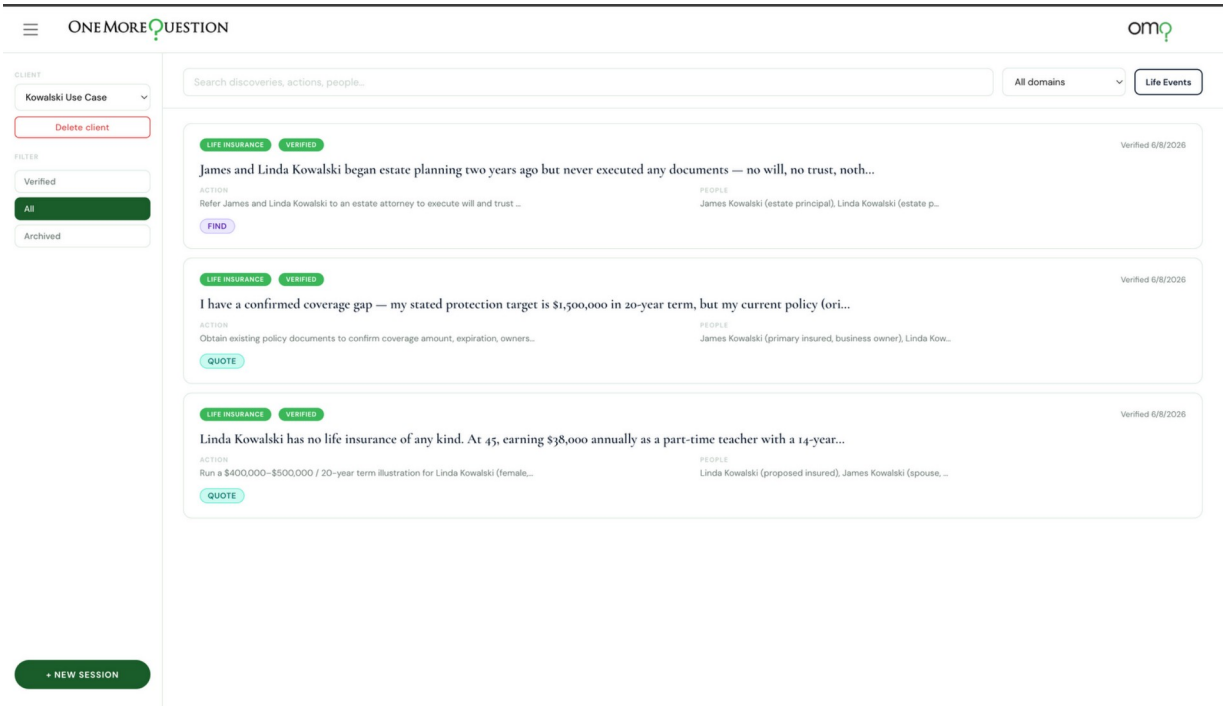

Exhibit 6. Discovery Agreement View. Verified Points of Discovery accumulate into a living advisory record that can support future action.

As verified Points of Discovery accumulate, they begin to form a Discovery Agreement.

A Discovery Agreement is a structured record of what has been learned, what actions may be required, who participated in the Discovery, and when those facts were established or updated. It is not a one-time meeting summary. It is a living record that can evolve as the client's circumstances change.

That distinction matters because advisory relationships unfold over time. Clients have children, sell businesses, update estate plans, increase coverage needs, approach retirement, and change priorities. A static note may preserve what was known at one moment, but a Discovery Agreement can preserve how the record changes while maintaining the provenance of the underlying facts.

This is where trust and state converge.

Trust determines whether information can be responsibly captured, processed, verified, and governed. State determines whether that information can remain useful as the client relationship continues.

Together, they allow AI to participate in advisory workflows without requiring the organization to surrender accountability. The advisor remains responsible, the client remains involved, the organization retains governance, and AI accelerates discovery without replacing judgment.

What the Case Demonstrates

This use case demonstrates the practical purpose of Trust Architecture.

Human inquiry remains the source of the highest-value information. Trust Boundaries protect the movement into and out of AI processing. DAPT gives each material fact structure. Clarify/Verify governs whether that fact can be relied upon. The Action component connects verified discoveries to responsible workflows without allowing the system to outrun human judgment.

OMQ is not valuable because it summarizes a meeting.

Many systems can summarize meetings.

The more important contribution is that OMQ shows how a discovered fact can move from conversation to action while preserving provenance, accountability, and trust.

That is the bridge this paper has been describing.

Trusted human inquiry produces the discovery. AI helps transform that discovery into structure. Trust Architecture determines whether the structured information is reliable enough to use.

In regulated industries, usefulness is not enough.

The information must be trustworthy enough to rely upon.

Section 10 — Recommendations

Capability creates the possibility of adoption. Trust determines whether adoption happens.

The central lesson of the Trust Gap is that capability alone will not drive adoption.

Trust must be designed into the workflow.

That requires a shift in how organizations evaluate AI-enabled advisory tools. The question should not be limited to whether a system produces a good summary or a faster workflow. The better question is whether the organization can trust the process by which information was discovered, protected, reviewed, verified, retained, and acted upon.

For Advisors

Advisors should treat AI as a tool for improving discovery, not replacing judgment.

The highest-value work in an advisory relationship still begins with the human conversation. AI can help preserve that conversation, identify material facts, and organize follow-up. But the advisor remains responsible for understanding the client, asking better questions, challenging weak outputs, and deciding what information is ready to support action.

That means advisors should avoid treating AI summaries as final records. A generated summary may be useful, but it should remain provisional until reviewed. If a discovery is wrong, incomplete, ambiguous, or stale, the advisor should clarify it. If the discovery is accurate enough to support action, the advisor should verify it.

This is not merely a compliance habit.

It is good advisory practice.

Advisors earn trust by showing that they understand the client's actual circumstances and concerns. AI can help capture those elements, but it cannot replace the human responsibility to confirm meaning.

The practical recommendation is simple: use AI to strengthen discovery, but do not let AI become the source of truth without human review.

For Firms and Enterprises

Firms should evaluate AI advisory workflows through the lens of Trust Architecture, not productivity alone.

A tool that summarizes meetings or drafts recommendations may be useful, but usefulness is not enough in regulated environments. Firms need to know what enters the system, what is sent to AI processing, what the model sees, how outputs are reviewed, which records are retained, and what evidence exists if the workflow is challenged later.

This requires more than a vendor questionnaire.

It requires a workflow-level review.

Before deploying AI into advisory processes, firms should map how information moves from conversation to record to action. That map should identify where consent is required, where information crosses Trust Boundaries, which outputs require human review, and which actions should never occur without approval.

A firm does not need to solve every enterprise governance problem before experimenting with AI. But it should avoid pilots that create hidden records, unclear data flows, unreviewed recommendations, or ambiguous responsibility.

The practical recommendation is to start with narrow use cases where the role of AI is clear. Summarizing a client-stated fact is different from recommending a product. Identifying missing information is different from determining suitability. Preparing a quote request is different from submitting an application. Those distinctions should be explicit before deployment.

For Carriers and Product Manufacturers

Carriers should evaluate AI adoption not only as an internal enterprise initiative, but as a distribution-facing strategy.

Many carrier AI initiatives naturally begin inside the enterprise. Underwriting, service, claims, and internal productivity are important areas for improvement, but they often begin after the most valuable information has already been captured poorly, inconsistently, or not at all.

In insurance distribution, the highest-value data is usually created before it reaches the carrier.

> DISTRIBUTION-FIRST IMPLICATION: AI strategy should start where trust and data begin: distribution.

It emerges in the advisor-client conversation.

That is where clients explain their needs, concerns, obligations, and planning constraints. If that information reaches the carrier only through incomplete forms, fragmented notes, or loosely structured submissions, the carrier's AI strategy begins downstream from the real source of value.

A distribution-first strategy starts earlier.

It uses AI to help advisors capture, structure, verify, and govern the information that originates in the field. The benefit to the carrier is not merely a better advisor experience, although that matters. The benefit is better upstream data entering carrier workflows.

Cleaner intake data can reduce rework, clarification loops, and underwriting delays. Verified discoveries can give carriers greater confidence in the source and status of submitted information. Structured fact-finding can help distinguish client-stated facts from advisor observations and AI inferences.

This matters because carriers do not merely receive data from distribution.

They rely on it.

A carrier should therefore distinguish between AI that produces unverified output and AI that produces structured, reviewable, evidence-backed information. Accuracy and efficiency are not enough. The carrier should ask whether a fact can be traced to the conversation where it originated. It should also ask whether protected information was controlled before AI processing, whether a licensed human reviewed the output, and whether the discovery was verified before use.

A distribution-first strategy also reduces adoption risk. Advisors are more likely to use AI that helps them do their work, preserve their relationships, and reduce administrative burden. If the system creates value for the advisor first, the carrier receives better data as a result. If the workflow is built only for the carrier's internal needs, adoption may stall before the enterprise ever receives the information it hoped to analyze.

This is why distribution should be treated as a strategic entry point for carrier AI.

It is not merely the channel through which products are sold. It is the point where the information powering underwriting, suitability, service, planning, and future enterprise workflows is first discovered.

The practical recommendation is that carriers should encourage AI innovation at the distribution edge while requiring evidence of consent, provenance, human review, and Trust Boundary controls before accepting AI-assisted outputs at scale.

For Compliance and Governance Leaders

Compliance and governance leaders should focus less on whether AI is present and more on what role AI plays in the workflow.

The compliance question is not simply, "Do we use AI?" It is, "What did AI touch, what did it change, who reviewed it, what was retained, and what action followed?"

That question moves governance closer to the actual risk.

An AI system that drafts an internal summary creates one kind of obligation. A system that supports a recommendation, application, or underwriting submission creates another. The governance model should reflect those differences.

Compliance leaders should therefore require workflows that distinguish source material from AI input, AI output from human-reviewed records, and provisional discoveries from verified facts. Those distinctions make supervision more practical because they show where human judgment entered the process.

The practical recommendation is to govern AI at the workflow level, not only at the policy level. Policies matter, but regulated organizations also need evidence that the policy was followed in the specific case.

For Regulators

Regulators should continue applying existing obligations to AI-enabled workflows while recognizing that advisory AI creates a more granular operational challenge.

Much of the current regulatory conversation properly focuses on governance, supervision, privacy, recordkeeping, fairness, and accountability. Those obligations remain essential.

But AI-enabled advisory workflows introduce the lifecycle of individual facts as a practical governance issue.

A material advisory fact may begin as a client statement. It may pass through AI processing, become a structured discovery, receive human clarification, become verified, and later support a recommendation, application, underwriting review, or service workflow. If regulatory expectations focus only on the model or the final record, they may miss the transformations that occurred along the way.

Regulators do not need to prescribe a single architecture for every firm. But they can help the market by encouraging evidence-based controls around consent, provenance, human review, retention, and data minimization.

The practical recommendation is to ask firms to demonstrate how trust is preserved across the workflow, not merely whether an AI policy exists. The most useful regulatory question may be: can the firm explain how a material fact moved from human inquiry to regulated action?

For Technology Providers

Technology providers building AI tools for regulated industries should resist the temptation to sell automation without accountability.

The market does not need another system that merely produces polished summaries. It needs systems that help organizations distinguish source conversations from AI inputs, AI outputs from human-reviewed records, and verified facts from provisional suggestions.

That requires design choices that may not look glamorous in a product demo but matter deeply in deployment. The workflow should build in consent, PII protection, de-identified processing records, human review, verification controls, audit logs, retention rules, and use-case boundaries from the beginning.

These controls should not be treated as compliance afterthoughts.

They are product features in regulated markets.

The practical recommendation is to build for evidence from the beginning. If a system cannot show what the model saw, what it produced, what a human reviewed, and what action followed, it will be difficult for regulated organizations to trust at scale.

Closing Recommendation

The common recommendation across all participants is this: move the center of AI governance closer to the point where information is discovered.

The industry already has frameworks for governing systems, organizations, vendors, models, records, and controls. Those frameworks are necessary, but AI-enabled advisory workflows require something more operational.

They require governance at the level of the material fact.

That is where trusted human inquiry becomes structured information. It is where AI can add value. It is also where trust can be preserved or lost.

A practical Trust Architecture should make five things visible: how information was discovered, how it crossed into AI processing, how the resulting fact was structured, how it was reviewed or verified, and what action it supported.

If those questions can be answered, AI adoption becomes easier to trust.

If they cannot, even powerful AI systems may remain trapped outside the workflows where they could create the most value.

Appendix A — Sources Reviewed

This appendix identifies the regulatory, governance, cybersecurity, cloud-adoption, and internal assessment materials reviewed for the working draft. The body of the paper remains practitioner-led; this appendix carries the source burden and shows the research base behind the framework discussion.

Source	Relevance to the Paper	Official Reference / Note
NIST AI Risk Management Framework (AI RMF 1.0)	Provides cross-sector AI governance language: govern, map, measure, manage; validity, reliability, accountability, transparency, privacy, security, resilience, and fairness.	National Institute of Standards and Technology, AI RMF 1.0, January 2023. https://www.nist.gov/itl/ai-risk-management-framework
NIST Generative AI Profile (NIST AI 600-1)	Extends the AI RMF vocabulary to generative AI risks and controls, including risks unique to foundation-model and GenAI use.	NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, July 2024.
NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers	Establishes insurance-specific expectations around governance, accountability, fair and ethical use, consumer protection, and documentation for insurer AI systems.	National Association of Insurance Commissioners, adopted December 4, 2023; state adoption map reviewed.
New York DFS Insurance Circular Letter No. 7 (2024)	Provides an insurance-specific signal on AI systems and external consumer data in underwriting and pricing, including scrutiny of unfair or unlawful discrimination.	New York Department of Financial Services, Circular Letter No. 7, July 11, 2024.
Colorado Regulation 10-1-1	Reviewed as an operational state-level model for governance and risk management of external consumer data, algorithms, and predictive models in insurance.	Colorado Division of Insurance, Regulation 10-1-1; reviewed as cited in the research report.
NAIC Insurance Data Security Model Law #668	Supports the cybersecurity and breach-response discussion for insurers, insurance producers, and other licensed entities.	NAIC Model #668, Insurance Data Security Law.
NAIC Suitability in Annuity Transactions Model Regulation #275	Provides a best-interest and supervision anchor for AI-supported annuity and retirement-income workflows.	NAIC Model Regulation #275; reviewed through NAIC materials.
FINRA Regulatory Notice 24-09	Confirms that existing member-firm obligations apply when firms use generative AI and large language models.	FINRA Regulatory Notice 24-09, June 27, 2024.
FINRA Annual Regulatory Oversight materials on GenAI and third-party risk	Supports the treatment of GenAI governance, supervisory controls, recordkeeping, model/vendor risk, and AI-generated communications as regulated workflow issues.	FINRA 2025 and 2026 Annual Regulatory Oversight Report materials reviewed.
SEC Regulation S-P amendments	Supports the privacy, customer-information safeguards, breach-notification, written-records, and incident-response themes relevant to advisor workflows.	Securities and Exchange Commission, Regulation S-P amendments, 2024.
SEC AI-washing enforcement materials	Supports the caution against overstating AI claims or presenting unsupported compliance and automation assertions.	SEC settled charges against Delphia (USA) Inc. and Global Predictions Inc., March 18, 2024.

FTC Safeguards Rule and AI enforcement materials	Supports customer-information safeguards, service-provider oversight, and the warning that AI claims must not be deceptive.	Federal Trade Commission, Safeguards Rule guidance; Operation AI Comply materials.
CFPB adverse-action and chatbot materials	Reviewed as cross-sector signals that AI opacity is not a defense and that consumer-facing automation requires accountability, accuracy, and escalation.	CFPB Circulars 2022-03 and 2023-03; CFPB chatbot report materials.
NYDFS Cybersecurity Regulation / 23 NYCRR Part 500	Supports cybersecurity, access control, audit trail, third-party service provider, and incident-notification themes for regulated financial institutions.	New York Department of Financial Services, 23 NYCRR Part 500 and amendments.
GLBA / FTC Safeguards and SEC Reg S-P privacy themes	Supports the treatment of nonpublic personal information, information-security programs, service-provider oversight, and incident response.	Federal Trade Commission and SEC privacy and safeguard materials reviewed.
EU AI Act	Reviewed as a future-facing signal that AI used in credit, life insurance, health insurance, and other high-impact settings will increasingly be treated as high-risk.	Regulation (EU) 2024/1689, Artificial Intelligence Act.
AWS Cloud Adoption Framework and cloud-security materials	Supports the analogy between cloud adoption and AI adoption: technology adoption accelerates when enterprises trust the surrounding control environment.	AWS Cloud Adoption Framework and related AWS security/adoption materials reviewed.
Trust Architecture Research Report	Internal research synthesis used to map regulatory themes to AI-enabled advisor workflows.	Research Report: AI-Enabled Advisor Workflows in Insurance and Financial Services, June 3, 2026.
OMQ Trust Architecture Assessment	Internal technical assessment used as a checklist for strengths, gaps, and the distinction between ordinary compliance components and the more differentiated fact-level governance layer.	OMQ Trust Architecture Assessment Against TAF v1, June 4, 2026. Used as an internal assessment, not as an external authority.

Appendix B — Regulatory Themes and Trust Architecture Relevance

The sources reviewed do not point to a single new AI rulebook. They point to a consistent regulatory pattern: existing duties continue to apply when AI participates in the workflow. The practical question is whether the organization can prove that those duties were honored as information moved from human inquiry into AI processing and regulated action.

Theme	Research Signal	Trust Architecture Relevance
Existing duties still apply	Regulators generally extend existing obligations into AI-enabled workflows rather than treating AI as a separate legal universe.	TAF is positioned as an operational layer that helps existing duties become visible inside the workflow.
Consent must be specific enough to matter	Recording, transcription, AI summarization, storage, sharing, and later use are different consent events.	Trust Architecture should distinguish consent to the conversation from consent to AI processing and downstream use.
AI opacity is not a defense	Guidance around complex algorithms,	Trust Boundaries and preserved AI-

	adverse action, and consumer protection shows that unexplained automation does not excuse accountability.	processing context help show what the model actually saw and produced.
Human oversight must be real	Human-in-the-loop language is weak if the human merely rubber-stamps AI output.	Clarify/Verify turns human oversight into a fact-level checkpoint before reliance.
AI outputs may become regulated artifacts	Transcripts, summaries, meeting notes, client communications, recommendations, and supervisory materials may require retention or review depending on use.	DAPT separates provisional AI output from verified advisory facts and connects each verified fact to action and provenance.
Vendor and model risk must be governed	LLM providers, transcription services, cloud hosts, and other subprocessors become part of the regulated data-flow chain.	TAF requires the workflow to show what data crossed which boundary, under what controls, and whether vendor/model-use restrictions applied.
Retention cannot be one-size-fits-all	Privacy expectations and recordkeeping obligations can point in different directions.	Trust Architecture should enforce the organization's regulatory interpretation by artifact type: source audio, transcript, AI input, AI output, POD, verified record, audit log, and client profile field.
Fairness and discrimination risks rise with use case	Insurance regulators focus heavily on unfair discrimination where AI affects underwriting, pricing, eligibility, and insurance practices.	TAF should distinguish client-stated facts, advisor observations, AI inferences, recommendation support, and underwriting/pricing decisions.
Audit evidence matters	Policies and vendor questionnaires are not enough if the workflow cannot reconstruct what happened.	DAPT, Clarify/Verify, Trust Boundaries, and audit trails create evidence at the level where information was discovered and acted upon.

Appendix C — Trust Architecture Assessment Notes

This appendix summarizes the internal OMQ Trust Architecture assessment. It is included as an implementation checklist and critical self-assessment, not as an external authority.

- TAF should not claim novelty merely because it includes consent, identity, data protection, retention, governance, fairness, or audit. Those areas map directly to existing compliance, cybersecurity, privacy, and enterprise-security frameworks.
- The stronger claim is narrower and more defensible: Trust Architecture adds value by applying provenance, human review, and lifecycle governance to individual advisory facts rather than only to systems, records, documents, or models.
- DAPT is best understood as an operational schema for the lifecycle of a material advisory fact. Its value is not that the four words are new; its value is that each Point of Discovery carries Discovery, Action, People, and Timestamp as a structured unit of trust.
- Clarify/Verify should be treated as a trust control, not a cosmetic edit function. The meaningful difference is the structured checkpoint between AI-generated output and information that may be relied upon.
- OMQ's strongest current architectural properties include consent-gated LLM processing, de-identified AI-processing context, a mandatory human Verify gate, immutable Discovered timestamp, fact-level POD structure, and the separation between provisional candidates and verified PODs.

- The most important open gaps include client-facing consent, formal retention rules, structured audit export, clarification history/versioning, vendor/model governance, and use-case boundaries. Those gaps do not defeat the Trust Architecture argument, but they should limit claims about production readiness.

Appendix D — Glossary of Key Terms

Term	Definition
Trust Gap	The distance between what AI systems are capable of doing and what regulated organizations can responsibly trust them to do.
Trust Architecture	The design of workflow controls that preserve accountability, provenance, privacy, oversight, and auditability as information moves from human inquiry into AI processing and regulated action.
Trust Architecture Framework / TAF	The framework proposed in this paper for organizing the controls needed to preserve trust inside AI-enabled advisory workflows.
DAPT	Discovery, Action, People, and Timestamp; a compact operational schema for governing the lifecycle of material advisory facts.
Discovery	What was learned, stated, observed, or inferred during the advisory process.
Action	What should happen next or what workflow may be triggered by the Discovery.
People	The participants in the Discovery event. People means those who were present for, contributed to, or can verify the Discovery as it was established; it does not mean every person affected by the Discovery.
Timestamp	When the information was discovered, updated, verified, or relied upon.
Point of Discovery / POD	A structured advisory fact that may matter to the relationship, support a next step, or become part of the working client record.
Clarify/Verify	The human review process by which AI-generated discoveries can be challenged, corrected, updated, and verified before reliance.
Trust Boundary	A point where information moves from one state, system, purpose, or control environment into another.
Source State	The state where real identities and full human context exist, usually in the original advisory conversation or source material.
AI Processing State	The state where limited or de-identified context is processed by the AI system.
Human Working State	The state where information is restored to the advisor workflow for review, clarification, verification, and action.
De-identification Boundary	The boundary where identity-linked fields are substituted, removed, or limited before AI processing.
Restoration Boundary	The boundary where AI output returns to the human workflow and is reconnected to the real client relationship.
Discovery Agreement	A living record of what has been learned, who participated in the Discovery, what action may follow, and when the facts were established or updated.
Identity-linked fields	Fields such as names, dates of birth, account identifiers,

	policy numbers, or similar data that link information to a specific person or account.
Verified Discovery	A Discovery that has passed through human review and is accurate enough to support responsible action.
Material advisory fact	A fact discovered through advisory inquiry that may affect planning, recommendations, service, underwriting, suitability, or another regulated workflow.