Sharpen

Whitepaper

Securing AI in the Contact Center

A Practical Guide to Trust, Compliance, and Control

How to evaluate AI governance, protect regulated data, and maintain compliance when your contact center platform adopts AI.

Written for security, compliance, and operations leaders in healthcare, financial services, and any organization that entrusts sensitive customer interactions to a third-party platform.

April 2026
Matt Grofsky, D.Eng, CISSP
Chief Information Security Officer, Sharpen

Executive Summary

AI adoption in the contact center has moved faster than most governance frameworks can keep up with. Regular generative AI use across organizations has more than doubled in two years, with customer service among the highest-adoption areas. [1]

The operational case is clear: lower cost per interaction, higher first-contact resolution, improved agent performance, and the ability to analyze 100% of interactions rather than a small sample. For regulated industries, there's a specific opportunity on top of that: monitoring compliance at a scale and consistency that manual review processes can't match.

But AI adoption also introduces new categories of risk. Customer data flows through AI systems in ways that differ from traditional software. AI models can be manipulated through adversarial techniques that most security programs were never designed to detect. The governance question is no longer theoretical — research from Vanta finds that 62% of business leaders report significant concern about AI compliance. [2]

This paper is for security, compliance, and operations leaders who need to evaluate whether their contact center platform's AI capabilities meet the governance standards their industries require. It addresses the specific concerns of healthcare and financial services organizations while remaining relevant to any organization that handles sensitive customer interactions. It is not a compliance attestation or product specification. It is a guide to what responsible AI governance looks like in the contact center, and what questions you should be asking your platform provider.

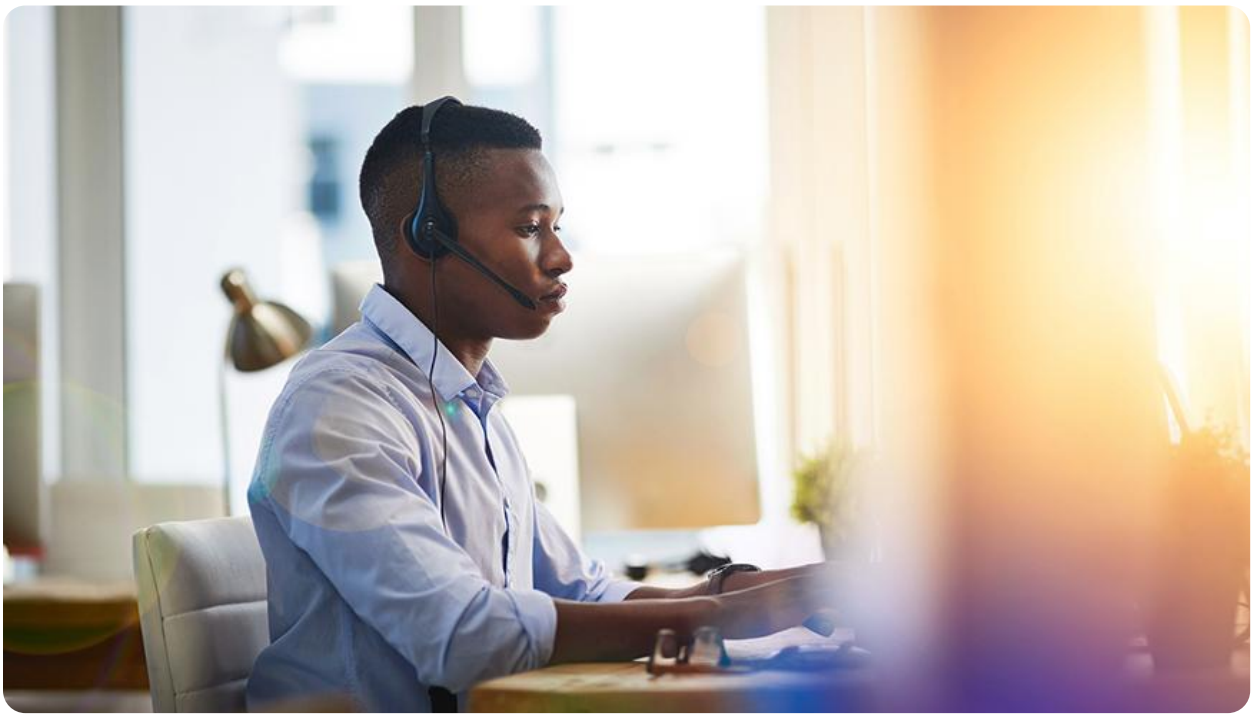

AI in the Contact Center: Why the Stakes Are Higher for Regulated Industries

Not all contact centers carry the same risk profile. A retail customer support operation and a healthcare payer's member services line operate in fundamentally different regulatory environments. When AI is introduced, those differences are amplified.

Healthcare: Warranted Caution

Healthcare organizations have reason to approach large language models with caution. ECRI, the nonprofit patient safety organization, identified risks from AI-enabled health technologies as the number one health technology hazard for 2025.[3] A 2025 IBM report found that the average security breach in the healthcare industry totaled over $7 million.[4] A Wolters Kluwer Health survey of over 500 healthcare administrators found that nearly 30% ranked data breaches as their top AI-related concern. Among large health systems with over 25,000 employees, that figure rises to 57%.[5]

The regulatory environment is also shifting rapidly. HHS OCR published a proposed HIPAA Security Rule update on January 6, 2025, the first update since 2013 and the most substantial revision since the rule's original 2003 publication, citing the rise in ransomware and the need for stronger cybersecurity.[6] At the state level, multiple states have enacted laws requiring AI disclosure in healthcare settings, mandating human oversight of clinical AI, and imposing transparency requirements beyond what HIPAA was designed to address.[7] For example, California AB 3030 requires disclosure when generative AI is used to communicate patient clinical information, and Texas TRAIGA, effective January 2026, requires written disclosure to patients when AI supports healthcare services.

The concern is not whether AI has value in healthcare contact centers. It is whether AI can be deployed in a way that maintains the integrity of the Business Associate Agreement chain, keeps PHI within controlled boundaries, and satisfies both federal and emerging state requirements. The answer depends on architecture and governance.

Financial Services: Compliance at Conversational Speed

Financial services contact centers face a different but equally demanding landscape. GLBA requires protection of nonpublic personal information. PCI DSS governs payment data. Financial services contact centers may also face communications recording, retention, supervision, and disclosure requirements under SEC, CFTC, and FINRA rules, including certain obligations implemented under Dodd-Frank. The TCPA and Telemarketing Sales Rule regulate outbound communications. The FTC Safeguards Rule was substantially revised in 2021 to expand the definition of covered financial institutions and was amended again in 2023 to add breach-notification requirements.[8]

The regulatory complexity continues to intensify. In February 2024, the FCC issued a declaratory ruling confirming that AI-generated voices are classified as artificial voices under the TCPA, requiring the same consent protections as traditional robocalls.[9] TCPA violations carry $500 per call in statutory damages, increasing to $1,500 per willful violation. State-level privacy laws modeled on CCPA/CPRA continue to roll out, and the FCC is actively considering additional rulemaking around AI-generated communications and consent revocation.[10]

For all regulated contact centers, the volume of interactions that must be monitored for compliance far exceeds what manual QA processes can cover. That reality shapes both the risk and the opportunity that AI represents.

Addressing the HIPAA Question Directly

For healthcare organizations evaluating AI in their contact center, the Business Associate Agreement chain is the first and most important question.

How the BAA Chain Should Work with AI

When a contact center platform uses a cloud-hosted AI service, the BAA chain must extend through the AI processing layer. This is a technical and contractual requirement. In practice:

- The customer (covered entity or business associate) has a BAA with the CCaaS provider.
- The CCaaS provider has a subcontractor BAA or equivalent agreement with the cloud infrastructure vendor.
- The AI service used for inference must be on the cloud vendor's official HIPAA Eligible Services list.
- The BAA must cover the specific services used for AI processing, not just general compute or storage.

If any link in this chain is missing, PHI that flows through the AI system may not be covered by HIPAA's contractual protections. Many organizations have a BAA with their CCaaS provider and assume that coverage extends to all platform features, including newly introduced AI capabilities. That assumption is worth verifying explicitly.

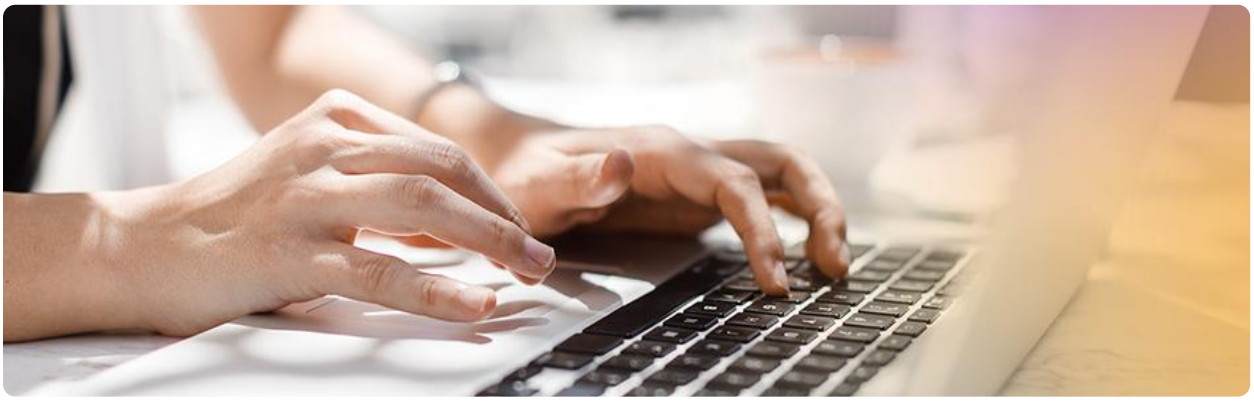

Beyond the BAA: Technical Controls That Matter

A BAA is necessary but not sufficient. The technical architecture must reinforce the contractual protections:

- Customer data must never be used to train, fine-tune, or improve AI models. This must be enforced both contractually and through platform-level controls on the AI service.
- AI inference must follow a bounded path: data is sent from the platform to the AI service, processed, and the response is returned. Nothing should be stored by the model after the interaction.
- All processing should remain within defined geographic boundaries (e.g., US-based regions only).
- Encryption must be maintained in transit (TLS 1.2 or higher) and at rest within the AI processing infrastructure.
- Access to AI systems must follow least-privilege principles, with logging sufficient to support security monitoring and incident investigation.

The Emerging Regulatory Landscape

HIPAA was enacted in 1996 for a fundamentally different technology environment. It provides a strong foundation for protecting PHI, but it was not written to address the risks of large language models or real-time AI inference. The proposed HIPAA Security Rule update is a step toward closing that gap, but the update process is ongoing. [6]

In the meantime, organizations should expect their vendors to go beyond minimum HIPAA compliance. A mature AI governance posture includes formal AI risk assessments, threat modeling for AI-specific attack vectors, documented human oversight models, and alignment to frameworks like the NIST AI Risk Management Framework. These are not HIPAA requirements today, but they represent the direction the regulatory environment is moving.

AI as a Compliance Tool, and Not Just a "Compliance Risk"

Most discussions about AI in regulated contact centers focus on the risks AI introduces. Those risks are real and deserve serious attention. But there is a second dimension: AI's potential to improve compliance outcomes at a scale that manual processes cannot achieve.

The Scale Problem

Contact centers in healthcare, financial services, and outbound operations generate thousands of interactions per day across voice, chat, and digital channels. Every interaction is a potential compliance event. An agent might fail to deliver a required disclosure. A sales representative might make a claim that violates state regulations. A debt collector might miss a consent requirement.

Traditional quality assurance processes review a small percentage of interactions. Industry estimates suggest that many compliance teams manually monitor only a small percentage of customer communications, often less than 5%. [11] Violations that occur outside the sample window are discovered only when they result in a complaint, a regulatory inquiry, or a lawsuit.

The structural limitation of manual compliance monitoring is not effort or intent. It is coverage. When the vast majority of interactions go unreviewed, compliance risk is a function of probability, not prevention.

AI Is Changing What's Possible at Scale

AI is changing which tasks can be performed at scale and who can perform them. Processes that once required dedicated teams of specialists — threat modeling, 100% interaction quality review, real-time regulatory monitoring — are becoming accessible to organizations that previously lacked the resources to do them at all. It lowers the expertise barrier and eliminates the sampling constraint, making comprehensive coverage possible where only spot-checking existed before.

Gartner projects that by end of 2026, 40% of enterprise applications will integrate task-specific AI agents, up from less than 5% in 2025. [15] Organizations that build governance frameworks now will have a structural advantage over those that adopt AI reactively.

This pattern is already visible in cybersecurity, one of the most resource-constrained fields in any organization. AI-enhanced tools are producing threat models that match expert quality at a fraction of the effort, making the process accessible to organizations without dedicated security teams. Compliance monitoring in the contact center is following the same trajectory.

The pressure is most acute in regulated industries. Financial services firms now face an average of 234 regulatory alerts per day, a 25-fold increase from just over a decade ago. [16]

Healthcare AI regulation is proliferating at the state level with no signs of slowing. Organizations that use AI to manage this complexity will have a structural advantage over those trying to keep pace manually. The question isn't whether AI gets adopted in regulated contact centers. It's whether the governance framework is in place when it does.

Where AI Changes the Equation

Large language models and speech analytics technologies can analyze every interaction, in real time or post-interaction, and evaluate it against compliance criteria. This has significant implications across several regulatory domains:

TCPA and Outbound Communications

The TCPA regulatory environment is complex and the penalties are severe. The FCC's 2024 ruling confirmed that AI-generated voices are artificial voices under the TCPA, requiring prior express consent.[9] The FCC continues to pursue rulemaking around consent revocation and caller identification requirements for AI-generated communications.[10] AI-powered monitoring can verify consent was obtained before each contact, confirm required disclosures were delivered, validate opt-out requests were honored, and verify Do-Not-Call registry compliance, across every interaction, not a sample.

Sales Practice and Disclosure Compliance

Financial services organizations must ensure that agents do not make misleading claims, omit required disclosures, or engage in practices that violate the FTC Act, Dodd-Frank, TILA, or ECOA requirements. State-level insurance and lending regulations add further obligations. AI can evaluate interactions for the presence or absence of required statements, flag language patterns that suggest compliance deviations, and surface systemic issues across agent populations that would be invisible in a random sample.

State-by-State Regulatory Variation

Contact centers operating across multiple states face a patchwork of regulations: varying consent requirements for call recording, different rules for outbound communications, and state-specific privacy laws. Increasingly, states are enacting AI-specific requirements including disclosure mandates when consumers interact with AI systems. AI monitoring can be configured to evaluate interactions against state-specific criteria, flagging potential violations that a generalized QA scorecard would miss.

HIPAA and PHI Handling

In healthcare contact centers, AI can monitor interactions for inadvertent PHI disclosure, verify that required privacy notices are communicated, and flag conversations where agents may be handling sensitive information outside approved workflows. This is particularly relevant as remote and hybrid work models make direct supervision more difficult.

Communications Oversight in Financial Services

Certain financial services organizations face large volumes of recorded interactions that must be monitored for missing disclosures, improper recommendations, and compliance deviations under applicable communications oversight and recordkeeping requirements.[11] AI offers the possibility of comprehensive, consistent analysis across every recorded interaction.

AI does not eliminate compliance risk. Human judgment, policy design, and organizational accountability remain essential. But AI transforms compliance monitoring from a sampling exercise into a comprehensive, consistent process. For regulated industries that face penalties measured in hundreds of thousands or millions of dollars per incident, that shift in coverage is significant.

A Framework-Driven Approach to AI Governance

Regulated industries need more than vendor assurances. They need to see that AI governance is built on recognized frameworks with documented controls, risk assessments, and accountability structures.

NIST AI Risk Management Framework (AI RMF 1.0)

Published by the National Institute of Standards and Technology in January 2023, the AI RMF provides a structured, voluntary approach to managing AI risks across the full system lifecycle. It was developed collaboratively with input from over 240 organizations.[12] The framework is organized around four core functions:

Function	Purpose
GOVERN	Establishes policies, roles, accountability, and a culture of AI risk awareness across the organization.
MAP	Identifies the context in which AI systems operate: intended use, data types, stakeholders, and potential impacts.
MEASURE	Assesses and tracks identified risks using methods appropriate to the deployment.
MANAGE	Prioritizes risks and implements technical mitigations, operational procedures, and governance actions.

NIST also released the Generative AI Profile (NIST AI 600-1) in July 2024, extending the AI RMF to address risks specific to large language models, including data privacy, content provenance, and information integrity.[13]

MITRE ATLAS

MITRE ATLAS (Adversarial Threat Landscape for Artificial-Intelligence Systems) provides the threat intelligence layer for AI security. Modeled after the widely adopted MITRE ATT&CK framework, ATLAS catalogs adversary tactics, techniques, and procedures that target AI and machine learning systems. As of early 2026, ATLAS documents 16 tactics, 84 techniques, and 56 sub-techniques across AI-specific attack categories including prompt injection, data poisoning, model extraction, and exploitation of AI service APIs. [14]

For regulated contact centers, where AI systems may access sensitive data and influence customer interactions, the threat coverage ATLAS provides is directly relevant and complements traditional security frameworks like MITRE ATT&CK.

NIST AI RMF provides the governance structure. MITRE ATLAS provides the adversarial intelligence. Together, they enable AI governance programs that are both policy-complete and threat-informed.

Five Controls You Should Expect from Your CCaaS Provider

The following controls represent what organizations in regulated industries, and any organization that values data protection, should look for when evaluating a CCaaS provider's AI governance posture.

1. Data Isolation and No-Training Guarantees

Customer data must never be used to train, fine-tune, or improve AI models. This must be enforced both contractually and technically. Data should follow a bounded inference path with no persistence after processing. For healthcare, the BAA chain must extend through the AI processing layer. For financial services, contractual protections must address GLBA's safeguard requirements for nonpublic personal information.

2. Human Oversight Proportional to Risk

AI used for employee productivity (drafting, research, analysis) should require human review before outputs are used. AI operating in customer-facing contexts requires layered oversight: pre-deployment validation, runtime guardrails, post-interaction review, and emergency rollback capability. This proportional approach aligns with the NIST AI RMF and with emerging state-level requirements for human oversight of AI. [12]

3. AI-Specific Threat Modeling

Traditional penetration testing and vulnerability scanning do not cover AI-specific attack vectors. Providers should conduct threat modeling that addresses data poisoning, prompt injection, data exfiltration through AI systems, and output manipulation. MITRE ATLAS provides the taxonomy for this work. [14] AI-related risks should be tracked through the provider's enterprise risk register alongside traditional security risks.

4. Compliance Integration, Not Compliance Exception

AI capabilities should be covered by the same control families that support the provider's existing compliance programs: access control, logging, change management, risk assessment, vendor management. AI should not create a gap in SOC 2 posture, should be scoped separately from PCI cardholder data environments, and for HIPAA customers, should be covered by the BAA chain through the inference layer.

5. Transparency and Customer Control

AI features should be optional. Governance documentation should be available upon request. Architecture walkthroughs should be available so customers can verify data flows. And the provider's security team should be accessible for direct engagement on questions that go beyond published materials. Transparency is not a liability. It is how trust is established in environments where trust matters most.

What to Ask Your CCaaS Provider

These questions are designed to help you evaluate AI governance maturity during vendor selection or ongoing vendor management. A provider that can answer them clearly and provide supporting documentation has moved beyond AI capability and into AI governance maturity.

Architecture and Data Protection

- Can you confirm, both contractually and technically, that customer data is never used for AI model training?
- Where does AI inference processing occur? Can it be restricted to specific geographic regions?
- What is the data retention behavior of the AI service? Is anything stored after inference is complete?

HIPAA and Healthcare

- Is your BAA chain documented through the AI processing layer, including the specific AI services used?
- Is the AI service on the cloud vendor's official HIPAA Eligible Services list?
- How do you address emerging state-level AI disclosure and oversight requirements for healthcare?

Financial Services and Regulatory Compliance

- How are AI systems scoped relative to PCI cardholder data environments?
- Can AI capabilities support monitoring for TCPA consent, required disclosures, and opt-out compliance?
- How do you address Dodd-Frank communication recording requirements when AI is involved in interactions?

Governance and Frameworks

- Does your AI governance program align to the NIST AI Risk Management Framework?
- Do you conduct AI-specific threat modeling using MITRE ATLAS or an equivalent methodology?
- Can I review your AI use policy, standards, and supporting governance documentation?
- Do you maintain an AI use inventory with documented risk assessments and approval records?

Sharpen's Perspective

Sharpen serves healthcare, financial services, and other organizations where compliance isn't a checkbox — it's a baseline. We wrote this paper because the industry needs more transparency about how AI governance actually works, not just what vendors promise.

Our approach is built on three principles: customer data stays protected, customers stay in control, and our controls are verifiable. That means a formal AI governance program overseen by the CISO, alignment to the NIST AI Risk Management Framework and MITRE ATLAS for threat modeling, and an architecture built around no-train/no-retain from the ground up — not retrofitted.

Our AI capabilities run on Amazon Bedrock, within AWS infrastructure, on services that are on AWS's official HIPAA Eligible Services list. Our BAA chain is documented through the AI processing layer. Cardholder data is architecturally excluded from the platform — not masked after the fact, never present to begin with. AI features are optional, our governance documentation is available upon request, and our security team is available for architecture walkthroughs when the standard documentation isn't enough.

AI governance is built through documentation, architecture, and verifiable controls — not vendor claims. If this paper raises questions about your current platform, or about ours, we're happy to get into the specifics.

For questions about AI governance, compliance documentation, or to schedule a security architecture walkthrough, contact your Sharpen account representative or reach out to our security team.

Sharpen

References

[1] McKinsey & Company, "The State of AI: Global Survey," 2025. mckinsey.com

[2] Vanta, "The State of Trust: AI Governance and Compliance," 2025. vanta.com

[3] ECRI, "Top 10 Health Technology Hazards for 2025." ecri.org

[4] IBM Security, "Cost of a Data Breach Report 2025." ibm.com/security/data-breach

[5] Wolters Kluwer Health, "AI Privacy and Security Concerns Survey," December 2025. wolterskluwer.com

[6] U.S. Department of Health and Human Services, Office for Civil Rights, "Proposed Updates to the HIPAA Security Rule," January 2025. hhs.gov/hipaa

[7] Akerman LLP, "New Year, New AI Rules: Healthcare AI Laws Now in Effect," January 2026. akerman.com

[8] Federal Trade Commission, "Gramm-Leach-Bliley Act: Financial Privacy and Safeguards Rules." ftc.gov/business-guidance/privacy-security/gramm-leach-bliley-act

[9] Federal Communications Commission, "Declaratory Ruling: TCPA Applies to AI Technologies that Generate Human Voices," February 2024. fcc.gov/document/fcc-confirms-tcpa-applies-ai-technologies-generate-human-voices

[10] The CommLaw Group, "FCC Extends Revoke All Consent Rule Effective Date," January 2026. commlawgroup.com

[11] ASC Technologies, "Dodd-Frank Compliance: AI-Based Monitoring of Conversations," 2025. asctechnologies.com

[12] National Institute of Standards and Technology, "AI Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, January 2023. nist.gov/itl/ai-risk-management-framework

[13] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative AI Profile," NIST AI 600-1, July 2024. nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[14] MITRE Corporation, "ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems." atlas.mitre.org

[15] Gartner, Inc., "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," August 2025. gartner.com/en/newsroom/press-releases/2025-08-26

Additional Resources

- NIST AI Resource Center (AIRC). airc.nist.gov
- Amazon Web Services, HIPAA Eligible Services Reference. aws.amazon.com/compliance/hipaa-eligible-services-reference
- PCI Security Standards Council, PCI DSS v4.0.1. pcisecuritystandards.org
- AICPA, SOC 2 Trust Services Criteria. aicpa.org

About Sharpen

Sharpen is a cloud-native Contact Center as a Service (CCaaS) platform serving healthcare, financial services, and other organizations where compliance and data protection are non-negotiable. With an AI governance program aligned to the NIST AI Risk Management Framework and MITRE ATLAS, Sharpen enables organizations to adopt AI in their contact center with confidence.