Pause-Resume Is Dead in the Age of AI: Why Manual Call Recording Controls No Longer Satisfy PCI DSS 4.0.1 in AI-Enabled Voice Environments

August 21, 2026

By William Placke, Certified Information Privacy Professional-US (CIPP-US), Juris Doctorate, Diploma EU Law, HBS Certificate in Data Privacy and Data Security, President, SecurePII Americas.

This publication is informational and is not legal or professional advice. Readers should consult their own legal counsel, Qualified Security Assessor, or Internal Security Assessor for guidance specific to their facts and circumstances.

Executive Summary

Pause-and-resume call recording control has served for more than a decade as the default method by which telephone-based merchants and call centers attempt to keep spoken cardholder data (“CHD”) and sensitive authentication data (“SAD”) out of the recorded environment. The PCI Security Standards Council (“PCI SSC”), from its earliest guidance, described Pause-Resume as a method whose effectiveness depends on an agent, or a system integrated with an agent's workflow, acting correctly at a precise moment in a live conversation. It is the author’s observation that over the past two years, upwards of 30% of PCI audit failures have arisen from the failure of Pause-Resume systems, even when there are automated controls to look for human error. Rightly so, QSAs do not report publicly on the causes of audit failure. Still, my observation stands and is supported through not just personal experience but from extrapolation from inferable error rates in relatable manual processes in section IIIB below. The author is open to any supporting or contrary input to establish the validity of this conclusion.

This paper examines pause-and-resume on its own terms, using the PCI SSC's own published guidance and the current PCI DSS v4.0.1 Requirements and Testing Procedures (June 2024), and finds that the control's known and documented failure modes are structural. The paper then examines the most common proposed backstop to a missed pause, after-the-fact redaction of the recording, and finds that redaction does not resolve the underlying problem so much as relocate it, since a system cannot redact data it has not first recognized, and recognition of SAD is itself a form of processing that the recording may not survive contact with intact. For these reasons, and given the availability of DTMF suppression technologies, Pause-Resume is dead technology. There is no reason that any organization that continues the use of Pause-Resume, when they are inevitably found deficient under PCI-DSS 4.0, should be surprised when they lose the right to take credit cards as a means of payment.

The introduction of artificial intelligence into the call center, in the form of transcription, sentiment analysis, and conversational summarization clearly removes the narrow set of conditions under which pause-and-resume's known weaknesses could be contained to an occasional, correctable incident. Once a recording is digitally transcribed, indexed, or fed into a model, a missed pause is no longer a discrete, findable, and deletable event. It is a data point inside a searchable corpus. The PCI SSC's own guidance is explicit that once cardholder data is digitally queryable, it must not be stored at all.

This paper does not name or evaluate any specific vendor, platform, or product. Its claims are addressed to any VoIP environment, regardless of the underlying technology stack, that relies on agent-initiated or

workflow-integrated pause-and-resume as its primary or sole control for SAD in voice channels.

I. Introduction and Thesis

Telephone-based payment acceptance remains common across retail, healthcare, utilities, financial services, and government. Where a customer speaks payment card information aloud to an agent, the PCI SSC's guidance is unambiguous that this act places the personnel, the technology, and the infrastructure receiving that data into the scope of PCI DSS.

Pause-and-resume technology was developed to narrow that scope by preventing the moment of spoken card data from ever entering the call recording in the first place. The premise is that if the recording is paused before the cardholder speaks the card number and resumed only after the number has been given, the recording itself never contains SAD. Accordingly, the recording and storage systems can be considered outside the boundary of the Cardholder Data Environment ("CDE").

That premise depends entirely on the pause occurring at the correct moment, every time, across every agent, every shift, and every transaction, years on end. This paper's thesis is that the premise was already on the worst of shaky ground before generative and analytic AI entered the call center, and that AI-driven transcription and analysis removes the conditions under which an occasional missed pause could be treated as a contained, correctable event rather than a permanent addition to a searchable, storable, and PCI-relevant data set.

II. How Pause-and-Resume Is Defined and Where It Sits in the Standard

PCI DSS itself does not certify "pause-and-resume" as a named technology. It is one recognized implementation method for satisfying Requirement 3.3.1, which provides that sensitive authentication data is not stored after authorization, even if encrypted, and that all SAD received is rendered unrecoverable upon completion of the authorization process.[1]

The PCI SSC's own supplemental guidance, Protecting Telephone-Based Payment Card Data, describes the goal of pause-and-resume as preventing account data from being recorded by temporarily halting the recording during the communication of that data and resuming only once the payment data transmission is complete.[2] The same guidance is careful to note that even a properly implemented pause-and-resume solution reduces the applicability of PCI DSS to the call-recording and storage systems specifically. It does not reduce the applicability of PCI DSS to the agent, the agent desktop environment, or any other system in the telephone environment.[3]

Requirement 3.3.1 is not eligible for the customized approach.[4] An assessor cannot substitute a demonstrated compensating outcome for the defined method the requirement specifies. Table 3 of the standard reinforces the same point in its data-element framework: sensitive authentication data cannot

---

[1]Payment Card Industry Data Security Standard: Requirements and Testing Procedures, v4.0.1 (June 2024), Requirement 3.3.1, at p. 79.

[2]PCI Security Standards Council, Information Supplement: Protecting Telephone-Based Payment Card Data, v3.0 (Nov. 2018), §6.5.1, at p. 36.

[3]Id.

[4]PCI DSS v4.0.1, Requirement 3.3.1, at p. 79 ("This requirement is not eligible for the customized approach.").

be stored after credit card authorization, and this restriction applies with equal force whether or not the data is encrypted.[5]

III. The Execution Problem, Manual and Automated

A. Manual Pause-and-Resume

The PCI SSC's own guidance catalogs the failure modes of manual pause-and-resume without qualification. Manual implementations, in the Council's words, rely completely on agent personnel to pause and restart the recording at exactly the right time.[6] The two failure modes the guidance identifies are the agent forgetting to pause before the card data is spoken, which results in the unintended capture of CHD and potentially SAD, and the agent forgetting to resume after the transaction, which can result in a breach of regional or local legal requirements and loss of other business data.[7]

The same guidance states plainly that manual pause-and-resume implementations require constant monitoring and regular verification, and that the degree of oversight and supervision required for manual solutions is materially greater than for automated solutions.[8] *This is the PCI SSC's own description of a control category it has already identified as failure-prone by design*. It has been my observation that upwards of 30% of failed PCI audits over the past two years have been caused by auditors finding cardholder data in enterprises due to the failure of Pause-Resume technology.[9] As my colleague, Jason Thals so well articulated in his paper based on his decades of Call Center operations experience, *The Platypus Problem: What PCI-DSS 4.0.1 Actually means for Contact Centers*, "human behavior is where Pause/Resume is most exposed. Security frameworks define controls. Humans execute workflows. Quality teams override Pause/Resume because agents misuse it, for example, pausing during difficult customer interactions to avoid being monitored."

B. What Independent Research on Manual, Repetitive Tasks Suggests

The PCI SSC's guidance does not publish a failure rate specific to pause-and-resume execution, and no PCI SSC document reviewed for this paper does either. In the absence of a task-specific figure, it is reasonable, and analytically honest, to look to independent research on comparable manual, repetitive, procedural tasks performed under time pressure, and to extrapolate an expected range rather than assert a measured one.

The most extensively cited body of research on manual data-entry accuracy is Raymond R. Panko's synthesis of decades of spreadsheet and data-entry error studies. Panko's research found that individual data-entry accuracy in manual tasks generally runs in the 96 percent to 99 percent range per entry, meaning a per-entry error rate of roughly 1 percent to 6 percent, and that this error rate compounds

---

[5]PCI DSS v4.0.1, Table 3, Account Data Element Storage Requirements, at pp. 6–7.

[6]Protecting Telephone-Based Payment Card Data, v3.0 (Nov. 2018), Diagram 9, at p. 37.

[7]Id.

[8]Id.

[9] See Jason Thals, "PCI DSS v4.0.1 Data Security Standards: Anomalous Cardholder Data, Scope Drift, and the End of Pause/Resume in PCI Compliance" (SecurePII, v.1.0.1, Apr. 10, 2026), colloquially the "Platypus Problem" paper, summarized at https://www.securepii.cloud/the-platypus-problem/.

materially as the volume or complexity of entries increases.[10] At the level of the completed work product rather than the individual entry, Panko's review of audited, in-production spreadsheets found that the probability a given spreadsheet already in active business use contained at least one material error ran between 80 percent and 90 percent across the studies he surveyed.[11]

A separate, peer-reviewed study of manufacturing personnel performing repetitive manual assembly tasks found that stress, repetition, fatigue, and work environment together statistically accounted for 48.8 percent of the variance in observed human error, based on a survey of 200 manufacturing workers across multiple facilities.[12] The same body of literature has separately found that human error accounts for the overwhelming majority, cited in that study as approximately 90 percent, of workplace accidents generally, underscoring that error rates of this magnitude are not an artifact of low-skill or careless personnel but a recurring feature of repetitive procedural work performed by trained staff.[13]

Neither study concerns call centers, and neither concerns pause-and-resume specifically. The extrapolation this paper draws is narrow and stated as such that a call center agent executing a manually triggered pause, under live conversational pressure, while simultaneously listening to a customer and operating a payment application, is performing a task that shares the defining characteristics of the tasks studied above, namely a repetitive, time-pressured, procedural action embedded inside a longer workflow. There is no principled basis in the available research to expect this task to be materially more error-resistant than manual data entry or repetitive manual assembly. A published, task-specific failure rate for pause-and-resume execution in live call center environments does not currently appear to exist in the public literature that the author has found.

Buttressing this data is the author’s own personal experience across the PCI compliance ecosystem. Over the past two years, the author has observed upwards of 30% of audit failures involve the QSA or ISA discovering cardholder data in recordings and the underlying data extrapolations. This is likely due to a review of data where a company is seeking large amounts of prior data to go into an LLM as source data to create better business outcomes. Upon the discovery, the Jenga block pyramid collapses all over the floor.

C. Automated Pause-and-Resume

The PCI SSC's guidance describes automated pause-and-resume as typically integrated with the desktop application an agent uses during the transaction, configured so that the recording pauses automatically when the agent clicks into a payment-entry field or launches a payment screen, and resumes automatically when the agent submits the transaction or moves to the next screen.[14] Industry implementations vary the trigger mechanism, including window or URL changes in the agent's desktop environment and custom triggers invoked through an application programming interface, but the

---

[10]Raymond R. Panko, “What We Know About Spreadsheet Errors,” Journal of End User Computing, Vol. 10, No. 2 (Spring 1998), pp. 15–21, revised May 2008. Panko's synthesis draws on multiple independent field and laboratory audits of operational spreadsheets and manual data-entry tasks.

[11]Id.

[12]Jian Ai Yeow, Poh Kiat Ng, Khong Sin Tan, Tee Suan Chin, and Wei Yin Lim, “Effects of Stress, Repetition, Fatigue and Work Environment on Human Error in Manufacturing Industries,” Journal of Applied Sciences, Vol. 14, No. 24 (2014), pp. 3464–3471.

[13]Id., citing A. DiDomenico and M.A. Nussbaum (2008).

[14]Protecting Telephone-Based Payment Card Data, v3.0 (Nov. 2018), §6.5.1, at p. 38.

underlying architecture, an external event triggering the pause rather than the agent's memory, is consistent across these approaches.[15]

Automation removes the agent's burden of remembering to act, and the PCI SSC's guidance credits this approach with requiring materially less ongoing supervision than a manual process.[16] The same guidance, however, does not describe automation as curing the underlying exposure. It states, without qualification, that the effectiveness of an automated solution relies largely on its integration with the agent's workflow process and the agent performing the correct steps at the correct time, and that if any ability exists for the agent to bypass the integrated process, the pause-and-resume technology could be circumvented and rendered ineffective.[17]

This is the PCI SSC's own acknowledgment, not an inference this paper is drawing, that automated pause-and-resume shifts the point of failure rather than eliminating it. The dependency moves from an individual agent's memory in a single call to the completeness of a workflow integration across every application screen, every call type, and every path a payment can take through the agent desktop. A gap in that integration, an unanticipated screen, a browser-based payment flow the pause trigger was not configured to detect, an agent taking payment information before formally opening the payment screen, produces the same result as a forgotten manual pause. The PCI SSC's guidance states this conclusion directly elsewhere in the same document where it states that recordings may still capture CHD or SAD where pause-and-resume is used, depending on the accuracy of the pause-and-resume process, whether manual or automated.[18] A concrete illustration of this dependency has been documented in industry commentary on this exact failure mode: an automated trigger configured to detect a payment-window URL can fail to fire when a network disruption delays or drops that detection event, again allowing card data into the recording notwithstanding the presence of automation.[19]

This same section of the PCI SSC's guidance contains the passage most directly relevant to the thesis of this paper. *The Council states that if SAD contained within audio recordings can be digitally queried, meaning if the SAD is easily accessible through search or indexing, it must not be stored at all.* The Council further states that where recordings cannot be data-mined, storage of SAD after authorization may be permissible only if appropriate validation has been performed.[20] ***This distinction, between a recording that merely exists and a recording that is digitally queryable, is the hinge on which the remainder of this paper turns.*** AI-driven transcription, indexing, and sentiment analysis convert every recording into the queryable category. A missed pause that once might have sat, undiscovered and effectively inert, inside an unindexed audio archive becomes, once transcribed, a locatable, searchable data point the moment the transcription pipeline runs. The 2018 guidance already anticipated this consequence in general terms.

---

[15]See, e.g., Dubber, "Keeping Call Recording PCI Compliant" (2023), describing automatic pause-and-resume triggered by window or URL changes or custom API triggers. Cited here as a description of an industry-standard mechanism category, not as an endorsement of any specific vendor implementation.

[16]Id., §6.5.1, at p. 38.

[17]Id. (emphasis added).

[18]Id., §6.6, at p. 40.

[19]Jason Thals, "PCI DSS v4.0.1 Data Security Standards: Anomalous Cardholder Data, Scope Drift, and the End of Pause/Resume in PCI Compliance" (SecurePII, v.1.0.1, Apr. 10, 2026), colloquially the "Platypus Problem" paper, summarized at https://www.securepii.cloud/the-platypus-problem/ ("Automatic pause processes fail silently when network disruptions mean a payment window URL goes undetected.").

[20]Id., §6.6, at p. 40 ("If SAD contained within audio recordings can be digitally queried—if SAD is easily accessible—it must not be stored. If these recordings cannot be data mined, storage of SAD after authorization may be permissible as long as appropriate validation has been performed.").

AI-enabled call platforms are the mechanism by which the general concern becomes the operating reality in a modern contact center.

IV. What a QSA or ISA Reviews in a Pause-and-Resume Environment

A Qualified Security Assessor or Internal Security Assessor examining a telephone environment that relies on pause-and-resume, whether manual or automated, will typically examine the following areas. This list draws on the applicable Requirements and Testing Procedures and the corresponding ROC Template reporting instructions.

- CDE boundary and call flow documentation. The assessor will require data flow diagrams establishing precisely where CHD and PAN travel through the telephone environment and which systems, technologies, or processes have the ability to affect the security of that data.
- Call recordings, screen recordings, and any transcription, indexing, or analytics capability layered on top of them. Because a recording that can be queried or copied is treated differently than one that cannot, the presence of any searchable index is a material fact for the assessment. Requirement 3.3.1's prohibition on storing SAD applies even where the data is encrypted, which forecloses an encryption-based defense of an unintended capture.
- Third-party service provider management, where any recording, transcription, or AI-analytics function is outsourced. Requirement 12.8.1 requires the entity to maintain a list of all third-party service providers with whom account data is shared or that could affect the security of account data, together with a description of the services each provides.[21] Requirement 12.9.1 separately requires the third-party service provider itself to furnish a written agreement acknowledging its responsibility for the security of the account data it stores, processes, or transmits on the customer's behalf.[22]
- Personnel screening and training. Requirement 12.7.1 requires interviews of Human Resources personnel to confirm that screening, within the constraints of local law, is conducted prior to hiring personnel who will have access to the CDE.[23] Requirement 8.3.8 is tested by interviewing users to confirm they are familiar with the entity's authentication policies and procedures. This requirement concerns familiarity with authentication policy specifically. It should not be cited, as it has sometimes loosely been cited in industry materials, as a general audit hook for pause-and-resume conduct; its scope is narrower and more specific than that.[24]
- The scope consequences of AI-enabled call processing. Where an entity has deployed AI transcription, summarization, or sentiment analysis on top of recorded calls, an assessor evaluating current PCI DSS 4.0.1 obligations has a reasoned basis to treat those tools as within the same scope analysis applied to any other system capable of capturing, processing, or storing CHD or SAD, since the standard's scoping principles turn on where the data can be heard, seen, processed, transmitted, or stored rather than on the label applied to the underlying technology.

V. Post-Call Redaction as a Backstop: Processing, Not Cure

[21]PCI DSS v4.0.1 ROC Template r3 (Jan. 2025), Requirement 12.8.1, at p. 305.
[22]Id., Requirement 12.9.1, at p. 310.
[23]Id., Requirement 12.7.1, at pp. 304–305.
[24]Id., Requirement 8.3.8, at p. 192.

Where an entity anticipates that pause-and-resume, manual or automated, will occasionally fail, the most common proposed remedy is a post-call redaction or masking layer applied to the recording or its transcript. The architecture is that the full call is recorded, a system subsequently scans the recording or transcript for patterns resembling a card number, and any match is masked before the recording is made available for playback or review.

This architecture does not resolve the problem Requirement 3.3.1 addresses. It relocates it. A system cannot selectively mask a card number without first identifying where in the audio or transcript that card number occurs. That identification step is itself an act of processing the sensitive authentication data, and it necessarily occurs after the full, unredacted recording already exists in some form, however briefly. Requirement 3.3.1 requires that SAD be rendered unrecoverable upon completion of the authorization process, and states expressly that this obligation applies even if the data is encrypted.[25] The same logic that forecloses an encryption-based defense of a retained recording forecloses a redaction-based one. Masking what a human reviewer can subsequently see or hear is a different act from rendering the underlying data itself unrecoverable at the moment the standard specifies.

The more consequential version of this problem concerns the original, unredacted recording rather than the redacted copy presented to reviewers. A redaction pass, by definition, produces an output that differs from its input. Whether the original, unmasked recording is deleted immediately upon redaction, retained temporarily in an intermediate buffer, or retained indefinitely for quality assurance, model training, or dispute-resolution purposes is a function of the specific system's architecture and retention policy, not an inherent feature of “redaction” as a concept. Where that original recording persists anywhere, including in intermediate storage, processing buffers, or backup systems, the entity is storing SAD notwithstanding the existence of a redacted copy layered on top of it. A masked display does not answer the separate question of whether the source data was ever rendered unrecoverable.

This is not merely a theoretical risk. The public documentation of at least one major cloud contact-center platform confirms the pattern directly where its administrative controls are explicitly designed so that certain users can retain access to unredacted, original recordings and transcripts even after automatic redaction has been applied elsewhere in the system, with access governed by permission settings rather than by the original ceasing to exist.[26] The same platform's own technical guidance separately confirms that a redaction pass can write the redacted artifact as an addition alongside the original recording rather than in place of it, so that both versions persist unless the original is independently locked down or deleted.[27]

---

[25]PCI DSS v4.0.1, Requirement 3.3.1, at p. 79.

[26]AWS, “Amazon Connect now provides granular permissions for conversation recordings and transcripts” (Oct. 23, 2025), available at https://aws.amazon.com/about-aws/whats-new/2025/10/amazon-connect-granular-permissions-conversation-recordings-transcripts/ (“The system also provides flexible download controls, enabling users to download redacted recordings while restricting downloads of unredacted versions. Administrators can also create sophisticated permission scenarios, providing access to redacted recordings of sensitive conversations while granting unredacted recording access for other conversations.”). This paper cites the passage as an illustration of a documented industry pattern rather than a critique of any specific platform.

[27]AWS re:Post, “How do I automatically redact PII from Amazon Connect call transcripts and recordings using Contact Lens?” available at https://repost.aws/articles/AR4WdXYxhhSx2YyoJRCeatQQ/how-do-i-automatically-redact-pii-from-amazon-connect-call-transcripts-and-recordings-using-contact-lens (“When you enable redaction in the contact flow, Contact Lens writes a redacted analyzed transcript (and, optionally, the original analyzed transcript)” to storage, and describing “how to lock down access to original (un-redacted) files, and the limits of what redaction can and cannot do.”).

Technical implementation guidance corroborates the same architectural point independent of any single vendor's documentation. One widely referenced implementation guide for cloud-based call-recording redaction describes a design in which an access-restricted storage location retains the unredacted source recording while a separate, more broadly accessible location holds the redacted output, precisely the two-copy architecture this paper's inference anticipated.[28] Industry commentary on redaction reaches the same conclusion by a different route, observing plainly that redaction, even where it improves on manual agent reliance, still leaves the unredacted original somewhere in the pipeline.[29]

This point is general to the architecture of after-the-fact redaction and is not a claim about the reliability of any specific product's pattern-matching accuracy. Even a hypothetically perfect redaction engine, one that never missed a card number and never produced a false negative, would still need to answer the antecedent question of what happens to the recording that existed before the redaction pass ran.

VI. Practical and Evidentiary Consequences

A. Compliance Cost

Completing a Report on Compliance under PCI DSS v4.0.1 requires evidence spanning close to 400 pages of requirements and testing procedures.[30] Industry reporting places the cost of third-party assessment for entities with complex, multi-location telephone environments from the tens of thousands of dollars into the several hundred thousand dollar range, depending on the breadth of in-scope systems.[31] A pause-and-resume architecture that keeps the call recording and storage environment within scope, whether because the pause is manual, because the automation depends on incomplete workflow integration, or because a redaction layer leaves an unresolved question about the source recording, keeps the full weight of that evidentiary burden in place year over year.

B. The Evidentiary Gap in Payment Disputes

A separate and distinct consequence of pause-and-resume architecture concerns what happens after the call, specifically in a payment dispute or chargeback scenario. Where the recording is paused during the portion of the call in which payment authorization occurred, the merchant has, by design, no recording of that segment of the call. If a cardholder later disputes the charge as unauthorized, the merchant's evidentiary record of the authorization moment itself does not exist, precisely because the control that kept the environment out of PCI scope also eliminated the merchant's own proof of what was said and agreed to at the moment of payment.

This is a business and evidentiary consequence rather than a PCI DSS compliance requirement in itself, and this paper does not represent it as a defined obligation under the standard. It is, however, a direct

---

[28]Redactor.ai, “Redact Call Recordings Stored in AWS S3” (2026), available at https://redactor.ai/blog/redact-call-recordings-aws-s3, describing an architecture in which an access-restricted bucket retains the unredacted source recording while a separate bucket holds the redacted output for broader access. Cited as an illustrative implementation pattern, not as an endorsement of this or any vendor's product.

[29]Paytia, “PCI Compliance and Call Recording Guide,” available at https://www.paytia.com/resources/blog/pci-compliance-and-call-recording-guide (describing post-call redaction as an improvement over manual agent reliance that nonetheless “still leaves the unredacted original somewhere in the pipeline”).

[30]PCI DSS v4.0.1 Requirements and Testing Procedures (June 2024).

[31]SecurityMetrics, “How Much Does PCI Compliance Cost,” by Gary Glover, available at https://www.securitymetrics.com/blog/how-much-does-pci-compliance-cost.

and foreseeable economic tradeoff of the pause-and-resume approach that merits attention alongside the compliance analysis above. That is, the same architectural choice that reduces PCI scope by removing the payment moment from the recording also removes the merchant's own record of that moment. An architecture that instead masks cardholder data at the point of entry, whether through DTMF capture, tokenization, or a comparable method, while allowing the recording to continue uninterrupted through the full transaction, preserves a continuous evidentiary record while still keeping the spoken or keyed card number itself out of the merchant's systems and recordings.

Mastercard's own published guidance for merchants confirms that phone records documenting a conversation with the customer belong among the categories of evidence a merchant should collect in advance and be prepared to submit in a chargeback dispute, alongside delivery confirmation, purchase history, and address or card-verification matches.[32] The same guidance states plainly that merchants are generally liable for all chargebacks arising from card-not-present transactions, a category that includes every telephone payment addressed in this paper, and that the strength of a merchant's evidence is what principally determines the outcome of a dispute the merchant did not initiate.[33]

One qualification is warranted for precision. Visa's newer Compelling Evidence 3.0 framework, which addresses disputes filed under reason code 10.4 for an allegedly unauthorized card-not-present transaction, does not rely on call recordings. It instead permits a merchant to establish legitimacy by matching device, IP address, and billing-descriptor data points against a prior undisputed transaction from the same cardholder.[34] A recording's evidentiary value accordingly sits within the broader, longer-standing representment process that applies across reason codes and card networks generally, rather than within this specific and newer mechanism, and this paper does not represent otherwise.

C. The Economics of Layering Compliance on Top of Native Recording

Multiple leading VoIP platforms include call recording in their standard offering. Redaction and other PCI-relevant data-loss controls are typically not included at that same base tier. The public documentation of at least one major cloud contact-center platform confirms that automatic sensitive-data redaction lives inside a distinct, separately enabled analytics capability rather than the base contact-center service, requiring its own affirmative configuration step and priced on its own dedicated pricing page rather than bundled automatically into base usage.[35] An organization adopting post-call redaction as its PCI control is

---

[32]Mastercard, “Disputing Chargebacks: Practical Guidance for Merchants” (Jan. 24, 2024), available at https://www.mastercard.com/us/en/news-and-trends/Insights/2024/how-can-merchants-dispute-credit-card-chargebacks.html.

[33]Id. (“Merchants are generally liable for all chargebacks in card-not-present (CNP) transactions.”; listing correspondence, including phone records, among the categories of evidence merchants should collect and be prepared to submit in a dispute).

[34]See, e.g., Chargeback.io, “The Ultimate Guide to Visa Compelling Evidence 3.0,” available at https://www.chargeback.io/blog/the-ultimate-guide-to-visa-compelling-evidence-3-0; Checkout.com, “Visa Compelling Evidence 3.0: New Rules Explained,” available at https://www.checkout.com/blog/visa-compelling-evidence-3-0. Visa Compelling Evidence 3.0 applies only to disputes filed under reason code 10.4 and requires two prior undisputed transactions from the same cardholder sharing specified data points.

[35]See, e.g., AWS, “Amazon Connect Contact Lens is now available in AWS GovCloud (US-West)” (Jul. 1, 2025), available at https://aws.amazon.com/about-aws/whats-new/2025/07/amazon-connect-contact-lens-govcloud-us-west/, and “Amazon Connect Contact Lens now provides free trials for conversational analytics and performance evaluations” (Jan. 8, 2025), available at https://aws.amazon.com/about-aws/whats-new/2025/01/amazon-connect-contact-lens-free-trials-conversational-analytics-performance-evaluations, both directing customers to a dedicated Contact Lens pricing page distinct from core Amazon Connect telephony pricing; Amazon Connect

therefore typically paying twice for two different reasons: once for the underlying telephony and recording platform it already operates, and again for the compliance layer added on top of it, a layer this paper has already shown does not fully resolve the exposure it is purchased to address.

The cost of the compensating architecture does not end at licensing. Where an organization discovers, whether through a QSA finding or an internal audit, that its existing recorded-call archive already contains unredacted cardholder data, remediating that archive after the fact is a significant expense in its own right. Industry commentary places the cost of scrubbing a mid-sized historical archive with audio-redaction tooling at a six-figure project, a cost entirely separate from, and additional to, the ongoing licensing cost of the redaction capability itself.[36] This retroactive cost can arise at any point in the architecture's operating life, each time a gap in coverage is discovered.

Layered against these figures are the compliance-scope costs already discussed in Section VI.A and the evidentiary exposure discussed in Section VI.B. An organization relying on pause-and-resume, automated or manual, with a redaction backstop, is carrying simultaneously the cost of the annual assessment scope those controls fail to reduce, the licensing cost of the redaction layer itself, the contingent cost of archive remediation if a gap is later found, and the dispute-resolution exposure of a chargeback it cannot fully evidence. None of these costs is hypothetical. Each is documented above with its own source.

An architecture that removes cardholder data from the voice channel before it reaches any recording, transcription, or storage system, of the kind SecurePII's SecureCall product is designed to provide, addresses each of these costs at its root rather than after the fact. Because the cardholder data never enters the recording, there is no unredacted archive to remediate and no ongoing redaction license to carry. Because the recording itself continues uninterrupted through the full transaction, including the authorization moment, the evidentiary record described in Section VI.B is preserved rather than sacrificed. Because the cardholder data never enters the merchant's network, desktops, or storage, the annual assessment scope described in Section VI.A is reduced at the architectural level rather than managed procedurally, year after year. The comparison, properly framed, is not between one compliance feature and another. It is between paying, on an ongoing basis, for a set of controls this paper has shown do not fully resolve the exposure they are meant to address, and adopting an architecture that removes the exposure itself.

VII. Scope Note and Transition to a Second Paper

This paper has confined its analysis to the pause-and-resume control as applied to human call center agents, and to the redaction architectures commonly proposed to backstop that control. A related and distinct question, whether routing the payment moment to an AI voice agent or conversational bot in place of a human agent changes this analysis, and what obligations arise from the data used to train or

---

Administrator Guide, “Enable conversational analytics in Amazon Connect Contact Lens,” available at https://docs.aws.amazon.com/connect/latest/adminguide/enable-analytics.html (confirming Contact Lens, whose capabilities include sensitive-data redaction, must be affirmatively enabled on an instance before it becomes available). This paper has no visibility into current or negotiated pricing for this capability and does not assert a specific dollar figure. The point drawn from this source is structural: the capability sits behind separate enablement and separate billing beyond the base contact-center service, not that any particular price applies.

[36]Paytia, “Call Center PCI Compliance: A Practical 2026 US Guide,” available at https://www.paytia.com/us/resources/blog/call-center-pci-compliance-guide (describing remediation of an existing recorded-call archive through audio-redaction tooling as “usually a six-figure project for a mid-sized archive”).

operate such a system, will be addressed in a future paper. That question rests on a different evidentiary footing, particularly regarding the distinction between real-time processing of spoken card data and the separate question of whether such data could enter a model's training pipeline. For now, it is intentionally treated outside the scope of this paper.

Conclusion

Pause-and-resume was never described, even in the PCI SSC's own founding guidance on the subject, as a self-executing or failure-proof control. It was described as a method whose reliability depends on precise human or system behavior repeated correctly across every transaction, and whose own documentation acknowledges that both manual and automated implementations can and do fail. The author’s own experience over the past two years with failed PCI audits coupled with independent research on comparable manual, repetitive tasks gives no reason to expect call center execution to outperform documented error rates in other procedural work. Post-call redaction, the most commonly proposed backstop, does not eliminate the underlying exposure; it relocates the moment of risk to the handling of the source recording the redaction pass itself required. The introduction of AI-driven transcription and analysis into this environment converts what was once an occasional, often undiscovered, and practically inert failure into a permanent, searchable, and PCI-relevant data point, precisely along the line the PCI SSC's own guidance already drew between recordings that merely exist and recordings that can be digitally queried. The compliance and evidentiary consequences that follow are not new risks introduced by AI. They are the consequences of a decade-old control's known limitations, made permanent and discoverable by the same technology that was meant to make call centers more efficient.