PII Redaction | Cresta Chat and Voice

Summary

As a conversation intelligence platform, Cresta handles large volumes of information between contact center agents and their end customers. Cresta's products serve some of the largest contact center organizations, with a majority focus on top-tier enterprise customers. Powering the core of Cresta's various product offerings is a real-time conversation engine that processes data across chat, messaging and voice channels.

In this document, we outline how Cresta handles sensitive personally identifiable information (PII) in a secure, compliant manner while aligning with industry best practices aimed at protecting the privacy of end users and mitigating business risks for customers.

What is Personally Identifiable Information (PII)?

Formal definitions of PII can vary, however generally speaking, it refers to information that can be used by organizations on its own or with other information to identify, contact or locate a single person, or to identify an individual in context.

Cresta follows the NIST guidelines and treats PII at the highest tier of confidentiality for the purposes of sensitive data protection. Additionally, Cresta recognizes that other sensitive data types exist which can be correlated back to a single person/identity, and also takes into account other standards/guidelines such as PCI-DSS v4.0 and the HIPAA Privacy Rule and treats these sensitive data types with due care.

For more information on Cresta's compliance certifications/frameworks, adherence to applicable data privacy regulations, and current security practices/mechanisms, please refer to Cresta's Trust page.

How Cresta Handles PII

Cresta leverages state-of-the-art entity recognition techniques to be able to automatically identify and redact PII within unstructured data in real-time and/or post-conversation. Unlike legacy solutions that tend to either over-redact or under-redact PII, Cresta has internally developed fine-tuned deep learning models that perform highly accurate PII detection and redaction where applicable.

Cresta's PII redaction pipeline supports the redaction of the following PII data types both within the Cresta UI and database:

- Names, Full Names
- Phone Numbers
- Email
- Addresses (i.e. Street Address)
- Date of Birth
- Credit Card Numbers (Partial or Complete)
- CVV / Card Security Codes
- U.S. Social Security Numbers (SSN)
- Drivers' License Number for U.S. States

However, Cresta allows organizations to specify policies around the handling of PII data recognizing that their needs regarding the use of such information can vary within different life-cycle moments (during/after the conversation) and user context (agent interacting with customer vs passive manager/coach observers).

For example, some organizations might want agents to be able to view a customer's email and phone number as it is extracted from a live call transcript to enable them to confirm their identity, but only store such data after first redacting when the call concludes. In addition to the standard list of PII entities listed in the previous section, Cresta also offers the ability to detect and redact custom entities that an organization identifies as sensitive information.

How Cresta Handles PII (Continued)

It should be noted that while real-time PII redaction can be performed on voice calls routed to the Cresta voice agent interface, the same is not supported for chat/messaging since those are handled by agents on an external agent platform that Cresta cannot directly control incoming/outgoing messages on.

Post conversation PII redaction in Cresta Director (Cresta's real-time coaching interface) however is supported on chat/messaging as well as voice. For chat/messaging, redaction is applied on the relevant sections in the conversation transcript, while for voice, redaction is applied to the transcript as well as the specific section within the recorded audio is replaced with a "bleep".

It is also important to note that all detected PII entities are redacted prior to being saved to the database across all channels.

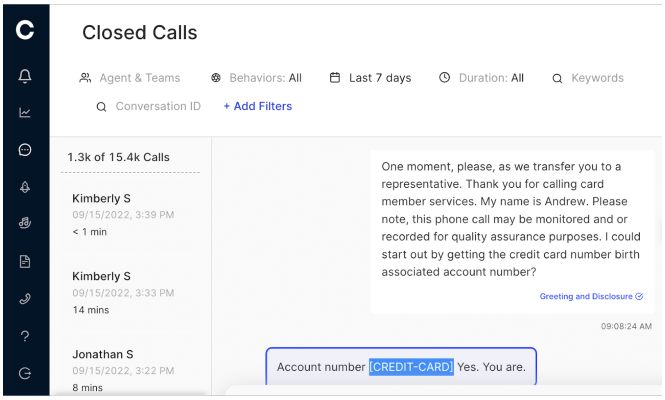

PII Redaction Performance Benchmarks

In a recent performance evaluation, Cresta's person name entity recognition performance was compared against a customer's own solution. As noted here, Cresta was able to detect >2X the volume of entities with only a small fraction of undetected entities.

** False Redactions by Cresta occurred in <5% of the entities*

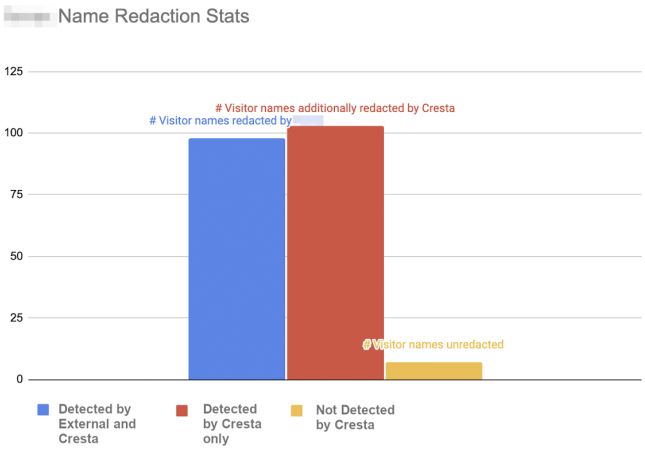

Network Architecture

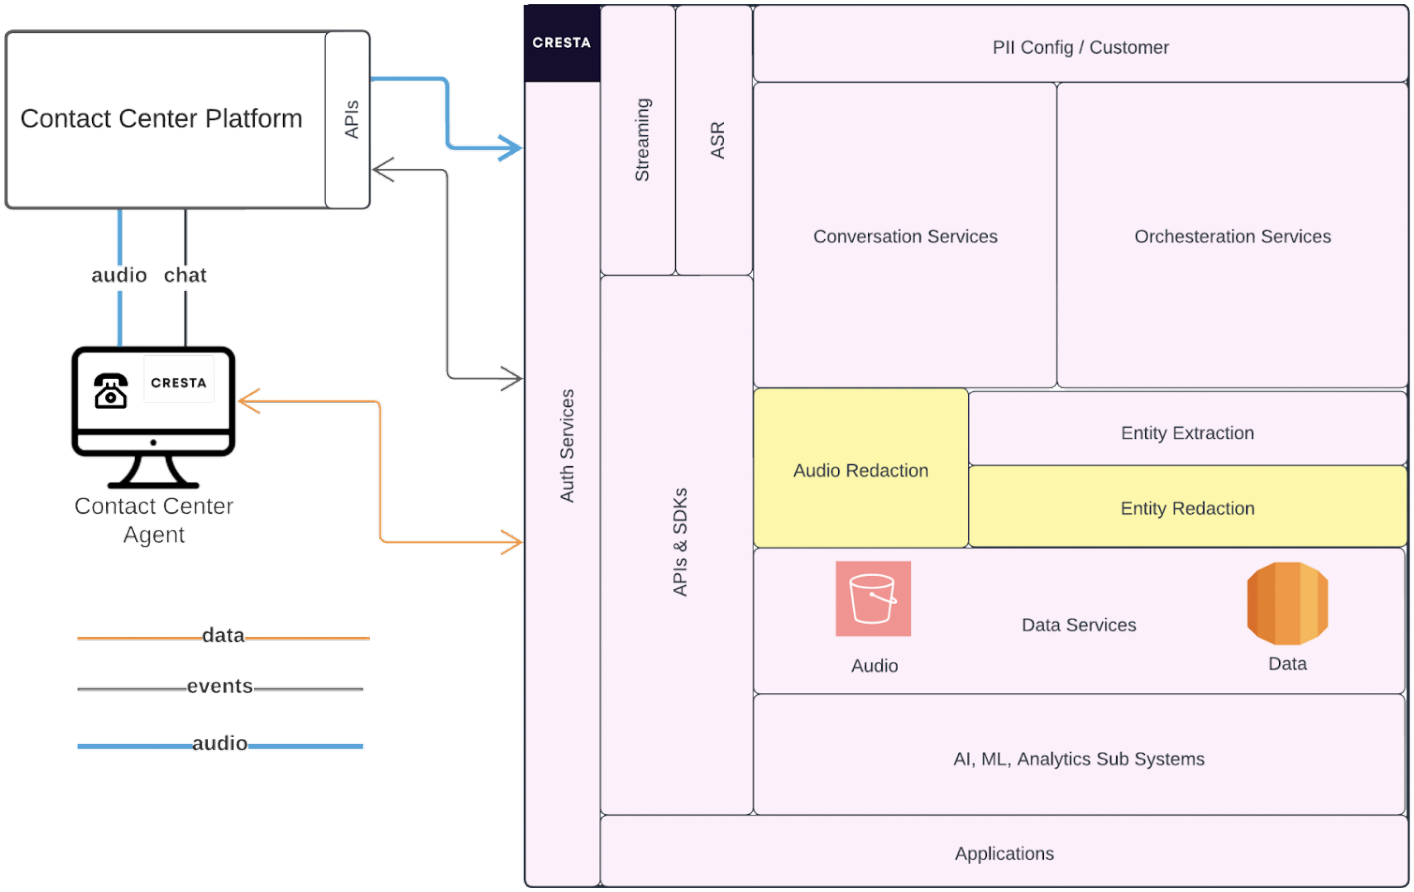

**The architecture diagram shown above illustrates the PII redaction pipeline for Cresta Chat and Voice.*

FAQ

Q: Does Cresta uses a third party service for PII redaction or is it built internally?

A: Cresta has built its own internal PII redaction service.

Q: Does PII get redacted prior to being saved in the database?

A: Yes, all support PII entities and any requested custom PII entities will be redacted prior to being saved to AWS RDS.

Additionally, the audio recording is also redacted prior to being saved to AWS S3.

Q: Is Cresta compliant with PCI-DSS?

A: Yes, Cresta is annually audited by an accredited third party Qualified Security Assessor (QSA) to maintain compliance with PCI-DSS Service Provider Level 2 requirements. Cresta's most recent PCI-DSS Attestation of Compliance (AoC) can be had upon request to security@cresta.ai. More information about Cresta's Compliance Certifications and Frameworks can be found here.

FAQ (Continued)

Q: Can Cresta please provide a list of all supported PII entities that are automatically redacted by default?

A: Please see the list below of supported PII entities that are automatically redacted via Cresta's internal redaction service.

- Names, Full Names
- Phone Numbers
- Email
- Addresses (i.e. Street Address)
- Date of Birth
- Credit Card Numbers (Partial or Complete)
- CVV / Card Security Codes
- U.S. Social Security Numbers (SSN)
- U.S. Driver's License Numbers

Please note that Cresta can turn off/on certain PII entities, and also supports the redaction of custom PII entities. Please reach out to your Customer Success or Sales Representative to turn on/off certain PII entities or request the detection and redaction of custom PII entities.

Q: Does Cresta perform real-time PII redaction?

A: Cresta performs real-time PII redaction for voice calls routed to the Cresta voice agent interface. After a voice call is completed, PII is redacted from the recorded audio and replaced with a "bleep" noise.

However, real-time PII redaction is not supported for chat/messaging since any chats/messages are handled by an external agent platform that Cresta cannot directly control incoming/outgoing messages on. However, supported PII entities are redacted after the chat/message has been closed and is redacted within the Cresta Director interface.

Q: Are supported PII entities redacted post conversation/call if using Cresta's Automatic Note Taking feature?

A: Yes, PII entities captured as part of Cresta's Automatic Note Taking feature are redacted in the Cresta Director interface post conversation/call, as well as in AWS RDS.

Q: Does Cresta support the redaction of electronic protected health information (ePHI)?

A: As of today, Cresta does not support the redaction of ePHI. If this is a feature you'd like to see implemented in the future, please contact your Customer Success Manager or Sales Representative.

Q: Do Cresta personnel have access to view or hear stored cardholder data?

A: Cresta personnel do not have access to view or hear cardholder data as it is redacted prior to it being stored in the database (and forever stored in a redacted manner unless deleted by the customer).

Q: I still have more questions about Cresta's PII detection and redaction pipeline. Who can I reach out to in order to discuss this further?

A: Please forward your questions to your dedicated Account or Customer Success Manager. Otherwise, please feel free to also route your questions to security@cresta.ai.