

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 5 November 2026

D. James
Marashlian & Donahue, PLLC
T. McCarthy-Howe
Strolid
4 May 2026

Privacy Primer for vCon Developers
draft-ietf-vcon-privacy-primer-01

Abstract

This document serves as a primer for technical professionals involved in the processing (which includes collecting, using, disclosure, and erasure) of personal data, including not only basic identifiers like name, age, and address, but also sensitive data contained in communications, including biometrics obtained from audio and video recordings. It outlines key concepts in data privacy and communications privacy, addressing the ethical and legal considerations surrounding the collection, processing, sharing, access, retention, and disclosure of individuals' data. The document covers fundamental privacy principles, defines important roles in data processing, and explains individuals' rights regarding their personal information. It also discusses the scope of protected personal information, including sensitive data categories, and explores techniques like data aggregation and anonymization. While referencing existing IETF work on privacy in Internet communications, this draft extends beyond to address individuals' data privacy in relation to organizations handling such data. The goal is to provide a comprehensive overview of privacy considerations, aligning with fair information practices and current regulatory frameworks, to guide professionals in implementing privacy-respecting data management practices.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://ietf-wg-vcon.github.io/draft-ietf-vcon-privacy-primer/draft-ietf-vcon-privacy-primer.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-vcon-privacy-primer/>.

Discussion of this document takes place on the Virtualized Conversations Working Group mailing list (<mailto:vcon@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/vcon/>. Subscribe at <https://www.ietf.org/mailman/listinfo/vcon/>.

Source for this draft and an issue tracker can be found at
<https://github.com/ietf-wg-vcon/draft-ietf-vcon-privacy-primer>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 5 November 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Purpose of this Document	3
1.2. Intended Audience	4
1.3. Goals of this Document	4
2. Conventions and Definitions	5
2.1. Privacy and vCon In General	5
2.2. Data privacy	5
2.2.1. Key Roles in Data Processing	6
2.2.2. What Data Rights do Data Subjects Have?	7
2.2.3. What Data Is Protected?	8
2.2.4. Sensitive Data	9
2.2.5. What Data Is Not Protected?	10

2.2.6. Deidentification/Anonymization	11
2.2.7. Aggregation/Anonymization	12
2.3. Communications Privacy	13
2.4. Key Privacy Principles	14
2.5. Artificial Intelligence-Specific Considerations	16
3. Security Considerations	17
4. IANA Considerations	19
5. References	19
5.1. Normative References	19
5.2. Informative References	20
Acknowledgments	21
Authors' Addresses	21

1. Introduction

The democratization of technology has led to a surge of new entrants in the growing market of personal data management. These entrants, driven by various motives ranging from commerce and regulation to fraud prevention and charitable causes, are increasingly engaging with conversational data across network boundaries. The vCon (Virtual Conversation) represents a significant step forward in this landscape, enabling the processing and ethical sharing of conversational data and fostering a rich ecosystem of services based on a novel concept: genuinely listening to what customers say.

However, many of these new entrants may not inherently understand the ethical and legal complexities surrounding crucial topics such as data minimization, legal basis for processing, redaction, the right to know, and the right to erasure. The design decisions behind the vCon framework directly address these concerns, incorporating features such as encryption capabilities, external data signing for change detection, and the creation of redacted versions that maintain a trail to the original data.

1.1. Purpose of this Document

This document serves as a primer for individuals and organizations grappling with the challenges of responsible management of personal data, including biometric information contained in audio and video recordings, or other sources of sensitive information, in messaging or other personal communications. It aims to provide a foundational understanding of key topics, explaining their importance and how they are addressed (or not) within the vCon framework. While the vCon is not a panacea, it offers a structure that enables well-intentioned actors to operate ethically and responsibly. Much like the distinction between HTTP and HTTPS, where HTTPS is trusted by default and HTTP is not, the vCon framework provides a basis for trust, with legal systems managing those who operate outside its principles.

IETF standards already address privacy in Internet communications, including the principle of data minimization [RFC7258]. However, those standards generally do not address the privacy of individuals' data privacy vis--vis organizations that collect, process, and disclose their data.

1.2. Intended Audience

This primer is designed to cater to three primary constituencies often present in IETF discussions:

- * Technologists and Engineers: Often immersed in technical details, these professionals may benefit from understanding the broader ethical and legal considerations that should inform their designs. This document aims to bridge the gap between technical implementation and important "non-technical" issues they need to consider.
- * Regulators, Lawyers, and Government Representatives: Responsible for tech policy, these individuals often approach discussions with the perspectives of their constituencies but are generally open to education. This document seeks to provide them with a clearer understanding of how their legal concerns are addressed within the vCon framework and what aspects fall outside its scope.
- * Non-Governmental Organizations (NGOs): Particularly those focused on privacy, security, and human rights, these organizations represent the intersection of policy and technology. Often skeptical of commercial and government interests, this audience will find information on how the vCon supports personal data privacy, transparency, and control.

1.3. Goals of this Document

The primary objectives of this primer are:

- * To educate an expanding audience on the fundamental concepts of responsible customer data management.
- * To foster a common understanding of the challenges involved in personal data handling.
- * To provide an informed perspective on what is currently addressed by the vCon framework and what remains outside its scope.
- * To encourage thoughtful consideration of ethical and legal issues in the design and implementation of systems handling personal data.

By achieving these goals, we aim to contribute to a more informed and responsible approach to personal data management across various sectors and disciplines.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.1. Privacy and vCon In General

Privacy in general can be understood as "the right to be let alone" [Warren1890]. It may be helpful to think of it in four aspects:

1. personal information (or data) privacy,
2. bodily privacy,
3. territorial privacy, and
4. communications privacy.

In the context of vCon, we will concentrate on **data privacy** and **communications privacy.**

2.2. Data privacy

Data privacy, also known as information privacy or data protection, refers to the practice of safeguarding individuals' personal information from unauthorized access, use, disclosure, alteration, or destruction. It involves ensuring that individuals have control over their own personal data and that organizations that collect, store, and process personal data do so in a manner that respects individuals' privacy rights.

Many countries and regions have enacted legislation to protect individuals' personal data, but their specific rules vary by jurisdiction. Examples of data privacy laws include the General Data Protection Regulation (GDPR) in the European Union, the Health Insurance Portability and Accountability Act (HIPAA) and California Consumer Privacy Act (CCPA) in the United States, and the Personal Data Protection Act (PDPA) in Singapore. The U.S. data privacy legal landscape is a particularly complex patchwork of federal and state laws. It includes comprehensive state-level data privacy acts, industry-specific federal laws, and intricate pre-emption relationships between them, creating a fragmented and multifaceted regulatory environment.

This document outlines common privacy rights and obligations as of the time of this writing in alignment with fair information practices (FIP), which are widely recognized principles but may or may not be legally required, depending on the jurisdiction. The framework presented here offers a general understanding of data privacy principles but does not guarantee legal compliance in any specific region. Readers are encouraged to seek legal or technical advice for their particular jurisdiction, industry(-ies), and situations.

2.2.1. Key Roles in Data Processing

The following terms are used by the GDPR and in the privacy industry in general to define the three key roles in data processing:

1. ***Data Subject***: The individual whose personal information is being processed (also referred to as "individual" in this RFC and "consumer" in some data privacy laws).
2. ***Data Controller***: An organization or individual with decision-making authority over data processing who has the legal basis to process data and determines the purposes and methods of data processing, bears primary responsibility under privacy laws and is the main target of most privacy and data protection regulations.
3. ***Data Processor***: Often a third-party service provider who processes data on behalf of the data controller. Under HIPAA, data processors are referred to as "business associates." Data processors may be hired for specialized tasks or to improve efficiency; can subcontract to other processors, creating a chain of responsibility; must operate within the scope defined by the data controller; and are expected to maintain trust and adhere to the controller's guidelines.

The relationship between these entities forms a hierarchy of responsibility and trust. The data controller sets the parameters for data use, while processors at various levels must operate within these boundaries. This structure ensures accountability and helps maintain data privacy throughout the information processing chain.

2.2.2. What Data Rights do Data Subjects Have?

Regarding individual rights in data privacy, organizations should focus on six key areas:

1. ***Notice of Data Processing***: Organizations must clearly communicate their privacy policies and practices. This includes explaining what personal data is collected and for which purposes, how it is used, stored, and shared, and how individuals can exercise their data privacy rights. This ensures transparency, holds data controllers accountable, and empowers individuals to make informed choices about their personal information.
2. ***Consent***: Organizations need to obtain data subjects' informed and freely given consent for collecting, using, storing, and sharing personal data. Different levels of consent may apply to different kinds of data or in different jurisdictions:
 - * Consent can be affirmative ("opt-in" consent) or presumed unless stated otherwise ("opt-out" consent). Opt-in consent is usually required for sensitive data and children's data (under 13 or 16 years old).
 - * Consent can be written, oral, or implied.
 - * In some jurisdictions, such as California, consent must be sought at or before the point of data collection. It should also be noted that consent can be layered (i.e. provided for one function but not others) and revoked at any time by the data subject, and data controllers/processors need to act on this expediently. Consent Management is therefore a key requirement when handling personal data.
3. ***Access***: Organizations should offer mechanisms for individuals to access and correct their personal data. This empowers people to ensure their data is accurate and up-to-date.
4. ***Data Choices***: In addition to rights to access and correct, data subjects often have the following data privacy rights:

- * right to have their information deleted (also referred to as the "right to be forgotten");
 - * right to download their data in a readily readable format and provide their data to a different data controller;
 - * right to opt out of certain data practices, such as sale of their data, profiling, targeted/cross contextual behavioral advertising, automated decision-making.
5. ***Non-Discrimination*:** Organizations must not discriminate against individuals who choose to exercise their data privacy rights.
6. ***Breach Notification*:** The large amounts of data held by organizations attract cyber criminals, increases the risk of data breaches. To mitigate the consequences of data breaches and incentivize advanced data security practices, most jurisdictions require reasonable security safeguards and prompt notification of affected individuals when breaches occur. Many regulatory bodies require notification from a data controller within 30 days of discovery. Controllers are also required to take prompt incident response measures to mitigate breaches and promptly inform and assist data subjects with breach mitigation. This holds organizations accountable for data security and allows individuals to take protective actions.

By addressing these areas, organizations can respect individual privacy rights and build trust with their customers or users. This approach aligns with many modern data protection regulations and best practices in privacy management.

2.2.3. What Data Is Protected?

Data privacy laws protect personal information, though its scope can vary across different laws. In general, the term "personal information" (also known as "personally identifiable information" or "PII") includes information that makes it possible to identify an individual or information about an "identified" or "identifiable" individual. Privacy laws may further extend the scope of PII: for example, the California Privacy Rights Act's definition of PII includes information about an individual and the individual's household, as well as employment data.

In general, examples of PII include:

- * Basic identifiers: Name, addresses (postal and email), government-issued identification numbers

- * Digital identifiers: IP address (in some jurisdictions like California).
- * Financial Data: financial account number or credit or debit card number, often in combination with any required security code or password that would permit access to a data subject's financial account.
- * Health data, including mental health or substance abuse information. Many healthcare identifiers are similar to other types of personal data (like names and addresses), but others may refer to specialized information like health insurance details and medical codes.
- * Characteristics protected under various civil rights and non-discrimination laws: Race, religion, disability, sexual orientation, national origin, etc.
- * Consumer behavior: Purchase history, product interests, consumption patterns.
- * Biometric data, including voiceprints, faceprints, fingerprints.
- * Online activity: Browsing and search history, website and app interaction.
- * Geolocation information
- * Sensory information: Audio, visual, thermal, and olfactory data.
- * Professional and educational information.

2.2.4. Sensitive Data

An important subset of PII to consider in designing data privacy practices is so-called "sensitive data" which is subject to a higher standard of protection and requires additional privacy and security limitations to safeguard its collection, use, and disclosure. For example, it may require an opt-in consent for data collection and processing. In various jurisdictions sensitive information may include the following:

- * Government-issued identifiers.
- * Physical or mental health information.
- * Genetic data.

- * Financial data.
- * Biometric information.
- * Precise geolocation data (information on an individual's location within a 1,850-foot radius).
- * Log-in credentials a customer's account log-in, password, or credentials allowing access to an account.
- * Citizenship and immigration status.
- * Sexual behavior.
- * Information revealing an individual's online activities over time and across websites or online services that do not share common branding or over time on any website or online service operated by a covered high-impact social media company.
- * Information about minors (17 years of age or younger, depending on the applicable law).
- * Certain communications data an individual's private communications, such as voicemails, emails, texts, direct messages or mail, or information identifying the parties to such communications, information contained in telephone bills, voice communications, and any information that pertains to the transmission of voice communications, including numbers called, numbers from which calls were placed, the time calls were made, call duration and location information of the parties to the call, unless the covered entity is an intended recipient of the communication. Communications with businesses may be awarded less protection.

2.2.5. What Data Is Not Protected?

The distinction between personal and nonpersonal information hinges on identifiability, meaning that personal data is identifiable and thus protected by most privacy laws when it can be reasonably linked to a particular person (or even computer or device). This boundary has varying interpretations across jurisdictions. For instance, IP addresses are considered personal information by the EU and FTC, but not by U.S. federal agencies under the Privacy Act. Moreover, sometimes, deidentified data can be reidentified, introducing challenges to personal data protection. When identifying elements are removed from data, it becomes nonpersonal information, generally falling outside the scope of privacy and data protection laws.

Some methods of transforming identifiable data into nonpersonal data are deidentification, anonymization, and aggregation. On the other hand, pseudonymization (replacing identifiable data with pseudonyms/unique codes) only temporarily removes identifiable data with the possibility of relatively easy reidentification. In many jurisdictions, pseudonymous data falls within the scope of personal data when used in conjunction with additional information that reasonably links the data to an identified or identifiable individual.

Many jurisdictions remove *publicly available information* from the scope of protected PI.

Pseudonymized data, where individuals are represented by unique codes, is temporarily nonpersonal but can be reversed to reidentify individuals. This reversibility can be crucial in scenarios like medical trials. In many jurisdictions pseudonymous data falls within the scope of personal data when used in conjunction with additional information that reasonably links the data to an identified or identifiable individual.

Many jurisdictions remove *publicly available information* from the scope of protected PI.

2.2.6. Deidentification/Anonymization

Deidentification is the process of removing identifiable data from the dataset/document. It may take the form of information suppression (direct removal of identifying information), generalization (replacing a data element with a more general equivalent), or noise addition (slightly altering select data). Noise addition is the primary mechanism of differential privacy a mathematical framework designed to ensure the privacy of individuals within a dataset while allowing for the extraction of useful statistical information by adding carefully calibrated noise to the data. This ensures that the inclusion or exclusion of any single individual's data does not significantly affect the outcome of the analysis and adds a level of personal data protection. Sometimes it is possible to reidentify the deidentified data using other available information. In this context, although it is often used interchangeably with the term "deidentification," the term "anonymization" refers to the more comprehensive irreversible removal of identifiable data. Rigorous deidentification or anonymization techniques are highly recommended to ensure that reidentification is either impossible or extremely difficult.

2.2.7. Aggregation/Anonymization

Data aggregation and anonymization are important techniques used in the context of data privacy to protect individuals' personal information while still allowing organizations to derive valuable insights. However, these methods are not without risks and limitations.

Data aggregation involves combining data from multiple sources or individuals into summary form. While this can obscure individual identities, there are still privacy concerns:

1. ***Re-identification risk*:** With enough granular data points, it may be possible to single out individuals even from aggregated datasets. Applying multiple specific filters to aggregate data could potentially identify a unique individual.
2. ***Inference attacks*:** Aggregated data can reveal patterns that allow inferences about individuals or small groups, even if direct identifiers are removed.
3. ***Unintended data exposure*:** As aggregated data is often shared between organizations, there's increased risk of unauthorized access or misuse.

Anonymization aims to remove or encrypt personal identifiers from datasets. However, true anonymization is challenging:

1. ***De-identification limitations*:** Simply removing obvious identifiers like names and addresses is often not sufficient to prevent re-identification. Indirect identifiers can still allow individuals to be singled out.
2. ***Data utility trade-offs*:** More thorough anonymization techniques tend to reduce the usefulness of the data for analysis.
3. ***Evolving re-identification techniques*:** As technology advances, previously anonymized data may become vulnerable to new re-identification methods.

To mitigate these risks, organizations should consider:

1. Implementing robust anonymization techniques beyond basic de-identification.
2. Carefully assessing the granularity and specificity of aggregated data released.

3. Combining anonymization with other privacy-enhancing technologies like differential privacy.
4. Conducting regular privacy impact assessments to evaluate potential risks.
5. Adhering to relevant privacy regulations and best practices for data handling.

While data aggregation and anonymization can enhance privacy protection, they should not be viewed as foolproof solutions. Organizations must remain vigilant and adopt a comprehensive approach to data privacy that considers the evolving nature of re-identification risks and the potential for unintended consequences when working with large datasets.

2.3. Communications Privacy

Communications privacy is a critical concern in our increasingly interconnected world, where various forms of communication including audio, video, text messages, and emails have become integral to both personal and professional interactions. Under the applicable laws, communications are protected when in transit and at rest.

Understanding the multifaceted legal and ethical frameworks around communications privacy is essential for anyone involved in capturing, storing, analyzing, or managing communications data. Communications privacy laws across various jurisdictions often share common elements, designed to protect individuals' privacy rights while balancing the needs of law enforcement and legitimate business interests. Here are some key provisions typically found in laws governing the recording, interception, eavesdropping, and storage of communications:

1. ***Notice:** Many laws require that parties be notified if their communications are being recorded or monitored, often through audible beeps, verbal announcements, or visible signage. The unauthorized surveillance or interception of an individual's private communications or activities is generally prohibited by law.
2. ***Consent:** Most laws stipulate that at least one party must consent to the recording or interception of a communication. Some jurisdictions require all parties to consent, known as "two-party" or "all-party" consent. The type of consent required (explicit or implied) may vary, but it has to be obtained prior to the recording.

3. ***Distinction Between Public and Private Communications:*** Laws often differentiate between communications where there is a reasonable expectation of privacy (e.g., private phone calls) and those in public spaces where such expectation may not exist.
4. ***Purpose Limitations:*** Regulations frequently specify permissible purposes for recording or intercepting communications, such as for security, quality assurance, or with court authorization for law enforcement activities.
5. ***Storage and Retention Limitations:*** Rules governing how long recorded communications can be stored, how they must be protected, and when they should be destroyed are common features of privacy laws.
6. ***Exceptions for Law Enforcement:*** Most laws include provisions allowing for authorized interception of communications by law enforcement agencies, typically requiring judicial oversight through warrants or court orders.
7. ***Technology-Specific Provisions:*** As technology evolves, laws may include specific provisions for different communication media, such as landlines, mobile phones, emails, instant messaging, video calls, and internet browsing activities.
8. ***Security Measures:*** Requirements for securing stored communications against unauthorized access, including encryption standards and access controls, are increasingly common. Moreover, using encryption may in some cases absolve the data processor from legal liability or at least mitigate it.

Understanding these common provisions is crucial for compliance with communications privacy laws, regardless of the specific jurisdiction. However, it is important to note that the exact implementation and interpretation of these provisions can vary significantly between different legal frameworks.

2.4. Key Privacy Principles

Data privacy and communications privacy are guided by similar principles, emphasizing consent, transparency, and data minimization while balancing privacy rights with societal interests. These areas aim to safeguard individuals' control over their personal information, whether stored or transmitted.

Key principles include:

1. ***Consent***: Organizations must usually seek individuals' consent for collecting, processing, and using their sensitive data (as defined under applicable laws) or recording private communications. When required, consent must be freely given, specific, informed, unambiguous, revocable, and documented. Consent may not be valid in situations with power imbalances and may not be required when PII processing is necessary to satisfy legal obligations or implement contracts. Many jurisdictions prohibit so-called "dark patterns," which are practices of seeking consent that effectively obscure, subvert, or impair the individuals' autonomy, decision-making or choice (for example, confusing user interfaces or hidden disclosures).
2. ***Notice/Transparency***: Organizations must clearly disclose their data handling practices. Privacy notices should be concise, transparent, and easily understandable. Changes to privacy practices must be promptly communicated.
3. ***Purpose Limitation***: Personal data should be collected for specific, explicit, and legitimate purposes, and it should not be used for purposes that are incompatible with those for which it was originally collected.
4. ***Data Minimization/Collection Limitation***: Organizations should collect only the minimum amount of personal data necessary to achieve the stated purpose. Excessive or irrelevant data should not be collected.
5. ***Storage Limitation***: Personal data should be retained only for as long as necessary to fulfill the purposes for which it was collected, consistent with legal limitations and requirements. Organizations should establish retention policies and securely dispose of data that is no longer needed.
6. ***Security***: Appropriate technical, physical and administrative measures must be implemented to protect covered data from unauthorized access and other risks. This may include encryption, access controls, and regular security assessments.
7. ***Individual Rights***: Individuals have certain rights regarding their personal data, including the right to access their data, the right to request corrections or deletions ("the right to be forgotten"), the right to object to certain uses of their data, and the right to data portability (the ability to transfer their data from one organization to another).

8. ***Data Integrity***: Personal data should be accurate, complete, up-to-date, and trustworthy throughout its lifecycle. The core principles of data integrity include consistency across systems, authenticity verification, and non-repudiation mechanisms.
9. ***Accountability***: Organizations are responsible for complying with data privacy laws and demonstrating compliance. Organizations are also accountable for any downstream entities with which they may share personal data for a specific defined purpose. Companies must also ensure they monitor and periodically audit third parties with which they share personal data. This may involve conducting privacy impact assessments, appointing a data protection officer, and maintaining records of data processing activities.
10. ***Recordkeeping***: Many laws require organizations to maintain accurate logs of consumers' profiles, data decisions, and data usage, including sales and marketing campaigns and instances of data disclosure to third parties.

While data privacy and communications privacy share many principles, there are some distinctions in their regulation. Communications privacy laws often focus more on real-time interception and communication confidentiality, while data privacy laws address a broader range of data handling practices.

As the digital landscape evolves, privacy laws must continually adapt to address emerging technologies and practices, ensuring the protection of personal information in our interconnected world.

2.5. Artificial Intelligence-Specific Considerations

As vCons are likely to be used in the context of artificial intelligence (AI) applications and services, either directly or tangentially, additional considerations are necessary due to the nascent regulatory environment regarding AI.

One of the most advanced regulations around AI is currently the EU AI Act. In the U.S., the Colorado AI Act contains similar principles. Various organizations, such as the National Institute of Standards and Technology (NIST), have also adopted guidance related to mitigating risks associated with AI use.

The EU AI Act provides a comprehensive legal framework for the development, marketing, and use of AI in the EU to promote human-centric and trustworthy uses of AI while ensuring a high level of protection of health, safety and fundamental rights. The legislation adopts a risk-based approach, distinguishing between tiers of AI use

cases: prohibited (such as behavioral manipulation and sensitive biometrics), high-risk (including critical infrastructure and medical devices), limited risk, and minimal risk, with each tier subject to regulatory requirements that scale from the strictest to the lightest, accordingly.

For example, key requirements for High-Risk AI use cases under the EU AI Act include: - stringent impact and conformity assessments and registration - risk and quality management - human oversight - strong data governance practices to not only mitigate bias but also ensure adequate controls of representative data used in training as well as production applications - transparency through technical documentation and instructions - privacy and data governance

As such, all of the previously discussed privacy and security measures apply to AI use cases under the EU AI Act and are supplemented with additional requirements. In particular, there is an increased emphasis on robust controls over data used for model testing and training, as well as the implementation of processes to ensure effective human oversight, including ongoing monitoring and auditing.

3. Security Considerations

vCons can contain sensitive personal and conversational data, which raises several data privacy and security concerns, particularly regarding data integrity and personal privacy. The following points outline the key security considerations for vCons:

1. Data Integrity and Immutability

- * vCons need to be protected against unauthorized modifications to ensure the authenticity of the conversational data.
- * Before a vCon leaves its original security domain, it should be digitally signed to prevent alteration, as specified in Section 5.2 (Signed Form of vCon Object).

2. Privacy Protection

- * vCons often contain personally identifiable information (PII) and sensitive data that must be safeguarded.
- * Different levels of redaction may be necessary, as outlined in Section 4.1.6 (redacted):
 - PII masking: Removing PII from text, audio, video, and transcripts.

- De-identification: Removing segments or whole recordings to prevent voice printing or facial recognition.

3. Encrypted Storage and Transmission

- * Unredacted versions of vCons must be encrypted to protect sensitive information, as described in Section 5.3 (Encrypted Form of vCon Object).
- * vCons transmitted over non-secure channels (e.g., email) must always be in encrypted form.

4. Access Control

- * Externally referenced files should be transported only over HTTPS, as specified in Section 2.4 (Externally Referenced Files).
- * Access to unredacted vCons and their referenced files should be restricted to authorized personnel only.

5. Version Management

- * Multiple versions of a vCon may exist (e.g., redacted versions, versions with added analysis).
- * Each version must maintain its own integrity while providing a secure reference to its predecessor, as described in Sections 4.1.6 (redacted) and 4.1.7 (appended).

6. Cross-Domain Security

- * vCons may be created and modified across different security domains, as discussed in Section 4 (Unsigned Form of vCon Object).
- * Each domain should sign the vCon before transferring it to maintain the chain of trust, using the method in Section 5.2 (Signed Form of vCon Object).

7. Redaction Processes

- * While methods exist for redacting text, audio, and video, the specific techniques are beyond the scope of the vCon standard, as noted in Section 4.1.6 (redacted).

- * Implementers must ensure that redaction methods effectively remove sensitive information without compromising the vCon's integrity.

8. Balancing Utility and Privacy

- * There is an inherent tension between maintaining the usefulness of vCons and protecting privacy, as implied throughout Section 4 (Unsigned Form of vCon Object).
- * Careful consideration is needed when deciding what information to redact or encrypt.

9. Encryption of Referenced Content

- * Externally referenced files that are part of a vCon should be encrypted if they contain sensitive information, as suggested in Section 2.4 (Externally Referenced Files).

10. Audit Trail

- * vCons should maintain a secure audit trail of modifications, especially for redactions and additions, to ensure accountability. This is supported by the structure described in Sections 4.1.6 (redacted) and 4.1.7 (appended).

By addressing these security concerns and following the guidelines in the vCon standard, implementers can help ensure that vCons protect the privacy of individuals involved in conversations while maintaining the integrity and utility of the conversational data.

4. IANA Considerations

This document has no IANA actions.

[Warren1890] Warren, S.D. and Brandeis, L.D., "The Right to Privacy", Harvard Law Review, Vol. 4, No. 5, pp. 193-220, December 1890.

5. References

5.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

5.2. Informative References

- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/rfc/rfc3552>>.
- [RFC6235] Boschi, E. and B. Trammell, "IP Flow Anonymization Support", RFC 6235, DOI 10.17487/RFC6235, May 2011, <<https://www.rfc-editor.org/rfc/rfc6235>>.
- [RFC6462] Cooper, A., "Report from the Internet Privacy Workshop", RFC 6462, DOI 10.17487/RFC6462, January 2012, <<https://www.rfc-editor.org/rfc/rfc6462>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/rfc/rfc6973>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/rfc/rfc7011>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/rfc/rfc7258>>.
- [RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <<https://www.rfc-editor.org/rfc/rfc7624>>.
- [RFC7844] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity Profiles for DHCP Clients", RFC 7844, DOI 10.17487/RFC7844, May 2016, <<https://www.rfc-editor.org/rfc/rfc7844>>.

- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/rfc/rfc7858>>.
- [RFC8165] Hardie, T., "Design Considerations for Metadata Insertion", RFC 8165, DOI 10.17487/RFC8165, May 2017, <<https://www.rfc-editor.org/rfc/rfc8165>>.
- [RFC8280] ten Oever, N. and C. Cath, "Research into Human Rights Protocol Considerations", RFC 8280, DOI 10.17487/RFC8280, October 2017, <<https://www.rfc-editor.org/rfc/rfc8280>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.

Acknowledgments

- * Thank you to Andy Newton for his review, suggestions and improvements

Authors' Addresses

Diana James
Marashlian & Donahue, PLLC
Email: daj@commllawgroup.com

Thomas McCarthy-Howe
Strolid
Email: thomas.howe@strolid.com