AI Preferences
Internet-Draft
Intended status: Standards Track
Expires: 7 August 2026

A. Bisht
Cisco Meraki

3 February 2026

AI Preferences for Real-Time Protocol Bindings
draft-altanai-aipref-realtime-protocol-bindings-00

Abstract

This document defines how Artificial Intelligence (AI) preference expressions are bound to signaling and media protocols used for real-time, session-based communications such as the Session Initiation Protocol (SIP) and associated Session Description Protocol (SDP) offers. It specifies a reusable binding model, concrete SIP header field conventions, and SDP attributes that allow endpoints, intermediary services, and AI assistants to advertise, negotiate, and enforce requirements about AI-driven processing of session metadata, media control events, and telemetry. The goal is to align real-time protocol behavior with the AI Preferences (AIPREF) vocabulary without disrupting existing call control semantics.

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://altanai.github.io/aipref-realtime-protocol-bindings/draft-altanai-aipref-realtime-protocol-bindings.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-altanai-aipref-realtime-protocol-bindings/.

Discussion of this document takes place on the AI Preferences Working Group mailing list (mailto:ai-control@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/ai-control/. Subscribe at https://www.ietf.org/mailman/listinfo/ai-control/.

Source for this draft and an issue tracker can be found at https://github.com/altanai/aipref-realtime-protocol-bindings.

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 7 August 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. Goals 3
  1.2. Scope 4
  1.3. Requirements Language 4
2. Terminology 5
3. Binding Requirements 5
4. Binding Model 6
5. SIP Signaling Binding 6
  5.1. AI-Pref Header Field 6
    5.1.1. Usage Rules 7
    5.1.2. Compact Tokens and URIs 7
    5.1.3. Error Handling 7
  5.2. SIP Body Considerations 8
6. SDP and Media Binding 8
  6.1. a=aipref Attribute 8
  6.2. RTP Control and Telemetry 8
7. Preference Discovery and Synchronization 9
  7.1. Retrieval via HTTPS 9
  7.2. Repository Interaction 9
  7.3. Conflict Resolution 9

8. Operational Considerations 9
9. Security Considerations 10
9.1. Privacy and End-User Impact Considerations 10
9.1.1. Impact on User Autonomy and Consent 10
9.1.2. Avoiding Over-Restriction of Beneficial Uses 11
9.1.3. Transparency to Participants 11
9.1.4. Intermediary Handling and Privacy Leakage 12
9.1.5. Compatibility with Archiving and Research 12
9.1.6. Avoiding Platform-Level Overreach 13
9.1.7. Interoperability and Open Ecosystem Considerations . 13
10. IANA Considerations 13
10.1. SIP Header Field Registration 13
10.2. SDP Attribute Registration 14
Acknowledgments 14
References 14
Normative References 14
Informative References 15
Author's Address 15

1. Introduction

Real-time communications (RTC) deployments are increasingly assisted by AI-driven components that evaluate signaling metadata, media control messages, and quality of experience (QoE) measurements. Examples include AI-based call routing, automated troubleshooting, adaptive encoding, or compliance monitoring. When these components operate on user or service provider data, the AI Preferences (AIPREF) working group requires that stakeholders can express and enforce preferences that describe what AI systems may inspect, retain, or export.

Existing AIPREF documents define preference vocabularies and repository behavior, but do not specify how those preferences are conveyed through session control protocols. This document fills that gap for SIP-based systems and applies the same pattern to other RTC bindings that reuse SIP constructs (including SIP events, PUBLISH, and WebRTC gateways). The binding guidance is protocol-agnostic where possible so that additional RTC protocols (such as XMPP Jingle or proprietary session controllers) can follow the same pattern.

1.1. Goals

The binding framework MUST:

1. Preserve backwards compatibility with SIP user agents, gateways, and middleboxes that are unaware of AI preference signaling.

2. Permit endpoints and administrative domains to advertise locally enforced AI policies and to consume peer policies before AI processing begins.

3. Support mid-dialog updates so that AI processing can adapt when session context changes (e.g., escalation from audio to video, invoking transcription services, or triggering automated remediation workflows).

4. Allow binding of AI preferences to both signaling-layer artifacts (message bodies, header fields, event packages) and media-layer descriptions (SDP attributes, record routes, and telemetry streams).

1.2. Scope

This document describes:

* A binding model that maps AIPREF vocabulary elements to SIP dialog state and SDP descriptions.

* A compact token and URI-based encoding carried in SIP header fields and bodies.

* Procedures for preference discovery, negotiation, and error handling across dialogs, subscriptions, and conferencing primitives.

* Operational recommendations for intermediaries, policy servers, and AI enforcement points.

This document does not standardize AI algorithms, privacy-preserving enforcement techniques, or the semantics of individual AIPREF vocabularies beyond referencing existing definitions in other working group documents such as [I-D.aipref-network-privacy-control].

1.3. Requirements Language

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. Terminology

The terminology defined in RFC3261, RFC3264, and the AIPREF framework documents applies. This document additionally uses the following conventions:

* *AI Preference Token (APT)*: A compact, possibly signed representation of one or more AIPREF statements, retrievable via URI or carried inline.

* *Preference Attacher*: The SIP entity that injects an APT into signaling or SDP (e.g., originating user agent, outbound proxy, or application server).

* *Preference Enforcement Point (PEP)*: A network element or AI component that validates and enforces APT requirements prior to processing protected data.

* *Real-Time Preference Channel*: Any mechanism that conveys updated APTs after a dialog is established (e.g., re-INVITE, UPDATE, INFO, SUBSCRIBE/NOTIFY pair).

3. Binding Requirements

The following requirements guide bindings for SIP and related RTC protocols:

* *Visibility*: APTs SHOULD be visible to the entities that are expected to enforce them, including downstream AI assistants. When hop-by-hop protection (e.g., TLS) is applied, intermediaries outside the trust domain MUST NOT rely on APTs they cannot validate.

* *Integrity*: APTs SHOULD be signed or integrity-protected when transported across untrusted proxies to prevent unauthorized downgrades.

* *Idempotency*: Repeated delivery of identical APTs MUST NOT cause state divergence. Implementations MAY treat the most recent valid APT as authoritative.

* *Granularity*: Bindings MUST allow preferences scoped to dialogs, media sections, or individual features (e.g., telemetry streams, AI feature lists). APTs therefore include target metadata identifying the scope (dialog-id, media mid, or subscription identifier).

* *Fallback*: Endpoints that cannot satisfy received preferences MUST reject or redirect the request using standard SIP procedures (e.g., 488 Not Acceptable Here) and SHOULD include diagnostic information.

4. Binding Model

The model in Figure 1 illustrates how preferences flow through RTC signaling.

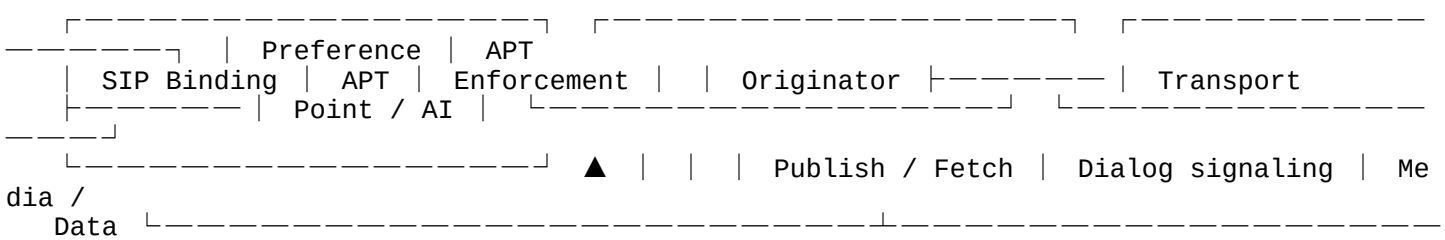

1. An originator (user agent, operator policy engine, or regulator) issues an APT identifier or token.

2. The preference attacher embeds the token using one or more bindings defined in this document.

3. Enforcement points inside AI workloads retrieve, validate, and apply the referenced constraints before consuming protected data.

Bindings MUST reference the same canonical APT identifier when expressing preferences across signaling and media layers to avoid ambiguity.

5. SIP Signaling Binding

5.1. AI-Pref Header Field

This document defines the AI-Pref SIP header field. Each instance conveys metadata that allows receivers to retrieve or validate an APT.

AI-Pref = "AI-Pref" HCOLON pref-value *(COMMA pref-value) pref-value = pref-id *(SEMI pref-param) pref-id = token / quoted-string ; identifier or opaque token pref-param = ("scope" EQUAL token) / ("type" EQUAL token) / ("version" EQUAL 1*DIGIT) / ("integrity" EQUAL token) / ("uri" EQUAL LAQUOT absoluteURI RAQUOT)

5.1.1. Usage Rules

1. *Initial INVITE*: The originating user agent server (UAC) SHOULD include an AI-Pref header referencing all preferences that govern AI handling of dialog metadata or media diagnostics. The header MAY appear multiple times when different scopes are advertised (e.g., scope=dialog, scope=media).

2. *Provisional Responses*: Proxies and user agent servers (UAS) MAY add AI-Pref headers to responses in order to enumerate additional policies that downstream AI components MUST accept before the session is confirmed.

3. *ACK and PRACK*: These messages MUST echo the latest accepted AI-Pref identifiers when the UAS requires confirmation. Absence of AI-Pref in ACK implies acceptance of the most recent set.

4. *Mid-Dialog Updates*: Re-INVITE and UPDATE requests MUST include AI-Pref whenever the preference scope of any media stream is modified. This allows AI systems that adapt encoders, perform transcription, or modify layouts to obey updated constraints.

5. *SUBSCRIBE/NOTIFY*: Event packages (e.g., dialog package, KPML, presence) MAY carry AI-Pref headers so that AI assistants consuming event payloads adhere to the advertised constraints.

5.1.2. Compact Tokens and URIs

pref-id values can represent:

* Inline tokens that encode the full preference statement (e.g., CBOR Web Token referencing vocabulary keys).

* Handles that require dereferencing via HTTPS using the uri parameter.

* Versioned identifiers that map to entries in a policy repository.

Receivers MUST treat unknown parameters according to RFC3261 rules (ignore them) and MUST NOT assume that the absence of AI-Pref implies permission to process data without AI constraints.

5.1.3. Error Handling

* If a UAS cannot comply with mandatory preferences, it SHOULD reply with 488 Not Acceptable Here and include a Warning header value of 399 aipref "Preference unsupported".

* When integrity verification fails, the recipient SHOULD respond with 403 Forbidden and MAY include diagnostic details in a Reason header referencing AI-Integrity-Failure.

* Gateways that strip AI-Pref MUST insert a History-Info entry explaining the removal so downstream entities understand why enforcement data is missing.

5.2. SIP Body Considerations

When APTs are too large for header fields, this document RECOMMENDS embedding them inside a multipart body part with media type application/aipref+json or application/aipref+cbor and referencing that part via Content-ID. This allows richer metadata (e.g., signed manifests) without overloading SIP header processing.

6. SDP and Media Binding

6.1. a=aipref Attribute

An SDP media description MAY include one or more a=aipref attributes, each binding a specific media stream to an APT identifier.

a=aipref:<scope> <identifier> [<parameter>=<value> ...]

Valid scopes include session, group, and mid. Parameters align with the AI Pref vocabulary, for example:

* features=media-metrics,rtcp-xr

* retention=24h

* export=aggregated-only

Endpoints MUST ensure that SDP attributes remain consistent with SIP-level AI-Pref headers. If SDP renegotiation removes an attribute, the corresponding AI processing MUST stop or transition to the remaining allowed scope.

6.2. RTP Control and Telemetry

RTP control protocols such as RTCP, RTCP XR, and RTP/RTCP extensions for feedback may carry AI-relevant telemetry. Implementations SHOULD map telemetry streams to the same APT identifiers declared in SDP by including the identifier in RTCP SDES items or header extensions defined for this purpose. When that is not feasible, implementations MAY rely on the dialog-level AI-Pref scope while documenting the implicit association.

7. Preference Discovery and Synchronization

7.1. Retrieval via HTTPS

Preferences referenced via uri MUST be retrievable over HTTPS with mutual authentication when sensitive. Servers SHOULD support conditional requests and caching so intermediaries can reuse validated APTs across multiple dialogs.

7.2. Repository Interaction

Policy repositories MAY expose a REST interface where clients submit dialog metadata (Call-ID, From-tag, To-tag) and receive the authoritative list of applicable APT identifiers. This pattern is especially useful for large conferencing services where centralized policy engines coordinate AI workloads.

7.3. Conflict Resolution

When multiple APTs apply to the same resource, the following precedence rules apply unless a policy repository states otherwise:

1. User-specific preferences override domain defaults.
2. Domain-level regulatory requirements override individual relaxations.
3. The most restrictive constraint wins when two preferences conflict on the same vocabulary key.

Endpoints MAY advertise their conflict-resolution policy through the policy parameter inside AI-Pref headers (e.g., policy=strictest-wins).

8. Operational Considerations

* *Logging*: Preference attachers SHOULD log emitted APT identifiers alongside Call-ID values to support audits and incident response.
* *Testing*: Interoperability testing MUST verify that dialogs proceed when AI-Pref is absent, ensuring that legacy devices remain compatible.
* *Scaling*: Implementations SHOULD compress header fields using SIP over HTTP/3 (RFC9397) or similar transports when large preference sets are common.

* *Federation*: Peering domains MAY translate local preference identifiers into a shared namespace. Translation MUST NOT weaken constraints without explicit consent from the originator.

9. Security Considerations

Preferences often contain sensitive information about user intent, regulatory exposure, or organizational policy. Therefore:

* Transport security such as TLS (for SIP over TLS, WebSocket, or HTTP/3) MUST be used whenever an AI-Pref header or body carries tokens that could be replayed or tampered with.

* APT signatures SHOULD be validated before AI systems act on the encapsulated instructions. Validation includes issuer authentication, expiration checks, and revocation status.

* Implementations MUST guard against downgrade attacks where a malicious intermediary strips AI-Pref headers. Techniques include SIPS-only routing, end-to-end integrity with S/MIME, or redundant signaling through policy repositories.

* Preference tokens SHOULD minimize personally identifiable information. Instead of embedding explicit user identifiers, use pseudonymous handles that map to access-controlled directories.

* Systems MUST treat AI enforcement failures as security incidents when they result in unauthorized data processing. Telemetry SHOULD be rate-limited to avoid revealing preference patterns to attackers probing the signaling fabric.

9.1. Privacy and End-User Impact Considerations

Real-time and session-oriented communication protocols (e.g., SIP, SDP, WebRTC, RTP/RTCP, QUIC-RTC) directly mediate human-to-human interaction. Introducing AI-usage preference signaling into these protocols therefore has immediate consequences for end users, including creators, participants, accessibility users, researchers, and the general public. This section outlines considerations necessary to ensure that AIPREF signaling in real-time environments does not unintentionally restrict legitimate uses, undermine user autonomy, or create new privacy risks.

9.1.1. Impact on User Autonomy and Consent

AI-usage preferences expressed at the signaling or media layer may be interpreted as binding restrictions by downstream systems. Implementations MUST ensure that:

* Preferences are treated as expressions of intent, not as access-control mechanisms.

* End users retain the ability to override platform-imposed defaults when they are the originators of the content.

* Intermediaries (e.g., conferencing platforms, SIP proxies, TURN servers, media mixers) do not silently substitute or modify user-provided preferences.

Because real-time sessions often involve multiple participants, systems SHOULD provide clear and accessible mechanisms for users to understand what AI-related preferences are being signaled on their behalf.

9.1.2. Avoiding Over-Restriction of Beneficial Uses

Real-time communication is frequently used for accessibility (captioning, translation), education, archiving, and research. Overly broad or ambiguous preference categories—particularly those related to “AI Input” or “AI Training” —may unintentionally block beneficial, lawful, or expected uses.

Implementations SHOULD:

* Distinguish between AI-assisted user features (e.g., live transcription) and model-building uses (e.g., training).

* Avoid treating a single preference (e.g., ai-input=n) as a blanket prohibition on all automated processing.

* Provide clear documentation on how preferences interact with accessibility features.

Preference categories defined in AIPREF vocabulary drafts (e.g., search, ai-input, ai-train) SHOULD be interpreted narrowly and consistently.

9.1.3. Transparency to Participants

Real-time sessions may involve dynamic negotiation (e.g., via SDP offer/answer). When AI-usage preferences are conveyed:

* Endpoints SHOULD surface these preferences to human participants in a clear and non-technical manner.

* Systems SHOULD indicate when preferences differ between participants or when intermediaries have applied additional constraints.

* If preferences affect session features (e.g., disabling transcription), participants SHOULD be notified.

Lack of transparency may create a chilling effect, where users avoid lawful or beneficial uses due to uncertainty.

9.1.4. Intermediary Handling and Privacy Leakage

Signaling AI-usage preferences at the session layer may inadvertently reveal information about user intent, content sensitivity, or organizational policy. For example, a preference of ai-train=n may imply that the content is proprietary or confidential.

To mitigate this:

* Intermediaries MUST NOT add, remove, or alter AI-usage preferences unless explicitly authorized by the originating endpoint.

* Preferences SHOULD be encrypted or integrity-protected when carried in protocols that support such mechanisms (e.g., DTLS-SRTP, QUIC).

* Implementations SHOULD avoid exposing preferences in logs, analytics, or telemetry unless necessary and with appropriate safeguards.

9.1.5. Compatibility with Archiving and Research

Real-time communication is often recorded for compliance, education, or archival purposes. AI-usage preferences SHOULD NOT be interpreted as prohibiting:

* lawful archiving,

* time-shifted review,

* accessibility processing, or

* research uses permitted by local law.

Where preferences do apply to stored recordings, systems SHOULD preserve the preferences alongside the stored media, consistent with the behavior defined in draft-ietf-aipref-attach for HTTP representations.

9.1.6. Avoiding Platform-Level Overreach

Large communication platforms may be tempted to apply AI-usage preferences globally on behalf of users. This can undermine user autonomy and distort the intent of AIPREF.

Platforms SHOULD:

* Allow per-participant and per-stream preferences.
* Avoid applying organization-wide defaults without clear user visibility.
* Provide APIs for endpoints to express their own preferences without mediation.

9.1.7. Interoperability and Open Ecosystem Considerations

Real-time protocols are used across diverse environments, including open-source clients, small organizations, and global platforms. To avoid fragmentation:

* Preferences SHOULD be optional and non-blocking.
* Absence of a preference MUST NOT be interpreted as consent.
* Implementations SHOULD follow the vocabulary and semantics defined in AIPREF drafts to ensure consistent interpretation across ecosystems.

10. IANA Considerations

IANA is requested to perform the following actions.

10.1. SIP Header Field Registration

Register the AI-Pref header field in the "Header Fields" sub-registry under "Session Initiation Protocol (SIP) Parameters" with the following values:

* Header Name: AI-Pref
* Compact Form: none
* Reference: This document

10.2. SDP Attribute Registration

Register the aipref attribute in the "att-field (media level only)" registry defined by RFC4566 with the following values:

* Attribute name: aipref
* Type of attribute: media / session
* Subject to charset: no
* Purpose: Associates SDP media sections with AI preference tokens
* Reference: This document

Acknowledgments

This work is informed by discussions within the AIPREF working group, including contributions on network privacy controls and media quality automation.

References

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>.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <https://www.rfc-editor.org/rfc/rfc3261>.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <https://www.rfc-editor.org/rfc/rfc3264>.

[RFC3311] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, DOI 10.17487/RFC3311, October 2002, <https://www.rfc-editor.org/rfc/rfc3311>.

[RFC6665] Roach, A.B., "SIP-Specific Event Notification", RFC 6665, DOI 10.17487/RFC6665, July 2012, <https://www.rfc-editor.org/rfc/rfc6665>.

[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>.

[RFC8830] Alvestrand, H., "WebRTC MediaStream Identification in the Session Description Protocol", RFC 8830, DOI 10.17487/RFC8830, January 2021, <https://www.rfc-editor.org/rfc/rfc8830>.

Informative References

[I-D.aipref-network-privacy-control] "*** BROKEN REFERENCE ***".

[RFC3265] Roach, A. B., "Session Initiation Protocol (SIP)-Specific Event Notification", RFC 3265, DOI 10.17487/RFC3265, June 2002, <https://www.rfc-editor.org/rfc/rfc3265>.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, DOI 10.17487/RFC4566, July 2006, <https://www.rfc-editor.org/rfc/rfc4566>.

[RFC7587] Spittka, J., Vos, K., and JM. Valin, "RTP Payload Format for the Opus Speech and Audio Codec", RFC 7587, DOI 10.17487/RFC7587, June 2015, <https://www.rfc-editor.org/rfc/rfc7587>.

[RFC8831] Jesup, R., Loreto, S., and M. Txen, "WebRTC Data Channels", RFC 8831, DOI 10.17487/RFC8831, January 2021, <https://www.rfc-editor.org/rfc/rfc8831>.

Author’s Address

Altanai Bisht
Cisco Meraki
Email: albisht@cisco.com