

Contact Center: Topic Analysis

AI Transparency Technical Note

March 2026

Contents

Feature Overview _____ 4

Model Overview _____ 4

 Introduction _____ 4

 Model Architecture _____ 5

 Usage Guidelines _____ 5

 Model Inputs and Outputs _____ 5

Data Sources for Training and Evaluation _____ 5

Model Evaluation and Performance _____ 6

Safety and Ethical Considerations _____ 6

Fairness _____ 6

Privacy and Security _____ 7

Updates and Maintenance _____ 7

References _____ 7



At Cisco, we appreciate that Artificial Intelligence (AI) can be leveraged to power an inclusive future for all. We also recognize that by applying this technology, we have a responsibility to mitigate potential harm. That is why Cisco adheres to our [Responsible AI Framework](#) (the “Framework”), which is based on six principles of Transparency, Fairness, Accountability, Privacy, Security and Reliability. Cisco translates these principles into product development requirements, which ultimately form part of the product development lifecycle alongside our Security by Design, Privacy by Design, and Human Rights by Design processes.

Accordingly, Webex by Cisco (Webex) features that leverage AI are built with transparency, fairness, accountability, privacy, security, and reliability at their core. Each feature powered by AI undergoes an AI Impact (AI) Assessment – a best-in-class review of how the technical underpinnings of the functionality measure against the Framework Principles.

Topic Analytics was built with the Framework principles at the center of how Cisco delivers the AI-powered technology. This Technical Note describes more information about the feature and how Cisco responsibly leverages AI to deliver the functionality. Topic Analytics is an optional feature; organization administrators can disable or allow individual users to choose whether and to what extent they use the feature.

Feature Overview

The Topic Analytics feature, enhanced with near real-time capabilities, empowers users to discover emerging topics as customer interactions occur. Topic Analytics provides instant visibility into evolving trends and concerns by automatically labeling each interaction and enabling immediate, actionable insights for faster data driven decisions. This dynamic approach ensures teams can quickly adapt to changing customer needs, optimize operations, and stay ahead of potential issues which move beyond historical batch analysis to provide continuous, automated insights.




Figure 1: Topic Analytics

Model Overview

Introduction

Topic Analytics is currently powered by Cisco AI technologies and by third-party large language models (LLM) hosted via Amazon Bedrock. The underlying LLM functionality utilizes Amazon's shared infrastructure, with certain configuration and enablement by Cisco. Topic Analytics is deployed on the Webex platform through the user interface within Webex Control Hub and is managed and viewed in Webex Contact Center Desktop and Analyzer.

Model Architecture

Topic Analytics utilizes commercially licensed large language models to aggregate and label the topics. This includes:

- GPT-4o mini for call driver extraction, hosted by the Azure OpenAI Service.
- Mistral's [Mistral-7B-Instruct-V0.2](#) open-source language model, hosted on [Amazon Bedrock](#) and accessed via API, for labeling call topics.

Usage Guidelines

Topic Analytics is not available for use outside of Webex Contact Center.

Model Inputs and Outputs

Input for Topic Analytics begins with contact center calls, which are first transcribed into accurate, high-quality text via Speech-to-Text transcription. From these transcripts, a LLM extracts a concise, single sentence contact reason from each call, ensuring precise identification of caller intent. These contact reasons are then transformed into high-dimensional representations for dynamic, high-speed clustering.

Output for Topic Analytics provides a continuously updated list of automatically labeled topics and call drivers, reflecting trends and emerging issues in near real-time. Users can define their own topics, which the system integrates by proactively matching contact reasons using Advanced Natural Language Processing (NLP) and Natural Language Inference (NLI) based matching.

For additional information regarding the call transcription process, reference the Cisco [In-House Transcription](#) and [Third Party Transcription and Real-Time Translation](#) AI Transparency Technical Notes.

Data Sources for Training and Evaluation

Cisco does not use customer content to fine tune or train the models used to deliver Topic Analytics. Amazon represents that it does not use customer data to improve its models or third-party models it provides access to. Amazon also represents that it does not store or retain Cisco customer data in AWS infrastructure.

For European Union (EU) customers, data is processed locally within the EU regions for transcription and topic analytics. For India customers, data is processed locally within India for

transcription and topic analytics. For other global customers, data is processed in the relevant US, EU, and APJC regions.

Model Evaluation and Performance

Cisco frequently assesses the Topic Analytics feature to maintain and improve both performance and accuracy of generated output. Humans are involved with the review, testing, and quality assurance associated with deployment of the underlying model.

Safety and Ethical Considerations

Topic Analytics has been assessed for safety and ethical risks through internal reviews and AI Impact Assessments. Cisco has ensured that the input data will be stripped of Personally Identifiable Information (PII) with the use of [Microsoft's Presidio](#) before processing transcripts for call driver extraction. While the model does not use or access geographical information, the outputs of the model can be used to support the analysis of trends in different geographies – this is by design.

To further support a safe and respectful user experience, Cisco has implemented a toxic hallucination filter that attempts to prevent the model from generating new toxic, harmful, or unsafe language in analytic outputs. However, if toxic, harmful, or unsafe language is present in the original transcript, it may still surface in the transcript and call driver results, as the filter does not alter the source input.

Cisco has also taken additional steps to mitigate safety risks and ensure ethical use, by requiring third-party vendors, including Microsoft, to undergo rigorous vendor reviews, which include security and safety assessments. Cisco also encourages customers to follow best practices for privacy and ethics when interpreting analytic outputs.

Please contact Cisco if you have concerns or feedback about your experience with the Topic Analytics feature.

Fairness

The third-party models leveraged by Topic Analytics, for its outputs, are trained and evaluated for fairness. Please reference [Transparency Note of Azure OpenAI](#), [Amazon Bedrock](#), and [Mistral AI](#) documentation regarding how the third-party models are trained and evaluated for fairness.

Privacy and Security

Cisco addresses processing of and security around personal data, including retention periods for that data, in its Privacy Data Sheets, found on the [Cisco Trust Portal](#). Transcripts are redacted of any Personally Identifiable Information prior to running the model. During inference, customer data is not stored by the model or retained. The model only generates call drivers and topics.

Amazon represents that it does not access, monitor, or store Cisco customer data. Regarding Cisco access to customer personal data, access rights are limited, and encryption is standard as indicated in the relevant Privacy Data Sheet.

Updates and Maintenance

Ongoing Webex updates and product release documents include product updates and any improvements in user experience. Webex may also use customer email notifications and Webex Help Center documentation to inform customers of product changes, as well as communicate product updates in blogs and videos.

References

[RAI Principles](#)

[RAI Framework](#)

[Cisco Trust Portal](#)

[Amazon Bedrock](#)

[Transparency Note for Azure OpenAI](#)

[Mistral AI](#)

[Mistral-7B-Instruct-v0.2 Model Card](#)

[Microsoft Presidio: Data Protection and De-identification SDK](#)

[In-House Transcription AI Transparency Technical Note](#)

[Third Party Transcription and Real-Time Translation AI Transparency Technical Note](#)