

{ Artemis } 1.5.0

Releasing August 22, 2026

Hi Everyone,

Artemis 1.5.0 will expand the platform's capabilities across AI, connectivity, voice, and analytics. A new customer-facing CLI will empower developers to automate the Agentic project lifecycle without Studio. Microsoft Foundry support and the Canonical Agent Definition will strengthen the AI foundation. New connectors will add productivity, ITSM, and cloud storage tools. Voice features will include configurable call control, deeper integrations with Genesys and Amazon Connect, and a new Voice Playground. Analytics will add public reporting APIs, better cost visibility, and custom dashboard filters.

AI, Agents, and Search

- A new customer-facing CLI will empower developers to automate the full Agentic lifecycle—including authentication, project sync, eval runs, and deployment—from external systems without Studio.
- Microsoft Foundry-hosted agents will be supported (proof of concept), laying the groundwork for running Artemis agents natively within Microsoft's Foundry hosting environment.
- The Microsoft Foundry model catalog will be supported as an LLM provider, enabling developers to select Foundry-hosted models alongside existing LLM providers.
- Search AI will support Docling OCR language detection and RapidOCR model selection, improving OCR accuracy when processing multilingual documents.
- Canonical Agent Definition will enable lossless sync between the DSL and Studio.
- Scheduled runs will be supported for Evals, enabling developers to automate recurring evaluation cycles instead of triggering them manually.

Connectors and Integrations

- New connectors will be added for productivity and ITSM tools, including MS 365, Salesforce, ServiceNow, SharePoint, Confluence, and Jira.
- New connectors will be added for cloud storage, including Box, Amazon S3, and Google Drive.

Voice and Telephony

- Call control parameters will be configurable at the agent level. This gives teams finer-grained control over call handling behavior without requiring platform-wide changes.
- Digital Channels will integrate with Genesys Cloud Open Messaging. This allows customers on Genesys Cloud to route messaging traffic directly into Artemis agents.
- Voice Bot will integrate with Genesys Cloud Audio Connector, enabling real-time voice interactions between Artemis Voice Bot and Genesys Cloud's audio infrastructure.

- Connect greeting audio for voice channels will be configurable, allowing developers to customize the initial greeting for each channel without engineering support.
- Amazon Connect will support live-agent transfer via SIP BYE, including a generated conversation summary.
- STT/TTS will be migrated from the legacy Emotech WS API to the v1 API. This removes dependency on a deprecated interface and aligns Artemis with Emotech's current API.
- Native multilingual handling will support per-language TTS/ASR configuration.
- The Voice Playground will be introduced to test and iterate on voice experiences, providing developers a sandbox to refine voice flows before deploying them to production.

Analytics and Insights

- Session summary reporting data will be accessible via an authenticated public API.
- LLM token and cost reporting will support breakdowns by custom session analytics dimensions.
- Custom dashboards will no longer cap results at 1,000 rows.
- Feedback records will include user information.
- Public analytics APIs (excluding recordings) will be available, enabling teams to integrate Artemis analytics data into their own reporting tools.
- Foundational storage for rebuilt analytics will be established. This lays the groundwork for the next generation of analytics features planned for future releases.
- Design-time and runtime LLM logs will be available in Insights. Teams will be able to trace LLM behavior across both build-time and live runtime scenarios in one place.
- Dashboards will support custom filters, enabling users to narrow dashboard views to the specific slice of data relevant to them.

Security and Configuration

- HTTP body-template array expansion (per-element secrets) will be supported for batch payloads, allowing secrets to be safely referenced per element within a batch request rather than only at the top level.
 - Authorization mechanisms will be enhanced.
 - Project Auth Profile configuration will support reading values from environment variables, allowing developers to avoid hardcoding sensitive values, such as client IDs and secrets, directly into profile configuration.
 - Project export will include portability hygiene advisories that flag hardcoded values rather than silently blocking them by name, reducing false positives during migration while still surfacing values that need review.
 - **BUG FIX** Tools and Auth Profiles will support per-environment/deployment scoping, resolving a conflict in which Prod-Stage and Prod-Production callback auth could not work simultaneously.
 - **BUG FIX** Secrets in Artemis configuration will no longer be exposed in plain view, closing a gap where secret values could be read on screen.
-

Platform and Reliability

- Redis resilience, capacity planning, and throttling will be improved, reducing the risk of Redis-related slowdowns or failures as usage scales.
 - A session termination trigger with follow-up automation will be added, allowing automated actions to run immediately after a session ends, without additional manual configuration.
 - ClickHouse database/table setup will be verifiable via apx-doctor, not just connectivity. This catches missing databases, tables, or outdated schemas that a simple connectivity check would miss.
 - **BUG FIX** A regression will be fixed where recent ABL changes caused previously working agents to fail.
 - **BUG FIX** Multi-agent wiring and runtime discrepancies affecting Autodoc will be resolved.
-

The release is planned for August 22, 2026, for Global Cloud in the US region and will be rolled out to other regions over the following 2-3 days. [Subscribe to Kore.ai status updates](#) to receive notifications for upcoming scheduled maintenance and deployments.

Thank you for your continued support and valuable feedback.

Thank you for your continued partnership with Kore.ai. We look forward to helping you leverage these enhancements to optimize customer and employee experiences.

Detailed information will be available on our [documentation portal](#) after the release.

Best regards,
Team Kore.ai

Kore.ai Inc, 7380 West Sand Lake Road,
Suite 390, Orlando,FL,32819,United States,

©2026 Kore.ai, Inc.

Connect with us on



[Unsubscribe or Manage Preferences](#)