

## RELEASE NOTES



# Product Releases

## Release 26.7 Notes

March 31, 2026

### What's New and Improved

#### LLM Gateway: Vertex AI Provider Support

The LLM Gateway now supports Google Vertex AI as a provider, enabling teams to route evaluator and LLM-as-a-Judge requests through GCP using service account credentials. Vertex AI provides access to a wide range of models from the [Model Garden](#):

- **Gemini:** `vertex_ai/gemini-*`
- **Anthropic (Claude):** `vertex_ai/claude-*`
- **Meta/Llama:** `vertex_ai/meta/{MODEL}`
- **Mistral:** `vertex_ai/mistral-*`
- **DeepSeek:** `vertex_ai/deepseek-ai/{MODEL}`
- **Qwen:** `vertex_ai/qwen/*`
- **AI21 (Jamba):** `vertex_ai/jamba-*`
- **OpenAI (GPT-OSS):** `vertex_ai/openai/gpt-oss-*`
- **ZAI (GLM):** `vertex_ai/zai-org/{MODEL}`

- **Self-deployed (OpenAI-compatible):** `vertex_ai/openai/<your-endpoint-id>`
- **Self-deployed (non-OpenAI-compatible):** `vertex_ai/<your-endpoint-id>`

Configure Vertex AI access in Settings > LLM Gateway by providing a GCP service account JSON key, project ID, and location.

## Multi-Target Event Updates for LLM Models

LLM and NOT\_SET models with multiple target columns now correctly preserve all target columns during event updates. Previously, only the first target column was retained and additional targets were silently dropped.

- *Benefit:* You can update multiple target columns (e.g., `comment` and `feedback`) in a single PATCH request without data loss
- *Scope:* Affects only `LLM` and `NOT_SET` model task types, which are the only tasks that support multiple targets. Classification, regression, and ranking models (single target) are unaffected

## Evaluation Dataset Management

Evaluation datasets can now be created, populated, and deleted directly from the UI. You can upload items via CSV with guided column mapping, including auto-mapping based on the existing dataset schema when adding to datasets that already contain data.

- *Benefit:* You can manage evaluation datasets entirely from the UI without writing code; smart auto-mapping reduces friction for repeated uploads
- *Note:* CSV uploads are limited to 1,000 rows per upload

## Fixes and Security Updates

- **LLM Gateway:** Provider delete and update operations now return `409 Conflict` instead of a generic `400` when credentials are referenced by evaluators