

API Changelog

All notable changes to the CliniQueue API are documented here.

Format follows [Keep a Changelog](#).

API versioning follows [Semantic Versioning](#).

[1.1.0] — June 2, 2026

Added

- **audit_id field in /classify response** — Every classification now returns a stable integer ID. Customers use this ID to submit label corrections via `POST /feedback`. Previously there was no way to reference a specific classification after the fact.
- **POST /feedback endpoint** — New endpoint for submitting label corrections. Accepts `audit_id` and `corrected_category`. Corrections are stored in the database and used to improve classification accuracy over time. Returns `{"accepted": true, "audit_id": ..., "corrected_category": ...}`.
- **phi_in_text_month in /analytics/summary** — New field alongside the existing `phi_flags_month`. Counts tickets where actual patient-identifying values (names, DOBs, MRNs) appeared verbatim in the ticket text during the current month. This is a staff hygiene signal distinct from PHI breach incidents.
- **PHI scrubbing pipeline** — Microsoft Presidio now runs on every ticket before it is sent to the AI model. Patient names, dates of birth, MRNs, SSNs, phone numbers, email addresses, and other identifiers are replaced with generic tokens (`[NAME]`, `[DATE]`, `[MRN]`, etc.) before classification. This change is transparent to customers — the classification result is unchanged, but PHI never reaches the model in raw form.
- **Embedding storage** — Scrubbed ticket embeddings are now stored in a pgvector database after every classification. This is a backend change with no customer-facing impact. Embeddings are used to power the shadow classifier and future fine-tuning pipeline.
- **Shadow classifier** — A kNN classifier now runs silently in parallel with the primary model on every request. It logs its predictions internally but never affects the customer-facing response. This is infrastructure for future model improvements.
- **confidence_score in usage logs** — Every classification now logs a proxy confidence score (0.0-1.0) based on uncertainty language in the model's reasoning. Available

in the database for analytics; not currently exposed in the API response.

Changed

- **AI provider migrated from Anthropic Claude to OpenAI GPT-4o** — The underlying model was changed to enable HIPAA BAA coverage. Classification accuracy is equivalent. The response format is identical. No customer integration changes are required.
- **/analytics/summary response** — Added `phi_in_text_month` field. The existing `phi_flags_month` field is unchanged.

Security

- **HIPAA BAA executed with OpenAI** — A Business Associate Agreement with zero data retention provisions is now in place with CliniQueue's AI model provider. Effective June 1, 2026.
- **Zero Data Retention provisioned** — OpenAI's Modified Retention (Zero Data Retention) feature is active on CliniQueue's organization. Ticket text submitted to the API is not logged or retained by OpenAI.

[1.0.0] — May 2026

Initial Release

- **POST /classify** — Core classification endpoint. Accepts `ticket_text`, returns `category`, `priority`, `phi_risk`, `escalate_immediately`, `compliance_flag`, `assigned_team`, `reasoning`, `suggested_reply`.
- **8 classification categories** — `clinical_urgency`, `phi_exposure`, `ehr_integration`, `billing_claims`, `access_login`, `compliance_request`, `regulatory_deadline`, `training_onboarding`.
- **4 priority levels** — `critical`, `high`, `medium`, `low`.
- **GET /analytics/summary** — Monthly usage summary including ticket count, plan limits, average response time, and PHI flag count.
- **GET /analytics/breakdown** — Category distribution for the current month.
- **GET /POST /rules** — Automation rules engine. Define trigger conditions and webhook/email actions that fire after classification.

- `GET /PATCH /draft/settings` — AI draft reply configuration. Enable suggested first responses with optional knowledge base context.
 - `POST /demo/classify` — Public demo endpoint requiring no API key. Powers demo.cliniqueue.com.
 - `GET /health` — Liveness check endpoint.
 - **BAA included at all plan tiers** — Every CliniQueue customer receives a Business Associate Agreement before their first API call.
-

Versioning Policy

CliniQueue follows semantic versioning for the API:

- **Patch (x.x.1)** — Bug fixes, accuracy improvements, performance improvements. No breaking changes. No migration required.
- **Minor (x.1.0)** — New fields added to responses, new endpoints, new optional parameters. Backwards compatible. No migration required.
- **Major (1.0.0)** — Breaking changes to existing endpoints, removed fields, changed field names or types. Migration guide provided. Advance notice of 90 days.

Breaking Change Policy

CliniQueue will never make breaking changes to the API without:

1. 90 days advance notice via email to all active customers
2. A migration guide with exact before/after examples
3. Running the old and new versions in parallel during the transition period

Deprecation Policy

When a field or endpoint is deprecated:

1. It remains functional for a minimum of 180 days after the deprecation announcement
 2. Deprecation warnings are added to the API documentation
 3. Customers are notified by email
-

Backwards Compatibility Guarantee

Fields present in any response will not be removed or renamed in a minor or patch release. New fields may be added to any response at any time — your integration should be built to ignore unknown fields gracefully.

The following are guaranteed stable across all 1.x releases:

- `category` values (the 8 category strings)
- `priority` values (`critical`, `high`, `medium`, `low`)
- `assigned_team` values
- All existing endpoint URLs
- The `X-API-Key` authentication mechanism