# AI Usage Policy

**Ready 4 Everything Training & Supplies Ltd.**
**Certivo Platform — Certivo Intelligence**

**Document ID:** Certivo-AIUP-001
**Version:** 1.6
**Effective Date:** May 2026
**Last Updated:** October 24, 2026
**Owner:** Aaron Hoyte, Privacy Officer / Security Officer
**Approved By:** Aaron Hoyte, Founder
**Review Cycle:** Annual (next review: May 2027), or upon any material change to model providers, model versions, or AI feature scope.

---

## 1. Purpose

This AI Usage Policy ("Policy") describes how Ready 4 Everything Training & Supplies Ltd. o/a Certivo (operating the Certivo platform, "Certivo," "we," "us," or "our") uses artificial intelligence and machine learning ("AI") in the Certivo Platform, what data we send to AI models, where that data is processed, what we do (and do not do) with the outputs, and how Customers and Authorized Users can opt out, audit, or override AI behavior.

This Policy is incorporated by reference into the **Master Subscription Agreement** and the **Privacy Policy**. In the event of a conflict between this Policy and either of those documents, the more privacy-protective provision controls.

---

## 2. Scope

This Policy applies to:

- **All AI-powered features** in the Certivo Platform, collectively branded as **"Certivo Intelligence"** (formerly "Certivo AI" during the May 2026 transition).
- **All Tenants and Authorized Users** of the Platform — Starter, Professional, Scale, and Enterprise — including users who do not actively invoke an AI feature but whose data may be referenced when another user in their organization does.
- **All AI providers** Certivo integrates with at the model-inference layer, currently limited to Google Cloud Vertex AI (see Section 5).
- **All AI outputs**, including text narratives, structured recommendations, drafted emails, categorization suggestions, expense classifications, expiry-renewal opportunities, and any other AI-generated content surfaced through the Platform.

This Policy does **not** apply to:

- Deterministic, non-AI features (e.g., tax calculators, scheduling algorithms, rate limiting). Those are governed by the Master Subscription Agreement and the SLA.
- A Tenant's own use of external AI tools (ChatGPT, Microsoft Copilot, etc.) operating outside the Platform. The Tenant remains responsible for those tools under the Acceptable Use Policy.

---

## 3. What "AI" Means at Certivo

"AI," "artificial intelligence," and "machine learning" in this Policy refer to **generative language models** ("LLMs") served by Google Cloud Vertex AI. Specifically:

- **Narrative tier** — used for narrative reasoning flows where output quality matters (Financial Health Report, Course Profitability narratives, the Ask Certivo chat agent). Runs **Gemini 3.5 Flash** ( `gemini-3.5-flash` ).
- **High-volume tier** — used for high-volume structured-output flows where each item gets a short narrative overlay (recert lead narratives, outreach email drafts, expense category suggestions, roster extraction). Runs **Gemini 3.5 Flash** ( `gemini-3.5-flash` ).

Both tiers run the same model today because no Pro model of that generation is served in `northamerica-northeast1` (Montréal), the Canadian Vertex AI region Certivo uses. The two tiers stay separate in our configuration and in the credit meter (§8.1), so the narrative tier can move to a Pro model when one is served in Canada; that move would be a sub-model upgrade under §14.3.
Both tiers are **explicitly pinned** to specific version ids in our Cloud Functions runtime configuration. We do **not** use floating tags ("gemini-latest") that could silently change model behavior between deploys. Model upgrades are deliberate code changes that go through our standard review process before shipping.

*Correction, August 15, 2026:* versions 1.0 to 1.2 of this Policy named **Gemini 3 Pro** and **Gemini 3 Flash**. That was accurate to the code's stated defaults at the time but not to what actually ran: the Gemini 3 model ids are not served in `northamerica-northeast1`, and both live environments overrode them back to the Gemini 2.5 family. The defaults were corrected in code on August 15, 2026 to the ids the region genuinely serves, and this Policy now matches. Gemini 2.5 has a published retirement date of October 16, 2026; the move to a newer generation will be reflected here at that time and is a sub-model upgrade under §14.3, not a sub-processor change.

*Correction, October 24, 2026:* versions 1.3 to 1.5 of this Policy named **Gemini 2.5 Pro** and **Gemini 2.5 Flash**. Both of Certivo's live environments moved to `gemini-3.5-flash` for both tiers on October 19, 2026, the day before Google retired the Gemini 2.5 models on Vertex AI (October 20, 2026; the October 16 date above was the Gemini Developer API's). Version 1.5, published October 18, still named Gemini 2.5. The provider (Google Cloud Vertex AI), the processing region and the data sent are unchanged; this was a sub-model upgrade under §14.3.

Certivo does **not** currently use:

- OpenAI GPT models, Anthropic Claude models, or any other non-Google LLM provider at the model-inference layer.
- Open-source LLMs deployed on Certivo-owned infrastructure.

- Image generation models (no DALL·E, Stable Diffusion, Midjourney, Imagen, etc.).
- Audio-generation or voice-cloning models.
- AI-driven facial recognition, biometric identification, or emotion analysis.

**Dependency-tree note (non-Google SDKs).** Certivo's Vertex AI client library — the official Google `@genkit-ai/google-genai` plugin (which replaced `@genkit-ai/vertexai` on August 13, 2026, after the older package's underlying class passed its stated removal date) — transitively includes third-party provider SDKs (for example, the Anthropic-on-Vertex SDK) because Google Cloud Vertex AI's Model Garden can serve non-Google models. Certivo does **not** invoke any of these; all inference is version-pinned to Gemini on Vertex AI (see §3 above). The presence of such an SDK in the dependency tree is a build-time artifact of the Google plugin — not an indication that a non-Google model is used, nor that any data is sent to a non-Google provider.
If Certivo adds a new model provider, a new model family, or a new modality (e.g., vision, audio), this Policy will be updated at least thirty (30) days before that change goes live to production. Material additions are also disclosed in the **Sub-Processor List** (Schedule G of the Master Subscription Agreement).

---

## 4. What AI Features Exist in the Platform

The following Certivo Intelligence flows are live as of the Effective Date. Each is described with: what it does, what data goes in, what comes out, and which tier(s) have access.

| Flow | What it does | Inputs (sent to AI) | Outputs | Tiers |
|---|---|---|---|---|
| **Financial Health Report** | Generates a quarterly narrative report on the Tenant's financial health, with cash position, AR aging context, top revenue clients, and trend commentary. | Aggregated financial metrics computed deterministically by Cloud Functions from the Tenant's own data (revenue totals, AR aging buckets, top-N client totals — all already aggregated, no individual invoice lines, no PII). | A narrative report (text, structured sections) plus an executive summary. | Professional, Scale, Enterprise |
| **Course Profitability** | Surfaces which of the Tenant's courses are most/least profitable, with narrative commentary on why. | Aggregated per-course metrics: revenue, instructor cost, material cost, student count, average margin. | Per-course narrative + ranked list. | Professional, Scale, Enterprise |

| Flow | What it does | Inputs (sent to AI) | Outputs | Tiers |
|---|---|---|---|---|
| **Recert Predictive Leads** | Identifies clients with cert expiries in the next 60/90/120 days and drafts a per-client narrative ("Acme Corp has 12 H2S certs expiring in October — likely renewal need of ~$3,600"). | Client name, count of expiring certs by course, expiry window. No employee names, no student PII. | Per-lead narrative + suggested outreach timing. | Professional, Scale, Enterprise |
| **Outreach Email Drafts** | Drafts a personalized outreach email to a prospect or recerting client. | Prospect/client name, industry segment, past course history (course names + dates, no student PII), Tenant's own brand voice settings. | A drafted email body for the operator to review and edit before sending. | Professional, Scale, Enterprise |
| **Expense Categorizer** | Suggests an expense category for an uploaded receipt or expense entry. | Merchant name, amount, date, free-text description. No payment method details, no card numbers. | A suggested category (e.g., "Equipment Repair", "Travel — Mileage"). | All tiers (subject to credit pool — see §8). |
| **Roster Parser** | Reads a regulator sign-in sheet or attendance roster that is **not** machine-readable — a PDF, a scan, or a phone photo — and extracts the student rows and course metadata. CSV and spreadsheet imports are parsed deterministically in the browser and **never reach a model**. | The uploaded PDF or image itself, in full. This file contains student names and, where the sheet has those columns, student emails, phone numbers and regulator ID numbers. It is the one AI flow whose input contains student personal information, because the document being read is the record. See §6.2. | Extracted rows + course metadata for the operator to review and correct before anything is saved. | All tiers |
| **QBO Invoice Analyzer** | Highlights anomalies in QuickBooks invoices that synced into Certivo (e.g., suspected duplicates, large variance from typical client invoice size). | Invoice totals + dates for the Tenant's own client(s). No client contact details, no payment method. | Per-anomaly narrative ("This invoice is 4× the client's usual size — confirm?"). | Professional, Scale, Enterprise |

| Flow | What it does | Inputs (sent to AI) | Outputs | Tiers |
|---|---|---|---|---|
| **Tenant Report** | Generates a monthly operations report for the Tenant's leadership. | Aggregated tenant metrics (class count, student count, payroll totals, revenue). | A narrative monthly report. | Scale, Enterprise |
| **Tenant Upsell Leads** | For Super Admin only — never visible to Tenants. Identifies Certivo customers who would benefit from an upgrade. | Aggregated platform-level usage metrics by tenant. | A ranked list with reasoning. | Internal Certivo use only |
| **Platform Financial Health** | Super Admin only. Platform-wide financial health snapshot. | Cross-tenant aggregated metrics. | A narrative report. | Internal Certivo use only |
| **Ask Certivo** | Chat agent for in-product questions about a Tenant's own data ("Which client owes me the most right now?", "When does Mike's H2S expire?"). | Each user turn + the records the agent's tools look up to answer it, such as a class with its roster, a certificate, a client or a student. A student's record includes the student's name, email address and phone number, so those are sent when a question needs that record (Section 6.1); accommodation flags and free-text notes about a student are removed first (Section 6.2). The agent runs under the user's auth context and respects TenantGuard, so it cannot access another Tenant's data. | A conversational response with citations to the source data. | Professional, Scale, Enterprise |

| Flow | What it does | Inputs (sent to AI) | Outputs | Tiers |
|---|---|---|---|---|
| **Client Assistant** | The same chat agent, offered to the Tenant's own clients in the Client Portal. **Off by default**; a Tenant switches it on in Settings ("Give your clients an assistant"). A client can ask only about their own classes, certificates, invoices and quotes. | Each user turn + the tools the agent calls, scoped to that client's own records. Runs under the client's auth context; it cannot read another client's data or the Tenant's internal data. | A conversational response about that client's own records. | Professional, Scale, Enterprise (Tenant opt-in) |
| **SMS Assistant** | Answers an inbound text message to the Tenant's Certivo number. **Off by default**; a Tenant switches it on in Settings. | The message text, and the sending phone number used to look up what Certivo already texts that number. | A reply limited to what Certivo already sends that number unprompted — class reminders and certificate-expiry information. A phone number is a weak identifier (SIMs are recycled, handsets are lost), so the reply is deliberately bounded to information the holder of that phone already receives. | Professional, Scale, Enterprise (Tenant opt-in) |
| **Email Reply Channel** | Lets a Tenant's operator reply to a Certivo email — the weekly digest, a report — and have the reply understood and acted on, rather than bouncing. **Off by default**; a Tenant switches it on in Settings. | The reply body and the thread it answers. | An answer or an action proposal, which follows the same confirmation rule as every other flow (§10.2). | Professional, Scale, Enterprise (Tenant opt-in) |
| **Feature Profitability** | Super Admin only. Identifies which Certivo features drive the most subscription value. | Platform-level aggregated usage. | A narrative analysis. | Internal Certivo use only |

| Flow | What it does | Inputs (sent to AI) | Outputs | Tiers |
|---|---|---|---|---|
| **Opportunities** | Surfaces growth opportunities for a Tenant (new course types, underutilized instructors, etc.). | Aggregated tenant metrics. | A list of ranked opportunities with rationale. | Professional, Scale, Enterprise |
| **Health Check** | Tenant-side health snapshot ("3 invoices overdue >60 days, 5 certs expiring this month"). | Aggregated tenant counts and date filters. | A snapshot summary. | All tiers |
| **Platform Health Check** | Super Admin only. Platform-wide ops snapshot. | Cross-tenant aggregates. | A platform-ops summary. | Internal Certivo use only |

Tenant-side AI credit consumption is visible in the Tenant Suite → **Intelligence Command Center**, which shows the monthly credit pool, how much of it has been used, and a per-flow breakdown of where it went. *Correction, August 15, 2026:* versions 1.0 to 1.2 said this screen also showed "which flow ran, when, who triggered it" and "a copy of the prompt + output for at least 30 days." It does not, and no screen does. The per-invocation record exists (see §11.1) but is server-side only, and it stores **hashes** of the input and output rather than copies — a deliberate minimisation choice, and one that means a Tenant-facing prompt viewer cannot be built from it. §11.4 describes the process for obtaining a rendered prompt on request.

---

## 5. Where Data Is Processed (Regional Compliance)

**5.1 Vertex AI region.** All Certivo Intelligence flows route to Google Cloud Vertex AI in `northamerica-northeast1` (Montréal, Canada). Vertex AI in Canada lives there; our Cloud Functions live in `northamerica-northeast2` (Toronto). Both are inside Canada. AI requests never leave Canada under normal operation.

**5.2 Google Cloud commitment.** Per Google Cloud's standard terms for Vertex AI in Canadian regions, model inference inputs and outputs are processed in the specified region. Google does **not** use Customer prompts or completions to train, fine-tune, or improve foundation models without an explicit opt-in. Certivo has not opted in and will not opt in without Tenant-by-Tenant consent.

**5.3 PIPEDA / PIPA alignment.** Routing AI inference inside Canada is a deliberate design choice for PIPEDA (Personal Information Protection and Electronic Documents Act) and Alberta PIPA compliance, and it matters because some flows do send personal information. Reports and the other narrative flows receive aggregated metrics, anonymized counts and operator-authored text. The chat agent sends the records a question needs, which can include a student's name, email address and phone number, and the Roster Parser reads the uploaded sheet itself. Section 6 sets out exactly what is and is not sent.

*Correction, October 18, 2026:* versions 1.0 to 1.4 said that no personal information about an identifiable individual is sent to AI in a form that identifies that individual, and that the inputs are aggregated metrics,

anonymized counts and operator-authored text. That was not true of the chat agent, whose lookups return student records, including names, email addresses and phone numbers.

**5.4 Cross-border.** If a Tenant explicitly enables a Beta feature that requires routing to a non-Canadian region (e.g., a future US-region model we evaluate), Certivo will disclose that in advance, require an explicit opt-in, and update the Sub-Processor List. As of the Effective Date, **no such feature exists or is planned for Q1–Q2 2027**.

---

## 6. What Data Is and Is Not Sent to AI

**6.1 Sent to AI:**

- **Aggregated metrics** the Tenant or Cloud Functions have already computed (revenue totals, AR aging totals by bucket, top-N client totals by spend, count of expiring certs by course, etc.).
- **Operator-authored free text** (a custom request to the chat agent, a free-text note attached to an expense, the body of an email draft the operator is iterating on).
- **Anonymized identifiers** (a client name only when the flow's job is to draft outreach to that named client, and only after the operator initiated that flow).
- **Schema metadata** (column headers from an uploaded roster, course type names from the Tenant's catalog).
- **Records the chat agent looks up to answer a question** (Ask Certivo, Section 4): a class with its roster, a certificate, a client, or a student record. A student record sent this way includes the student's name, email address and phone number. Accommodation flags and free-text notes about a student, whether on the student record, on a class roster entry or on a certificate, are removed before anything is sent (Section 6.2).

**6.2 NEVER sent to AI:**

- **Student personally identifiable information** — student names, emails, phone numbers, dates of birth, certificate numbers — is not sent to AI by any analytical or narrative flow. Those flows receive aggregated counts (e.g., "12 students completed H2S in October").

**One documented exception: the Roster Parser (§4).** When an operator uploads a sign-in sheet as a PDF, scan or photo, the file is sent to the model as-is, and that file contains the students' names — that is what a sign-in sheet is. There is no way to extract a roster from an image without the model seeing the roster. The exception is bounded: the operator initiates it, the file is processed in `northamerica-northeast1` and not retained by the model provider, only hashes of it reach our audit log (§11.1), and CSV/spreadsheet imports never use a model at all. *Correction, August 15, 2026:* versions 1.0 to 1.2 stated this as an unqualified "NEVER" while the Roster Parser was live and vision-based. The rule and the product disagreed, and the product was right about what it does. The exception is now stated rather than the absolute overclaimed. **A second documented exception: the chat agent (Section 4, Section 6.1).** When a question needs a student's record, the agent's lookup returns it, and it includes the student's name, email address and phone number. This is a deliberate choice: the agent exists to answer questions about the Tenant's own students and classes, and the processing stays in `northamerica-northeast1` (Section 5). The student's accommodation flags and free-text notes are removed from the record before it is sent (see Health information below). *Correction,

October 18, 2026:* versions 1.0 to 1.4 listed student names, emails and phone numbers under this heading without saying that the chat agent sends them. It does, and this version says so.

- **Financial account details** — credit card numbers, bank account numbers, Stripe customer IDs, individual transaction amounts at the per-payment level. Aggregates only.
- **Authentication credentials** — passwords, PINs, MFA secrets, API tokens, OAuth tokens. Hard barrier.
- **Government identification** — SIN, driver's license, passport numbers, ID photos. The Platform does not collect these at all, so they are not present in any data Certivo assembles into a prompt. See §6.3 for what happens if a user types one into a free-text field.
- **Health information** — medical conditions, accommodation notes, anything that could be a personal health record. Student safety flags (e.g., "needs interpreter") are stored in Firestore but not sent to AI.

*Correction, October 18, 2026:* until this date the chat agent returned whole records, so a student's accommodation flags (with their notes) and the free-text notes kept on a student record, a class roster entry or a certificate reached the model when the agent listed students, looked up classes or searched certificates. They are now removed before anything is sent, and an automated test calls every chat-agent tool and fails if any of them comes back. The rule in this bullet is unchanged; the product now follows it.

- **Incident details with named victims** — incident reports are restricted to operators; aggregate counts of incidents per month may be sent, but the body of an incident report is not.
- **Other tenants' data** — TenantGuard prevents cross-tenant queries from the user's session, and the AI agent runs under that session, so a Tenant's data cannot reach another Tenant's AI flow.
- **Customer Data of any kind for model training** — see §7.

**6.3 Data excluded from AI processing (sensitive identifiers).** As a matter of design, Certivo builds its AI features so that the following sensitive personal identifiers are **excluded from the prompt/input sent to the AI model** — they are minimized, redacted, or withheld at the Cloud Function layer before the prompt is assembled, and **none of them is required for any AI feature to function**:

- **Social Insurance Numbers (SINs)** and other government identifiers (driver's licence, passport numbers, ID photos);
- **Full financial account, bank account, or payment card numbers**, and any other individual financial account details; and
- **The sensitive contents of health, injury, or safety incident reports** — including medical conditions, accommodation notes, and any narrative body identifying an injured or affected individual.

Aggregated, de-identified counts derived from this data (e.g., "12 students completed H2S in October," "3 incidents this month") may still be sent, as those contain no identifiable individual.
**What is and is not automated, stated precisely (corrected August 15, 2026).** Versions 1.1 and 1.2 said that if a user pasted a SIN into a free-text field, "it is redacted at the Cloud Function layer before the prompt is built." That described a scanner that does not exist. What does exist:

- **Certivo does not collect any of the identifiers above,** so they cannot enter a prompt through the data Certivo assembles. This is the control that does the real work, and it is structural: you

cannot redact a field you never had.

- **Tool results returned to the chat agent are field-level scrubbed.** A defined set of personal fields is replaced with `[redacted]` before results are handed back to the model, recursively through nested objects. This governs data the agent looks up, not text a user types.
- **Text a user types into the chat agent is passed through as written.** If an Authorized User types a SIN into a question, that text reaches the model. There is no automated detector for free-text government identifiers, and this Policy no longer implies otherwise.

Tenants should treat the AI chat box the way they would treat any outbound message field: do not paste identifiers into it that you would not send. If such data is pasted in error, that is a **Security Incident** under the Information Security Policy and must be reported to `aaron@certivo.ca` immediately for investigation and (where required) breach notification under PIPEDA and PIPA.
If a Tenant believes any of these classes of data have been sent to AI in error, this is a **Security Incident** under the Information Security Policy and must be reported to `aaron@certivo.ca` immediately for investigation and (where required) breach notification under PIPEDA and PIPA.

---

## 7. Model Training — Customer Data Is Never Used to Train Models

**7.1 The hard rule.** Certivo does **not**, and **will not**, use Customer Data — including any data sent to AI through the Platform — to train, fine-tune, retrain, distill, or otherwise improve any foundation model, embedding model, or other machine learning model, whether for Certivo's own benefit or any third party's.

**7.2 Vertex AI commitment.** Google Cloud Vertex AI's terms specifically exclude customer prompts/completions from foundation-model training without an explicit opt-in. Certivo has not opted in.

**7.3 Internal evaluation use.** Certivo intends to maintain a small **golden evaluation set** — synthetic prompts and expected outputs Certivo authors itself — to detect regressions when model versions change. This set never contains Customer Data. If a real Customer Data sample would be useful to add to the eval set (e.g., a parsing failure on an unusual roster), Certivo will request explicit written consent from the Tenant before doing so, will redact any PII, and will document the consent in the audit log.

> **Open item, tracked separately.** The golden-set claim in versions 1.0 to 1.2 was written in the present tense before the harness existed. The commitment that matters to a Tenant — no Customer Data used for evaluation without written consent — is and remains true, because there is nothing to put it in. The build-out of the evaluation harness itself is tracked as its own work item and this paragraph will move back to the present tense when it lands.

**7.4 No reverse engineering of model behavior.** Certivo does not run statistical analyses across Tenant outputs to "learn" Tenant patterns and feed them back to models. Aggregated platform metrics (number of flows run per month, average credit consumption, success/failure rates) are tracked for operational and billing purposes only.

**7.5 No cross-tenant benchmarking — a deliberate product decision.** Certivo does **not** compare one Tenant's operational or financial performance against another's, and does not surface peer benchmarks, percentile rankings, industry-average comparisons, or "companies like yours" metrics derived from other Tenants' data. This is a choice, not a gap in the roadmap.

The reason is straightforward: **Certivo's Tenants compete with one another.** Two safety-training companies bidding the same contract in the same region are both customers of this platform. A benchmark is not a neutral statistic when the comparison set is your competitors — it is a disclosure of their pricing, their utilisation, their close rates and their margins, aggregated well enough to be useful and therefore well enough to be inferred from. At the scale a regional vertical actually operates, "anonymised industry average" across a handful of local competitors is not anonymous.

Certivo will not put a Tenant in the position of having to weigh whether using the product leaks their commercial position to a rival.

What this rules out, explicitly:

- Peer or cohort benchmarking of any operational or financial metric.
- Percentile or ranking displays positioned against other Tenants.
- "Top performers in your industry do X" guidance derived from other Tenants' data.
- Any AI flow whose prompt or grounding context contains another Tenant's data.

What Certivo does instead: **every metric is compared against the Tenant's own history.** Trends, seasonality, targets the Tenant sets, and thresholds calibrated from the Tenant's own operating profile. A Tenant's own past is a fairer and more actionable comparator than a competitor's present, and it requires no one else's data to compute.

This constraint is enforced structurally, not by policy alone: AI flow context is assembled server-side and scoped to the invoking Tenant's `orgId`, the same isolation boundary that governs all Platform data access. Cross-tenant aggregate data exists only in Certivo's own operational and billing metrics (§7.4), which are never returned to a Tenant surface.

If Certivo ever proposes changing this position, it will require **explicit, informed, per-Tenant opt-in** — never a default, never a settings toggle buried in an update — and thirty (30) days' advance written notice.

---

## 8. AI Credits, Quotas, and Cost Control

**8.1 Credit pools.** Each tier has a monthly AI credit pool. Credits are consumed per flow invocation, with cost depending on the flow's tier (narrative or high-volume, see §3) and the flow's weight. Current pool sizes, as enforced by the server:

| Tier | Monthly Credit Pool |
|---|---|
| Free | 0 credits (AI features off) |
| Starter | 100 credits |
| Professional | 500 credits |
| Scale | 1,000 credits |

| Tier | Monthly Credit Pool |
|---|---|
| Enterprise | 2,000 credits |

Capacity in "flow invocations" depends entirely on which flows: a high-volume expense categorization costs 1 credit and a full narrative report on the narrative tier costs 50, so a single figure would be misleading. The per-flow credit costs are shown in the Intelligence Command Center breakdown.
*Correction, August 15, 2026:* versions 1.0 to 1.2 listed Scale at 2,000 credits and Enterprise at "typically 5,000+", and advertised a "**Certivo Intelligence Boost** — +1,000 credits per month for $29 CAD/month." None of the four figures matched what the platform enforces. Scale is 1,000 and Enterprise is 2,000. There is no product called Intelligence Boost; the bonus-credit mechanism that exists is the **Intelligence Command Center** entitlement, which grants +5,000 credits per month and is bundled into paid plans rather than sold at $29. The approximate-capacity column has been removed rather than corrected, because it was never a meaningful number. **Authoritative figures live in the plan catalogue and the server quota table, not in this document**; where they differ, the server governs and this document is wrong and should be reported.

**8.2 Quota warnings.** When a Tenant has 20% of its monthly credit pool left, the platform surfaces a warning; at 5% remaining the warning escalates. At 0 remaining, AI flows are refused for the rest of the calendar month with an explicit message saying so. Deterministic flows (scheduling, billing, invoicing, certificates) are never affected.

**8.3 Per-flow cost visibility.** Every AI flow records its credit cost, and its measured model cost, against the Tenant's monthly usage record. Operators can view the pool, the balance and the per-flow breakdown in the **Intelligence Command Center** at any time.

**8.4 No surprise overages.** Certivo does not auto-bill overage charges. When the pool is exhausted, flows pause until the next monthly reset or until the Tenant moves to a plan with a larger pool.

**8.5 Automatic model escalation (roster reading only).** The Roster Parser reads a document on the cheaper model first. If that read comes back with low confidence — a skewed photo, handwriting, a faint scan — it re-reads once on the stronger model, because the alternative is handing the operator a roster they have to retype. That second read costs additional credits from the same pool, it happens at most once per upload, it is refused outright if the Tenant cannot afford it, and **the result tells the operator it happened**. No other flow escalates itself.

---

## 9. Customer Opt-Out

**9.1 Per-tenant opt-out.** A Tenant may disable Certivo Intelligence entirely at any time via **Settings** → **Intelligence** → **Enable Certivo Intelligence**. While disabled, no AI flow can be invoked by any user in the Tenant's organization; the refusal is enforced server-side, ahead of any request reaching a model, not merely by hiding buttons. The Tenant continues to pay its subscription fee and recovers none of the credit pool — the pool renews and expires on its normal schedule whether or not it is used. The Tenant can re-enable at any time from the same screen. Certivo support can also disable AI for an organization on request; a support-initiated disable cannot be reversed from the Tenant screen and requires contacting support.

*Correction, August 15, 2026:* this control was promised from version 1.0 and **did not exist in the product until August 15, 2026**, when it was built at the shared access-control chokepoint every AI flow passes through. The screen name in earlier versions ("Settings → AI Usage → Disable AI features") pointed at a screen that has never existed. A Tenant who relied on this Policy between May and August 2026 and wanted AI switched off should have contacted `aaron@certivo.ca`, which remains a supported route; any Tenant who did so and was not accommodated should raise it at that address.

**9.2 Per-feature opt-out. Not yet available.** Versions 1.0 to 1.2 stated that Tenants on Professional or above could disable individual AI flows. No such control exists in the product. Today the choice is all Certivo Intelligence (§9.1) or none. Per-feature control is on the roadmap; this Policy will be updated when it ships and not before.

**9.3 Per-user access.** Access to AI features is governed by role. Administrators and managers can invoke AI flows; the expense categorizer is additionally available to instructors and staff; the roster parser and the reporting flows are administrator-only. This is enforced server-side on every call. There is not currently a separate per-user AI toggle beyond these role gates, and versions 1.0 to 1.2 overstated the degree to which a Tenant administrator can tune it.

**9.4 Client Portal users.** Client Portal Users (the Tenant's customers' employees) have **no AI feature unless their training provider switches one on.** The Tenant controls this in Settings ("Give your clients an assistant"), it is **off by default**, and when it is on a client may ask only about their own classes, certificates, invoices and quotes — enforced server-side under that client's own auth context, not by hiding a menu. Outputs a Tenant sends TO a client (e.g., a renewal-suggestion email) remain operator-reviewed and operator-sent; nothing is auto-delivered to a client. *Correction, September 19, 2026: versions 1.0 to 1.3 said Client Portal Users have no access to any AI feature. That has not been true since the client assistant shipped. The setting has always defaulted to off and has always been the Tenant's choice, but a Tenant who read this section and left the switch alone was relying on a sentence that no longer described the product.*

---

## 10. Outputs Are Advisory — Not Authoritative

**10.1 Outputs are recommendations, not decisions.** Every AI-generated output in the Platform is **advisory**. A narrative report describes what the deterministic numbers show; a draft email is a starting point the operator edits; a categorization suggestion is reviewable and overrideable. The Tenant's operators retain full decision authority.

**10.2 No automated regulatory actions.** AI does not auto-issue certificates, auto-approve payroll, auto-send invoices, auto-send certificate renewals to clients, auto-process refunds, or take any other state-changing action without explicit operator confirmation. Every state change goes through the existing deterministic Cloud Function with the existing audit log.

**10.3 No diagnosis, no legal advice, no financial advice.** Certivo Intelligence does not make medical/clinical determinations, give legal advice, give individualized financial advice, or evaluate fitness for a specific occupation. The reports it generates are operational summaries of the Tenant's own data, not professional advice. Tenants requiring professional advice should consult qualified counsel, accountants, or regulators.

**10.4 Hallucinations and errors.** Generative models can produce confident-sounding but factually wrong output ("hallucinations"). Certivo mitigates this by: (a) using deterministic Cloud Functions to compute the numbers AI narrates, so the numbers themselves come from the database and not the model; (b) showing the source data alongside the narrative wherever feasible; (c) including a "verify before acting" notice on every AI-drafted email and report. **The Tenant remains responsible for verifying AI output before relying on it for any business decision.**

**10.5 Advisory only — human review required.** All AI-generated output is **advisory and informational only**. It is **not** professional, legal, safety, medical, tax, accounting, or regulatory-compliance advice, and it may be inaccurate, incomplete, or out of date. No AI output should be relied upon or acted on until a qualified person at the Tenant has **reviewed and verified** it. The Tenant is solely responsible for that review and verification before any AI output is used, sent, filed, or otherwise relied upon.

**10.6 Tenant responsibility and indemnity.** The Tenant assumes responsibility for the consequences of relying on AI output that it has not adequately reviewed. To the extent set out in, and subject to the limits and procedures of, the **indemnification provisions of the Master Subscription Agreement**, the Tenant will indemnify and hold Certivo harmless from losses, claims, and liabilities arising out of the Tenant's reliance on an AI-generated draft or output without adequate human review. This Section does not create a separate or broader indemnity than the Master Subscription Agreement; where they differ, the Master Subscription Agreement governs.

**10.7 Limitation of liability.** Section 17 (Limitation of Liability) of the Master Subscription Agreement applies in full to all AI-generated outputs.

---

## 11. Logging, Audit, and Transparency

**11.1 Per-flow audit log.** Every AI flow invocation writes one row to the `ai_audit_logs` collection with: timestamp, organization id, user UID, flow name, pinned model version, run id, duration, input and output token counts, estimated cost, success / failure / cached status, and **SHA-256 hashes of the input and the output**. The input and output themselves are not stored in this record — only their hashes, so a specific run can be identified and correlated but its content cannot be read back out of the log.

The separate `ai_usage` record holds the monthly credit meter: credits consumed, per-flow breakdown, and measured model spend. Versions 1.0 to 1.2 described `ai_usage` as holding the per-invocation rows; that was the wrong collection, and the correction matters because the two have different contents and different sensitivity.

**11.2 Report retention.** For narrative flows (Financial Health Report, Course Profitability, Operations Report, Business Overview), Certivo retains the generated **report** — the output — in an internal `ai_reports` history, together with the figures it was computed from and the model and prompt versions in force at the time. This supports operator review, regression debugging, and incident response.

*Correction, August 15, 2026:* versions 1.0 to 1.2 said this history held "the rendered prompt and the rendered output" for "thirty (30) days" in `ai_drafts`, purged by a daily sweep. Three things were wrong. The **prompt is not stored** — only the output and the figures behind it. The collection is `ai_reports`, not `ai_drafts` (which holds AI suggestions awaiting operator approval, a different thing). And there is **no 30-day purge**:

report history is retained indefinitely as an operational and methodology record. Deletion on a Tenant's request remains available — every record carries the organization id, so erasure by organization is supported under the Data Retention Policy and PIPEDA — but it is request-driven, not automatic, and Tenants who understood otherwise from earlier versions of this Policy should assume their report history is still held and may request its deletion at `aaron@certivo.ca`.

**11.3 Customer-visible transparency.** Tenant administrators can see the credit pool, the balance and the per-flow usage breakdown in the Intelligence Command Center. A self-serve **export** of the usage log is **not currently available** — versions 1.0 to 1.2 described a "Settings → AI Usage → Export" control that does not exist. Until it does, a Tenant administrator may request their organization's AI usage log in writing at `aaron@certivo.ca` and Certivo will provide it within five (5) Business Days.

**11.4 Operator review SLA.** Certivo does not currently surface real-time prompt logs to Customers (we keep the prompt-rendering layer internal so we can evolve it without breaking external contracts). However, on written request, Certivo will provide a tenant administrator with the actual rendered prompt for any specific flow invocation in the last 30 days within five (5) Business Days.

---

## 12. Bias, Fairness, and Discrimination

**12.1 Avoidance of high-risk decisions.** Certivo Intelligence does not make hiring decisions, lending decisions, insurance decisions, housing decisions, or any other decision class subject to anti-discrimination law. The flows surface aggregated patterns and draft text; people make decisions.

**12.2 Outreach email caution.** AI-drafted outreach emails are operator-reviewed before sending. Tenants must not use AI-drafted text in a way that creates the appearance of personalized correspondence without operator review — that risks CASL violation (see Acceptable Use Policy §7) and undermines the editorial human-in-the-loop posture of this Policy.

**12.3 Bias reporting.** If a Tenant believes a Certivo Intelligence output reflects bias (e.g., consistently undervalues certain courses, demographic patterns in expense classification, etc.), the Tenant should report it to `aaron@certivo.ca` with the specific flow invocation ID. Certivo will investigate within ten (10) Business Days and may pause the affected flow while investigating.

---

## 13. Security Controls Specific to AI

Beyond the general controls in the Information Security Policy, the following apply to AI:

- **Prompt construction is centralized.** All prompts are constructed server-side by Cloud Functions, not by the client. Clients cannot send arbitrary prompts to AI — they invoke a named flow with structured inputs, and the Cloud Function builds the prompt.
- **Input handling.** Prompts are assembled from structured, server-computed context plus the operator's own text. Free-text is length-capped and structurally constrained by the flow that carries it. It is **not** scanned for prompt-injection patterns or for sensitive identifiers — versions 1.0 to 1.2 claimed both, and neither is implemented. The controls that carry the weight here are the

ones below: the model cannot take actions, its output is treated as untrusted, and the numbers it narrates come from the database rather than from the model.

- **Output isolation.** AI outputs returned to the Platform are treated as untrusted strings. They are HTML-escaped before being rendered in any UI surface. AI outputs are never executed as code, never used to construct database queries, and never used to determine access permissions.
- **Rate limiting.** Per-flow rate limits prevent runaway invocation costs, alongside the credit pool in §8.
- **No autonomous actions.** AI flows return data to the calling Cloud Function; they cannot call mutating functions on their own. State-changing tools available to the chat agent (e.g., "draft an email", "create a quote") return drafts for operator confirmation, never auto-execute.
- **Emergency stop.** Certivo operators can pause every generative AI surface across the entire platform with a single switch, independently of any Tenant's own setting. This exists for cost runaways, provider incidents, and model misbehaviour. As of August 15, 2026 an automated check runs before every code change and fails it if a new AI entry point is added that does not honour that switch — previously this was a convention, and conventions are missed.
- **Operational monitoring.** Per-flow latency, error rate and spend are exported to Certivo's monitoring, and a weekly internal digest summarises AI failure rates, cache-hit ratios, real spend and the most common output-verification warnings. This digest is internal operator telemetry and is never sent to a Tenant.

---

## 14. Subprocessors and Sub-Models

**14.1 Sub-processor list.** Google Cloud (Vertex AI, in `northamerica-northeast1`) is the sole AI sub-processor as of the Effective Date. The full Sub-Processor List is Schedule G of the Master Subscription Agreement and is also published at `certivo.ca/legal/sub-processors`.

**14.2 New sub-processor notice.** Certivo will provide thirty (30) days' advance written notice before adding a new AI sub-processor (e.g., switching to or adding a non-Google provider). During the notice period, Tenants may terminate this Policy's authorization and rely on the Sections 9 opt-out controls; persistent objection that cannot be remediated is grounds for termination under the Master Subscription Agreement.

**14.3 Sub-model upgrades.** Moving a tier from one Gemini model version to a later one on Vertex AI (for example the October 2026 move from Gemini 2.5 to Gemini 3.5 Flash, or a future move of the narrative tier to a Pro model) is not a sub-processor change and does not require thirty-day notice. We do, however, communicate model version pins in our public changelog and update this Policy on the next regular review.

---

## 15. Tenant Obligations

**15.1 Authorized use only.** Tenants must not use AI features:

- to draft communications that violate CASL (Canada's Anti-Spam Legislation) — see Acceptable Use Policy §7;

- to generate certificates, regulatory records, or any artifact that misrepresents the qualifications of a student or instructor;
- to draft content that would, if sent by a human, violate the Master Subscription Agreement's prohibited-use restrictions;
- to circumvent any access control, rate limit, or quota in the Platform;
- to reverse-engineer the Platform's prompt construction or model behavior for the purpose of building a competitive product.

**15.2 Disclosure to Tenant's end users.** Tenants who incorporate Certivo Intelligence output into communications to their own clients or students (e.g., AI-drafted renewal nudges) are responsible for disclosing the use of AI in those communications where required by applicable law or industry guidance. **15.3 Indemnification.** Tenants indemnify Certivo for misuse of AI features outside the authorized scope of this Policy and the Acceptable Use Policy.

---

## 16. Changes to This Policy

Certivo may update this Policy from time to time. Material changes (new model providers, expanded data classes sent to AI, removed opt-outs, new mandatory features) will be notified to Tenants at least thirty (30) days in advance via email to the Tenant administrator and via in-Platform notification. Non-material changes (typo fixes, clarification of existing scope, new flows that follow the existing data-handling rules) will be reflected in the version date at the top of this document and disclosed in our public changelog.

---

## 17. Contact

**Ready 4 Everything Training & Supplies Ltd.**
4952 Westbrooke Rd.
Blackfalds, Alberta, Canada T4M 0L1

- **AI / Privacy questions:** `aaron@certivo.ca`
- **Subject line for AI-related requests:** `[AI USAGE]` (e.g., `[AI USAGE] Opt-out request for tenant orgId-XXXX`)
- **AI incident reports:** `aaron@certivo.ca` with subject `[AI INCIDENT]`
- **Subprocessor list:** `certivo.ca/legal/sub-processors`

---

*This AI Usage Policy is part of the Certivo Master Subscription Agreement. It does not constitute legal advice. The Tenant remains responsible for ensuring its own use of Certivo Intelligence outputs complies with applicable law in its jurisdiction.*

**Revision history**

- *v1.0 — May 2026 — initial publication.*

- *v1.1 — July 5, 2026 — counsel review: added §6.3 (sensitive identifiers excluded from AI prompts — SINs, full financial account/payment card numbers, and health/injury incident report contents); added §10.5–§10.6 (AI output is advisory only, requires human review, and Tenant indemnifies Certivo for reliance without adequate review, subject to the Master Subscription Agreement); former §10.5 renumbered to §10.7.*
- *v1.2 — July 11, 2026 — clarified in §3 that the Google Vertex AI client library ( `@genkit-ai/vertexai` ) transitively includes non-Google provider SDKs (e.g., Anthropic-on-Vertex) that Certivo does not invoke; all inference remains version-pinned to Gemini on Vertex AI. No change to data handling.*
- *v1.3 — August 15, 2026 — **accuracy revision.** A line-by-line reconciliation of this Policy against the implementation found nine statements that did not match the product. Each is corrected in place with a dated note saying what the earlier text claimed, rather than being quietly rewritten, because a Tenant who relied on the old text needs to be able to find out that they did. Summary:*

- *§3 — models are Gemini **2.5** Pro / Flash, not Gemini 3; client library is `@genkit-ai/google-genai`.* - *§4 — the Roster Parser reads PDFs, scans and photos via vision, not CSV column headers; the "prompt + output visible for 30 days" screen does not exist.* - *§6.2 — the "student PII is NEVER sent" absolute now states its one real exception: the roster document itself.* - *§6.3, §13 — no automated free-text redaction or prompt-injection scanning exists; the controls that do exist are stated instead.* - *§7.3 — the golden evaluation set is an intent, not a shipped harness; tracked separately. The no-Customer-Data commitment is unchanged.* - *§8.1 — Scale is 1,000 credits (was stated 2,000), Enterprise 2,000 (was "5,000+"); the "$29 Intelligence Boost" product does not exist. §8.5 added for roster model escalation.* - *§9.1 — the per-tenant opt-out promised since v1.0 **did not exist until August 15, 2026**; it now does, enforced server-side. §9.2 per-feature opt-out is withdrawn as not built. §9.3 restated to match the actual role gates.* - *§11.1–11.3 — per-invocation rows live in `ai_audit_logs` with hashed input/output, not `ai_usage`; report history is the output only, in `ai_reports`, retained indefinitely rather than purged at 30 days; no self-serve usage export exists.* - *§13 — added the platform-wide AI emergency stop and its new automated enforcement, and the internal AI health monitoring.* - *Two controls were also **added** in the same change rather than merely documented: the per-tenant opt-out (§9.1) and the pre-merge check that a new AI entry point honours the emergency stop (§13).*

- *v1.4 — September 19, 2026 — **scope revision.** Three AI surfaces that exist in the product were missing from this Policy, and one statement in it was no longer true:*

- *§4 — added the **Client Assistant**, the **SMS Assistant** and the **Email Reply Channel**. All three are Tenant opt-in and off by default, and all three were shipped without being listed here.* - *§9.4 — rewritten. It said Client Portal Users have access to no AI feature; they have the client assistant when their provider switches it on. The correction is stated in place rather than quietly replaced, for the same reason v1.3 gave: a Tenant who relied on the old sentence should be able to find out that they did.* - *No change to what data is sent (§6), to the advisory-only rule (§10), or to the retention and audit commitments (§11).*

- *v1.5, October 18, 2026: **accuracy revision.** This Policy said student names, email addresses and phone numbers are not sent to AI. The chat agent's lookups send them:*

- *Sections 4, 5.3, 6.1 and 6.2: a student's name, email address and phone number can be sent to the model when a chat-agent question needs that student's record. Stated in place with dated correction notes, for the*

reason v1.3 gave.* - *Section 6.2: the chat agent was also sending a student's accommodation flags and the free-text notes kept about a student (on the student record, a class roster entry or a certificate), which this Policy has always said are not sent. The product was corrected rather than the Policy: they are now removed before anything is sent.* - *No change to model providers, the processing region, training, opt-out, the advisory-only rule, or retention.*

- *v1.6, October 24, 2026: **accuracy revision.** Section 3 named Gemini 2.5 Pro and Gemini 2.5 Flash, which Google retired on Vertex AI on October 20, 2026. Since October 19 both tiers run Gemini 3.5 Flash (`gemini-3.5-flash`):*

- *Section 3: names the model each tier runs, says why the narrative tier runs a Flash model, and adds a dated correction note, for the reason v1.3 gave.* - *Sections 8.1 and 14.3: refer to the narrative and high-volume tiers rather than to Gemini 2.5 Pro and Flash.* - *No change to the provider, the processing region, the data sent, training, opt-out, the advisory-only rule, or retention.* *Reviewed annually.* *Ready 4 Everything Training & Supplies Ltd.*