

Scoped-Key Creation Guide

Step-by-step instructions for IT administrators

Reflects Global Barry's current security posture as of **October 2, 2026**. Items not yet in place are marked **In progress**; nothing in this document asserts a certification Global Barry does not hold.

Why scoped credentials

Global Barry API credentials are issued with a **read-only scope** (`read`) and a per-key rate limit (default 60 requests/minute). Global Barry API credentials (`bgk_<prefix>_<secret>`) are stored only as a one-way SHA-256 hash plus a short public prefix. The secret is shown once at creation/rotation and cannot be recovered or reversed. Each credential can be rotated or revoked independently from the developer portal.

Part A — Create a restricted, read-only Global Barry API credential

1. Sign in at <https://barry-intel.abacusai.cloud/login> with an account whose plan includes API access (API access depends on plan; see <https://barry-intel.abacusai.cloud/pricing>).
2. Open the developer portal: <https://barry-intel.abacusai.cloud/developer>.
3. Enter a descriptive name tied to one integration (e.g. `risk-dashboard-prod`). Use one credential per system so each can be revoked without affecting others.
4. Select **Create credential**. The credential is issued with the read-only `read` scope.
5. Copy the full credential immediately into your organisation's secret manager (e.g. AWS Secrets Manager, Azure Key Vault, HashiCorp Vault). **It is displayed only once**; Global Barry keeps only its hash and cannot show it again.
6. Send it from your integration in the request header documented at <https://barry-intel.abacusai.cloud/api/docs>. Never embed it in client-side code or source control.
7. Review usage in the portal. **Rotate** on a schedule (e.g. every 90 days) or on staff change; **Revoke** immediately if exposure is suspected.

Part B (optional appendix) — Restrict an external model provider key (OpenAI example)

Global Barry does not require you to supply any third-party API key: its AI features run on Global Barry's own managed inference endpoint. This appendix applies only if your team pairs Global Barry with its own model-provider key in its own workflows, or arranges a Bring-Your-Own-Vault enterprise deployment.

1. In the OpenAI platform, create a dedicated **Project** for the integration (e.g. `barry-global-research`).

2. Under the project's model permissions / limits, allow only a **single low-cost model**, for example `gpt-4o-mini`, and block all others.
3. Set a strict **monthly budget / spend limit**, for example **\$50**, with an alert threshold below it (e.g. 80%).
4. Create a **project-scoped API key** with restricted permissions (only the endpoints required, e.g. model requests) rather than an all-access user key.
5. In data controls, **disable sharing of data for model training** and apply the shortest available data-retention option your agreement permits (zero data retention may require an enterprise agreement with the provider).
6. Store the key in your vault and inject it at request time; rotate and revoke from the provider console.

Provider console labels change over time; follow the provider's current documentation for exact menu names.

Administrator checklist

- ☐ Global Barry credential created per integration, with a descriptive name
- ☐ Scope confirmed as read-only (`read`)
- ☐ Secret copied once into the corporate secret manager; not stored in code, tickets or chat
- ☐ Rotation schedule defined (e.g. 90 days) and owner assigned
- ☐ Revocation procedure documented for incident response
- ☐ (If applicable) External provider key limited to one model (e.g. `gpt-4o-mini`)
- ☐ (If applicable) Monthly spend cap set (e.g. \$50) with alerting
- ☐ (If applicable) Training data sharing disabled and minimum retention applied