Logging

- Default configuration
- Per-request logging
  - Collect logs (cf-aig-collect-log)
  - Collect log payload (cf-aig-collect-log-payload)
- DLP fields in logs

import { Render } from "~/components";

Logs help you monitor requests and troubleshoot application issues. Each log can include the prompt, response, provider, timestamp, status, token usage, cost, duration, and user agent.

When a request matches a Data Loss Prevention (DLP) policy, its log can also include the action and matched policy details.

When Guardrails are turned on, Guardrails evaluation calls are also recorded in your logs.

:::note[Logging eligibility] New AI Gateway customers who create their first gateway on or after September 24, 2026 follow Workers Logs pricing and retention.

AI Gateway customers who created a gateway before September 24, 2026 use Legacy Logs. :::

Use an authenticated gateway to prevent unauthorized access. Authentication also prevents invalid requests from increasing log volume.

Default configuration

Logs are enabled by default for each gateway. This setting applies to all requests unless a request overrides it.

You can turn off collection for privacy or compliance requirements.

Per-request logging

To override the default logging behavior set in the settings tab, you can define headers on a per-request basis.

These headers only control the logs that AI Gateway stores. They do not change what Logpush exports.

Collect logs (`cf-aig-collect-log`)

The `cf-aig-collect-log` header allows you to bypass the default log setting for the gateway. If the gateway is configured to save logs, the header will exclude the log for that specific request. Conversely, if logging is disabled at the gateway level, this header will save the log for that request.

In the example below, we use `cf-aig-collect-log` to bypass the default setting to avoid saving the log.

```
# Run `wrangler whoami` to get your account ID to replace $CLOUDFLARE_ACCOUNT_ID,
# and `wrangler auth token` to get an auth token to replace $CLOUDFLARE_API_TOKEN.
curl -X POST
"https://api.cloudflare.com/client/v4/accounts/$CLOUDFLARE_ACCOUNT_ID/ai/v1/chat/completions"
\
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --header "cf-aig-collect-log: false" \
  --data '{
    "model": "openai/gpt-4.1-mini",
    "messages": [
      {
        "role": "user",
        "content": "What is the email address and phone number of user123?"
      }
    ]
  }'
```

Collect log payload (`cf-aig-collect-log-payload`)

The `cf-aig-collect-log-payload` header allows you to control whether the raw request and response bodies (payloads) are stored for a given request. Unlike `cf-aig-collect-log`, which controls the entire log entry, this header only affects payload storage — metadata such as token counts, model, provider, status code, cost, and duration will still be logged.

This is useful when you want to maintain visibility into usage metrics and request metadata without persisting sensitive prompt or completion data.

Header value	Behavior
`true`	Request and response payloads are stored.
`false`	Payload storage is skipped. Metadata-only log entries are still saved.

In the example below, we use `cf-aig-collect-log-payload` to skip storing the request and response bodies while keeping the metadata log.

```
# Run `wrangler whoami` to get your account ID to replace $CLOUDFLARE_ACCOUNT_ID,
# and `wrangler auth token` to get an auth token to replace $CLOUDFLARE_API_TOKEN.
curl -X POST
"https://api.cloudflare.com/client/v4/accounts/$CLOUDFLARE_ACCOUNT_ID/ai/v1/chat/completions"
\
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --header "cf-aig-collect-log-payload: false" \
  --data '{
    "model": "openai/gpt-4.1-mini",
    "messages": [
      {
        "role": "user",
        "content": "What is the email address and phone number of user123?"
      }
    ]
  }'
```

:::note If `cf-aig-collect-log` is set to `false`, the entire log entry (including metadata) is skipped regardless of the `cf-aig-collect-log-payload` value. Use `cf-aig-collect-log-payload: false` on its own if you only want to suppress payload storage while retaining metadata logs. :::

DLP fields in logs

When Data Loss Prevention (DLP) policies are enabled on a gateway, log entries for requests that trigger a DLP policy match include additional fields:

□□□ #2
□ □□□ □□□ 2026-10-31 23:40:34 UTC
□ □□□ □□□ 2026-10-31 23:40:35 UTC