

Reviewing AI Usage and Request Activity

August 25, 2026

Opening the AI usage dashboard and understanding what each metric means

In Atloria, open the **Admin** workspace and select **Analytics** to reach the **Analytics & Insights** area. This screen is where administrators and team leads review usage statistics and performance metrics. Start by checking the date range shown at the top of the page or in the reporting controls if your workspace includes them. That date range determines which activity appears in every chart, summary card, and table you review, so confirm it before drawing conclusions from the numbers.

The first items to review are the top summary cards. In most AI usage reviews, these are the fastest way to understand overall demand. Focus on values such as:

- **Total AI requests**
- **Total tokens or credits consumed**
- **Estimated cost**
- **Active users** or active teams during the selected period

These numbers help you answer basic questions quickly: how much AI activity happened, how expensive it was, and how widely the usage was spread across your organization.

Grouped views are useful when you need to compare behavior instead of just reading totals. Depending on the reporting options available in your Atloria workspace, review usage grouped by:

- **Project**
- **Team member**
- **Model**
- **Feature area**

This makes it easier to spot whether one project is driving most of the cost, whether one user is generating unusually heavy traffic, or whether a specific feature is responsible for a usage increase.

Trend charts help you judge whether activity is normal or unusual. A steady line usually points to predictable day-to-day use. A sudden spike can indicate a large documentation run, repeated retries, or a short-term burst of support activity. Low-activity periods may simply reflect quiet workdays, but they can also help you confirm that a rollout or team change reduced usage as expected.

If you need a broader introduction to usage reporting before reviewing activity patterns, see [Monitoring AI Usage and Request Activity](#).

Filtering request history to find the activity you need to review

Once you move from summary metrics into detailed review, open the request history view connected to your AI usage reporting area. This is where you narrow a large activity list into the exact requests you need to inspect. Begin with the same **date range** used on the dashboard so your detailed review matches the totals and charts you just examined.

Use the available filters to reduce the list. In Atloria, common review filters include:

- **Date range**
- **User**
- **Team**
- **Project**
- **Model**
- **Status**
- **Feature area**

These filters are especially helpful when you are investigating one issue at a time. For example, if a team lead wants to review failed requests from a single project, you can combine **Project** and **Status** filters to isolate only those records.

Search tools help when you already know what you are looking for. Use the search field or column-level filtering to find a specific prompt, a response snippet, a request identifier, or a user session. This is often the fastest way to locate one request mentioned in a support conversation or internal review.

Sorting changes the order of results so you can match the list to your goal:

1. Sort by newest activity when you are checking recent behavior after a spike.
2. Sort by highest cost when you want to identify expensive requests first.

3. Sort by failed or incomplete activity when you are investigating reliability issues.

If your Atloria workspace supports reusable filtered views, save combinations you use often, such as a weekly failed-request review or a monthly project cost review. Shared filtered views are useful for support leads and administrators because everyone can return to the same scope without rebuilding the filters each time.

[SCREENSHOT: Request history screen showing date range, user, project, status, and model filters]

Inspecting individual requests to review quality, cost, and outcomes

After narrowing the request list, open a single request to review its details. Depending on your Atloria workspace, this may appear in a side panel or on a full detail page. Start with the top-level fields so you understand the context before reading the content itself. The most important fields to confirm are the **timestamp**, **requesting user**, **project** or workspace, **model used**, and **request status**.

A request record usually includes both the prompt and the response. Review these together rather than separately. The prompt shows what the user asked Atloria to do, while the response shows what was returned. If your organization uses privacy controls, you may see signs that some content has been limited for review. Typical signs include redacted sections, shortened text, or controls that let you expand truncated content. When that happens, review what is visible first and avoid assuming the missing text indicates a broken request.

The usage section is where cost and performance become clearer. Look for fields such as:

Field	What to review
Input tokens	How large the prompt or submitted content was
Output tokens	How long the generated response was
Latency	How long the request took to complete
Retries	Whether Atloria had to try again before finishing
Estimated cost	The likely cost of that single request

Outcome details help you judge reliability. A **Success** result usually means the request completed normally. **Error** or **Timeout** results point to failed or incomplete processing. If you see fallback model usage, that can explain why a request

completed with different cost or quality than expected. If a request appears abandoned, review whether the user stopped the workflow before the response finished.

When you need to compare request-level findings with broader patterns, return to [Monitoring AI Usage and Request History](#).

Using activity patterns to manage spend and team behavior

Request review becomes more valuable when you compare patterns over time. In Atloria, grouped reporting lets you compare usage by **user**, **team**, and **project** so you can see who is driving the highest request count or token consumption. This is often more useful than looking at organization totals alone, because one expensive project or one heavy user can change the overall picture.

Start by comparing high-volume users and projects against the work they are expected to perform. A project with frequent documentation generation or support-agent updates may naturally produce more requests than a smaller team space. What matters is whether the usage pattern matches the work being done. If it does not, inspect the request history for repeated prompts, unusually long responses, or frequent retries.

Certain patterns deserve closer attention:

- **End-of-week spikes** can reflect scheduled content generation or review cycles.
- **Repeated retries** can point to unclear prompts or unstable workflows.
- **Very long responses** can increase cost quickly, especially when they happen often.
- **Heavy use from one project** may suggest a team needs better guidance on when to run AI-assisted tasks.

Model comparisons are also important. If one model consistently produces higher cost with little visible improvement in outcome, team leads may decide that a lower-cost option is enough for routine work. On the other hand, a higher-cost model may be justified for complex generation, sensitive review tasks, or high-visibility documentation output.

Use these findings to make practical decisions. You might update prompt guidance, review shared templates, or set clearer expectations for when teams should use AI-assisted features. This is not only about reducing spend. It also helps improve

consistency, reduce failed attempts, and make AI usage easier to explain during regular operational reviews.

Sharing findings with administrators, managers, and support leads

After you identify the activity you want to discuss, share the same scoped view so others can review the exact data you used. In Atloria, this usually starts by exporting the filtered request history or a usage summary from the reporting area. Before exporting, confirm the **date range**, **project**, **team**, and **status** filters so the report reflects the discussion you are preparing for.

For recurring reviews, use any sharing options available in your workspace, such as a saved dashboard view, a scheduled report, or a copied filtered view. These are especially useful when project administrators, department managers, and support leads need repeatable visibility into the same slice of activity each week or month. A shared filtered view avoids confusion because everyone sees the same scope instead of rebuilding filters manually.

In regular review meetings, keep the focus on a small set of metrics that are easy to compare over time:

- **Request count**
- **Cost trend**
- **Failure rate**
- **Top users**
- **High-cost projects**

These measures are usually enough to show whether usage is stable, growing, or becoming more expensive than expected. If a spike appears, attach a filtered request list so reviewers can move from summary numbers into individual examples without starting over.

A lightweight review rhythm works well for most teams:

1. Review **weekly** failed requests and unusual activity spikes with support leads.
2. Review **monthly** cost and volume trends with administrators or managers.
3. Run an **ad hoc** review after any unexpected increase in spend or request failures.

If your review needs to connect AI activity with broader admin reporting, see [Monitoring Administrative Analytics and Activity](#).

Resolving common issues when usage data looks incomplete or unexpected

If the usage dashboard or request history does not look right, begin with the filters. The most common reason for missing activity is the selected **date range**. A narrow date window can make it appear that no requests were recorded when the activity actually happened outside the current range. Also check whether a **project**, **team**, or **user** filter is still active from an earlier review.

If no requests appear at all, confirm that you are viewing a scope you are allowed to see. In Atloria, some users may only see activity for specific projects or teams rather than organization-wide usage. If your expected records are missing, compare your current view with the permissions and workspace access available to you.

When totals do not match what you expected, check how the dashboard is grouped. A view grouped by **user**, **project**, or **model** can change how totals are presented. Also confirm whether failed requests and retried requests are included in the numbers you are reviewing. A cost total that includes retries may look higher than a simple completed-request count suggests.

Sometimes a request detail opens without full prompt or response text. That does not always mean the record is broken. Atloria may limit what you can see because of retention settings, redaction rules, or privacy controls. If the request metadata is present but the content is shortened or hidden, review the visible status, model, timing, and usage details before deciding whether further follow-up is needed.

If cost appears unusually high, inspect the most expensive requests first. Look for:

- Large prompt sizes
- Very long responses
- Repeated retries
- Use of a higher-cost model
- Sudden increases in request volume

These checks usually reveal whether the increase came from more activity, larger requests, or a more expensive model choice.

Overview

This guide focuses on the review side of AI activity in Atloria. It assumes you already know how to reach the reporting area and read the basic usage screens. Here, the goal is to move beyond simple monitoring and into practical review work: narrowing

request history, opening individual records, comparing activity patterns, and sharing findings with other decision-makers.

Use this guide when you need to answer questions such as:

- Which projects are driving the most AI usage?
- Why did estimated cost increase this week or month?
- Which requests failed, timed out, or retried?
- Are certain users or teams relying heavily on one model or feature area?
- Does a spike reflect normal work or something unusual?

The sections above walk through the main review flow most administrators and team leads follow:

1. Open the **Analytics** area in the **Admin** workspace.
2. Confirm the reporting date range and top-level metrics.
3. Filter request history to isolate the activity you need.
4. Inspect individual request details for quality, cost, and outcome signals.
5. Compare patterns across users, teams, projects, and models.
6. Share filtered findings for weekly, monthly, or ad hoc review.

This document does not repeat the setup and monitoring basics covered earlier in the AI Usage set. If you need a refresher on general usage tracking, start with [Monitoring AI Usage and Request History](#). If you want a broader view of dashboard-based monitoring before doing detailed request review, use [Monitoring AI Usage and Request Activity](#).

[SCREENSHOT: Analytics & Insights page in the Admin workspace with summary cards and trend area]

Prerequisites

Before using this review workflow in Atloria, make sure the following conditions are met:

- You can sign in to Atloria and open the main authenticated workspace.
- You have access to the **Admin** workspace.
- You can open **Analytics** from the admin area.
- You have permission to view the projects, teams, or organization activity you are expected to review.

- AI usage data already exists for the date range you plan to inspect.

It also helps if you already understand the basic navigation patterns for the admin area. If you are still getting familiar with that workspace, read [Using the Admin Workspace](#). If your role includes reviewing user access and visibility, [Managing User Access and Administrative Permissions](#) explains how access scope can affect what appears in reporting screens.

For the most useful review sessions, gather a clear question before you open the dashboard. Examples include:

- Reviewing a recent cost increase
- Checking failed requests for a support issue
- Comparing AI usage across projects
- Investigating one user's recent request activity
- Preparing a weekly or monthly usage summary

You do not need to prepare external tools to complete the steps in this guide. Everything described here is based on the reporting and request-review views available inside Atloria. If your team also reviews security-related records alongside AI activity, pair this guide with [Reviewing Security and Audit Controls](#).