Troubleshooting

What a given error means and what to do about it.

- Nothing appears in Activity
- A tool stopped working after connecting
- Sign-in and passkey problems
- Endpoint Suite problems
- Policy is not doing what I expect
- SIEM and alerts are not arriving

Nothing appears in Activity

The tool works, the developer is using it, and the feed stays empty.

Check the restart first

This is the answer roughly four times out of five. Every one of these tools reads its configuration once, at startup. Closing the window is not enough — the application must be fully quit and reopened, and a terminal must be closed and reopened.

Then check the tool is actually governed

In the Endpoint Suite, the tool's card should say **routed through Sentilai**, not **detected**. If it says detected, the connect step did not complete — try it again and read the error.

Check the device is not revoked

Diagnostics in the console. A revoked device's tools fail closed; they do not silently revert. If the developer says "it stopped working entirely", this is more likely than an empty feed.

Check the filters

Activity's filters persist. A developer filter left set to somebody else, or a time window of 24 hours when the request was yesterday, produces an empty feed that looks like a failure.

Check for drift

A **Non-compliant** badge on the device in Diagnostics means a governed tool no longer points at Sentilai — its configuration was changed outside the app. With automatic governing on, it repairs itself on the next sweep.

Cursor and Copilot special cases

- **Cursor**: Tab autocomplete is never routed. Only chat appears.
- **Copilot in VS Code**: inline completions are never routed. Only chat appears, and only if the custom endpoint was set up manually.

A tool stopped working after connecting

It worked before governance and now returns an error. Almost always one of the things below.

No provider key — error 424

The tool is in **managed API key** mode and there is no matching key. Claude Code needs an Anthropic key; Cursor and Copilot need OpenAI.

Fix it on **Providers**, or move that tool to subscription mode in **Policy ? Connection modes**.

Ambiguous provider key — error 409

Two keys of the same type both list the requested model in their enabled models. Sentilai will not guess which one to bill. Make the model lists disjoint.

Cursor needs its model selected

After connecting Cursor, open its **Settings ? Models**, make sure the OpenAI key is enabled and a `gpt-…` model is chosen. Otherwise Cursor quietly uses its own models and chat may behave oddly.

Copilot BYOK is not enabled on your GitHub organization

If VS Code does not offer a **Custom Endpoint** option at all, that is a GitHub organization setting. Sentilai cannot enable it.

An ungoverned agent got the device blocked

If the error reads *"access blocked: an ungoverned AI agent (…) was detected on this device — remove it to restore access"*, nobody revoked anything. Your organization's ungoverned-agent policy is set to block, and that machine has an autonomous AI agent on it.

The developer fixes this themselves: uninstall the agent, and access returns within about half a minute of the next check. No ticket, no admin action. **Diagnostics** shows which device and which agent; Ungoverned-agent policy explains the setting.

It is a policy block, not an error

Check **Activity** and filter by outcome **Blocked**. A request stopped by your own policy is not a malfunction. The Signals column names the rule that stopped it.

Sign-in and passkey problems

"This account can't open the admin console"

The identity is valid but has no admin record. Usually **passkey autofill chose the wrong saved account** — common for anyone with a work and a personal identity on the same machine. Sign out and sign in again, choosing deliberately.

Or the person is a developer, who correctly cannot open the console.

The invite link does not work

Invites are single-use and expire. **Users & Teams ? Pending & past invites** shows the status. **Resend** issues a fresh link and kills the old one.

Lost the device with the only passkey

Another admin uses **Resend passkey** on that person's row. If the locked-out person is the only admin, contact us.

Prevent it: register a second passkey on **Account & Profile**. The console warns you until you do.

"Cannot remove your last passkey"

Working as intended. Register the replacement first, then remove the old one.

The session keeps expiring

If you are being asked to sign in repeatedly in a short period, tell us — that is a bug rather than a policy, and we have fixed one like it before.

Note that **Policy ? Device sessions** can legitimately end *device* sessions by idle timeout or maximum age. That affects developers' machines, not your console session.

Endpoint Suite problems

SmartScreen warns on Windows

The installer is Authenticode-signed, but SmartScreen also scores *reputation*, which a young signing certificate earns gradually — so the warning can still appear. Click **More info** and verify the publisher is **BRIGHTOPS SMART SOLUTIONS SOCIEDAD LIMITADA** before **Run anyway**. No publisher shown = not our installer; stop there.

A tool is installed but shows "not detected"

The app looks for tools on your login shell's PATH. Claude Code's native installer puts the binary in `~/.local/bin` (`%USERPROFILE%\.local\bin` on Windows) and **does not add it to PATH** — if `claude` doesn't run in a fresh terminal either, add that directory to your PATH (the installer prints the exact instruction), then fully quit and reopen the Endpoint Suite. It reads the tool list at startup.

The app cannot create a device identity

On macOS, this is nearly always **running it from the disk image**. Drag the app into Applications and open it from there — macOS otherwise runs it from a temporary read-only location where it cannot use the Keychain.

On Windows, it means Credential Manager refused. Signing out of Windows and back in usually clears it; if not, send us a log bundle.

"Can't reach {host}"

The banner names the host. Give that name to your network team — it is normally a VPN that has not come up yet or an egress rule that has not been widened. The app keeps retrying and recovers by itself once the host is reachable.

Buttons are missing

If your organization deploys a managed configuration, it can prevent signing out, disconnecting a tool, or ungoverning MCP servers. The card says the setting is managed by your organization. This is your own policy, not a fault.

Signing out did not ungovern my tools

Deliberate. Tools stay pointed at the Gateway and fail closed. Silently returning a machine to ungoverned direct access because someone signed out would be the worst possible failure for a governance product. An admin removes the device from the console.

Sending us logs

Account ? Export logs, or ask your admin to request them from **Diagnostics**. Event logs only — no prompt content, no secrets, 14-day retention.

Policy is not doing what I expect

An MCP server is blocked that I allowed

Check the precedence. A **per-tool** rule beats a per-server rule. If you allowed the server but blocked one of its tools, that tool stays blocked.

Also: when a single request names several MCP servers, **the most restrictive decision applies to the whole request**. The Gateway cannot partially block one API call. One blocked server in the set blocks it all.

A server I never approved is working

Approval only applies if **Require approval for new MCP servers** is on. Without it, the tenant default applies — and if that is Allow, everything works.

Remember that "reviewed" means *any explicit rule*. There is no separate approve action.

Nothing is being blocked at all

Check the detector actions on **Policy ? Risk detectors**. The defaults are mild by design: secrets warn, credentials and personal data report, lethal-trifecta warns. Report means recorded and not interrupted.

Filter Activity by **Risk events** to confirm the detectors are firing — if findings appear but nothing is blocked, the detectors are working and the actions are set to observe.

Too much is being blocked

Usually the classifier sensitivity is one step too high. Move **High** to **Medium**. The clearest cases are caught at every setting; High mostly adds the ambiguous ones.

Check **Blocked conversations** at the bottom of Policy — you can unblock individually.

A policy change has not taken effect on a developer's machine

Local MCP decisions are cached for as long as **Policy ? Local MCP decision cache** allows. Set it to 0 while you are actively tuning, then put it back.

SIEM and alerts are not arriving

Push delivery stopped

SIEM Export shows delivered count, last success, last attempt and **last error**. Start there — it usually names the cause: an expired certificate, a moved collector, a closed port.

The status badge reads **Active** or **Paused**.

Push was never configured correctly

Private and internal hostnames are rejected — the collector must be reachable from Sentilai, so a `10.x` address will not work. Either publish the collector on a name we can reach, or use the pull API from inside your network.

Send test event delivers a `siem_push_test`. If that does not arrive, nothing else will.

Pull returns nothing

Check the credential has **`audit:read`**, has not expired, and is not blocked by its own IP allowlist. The **Last used** column on API Credentials tells you whether your calls are arriving at all.

Check you are passing the **cursor** back. Without it you re-read the same page forever.

Alerts stopped arriving

Alerts are **fire-and-forget** — delivery is attempted alongside the request and never blocks it, which means an alert can be lost if the destination is briefly unreachable.

Use **Send test** on the channel. If the test arrives and real alerts do not, check the channel's **minimum severity**: "High only" will stay silent through a great deal of genuine medium-severity activity.

If you need guaranteed delivery of every event, that is the SIEM path, not alerts.