All policy settings at a glance

Everything on the **Policy** screen, with its default.

MCP

Setting	Default	Notes
MCP default action	Warn	Allow / Warn / Report / Block. Per-server rules override it
Require approval for new MCP servers	Off	Unreviewed servers get the pending action
Pending action	Warn	Applies only when approval is required
Local MCP decision cache	5 seconds	0–3600 seconds. 0 means ask every time

Risk

Setting	Default
AI risk classifier sensitivity	Medium
Secret scanning	Warn
Credentials in context	Report
Personal data (PII)	Report
Lethal-trifecta enforcement	Warn

Detected secrets are handled by their detector's action regardless of the sensitivity dial.

Data

Setting	Default	Notes
Prompt capture	Off	Redacted before storage; follows retention
Log retention	Set by support	Open a ticket to change it

Devices

Setting	Default	Notes
Idle timeout (hours)	Off	Blank disables
Max age (days)	Off	Blank disables
Ungoverned AI agents	Report it only	Report / Mark non-compliant / Block the device's access
Grace period (hours)	24	Blocking only. 0 blocks immediately

All of these take effect within about 30 seconds, like manual revocation. Blocking on an ungoverned agent is the one setting here that can withdraw a developer's access without an admin acting at the time — see Ungoverned-agent policy before changing it.

Per tool

Connection modes — subscription or managed API key, per tool. Cursor is locked to managed.

Also on this screen

Per-server and per-tool MCP rules, real-time alert channels, the egress blocklist, and the blocked-conversations list.

Revision #7
Created 2026-08-02 10:20:10 UTC by Sentilai Docs
Updated 2026-08-04 08:40:46 UTC by Sentilai Docs