

Mitra administrator settings

Settings → Mitra is where an administrator sees and controls Mitra's autonomous operations. It has seven tabs. This page says what each one shows, so you can find the right one and read it correctly. [What Mitra is](#) explains the idea; this page is the map of the screens.

Mitra here is a rule engine that watches your records and raises Help Desk cases. It is not a chat assistant and it does not hold conversations.

<https://www.youtube-nocookie.com/embed/npuu3LF3IPc>

[Mitra Operations: where an administrator sets the rules \(a tour\)](#)

Before you start

- Mitra must be switched on for your subscription. If it is not, opening the page says "Mitra is not enabled for this subscription."
- Seeing the tabs needs the permission to view Mitra's autonomous operations. Changing anything needs further permissions, and they are granted to a group by an explicit administrator command, not by an upgrade. In the test workspace used for this page, HR Head could see and configure everything below; HR Ops could see the tabs and change scopes and actions but not the master switch or the policies; and the Manager user group had none of these permissions.
- **Do not use the buttons described here on a real workspace just to see what happens.** Some of them make Mitra act.

The seven tabs

They share one tab row under **Settings → MITRA: Autonomous Operations, Policies, Activity, Action Review, Simulation, Engine, Audit.**

Autonomous Operations

The control page. It answers "what is Mitra allowed to do here?". The top shows a status line such as "Mitra is watching 9 scopes and may act in 1 of them", then three stages in order: **Detection**,

Recommendation, Execution.

- **Operational controls:** **Autonomous execution** with **Turn execution off**, **Pause operations** (1 hour, 4 hours, until end of day, until manually resumed) and **Detection** with **Stop detecting**. Beside them are today's counters: detected today, cases raised, escalations, open situations, minutes since the last run.
- **Operational scopes:** nine cards (Assets, Attendance, Compliance, Employee Records, Help Desk, Leave, Offboarding, Onboarding, Payroll). Each has its own **Detection**, **Recommendation** and **Execution** switches, the number of policies in it, and its open situations.
- **Autonomous actions:** eleven actions grouped as **Case management**, **Communication and follow-up** and **Resolution**. The page states: "An action that is off is never performed, even when an enabled policy asks for it."
- **Defaults and safety limits:** **Default Help Desk ticket type**, **Fallback owner group** and **Maximum autonomous actions per run**, with **Save settings**.
- **Explore Mitra:** links to Activity, Simulation, Engine and Audit.

See [What Mitra is](#) for the checklist to read before turning execution on.

Policies

Free trial. Your workspace is fully available.

UpgradeTalk to sales

BlueOneHR

Admin viewEmployee view

Signed in as Ananya Iyer · HR Head

Search modules...

Dashboard

Recruitment

Onboarding

Employee KYC

Employee

Attendance

Leave

Assets

Performance

Offboarding

Geo-Fence

Field Ops

Payroll

HR Letters

Project

Help Desk

Integrations

Settings

Dashboard / Settings / Mitra

MITRA

24

Simulate

Autonomous Operations

A policy is half configuration and half code. The thresholds, severity, ownership and permitted actions are yours to change. The rule that decides whether a situation exists is shipped code, because a mistyped threshold detects the wrong number of things and is obvious, while a mistyped rule detects nothing at all and looks like a quiet week.

GENERAL

BASE

RECRUITMENT

EMPLOYEE

ATTENDANCE

LEAVE

PAYROLL

PERFORMANCE

HELP DESK

CONFIGURATION

MITRA

Autonomous Operations

Policies

Activity

Action Review

Simulation

Engine

Audit

POLICY	CONDITION	SEVERITY	PERMITTED ACTIONS	SITUATIONS STATE
ASSETS				
Asset request awaiting a decision v1	Asset request status is Requested AND it is more than N days old	Low	10 ACTIONS Create Help Desk ticket Add evidence and context Determine category +7 Show every action	1 Enabled
Asset return requested but not completed v1	Return was requested AND no return date is recorded AND the request is more than N days old	Medium	10 ACTIONS Create Help Desk ticket Add evidence and context Determine category +7 Show every action	0 Enabled
Asset still assigned to an inactive employee v2	Assignment has no return date AND the assigned employee is inactive	High	10 ACTIONS Create Help Desk ticket Add evidence and context Determine category +7 Show every action	0 Enabled

The list of the situations Mitra watches for. The workspace used for this page had 24, grouped by scope. Each row gives the policy name and version, its **Condition** in words, its **Severity**, the

Permitted actions, how many **Situations** it has found, and its **State** (**Enabled**, or **Not evaluated**).

A policy that looks for something being absent will not run while the data it depends on is stale. For example, the policy for a missing attendance record on a working day shows **Not evaluated** when the data is stale, with the reason: the newest attendance record is older than "the 3 days this policy needs", and missing punches could not be told apart from a feed that had stopped, so nothing was reported. That is the intended behaviour.

Opening a policy shows its condition, severity, whether it raises new cases, its version and who saved it, the actions that would run **right now** and the ones **asked for but blocked** (with the reason, for example "This action is switched off"), and then the settings that can be changed: thresholds, follow-up and escalation timing, severity, whether a person must always approve, and who owns the cases it raises. Each change is versioned and asks for a reason. Editing a policy needs the policy-editing permission.


Activity

Every situation Mitra has detected, newest first, with its severity, the policy that found it, who it concerns, how many times it has been seen, its status and the case raised. The filters are **Scope**, **Policy**, **Status** (**Detected**, **Recommended**, **Acted**, **Skipped**, **Resolved**, **Failed**), **Action**, employee and **Seen since**. Opening an entry shows the evidence captured at the time and the version of the policy that was in force. It is assembled from stored records and is not recomputed when you open it.

Action Review

Actions Mitra has proposed and is waiting for a person to approve or refuse. The page states the rule: "Mitra proposes; a human decides; the approved action then runs as that human." It has status filters (**Open**, **Proposed**, **Awaiting approval**, **Executed**, **Refused**, **Failed**). When there is nothing to decide it reads "Nothing to review". A person who belongs to a company's Mitra approvers group can see this tab even without the wider permission; it grants sight only.

Simulation



Runs every check and every decision against today's data and changes nothing. You tick the policies to test, or leave all unticked to test every enabled policy, and select **Run simulation**. The page says: "No ticket is raised, no owner is assigned and no case is closed." The result lists what would be created, escalated and assigned, and to whom. Read it before you allow Mitra to do anything. A simulation was not run while writing this page.

Engine

Free trial. Your workspace is fully available.

UpgradeTalk to sales

BlueOneHR

Admin viewEmployee view

Signed in as Ananya Iyer · HR Head

47

Ananya

Search modules...

DashboardRecruitmentOnboardingEmployee KYCEmployeeAttendanceLeaveAssetsPerformanceOffboardingGeo-FenceField OpsPayrollHR LettersProjectHelp DeskInformationSettings

Dashboard / Settings / Mitra / <Page 1 of 84>

+ MITRA

Engine

Whether the scheduled scan is really running, and what every pass of it did.

Run nowSimulateActivity

GENERALBASERECRUITMENTEMPLOYEEATTENDANCELEAVEPAYROLLPERFORMANCEHELP DESKCONFIGURATIONMITRA

Autonomous OperationsPoliciesActivityAction ReviewSimulationEngineAudit

Engine status

Read from the platform's own Celery Beat schedule, not from this app's opinion of it.

SCHEDULE every 900 seconds	LAST COMPLETED RUN 21 Sep 2026, 15:44 8959 ms	NEXT RUN 21 Sep 2026, 15:59	SCHEDULER LAST RAN IT 21 Sep 2026, 15:44	TOTAL SCHEDULED RUNS 2070	OPEN SITUATIONS 493
-------------------------------	---	--------------------------------	---	------------------------------	------------------------

Run history

Every pass of the engine, real and simulated. A simulation evaluates everything and writes nothing.

Real runsSimulationsAll

STARTED	TRIGGER	DURATION	POLICIES	SITUATIONS	OUTCOME	STATUS
21 Sep 2026, 15:44:41	Scheduled	8959 ms	23	278 0 new	0 acted 0 recommended 0 no change 1832 skipped 0 failed	Completed

No attendance on a scheduled working day was not evaluated. The most recent attendance record is from 06 Sep 2026, 15 days ago, which is beyond the 3 days this policy needs. Missing punches cannot be told apart from a feed that has stopped, so nothing is being reported.

215 open detections could not be re-checked because the records they refer to no longer exist. They were left open.

Whether the scheduled scan is really running. It says it reads the platform's own scheduler, not the app's opinion of it. It shows the state (**Running**), the schedule (for example **every 900 seconds**), the last completed run and how long it took, the next run, when the scheduler last started it, the total scheduled runs and the open situations. Below is the **Run history**, real runs and simulations, each with its trigger, duration, number of policies, situations found, outcome counts (acted, recommended, no change, skipped, failed) and status. Run notes appear under a run, such as a policy that was not evaluated. **Run now** starts a real run and **Simulate** opens Simulation.

Audit

Every configuration change and every autonomous decision, with who made it, when, and what changed from and to. The filters are **Everything**, **Configuration changes** and **What Mitra did**. The page says entries are written as they happen and are never edited.

What you should see

Each tab opens on its own address under `/mitra-operations/`, marks itself in the tab row, and states in its subtitle what it is for. Counters and run times change with every scan, so treat any number, including the ones in the captures, as one moment in time.

If it does not work

- **You see "You do not have permission to manage Mitra's autonomous operations."** You lack the view permission. Ask an administrator.
- **You can read a tab but its buttons do nothing or are missing.** Viewing and changing are separate. Changing scopes, actions, policies and the master switch each need their own permission.
- **A policy says Not evaluated.** That is the guard described under Policies, not a fault.
- **The Engine tab notes that some open detections could not be re-checked because the records they refer to no longer exist.** They are left open on purpose.

Not confirmed

- **Toggling, saving or running** anything was not tried live. In particular a simulation and **Run now** were not run.
- **The detail pages** of an Activity entry and of an Action Review proposal were not tried.
- **Which other roles** can see or change each tab, beyond what the sample user groups showed.

Related

- [What Mitra is](#)
- [The Control Center dashboard](#)
- [Help Desk](#)

Created 2026-09-21 15:15:50 UTC by Docs Publisher (CI)

Updated 2026-10-03 03:28:47 UTC by Docs Publisher (CI)