



SYNAPSE CORTEX

User Guide

Day-to-day ticket, asset, and AI Remediation workflows for agents and requesters.

Version **2.0.0** · **Confidential** · 6 July 2026

Synapse IronNode product family

This guide covers day-to-day use of Synapse-Cortex v2 for agents and requesters: logging in, working tickets, browsing assets, and using the new AI Remediation panel.

1. Logging In

Go to your organization's Synapse-Cortex URL and sign in with the email and password your admin gave you.

If you have MFA enabled, after entering your password you'll be asked for a 6-digit code from your authenticator app before you're let in.

If your admin has required MFA for your account but you haven't set it up yet, you'll be sent straight to the enrollment screen the moment you log in — you can't reach the dashboard until you finish it. To enroll:

1. Open your authenticator app (Google Authenticator, Authy, 1Password, etc.) and scan the QR code shown, or type in the manual entry code below it.
2. Enter the 6-digit code your app is currently showing and click **Confirm & Enable**.

You can also turn MFA on yourself any time from **Settings** → **Profile & MFA**, even if it isn't required.

2. The Dashboard

The dashboard opens with an **AI Enabled** or **AI Disabled** badge at the top — this reflects whether your tenant's admin has turned on the AI Remediation Module (Section 6). Below it are four **queue cards**:

- **Unassigned** — open tickets nobody is working yet.
- **Awaiting User** — tickets waiting on a reply from the requester.
- **SLA Breaching** — open tickets past their resolution deadline. This card glows red when it's non-zero — check it first.
- **Resolved Today** — tickets closed out today.

Click any card to jump straight to the Tickets list, pre-filtered to that exact set — no need to build a filter yourself.

Below the cards is a table of the most recent tickets tenant-wide.

3. Working Tickets

Searching and sorting

Above the ticket table is a search box and a sort dropdown. Type one or more keywords and click **Search** — it matches against the ticket number, title, and description, and requires every keyword you type to appear somewhere (so "printer lobby" narrows to tickets mentioning both words, not either one). Combine it with a dashboard queue-card filter if one's active. The sort dropdown reorders the current list by newest, oldest, priority, or status, and re-applies immediately when changed. Click **× clear search** to drop the keyword filter.

Creating a ticket

From **Tickets**, click **+ New Ticket**. Fill in a title, optional description, type (Incident or Request), priority, and optionally link it to an asset or assign it to someone. Priority determines the SLA clock — Critical tickets have the tightest deadlines. Linking an asset now also matters for AI Remediation (Section 6): the asset's discovered OS is what the AI checks a playbook's target-OS guardrail against.

Viewing and updating a ticket

Click any **ticket reference** (e.g. TCK-000042) — in the dashboard table, the Tickets list, or an asset's linked-tickets panel — to open its detail page. From there you can edit:

- **Title and description** — the main text fields at the top.
- **Status** — New, In Progress, Awaiting User, Resolved, Closed.
- **Priority** — changing this recalculates the SLA deadline based on the new priority's target.
- **Assigned To** — pick any agent, tenant admin, or global admin in your organization.

Click **Save Changes** — each field updates independently, and you'll see a confirmation banner when it's saved. The sidebar also shows the linked asset (if any), requester, resolution deadline, and whether the ticket is currently breaching SLA.

Logging actions taken

Below the title and description is an **Actions Taken** section — a running, timestamped log of what's actually been done on the ticket, separate from the description. Type what you did into the box and click **Log Action**; it's recorded with your name and the current date/time, newest entry first. There's no edit or delete for a logged action — it's a permanent record, so write it once it's actually done. Anyone who can view the ticket can see its action history, which is useful for handoffs between agents. If AI Remediation is used on the ticket (Section 6), its proposals and outcomes are recorded here too, alongside your own manual entries.

If you're a `global_admin` or `tenant_admin`, you'll also see a **Delete Ticket** button on this page. Deleting a ticket is permanent (it removes its action history and remediation history too) and can't be undone.

Deleting multiple tickets at once

Admins also get a checkbox in the leftmost column of the Tickets list — check the box on any tickets you want to remove (or the header checkbox to select every ticket currently shown), then click **Delete Selected** and confirm. This is just as permanent as deleting one ticket at a time from its detail page.

4. Browsing Assets

Assets (CMDB) lists every registered device: its Asset ID, name, type, MAC/IP, status, location, and where it came from (**Manual** or **NetscanXi**).

Click any **Asset ID** to open its detail page, which is organized into three independently-editable sections:

- **Core Info** — name, type, status, location, and who it's assigned to.
- **Network** — MAC and IP address.
- **Discovered OS & Software** — operating system and installed software. For assets imported from NetscanXi, you'll also see the raw device-type label NetscanXi reported, even if Cortex categorized it differently. The **OS** field here is what AI Remediation checks a playbook's target-OS restriction against, so keeping it accurate matters if your tenant uses that module.

Each section has its own **Save** button, so you can update one part without touching the others. Any tickets linked to the asset are listed at the bottom — click through to jump straight to them.

To register a new asset yourself, click **+ Register Asset** from the Assets list and fill in at least a name and MAC address (MAC is required for manually-created assets; it must be unique within your organization). Assets imported automatically from NetscanXi don't have this restriction, since some devices are discovered without a usable MAC address. The options in the **Type** dropdown are set by your admin (**Admin** → **Asset Types**), so the list you see may differ from another organization's.

5. Settings

Settings → **Profile & MFA** shows your name, email, role, and tenant, lets you change your password, and manage your own MFA enrollment (see Section 1).

6. Using AI Remediation on a Ticket

This section only applies if your tenant admin has turned the AI Remediation Module on (check the badge on the Dashboard). When it's on, every ticket detail page has an **AI Remediation** panel below Actions Taken.

Step 1 — Link a credential (if the ticket needs execution)

If you want an AI-selected playbook to actually be able to run against a host, the ticket needs a linked credential first. In the panel's **Linked Credential** dropdown, either pick an existing credential (shown as `label (username@host)` — the password itself is never displayed) or choose **+ Vault new credential...** to add one on the spot: fill in a label, username, host, port, and the secret, then click **Vault & Link**. It's linked to this ticket automatically the moment it's created, and stays available to link to other tickets later from the same dropdown.

You can skip this step if you just want to see what the AI would propose — the investigation itself doesn't need a credential, only the actual execution does.

Step 2 — Investigate

Click **Investigate with AI**. Claude reviews the ticket's title, description, priority, and linked asset, and picks the single best-matching playbook your tenant currently has enabled — or decides none of them fit. It never invents or writes commands itself; it only selects from what your admins have already built and approved as playbooks.

Step 3 — Review and respond

What you see next depends on what happened:

- **Blocked** — the playbook the AI picked failed a safety check (wrong target OS for the asset, or a forbidden command) and was stopped automatically before anything could run. No approval is possible here — talk to your admin if you think a playbook's guardrails need adjusting.
- **Pending Approval** — this is the normal case. You'll see the AI's reasoning, the playbook it selected, a summary of its guardrails (allowed OS, approval level, forbidden-command count), and the exact ordered list of commands it's proposing. Read it, then:
 - Click **Approve & Execute** to run it for real. The panel will show "Executing..." briefly, then the completed execution log (each command with its exit status) and a plain-language outcome summary.
 - Click **Reject** if it's not the right call — nothing runs, and the rejection is recorded with your name.
- **Succeeded / Failed** — some playbooks are configured by a global admin to skip the approval click entirely (Auto-Approve); if the AI selects one of those, you'll see the completed result directly, with the same proposal and execution details as the approved case above.

Whatever happens, it's also written to the **Actions Taken** log automatically, so the record survives even if you navigate away. You can click **Investigate with AI** again at any time — for the same issue recurring, or a different problem on the same ticket — and it starts a fresh cycle without touching the history of previous ones.