CyberArmor.AI Customer Portal — User Guide

Last verified: 2026-11-21, against the portal's code only, on `main` at commit 3310fef64. About 150 statements were checked; the menu, Mission Control, Endpoints, Onboarding & Downloads, Settings, Proxy Controls, Reports and the sign-in page were brought up to date. Nothing was re-run in a browser for that pass, and the pictures do not show the Simple / Advanced switch, the Protection page, or the AI off switch and Endpoint response cards on Endpoints.

Earlier: 2026-10-29, every section was checked against the portal's code on `main` at commit 1f4497db6. The screenshots were taken on 2026-10-24 on the live Singapore portal (am-app.cyberarmor.ai), release 1.0.13-rc.12, in the demo organization "CyberArmor.AI Asia"; several screens have changed since (details views, Identity / SSO, Settings, Reports). Where a picture and the text differ, the text describes what you will see.

This guide shows you, screen by screen, how to use the CyberArmor.AI Customer Portal. Every picture has red boxes and arrows showing exactly where to click. Numbered red circles match the numbered explanations under each picture.

Which address do I use?

Region	Customer Portal address
Asia and Middle East (including Pakistan)	**https://am-app.cyberarmor.ai**
United States	**https://app.cyberarmor.ai**

Both portals work the same. A few screens can differ while a feature is being rolled out to one region first; where that is so, this guide says "when shown". This guide uses the Asia and Middle East portal. If your organization runs its own CyberArmor.AI server, use the address your administrator or partner gave you.

Prefer to watch? Six training videos, *Customer portal training, chapters 1 to 6*, walk through every view of this portal. They are on the documentation site at **https://docs.cyberarmor.ai/training/**, on **https://cyberarmor.ai/training**, and inside the portal under **Support → Help & FAQ → Training videos** (section 15).

Contents

1. Sign in
2. Find your way around the screen (includes "Every row opens its details")
3. The dashboard (Mission Control), item by item
4. Download the software and get an install token
5. Add policies from a compliance pack (example: Pakistan SBP)
6. Change a policy (includes "Priority decides which policy applies")
7. Create your own policy (includes the card-PIN example)
8. Turn a policy off or on
9. See your devices (Endpoints)
10. See what is happening (Policy Events, Telemetry, Incidents, Shadow AI)
11. Reports and the Audit log
12. Add a colleague (Users)
13. Sign-in security (Identity / SSO)
14. Settings for every device (automated response, private network addresses, local page check, Local Models)
15. Get help and training videos
16. Quick reference: every menu item
17. Every other view, one by one

---

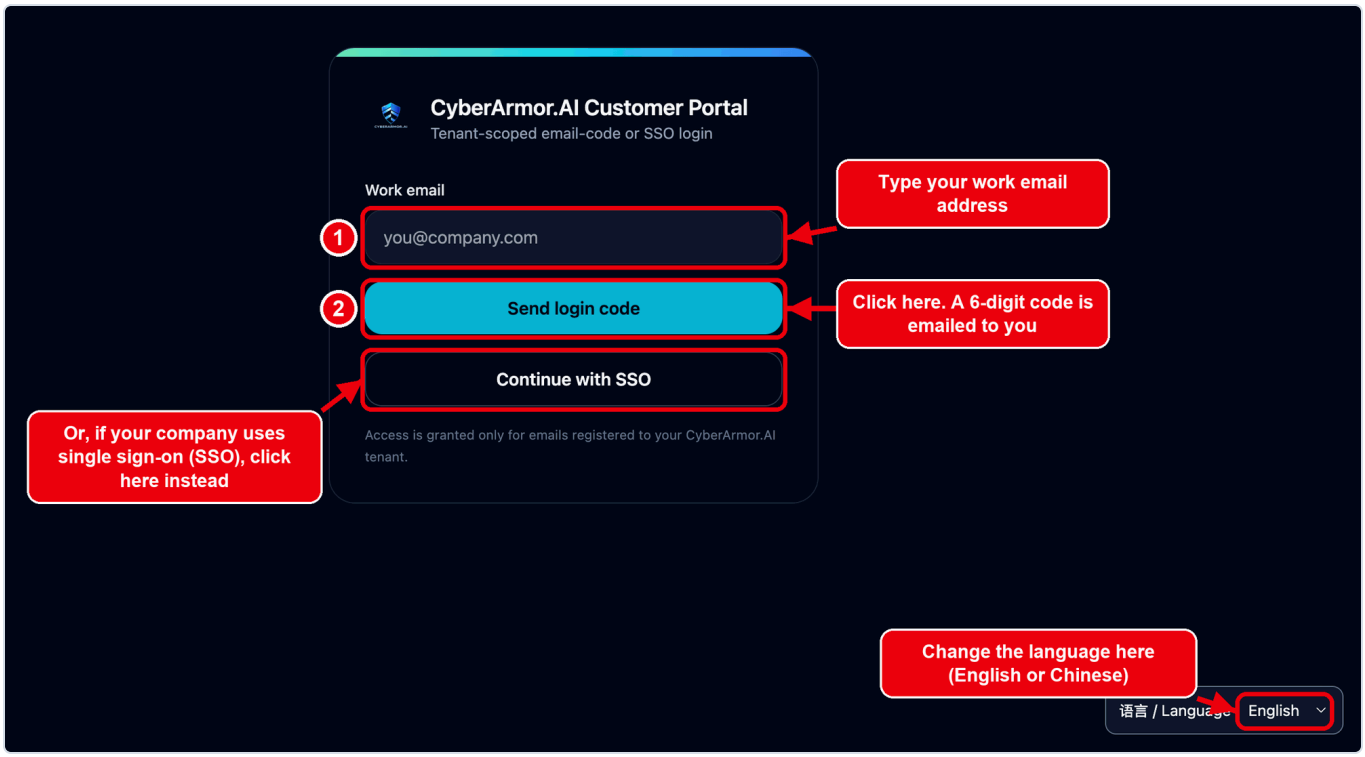
1. Sign in

Step 1. Open **https://am-app.cyberarmor.ai** in Chrome, Edge or Safari.

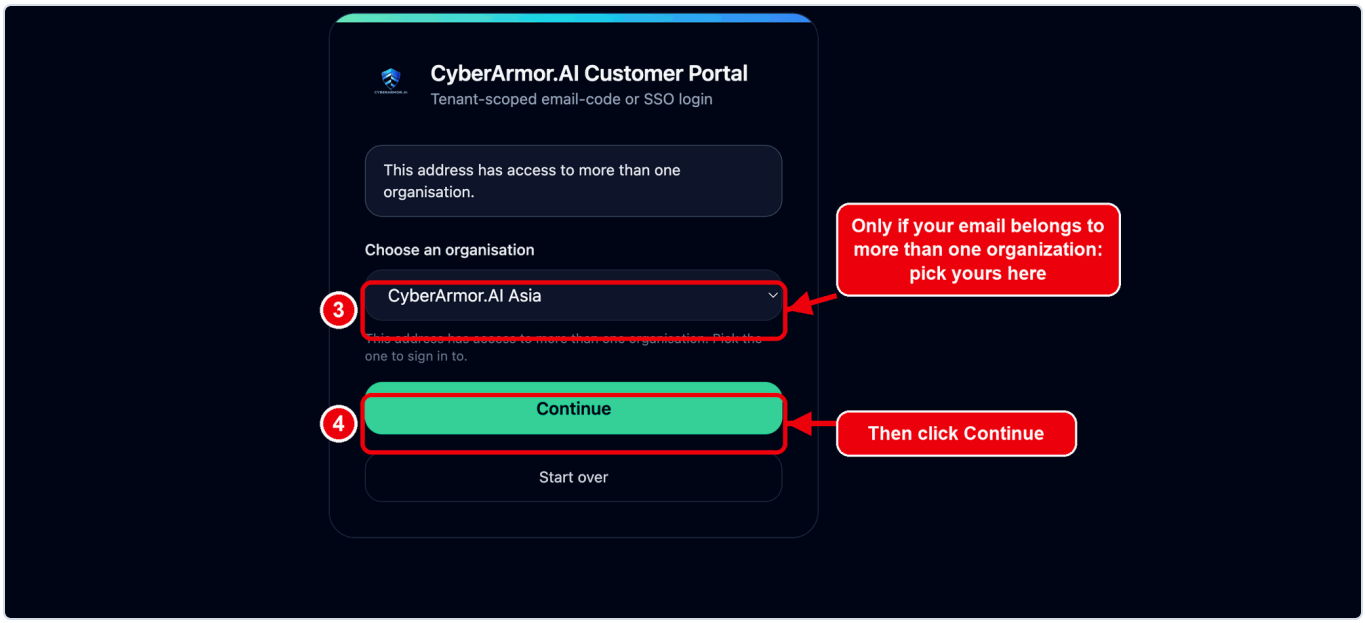
Step 2. Type your work email address (1) and click **Send login code** (2). Within a minute you receive an email with a 6-digit code. Type the code on the next screen and click **Verify and enter**. (**Resend code** sends another; **Use a different email** starts again.) The next screen says which address the code was sent to. **Resend code** is available after 35 seconds; a new code replaces the old one. If you opened the wrong region's portal, click **Not your region? Choose again** at the bottom.

You do not need a password. If your company uses single sign-on (Microsoft Entra ID, Okta and similar), click **Continue with SSO** instead.

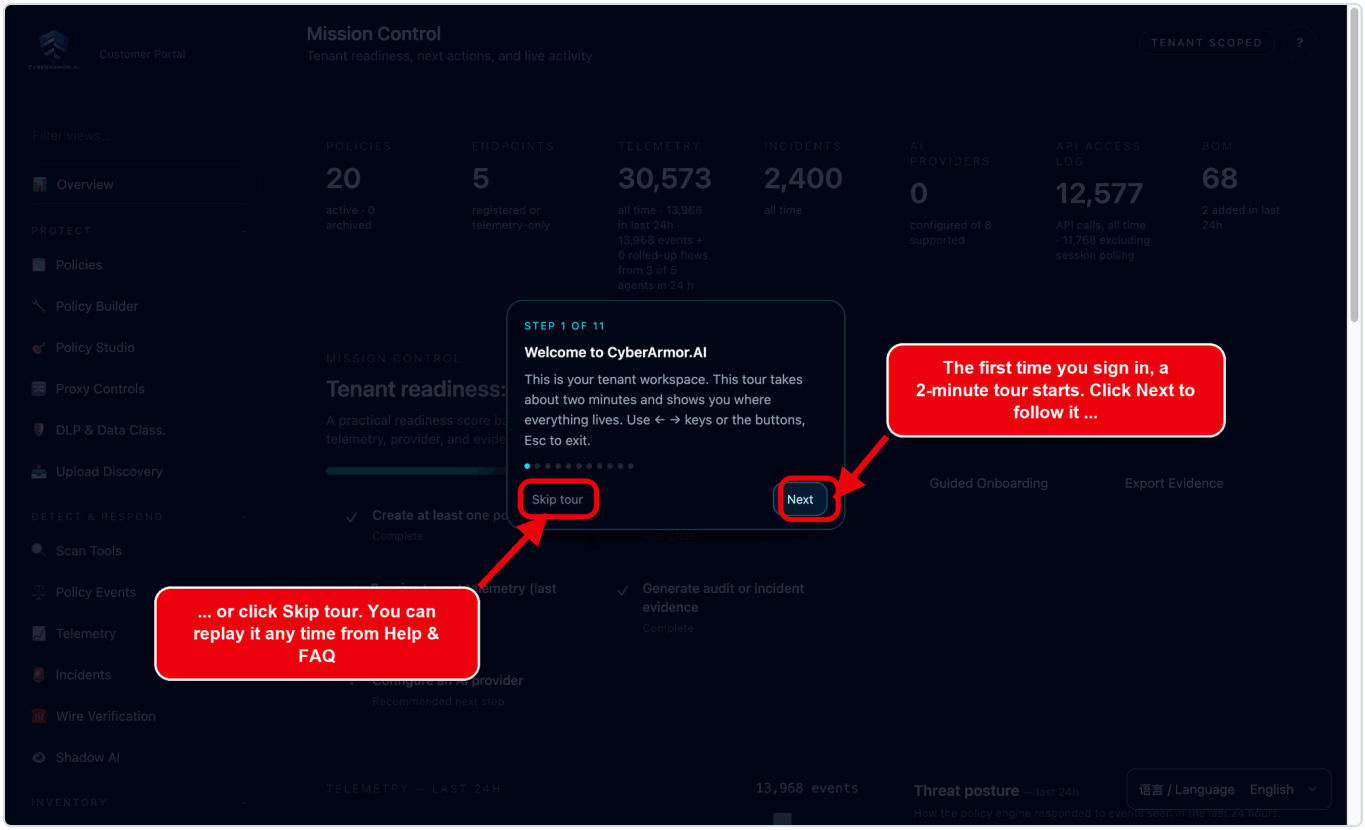
If you have set up an authenticator app (section 13), the portal then asks for the **Authenticator code**: type the 6-digit code from the app, or one of your backup codes, and click **Verify and enter**.



Step 3 (only for some people). If your email address belongs to more than one organization, the portal asks you to **Choose an organisation**. Pick it from the list (3) and click **Continue** (4).



Step 4. If you are the first tenant admin to sign in to a new organization, the portal opens **Protection** and asks you to **Choose a protection level**: three questions, then **Recommend a level** and **Apply this level**. **Not now — keep my settings** skips it. After you apply a level, the menu switches to **Simple** (section 2). The guided tour then starts on your next visit. For everyone else, a short tour starts the first time you sign in. Click **Next** to follow it (about two minutes, 11 steps), or **Skip tour**. You can replay it later from **Help & FAQ**, or from **?** → **Start guided tour**.

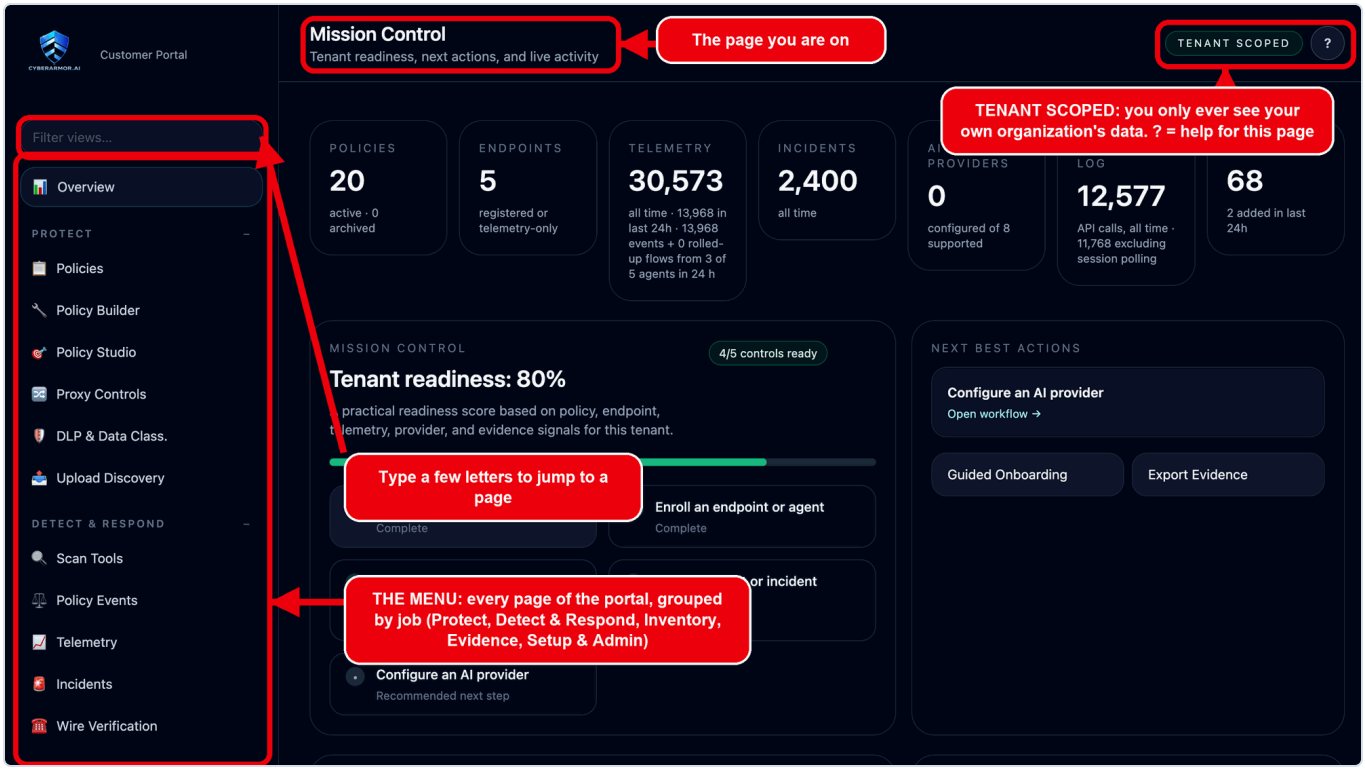


2. Find your way around the screen

Every page of the portal has the same layout:

- **The menu (left).** Every page of the portal, grouped by job: *Protect*, *Detect & Respond*, *Inventory*, *Evidence*, *Setup & Admin* and *Support*. Click a group's title to fold it away.
- **Simple / Advanced (top of the menu). Advanced** shows every page in this guide. **Simple** shows five: **Protection**, **Value & Friction**, **Onboarding & Downloads**, **Settings** and **Help & FAQ**. If a page in this guide is missing from your menu, click **Advanced**. The choice is remembered in your browser.
- **Filter views (top of the menu).** Type a few letters, for example `pol`, to find a page quickly.
- **The page title (top).** Tells you which page you are on, with a one-line description under it.
- **TENANT SCOPED (top right).** A reminder that you only ever see your own organization's data. Nobody in another organization can see yours.
- **? (top right).** Opens a menu: **Start guided tour**, **Help & FAQ**, **Documentation** (this documentation site) and **Support portal**.
- **Your organization, your email and your role (bottom of the menu),** with the **Log out** button.
- **Language (bottom right).** Switches the portal between English and ten other languages: 简体中文, Português (Brasil), Español (Latinoamérica), 한국어, 日本語, ไทย, Tiếng Việt, Bahasa Indonesia, Bahasa Melayu and Filipino. The translations are machine translated and say so; English governs. The Asia and Middle East portal opens in 简体中文 the first time; choose **English** here to match this guide.

If a partner manages your organization and opened your portal for you, a **Partner session** banner across the top says so, with ← **Back to Partner Console**.



Every row opens its details

Click any row in any list — an event, an incident, a device, a policy, a user, an artifact, an API key, a SIEM connector — and a details window opens with every field of that record in plain words. Click anywhere on the row except its own buttons; with the keyboard, move to the row and press Enter.

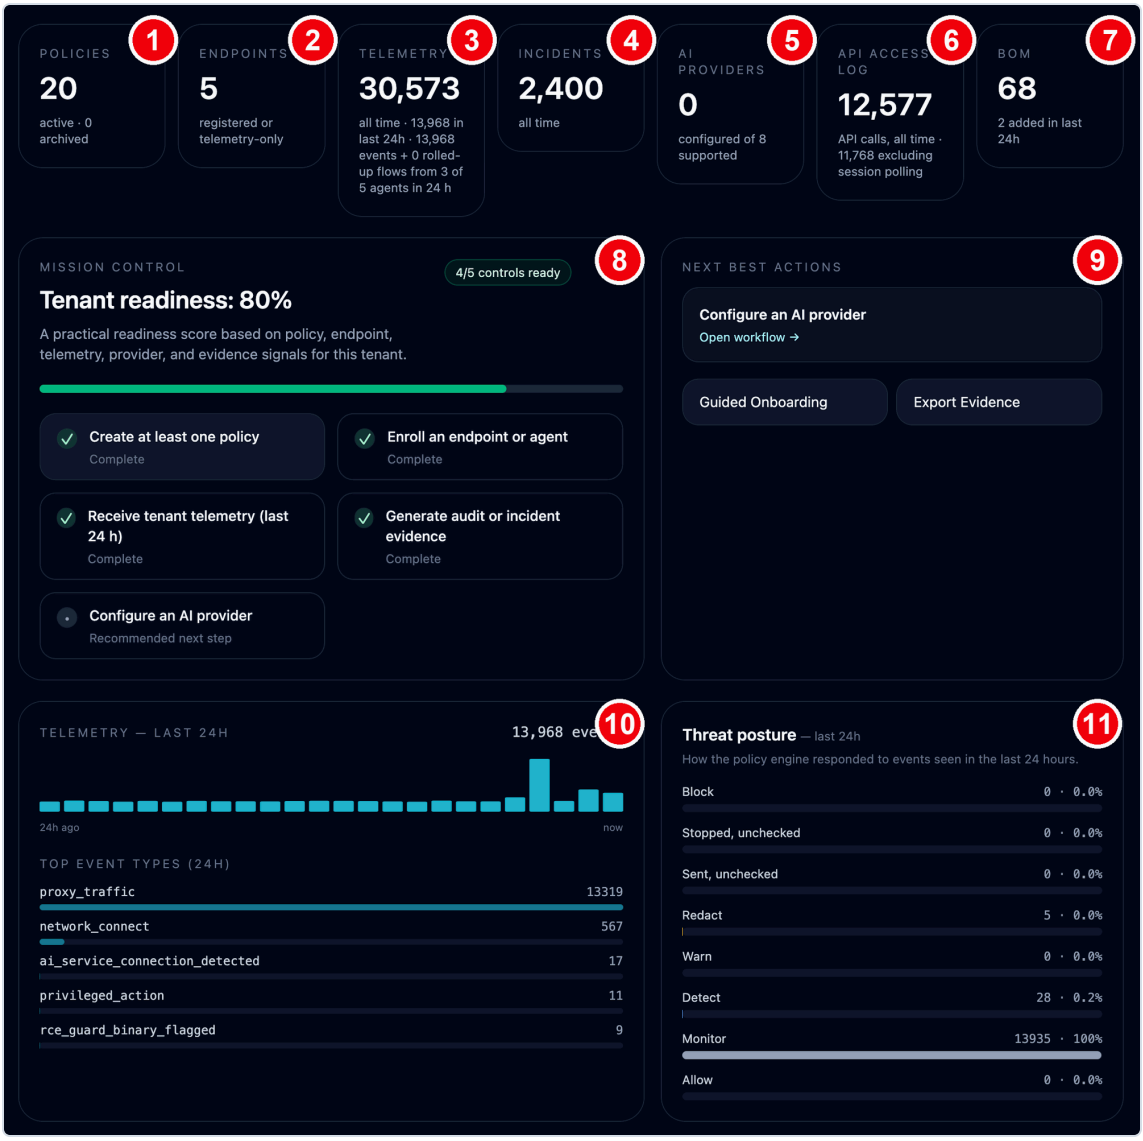
- At the bottom of the window, ▸ **View raw JSON** opens the record exactly as the server sent it, for your technical team. Fields that hold a credential are removed from that copy before it is shown.
- **Policy.** A record that a policy decided names the policy, its action and its priority, and offers **Open policy**, which opens that policy's own details. If the policy has been deleted or archived, or the name could not be checked, the row says so instead of leaving it blank.
- **Content scan.** Says whether the content was scanned and what the scan found, for example "Scanned: no sensitive data found", or "Not scanned — blocked by destination policy". An empty list of findings never stands in for "not scanned".
- **Correlation.** On a record from a browser extension or an agentless sensor, says whether it was joined to an endpoint agent on the same computer: **Correlated** (and how, with the computer's name and operating-system account), **Not correlated** (and why), or **Not checked** (and why).

Two details windows have no **View raw JSON**: a policy's window (its stored conditions are shown in full, and the whole policy is under **Edit → JSON**) and an audit record's window (the original signed record is printed in full).

---

3. The dashboard (Mission Control), item by item

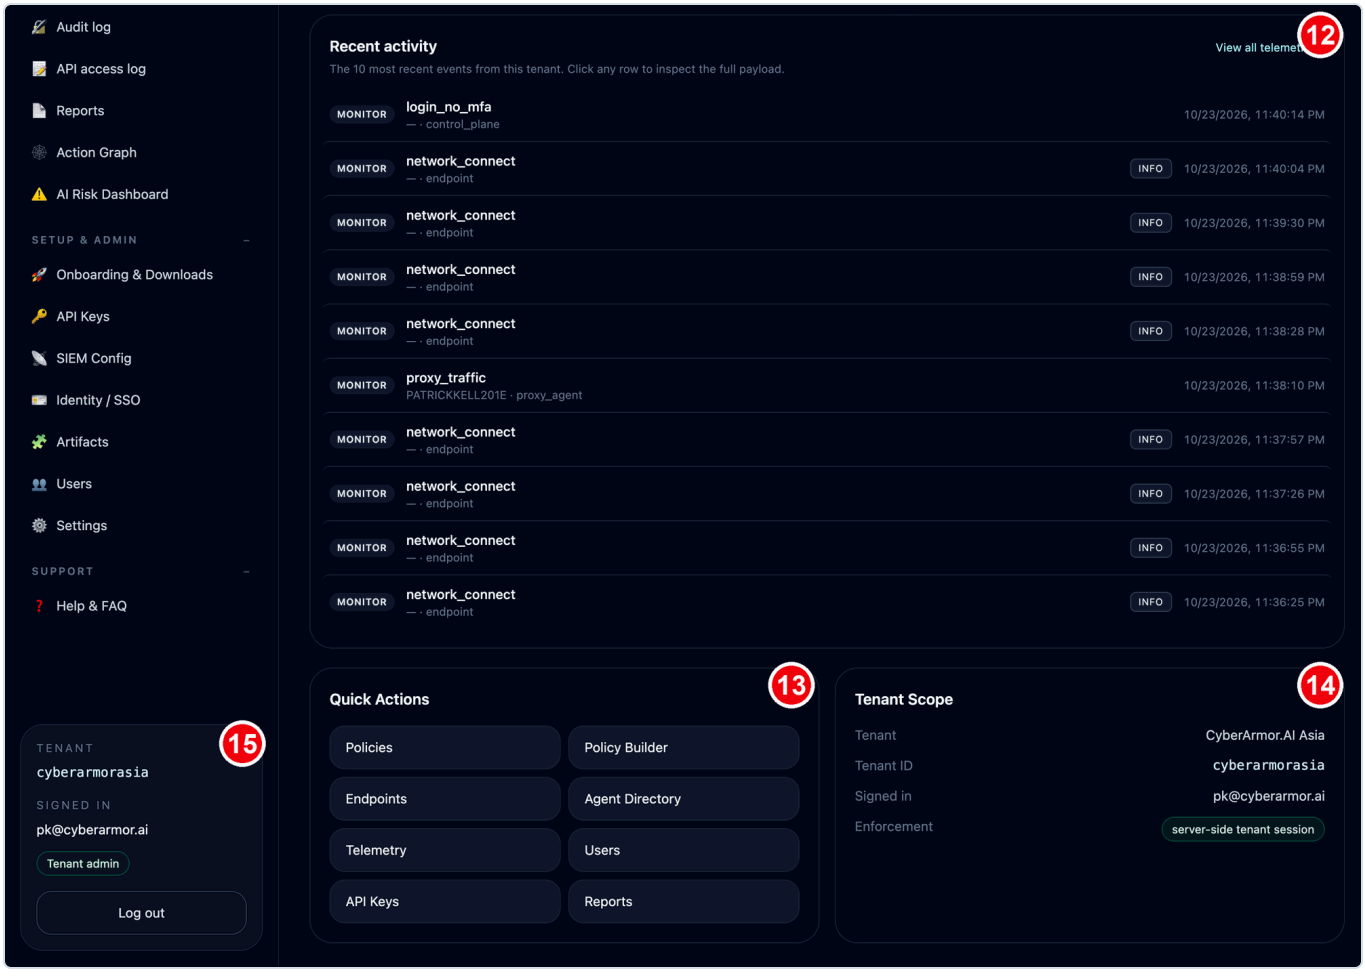
Overview opens **Mission Control**, the page you see after signing in. It answers three questions at a glance: *Is everything set up? Is data coming in? Did anything happen?*



#	What it is	What it tells you
1	**Policies**	How many security rules your organization has, and how many are archived.
2	**Endpoints**	How many computers, phones and browsers are connected to CyberArmor.AI.
3	**Telemetry**	How many activity records CyberArmor.AI has received, in total and in the last 24 hours, and from how many devices.
4	**Incidents**	How many times something was blocked, redacted or flagged as a finding that needs a look.
5	**AI Providers**	How many AI services (for example OpenAI or Anthropic) you have connected for CyberArmor.AI to manage.
6	**API Access Log**	How many times your portal and devices talked to the CyberArmor.AI server.
7	**BOM** (Bill of Materials)	How many AI tools, models and libraries were found on your devices, and how many are new today.
8	**Tenant readiness**	A score of how completely your organization is set up, from a five-step checklist. Each tick is one finished step; each step opens the page that finishes it.
9	**Next Best Actions**	Shortcuts to the most useful thing to do next. **Guided Onboarding** walks you through setup; **Export Evidence** opens **Reports**, where the evidence exports are.

#	What it is	What it tells you
10	**Telemetry, last 24 hours**	A bar for each hour of the last day, and the most common kinds of activity. A tall bar means a busy hour.
11	**Threat posture**	Two parts. **Enforcement decisions** counts every block, redaction and warning in the last 24 hours exactly as **Incidents** counts them, says which store each block was counted from, and links to **Open Incidents** →. Below it, the telemetry rows of the same day by what CyberArmor.AI did: **Block** (stopped by a policy), **Stopped, unchecked** (stopped because it could not be checked; no policy was involved), **Sent, unchecked** (went out without being checked, so look at these), **Redact** (sensitive data was removed before it left), **Warn** (the user was warned), **Detect** (noted as a finding), **Monitor** (recorded only), **Allow** (allowed by a policy). A small share reads "<0.1%", never "0.0%".

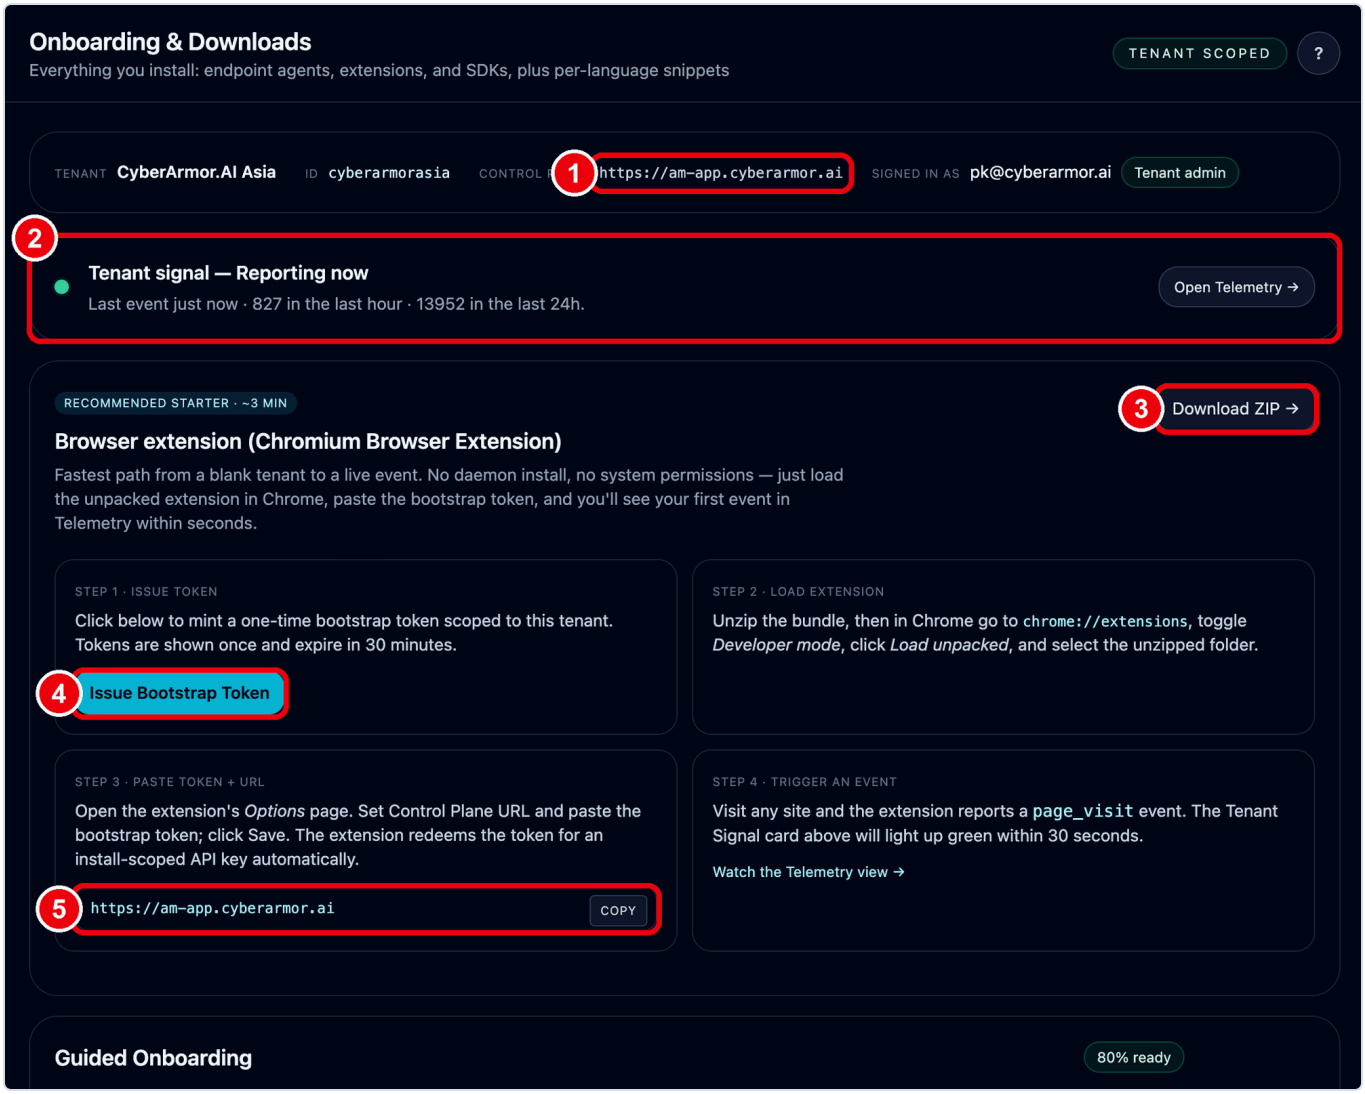
Every view of the portal counts a block the same way and opens on the last 24 hours, so the numbers on Mission Control, Telemetry, Incidents, Policy Events, Agent Directory, Action Graph and the AI Risk Dashboard agree, and each says which store a number came from.



#	What it is	What it tells you
12	**Recent activity**	The 10 newest events. Click any row to see every detail (section 2). **View all telemetry** opens the full list.
13	**Quick Actions**	Eight one-click shortcuts: Policies, Policy Builder, Endpoints, Agent Directory, Telemetry, Users, API Keys and Reports.
14	**Tenant Scope**	Your organization's name and ID, who is signed in, and confirmation that the server enforces the separation between organizations.
15	**Your account** (bottom of the menu)	Your organization ID, your email, your role, and **Log out**.

4. Download the software and get an install token

Everything you install comes from one page: **Setup & Admin** → **Onboarding & Downloads**.

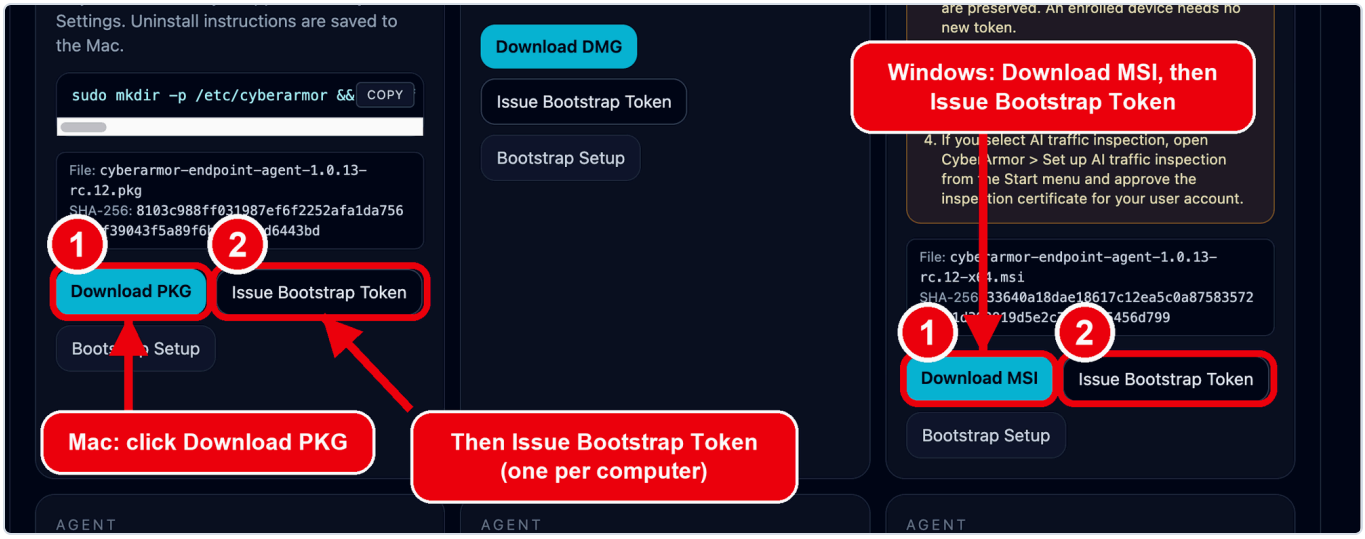


1. **Control plane address.** The portal address your devices connect to. You choose or type it during installation. For Asia and the Middle East it is `https://am-app.cyberarmor.ai`.
2. **Tenant signal.** A green dot means your devices are sending data right now.
3. **Download ZIP. CyberArmor.AI Protect**, the browser extension for Chrome, Edge, Brave and Opera, the fastest way to get started (about 3 minutes).
4. **Issue Bootstrap Token.** Creates the one-time token the extension needs to join your organization.
5. **Copy** the control plane address from here.

Scroll down to the **Endpoint Agent** group for the Mac and Windows installers: **CyberArmor.AI Endpoint Agent — macOS installer (PKG, signed and notarized)** and **CyberArmor.AI Endpoint Agent — Windows installer (MSI, beta)**. The **Browser, IDE & Office** group holds **CyberArmor.AI Protect**, the browser extension, and the **SDKs & RASP** group the developer packages. On each card, click the download button (1), then **Issue Bootstrap Token** (2). Issue one token for each computer.

The **Endpoint Agent** group also has a macOS disk image (DMG), a Windows offline bundle, the **CyberArmor.AI Trust Gateway**, and **CyberArmor.AI for Android — APK** (Android 8.0 or later; the card shows the file's SHA-256 to compare before you install, and this copy and the Google Play copy cannot update each other). The macOS installer is for Apple silicon Macs (M1 and later). For an Intel Mac, use

Download PKG for Intel Macs on the same card when it is shown; Apple menu › About This Mac says which chip you have. The **Browser, IDE & Office** group also has the extension for Firefox and for Safari, the Office add-in, and the extensions for VS Code, Cursor and Kiro. Those three download as a `.vsix` file (**Download VSIX**), and the card gives the command that installs it. Every card also has **Bootstrap Setup**, which explains the three steps.



After you click **Issue Bootstrap Token**, the token appears at the top of the page, shown masked. **COPY** copies the real value; **Reveal** shows it. The install commands on the same card now contain the token, so treat them as private too. The token:

- works only once, for one computer;
- expires after **30 minutes**, so issue it just before you install;
- is private. Do not email it or paste it in a chat.



At the bottom of the page, **Local model pack** lists this organization's on-device model pack for manual and air-gapped installs, once Local Models are on (section 14).

The step-by-step installation for Windows, Mac and the browser extension is in the companion **Installation Guide: Windows, Mac and Chrome**.

---

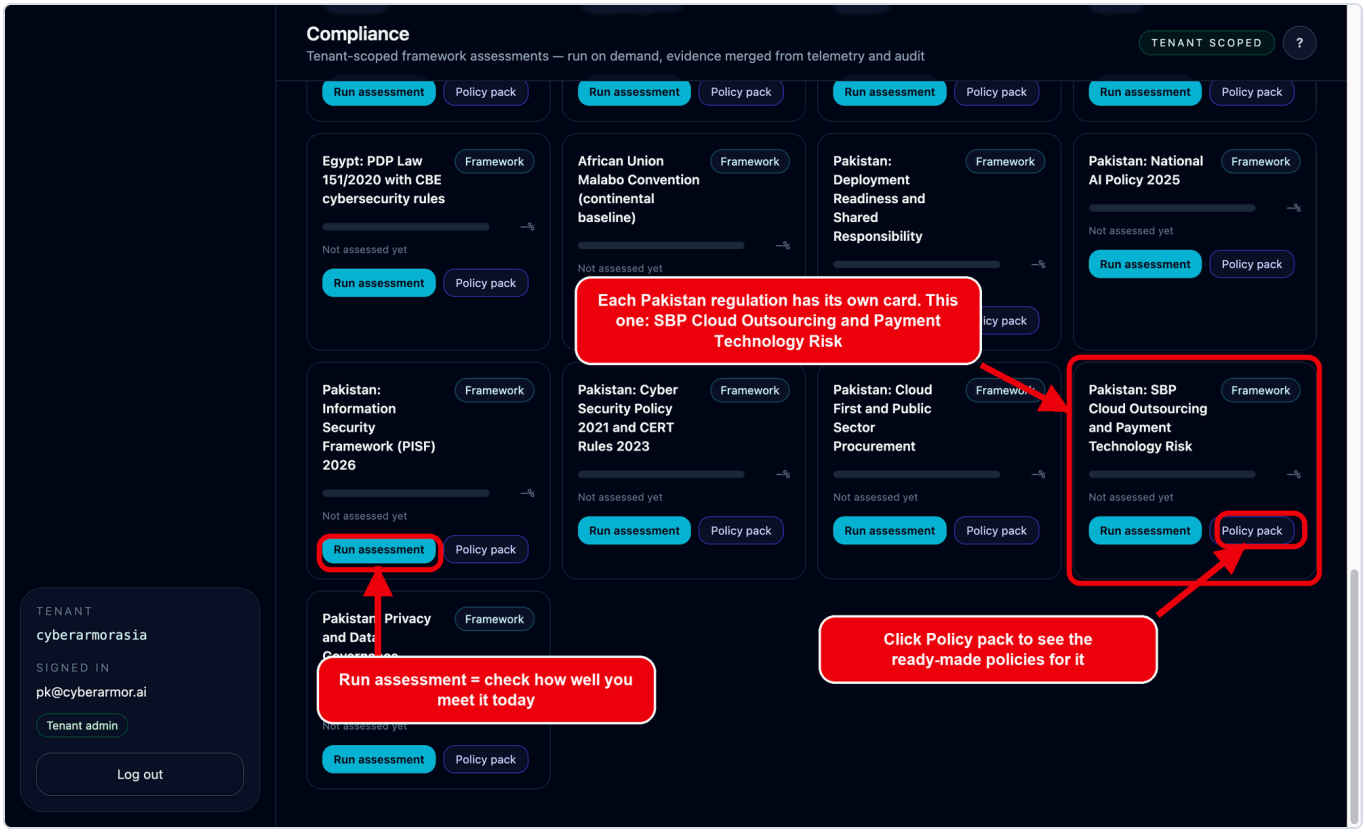
5. Add policies from a compliance pack (example: Pakistan SBP)

A **policy** is a rule, for example "record every conversation with an AI service" or "remove national ID numbers before they reach an AI service". You do not have to write them yourself: the **Compliance** page has a ready-made **policy pack** for most frameworks.

The Compliance page lists 45 frameworks, and 44 of them have a policy pack (SANS/CWE Top 25 has none on purpose: its controls are software weaknesses, not AI-traffic rules). Seven are for Pakistan: Deployment Readiness and Shared Responsibility, National AI Policy 2025, Information Security Framework (PISF) 2026, Cyber Security Policy 2021 and CERT Rules 2023, Cloud First and Public Sector Procurement, **SBP Cloud Outsourcing and Payment Technology Risk**, and Privacy and Data Governance Preparation (Draft).

Step 1. You need to be a tenant admin for this. In the menu, click **Evidence** → **Compliance**. Scroll to the Pakistan cards. On **Pakistan: SBP Cloud Outsourcing and Payment Technology Risk**, click **Policy pack**.

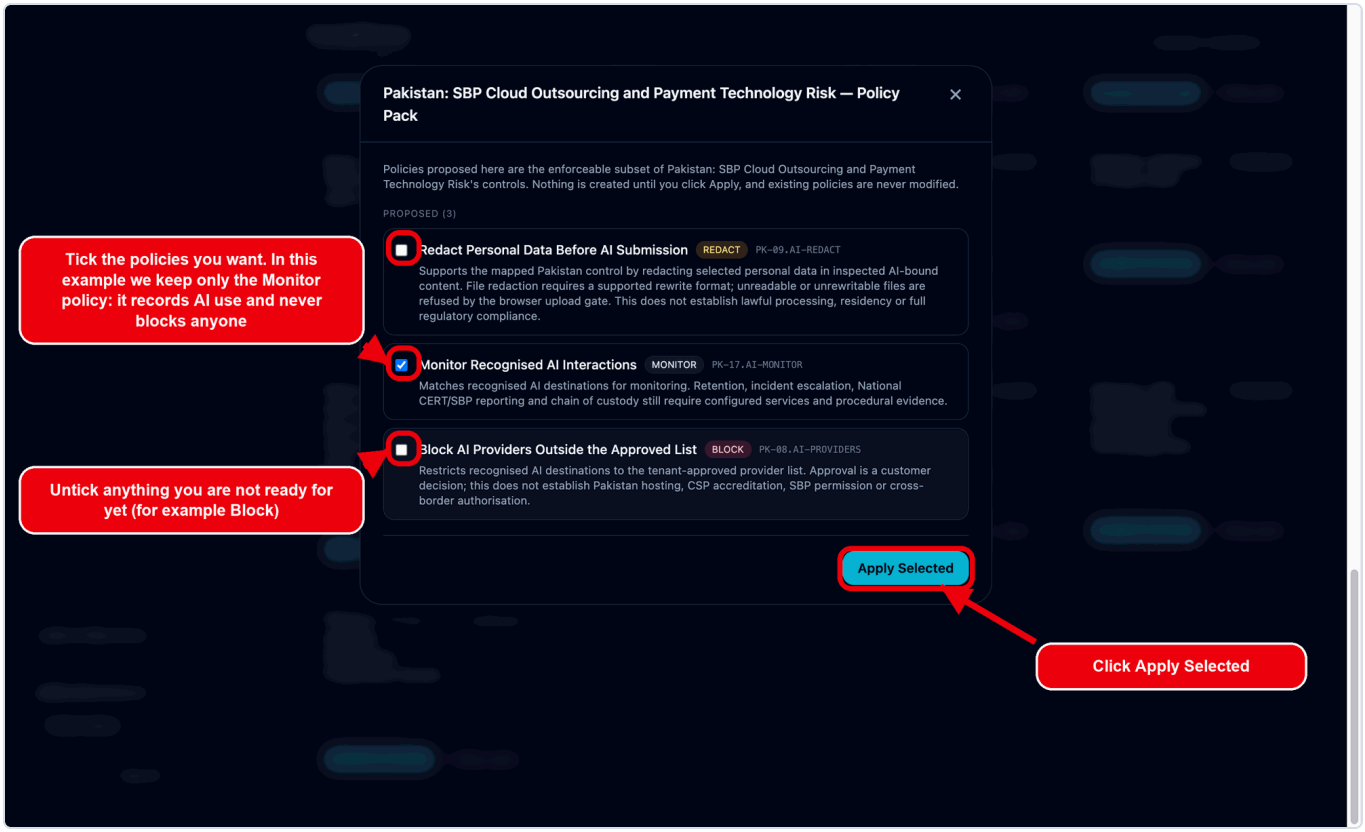
(**Run assessment** on the same card checks how well your organization meets that framework today, using your real activity as evidence, and stores the result; the last stored result loads when you open the page, and **View details** opens it control by control.) Under the summary the page says when the results were assessed ("Assessed 2 hours ago") and whether a refresh is running. If the newest scheduled check could not complete, the card says so and keeps showing the last complete assessment.



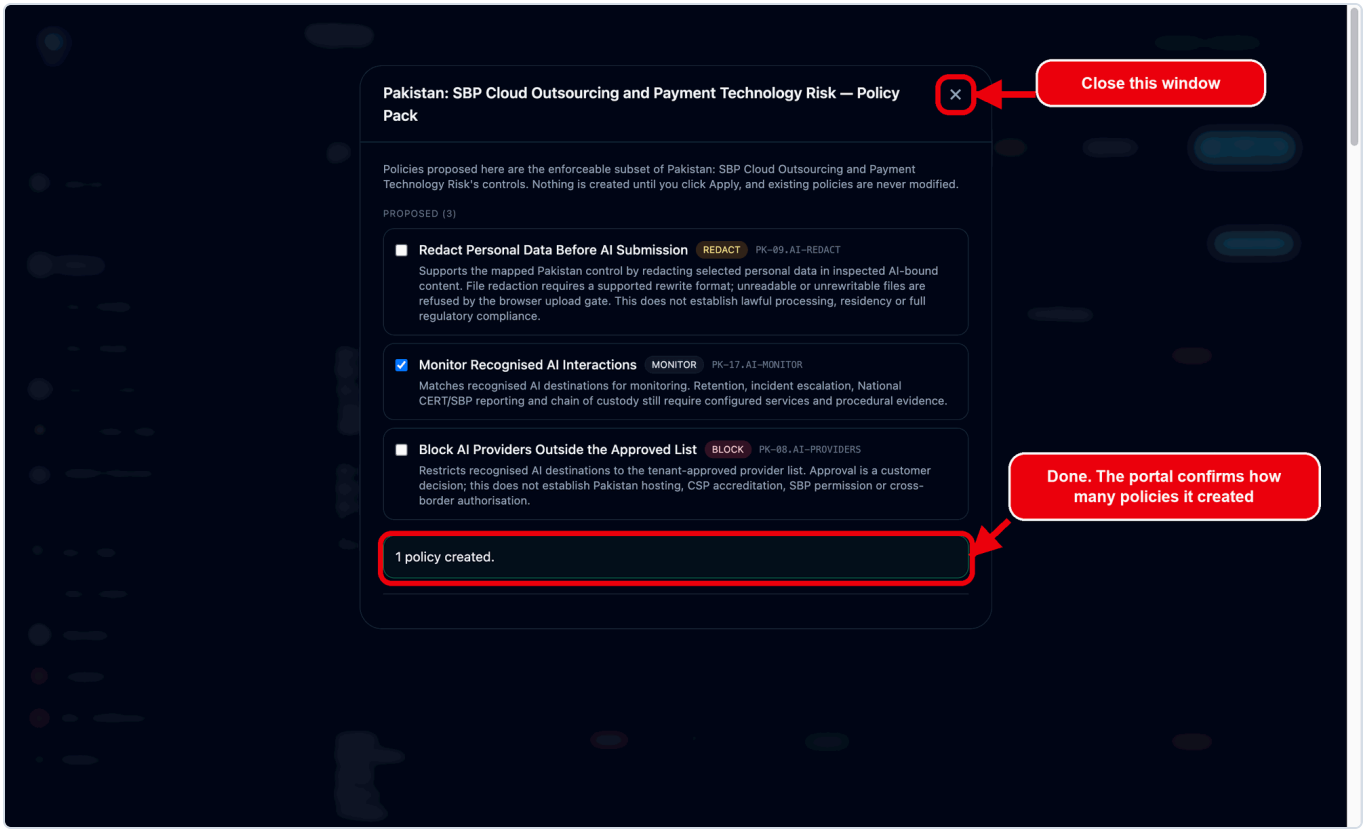
Step 2. The pack lists the policies it proposes, each with a coloured label showing what it does (REDACT, MONITOR, BLOCK) and a plain explanation. **Every proposed policy starts ticked.** Untick the ones you do not want. In this example we keep only **Monitor Recognised AI Interactions**: it records use of AI services and never blocks anyone, which makes it a safe first policy. Click **Apply Selected**.

- Policies already applied from this pack are listed under **Already applied**, locked.
- If a control needs a list that your organization does not have yet (for example `approved_providers`, the approved AI providers), a box under it says so, with **Create this artifact**, which opens **Artifacts** with the name filled in. The policy cannot work until the list exists.
- Nothing changes until you click **Apply Selected**, and your existing policies are never modified.

If you keep more than one policy, check their priorities next. A pack creates every policy at priority **100**. See section 6, "Priority decides which policy applies".



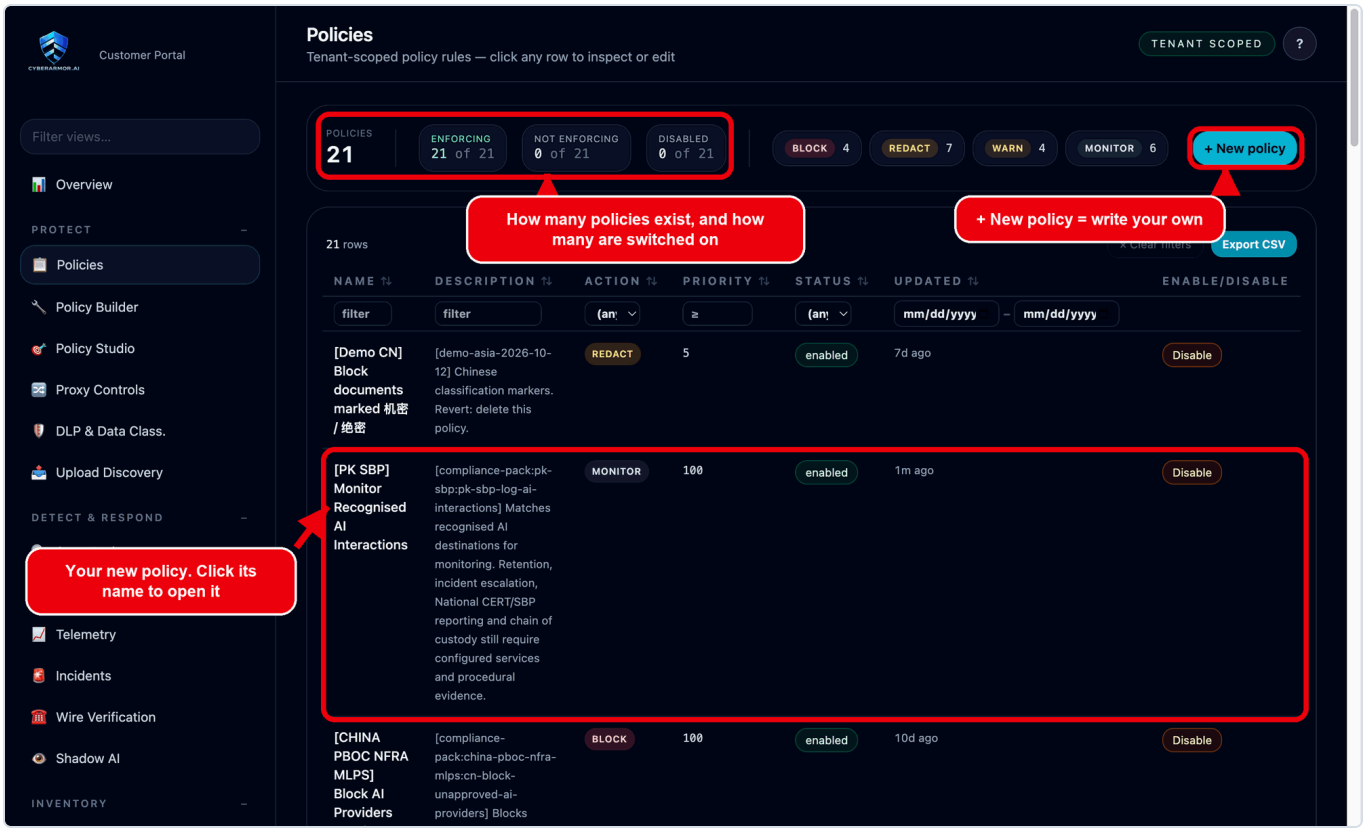
Step 3. The portal confirms "1 policy created" and runs the framework's assessment again, so the score moves on the same screen. Close the window with **X**.



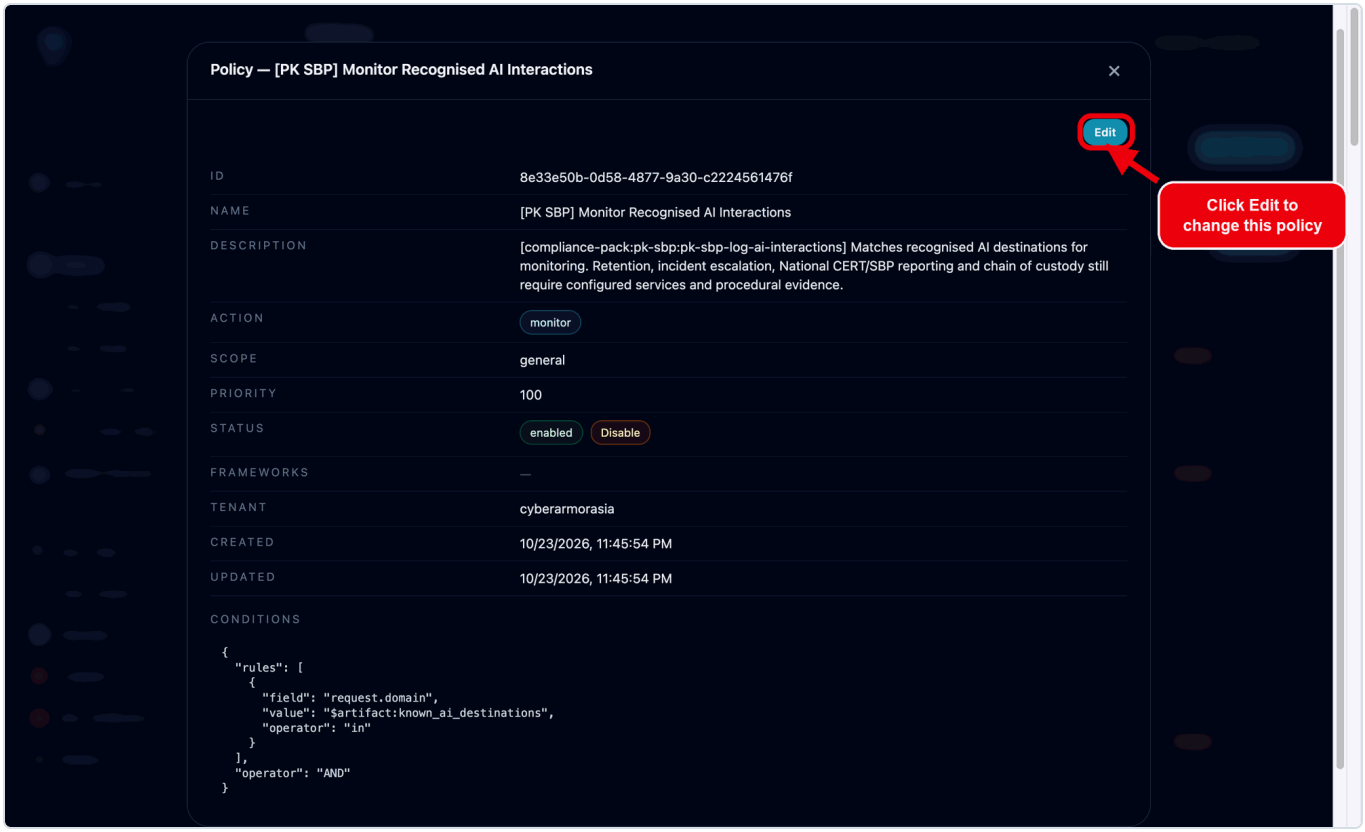
To remove a pack later, open **Policy pack** again and click **Disable pack**: the policies it created are archived, not deleted.

6. Change a policy

Step 1. In the menu, click **Protect** → **Policies**. The new policy is in the list, named with its pack's prefix: **[PK SBP] Monitor Recognised AI Interactions**. Click its name.



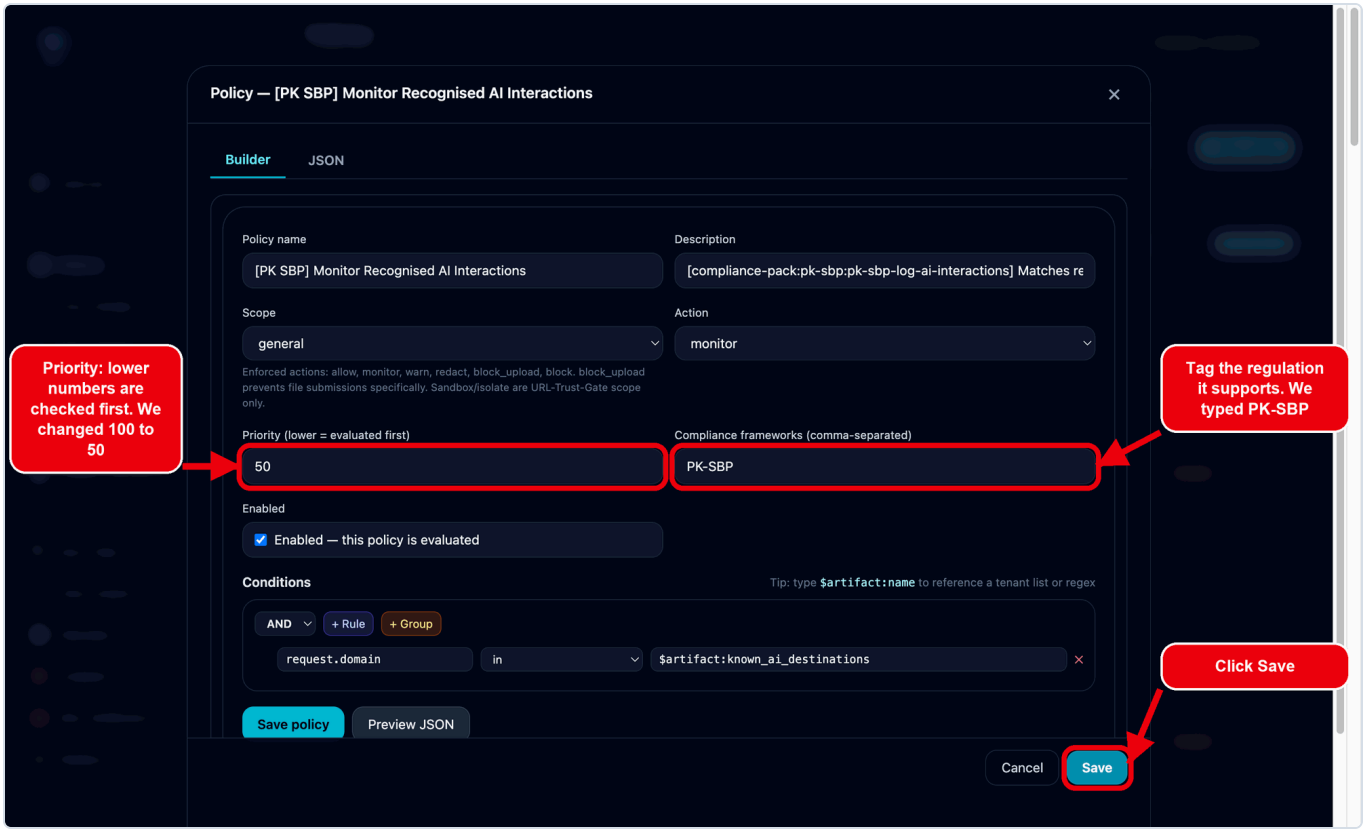
Step 2. The policy opens and shows all its settings: name, description, action, scope, priority, status, frameworks and its conditions. Click **Edit**. Edit has two tabs: **Builder** (the form) and **JSON** (the whole policy as text).



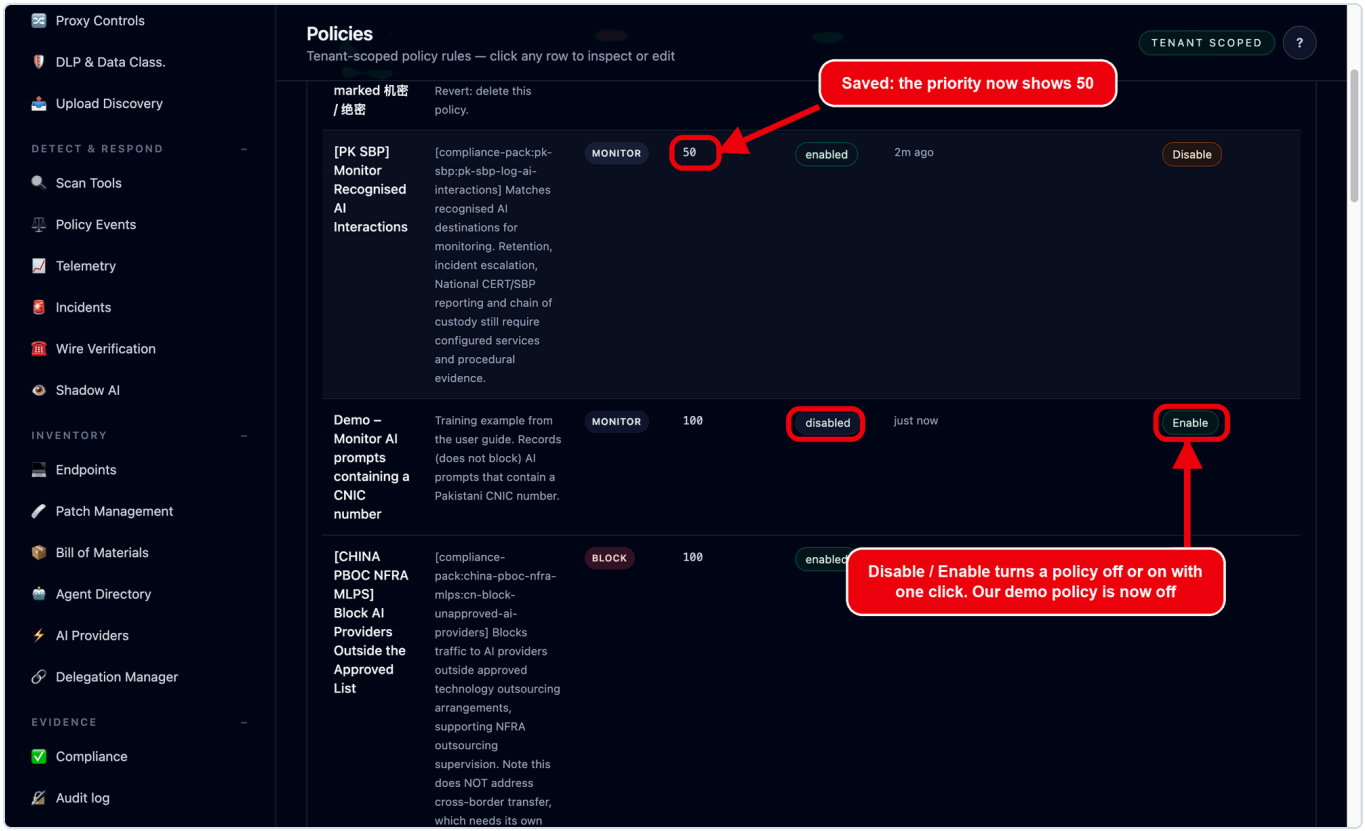
Step 3. Change what you need. In this example we changed two things:

- **Priority** from 100 to **50**. Policies with lower numbers decide first — see the next subsection.
- **Compliance frameworks**: we typed **PK-SBP**, so reports show which regulation this policy supports.

Click **Save**.



Step 4. The list now shows priority **50**. Devices pick up the change automatically the next time they check in with the portal; you do not need to touch them.



Priority decides which policy applies (please read this)

Priority is not only a checking order. It decides **which policy actually takes effect**.

For each request, the CyberArmor.AI service and the endpoint agent find every policy that matches. **The matching policies with the lowest priority number decide.** If more than one matching policy has that same number, **the strictest action among them decides**: block (and block_upload) over warn, warn over redact, redact over monitor (and audit-only), and monitor over allow. The decision then names the policy that was set aside, so the tie is visible.

The browser extension decides the same request more strictly: it applies the strictest matching action it can carry out on that page, whatever the priorities, and takes the classes to redact from the matching redact policies with the lowest priority number.

The consequence: **a weaker policy with a lower number stops a stronger one with a higher number on the service and the agent**, but not in the browser. To get the same answer everywhere, give every policy its own priority number, and put the strongest actions at the lowest numbers.

A compliance pack creates every policy at priority 100. If you keep more than one policy from a pack — say both *Record Every AI Interaction* and *Block AI Providers Outside the Approved List* — the block now wins their tie, but reorder them anyway so that the order is yours and written down: move the **strongest** action to the **lowest** number. Check your own policies too: a *monitor* policy of yours at priority 50 decides ahead of a pack's *block* at 100.

A sequence that works (lower numbers take effect first):

Priority	Action	Why here
8	`block`, `block_upload`	the strongest actions, so they decide first
85	`warn`	shows the user a warning
90	`redact`	removes the sensitive part before sending
100	`monitor`, `audit-only`	records only — the catch-all, so it decides last

The gaps between those numbers are deliberate: you can insert a new policy between two of them later without renumbering the rest. Policies you created for a specific purpose (at priority 5, 6 or 10, for example) still take effect ahead of the pack's policies, which is what you want.

Priority is a whole number and defaults to 100. `0` is accepted and would sort ahead of `1`, but do not use it — leave yourself room above whatever you set.

---

7. Create your own policy

Step 1. On the **Policies** page, click **+ New policy**. The **Policy Builder** opens.

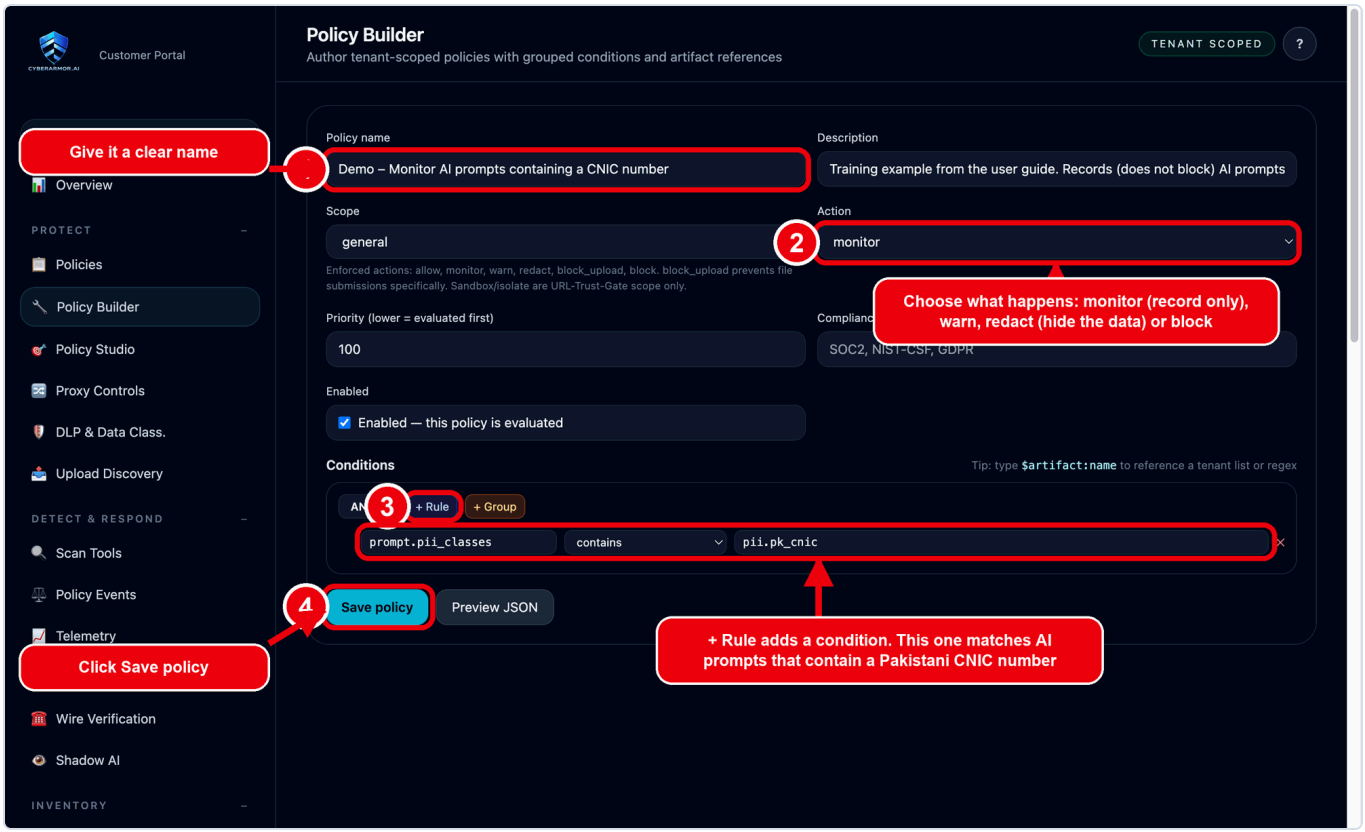
Step 2. Fill it in:

1. **Policy name.** A name anyone will understand. Example: *Demo – Monitor AI prompts containing a CNIC number*.
2. **Scope.** Where the policy applies: `general`, `proxy`, `endpoint`, `identity` or `url-trust-gate`. Leave `general` unless your technical team says otherwise; the hint under the box says which actions each scope enforces.
3. **Action.** What happens when the policy matches: - **monitor**: record it only; - **warn**: show the user a warning; - **redact**: remove the sensitive part before it is sent (choose which kinds of data under **DLP classes to redact**); - **block_upload**: stop file uploads only; - **block**: stop it completely; - **allow**: let it through.

Start with **monitor**, check the results in **Policy Events**, then make it stronger. 4. **Priority.** Defaults to 100. If this policy needs to take effect ahead of the ones a compliance pack created, give it a lower number, following the sequence in section 6, "Priority decides which policy applies". 5. **+ Rule** adds a condition, and **+ Group** a group of conditions joined by AND, OR or NOT. Our example matches AI prompts

that contain a Pakistani national ID number: field `prompt.pii_classes`, condition **contains**, value `pii.pk_cnic`. Start typing in each box and the portal suggests the choices (other examples: `pii.pk_ntn` for an NTN, `pii.pk_mobile` for a Pakistani mobile number, `pii.credit_card`). A value can also name a list from **Artifacts**, written `$artifact:` and the list's name. 6. Click **Save policy**.

Preview JSON shows the finished policy as text, if your technical team asks for it.



Example: keep a card PIN off every site but the bank's

A debit-card PIN belongs on a card terminal, an ATM, or at most the bank's own website. A card-PIN policy redacts or blocks a PIN written next to a card number anywhere else. It has three parts:

1. **The bank lists.** Open **Setup & Admin** → **Artifacts**. The **List library** at the top holds ready-made lists of banking websites: **US banks**, **US credit unions**, **Chinese banks** and **Pakistani banks**, as well as **Card networks** and **Payment processors**. Each domain in them was checked on the institution's own website or an official regulator list, and **View sources** shows where. Click **Add to this organization** on the bank lists for your country. Adding a list never removes an entry you added yourself, and a list does nothing until a policy names it. When a list has new domains, its button reads **Add new domains**.
2. **The exception is banks only.** Name only bank lists as the places a PIN may go. Never use **Payment processors** or **Card networks** as an exception: a processor, wallet or merchant platform never asks for a card PIN on a web page, and anyone can publish a payment page on such a site, so a PIN typed there is a warning sign.
3. **The policy.** Action **redact** (with `pii.card_pin` and `pii.credit_card` among the classes to redact) or **block**, at a priority lower than your other redact policies, with two rules joined by AND: the text contains a card PIN next to a card number, and `request.domain` is **not_in** your bank lists (for example `$artifact:us-banks`). CyberArmor.AI applies this policy for an organization with a script that writes it and then checks the result; ask your CyberArmor.AI contact.

With the policy in place, the browser extension removes a PIN as it is typed into an AI chat (redact) or stops the message (block), and holds a web form that has a PIN next to a card number: a block sends nothing and names the policy; a redact empties the PIN, masks the card and holds the form so the person sees why. On the bank's own website nothing changes.

---

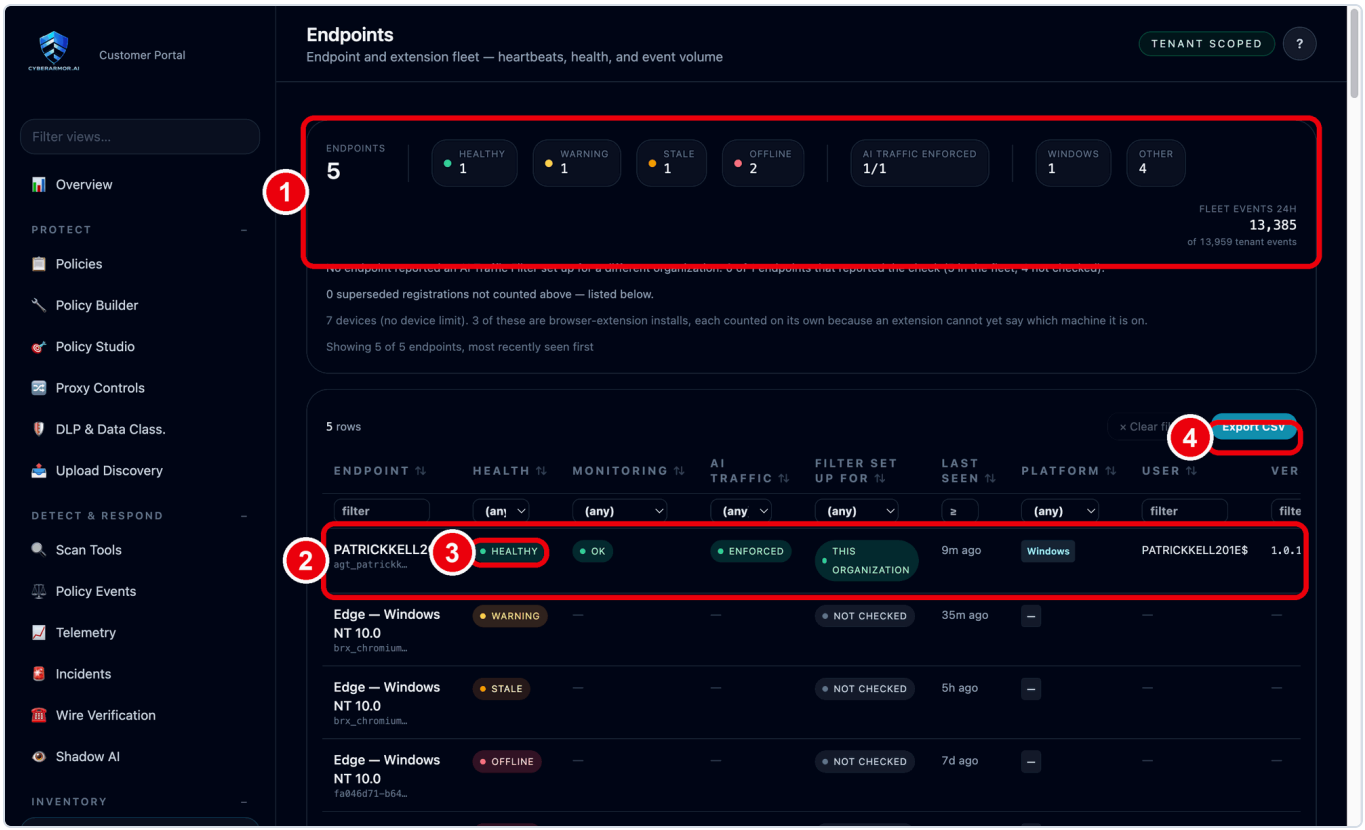
8. Turn a policy off or on

Every row on the **Policies** page has a **Disable** button (or **Enable**, if the policy is off). One click switches the policy without deleting it. The same switch is in the **Status** row of a policy's details. The **Status** row also has **Archive**. An archived policy is not evaluated anywhere and cannot be edited until a tenant admin clicks **Unarchive**. If you save a new policy under the name of an archived one, the Builder tells you and does not save it as an active policy. In the picture in section 6, step 4, our demo policy is **disabled** and shows an **Enable** button.

---

9. See your devices (Endpoints)

Inventory → **Endpoints** lists every computer, browser and phone connected to your organization.



1. **Summary.** How many devices are **Healthy**, **Warning**, **Stale** (has not reported recently) or **Offline**, how many enforce AI traffic inspection, totals by platform, and notices when something needs attention (for example a device whose AI Traffic Filter was set up for a different organization, or the device limit).
2. **One row per device:** its name, health, monitoring status, whether **AI traffic** inspection is enforced, which organization its **Filter** is **set up for**, last seen time, platform, user, version, events in the last 24 hours and active monitors.
3. **Healthy** (green) means the device is protected and reporting normally.
4. **Export CSV** downloads the list as a spreadsheet.

Click a row for the device's details: health, credential, AI traffic, last heartbeat, operating system, agent version, events, active monitors, registration time, and the raw record. A tenant admin can **Rename**, **Disable**, **Enable**, **Retire** or **Archive** a device there. Older registrations of a reinstalled computer are listed separately as **superseded**, and archived devices can be restored to the list.

AI off switch. A red-bordered card above the device list. A tenant admin chooses what it **Reaches** (**Whole organisation**, **One platform** or **One directory group**) and clicks **Turn AI off**. The portal counts the devices, asks you to confirm and takes a reason, which is kept in the audit log. The card then reads "AI off on N of M devices" and lists the devices that have not confirmed yet. A device applies it at its next check-in, and the person using it cannot turn AI back on. **Turn AI back on** reverses it. The same switch for one device is in that device's

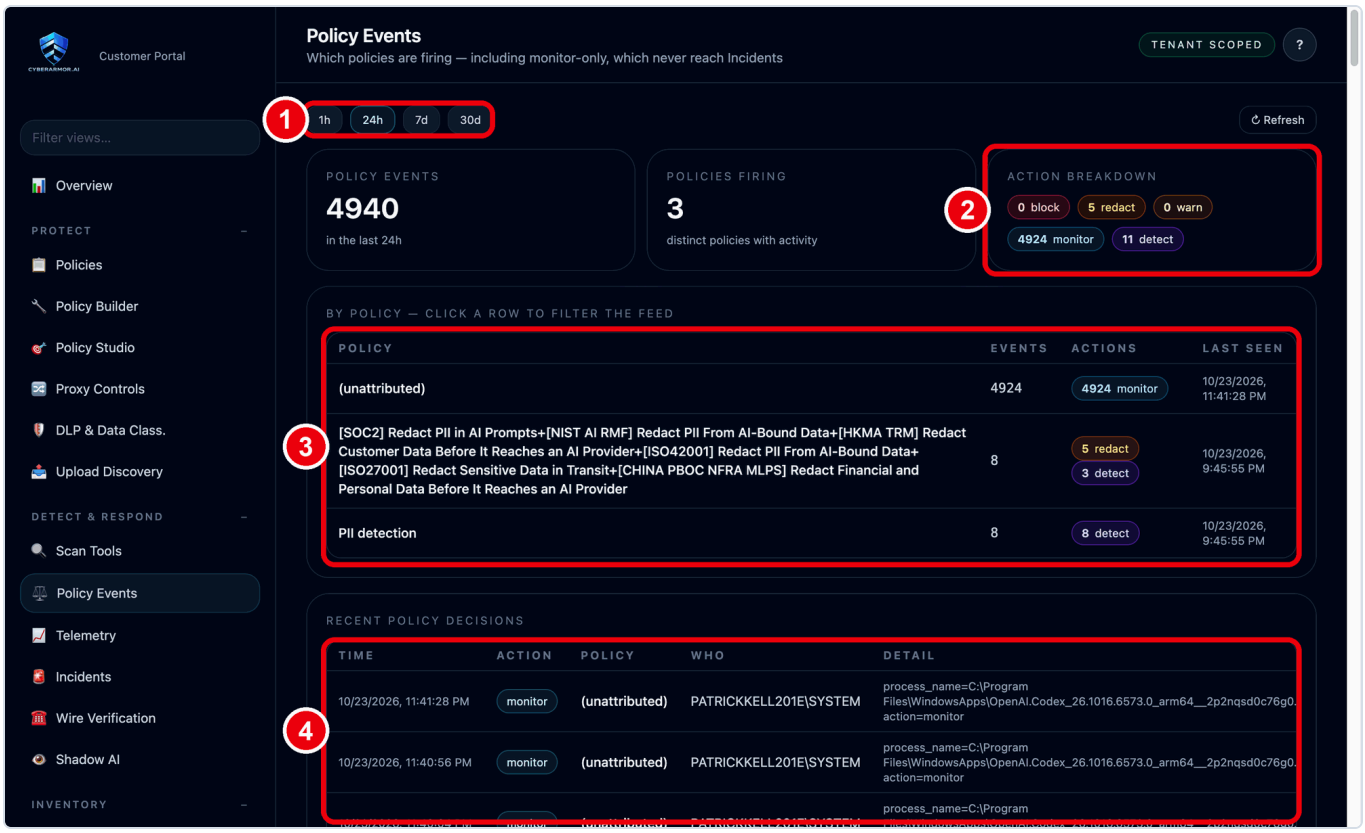
details. CyberArmor.AI protection stays on throughout. On a Mac or Windows computer, off blocks AI service addresses and stops running AI apps. **Do not rely on it for iPhone, iPad or Android yet:** on those devices it does not reliably stop AI today, even though the list offers them. The AI off switch is at the **Configurable** tier on the capability status page, and on phones it is **Roadmap**.

Endpoint response. Below the summary: computers that are isolated from the network, and files an agent has quarantined, as each agent measured them at its last check-in. A tenant admin clicks **Release** to end an isolation or **Restore** to put a quarantined file back; both ask you to confirm and are carried out at the computer's next check-in. A restored file never overwrites anything.

---

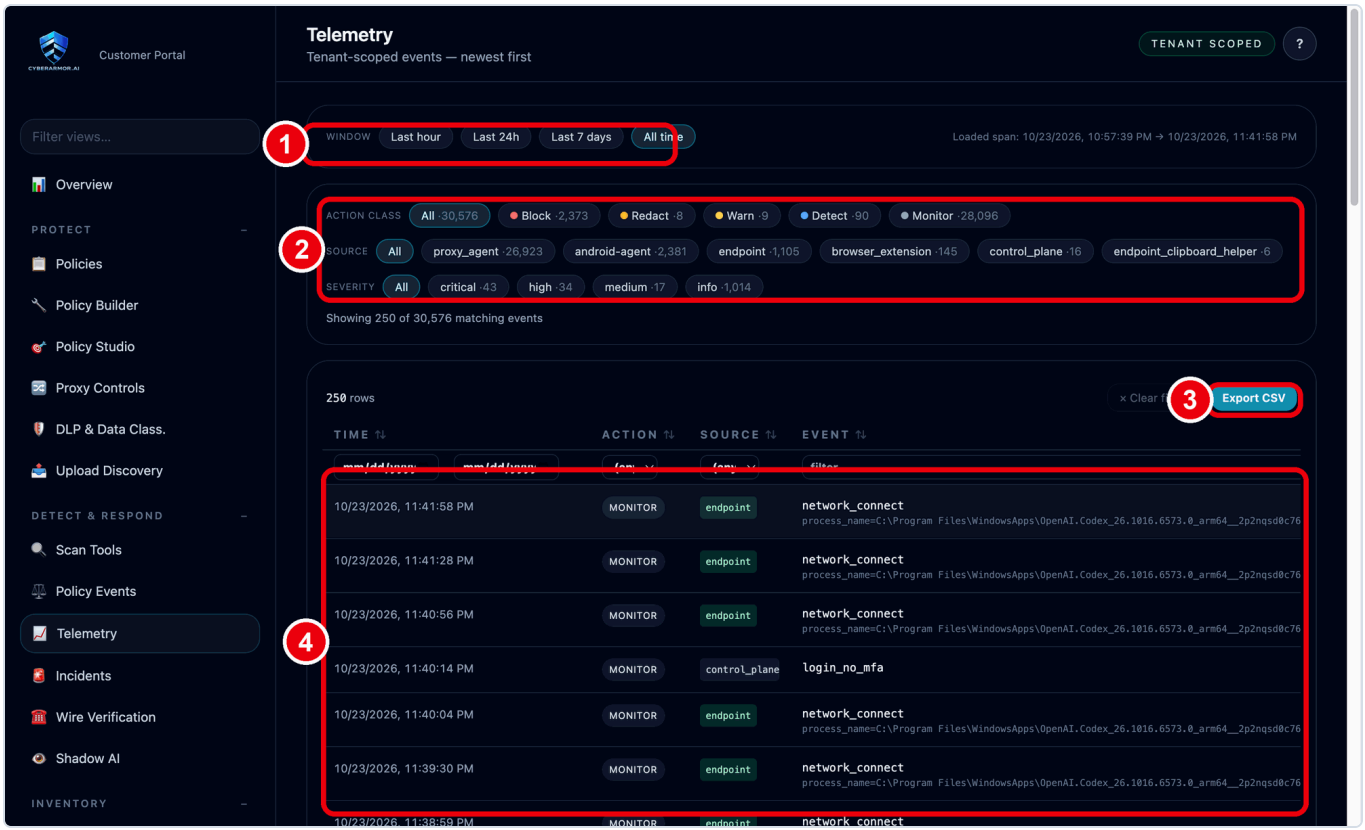
10. See what is happening

Policy Events: which policies are firing



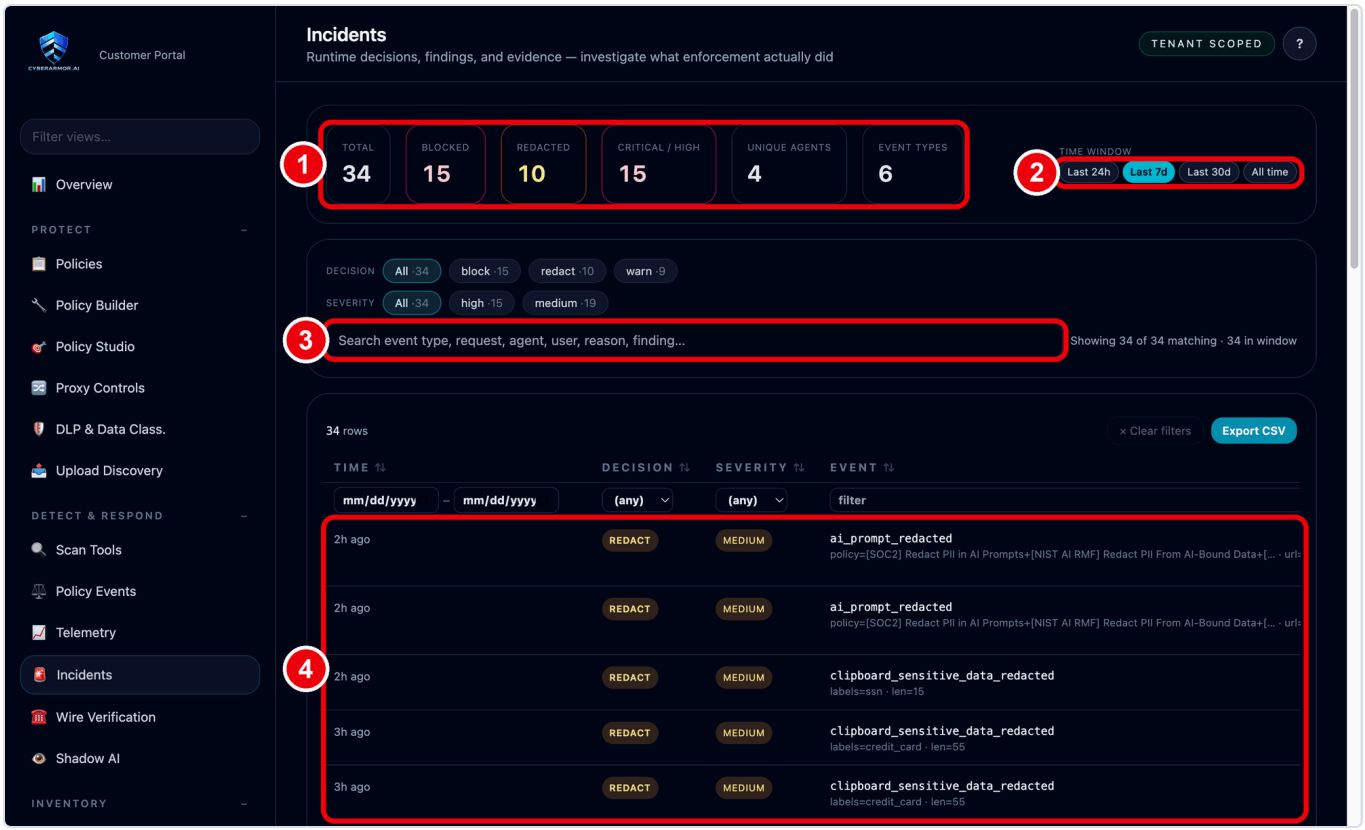
1. **Time window:** the last 1 hour, 24 hours, 7 days or 30 days, with **↻ Refresh**.
2. **Policy events, policies firing, action breakdown:** how many events were blocked, redacted, warned, detected or monitored.
3. **By policy:** how often each policy fired. Click a row to see only that policy's events.
4. **Recent policy decisions:** each decision, with time, action, policy, user and details. Click one to see its details: the policy by name with **Open policy**, the content scan, and the decision's context.

Telemetry: every event



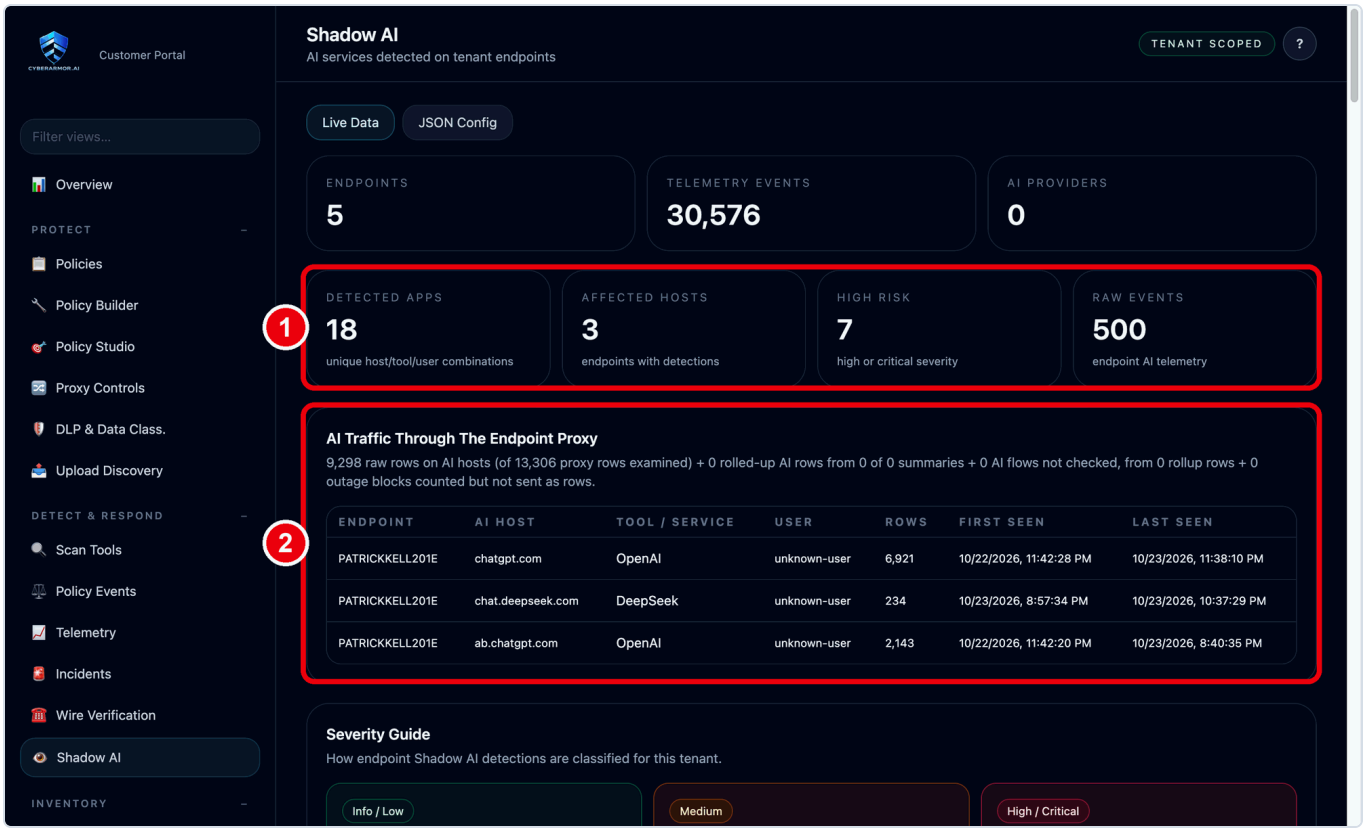
1. **Window:** the last hour, 24 hours, 7 days, or all time.
2. **Filters:** click to show only one action (Block, Redact, ...), one source (endpoint, browser extension, ...) or one severity.
3. **Export CSV** downloads what you see.
4. **The events, newest first.** Click any row to see its full details, including the policy, the content scan and the correlation (section 2).

Incidents: what was stopped or changed



1. **Totals:** total, blocked, redacted, critical or high severity, unique agents and event types.
2. **Time window**, with decision and severity filters.
3. **Search** by event type, user, device or reason.
4. **The incidents.** For example `ai_prompt_redacted` means personal data was removed from an AI prompt before it was sent. Click one for its details: decision, reasons, the policy with **Open policy**, the content scan, the findings, the user, device and correlation.

Shadow AI: AI tools in use



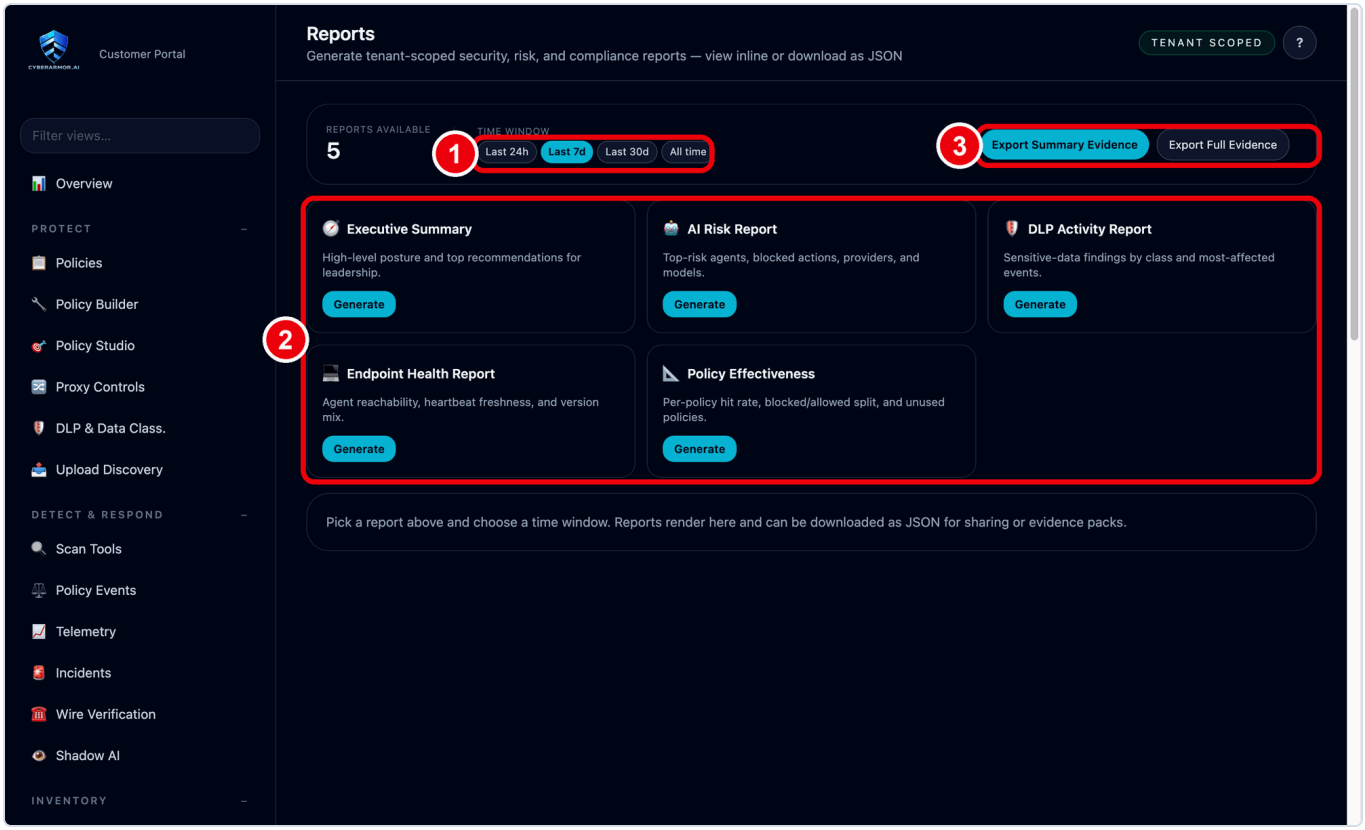
1. **Detected apps, affected hosts, high-risk findings:** a summary of the AI services and tools found on your devices.
2. **AI traffic through the endpoint proxy:** which computer used which AI service (for example ChatGPT or DeepSeek), how much, and when it was first and last seen.

Click a traffic row or a detection to see its details and its **Actions**. The window shows what is set for that AI service today (on the approved list or not, its policy, whether it was reviewed), and a tenant admin can act on it from there: **Approve** (adds it to the approved AI list; writes no policy), **Block**, **Monitor**, **Warn**, **Redact** (each writes a policy for that service; Monitor, Warn and Redact are placed after every block policy, so nothing blocked today stops being blocked) or **Mark reviewed** (records your name and the time). Each asks before it changes anything. **Open incidents for this host** and **Open policy events** take you to the matching records.

---

11. Reports and the Audit log

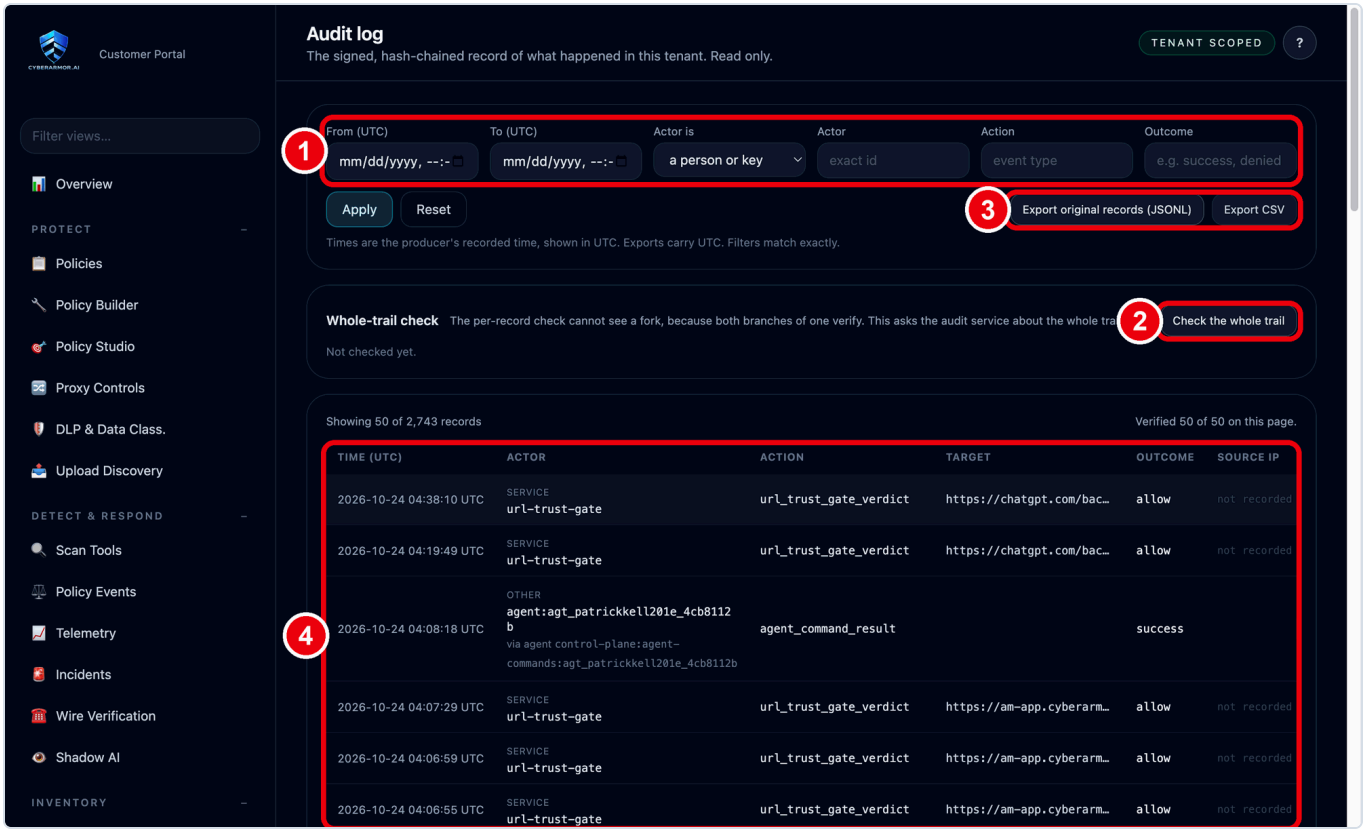
Reports



1. **Time window** for the report: the last 24 hours, 7 days, 30 days, or all time.
2. Ten ready-made reports: **Executive Summary**, **Compliance Posture**, **Incidents and Response**, **AI Usage**, **AI Risk Report**, **DLP Activity Report**, **Wire Verification**, **Endpoint Health Report**, **Policy Effectiveness** and **Audit Trail Integrity** (the whole-trail check of the signed audit log). Click **Generate**. The report appears on the page. **Download PDF** opens your browser's print dialog: choose **Save as PDF**. **Download JSON** saves the data itself.
3. **Export Summary Evidence** / **Export Full Evidence** download an evidence pack for auditors and regulators.

Audit log

The audit log is the signed record of what happened in your organization: who signed in, which policies changed, which commands ran. Each record is signed with HMAC-SHA256 and hash-chained to the record before it, and the chain is sealed at checkpoints signed with ML-DSA-87, a post-quantum signature that anyone holding the public key can verify. The log is read only.

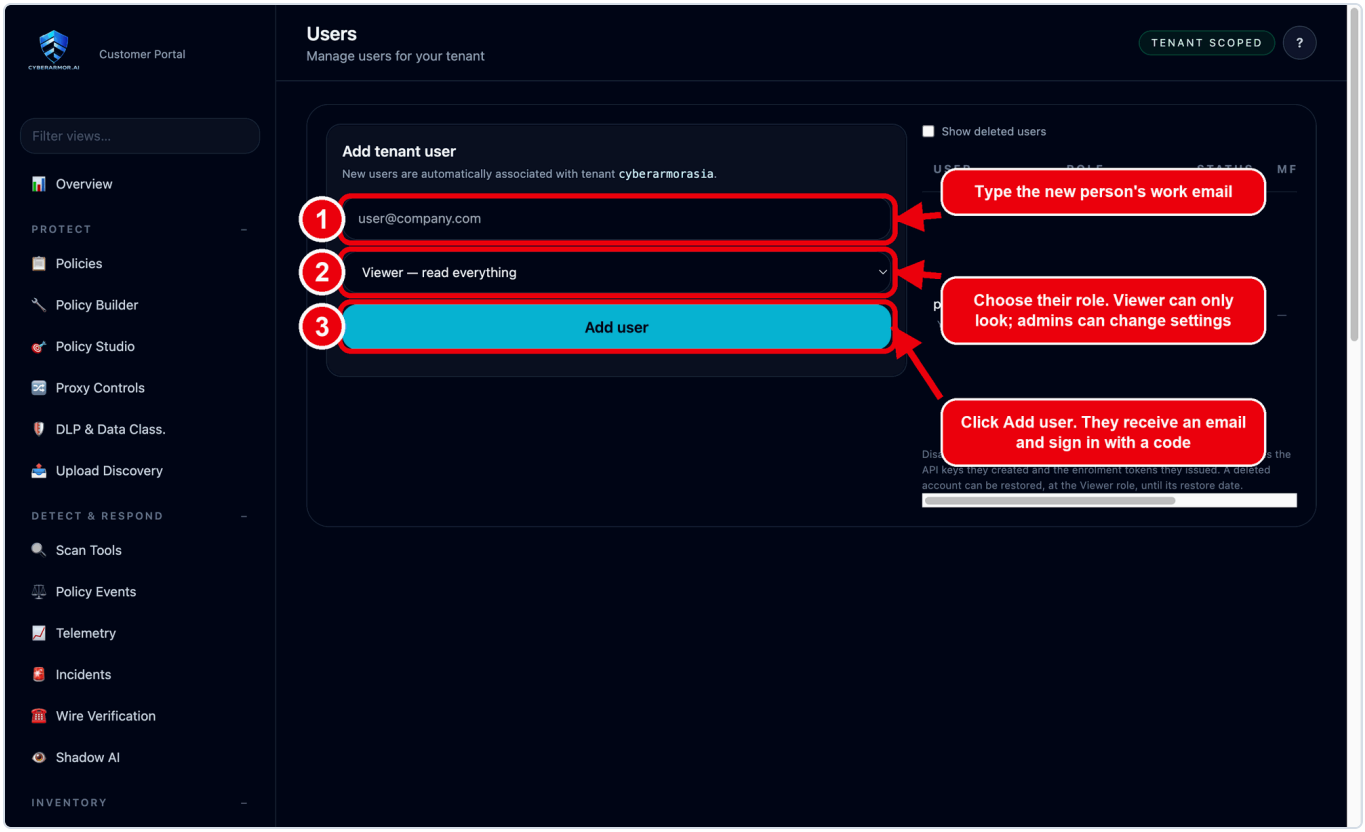


1. **Filters:** date range, actor, action and outcome. Click **Apply**.
2. **Check the whole trail** asks the audit service to verify every chain and its signed checkpoints, and to report any fork.
3. **Export original records (JSONL)** downloads the signed records exactly as the audit service holds them; this is the evidence file. **Export CSV** gives a spreadsheet.
4. **The records.** "Verified 50 of 50 on this page" means every record shown passed its signature check. Click a record to see it in full, with its verification.

---

12. Add a colleague (Users)

Setup & Admin → Users. Only a tenant administrator can add users.



1. Type the person's work email.
2. Choose a role: - **Viewer**: can read almost everything, change nothing; - **Deployer**: can also issue install tokens and download installers, but cannot open the **Audit log**; - **Analyst**: can also approve actions; - **Tenant admin**: full control. Only a tenant admin can open **Proxy Controls** and **Users**.
3. Click **Add user**. They receive an email and sign in with an emailed code.

Each row has the person's role (a tenant admin changes it in place; it saves at once) and **Rename**, **Disable**, **Sign out**, **Resend invite**, **Delete** and **Change email**. Click a row for the person's details, including whether they have set up an authenticator app (MFA) and when they last signed in.

---

13. Sign-in security (Identity / SSO)

Setup & Admin → **Identity / SSO** controls how people sign in.

- **Single Sign-On.** A tenant admin enters the organization's OIDC settings (for example Microsoft Entra ID or Okta), turns on **Enable SSO**, and can let people be created, as Viewers, on their first single sign-on (**JIT provisioning**; off by default). Click **Save SSO configuration**. The email-code sign-in keeps working.
- **Multi-factor authentication.** A tenant admin ticks **Make MFA available to users in this tenant**. This makes MFA available; it does not turn it on for anyone. Each person then sets up their own authenticator app on this page: **Set up authenticator app**, scan the QR code with an app such as Microsoft Authenticator or Google Authenticator, type the 6-digit code, and save the backup codes, which are shown once. From then on, that person's sign-in asks for a code from the app. Turning the setting off stops the challenge for everyone but keeps their enrollments, which come back if it is turned on again.
- **Sign out everywhere** ends every session of your own account in this organization, this one included — for example after using a shared or lost computer.

The tick box saves the moment it changes.

---

14. Settings for every device

Setup & Admin → **Settings** shows your organization ID and name (set by CyberArmor.AI; they cannot be changed here) and the settings that apply to every device. Only a tenant admin can change them; everyone else can read them. Sign-in settings are on **Identity / SSO** (section 13).

Protection level

The same control as on the **Protection** page (section 17): the level your organization is on, whether your settings still match it, and **Apply** for a tenant admin. Changing any setting a level wrote turns the badge to **Custom**.

Billing

Shown only to a tenant admin of an organization that pays CyberArmor.AI directly by card. **Manage billing** opens Stripe's billing page in a new tab, where you cancel, change the number of users or devices, update the card or download invoices.

Automated endpoint response

Lets endpoint agents act on a destructive detection without waiting for a person. **A new organization starts with this off.** With it off, detections are still reported; nothing is acted on.

There are two separate choices, and both are needed before anything happens:

1. **Allow agents in this tenant to respond automatically** (the master switch).
2. An action for each detection: **Ransomware suspected** and **Mass deletion**. Each starts at **Report only**. The other choices are **Pause it and ask the person** and **Kill the process**.

Click **Save response settings**. Saving an action other than **Report only** asks you to confirm, naming what will be armed.

Pause it and ask the person suspends the program the computer's sensor saw make the change and asks the person at the keyboard to allow it, keep it paused or stop it. Under **When the endpoint pauses a program and asks** you set how long to wait for an answer (120 seconds to start with; 15 to 900), what happens with no answer (**Keep it paused**, the default; **Stop the program**; or **Let it continue**), and how long not to ask again after the person allows it (10 minutes to start with; 0 asks every time). On macOS the agent also refuses that program's further deletions; on Windows the program is suspended. Linux computers are reported only. Files already deleted cannot be recovered.

Kill the process acts only on a process that the computer's own sensor saw make the change (Endpoint Security on macOS, ETW on Windows), and only while it is still that same process. Linux computers, and any detection the sensor could not attribute, are reported only. The ransomware kill is at the **Configurable** tier on the capability status page (https://cyberarmor.ai/status): it has been run on a Windows 11 test computer, it is built but not yet validated on Mac hardware, and it has not yet run on a customer's computer. Detection alone does not stop encryption, and files already changed stay changed. Isolating a computer from the network is offered only on a server where CyberArmor.AI has turned it on; otherwise the page says "Isolation is not available on this server yet". Where it is on, an isolated computer is released from the **Endpoints** page (section 9).

Private network addresses and the local page check

- **Private network addresses (browser extension).** What the browser extension does when someone opens a router, printer, NAS or intranet app by its private IP address: **Allow (default)**, **Warn** (users can continue) or **Block** (users cannot continue). Whatever you choose, the extension never sends these addresses to CyberArmor.AI. It saves the moment you choose; browsers apply it at their next policy sync.
- **Check pages on the device first (local Trust Gate).** On by default. On a computer with the CyberArmor.AI agent and its model pack, pages are checked on the computer and only risky or unreadable pages go to the server. Off: every page is checked on the server. It saves the moment it changes.

Local Models

Runs CyberArmor.AI's detection models on each computer, from a signed model pack encrypted to each device, instead of sending the content to the server to be checked. It runs on Macs with Apple silicon (macOS 14 or later) and on Windows; other devices decide remotely, as before. Local Models are at the **Configurable** tier on the capability status page.

A new organization starts with Local Models off — the status reads **off (not yet rolled out)** — unless CyberArmor.AI's staged rollout has reached it. **To turn them on**, a tenant admin chooses **On** and clicks **Save**. The choices are:

- **Follow the CyberArmor.AI rollout** (shows whether the rollout has it on or off right now), **On**, or **Off**. An admin's **On** or **Off** is kept until an admin changes it.
- **Include the PHI model (healthcare)**.
- **Model set: Standard** (prompt injection, PII, toxicity, output safety) or **Full** (all five models), each with its download size.
- **Pack version: Newest published**, or pin a version.
- **Update schedule** (in each device's local time): **Daily, off-hours** inside a window that starts as **02:00 to 05:00**, **Tenant schedule** (only on the days you tick, inside the window), or **Manual only**. A device checks daily and downloads only what changed; a laptop closed through the window catches up within 36 hours.
- **Push update now (all devices)** asks every device to update now; **Push** on a device's row does it for one.

The headline reads "N of M eligible devices on pack vX". A device that has not reported is shown as **unknown — not reporting**, never as current, and devices outside the count (phones, tablets, older reinstalls) are listed under **Not counted in M**. When an admin turns Local Models off, each device deletes its pack at its next daily check and its checks go to the server, as before.

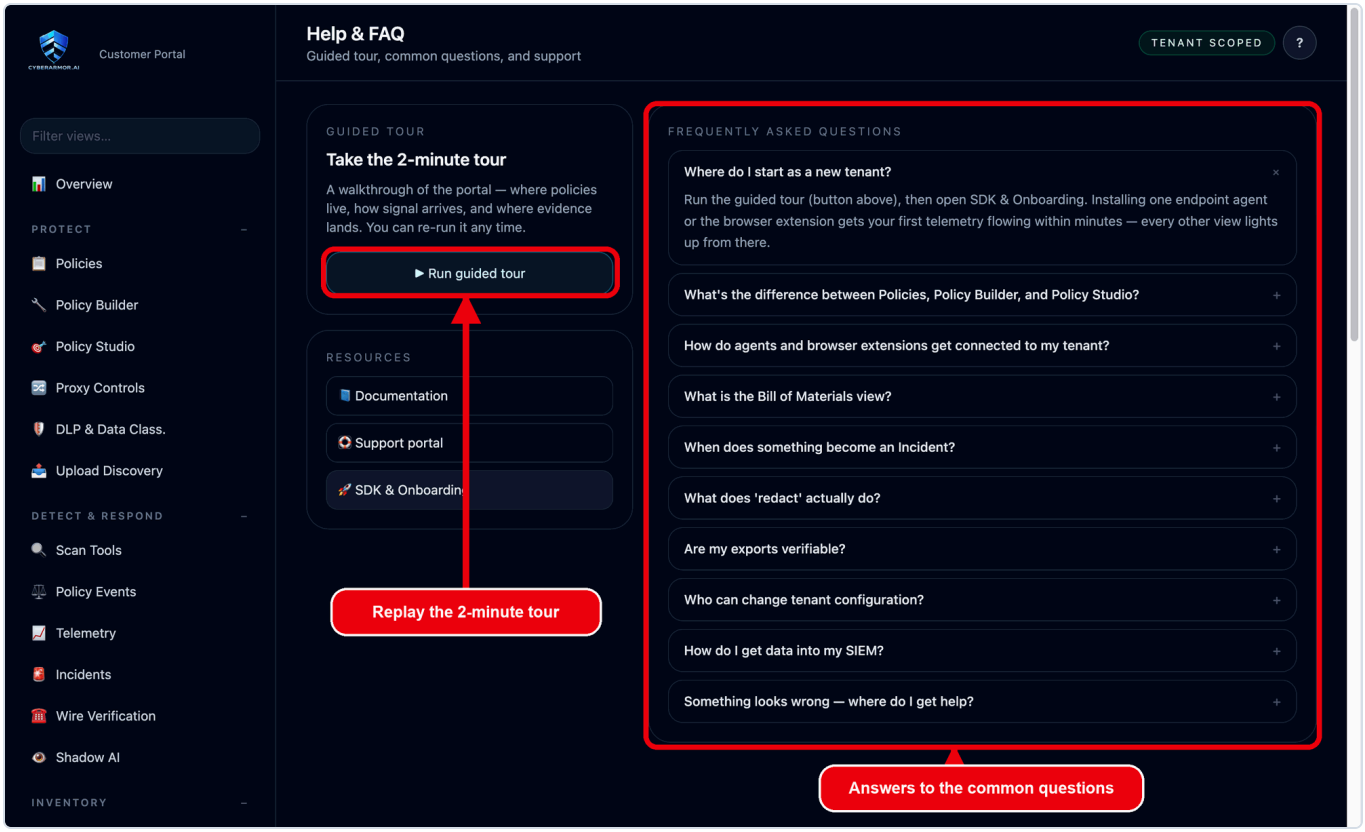
---

15. Get help and training videos

Support → **Help & FAQ** has:

- **Take the 2-minute tour** (**▶ Run guided tour**);
- **Resources: Documentation**, **Support portal** and **Onboarding & Downloads**;
- **Training videos:** the eight customer videos — *Installation: from sign-in to a protected computer*, *Protection in action*, and *Customer portal training*, chapters 1 to 6. **Watch** plays a video in a new tab; **Chapters** opens its page on the documentation site with a chapter list. **All videos with their chapters** opens the full list;
- **Frequently asked questions**, such as *How do agents and browser extensions get connected to my tenant?* and *What does 'redact' actually do?*

The same videos are at **https://docs.cyberarmor.ai/training/** and **https://cyberarmor.ai/training**. They are the English (Pakistan) edition, recorded on demonstration machines. This guide and its companions are on the documentation site under **User guides**, as web pages and as PDF and Word downloads.



16. Quick reference: every menu item

Menu group	Page	What it is for
—	**Overview**	Mission Control dashboard (section 3).
Protect	**Policies**	List, open, edit, enable and disable policies (sections 6 and 8).
Protect	**Policy Builder**	Create a policy with a form (section 7); for the priority field see section 6.
Protect	**Policy Studio**	Your policies as AI decisions, with a **Test** for each (section 17).
Protect	**Proxy Controls**	Fail mode, sites never decrypted, and the policies used by AI traffic inspection (section 17).
Protect	**DLP & Data Classification**	Data classification labels, custom rules, detection patterns and a live scan (section 17).
Protect	**Upload Discovery**	Upload websites the browser extension found that are not yet in the catalog; add them so upload blocking covers them.
—	**Protection**	Your protection level, what is protected right now (desktop agent, browser extension, phone or tablet, AI traffic proxy) and the one next step (section 17).
—	**Value & Friction**	What CyberArmor.AI did, with the count behind every figure, and what it cost the people using it (section 17).
Detect & Respond	**Scan Tools**	Paste any text and see what CyberArmor.AI detects in it.
Detect & Respond	**Policy Events**	Which policies fired, and how often (section 10).

Menu group	Page	What it is for
Detect & Respond	**Telemetry**	Every event from every device (section 10).
Detect & Respond	**Incidents**	What was blocked, redacted or flagged (section 10).
Detect & Respond	**Wire Verification**	Payments held until someone calls back the number on record (section 17).
Detect & Respond	**Reported Blocks**	Blocks your own people say were wrong, with the policy that decided (section 17).
Detect & Respond	**Privileged Changes**	Holds a change to who controls a wallet, a signer or a privileged deploy target until a second named person approves it (section 17).
Detect & Respond	**Payment Path**	Checks that the destination about to be paid is the one that was approved (section 17).
Detect & Respond	**Shadow AI**	AI tools and services found in use, and what to do about each (section 10).
Inventory	**Endpoints**	Your connected devices and their health (section 9).
Inventory	**Patch Management**	Keep software on your devices up to date, with approval and maintenance windows (paid add-on).
Inventory	**Bill of Materials**	AI models, tools and libraries found on devices, with drift, vulnerabilities and an IOC scan.
Inventory	**Agent Directory**	AI agent identities seen in your audit records and activity, plus your devices.
Inventory	**AI Providers**	Credentials for the AI services (OpenAI, Anthropic, ...) your own applications use.
Inventory	**Delegation Manager**	What each person has delegated to which AI agent: issue, limit and revoke.
Evidence	**Compliance**	Framework assessments and policy packs (section 5).
Evidence	**Audit log**	The signed, read-only record (section 11).
Evidence	**API access log**	Every call made to the CyberArmor.AI server (unsigned).
Evidence	**Reports**	Ready-made reports, PDF and JSON, and evidence exports (section 11).
Evidence	**Action Graph**	Chains linking each AI agent to the AI services, models and tools it used.
Evidence	**AI Risk Dashboard**	Overall risk score, riskiest agents and recommendations.
Setup & Admin	**Onboarding & Downloads**	Installers and install tokens (section 4).
Setup & Admin	**API Keys**	Keys for connecting your own systems.
Setup & Admin	**SIEM Config**	Send events to your security monitoring system (Splunk, Sentinel, ...).
Setup & Admin	**Identity / SSO**	Single sign-on, multi-factor authentication, sign out everywhere (section 13).
Setup & Admin	**Artifacts**	Reusable lists and patterns used in policy rules, and the List library of bank and payment lists (section 7).
Setup & Admin	**Users**	Add and manage colleagues (section 12).
Setup & Admin	**Settings**	Automated endpoint response, private network addresses, local page check and Local Models (section 14).
Support	**Help & FAQ**	Tour, training videos and common questions (section 15).

17. Every other view, one by one

This section covers the pages sections 1–16 do not walk through. Each entry says what the view is for. Every list on these pages opens its details on a click (section 2).

Protect

Policy Studio — your policies as AI decisions: **Allow**, **Deny**, **Allow with redaction** and **Require user decision**, with the AI-aware policies in a table. **Test** on a row opens **Evaluate Policy**: choose a provider and model, type a test prompt and click **Evaluate**. It shows the decision, the reason, the deciding policy, which policies were checked and any problems, and changes nothing.

Protection — the level your organization is on, measured against your real settings: a level's name, **Custom** (a setting was changed since), **Could not verify** or **No level chosen**. Four cards show what is **Protecting**, **Installed, not reporting** or **Not installed**, and **Next step** names one thing to do. A tenant admin picks a level with **Apply**; **What it sets** lists every setting a level writes. The five levels are **High AI Security & Trust**, **Strong Protection**, **Balanced Protection**, **Light Protection** and **Monitor Only**, and each says how much it interferes. On a level that allows only approved AI services, **Apply** first shows your approved list and the AI services your devices were seen using that are not on it. The same control is on **Settings**.

Value & Friction — the last 30 days: what was blocked, redacted or held, and beside it the appeals, reported blocks and pauses the same policies cost. A tenant admin sets what a record is worth under **What a record is worth**; amounts are estimates and say so. When shown, a saved-amount tile on Mission Control opens this page.

Reported Blocks — a person who is blocked can say the block was wrong; the report lands here with the policy and the evidence reference. **Open** and **All** tabs. Accepting a report records your judgement only. It changes no policy; change the policy on **Policies**.

Privileged Changes — set-up and operation of the second-person approval for wallet, signer and deploy-target changes. It shows nothing until the control is enabled and connected.

Payment Path — approved destinations, certificates, what was observed in the last 25 hours and how many payment surfaces reported. One person proposes an approved destination and a different person approves it.

Proxy Controls — how computers treat AI traffic when CyberArmor.AI cannot check it (**Fail closed** or **Fail open**; this saves the moment you choose), the hosts whose encrypted traffic is never decrypted (TLS interception bypass), and the policies used by AI traffic inspection, with **Show archived**. **Files a redact policy cannot redact** — an image, a PDF or an Office document cannot be rewritten without corrupting it, so you choose **Warn and let the user decide (default)**, where the person sees the policy and the data class and picks "Upload anyway" or "Don't upload", or **Block**.

DLP & Data Classification — the four data classification labels (Public, Internal, Confidential, Restricted), your **Custom Classification Rules** (the first matching rule wins), the **DLP Detection Patterns** with which ones your redact policies use, and **Scan & Classify Content**, where you paste text and click **Classify** to see the label and findings.

Upload Discovery — upload addresses the browser extension saw that the built-in catalog does not cover yet, with how many uploads, which file types and who used them. A tenant admin clicks **Promote** to add one, so upload blocking covers it.

Detect & Respond

Scan Tools — four scans you run on pasted text: **Prompt Injection Scan**, **Sensitive Data Scan**, **Output Safety Scan** and **Full Pipeline Scan**. Each shows its findings and a ▸ **Raw JSON** expander. Useful to see what a policy would see before you write it.

Wire Verification — holds a payment above the threshold your firm sets until someone phones the counterparty back on the number in your own **directory of record**, never a number supplied with the instruction, and the number of people your firm requires confirm it — never the person who started the payment. The queue shows each held payment, who can clear it and what happens if nobody does; **I am calling now**, **Verified on the call** and **Refused** record the call. CyberArmor.AI holds no funds and does not watch your payment systems: the system that holds the payment must send each payment to CyberArmor.AI, and until that one integration is built, this view shows nothing. Wire Verification is at the **Pilot-ready** tier on the capability status page.

Inventory

Patch Management — keeps software on your devices up to date (winget, Homebrew, apt, yum/dnf) with approvals, maintenance windows and per-app auto-approve. A paid add-on: without it, the page says it is not enabled.

Bill of Materials — a continuously updated inventory of software, hardware, ML models and collectors, with tabs for **Components**, **Drift** (added, removed, version changed), **Loaded vs Installed**, **Sources** (repositories, artifact registries, cloud inventory), **Vulnerabilities** (OSV, CISA KEV and EPSS, with VEX status) and **IOC Scan**, which says where a match was seen and what to do about it. **Export CycloneDX** and **Export Signed** produce evidence files.

Agent Directory — the AI agents seen in your activity (with events, blocks, redactions, providers, models, tools, users and hosts) and, on its second tab, your endpoint agents.

AI Providers — one card per AI provider your own applications can use through CyberArmor.AI, with its models and status. **Configure** stores the provider's API key and settings; **Test** checks it.

Delegation Manager — the record of which person delegated what to which AI agent: issue a delegation with a scope and an expiry, and **Revoke** it (this cannot be undone).

Evidence

API access log — every call made to this organization's part of the CyberArmor.AI server, with method, path, status, duration and who made it. It is not signed; the signed record is the **Audit log** (section 11).

Action Graph — a picture of which agents and people used which AI services, models and tools, with blocked links marked. Click a node to inspect it and walk to its neighbours; click an action for its details.

AI Risk Dashboard — average risk score, the riskiest agents, high-risk events and recommendations, over a time window you choose.

Setup & Admin

API Keys — keys for connecting your own systems, each with a role. A new key is shown once, when it is created. **Disable** turns a key off.

SIEM Config — one connector each for Splunk, Microsoft Sentinel, QRadar, Elastic, Google SecOps and Syslog CEF. A tenant admin clicks **Configure**, saves the connector, then **Test** sends a test event. Secret fields are never shown again.

Artifacts — reusable lists and patterns that policies name as `$artifact:` and the list's name (for example the approved AI providers). The **List library** of bank and payment lists is at the top (section 7). Each artifact can be edited, disabled, archived and restored. **Block inappropriate websites** — a ready-made acceptable-use policy at the top of **Artifacts** (also linked from **Policies**). A tenant admin ticks adult, gambling or piracy and clicks **Apply**: each list is added to your artifacts and one block policy names them. The block page tells the person the site is blocked by your organization's acceptable use policy. It is enforced by the browser extension, by the desktop agent once AI traffic inspection is set up, and by the Android app; not by the iOS app.