AI CCTV Monitoring Platform

White-label ready

User manual

Set up and run the platform on your site

AI that watches every camera, on your own server. Install the platform, connect the CCTV you already have, switch on the checks you need, review the evidence behind every alert, and train it on your own footage.

Version 1.1.804 Updated 8 November 2026

For administrators, supervisors and reviewers

Contents

Get started

01 How the platform works

02 Before you install

03 Install

04 First sign-in and first hour

05 Finding your way

Operate

06 Cameras and zones

07 Rules and PPE types

08 PPE and security modes

09 Watching live

Investigate

10 How a violation is decided

11 Events and review

12 Notifications

13 People

14 Overview and Reports

Improve

15 Test bench

16 Improving accuracy

Administer

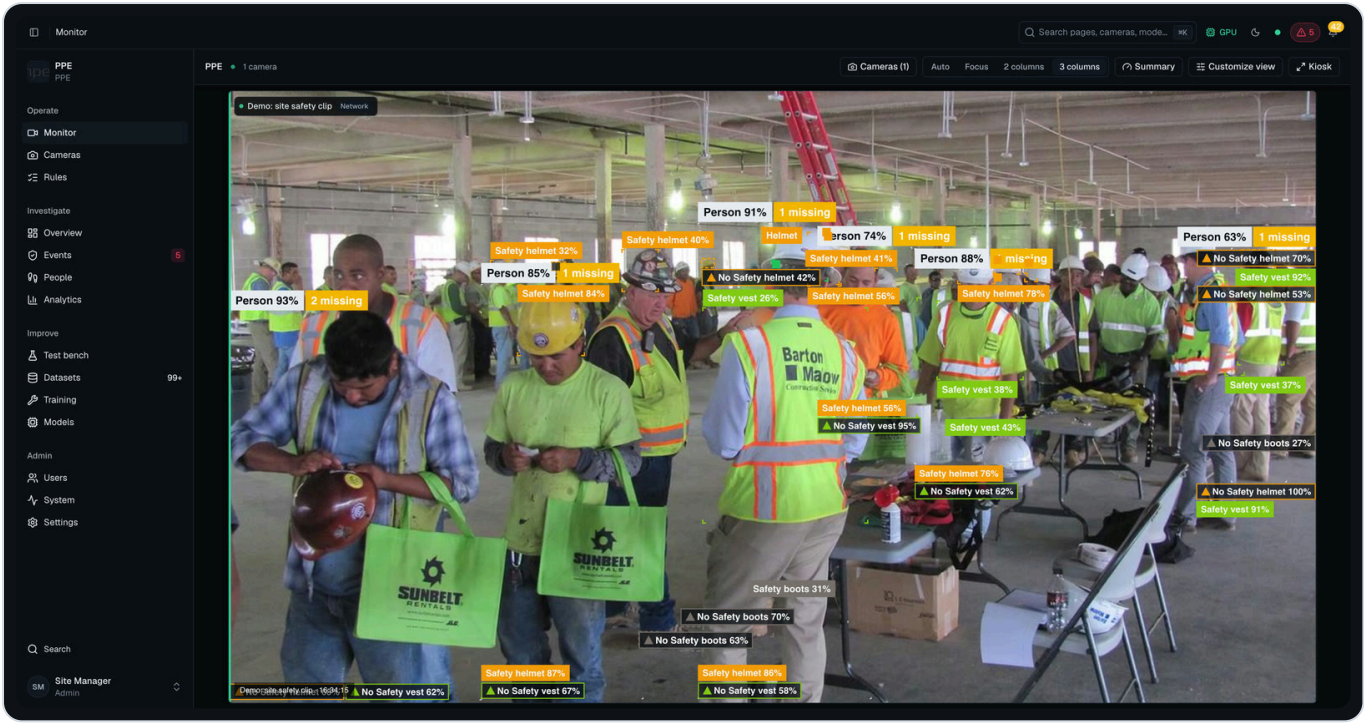
17 Administration

Reference

18 Troubleshooting

19 Glossary

20 How the finished system will look



The Monitor wall in the Control room look. Demo footage built from public-domain photos.

1 How the platform works

The platform is AI video analytics for the CCTV you already have. It runs on one server at your site, with no cloud, and watches every camera for missing protective equipment (PPE), fire and smoke, people in restricted zones, risky moves near forklifts, people and vehicle counts, number plates, visitors and more. You switch the checks on per camera. Each alert carries a picture and a clip, reaches the right people, and can become an incident to assign, resolve and hand over at the end of the shift. Video, faces and site data stay on the site.

The steps below follow the PPE check, the platform's core check. When someone lacks a required item for long enough, the platform records an event with evidence and tells the people who are watching. The other checks follow the same path from camera to evidence to review. People decide: an event is evidence for a person to review, and a new model never goes live until a person activates it.

1. **Cameras send video.** The platform reads each camera's stream and records it to the server's disk.
2. **People are detected and tracked.** Each person keeps one id across frames and cameras (chapter 13).
3. **PPE is checked against the zone's rules.** The zone decides which items a person needs and how serious a miss is (chapter 7).
4. **A violation is confirmed over a few seconds.** An item must stay missing for about 8 seconds by default, so one blurry frame raises no alarm (chapter 10).
5. **An event is created with evidence.** The platform saves two snapshots and a short clip, and the event appears in the bell and in Events.

6. **A person reviews it.** A reviewer marks the event as a real violation or not a violation (chapter 11).
7. **The model learns from reviewed frames.** Approved frames train a new candidate model, and a person decides whether to activate it (chapter 16).

The parts, for the person who installs it

Part	What it does
Web app	The screens you use in the browser. Listens on port 3000.
API	Sign-in and roles, settings, events, datasets, models, the training queue, and live updates to browsers. Listens on port 4000 and updates the database layout itself when it starts.
Vision worker	Reads the cameras, detects and tracks people, checks PPE, applies the rules, writes events, records video and picks frames for review. Uses the GPU.
Trainer	Builds dataset versions, trains candidate models when the site is quiet, and scores them on the golden set. Uses the GPU.
Postgres	Cameras, zones, rules, events, people, users, settings and model records.
Redis	Short-lived data only: heartbeats, live frames and the training queue. You can flush it at any time without losing records.
go2rtc	Relays live video to browsers over WebRTC.
Caddy	Optional. Puts HTTPS in front of the web app and the API.
Helper server	Optional. Another computer of yours that runs the checks of some cameras (Settings → Helper servers). It connects to the main server on port 7443.

Everything you would want to keep lives in Postgres and in one data folder on the SSD (`/data/ppe` in the Docker install). Settings such as thresholds, rules and zones are data edited in the screens. They reach the vision worker within 15 seconds and need no restart.

How accurate is it?

> ⚠ **Important.** Live cameras run a fast profile so that many streams fit on one GPU. On wide, crowded views that profile can miss small items such as helmets and vests, so a fully equipped worker can show as missing. The Test bench in its Accurate and Max modes finds far more of the PPE on the same footage. The biggest gain comes from training on labelled footage from your own cameras (chapter 16). The platform never activates a new model by itself.

Treat the first weeks as a tuning period. Review events, mark the wrong ones as not a violation, and let the review queue fill. Camera placement matters too: people who appear small, far away or seen from high above are the hardest to judge.

> **Note.** The platform is fully white-label. Your licence sets the product name, logo and accent colour (a licence that includes Own branding lets an administrator change them), and each site sets its own site name in Settings (chapter 4). The sign-in page, every screen and the browser tab use them. This manual calls it "the platform".

Where to start

Administrators
Chapters 4 to 8, then 16 and 17.

The person who installs it
Chapters 2 and 3, then 17 and 18.

Supervisors
Chapters 5 to 7 and 9 to 15.

Reviewers
Chapters 9, 11, 12 and 13, and the Review queue and Accuracy checks in chapter 16.

2 Before you install

Check the server, the cameras, the network and the browser before you start.

The server

Part	What you need
Operating system	Ubuntu 24.04 LTS (x86-64, with `sudo` access), Windows 10 or 11 (64-bit), or macOS on Apple Silicon.
GPU	On Ubuntu and Windows: NVIDIA, Turing generation (RTX 20 series, T4) or newer. Pascal (GTX 10 series) and Volta are not supported; on those, detection falls back to the CPU. On macOS: detection uses the Mac's own machine-learning chip.
Driver	Ubuntu and Windows: NVIDIA driver 580 or newer. The detection libraries use CUDA 13, which needs it.
Memory	16 GB of RAM or more.
Disk	An SSD for recordings, events, datasets and models, mounted at `/data/ppe`. You can mount it elsewhere and set `DATA_DIR`.
Address	A static address on the camera network.

Part	What you need
Internet	Needed while you install (Docker, NVIDIA packages, Python packages, container images). Afterwards the site can run offline, with two exceptions: the ready-made model catalog and the first start of face recognition download files.

Choose a server size

	Minimum	Recommended Best value	Best performance
For	Trials, demos and a few cameras	A typical site, training in idle hours	Many cameras, training while live, the thorough Test bench mode
Processor	6 cores	8 cores	16 cores or more
Memory	16 GB	32 GB	64 GB
GPU	NVIDIA RTX 2060 6 GB or T4 (driver 580+), or an Apple Silicon Mac with 16 GB	NVIDIA RTX 4060 Ti 16 GB, RTX 4070 12 GB or better	NVIDIA RTX 4090 24 GB, RTX 5090 32 GB or RTX 6000 Ada 48 GB
Storage	512 GB SSD	1 TB NVMe for the system, models and evidence, plus a recordings drive	2 TB NVMe plus a RAID recordings array
Network	1 Gb Ethernet	1 Gb Ethernet, cameras on their own network	2.5 or 10 Gb Ethernet
Operating system	Ubuntu, Windows or macOS	Ubuntu 24.04 LTS with Docker, or Windows 10 or 11	Ubuntu 24.04 LTS with Docker

Recordings take about 43 GB per camera per day at 4 Mbit/s, so size the recordings drive for the days you keep. The tiers name a use, not a camera count: **Run a load test** on the System page says how many cameras fit on your server (chapter 17).

Example: 4 cameras kept for 4 weeks

Recording bitrate per camera	4 cameras per day	4 weeks (28 days)	On a 4 TB drive
4 Mbit/s (typical 1080p, H.264)	173 GB	about 4.8 TB	Does not fit: a 4 TB drive keeps about 22 days
2 Mbit/s (typical 1080p, H.265)	86 GB	about 2.4 TB	Fits: a 4 TB drive keeps about 45 days

Estimates. They leave the default 50 GB of free space and assume continuous recording. Recording the sub stream or a lower bitrate stretches the days; the Space saving mode on the Storage page does that automatically before the disk fills (chapter 17).

Tip. Recordings use the most space. A 4 Mbit/s main stream fills about 43 GB a day, so four cameras kept for the default 14 days need about 2.4 TB. After the first days, read the storage numbers on the System page and adjust.

Note. The vision worker keeps a working copy of the camera addresses, passwords included, in `cache/vision_config.json` so it can start while the API is down. Only the service user can read the file. Consider encrypting the data disk (LUKS).

The cameras

Each camera needs an RTSP address and its login, reachable from the server. Use the camera's low-resolution sub stream for detection and its main stream for recording. The platform decodes the sub stream, and it records the main stream without re-encoding it. A camera that must find small or far people can read its main stream for detection instead (chapter 6). A Hikvision camera, for example, offers both:

```
Main stream   rtsp://user:password@192.0.2.10:554/Streaming/Channels/101
Sub stream    rtsp://user:password@192.0.2.10:554/Streaming/Channels/102
```

Before you add a camera, test the address from the server. If this command ends without an error, the address works:

```
ffmpeg -rtsp_transport tcp -i URL -t 5 -f null -
```

The platform does not judge people who are shorter than 80 pixels in the picture (a setting you can change). Aim and zoom each camera so workers appear large enough to show a helmet and a vest.

Network ports

Port	Used for	Needed when
`443` and `80` (TCP)	HTTPS through Caddy. Port 80 is used to obtain certificates.	You run Caddy, with a DNS name or with Caddy's own certificate authority.
`3000` (TCP)	The web app, straight on the LAN.	You run the `web-lan` service without Caddy.
`4000` (TCP)	The API, straight on the LAN.	The same LAN-only setup.
`8555` (TCP and UDP)	Live video on the wall (WebRTC).	You want live video on the tiles.

With a DNS name in `PUBLIC_HOST` that points at the static address, Caddy obtains a Let's Encrypt certificate. With an IP address or `localhost` it uses its own certificate authority, and you import `caddy_data/pki/authorities/local/root.crt` on each viewing computer. Viewers outside the LAN also need port 8555 reachable from outside, and the server's public address set as `webrtc.candidates` in `infra/go2rtc.yaml`.

The browser

The platform runs in the browser, so nothing is installed on viewing computers. Use Chrome 111, Edge 111, Firefox 128 or Safari 16.4 or later, on a computer, an iPhone or an iPad. Live video on an iPhone through the server needs iOS 17.1 or later. Three features ask for something extra. Live video on the wall needs WebRTC and port 8555. Kiosk needs the browser's full-screen mode. Browser notifications need you to allow notifications for the site.

3 Install

The platform runs on Ubuntu, Windows or macOS. Choose the route for the server you have. Other devices on the local network open the platform at the server's addresses or its `.local` name, found over Bonjour.

Windows | **macOS** | **Linux** | **Docker**

Route	Use it for	Full guide in the repository
Ubuntu with Docker Compose	A dedicated Ubuntu server with an NVIDIA GPU	`docs/INSTALL-LINUX-GPU.md`
Ubuntu with systemd	An Ubuntu server without Docker	`docs/RUNBOOK.md`, section 9
Windows or macOS, download installed as background services	A PC or an Apple Silicon Mac that runs the platform all day, with nothing else to install	`docs/UPGRADE.md`, and the `README.txt` inside the download
Windows, `Run Me.bat`	A Windows 10 or 11 machine with an NVIDIA GPU	`README.md`
Ubuntu desktop or VM, `Run Me (Linux).desktop`	A fresh Ubuntu or Debian machine without a GPU	`README.md`
macOS, `Run Me.command`	A Mac with Apple Silicon	`README.md`

Ubuntu: Docker Compose with an NVIDIA GPU

Plan one to two hours, most of it downloads. The steps below summarise `docs/INSTALL-LINUX-GPU.md`, which gives every command in full. First, know the five secrets you will create in step 4. Compose refuses to start while any of them is empty, and in production the API refuses placeholders and short values.

Variable	What it protects	How to make it
`POSTGRES_PASSWORD`	The database user `ppe`.	`openssl rand -hex 24`. Use hex only: the value becomes part of a web address, and characters such as `/ @ : # ?` break it.
`ADMIN_PASSWORD`	The first administrator login. Read only when the database is empty.	Eight or more characters, not `admin` or `admin1234`.
`SERVICE_TOKEN`	Calls from vision and the trainer to the API, and the health check.	`openssl rand -hex 32`
`GO2RTC_PASSWORD`	go2rtc's HTTP interface, which lists every camera's address.	`openssl rand -hex 16`
`APP_SECRET`	The key that encrypts camera passwords in the database.	`openssl rand -hex 32`

⚠ **Important.** Copy `APP_SECRET` to a password manager or the site safe as soon as you create it. A backup does not hold it, and a different value makes every stored camera password unreadable. Keep it away from the backup files.

1. **Install the NVIDIA driver, version 580 or newer.** Reboot, then check that `nvidia-smi` prints your GPU and a driver version of 580 or higher.

2. **Install Docker Engine, the Compose plugin and the NVIDIA container toolkit.** Then run this test. It must print the same table as the host. If it fails, nothing below can see the GPU, so fix it first.

```
docker run --rm --gpus all ubuntu:24.04 nvidia-smi
```

3 **Get the code and create the data folders.**

```
sudo apt-get install -y git jq postgresql-client
sudo install -d -o "$USER" /opt/ppe /data/ppe /var/backups/ppe
git clone <the repository url> /opt/ppe
```

4 **Create the settings file and fill in the five secrets.** Make each value with the command in the table. Also set `DATA_DIR`. For a LAN site served on port 3000, set `EXTRA_WEB_ORIGINS=http://<server address>:3000` so the API accepts sign-ins from that page. For a public name with HTTPS, set `PUBLIC_HOST` and `ACME_EMAIL` instead.

```
cd /opt/ppe
cp infra/.env.compose.example infra/.env && chmod 600 infra/.env
```

5 **Check the file before go-live.** The check flags weak or missing secrets and prints no values.

```
scripts/prod-check.sh infra/.env
```

6 **Start the platform.** The first build downloads PyTorch and the CUDA libraries, several GB, and takes 10 to 30 minutes. A new site downloads the pinned D-FINE small person detector by itself. To use your own person model instead, put it in `/data/ppe/models/base/`, or install one from the Models page later.

```
mkdir -p /data/ppe/models/base
scripts/compose.sh up -d --build postgres redis go2rtc api vision trainer web-lan
```

For HTTPS with Caddy, name `web caddy` instead of `web-lan`. Running `scripts/compose.sh up -d --build` with no names starts every service in the file. If `ufw` is on, open the ports from the table in chapter 2.

7 **Verify.** Every service should show `running` and Postgres `healthy`. The second command answers with `{"ok":true,...}`.

```
scripts/compose.sh ps
curl -fsS http://localhost:4000/api/health
```

Then sign in (chapter 4) and open System. Vision shows a recent heartbeat and the provider `torch-cuda`. If it shows `torch-cpu`, the driver is older than 580, the GPU is Pascal or older, or the container toolkit is missing. The vision log says which: `scripts/compose.sh logs vision`.

8 **Turn on the nightly backup.** It runs at 02:30 UTC by default (chapter 17).

```
scripts/compose.sh --profile backup up -d backup
```

> **Tip.** Measure before you promise numbers. Add cameras until a camera's checks per second fall below the rate you set, then note how many the GPU carried. The platform keeps up while each camera holds its rate, the Detection load meter on the System page stays below about 80 percent, and latency stays steady. Section 9 of the install guide has the commands.

If Compose stops with `bind source path does not exist: /etc/machine-id`, the server has no machine id. Run `sudo systemd-machine-id-setup` and start again.

Windows and macOS: background services from the download

The download for Windows or for a Mac with Apple Silicon installs the platform as background services that start with the computer. It has its own built-in database and key-value store, so nothing else is installed. On a Mac, detection uses the Mac's own machine-learning chip. Unpack the download, open an administrator PowerShell (Windows) or Terminal (macOS), go to the unpacked folder and run `.\bin\platformctl.exe install --from` (Windows) or `sudo ./bin/platformctl install --from` (macOS) followed by the path of the downloaded `.zip`. It checks the package, asks before it installs, and ends by showing the first administrator password once. On Windows, other computers can open the platform only after an administrator allows ports 3000 and 4000 in Windows Firewall; the `README.txt` has the one line for it.

Each system has an install file as well as its zip: the Windows setup (.exe), the macOS installer (.pkg) and the Ubuntu and Debian package (.deb). Open it for a new computer, or for a site installed from the zip. The Windows setup needs no Unblock step. These downloads are not signed. The first time, Windows shows "Windows protected your PC": choose More info, then Run anyway. If there is no Run anyway, Smart App Control is on and blocks unsigned programs. On a Mac, remove the download mark with `xattr -dr com.apple.quarantine` on the unpacked folder, or choose Open Anyway in System Settings, Privacy & Security. The `README.txt` in the download has each step. To move from the earlier one-click install, add `-`

`-adopt-r0`: it keeps the database and passwords. Later versions install from Settings, Updates, or with `platformctl update --from`, and `platformctl rollback` goes back for 7 days.

Windows: Run Me.bat

First unpack the download into a short folder outside OneDrive, for example `C:\ppe`: Downloads, Desktop and Documents are often inside OneDrive, and a deep folder makes Windows refuse the longest file names. The installer checks this in its first minutes and says the exact move if the folder is wrong. Then double-click `Run Me.bat`. It runs `scripts/dev.ps1`, which uses winget (or Chocolatey where winget is missing) to install bun, Node.js, uv, ffmpeg, PostgreSQL 17 and a Redis-compatible server, creates the settings file and the `ppe` database, installs the dependencies, and starts the API, the web app, the vision worker and the trainer. Logs go to `data\logs`. The first start compiles the web app and takes a few minutes. If a step fails, for example because a service needs Administrator rights once or winget is missing, the window stays open with the error. Fix it and run the file again. The installer generates the administrator password, shows it once in the window and keeps it in the settings file (the `ADMIN_PASSWORD` line).

macOS: Run Me.command

Double-click `Run Me.command`. A downloaded copy is blocked the first time: right-click it and choose Open, or run `xattr -d com.apple.quarantine "Run Me.command"` in Terminal. The script checks for Homebrew, bun, uv, ffmpeg, PostgreSQL 17 and Redis, offers to install what is missing, starts everything and opens the browser when the API is healthy. `bash "Run Me.command" --check` only reports what is missing.

Ubuntu desktop or VM: Run Me (Linux).desktop

Right-click `Run Me (Linux).desktop` and choose Allow Launching (once), then double-click it. It asks for your password once, installs ffmpeg, PostgreSQL, Redis, bun and uv, creates the settings file and the `ppe` database, and starts the platform. The first start downloads several gigabytes of packages and can take 10 to 20 minutes; later starts take seconds. It runs on the processor, so no GPU is needed.

Ubuntu: system services without Docker

Clone to `/opt/ppe`, install the dependencies, build the web app, copy `.env.example` to `/opt/ppe/.env`, install the units from `infra/systemd` and enable `ppe-api`, `ppe-web`, `ppe-vision` and `ppe-trainer`. Set `STORAGE_ROOT` to an absolute path, because the units start in different folders. Section 9 of `docs/RUNBOOK.md` has the details.

Updates

- **Upload a package in Settings.** An administrator opens **Settings → Updates** and installs a new version from its package file (chapter 17, Backups and upgrades).
- **Re-run a newer installer.** It upgrades the platform in place and keeps your data and settings. It also keeps the previous version beside the new one, so you can go back, and removes only older versions. On a computer that already had the platform's database from an earlier setup, the installer takes that

database over with everything in it; running it again repairs a computer where every page said it could not reach the server.

- **Roll back in Settings → Updates.** The earlier version is kept for 7 days. If a security update has set a lowest version, the Roll back button is locked and the page says why.
- **Online updates.** An administrator opens **Settings → Updates** and turns on **Check automatically** to look in an update folder every few hours, or presses **Check now**. A newer version for this server is downloaded, and its signature is checked before **Install** is offered. Install asks for the administrator's password. With **Install automatically** on, a checked update installs by itself in the window the administrator chooses, when no training, analysis or export is running. Both are off until an administrator turns them on.

Release packages ship no source code.

More guides in the repository

- `docs/INSTALL-LINUX-GPU.md`: a new GPU server from scratch.
- `docs/RUNBOOK.md`: operations, settings, models, storage and troubleshooting.
- `docs/DEPLOYMENT.md`: remote access through a Cloudflare Tunnel, and the security checklist.
- `docs/BACKUP.md` and `docs/UPGRADE.md`: backups, restores, upgrades and rollback.

4 First sign-in and first hour

This chapter takes a fresh install to a camera that checks PPE. Plan about an hour.

> ⚠ **Important.** A fresh install has one account, `admin@local`, with the password `admin1234`, unless `ADMIN_PASSWORD` was set before the first start. Change the password at once. Changing `ADMIN_PASSWORD` in the settings file later does nothing, because the account is created only once, when the database is empty.

1. **Sign in.** Open `http://<server address>:3000` on a LAN install, or `https://` and your `PUBLIC_HOST` name when Caddy is in front. Sign in with the administrator account. Tick **Keep me signed in** on your own computer so the platform does not ask again for the number of days it names; leave it unticked on a shared one. The password field warns you when Caps Lock is on, and a failed sign-in says what went wrong.
2. **Change the administrator password.** Open the account menu at the bottom of the sidebar and choose **Change password**. Enter the current password and a new one of at least 10 characters twice, then choose **Change password**. Your other sessions are signed out.

3. **Set your site name.** Open **Settings → General & branding**. Enter a **Site name** (shown in the top bar, hidden when empty) and choose **Save changes**. The **Product name**, **Accent colour** and **Logo** come from your licence. If the licence includes Own branding, an administrator can change them here (the logo is a PNG or SVG of up to 64 KiB) and **Reset branding** returns to the licence's own. Otherwise they show read-only, and you ask whoever supplied the licence. A **Brand preview** shows the sign-in page.

4. **Create users.** Open **Users** and choose **Invite**. Enter a name, an email and a role (chapter 5), choose **Create invitation** and send the person the link it shows, so they choose their own password. **Add user** sets a password for them instead. Give each person the lowest role that lets them do their job.

5. **Activate a PPE model.** Open **Models**. In the Model library, **Download** a ready-made PPE detector, then choose **Activate**. The first activation needs an override because no golden-set score exists yet: enter a reason such as "first install" (chapter 16). Without an active PPE model, the platform tracks people but judges nothing, and the System page shows the `no_ppe_model` warning.

6. **Add your first camera.** Open **Cameras** and choose **Add camera**. Enter a name, the main stream address and the camera's login. Add the sub stream address if the camera has one. Keep **Checks** at every check, **Checks per second** at 5 and **Record** on, then choose **Add camera**. The Status column changes to Online within moments. The platform creates a zone called "Whole frame" for the camera.

7. **Draw a zone** if parts of the picture need different rules. Open the camera, choose the **Zones** tab and choose **Add zone** (chapter 6). Skip this step if the same rules apply everywhere.

8. **Tick the required PPE.** Open **Rules**. In the camera's row, click the cells for the items that zone requires, such as Safety helmet and Safety vest. Use the arrow in a cell to set the severity (chapter 7).

9. **Check the Monitor.** Open **Monitor**. You should see the live picture with a tag on each person. To test without waiting for a real miss, step in front of the camera without a required item for several seconds, or add a stored video as a camera (chapter 6). An event appears in **Events** a few seconds after the item has been missing.

My account: two-step sign-in and signed-in devices

Choose **My account** in the account menu. **Two-step sign-in** asks for a 6-digit code from an authenticator app on your phone after your password. Choose **Set up two-step sign-in**, enter your password, scan the QR code (or type the key under it), enter the code your phone shows and choose **Turn on**. The platform then shows ten recovery codes: copy them or download them as a text file, keep them somewhere safe and tick the box before you choose **Done**. Each recovery code signs you in once if the phone is lost. **Make new recovery codes** and **Turn off** ask for your password and a code. An administrator can require two-step sign-in for administrators in Settings → Sign-in & access, and the platform then asks them to set it up before anything else. Wall-screen accounts sign in without a code.

Signed-in devices on the same page lists every browser signed in as you, with its address and device, and lets you sign one out or all the others. A one-time link for a reset or an invitation can be shown as a QR code, so it reaches a phone without typing.

If someone loses their phone and their recovery codes, an administrator opens **Users** and chooses **Turn off two-step sign-in** on their row. Every device of theirs is signed out, and they sign in with their password only until they set it up again.

If you cannot sign in

- **Forgotten password.** Choose **Forgot password?** on the sign-in page and enter your email. The page gives the same answer whether or not the address has an account, and your administrators are told. One of them hands you a one-time link: it lets you choose a new password and signs you out everywhere else. An expired or used link says so and lets you ask again. It needs no email server.
- **No account yet.** When the site allows it, choose **Request access** and enter your name, email and a password, with an optional note on who you are. A strength meter helps you pick a good password, and each problem shows under its own field. You see nothing until an administrator approves the request and picks your role. Until then, signing in shows **Waiting for approval** with the date you asked; the page checks every few seconds and opens the platform the moment you are approved.
- **A lost phone.** At the code step choose **Use a recovery code instead** and enter one of the ten codes you saved. If you have none left, ask an administrator to turn off your two-step sign-in.
- **A temporary password.** If an administrator gives you one, sign in with it. The platform then asks you to **Choose a new password** before anything else.

5 Finding your way

The sidebar groups pages by the job they do. You see only the pages your role can use.

Group	Pages	Use it to
Operate	Monitor, Floor map, Playback, Share safely, Cameras, Expected visitors, Counting, Vehicles, Conveyor, Rules, Sound	Watch live video, see where cameras are on the floor map, look back at recordings and export clips, check a file someone sends back, manage cameras and zones, count people and vehicles, see expected visitors and what the site hears, and choose the checks each zone needs. Counting, Vehicles, Conveyor, Expected visitors and Sound appear with their apps.
Investigate	Overview, All sites, Events, Incidents, Visits, People, Attendance, Reports, Video search, Ask the site	See what happened: a summary, every violation, the cases worked on, where people stayed, who was involved, who clocked in, a report on the trends, recorded moments found by description and plain questions about the site. All sites shows every site on one screen at head office.
Improve	Test bench, Accuracy checks, Datasets, Training, Models	Check detection, measure how often it is right, label frames, train models and choose the active one.
Admin	Users, Privacy, System, Storage, Licence, Settings, Apps	Accounts, privacy requests, health, disk space, the licence, site settings and which apps the site uses.

The menu and the command palette show only the apps that are switched on (Apps, in chapter 17). The AI pages (Rules, Events, Models) leave the menu when no AI app is on. A page of an app that is switched off says **This app is switched off**: administrators get an **Open Apps** button, and everyone else is asked to ask an administrator.

Press `⌘B` on a Mac or `Ctrl` `B` elsewhere to collapse the sidebar to icons. The Events entry carries a count of open alerts, and the Datasets entry carries the number of samples waiting for review.

Near the bottom, **Help** opens **Documents** (this manual, the quick reference and the other documents) in a new tab. It opens the copy your installation serves at `/docs`, so it works offline. A development build also shows a **Live demo**, a clickable preview of the finished system with demo data; an installed site shows none. **Search** opens the command palette, like `⌘K`.

What's new at the very bottom shows the version you are running, such as `v1.0.6`. Choose it to see what changed in each update, newest first. A release with a long title shows a short line while it is folded, ending in "and more"; open it to read the whole title. A dot on the icon means unread news.

The top bar

Item	What it does
Search (`⌘K` or `Ctrl` `K`)	Type to jump to a page, a camera, a model or a named person, or paste an event id or a person id such as `P-1A2B3C4D` to open it. Type a setting's name, or a word from its description, to open its section of Settings with the setting pointed out.
GPU or CPU chip	Where detection runs. It reads Offline when vision is off and Switching during a change. Administrators open it to choose Auto, GPU or CPU, and the Test bench mode, Fast, Accurate or Max (chapter 15).
Theme button	Switches between light and dark in one click. Kept with your account. On a phone it sits in the header on every page.
Status dot	Green: live updates are connected and vision is online. Amber: connecting, or vision is online without a working PPE model. Red: live updates are disconnected or vision is offline. After 5 seconds without live updates, a banner offers **Retry now**.
Open count	"N open": the open violations nobody has acknowledged or reviewed, the same number as the Events badge, the browser tab and Overview's Needs attention. Choose it to open Events on those rows. "Not checking" means nothing is being checked.
Alarm chip	A red "N alarms" chip appears while an alarm is open. Choose it to open the alarm's event (chapter 12).
Not enough compute	Appears when detection cannot keep up. It links to System (chapter 17).
Bell	The unread count and the list of live events (chapter 12).

Note. The alert count, the bell and the badge on Events count only what reaches your browser while a page of the platform is open. After a reload they start again from zero. For the full picture, use the **Open violations** and **Awaiting review** tiles on Overview, or the filters on Events.

Roles

Role	Home page	What the role can do
Admin	Overview	Everything: users, settings and branding, apps, models, cameras, zones and rules, events, datasets, training and system.
Supervisor	Monitor	Everything except Users, Settings, Apps, Licence, Models and Privacy. Manages cameras, zones and rules, uses the Test bench, reviews events and dataset samples, runs training, manages storage and reads the audit log.
Reviewer	Monitor	Watches, reviews events (real violation or not a violation, notes, delete), names and merges people, reviews dataset samples and answers the weekly accuracy check. Reads Monitor, Floor map, Playback, Share safely, Cameras, Expected visitors, Counting, Vehicles, Conveyor, Rules, Sound, Overview, Events, Incidents, Reports, People, Visits, Video search, Ask the site, Accuracy checks, Datasets, System and Storage, but cannot change cameras, rules, datasets or training.
Viewer	Overview	Reads Monitor, Floor map, Playback, Cameras, Counting, Vehicles, Conveyor, Rules, Sound, Overview, Events, Incidents, Reports, Video search, Ask the site, Storage and System, with no names or faces: no People, Visits or Expected visitors page, and no person on an event. Changes nothing.

Reviewers, supervisors and administrators can also export clips and snapshots from recorded footage and move PTZ cameras. Viewers can watch recordings but cannot export them or move a camera. Licence, Settings, Apps, Users and Models are for administrators. A Display account is always a Viewer (chapter 17).

Custom roles. When the four built-in roles do not fit, an administrator can make a custom role: open **Users**, choose the **Roles** tab, choose **Copy** on a built-in role or on one of your own, give it a name (1 to 40 characters), an optional description and a home page, then switch on exactly what it may do. A role can never give more than you can do yourself, and a site can have at most 20 roles of its own. Built-in roles cannot be changed, only copied. **Edit** changes a custom role, and you cannot edit or delete your own role. To **Delete** a role that people hold, choose the new role they move to first. A custom role is picked for a person in the same lists as the built-in ones.

Themes, density and home page

The account menu at the bottom of the sidebar switches the theme and the row density (Comfortable or Compact), opens **Your preferences** and **Notification settings**, and signs you out. After sign-in you land on a home page that depends on your role: Overview for administrators and viewers, Monitor for supervisors and reviewers. **Your preferences** sets your own theme, row spacing, **Date format**, **Clock** (automatic, 24-

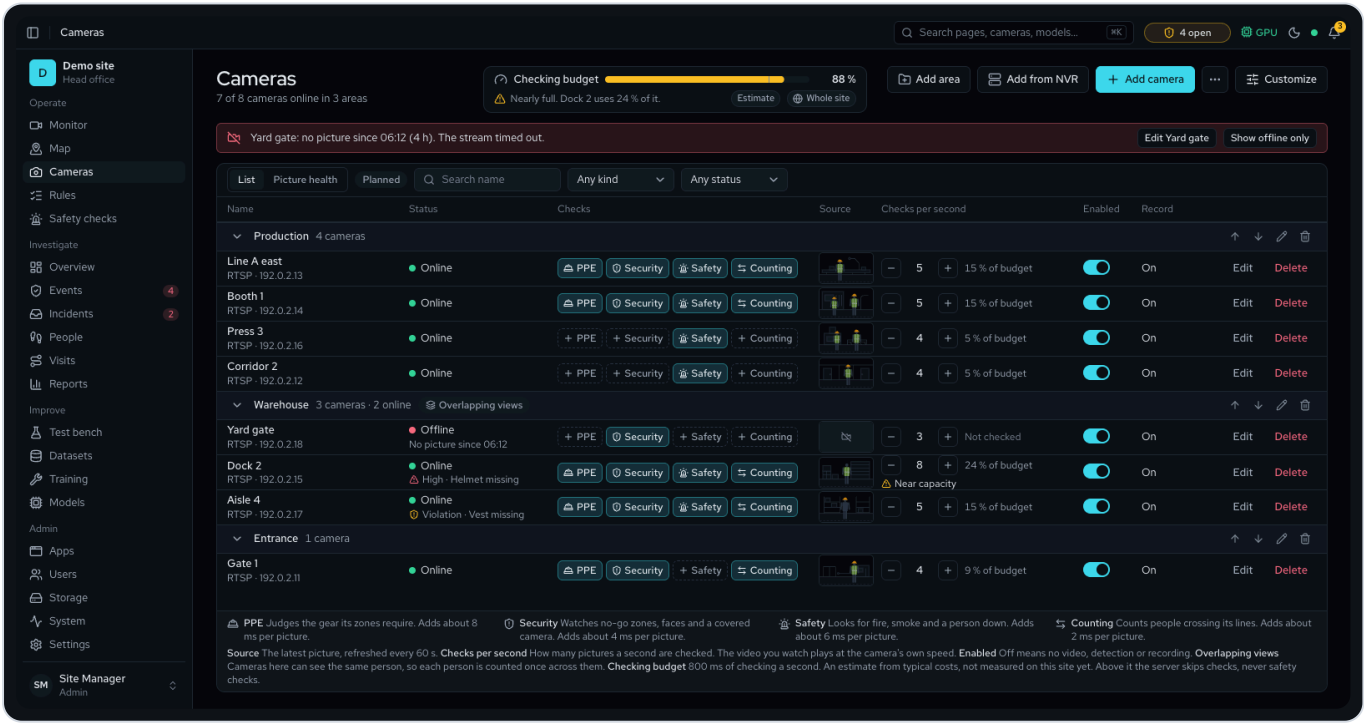
hour or 12-hour) and **Home page**. A new date or clock choice applies to each page as you open it. They are kept with your account, so they follow you to any computer, and so does your Table or Cards choice on Events. An administrator also reaches it from the **Your preferences** panel in **Settings → General & branding**.

The screens move only when something changes. A page fades in the first time you open it and does not replay on a second visit or when you change a filter. Lists, charts, markers and counts move only when their data changes. If your system asks for reduced motion, a change shows as a short fade instead. A page opens already filled in after you rest the pointer on its menu entry. During a burst of events the lists on screen refresh together about every 1.5 s, while the alert count and pop-ups update at once.

Every page's help opens with a map of the page: a small sketch with numbered parts, what each part does, and what changes when you use its main controls. Thirty words the pages use, such as zone, grace, golden set and legal hold, show a one-sentence meaning when you point at them. Most pages, from Overview, Events and People to System and each camera's page, have **Customize** for the panels they show and what they open on, and every control says what it does, its default and its range. Press ? on any page, outside a text field, for a map of the keys that work there and everywhere; the account menu has it as well.

6 Cameras and zones

A camera is a video source. A zone is a part of that camera's picture where rules apply. Cameras are grouped into **areas**, such as a warehouse or a production hall.



Cameras and areas Available Cameras grouped into areas, with the checks each one runs, how many pictures a second it is checked and what that costs, plus camera presets, a test connection and an NVR's channels added in one step.

Find cameras on the Cameras page searches the networks the server is on by itself, and lists the cameras that answer. A camera on another network or VLAN does not answer, so type its range in **Look in this range**, such as `192.168.1.0/24`, and choose **Look**. A range runs from `/24` to `/32` (at most 256 addresses). In the Docker install the search finds nothing by itself, because the server sees only Docker's own network: type the range. Choose a camera from the list and type its login: the platform reads its streams (the password is never shown again) and opens the camera form filled in. A recorder shows **Add all channels**, which fills the **Add from NVR** form. A camera you already added is marked **Already added**.

Add or edit a camera

Open **Cameras** and choose **Add camera**. To change a camera later, choose **Edit** on its row, or open it and use its **Settings** tab. Administrators and supervisors can add and edit cameras. The platform stores camera passwords encrypted with `APP_SECRET`.

Field	What to enter
Name	A name people recognise, such as "Gate 2".
Enabled	Switch it off to stop a camera without deleting it.
Source	Network stream, Webcam or Stored video.
Main stream URL	The camera's main stream, which the platform records. Its login goes in **User name** and **Password**: saving tells you if it is missing or refused.
Local address	Optional. Its address on the site network, used when the server can reach it.
Sub stream URL (optional)	The camera's low-resolution stream. The platform decodes it for detection.
Vendor preset	Hikvision, Dahua, EZVIZ or Generic. A preset fills the main and the sub stream address from the camera's address, port and channel. For EZVIZ the password is the 6-letter verification code on the camera's label unless you set your own, and RTSP (LAN live view) must be on in the EZVIZ app first.
Area	The group the camera belongs to. A camera with no area sits in a last group, "No area".
Checks	What the camera runs: a chip each for PPE, Security, Safety and Counting, all on by default. Switching one off keeps its rules (Checks on each camera).
Checks per second	Pictures a second the model checks, from 0.5 to 30. The default is 5, and 3 to 8 is enough.

Field	What to enter
Read the main stream for detection	Checks the full-size picture instead of the sub stream, so smaller and farther people are found. Off by default; it costs more processing.
Find far people and vehicles	Looks again at the far part of the picture at higher detail, for people and for vehicles. Off by default; a camera that uses it costs about 3 times as much work as one that does not (measured on a CPU server). Under a long overload the platform may run it on every second check until the load falls. It needs the person model's licence. On a camera with a sub stream it needs **Read the main stream for detection** first.
Record	Keep a continuous recording. Event clips are cut from it.
Retention	How many days recordings are kept, from 1 to 365. The default is 14. The site's privacy profile caps it: at most 30 days under the Strictest profile, and a save above the cap says what the cap is (Settings, Privacy).
Private areas	Parts of the picture that are always covered. On a saved camera, choose Add a private area, name it and click its corners on the picture. It is filled in solid in the checked picture with its boxes, in snapshots and event evidence, in new recordings and in anything shared. Live video played straight from the camera is not covered yet. Areas are saved as you add or remove them, without the form's Save. Because a stored recording has the cover burned in, the camera records with a video encoder instead of copying the stream, which uses more of the server. On a server with no video encoder such a camera cannot record: its health shows a recording error saying so, until the area is removed or an encoder is installed. In a recording a slanted area is covered by the rectangle around it, a little wider than in a picture. Removing an area shows that part again from then on; footage already recorded stays as it was.
Critical camera	Switch on for a camera whose recording must never be reduced or paused. Space saving and automatic clean-up leave it alone.
Record sound	Keeps this camera's sound with its recordings, evidence clips and exports. Off by default, and shown for a saved camera but not a USB webcam. Live view and Playback start muted. Only an administrator can switch it on. The window **Record sound from** the camera's name asks you to tick **I understand. Everyone in view of this camera will be told that sound is recorded.** and to give a reason of 4 to 500 characters under **Why this camera needs sound**, which goes in the audit log with your name. A privacy profile whose law needs everyone recorded to agree keeps the switch locked off, and so does the Strictest profile that applies until a country is picked. It acts at once, not on Save, and switching it off needs only permission to edit cameras.

Field	What to enter
Where to run its checks	Shown only when the site has a helper server. **Automatic** lets the system choose; pick a server to pin the camera's checks to it. The camera is still recorded on the main server. Administrators only (Helper servers).

The Cameras page groups each camera under its area, and lists its source, status, checks and checks per second. Status reads Online, Starting, Offline, Disabled or No heartbeat: hover it for the last frame or error. It reads **Password needs re-entering** when the saved password can no longer be read (for example after the server's secret changed): the platform stops calling that camera until you type the password again. Changing a camera's port, or switching between rtsp and rtsps, also asks for the password again. If someone else saved the camera while you edited it, you get a "changed meanwhile" message. Deleting a camera removes its zones, rules and events, so you type `delete` first.

Test connection in the camera form asks the vision service to open the stream. It reports the codec, size and frames per second of each stream, or why it failed, and says whether the local or the outside address answered. When a sub stream is set, it compares the main and the sub stream. If they show different areas, the camera keeps its pictures and the form says why. The password is never shown. If the vision service is not running, it says "Vision worker offline".

A sub stream uses a second session on the camera, which the form notes beside the field. **Use the sub stream on busy walls** (**Settings → Monitor defaults**) keeps the walls off that second session.

Checks on each camera

Each camera row has a chip for every check whose app is on: **PPE**, **Security**, **Safety** and **Counting**. A solid chip is on and a dashed one is off. A check whose app is off, or not in the licence, has no chip. Switching a check on says what it adds to each picture, such as "Safety is on for Dock 2. Adds about 8 ms per picture (estimate)." Switching one off asks first, and says its rules are kept and apply again when it is switched back on. A viewer sees the chips locked, with the reason under them. Switching the Safety app off on the Apps page removes the Safety chip; the camera's stored Safety check comes back with the app. The Add or edit camera form has the same chips, under **Checks**.

Checking budget and slowed cameras

The **Checking budget** bar above the list shows how much of the server's checking time is in use, measured over the last 30 seconds. It turns amber at 80 % and reads **Nearly full**, and red over 100 % and reads **Over budget**. Its sentence says how much room is left, such as "Room for about 3 more cameras like these." The tag reads **Measured** once the server has measured its checks, and **Estimate** until then.

A camera's rate shows what it is set to. When the server cannot check it that often, the rate reads **Checked 3.4 of 5 per second** in amber, with the reason under it: the server is over its checking budget, or the camera sends fewer pictures than it was set to. **Set to 3.4** beside it lowers the rate to what the server manages. The banner above the list covers every slowed camera at once with **Set to what fits**. Either change can be undone: a message offers **Undo** for a few seconds and puts the old rates back.

Running now

Running now (in the Cameras header; the same panel is on Apps and Models) lists each model the AI engine holds in memory: its milliseconds a picture, the cameras it serves, **Change** to pick another model, and its memory. A bar shows the server's memory, or the graphics card's memory where there is one, split into models, other programs and what is free. A model that does not fit in the memory kept free (`scheduler.memoryReserveMb`, Memory kept free under Settings → Detection, advanced; 512 MB by default) is not loaded. The panel says why in one sentence, such as "Pose needs about 0.3 GB, 0.2 GB free. Turn off Smoking.", and a check that needs that model shows the same sentence on the Rules page.

The licence limit

The licence sets how many cameras may run (Licence). At the limit, **Add camera** and **Find cameras** switch off and say "Camera limit reached". A camera beyond the limit is paused, not deleted: it shows **Paused — licence limit** in the list and on its page, and its Monitor tile says "The camera limit of the licence is used up, so this camera is not being watched or recorded." No live feed opens, and a paused camera raises no "is offline" alert. The newest cameras pause first, and the oldest keep running. An administrator can switch on **Keep active** on a paused camera in the Cameras list: it then runs, and the newest of the others gives up its place. Switch it off to return to the usual order.

Areas

Choose **Add area** on the Cameras page to create one. You can rename, reorder, fold and delete areas there. Deleting an area keeps its cameras and its rules: the cameras move to "No area". Two areas cannot have names that differ only in capital letters. While a search or filter is on, folding and moving areas is switched off.

Switch on **Overlapping views** for an area whose cameras can see the same people. The Monitor then counts people from the camera that sees the most and shows the most there could be, such as "3 up to 5"; it does not guess who is the same person. A violation both cameras see becomes one event. The event page shows the second camera's pictures under **Also seen by**. A person with no recognised id, or missing different items, still gets a separate event. The **Area** filter on the Monitor, the Overview and Reports narrows each page to one area.

An area also has **Most people allowed**, from 1 to 5000: a warning opens when more people than that are in the area for a minute (empty means no limit). **Cameras see the whole area** (off by default) says the cameras between them cover all of it, so the number inside is counted from the pictures, not from people crossing lines.

The Cameras page shows a banner for an area or camera that needs attention, and **Customize** changes how the groups show. Each camera's form has an area picker. On the Monitor, a camera you pinned to the wall stays pinned when you change the area, and **Needs attention** ignores the area.

Add cameras from an NVR

Choose **Add from NVR**. Enter the recorder's address, port and login, a vendor preset, the area and the first and last channel (1 to 64), then choose **Find channels**. The platform tests four channels at a time and

unticks those that do not answer. **Add** creates every ticked channel, each with a whole-frame zone, in one step, or none of them: if a channel was already added, the dialog names it. This needs the vision service to be running.

Sources

Source	Address form	Notes
Network stream	`rtsp://` or `rtsps://`. Also `http://` and `https://` for MJPEG or HLS-style streams.	The normal case. Only network streams play as live video; the others show pictures.
Webcam	A device index such as `webcam://0`. Under Advanced, a device path such as `/dev/video0` on Linux.	macOS asks for camera permission the first time. On Linux the vision user needs access to the device.
Stored video	A file uploaded to the platform. **Loop** replays it when it ends.	Without Loop, the camera goes offline at the end of the file.

Upload stored videos under **Cameras → More camera actions → Stored videos**. The platform accepts mp4, mkv, mov, avi and webm files of up to 4 GiB, and keeps a file a camera still uses. An upload shows a bar and a percentage. A stored video behaves like a live camera: it raises real events and records, so use it to test detection and rules.

Zones

Every camera gets a zone called **Whole frame** that covers the entire picture. Add more zones where the rules differ, for example a loading bay that needs a vest and a workshop that needs goggles as well. Zones can be drawn whenever an AI app is on, not only PPE Compliance.

1. Open the camera and choose the **Zones** tab.
2. Choose **Add zone**. Click on the snapshot to place the corners of the area. A zone needs at least 3 points and can have up to 64.
3. Press `Enter`, double-click, or choose **Close polygon** to finish. Press `Esc` to cancel. The platform creates the zone at once and names it "Zone 1", "Zone 2" and so on.
4. Or choose **Click to outline** and click once on a floor, walkway or door (or move the cross with the arrow keys and press `Enter`). A dashed outline is suggested. **Accept** creates the zone, **Adjust** lets you drag its corners first, and **Reject** or `Esc` drops it. Nothing is saved before you accept. **Click to outline** uses the SAM, SAM 2, MobileSAM or FastSAM model you activate on the Models page, from the next click and with no restart; MobileSAM is the light choice on a server without a GPU.
5. Select the zone in the list on the right. Type a new name in the **Zone name** field and click away to save it. Set its **Grace** there too.

6 To change the shape, drag the corners, or choose **Redraw**. Then choose **Save**, or **Discard** to drop the change. **Refresh snapshot** fetches a new picture from the camera.

Where zones overlap, the smallest zone that contains a person's feet decides which rules apply. **Grace** is how many seconds a person must be inside a zone before the platform judges them. The default is 2 seconds, and the range is 0 to 60. A zone you switch off with **Enabled** shows an Off badge. Deleting a zone deletes its rules too. If the picture area shows "NO SIGNAL", vision has no frame for that camera yet (chapter 18).

Fixed items

With the Safety Checks app on, a camera has a **Fixed items** tab. A fixed item is something that looks like fire, smoke or a forklift but never changes, such as a red lamp or a steam vent. The platform learns one when it keeps triggering a check in the same place; you can also choose **Draw manually**, type a **Label** such as "steam vent" and click the corners of the area to ignore. Each item shows a thumbnail and a status. **Confirm** keeps it, **Reject** says this place may really have a fire, so the platform never learns here, and **Remove** deletes it. Editing needs permission to edit cameras. The tab opens only while the app is on.

An item hides only a still copy of itself. While the platform is still learning it, and for the first 10 minutes, it raises only a low-severity event with no alarm. After that it ignores the item while it stays still. A check alarms again when the box grows, another box overlaps it, it flickers, or, for smoke, its top edge rises. An item not seen for 7 days expires. Reports show how many detections each camera's fixed items ignored (chapter 14).

Floor calibration

Every camera's **Settings** tab holds a **Floor calibration** panel, which turns positions and distances on that camera's floor into metres. The **Person near forklift** check needs it to measure in metres. Click four to eight spots on the picture that you can find on the floor, then give each its position in metres; **Suggest points** can place them for you.

Start from a rectangle

Measure a floor rectangle is one way to start: click the four corners of a floor rectangle you can measure, type its **First side (corner 1 to 2)** and **Second side (corner 2 to 3)**, and choose **Use these points**.

Check the points

Each point shows its error, and a **Tape check** between two spots you measured gives a plain accuracy figure. Points that cannot be a calibration are refused in plain words.

Remove the calibration

Remove calibration says "Calibration removed: distances are approximate", and the forklift check goes back to the gap between boxes. Calibrations made before this update carry over.

Accuracy and size limits

A camera's own page (open it from **Cameras**) holds an **Accuracy and size limits** panel. Its first line gives the camera's accuracy grade and the day it was checked, such as "Accuracy: Counting": **Counting** (good enough for counts, presence and PPE), **Presence and PPE only**, **Unreliable** or "not yet checked". A grade taken from picture size alone says "(from picture size)". The grade comes from the accuracy checks (chapter

16). A second line says how much of the camera's zones is big enough to judge, such as "Usable area: 62 % of its zones. People are under 40 px in the rest."

The panel also holds size limits for that camera alone: the smallest person, in pixels tall, each check accepts. When the camera's own check set suggests a presence limit, "Suggested from this camera's check set" offers **Use it**. **Save size limits** needs permission to change cameras.

Limit	Default	Range
Presence	40 px	10 to 400
Counting line	60 px	10 to 400
Head and torso items	80 px	10 to 2000
Hands and feet items	140 px	10 to 800
Identity	100 px	10 to 800
Face	40 px	16 to 200

Each day the platform compares a camera's picture with a saved one. If the view has shifted, or the camera's stream changed, its accuracy reads "not yet checked" and an alert asks you to check its zones and collect a fresh check set. Before a new model, a counting setting or a camera's size limits takes effect, it is replayed on the camera's saved check sets. A camera whose new figures do not hold up loses its grade at once, and saving without that check also leaves the camera ungraded until it is checked again.

Counting lines

With the People Counting app on, a camera has a **Counting lines** tab. Give a line a **Name**, such as Main gate, and click 2 to 8 points along where people walk through; the arrow points the way people are counted in. **Counts into area** and **Counts out of area** name the areas people enter and leave when they cross in. Each line shows today's counts, such as "Today: 42 in, 39 out", and can be switched off. The error of these counts is not measured yet. The line settings, such as **Smallest person counted at a line**, are under **Settings → Counting & coverage**.

Counting page

With the People Counting app on, **Counting** (Operate; the page title is People counting) shows who came in and went out, and how many people are inside each area now. Where the cameras of an area overlap and are joined on the floor map, a person two cameras see counts once and the figure is one number; if they are not joined, it stays a range and the row says why. Four tiles lead: **Counted in**, **Counted out**, **Inside now** and **Busiest hour**. Below them are a bar for each hour of people coming in and going out (a dash is an hour no counting camera was watching), **Inside now** for each area against its **Most people allowed** limit, and a table of **Counting lines**. An area over its limit for a minute raises an alert.

- **Period and Area.** Choose **Today**, **Yesterday**, **Last 7 days** or **Last 30 days**, and one area or **All areas**.

- **Download. CSV, per line**, **CSV, per hour** or **PDF report**, for the period and area on screen.
- **Limits of the figures.** A line checked against a careful count shows how far off it can be, and its tile reads like "1,240 ± 74" with how many lines were checked. A line not yet checked keeps its count and says "Error not measured", and so does any total it enters. Crossings the platform could not place show as a higher end ("up to 1,262 with 22 unsure crossings"), a line that tends to read low or high says so, and a line off by more than 20 % stays in the table but leaves the totals and is named under them. **Count crowds** is Planned. Until a line exists the page says "No counting lines yet" and offers **Open Cameras**.

Vehicles

With the Vehicles & Plates app on, **Vehicles** (Operate; the page title is Vehicles and plates) counts cars, motorbikes, buses, trucks and bicycles crossing a line on a camera, in and out. It works with or without PPE detection, but it needs a person model that can see vehicles: if the one in use cannot, the page says "Vehicles cannot be counted yet" and links to Models.

- **What it shows. Vehicle crossings**, **Wrong-way alerts** and **Vehicles by type**, for **Today**, **Last 7 days** or **Last 30 days**, for every line or the one chosen in **Line**. Each pass over a line counts once, so in plus out is crossings, not vehicles.
- **Draw a line.** Administrators and supervisors choose **Set up lines and bays**, pick a camera and draw a line across a gate or road on its picture. **Allowed direction** is **Both ways**, **Only in** or **Only out**, and a vehicle going the other way raises a wrong-way alert. Switch **Ring on wrong way** off and such a vehicle is still logged and counted, but no alarm rings. Parking bays are drawn in the same place: choose the **Parking bays** tab, and the page shows each bay taken, free or reserved.
- **Boxes and reports.** Playback draws each vehicle's box and type on video from a camera with a vehicle line. **Reports → Vehicles** gives crossings by type and hour on each line, and every wrong-way alert, as a landscape PDF or a CSV.
- **When it cannot count.** A dash is not zero: it means nothing was checked, and the page says why, for example "Camera offline since 10:12". "No vehicles in this period." appears only for a period the platform watched and saw none.
- **Events.** A wrong-way crossing appears in Events as **Wrong way** (Medium severity), with an alarm unless it is switched off for that line.
- **Settings. Settings → Vehicles** sets **Smallest vehicle counted** (60 px wide). **A bay is taken after** (20 s) and **A bay is free after** (60 s) set how fast a bay changes. **Settings → Vehicles** also holds the plate layouts: add the layouts your site uses, such as AA AAA-9999 (A is a letter, 9 a digit, # either), and each read is checked against them; with none listed, any plate text is accepted.
- **Number plates.** The platform can read the plate of each vehicle that crosses a gate line, on the computer at your site. First install the **Plate reader** card on the Models page and activate it (Latin letters and digits only); until then **Read plates here** cannot be switched on. Then, on a vehicle line in **Set up lines and bays**, switch on **Read plates here** (off by default; it uses the camera's main stream, so it needs more detail and more computer time). A plate is accepted only when several pictures of the vehicle agree, and a faint or too small plate is shown as **Plate not read**.

- **Gate events.** Each pass with a plate appears under **Gate events**. A plate on a deny list, an unknown plate or a near match shows **Plate waiting for a guard** until someone chooses **Let in** or **Turn away**. That choice is recorded and moves no barrier. A reviewer can correct a misread plate. While a plate waits, the Monitor's automatic layout brings that gate camera forward.
- **Plate lists. Plate lists** holds an **Allow list**, a **Deny list** and a visitor list, with groups, end dates and reasons, and can be imported or exported as a spreadsheet (CSV).
- **Who sees plates, and privacy.** The permission **See number plates** (administrators, supervisors and reviewers, not viewers) shows plate text and pictures and allows the plate search; without it a plate shows as **Hidden**. **Edit plate lists** changes the lists. Plate pictures are kept for 30 days by default, and a plate on no list is kept as text unless you turn that off. A plate search, an opened plate picture, a list change, an import and an export are written to the audit log as a hash of the plate, never the plate itself.
- **Erase a plate.** An administrator can erase one plate everywhere: its passes with their pictures, its list entries and the plate on events, where the text becomes "Erased". An event under a legal hold keeps its plate, and the answer says how many were held back. A privacy request about a plate runs the same erase and exports that plate's data. The event's own snapshot may still show the plate.

Conveyor

With the Conveyor Counting app on, **Conveyor** (Operate) counts the items on a belt with the camera above it. Each belt line shows the **Batch** count, the **Total** and the batch number. The batch count goes back to 0 when a new batch starts. The total keeps going.

- **Draw a belt line.** Administrators and supervisors open the camera, choose its **Counting lines** tab and, under **Belt counting**, draw a line across the belt, then a belt area that leaves out the belt edges. The line needs an Object model that can find the item; the page says so when it has none.
- **New batch.** Press **New batch** when a new lot starts. A second press within 10 seconds is refused, so a double tap starts one batch. The button needs permission to review events. The operator panel's own New batch button does the same.
- **Open count screen.** The button opens the count in large numbers for a monitor. A display account sees it with no New batch button. **Not live** means the number may be out of date. **Not counting** means the camera or the analysis is stopped. Do not trust the number until the label goes away.
- **Batch report.** The **Batch report** menu gives the batches of a period as **CSV data** or a **PDF report**. **Recent batches** lists the latest ones. Enter the items counted by hand for a finished batch with **Enter** (or **Change**), and the page shows the measured error. Until batches have been hand counted it says "Error not measured".
- **Operator panel. Settings → Operator panel** (administrators only) sets the **Panel link** (**Off** or **Serial**), the **Serial port**, **Speed**, **Device number** and the **Register map**: the registers that hold the batch count, the total, the batch number, the status and the New batch request. Check the map on the panel before you switch the link on. The panel keeps its last number while the computer or the analysis is off; it is not live then. A cable problem never stops counting, and the page says the link has a problem.
- **What is kept.** Every counted crossing is kept for 7 days, so a resent batch of counts is never added twice. Batches are kept. Neither holds pictures.

Floor map

Floor map (Operate) shows the site as one place: where the cameras are and what they see. It opens in 3D, with 2D and a flat plan when you need them, and a **Floor** list when the site has more than one floor. With no map it says "No floor map yet" and offers **Start a map**. Administrators and supervisors also see **Edit in the studio**, which opens five steps. Until a camera is calibrated, people show as a count on that camera.

1. **Plan.** Choose **Choose a file** for a picture or a PDF page of the floor plan, or **Start drawing**. The tools are Select, Wall, Room, Stairs, Lift, Fixed item, Door and Label. Mark a line you know the length of to set the scale. Measure again at any time and the whole draft is rescaled. A draft with walls, doors or cameras already drawn on an unscaled picture is rescaled with the picture when you set its first scale.
2. **Cameras.** Pick a camera and click the plan to drop it, choose **Aim: click where it looks** and click, and set how far it sees. A camera already on another floor says so and cannot be dropped twice. **Take off the map** removes one.
3. **Calibrate.** The camera's picture sits beside the plan, so clicks on the plan give floor points in metres. A camera calibrated on its own can be fitted to the plan with **Fit to the plan**, by dragging its outline or clicking two matching points. **Calibrate again on the plan** starts over.
4. **Zones.** Choose **Open the Zones screen** to draw a zone once and have it apply to every calibrated camera that sees it.
5. **Publish.** Choose **Review and publish**, then **Publish to the live map**. Nothing on the live map or in detection changes until you publish.

A camera on no floor map shows **Not on the map** in the camera list, with a link to place it. Editing the map needs a screen wider than a phone's. Maps made from a photographed sketch and a live map of people are In progress.

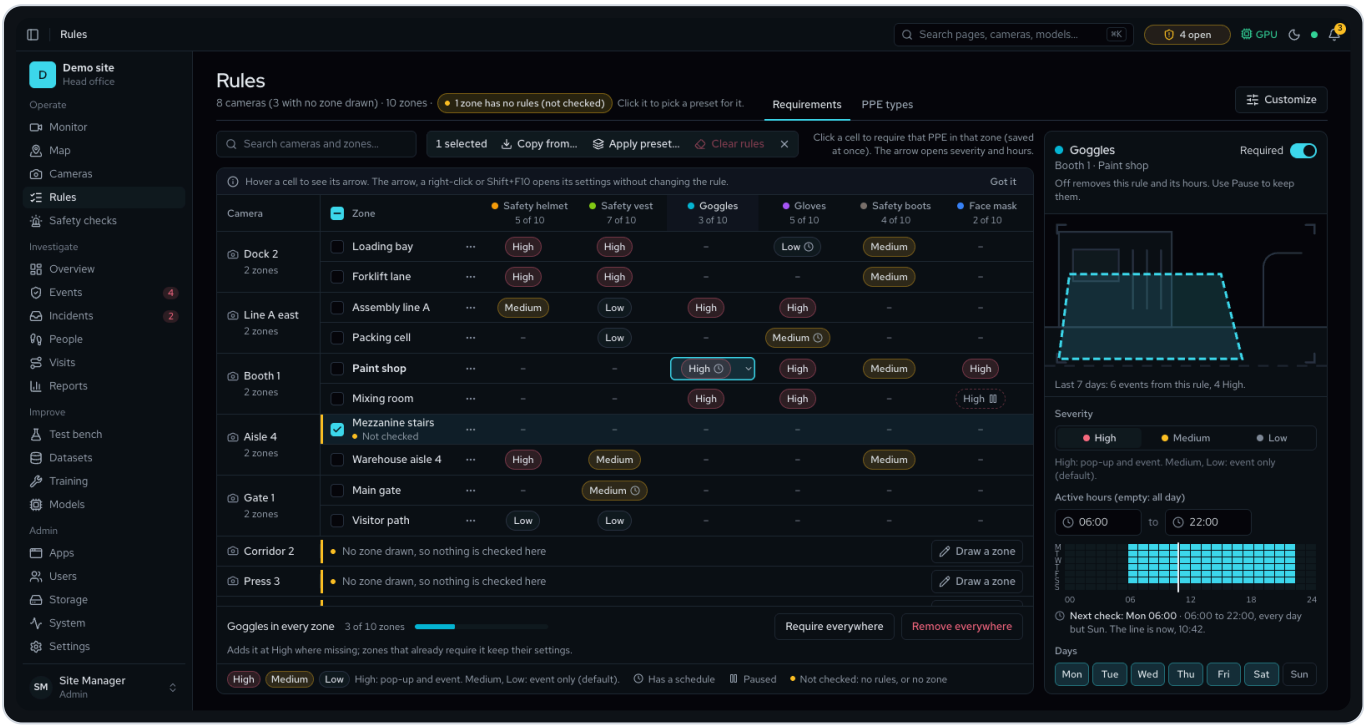
Picture health

Cameras, then **Picture health**, shows each camera's picture against the reference picture a person saved for it. Each camera has one state: **Good**, **Blurred**, **Blocked**, **Moved**, **Dark** or **Frozen**. The worst cameras come first. Checks run once a second, and a state is reported once it has held for 10 seconds.

- **Save a reference.** Choose **Set a new reference** and confirm. Only a person can save one, and it never refreshes itself. A blurred, blocked, dark or frozen picture cannot be saved.
- **Thresholds.** The **Day threshold** (30 % by default) and the **Night threshold** (45 %) say how far the picture may differ from its reference before an alert. Each one is set between 10 and 80 %, in steps of 5.
- **Infrared.** The night threshold applies while the camera is in infrared.
- **Alerts.** A picture problem shows on Overview, on the Monitor and in the bell, as "camera name picture: state".

7 Rules and PPE types

A rule says that one PPE type is required in one zone. Each rule has a severity and an optional schedule. A zone with no rules checks nothing, and the Rules page marks it "Not checked".



Rules Available The PPE each zone requires, with a severity and a schedule per cell, edited in one matrix.

Set requirements

1. Open **Rules** and stay on the **Requirements** tab. It lists every zone of every camera, with one column for each enabled PPE type.
2. Click a cell to require that item in that zone. Click again to remove the rule.
3. Use the arrow in a cell, or right-click it, to open its settings beside the table: **Required**, **Severity**, **Active hours**, **Days**, **Warn after (seconds)**, **Violate after (seconds)** and **Paused**, which keeps a rule but stops checking it.

Many zones at once

- A cell's settings offer **Require everywhere**, **Set all to** its severity and **Remove everywhere** for that PPE type.
- Tick rows, then use **Copy from...**, **Apply preset...** or **Clear rules**. The presets are Basic site (helmet, vest, boots), Battery / chemical area (helmet, goggles, face mask, gloves, boots, coverall) and All enabled PPE types. A preset applies only the types that are enabled.

- The menu on a zone offers **Copy this zone's rules to...**, **Import rules from...**, **Edit details...**, **Open camera** and **Clear all rules**.
- Copying, importing and presets replace a zone's rules. They do not merge with them.
- Arrow keys move between cells, `Space` toggles a cell, `Shift` `F10` shows its settings and `Esc` returns to the cell.

Each camera's **Rules** tab shows the same rules one zone at a time, with a **Required** switch per type. Reviewers and viewers see the page but cannot change it.

Severity

Every rule is High, Medium or Low. An event takes the highest severity among its missing items. By default only High pops up (chapter 12), and Medium and High count in the alert number.

Schedules

A rule can apply only at certain times. In the cell's settings, set **Active hours** and the **Days** of the week. Set both a start and an end time, or leave both empty to check all day. Hours may cross midnight, such as 22:00 to 06:00. Leaving all seven days on means the rule never pauses.

> **Note.** Schedules follow the site's time zone, set under **Settings → Site**, not the clock of the machine that runs the vision worker. If that machine cannot read the zone, the System page says so and the schedules follow its own clock.

Warn after and violate after

Each rule can set two times. **Warn after** (0 seconds by default) is how long a person may miss the item before they show as a warning. **Violate after** (the site's violation time, 8 seconds by default) is how long before an event is raised. Leave a field empty to use the default. Warn after cannot be later than violate after. The Rules page groups zones by the area their camera is in.

PPE types

Open **Rules → PPE types** to see the gear the platform can check for. A disabled type is ignored in every zone. The first six types are on by default. The last six are available but off, until you switch them on and have an active model that detects them. **Start from a preset**, beside **Add type**, applies a ready-made set of types. It names what it will switch on and off before you confirm. **Construction and industry** is Safety helmet, Safety vest, Safety boots, Gloves and Goggles. **Hygiene & food safety** is Hairnet, Face mask, Gloves and Apron, for kitchens and food plants. A preset re-creates any of its types you deleted, keeps every zone rule, and is recorded once in the audit log. If the PPE Compliance app is off, it answers "The PPE app is

switched off." No ready-made model detects hairnets or aprons yet, so train one or map a model's classes to them.

PPE type	Key	Body region	Colour	On by default
Safety helmet	`helmet`	Head	`#f59e0b`	Yes
Safety vest	`vest`	Torso	`#84cc16`	Yes
Goggles	`goggles`	Face	`#06b6d4`	Yes
Gloves	`gloves`	Hands	`#a855f7`	Yes
Safety boots	`boots`	Feet	`#78716c`	Yes
Face mask	`mask`	Face	`#3b82f6`	Yes
Coverall	`coverall`	Torso	`#10b981`	No
Ear protection	`ear_protection`	Head	`#ec4899`	No
Face shield	`face_shield`	Face	`#6366f1`	No
Safety harness	`harness`	Torso	`#ef4444`	No
Hairnet	`hairnet`	Head	`#d946ef`	No
Apron	`apron`	Torso	`#0d9488`	No

Choose **Add type** to create a custom type. Give it a key (lowercase letters, digits and underscores, starting with a letter), a label, a colour and a body region. A new type starts enabled and is off in every zone until you require it. Pick a key that matches a class your model predicts, or map that class to the type in the model's class map (chapter 16). Each type's menu also offers **Edit**, **Require in every zone** and **Delete**. Deleting a type stops every rule that requires it.

Safety checks

Rules has four tabs: **Requirements**, **Safety checks**, **Security checks** and **PPE types**. Open **Rules → Safety checks** to switch other checks on for each zone. Zones run down the side and the checks across. A check shows only when its app is switched on and in the licence (Apps); when none is, the tab says **No safety checks are available**. Switch a cell on, then use its sliders button for **Severity**, **Alarm (siren and alarm channels)**, the check's own values and **Active (empty = always)** hours and days. A refused change reads

as a plain sentence. A check that needs a model nothing has loaded shows **Needs a model · Get model**, a link to **Models**. Each check has its own colour, used on the Monitor, in Events and in Reports.

Check	App	What it finds	Starts as
Fire	Safety Checks	Flames in the zone. Needs a fire and smoke model. It must hold for 3 seconds.	High, alarm
Smoke	Safety Checks	Smoke in the zone; steam can look like smoke. It must hold for 10 seconds and move or drift in at least half of the pictures. Smoke that never moves is logged as possible smoke, with no alert.	High, alarm
Person down	Safety Checks	A person lying or fallen for 2 seconds. The **Method** is a pose model with a rule (a lying box at least 1.0 times as wide as it is tall) or a person-down model.	High, alarm
Person near forklift	Safety Checks	A person within 2 metres of a forklift. Metres need a floor calibration; without one the gap is approximate.	High
Phone use	Safety Checks	A phone held in a person's hand. It needs no extra download.	Low
Smoking	Safety Checks	*Experimental.* A cigarette at a person's mouth or hand. A phone or food held there cancels it. Test it on your site before you rely on it.	High
Occupancy limit	People Counting	More people in the zone than **Most people** (5 by default) for 10 seconds.	Low
Left behind	Safety Checks	*Experimental, In progress.* Something new that stays in the zone, compared with the zone's reference picture. It is listed on this tab but raises nothing yet.	Low
Removed	Safety Checks	*Experimental, In progress.* Something that should stay is gone, compared with the zone's reference picture. It is listed on this tab but raises nothing yet.	Low
Blocked exit or walkway	Safety Checks	*Experimental, In progress.* An exit or lane covered by things that stay, compared with its empty reference picture. It is listed on this tab but raises nothing yet.	Medium
Queue point	Retail & Area Analytics	People in the queue would wait longer than the target, 300 seconds by default.	Low

A check raises an event only after it holds through its confirmation window, never on a single frame. Where a check lists it, **Confirm for** and **Seen in at least** set that window. A pose model needs the PyTorch engine, except RTMPose.

> ⚠ **Important.** Camera-based fire and smoke detection is an early warning that supplements the certified fire alarm system. It is not a certified life-safety function.

Events from these checks appear in Events with a **Check** filter (chapter 11). Fire and person down raise alarms (chapter 12). A lamp, a vent or a sign that looks like fire keeps triggering a check: the camera's **Fixed items** tab teaches the platform to ignore it (chapter 6).

Security checks

Open **Rules → Security checks** for the checks of the Security app. It works like the Safety checks tab, and says **No security checks are available** when the app is off or not in the licence. A check that cannot run yet says why under its name, for example "Needs the People & Visitors app. Switch it on in Apps."

Check	What it finds	Starts as
Restricted area	Anyone in the zone while the check is active.	Medium
Loitering	A person staying in the zone longer than 300 seconds.	Low
Unknown face	A face seen clearly that matches nobody you enrolled. One alert per person; it closes as recognised if they are recognised later.	Medium
Watchlist match	A face that looks like someone on the watchlist. Always a possible match, never proof of who it is.	High, alarm
Authorised people only	Anyone in the zone who is not one of the **Allowed groups** or **Allowed people**. Unknown faces count as not allowed.	High
Camera blocked or moved	The picture is covered or very dark for longer than **Hold for** (10 seconds), or blurred or pointing somewhere else for longer than **Moved or blurred for** (1 minute). Not for cameras that pan or patrol.	High

The three face checks need the People & Visitors app and **Recognise named people by face** (chapter 8). **Most security alerts per camera per hour** (**Settings → Detection**, 60 by default) stops a camera flooding the bell: over the limit, the camera keeps its further alerts that hour as sightings and shows one notice. Alarms still ring, and 0 means no limit.

A face alert shows the face that was seen beside the enrolled photo, how wide the face was against the rule's smallest face, the reason a person is on the watchlist, and the same person's earlier alerts today. Alerts raised

before the width was recorded show no width.

8 PPE and security modes

The platform can check PPE, or it can only track people. You choose per camera, and you can switch PPE checking off for the whole site.

PPE on each camera

Whether a camera checks PPE is its **PPE** chip, set on the Cameras page or in the camera's form (Checks on each camera).

- **PPE on** is the default. The camera checks the PPE each zone requires and raises PPE events.
- **PPE off** makes the camera track people only. It keeps its zones and PPE rules, which apply again when PPE is switched back on, and no PPE events are raised. Its other checks, Security, Safety and Counting, carry on.

The site-wide PPE checking switch

Administrators can switch PPE checking off under **Settings → Detection → PPE checking**. It is on by default. The platform asks you to confirm before it turns the switch off. A camera judges PPE only when the site switch is on and the camera's PPE chip is on.

What happens to a camera that does not judge PPE

- Tracking, the live view and recording carry on. There is no PPE check, no compliance verdict, no frame mining and no PPE event.
- An event that was open when the PPE chip or the switch changed is closed. It stays on record.
- When no enabled camera judges PPE, the platform unloads the PPE model and frees its GPU memory. It loads the model again when a camera needs it.
- A change reaches the vision worker within 15 seconds.
- Nothing is deleted. Zones, rules and past events stay, and switching PPE checking back on restores every camera's own PPE chip.
- The Test bench still works while the site switch is on, even if no camera has its PPE chip on. With the switch off it answers "PPE checking is switched off". Model downloads and training are not affected.

Face recognition

Settings → Detection also holds **Recognise named people by face**. When it is on, the platform recognises the people you have named, on any camera and on later days, whenever their face is visible. Masks and respirators hide it. It works only when **Follow people across cameras** is on (**Settings → Re-**

identification). The first start downloads two small face models, about 38 MB. A site without internet then matches people by appearance only.

Keep unknown faces for sets how many days the platform keeps the face templates and thumbnails of people who have no name, and 0 keeps none. Under the strictest privacy rules it is 0, so the face of a person with no name is not kept; a profile that allows more sets the limit. The face data of a named person is kept until no camera has seen them for a year: **Settings → Privacy → Delete face data of people unseen for** (Advanced) sets the days, 365 by default, and then the face photos, templates and consents are deleted and the name stays. Two weeks before, the System page warns how many people will be deleted and on which day, and a person on legal hold is skipped.

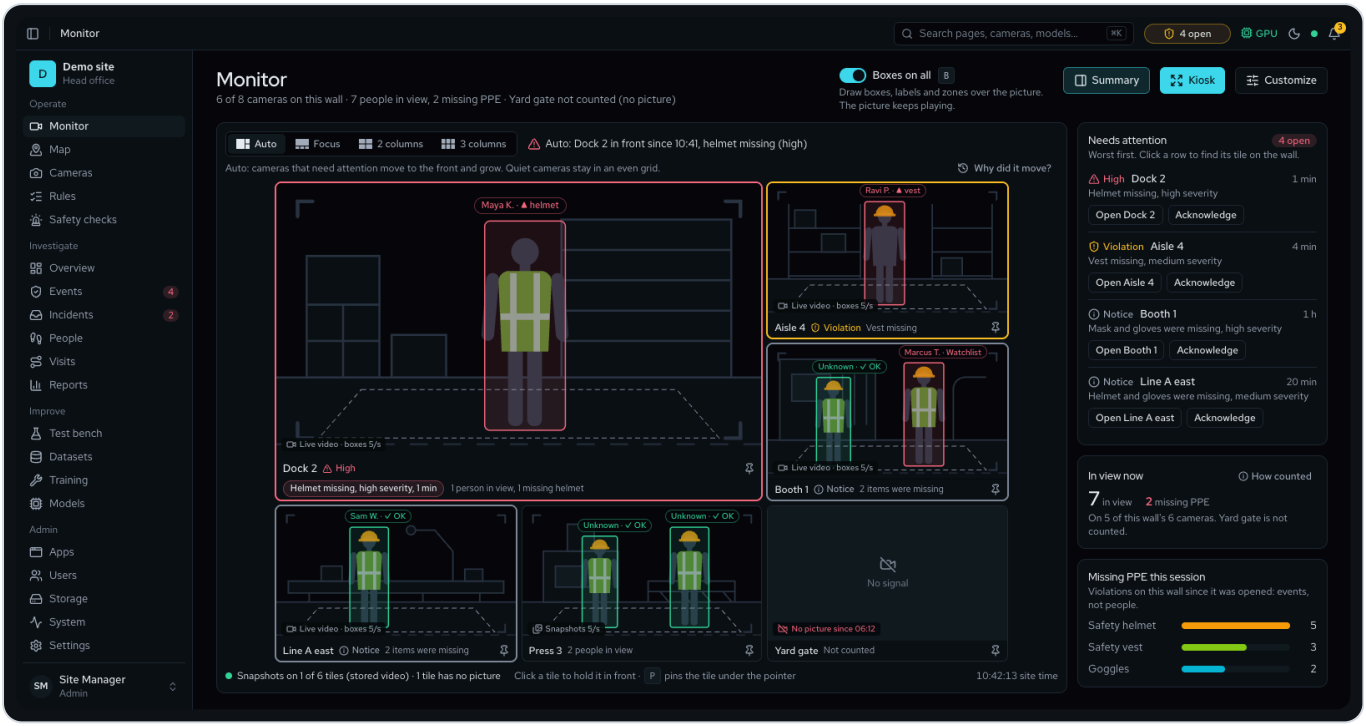
Settings → Privacy (administrators) holds the site's **privacy profile**. Until an administrator picks the site's country, the strictest rules apply. Picking a country applies its profile at once; the dialog first lists every value that goes down, and values never go up by themselves. Each setting with a limit shows it beside the value, and a value above the advised one asks for a reason, which goes into the site's privacy paperwork and never locks anything. Cameras record no sound unless an administrator switches on **Record sound** for one camera, and a profile can keep that switch locked off. Fill the privacy contact and the sign-in page names it.

> ⚠ **Important.** Faces are biometric data. Use face recognition only where you have a lawful basis, tell the people on site with signage, and keep the retention short.

9 Watching live

Open **Monitor** to watch your cameras as a wall of tiles. Supervisors and reviewers land here after they sign in. A line under the title gives the cameras on screen, the people in view and the compliance right now. The

summary's compliance card shows **Time in compliance, last hour** with its coverage and best and worst case, compared with the hour before (how the numbers are counted).



Monitor Available The Monitor: cameras that need attention move to the front with the reason shown, one switch draws or hides the boxes over live video, and every tile says how fresh its picture is.

The toolbar

Control	What it does
Cameras	Chooses up to 16 cameras; until you do, the first four show. **Include test runs** adds the Test bench's.
Area	Shows only the cameras of one area. The choice is in the page address. An alarm in another area can still come to the wall.
Auto, **Focus**, **2 columns**, **3 columns**	The layout. **Auto** brings the camera that needs a person to the front with its reason, and lets it go 15 s after the event ends; while you use the wall it waits and offers **Show now**. **Focus** shows one large tile; arrows or 1 to 9 change it. The other two fix the columns. A pinned camera keeps its place, and pins follow your account to every machine you sign in to. In Auto, the **Pins** menu lets an administrator choose **Pin for every screen** or **Clear pins for every screen**; **Use the pins for every screen** drops your own pins.
Summary	Shows or hides the summary beside the wall. Its first card, **Needs attention**, lists what is open, worst first; reviewers and above can **Acknowledge** a row (with Undo), and a click on a row rings its tile.

Control	What it does
Customize view	Changes how the wall draws people (below).
Wall screen	Opens the wall in its own window for a second display: the whole site, one area, or **Rotate areas...**, which shows the areas you pick in turn, every 30 seconds, 1, 2 or 5 minutes. Rotation needs two or more areas.
Kiosk	Fills the screen with the wall, for a control-room display. `Esc` or the small exit button leaves it.

Your camera choice and layout are saved in your browser and in the page address, so you can bookmark a wall or send its link. For a second display, open it with `?display=wall`.

Live video and Boxes

Each tile plays the camera's own video at its own speed, with the platform's boxes, labels and zones drawn over it.

Boxes

Boxes follow every picture the video shows and glide between checks (5 a second by default). When the detector misses a person or vehicle for a moment, the box stays on its predicted spot for up to a second; a held box is only drawn, never counted or judged. **Boxes** shows or hides them at once, without reloading; **Boxes on all** above the wall, or `B`, does every tile.

Streams

Small tiles and big walls play each camera's sub stream, and Focus and large tiles play its main stream, so a wall of many cameras does not open the main stream for every tile. Tiles change without a black flash. Stored videos and webcams show pictures.

Outside the local network

A viewer whose browser cannot reach the camera's video directly, such as someone outside the local network, gets the video through the server instead, about a second behind, and only then falls back to pictures. On an iPhone this needs iOS 17.1 or later. A tile whose video cannot play says why and retries.

Reading a tile

A badge on every tile says what you see, such as "Live video · boxes 5/s"; point at it for the stream and picture size, or for "Video through the server". The line under the wall sums it up. Point at a person's box and their row lights in the summary and the track list.

Boxes on recorded video

Playback draws the same layer over a recording. Above the tiles, **Boxes** shows or hides it, **Labels** names each box and says its verdict, and one button per PPE type shows or hides that type. A chip says how often the boxes update, such as "Boxes update 5 per second". The recording plays at its own picture speed, and the playhead moves at the speed you choose, from 0.25× to 8×. A camera whose video was not checked has no boxes, and the chip says so. The system saves what it found beside each recording, about 2 MB an hour for each camera, and removes it with the recording, whether the automatic clean-up or a clean-up you start by

hand in Storage deletes the video. A camera with a vehicle line also shows each vehicle's box and type. Names and person ids stay hidden from anyone without access to people. An export shows only what the system found at that moment, never the smoothed boxes shown while playing.

Reading a tile

A thin strip on the tile's left edge shows the camera's state: Online, Starting, Offline or Disabled. The look of the labels depends on the view you choose. This table describes the default look, called Reference. The colour of an item box is the colour of its PPE type, listed in chapter 7.

On the picture	What it means
`Person 91%`	Every label carries an icon, the item's name in its colour and a percentage. A person shows Known person or Unknown person, or their name on security cameras. A person never shows an item and its absence at once, such as Face mask and No face mask: the more confident verdict stays.
`Car 84%`	A vehicle followed on a camera with a vehicle line or bay, with its type and confidence.
A green ✓ **OK** tag	Every item required in this person's zone is present.
A **2 missing** tag in red or amber	That many required items were not found on the person.
A box such as `Safety helmet 84%`	A detected item, drawn in its PPE type's colour with the detector's confidence.
A dashed box such as `No Safety helmet 42%`	The model detected that the item is absent.
A **Can't see safety helmet** tag	The helmet is out of sight, not missing. Only a bare head starts the violation clock; while the item is hidden, the clock pauses.
A name, **Unknown** or **Possible match:** on a person	On security cameras: the person's name, Unknown when their face matches nobody enrolled, or a possible watchlist match in the alert colour.
A red outline and a **Person down** tab	The Person down check flagged this person.
A box such as `Fire 72%`	A hazard found by a safety check (fire, smoke or a forklift), in the check's colour with its confidence.
A faint, dotted box	A detection ignored as a fixed item. It shows when zones are shown.

On the picture	What it means
A **Phone** chip on a person	The Phone use check saw a phone in their hand.
No tag	The person is not judged yet. They were tracked for less than 2 seconds, are too small, or stand in a zone with no rules.

A person who is missing an item but has not yet reached the violation time is a **Warning**: the tile's tab, the track list and Active alerts say so in amber, and the counts card has a Warning row. The words always name the item, so a warning is never told from a violation by colour alone.

By default the Summary panel shows severity as a pill: ALERT for High, WARN for Medium, NOTICE for Low and `OK` for a compliant person. A chip on each tile counts its people: 1 missing, or 2 OK. Where cameras overlap, the people count reads as a range, such as "3 up to 5": the camera that sees the most, then every view added together. Where the overlapping cameras are all placed on the floor map and can be joined, a person seen by two of them counts once, and the range closes to that one number.

Customize view

On the Monitor, Customize opens beside the wall on the left (from the bottom on a phone), so each change shows at once; on Settings and the Test bench it opens as a window. A few parts always stay on, such as Needs attention and the line saying how a figure is counted, and each says why. It has seven tabs: **Presets**, **Boxes**, **Labels**, **Colours**, **Summary column**, **Auto layout** and **PPE wording**. Every option has a plain name and one line saying what it changes. Your look follows your account to any computer. The presets are the quickest way in:

- **Reference**: full detections with a compliance tag on each person.
- **Control room**: amber violations on corner boxes, with no icons. Calmer for a wall display.
- **High contrast**: thick boxes and large labels for a screen viewed from across the room.
- **Minimal**: thin corner boxes, no OK labels, no clock and no Summary. Only the violations.
- **Detections**: every detection in its own colour with its confidence.

The Boxes tab also sets the overlay style: **Detections**, **Compliance** or **Both**. **Skeleton (pose models)** is off by default; switched on, it draws the body outline on each person when the camera runs a pose model. A change shows at once and is kept with your account, so it survives a reload and another browser. **Reset to default** puts back the site default, with Undo for 7 seconds. **Save as site default** is for administrators only.

The **Auto layout** tab has a separate switch for each thing Auto may do by itself: move tiles, bring in a camera from outside the wall, and fill a wall screen for a high event, for a number of seconds you choose. It also sets which events move a tile and how long a camera stays in front; those rows need **Move tiles** switched on. Beside the wall, the panel also sizes **Needs attention** Standard or Narrow (it always stays on, and says why), sets **Row spacing**, and has three switches for the lines around the wall: the layout caption, **Why did it move?** and the tips.

> **Tip.** If a fully equipped worker shows as missing on a wide, crowded view, that is the live profile reaching its limit. Check the same footage in the Test bench in Accurate or Max mode (chapter 15), then see chapter 16.

Playback and export

Open **Playback** (Operate) to look back at what the cameras recorded. It shows up to four cameras at the same moment: one main picture and up to three small ones. Choose a tile to make it the main picture. Double-click the main picture, or choose **Fill**, to fill the grid with it; **Show all four** brings the others back. A tile says so where its camera recorded nothing at that moment. A recording appears about a minute late, because the minute being recorded is not ready until it closes. Anyone who can read cameras can watch, viewers included. On a phone Playback shows one camera: swipe sideways to change camera.

The play bar has **Play** (then **Pause**), a **Speed** list from 0.25× to 8×, where 8× plays **8× key frames only**, buttons that step to the previous and next event, and the playhead's clock on the site's time. `Space` plays and pauses, the left and right arrows move 1 second, `Shift` with an arrow moves 10 seconds, and `B` switches the boxes. **Sync cameras** keeps every picture on the same moment; switched off, only the main picture moves. **Pause at events** stops playing at the next event mark. Customize sets whether the page opens with it on. **Boxes**, **Labels** and a button for each PPE type draw the system's boxes over the recording (Boxes on recorded video).

The timeline shows the last half hour, with one lane for each camera. A solid bar is time that was recorded and a hatched bar is time that was not. Marks show motion and events: point at an event for its picture, or click the timeline to move the playhead. The **Day** strip counts today's events for each hour and boxes the 30 minutes in view; choose another day to move there. Type a time such as 10:38 in **Jump to time** to move the playhead. The **Legend** explains the marks.

Customize on Playback shows or hides **Layer controls**, the **Day strip**, **Jump to time** and the **Legend**. It also sets what the page opens with: **Boxes** and **Labels** (On or Off), **Speed** (1× unless you choose another), **Sync cameras** and **Pause at events**. The recorded and not recorded bars, the chip that says how often boxes update and the manifest line in the export panel always stay, because they say what exists as evidence and how exact it is.

An **Open in Playback** link on an event page and in an incident opens Playback at that moment, and a **Look back 5 minutes** link in the Monitor's camera panel opens it five minutes before now. A closed visit has **Play back this visit**. Only roles that may watch recordings see these links. When a camera has saved views, Playback lists them under **Saved views**, each as a link to the live picture.

To export, choose **Clip, MP4** or **Snapshot** in **Export type**, set **From** and **To** (or drag the handles on the timeline), tick the cameras and choose **Export clip** or **Export snapshot**. A camera with nothing recorded in the range is unticked and says **Nothing recorded in this range**. With Boxes on, the export holds the boxes the system found at that moment, not the smoothed boxes shown while playing. If people may be in view, a **Reason for this export** of at least 4 characters is asked and logged. No reason is asked when nobody is in the pictures. On a phone the export panel opens as a sheet from the bottom.

The zip holds one video (or picture) per camera, and a `manifest.json` with the time range, who exported, the reason and a SHA-256 hash of every file. With Boxes on it also holds each camera's boxes. **Verify** hashes the files again and says how many still match. Choose **Download** within 14 days: after that the files are deleted and the export says "This export's files were deleted." Exports use the **Exports** category on the Storage page. Only you or an administrator can see your exports. Exporting needs the export permission, which viewers do not have.

To add a microphone on the Sound page, pick a sound card or a camera whose video carries sound. A camera's own microphone needs no address or password and also sets the camera that sees where it hears. Its sound is read in memory for alarms only; recordings and live views still carry no sound unless Record sound is on.

Cameras keep pictures only unless an administrator switches on **Record sound** for one camera (chapter 6). That camera's sound is then kept in its recordings, evidence clips and exports, and live view and Playback play it muted until someone unmutes it with the button on its tile. Every other camera's recordings, clips, exports and live view carry pictures only, even from a camera with a microphone. Recordings made before sound was dropped keep their sound until they age out, at most 30 days later.

On the Monitor, in Focus, a camera that turns shows a **Live camera control** pad: hold a direction button to move (a tap moves it a moment), let go to stop, use the zoom buttons, and choose **Home position** or **Go to this view** under **Saved views**. To save a view, type a name of up to 32 characters; the view's **Remove** button offers Undo. Moving a camera needs the move-cameras permission, which viewers do not have; their pad is disabled and says "Your role cannot move a camera. Ask an administrator." Saving a view needs permission to edit cameras.

Share safely

To share a clip outside the site, choose **Cover faces and screens** on Playback. Share safely finds the faces, screens and people for one camera and a range; you watch the whole clip, pick what to show, say why you are sharing it, and export. Every face, screen and bystander you do not show is covered with a solid fill, and the export holds only that covered, silent copy, signed. The original recording stays on site and nothing is recognised or matched.

Personal text. The text that travels with a clip can name people too. Names the site already knows (users and named people), emails, international phone numbers and the local ID card and phone formats of the site's country are found in the reason and notes you write and replaced by a mark: `[NAME]`, `[PHONE]`, `[EMAIL]` or `[ID]`. The export panel has a **Personal text** group that lists each find with **Keep** or **Mask**, so you can keep one as written when it is meant to be seen; a role that cannot see people cannot keep any. The same applies to the notes and comments in an incident's record. Downloads from Reports (PDF and CSV) have a **Hide personal details** switch for their free text, and after a download that replaced something a line says how many personal details were replaced by a mark. The audit log keeps the reason as written, because it stays on site. Names are matched as whole words of at least three letters, so a name can be missed in a different spelling; read the text before you send it.

Every export is signed with this site's own key unless signing is off, so whoever receives it can check that nothing was changed. **Settings → Sharing** holds **Sign exports** (on by default), this site's key with its

fingerprint to give recipients, and **Make a new site key** (administrators). In the export panel, **Sign this export** can leave one export unsigned: a ready export says "Signed with this site's key" and the fingerprint, or "Not signed. The file carries none of this." A signed zip also holds a list of check values, the signature, the public key and a `HOW-TO-CHECK.txt` with the steps and commands to check it offline with OpenSSL, so a recipient needs no account. An export made before signing was added stays as it was.

Open **Share safely** (Operate; it needs the right to export recordings). Under **Check a file** you can test a file someone sends back. Choose **Choose files**, or drop one file, several files or a whole export zip: nothing is uploaded, only the check values are sent. The verdict reads **Unchanged**, **Changed** (naming the file that is not the one exported), **Signature does not match**, **Signed by another site** or **Not known here**, and ends: "This shows the file is unchanged, not that the scene was real." **Recent exports** lists this site's exports: their files are deleted after 14 days, but their record stays, so a copy can still be checked.

10 How a violation is decided

Nothing fires instantly, on purpose. For each tracked person and each required item, the platform keeps a smoothed probability that the item is present. It opens a violation only when the evidence stays bad for a few seconds.

1. The person has been tracked for at least 2 seconds, has been inside the zone for longer than its grace period, and is taller than 80 pixels.
2. The item's probability stays below 0.35 for the violation time (8 seconds by default) without a break. The violation opens, and so does an event. When the model can tell a bare head from a hidden one, only a bare head counts: a helmet out of sight pauses the clock, and the live view says **Can't see** and the item.
3. The violation closes only after the probability stays above 0.6 for 2 seconds. The gap between 0.35 and 0.6 stops a borderline item from opening and closing events again and again.
4. The same person cannot open another event for 60 seconds after one closes.

Setting	Meaning	Default
`minTrackAgeS`	Seconds a person must be tracked before rules apply.	2
Zone **Grace**	Seconds a person must be inside a zone before they are judged.	2
`minPersonHeightPx`	Smallest height, in pixels, to check helmet, vest, goggles or mask. Gloves and boots need 140; under 40 a person is not counted.	80
`absentBelow`	Weakest sighting that counts as worn. A weaker one is too weak to tell and counts neither way.	0.35

Setting	Meaning	Default
`violationDwellS`	Seconds over which an item must be missing, in the share below, before a violation opens.	8
`missingShare`	Share of pictures in that time that must show the item missing. Pictures where it could not be told do not count.	0.8
`presentAbove`	Score above which an item counts as clearly worn.	0.6
`wornHoldS`	Seconds a clearly worn item stays worn while out of sight. An item seen missing still counts at once.	15
`clearDwellS`	Seconds the item must be seen worn again, in at least half the pictures, before a violation closes.	4
`eventCooldownS`	Seconds before the same person can open another event.	60
`emaAlpha`	Smoothing factor. A higher value reacts faster and is noisier.	0.3

Two more values set how sure the detector must be before it counts a box: `personConf` for people (0.4) and `ppeConf` for PPE items (0.25). Administrators edit all of these under **Settings → Detection → PPE evidence**. The zone grace is set in the zone list (chapter 6).

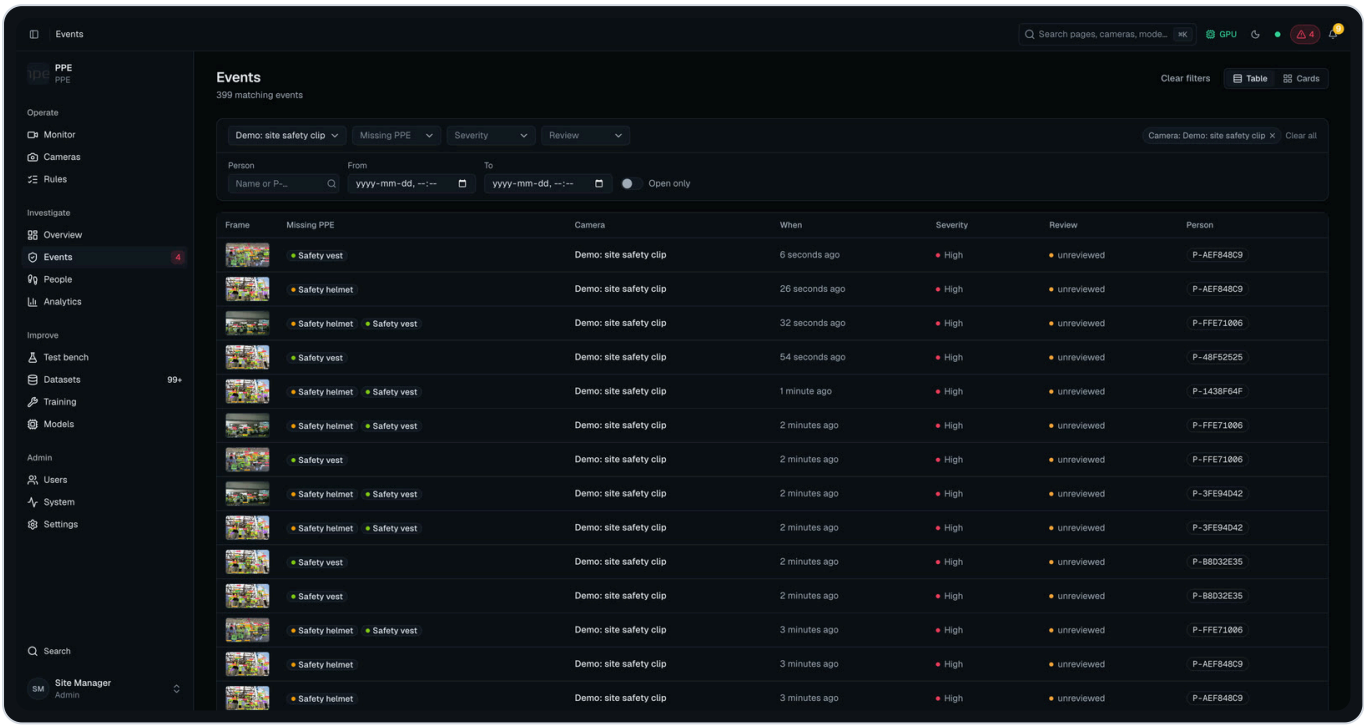
Between a rule's **warn after** and **violate after** times, the person shows as a **Warning** in amber, with the item named in words. Nothing is saved until the violation time is reached.

One event covers one episode: the stretch of time in which one person is missing PPE. If more items go missing during the episode, they are added to the same event, and an event that is already open adds no second violation. A snapshot taken in the warning stage says **Warning**. An event records snapshots when it opens, and a clip of the episode with 5 seconds before and after (`clipPreS` and `clipPostS` under Settings → Retention). The clip is ready about 65 seconds after the event ends, once the 60-second recording segment that holds the end has closed. If recording was off, the event has snapshots and no clip.

A person hidden for a moment keeps their open event, so a pillar or a passing truck does not create a second one. When a camera loses a person briefly and finds them again, it joins the two pieces of their path into one person with one violation. Pieces it is not sure about stay apart and count as possible repeats; the joining scores, such as **Score to join two pieces**, are under **Settings → Counting & coverage**. When two people pass close enough to be mixed up, both are checked again for 2 seconds. An event ends when its violation was last checked, not when the person leaves view.

11 Events and review

Every confirmed violation becomes an event. Open **Events** to see them all.



The Events list with its filters. Each row shows what was missing, where, when, how serious, and who.

The list

Switch between **Table** and **Cards**. The columns are Frame, **What**, Camera, When, Severity, Alarm, Review and Person. **What** names each event by its check and subject, in the check's colour; a PPE event still lists the missing items.

Filters

The filters are **Camera**, **Check** (starting at **Any check**), **Missing PPE**, **Severity**, **Review**, **Person** (a name or an id), **From**, **To**, **Open only**, **Not acknowledged** and **Held only** (events under a legal hold).

Export

Export CSV saves the list as it is filtered, the Check filter included.

Times

Every time on Events, and the **From** and **To** filters, follow the site's clock. An open event of a camera silent for 3 minutes is closed at its last sign of life and left out of duration figures.

Who sees the Person column

The Person column and filter are for admins, supervisors and reviewers; a viewer's list has neither.

Quick filters

Review set to unreviewed lists what still needs a decision. **Open only** lists events still in progress. **Clear filters** resets them.

Bookmark and page size

The filters and the chosen event are part of the page address, so you can bookmark or share them. The list shows 25, 50 or 100 events per page.

Open and closed

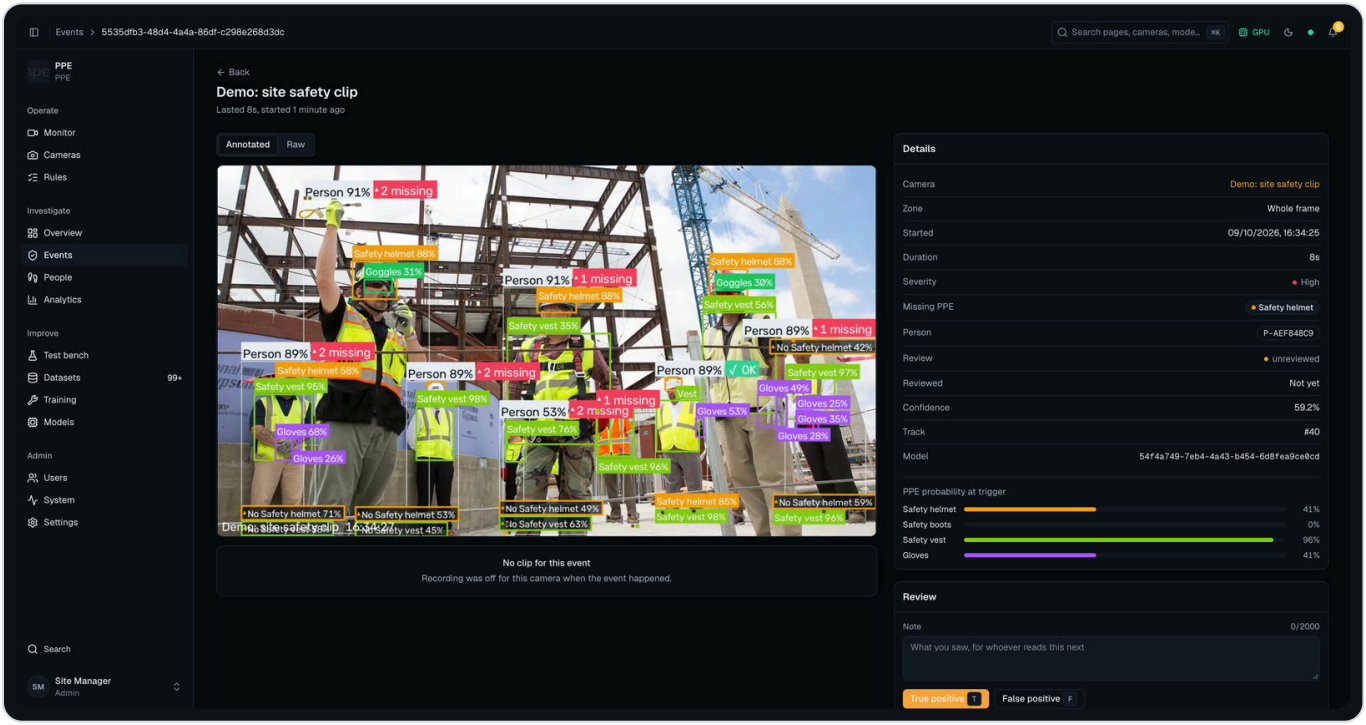
An event is open while the person is still missing PPE, and closed once they are compliant again or have left the picture.

One event

Choose an event to open it, beside the list on a wide window or on its own page on a narrow one. Above the picture, two lines say what the event is and where: for example, "Something new in Whole frame differs from its reference picture" and "Outlined: where the picture differs from the reference." Toggle **Annotated** and **Raw** to see the snapshot with or without the platform's boxes. Pictures of old events show only the items the camera's rules asked for. A clip plays in the same frame, with its marks moving with the video, or the page says why there is none: not cut yet, recording off, extraction failed, or unplayable. A clip made before this change has no saved marks, and the page says so. **Open in Playback** opens the recording at the event. If the disk was full and the snapshot could not be saved, the picture says **Snapshot not saved (disk full)**.

Open full screen (on the snapshot or on the clip bar) opens the snapshot or the clip over the whole screen, with the same marks, and zoom, pan and pinch on a phone. On a desktop, `Esc` closes it, `←` and `→` switch between snapshot and clip, `Space` plays the clip, `B` switches Annotated and Raw, and `+`, `−` and `0` zoom. On a phone, swipe left or right to switch between snapshot and clip. Closing it returns you to the button you opened it from.

When the camera kept no recording of the event's start, the clip is not there. The page then says: "No recording of that time is kept." It does not claim that recording was off. A supervisor or administrator also sees a link to the camera's recording setting.



An event page: the snapshot, the details, the probability of each item when the event opened, and the review buttons.

When an area has overlapping views and two of its cameras saw the same person miss the same items, the event page has an **Also seen by** section with the other camera's picture. **Delete event** removes those pictures too.

The Details panel lists the camera, zone, start time, duration, severity, missing items, person, review state, absence score, track number and model. **Absence score** is how clearly the missing items were absent. It is a score, not the chance that the event is right. **PPE probability at trigger** shows each item's probability when the event opened.

Alarms

An event from a check that raises an alarm, such as fire or person down, has an **Alarm** column in the list and an **Alarm** card on its page. Choose **Acknowledge** (on the page, **Acknowledge alarm**) to show that someone is dealing with it. It needs permission to review events. An alarm is acknowledged once; afterwards the list and the page show **Acknowledged**, with when and the note. Acknowledging switches the siren off (chapter 12). Evidence cards, the bell and its toasts name the check too.

Reviewing an event

Reviewing needs the Admin, Supervisor or Reviewer role. Read the snapshot, then choose one of two buttons or press its key. The event shown stays put while new ones arrive at the top. You can add a note of up to 2000 characters first.

- **Real violation** (`T`): the person really lacked the item.
- **Not a violation** (`F`): the platform was wrong. The person wore the item, or the box was not a person.

Marking an event Not a violation also puts the raw frame, with the model's boxes, into the Datasets review queue, so the next model learns from the mistake (chapter 16). If an event is wrong because of the zone or the rule, fix the zone or the rule instead of marking it false. Your reviews feed the **Reviewed so far** panel on Reports. It counts the events a person has looked at, not accuracy.

Deleting events

Only an administrator can delete events. **Delete event** is on the side panel of Events (on the event page on a narrow window). It asks you to type `delete`, then removes the snapshots, the clip and its marks for good. It does not run for an event on legal hold, for a person on legal hold, for an event still running, or for an event whose pictures a training sample uses; the page says why.

Clear events, in the Events header, deletes many events at once. It works on the ticked rows, or on everything the filters match. Events on legal hold, still running, confirmed as violations, open incidents or used by a training sample are skipped and counted, and the message gives each count. Events that arrive after the dialog opened are not deleted. Every request is written to the audit log.

Legal hold

Put an event on **legal hold** from its page to keep its evidence. A held event is never deleted automatically, whatever the storage age rules say, and neither Delete event nor Clear events removes it. The **Held only** filter on Events lists the held ones.

Incidents

Incidents works violations as cases. Each incident has one status: **New** (nobody has it), **Acknowledged** and **In progress** (someone owns it), **Resolved** or **False alarm**. The tabs are **All**, **New**, **Being handled** and **Closed**, and the filters narrow the list further. Select several incidents to change their status or **Assign to…** them in one step.

Open an incident to read its **Notes** and **Status history**, add notes, open its recording with **Open in Playback**, and export its incident PDF or an evidence pack as a zip file. Incidents assigned to you also appear under **Assigned to you** in the bell.

Video search

Open **Video search** (Investigate) to find recorded footage by describing it or by showing an example. It needs the Video Search app, switched on and in the licence (Apps), recorded footage, and the AI engine running. If the engine is off, the page says "Search runs on the AI engine, which is not running. Recordings still play in Playback." With no footage indexed it says **No footage is indexed yet**; turn recording on for a camera and it is searchable within minutes. Two badges say how it works: **Runs on this server**, so nothing leaves the site, and **Searches are logged**, so who searched for what is kept. Each person sees only their own searches.

1. **Say what you look for.** Choose **Describe it** and type a few words, such as "a person in a red jacket near the loading bay", or choose **Example image** and pick a picture of up to 25 MiB. A role without access to people has faces covered on the picture, and the picture is deleted after the search.
2. **Pick the cameras and the period.** A camera that cannot be searched is struck through and says why: switched off, not recording, no picture, or not indexed yet. The period is Last hour, Today, Yesterday, Last 7 days or Custom. Only the first 64 cameras are searched. A line under the page says how far each camera is indexed.
3. **Choose how sure a match must be. Show** sets **Strong only** (80% and above), **Strong and good** (70% and above) or **Include possible** (60% and above). Sort by **Best match** or **Newest**. Matches that are close in time fold into one moment.
4. **Read the matches.** Each card is a still from the recording with its camera, the time on the site's clock and a score such as "82% match". Point at a card, or focus it, to play its 30-second clip (10 seconds before and 20 after) with the detection boxes drawn on it; with reduced motion, a **Play 30 s clip** button does this instead. **Boxes** (key `B`) switches the boxes. A card says in plain words when there is no still: not recorded yet, no recording at that moment, deleted, or could not load.

5 **Open a moment. Selected moment** shows the frames 10 seconds before, at and 20 seconds after the match, and a day strip of that camera's other matches. **Open in playback** opens the recording there. **Export 30 s clip** needs the export permission. **Search with this frame** searches for more like it. On a phone it opens as a sheet.

A score for words and a score for a picture are measured differently, so compare scores within one search. Deleting footage also removes it from the index. Searching for vehicles is Later.

12 Notifications

Notifications reach you while the platform is open in your browser. Their settings are saved in the browser, so each browser you use has its own.

- **The bell.** It shows the number of new violations since you last closed the list. Choose it to open the list, then an event to open it. The filters are **All**, **Open** and **High**, and **Mark all read** clears the count. The list holds up to 200 events that arrived while the platform was open; after a lost connection it adds what started meanwhile as unread, without pop-ups. For history, use Events.
- **Pop-ups.** By default a toast pops up for each new High violation and for account and system notices; Medium and Low violations reach only the bell. A camera raises at most one toast every 30 seconds, and further events in that window are grouped ("3 more on Gate cam"). A toast names the person when they have a name.
- **Mute.** Choose the bell-off icon on a notification, then **Mute for 10 minutes**, **Mute for 1 hour** or **Mute for today**. A muted camera raises no toasts, sounds or browser notifications. Its events still reach the list. The list shows "muted until" with an **Unmute** button.
- **Browser notifications.** Optional. The browser asks for permission when you switch them on. The platform posts them only while the tab is hidden.
- **Assigned to you.** The bell lists the incidents assigned to you, each linked to its incident.
- **System alerts.** A camera that goes offline, a stopped vision worker or trainer, lost live updates, and low disk raise an alert. It shows in the **System** group of the bell, above the events, and a camera offline alert appears within about 70 seconds of its last picture. Choose an alert to open the camera or the System page. Administrators acknowledge or reopen alerts under **System > System alerts**, and an alert closes itself when the cause is gone.
- **Alarms.** A check that raises an alarm, such as fire or person down, shows a full-width red banner at the top of every page, above the top bar so the bell stays within reach, and repeats a siren until someone acknowledges it or it clears. Each row reads, for example, "Fire at Dock 2", with the time, an **Open** button and, for people who may review events, an **Acknowledge** button. **Mute the siren** and **Unmute the siren** control the sound on your screen. A browser plays sound only after one click on the page, so the banner offers **Click to enable the alarm sound**. A red chip in the top bar counts the open alarms, such as "2 alarms". An alarm that waited while the server was down says "delayed by N min" and never sounds the siren. Alarms reach every open screen. Settings → Notifications sends them elsewhere too (chapter 17).

- **Sound.** Choose **Pop-up with sound** for a kind of alert and its pop-up plays a tone: an urgent tone for a High violation and a soft one for other alerts, so a High stands out. Nothing plays until you choose it. A browser can block the sound until you have clicked in the page once.
- **Repeat for High violations.** Off until you choose **Every 30 seconds**, **Every minute**, **Every 2 minutes** or **Every 5 minutes**. It needs High violations on **Pop-up with sound** and the right to acknowledge violations. The urgent tone then plays again, one interval after the page loaded and then once per interval however many violations wait, until someone acknowledges the violation, reviews it as a true violation, or it ends. The count on the bell bumps once for each repeat. A tab that stays hidden for minutes may repeat later than you chose, up to about a minute apart, so keep a control-room tab visible. A muted camera and Do not disturb stop it, and it starts again when Do not disturb ends. When a person cannot use it, the row says why.
- **Do not disturb.** In **Notification settings**, the **Do not disturb** row offers 15 minutes and longer, up to the longest the site allows. While it is on, your pop-ups and tones stop, repeats stop and the bell shows a bell-off icon. The bell still counts every alert, and its list says "Quiet until" with the time and an **End now** button. Alarms always show the red banner and sound the siren. A wall screen (`/monitor?display=wall`) keeps its tone and shows no pop-ups, and a display account cannot turn it on: its row says "A wall screen always shows alerts." It follows you on every device, so starting it or ending it in one browser changes the bell in your other browsers within a second. When it ends, a message says how many alerts came in, with an **Open bell** button. A server without it shows no Do not disturb row.

Under **Notification settings** in the account menu or the bell, choose for each kind of alert (**High violations**, **Medium violations**, **Low violations** and **Account and system notices**) whether it is **Bell only**, **Pop-up** or **Pop-up with sound**. **Play** tries the tone. Every violation still reaches the bell. The defaults are Pop-up for High violations and for notices, and Bell only for Medium and Low. **Desktop notices** is a separate switch, and the page lists the **Muted cameras**. Alarms always sound with the siren and a red banner, whatever you choose. Every role has these choices, and they apply to this browser only.

13 People

The platform gives each person it tracks an id such as `P-1A2B3C4D`. The same id follows a person from camera to camera, so you can see a journey and everything that person missed. People, names and face pictures are for admins, supervisors and reviewers. A viewer has no People entry in the sidebar, its address answers **No access**, and the Monitor shows a track number where the person id would be.

The People page

Open **People**. Switch between cards and a table. Search by name or id.

Filters and sorting

Show filters by All, Named or Unnamed, **Seen** by 24h, 7d, 30d or All time, and you can filter by camera and by **Group**.

On site now

On site now lists the people a camera saw within the last few minutes (**Settings → Detection**, "On site now means seen in the last"); a green dot marks them. A person who was enrolled but never seen reads **Not seen yet**.

Faces and readability

A short notice on the page reminds you that faces are biometric data. Sort by Last seen, Most violations or Name. People under 100 px tall are counted but not listed. **Face readability** shows, for each camera over the last 7 days, how often a face was big and clear enough to recognise; a low figure means recognition will rarely work on that camera.

A person's card

A card shows the sightings, the cameras and the violations of one person. Open a person for their activity timeline, the cameras they were seen on, the PPE they missed most and links to their events.

Naming people

Reviewers, supervisors and administrators can name a person: choose **Name this person** on their page or menu, type a name and add optional notes. The name belongs to the id. It shows next to the id on Events, in notifications and in the Summary, and it returns whenever the cameras recognise that id again. With face recognition on (chapter 8), a named person is also recognised by face on later days.

Enrolling people and groups

Administrators and supervisors can add people before a camera has seen them. Choose **Enrol person**, then **Enrol a person from photos** opens: add one to five clear photos, each with one face (JPEG or PNG, 10 MB at most), a name, and optionally a group. The dialog checks the count, size and type before it sends, warns when face recognition is off, and shows the server's answer for each photo, so you can try better ones. A person keeps at most ten enrolment photos. **Add photos** on a named person's menu or page adds more (the title reads "Add photos for" and the name). A failed enrolment leaves nothing behind.

Groups (administrators) opens a dialog to add, rename and remove groups such as staff, contractors and visitors. A person has one group and can be on the **Watchlist**, with an optional reason; set both when you name them. The page of a person shows their group, the watchlist reason and how many enrolment photos they have. People who can see people can see groups.

From a person's menu, **Delete face data...** removes their enrolment photos, face templates and face picture; the name, the sightings and the events stay, and they are no longer recognised by face until you enrol them again. **Delete person...** removes that and also the name, group and watchlist entry; sightings and events stay under the bare id. Both ask first. Someone who is forgotten and then seen live is not learnt again.

Merging two ids

When one person ends up under two ids, open one of them and choose **Same person as...**. Pick the other id, then choose which id to keep. The other id's sightings, events and enrolment photos move to it, and the cameras treat both as one person from then on. The dialog is titled **Merge into one person?**. Naming and merging are recorded in the audit log.

Visits

A visit is one stay of a person in a zone. The **Visits** tab on a person's page lists their visits day by day, and each sighting names its zone. A stay is one sighting, however many frames it spans. When a person has been out of sight in a zone for longer than **Time out of sight before a new visit** (**Settings → Re-**

identification), seeing them there again starts a new visit. The default is 60 seconds. Merging two ids moves their visits to the id you keep. Visits are kept for 90 days by default. A visit whose camera stopped watching near its start or end reads "at least" before its duration, and the page marks its figures as estimates.

The **Visits** page lists every visit, newest first, with a summary of the visits, their average duration and the longest one for the period and filters you chose. Filter by period, area, zone, person and duration, or show only visits missing PPE or only people still in a zone. Choosing a visit opens its side panel. The list exports as CSV (**Export N visits (CSV)**), with a column for the area and a note of the site's time zone. Only roles that can see people can open it.

Limits you should know

- Matching uses appearance: clothing, build and colour. People in identical uniforms can be merged into one id, and a person who changes clothes can look new.
- Face recognition helps only for named people, and only when the face is visible.
- An id is a strong lead, not proof of who someone is.

Administrators can tune matching under **Settings → Re-identification**. If one person gets many ids, raise `reid.maxGapS` or lower `reid.threshold`. If different people merge, raise `reid.threshold` and switch off `crossCameraOverlap`. The defaults are a threshold of 0.72 and a gap of 180 seconds.

14 Overview and Reports

Overview

Overview is the home page for administrators and viewers. Under the title it says how many cameras there are and how many are reporting, with the time and its time zone.

Tiles

The tiles are **Open violations** (events still in progress), **Awaiting review** (events nobody has reviewed yet), the **Active model** and **System** (the vision service's state, its provider and its load, with a link to the System page). Reviewers and above also see **Pending samples**.

Needs attention

Below the tiles, **Needs attention** lists the worst open violations nobody has reviewed first, without Test bench runs, then cameras that stopped sending pictures (one that never sent any reads "No video yet"), events awaiting review, and a vision service or PPE model that is not working. It never shows a tick while a camera is unchecked. Each row opens the page where it is dealt with.

Trend and cameras

Violations charts the last 24 hours by hour, or the last 7 days by day, against the period before. **Cameras** shows each camera with its coverage over the last 24 hours. The six latest violations follow.

Clock changes and coverage

On a day the clocks change, today starts at the real local midnight and the 7-day trend gives a 23-hour or 25-hour day its own bar. Violations and the Cameras strip carry the coverage and rate lines described under How the numbers are counted.

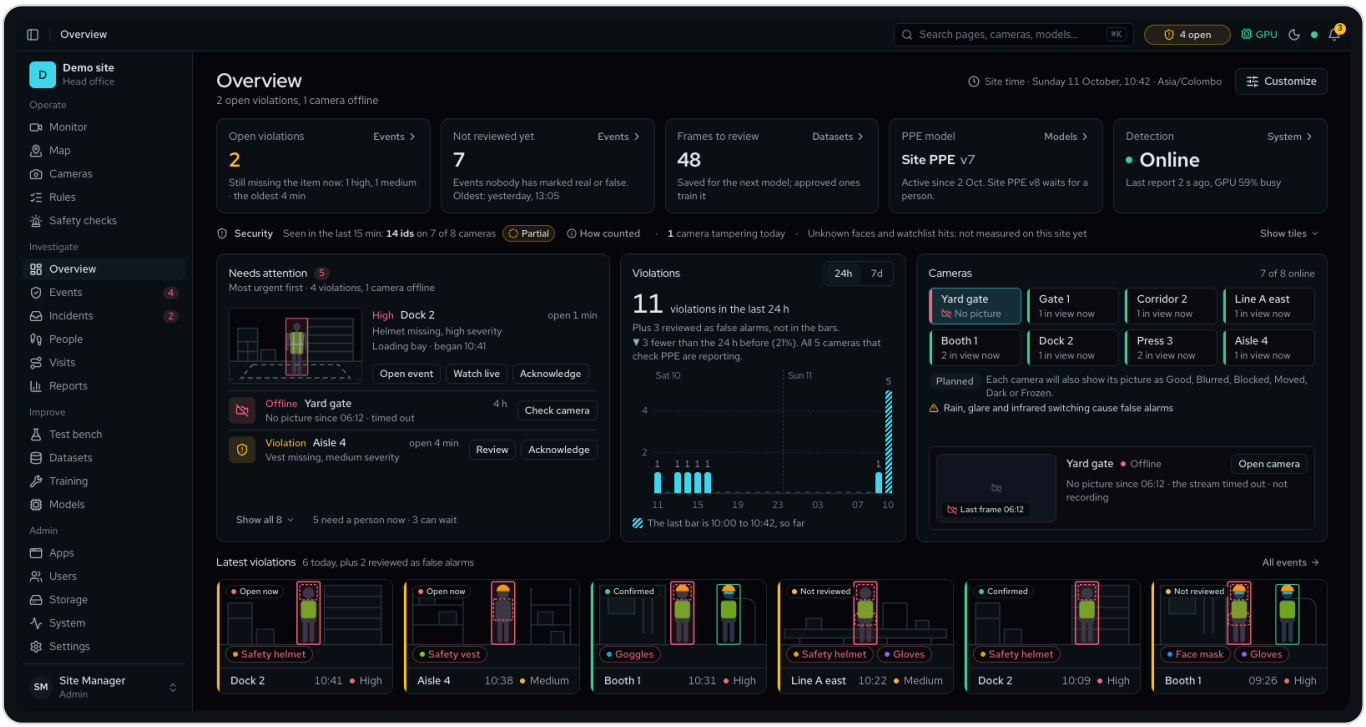
Area select

The **Area** select in the header makes the counts, the trend and the latest violations about one area;

Needs attention and the coverage figures still cover the whole site, and the page says so.

A new site shows **Set up this site** above the tiles: a short list of steps from replacing the install password to adding the people who will use the platform, then requiring two-step sign-in for administrators. Each step ticks itself when the work is really done, and its button opens the page where it is done. To leave a step for later, choose **Skip** and say why (5 to 200 characters), so the next administrator knows; **Undo** brings it back. Someone whose role cannot do a step sees why, in words, beside it. When every step is done or skipped the list says **Set-up is finished.** and hides itself a week later, or at once with **Hide it now**. An administrator can switch the list off under **Settings → Site**, **Set-up list on Overview**.

The first time an administrator opens Overview on a new site, a short welcome asks where the site is. Its time zone sets the clock for reports, rule schedules and every time shown, and the welcome says which privacy profile applies. It shows once for each administrator. The list also asks for the site's privacy country, two-step sign-in for administrators and the nightly backup. Two-step must be on for your own account before it can be required for administrators, and a site that does not want it skips that step with a reason.



Overview Available Camera health, open violations and the day's trend at a glance, with what needs attention first at the top.

How the numbers are counted

Every figure on Overview and Monitor follows the same definitions.

- **Time in compliance** is the share of *judged* person-time that was compliant. Every second a person is present is exactly one of compliant, violating, could not decide (too small, too new, or no PPE model) or out of scope. The seconds before a violation is declared count as violating. The figure always shows a best and a worst case, with all the undecided time one way and then the other; when less than 70% of the time was judged, only those two bounds are shown.

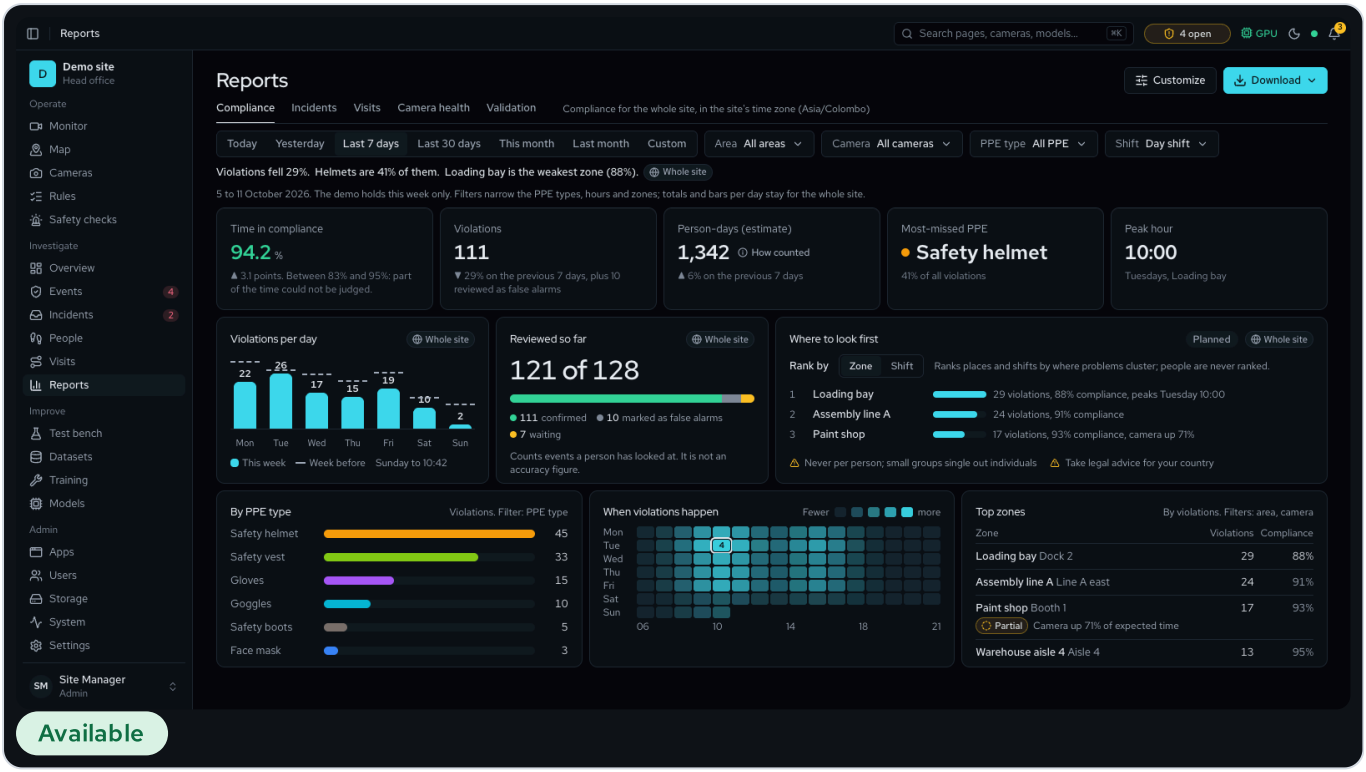
- **Coverage** is the camera time observed over the camera time expected inside the working hours set under **Settings → Counting & coverage** (no windows means all hours). Under 90% a figure says "partial"; under 50% it shows "—" and names the camera. If a camera is online but the server checks it less often than it is set to, the page says so ("The server checked 30% of what this camera is set to. Lower its checks per second or add capacity.").
- **Violations** show "N per 100 judged person-hours" once there are 10 judged person-hours. An arrow compares with the previous period on the cameras observed in both ("compared on 12 of 14 cameras"). Events reviewed as false alarms leave the count and are shown beside it ("plus 4 reviewed as false alarms"). Test-bench runs are in no figure.
- **No data is not zero.** A period with nothing observed shows "—" and says why.
- **Estimates.** "People seen (estimate)", "Compliant so far" and "Non-compliant so far" in the Monitor's session view count confirmed tracks since the stream started, and can count a person again if they leave and return. A person is counted only after being tracked for a short while at a usable size. Every people and visit figure carries **How counted**, which says what one unit is and that a person may count more than once.
- **Silent cameras.** A camera that stops reporting for 3 minutes has its open violations closed at its last heartbeat; such an event has no duration for statistics.

> **Note.** The size limits and tracking times are starting values that have not been measured on your cameras. The weekly accuracy check (chapter 16) collects reviewers' answers on each camera. Camera grades built from those answers are on each camera's page (chapter 6), and a camera whose checks show it cannot be trusted for a figure is left out of it and named under the figures. Each accuracy figure is marked current, recomputing, stale or aged, so one that no longer matches the camera's settings, or whose check set is too old, is not read as fresh. Measured error ranges on every report figure are In progress, so no report figure has a measured error yet.

Reports

Reports (**Investigate → Reports**) replaced the Analytics page; the old Analytics address still opens it. Every role that can read Events can open it. The Compliance tab carries no names, faces or person ids. Reports has up to five tabs. **Compliance** summarises a period for the whole site, **Visits** is the visit log, **Incidents** is the incident log and **Vehicles** reports crossings by type and hour on each line and every wrong-way alert

(chapter 6). **Validation**, for administrators, is the model validation report (chapter 16). The Visits tab lists people's visits, so it needs access to People. A camera health report is In progress.



The finished design of Reports, with sample figures. The page says "Ids seen (estimate)" where this picture says "Person-days (estimate)".

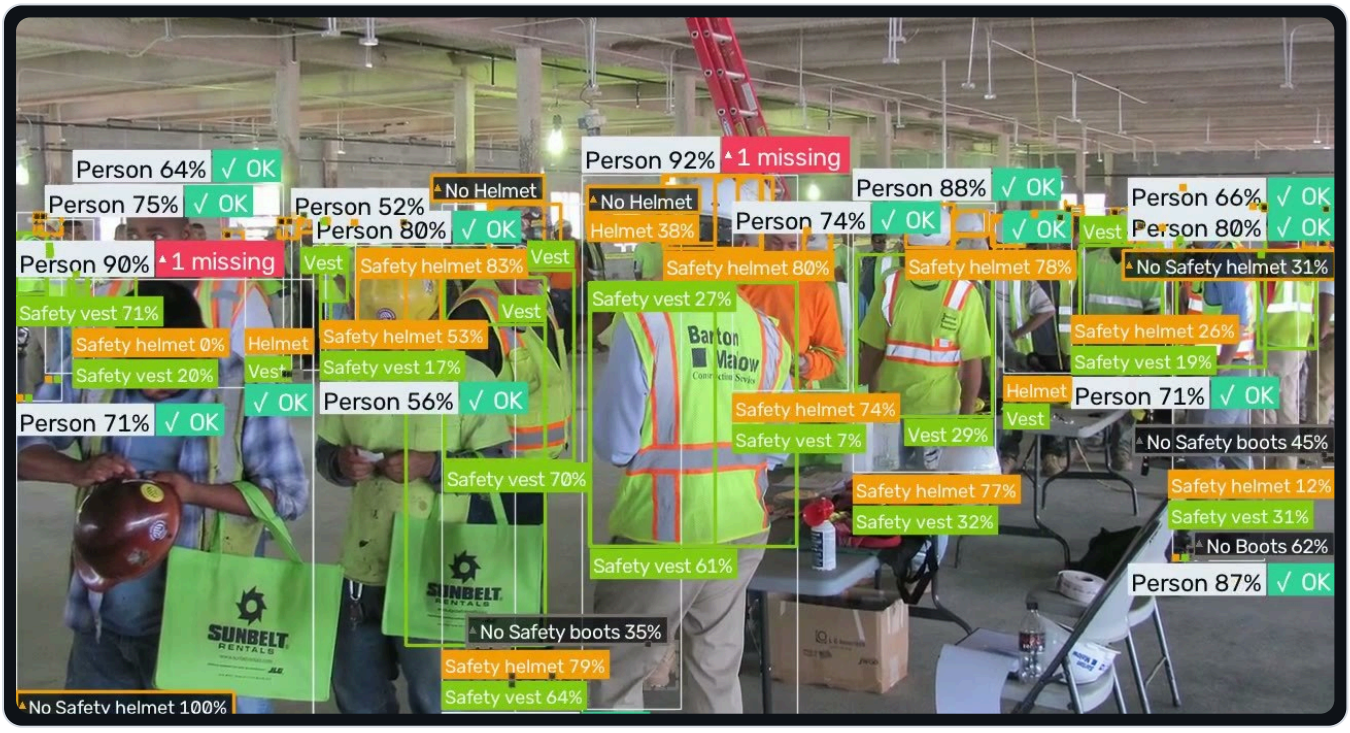
- **Period.** Pick **Today**, **Yesterday**, **Last 7 days**, **Last 30 days**, **This month** or **Last month**, or choose **Custom** and set the two dates (at most 366 days). The site's time zone is shown beside the period ("Times in Asia/Colombo"): every day, hour and time in a report follows it. Set it, and the work shifts, under **Settings → Site**.
- **Filters. Area**, **Camera**, **PPE type** and **Shift** (the Shift filter appears once a shift exists). The camera filter narrows everything. A PPE type or a shift narrows the tables and the hour grid only; the five tiles still cover every PPE type and shift, and the page says so. **Clear all** keeps the period.
- **The tiles. Time in compliance** (with the range it lies in, or only the range when too little time could be judged), **Violations** (with the rate per 100 judged person-hours), **Ids seen (estimate)** (a person who leaves and returns can count again), the **most-missed PPE** and the **peak hour**, each against the previous period of the same length. **How counted** on a tile opens its definition, the same one Overview uses. A sentence under the filters says how much of the expected camera time was observed.
- **The panels. Violations per day** against the days before, **Reviewed so far** (how many events a person has looked at; it is not an accuracy figure), **By PPE type**, **When violations happen** (hour by weekday; press a cell for its count), **Top zones**, **By area** (each area's violations and time in compliance) and **By camera** with how long each camera was up. **By check** is a table of events, alarms and the events confirmed in review for each check, counts rather than an accuracy figure. **Ignored at fixed items** shows how many detections each camera's fixed items ignored. Both follow the filters and appear in the

PDF and CSV too. **Where to look first** ranks the zones, or with **Shift** the shifts, where problems cluster; it never ranks people.

- **Download. PDF report** is a branded A4 file with the licence's logo and who made it and when; **CSV data** opens correctly in Excel, one row per day, camera and PPE type. Both carry the period and filters you have on screen and are named for them. Camera, zone and person names in Sinhala or Tamil print in report and incident PDFs; the screens are in English only. A person can download 30 files a minute, and each download is in the audit log.
- **Incident and visit logs.** The Incidents tab downloads the incident log as CSV, JSON or PDF. The Visits tab downloads the visit log as a landscape PDF or as CSV. The CSV files carry an Area column. Zones that share a name show their camera, so they can be told apart.
- **Customize.** Choose which panels show, in what order, whether the marks for the period before appear, the row spacing (Comfortable or Compact), and the area, PPE type and shift a report opens on. It is kept for you and **Reset to default** undoes it.

Use it to spot a shift or an hour with more misses, the item people forget most, and a camera that raises many false alarms. A camera with many false alarms needs attention: look at its view, its zones and its rules, and feed it training data.

15 Test bench



A stored video checked on the Test bench in Accurate mode. Demo footage built from public-domain photos.

The Test bench checks PPE on an image or a stored video without a live camera. Use it to see how the active PPE model performs on your own footage, and to compare views side by side. Administrators and supervisors can use it. It is under **Improve → Test bench**.



One photo checked in the Test bench. The side panel lists what each person wears or is missing.

A pill at the top names the active PPE model. It warns when no PPE model is active or the vision worker is offline.

Set up a run

The setup card holds one run, **A**. Choose a **Person model** and a **PPE model** for it. Only the test bench uses them; the live cameras keep the active model. Turn on **Compare** to add run **B**, which checks each file a second way, so two models or modes can be compared. **Required PPE** sets the items that count. **Required checks** names the safety checks a video run must exercise, such as Fire or Phone use; it lists only checks whose app is on, and says "No safety checks are available." when there are none. The run card shows them, and a picture you drop shows those checks' boxes, so you can try a check before you switch it on for a camera. **Second opinion** makes another model look again at each person for items the first missed. On a picture its finds are drawn dashed, and the card says how many verdicts it would change. On a video it is a third PPE model: it is shown, never counted, and the finished run says how many people it found extra items on. A looping video uses the live models only, so it takes no third model.

Add files by dropping them on the page, choosing them from your computer, picking a stored video from the library, or pasting an image. On an empty bench, **Load the sample runs** runs three sample photos and one short video through your models, so you can see a result before you add your own files. Every result names the models and the mode that made it, and a required item the model cannot detect is reported as not checked, not as missing. A person who could not be checked for every required item reads as not decided, never as compliant.

An image is one frame. There is no tracking and no time evidence, so a result is one guess. Do not read a single image as proof of how a live camera will behave.

Results

A photo or video shows a summary with the totals, and a card per scene. Its tabs are **People** (an OK tag or the items each person misses), **Detections** and **Info**. A colour key explains the boxes, and **Layers** chooses what is drawn. Under a finished video, **People in view now** follows the playhead, and the run says how long processing took, in GPU seconds on a GPU or CPU seconds otherwise. **Side by side** and **Slider** compare A and B, and the **Differences** panel lists where they disagree. PPE boxes take the colour of their PPE type (chapter 7).

Drop a video (mp4, mkv, mov, avi or webm), or choose **Add from library** for one you uploaded before. Before you start, set **Required PPE**, **Checks per second**, **Loop** and **Record clips**. A video run goes through the full pipeline: tracking, zones, time-based rules and real events. Each run has a timeline, its tracks and a report, and its results download as CSV. A run says how many of its frames were actually checked, such as "Checked 1,180 of 1,200 frames (98 %)", so a run that skipped frames is easy to spot. A run raises real events that you can open and review like any other. Removing a run deletes its events and keeps the uploaded video.

Download zip exports the results as one zip file. A zip that is too large is refused, so export fewer runs at a time. On a phone the setup is a read-only summary, and the results switch between A, B, Slider and Diff.

Replay with boxes

Open a video run to replay the original video with the boxes the run found drawn over it, with **Boxes** and **Labels** switches and a timeline strip with a tick at each violation. After **Reanalyze**, **Before** and **After** replay the run as it was and as it is now. Reanalyze replaces only the saved boxes and keeps the video. If a run's original video is no longer stored, the page says so.

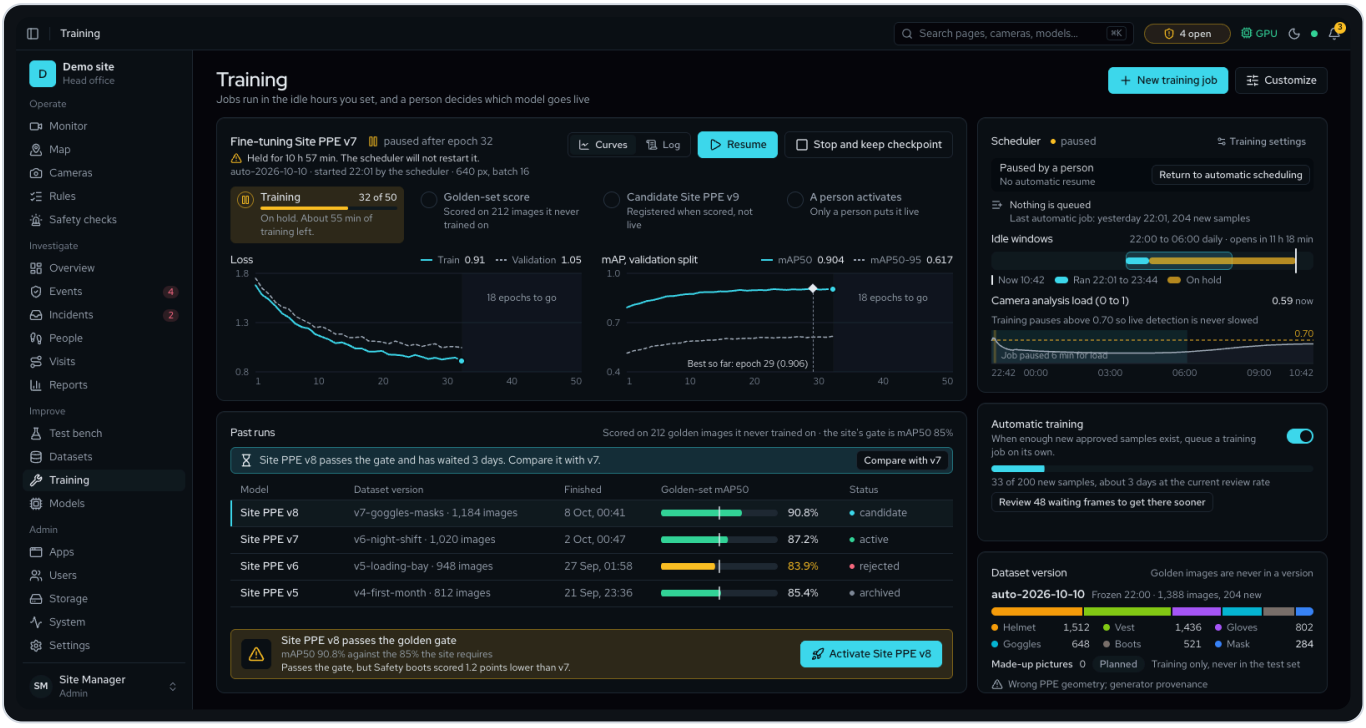
Fast, Accurate or Max

The mode is set per run in the setup card. The GPU or CPU chip in the top bar holds the same three modes. **Fast** is the live cameras' setting and finds the least. **Accurate** doubles the resolution, adds extra passes, runs slower and finds more. **Max** also searches tiles for far-away people, and on a 1080p picture it takes about 5 times as long as Accurate (measured on a CPU server).

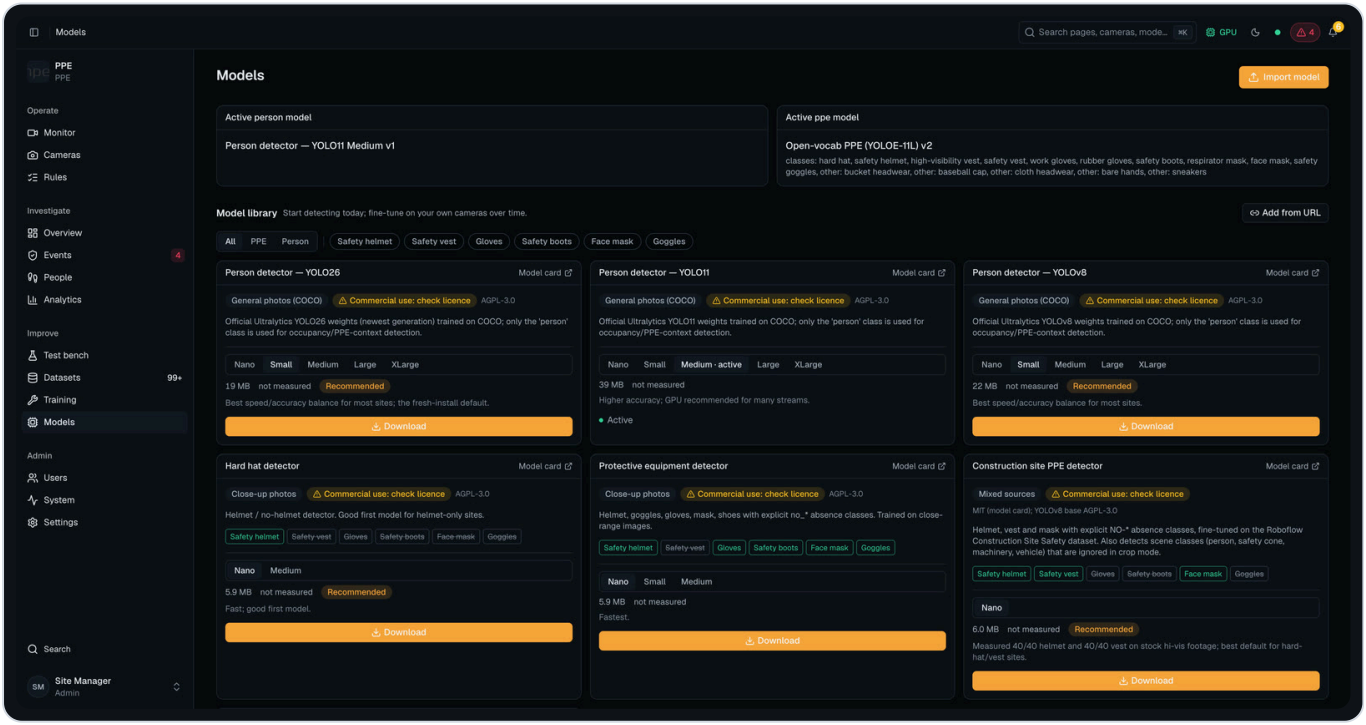
> **Tip.** Accurate and Max show what the model can do on a good day. To see what live cameras will do, run the same footage in Fast, which is their setting. The gap between the two is the accuracy you can recover with training.

16 Improving accuracy

The platform gets better when you teach it what your site looks like. This chapter covers the models and the loop that improves them.



Training Available Training runs in the idle hours you set, with live progress, the golden-set score and activation only by a person.



The Models page: the active models at the top and the ready-made model library below.

Models

Only administrators see **Models**. Two models work together: a **person** model that finds people, and a **PPE** model that finds the items. The panels at the top show which is active.

- **Model library.** A catalog of ready-made detectors. Pick a family and a size (the Recommended one is a good start), choose **Download**, then **Activate** once it shows Installed. The platform reads the model's classes and suggests a class map. Ultralytics models run in every edition. Filters above the library narrow it to All, PPE, Person and, with the Safety Checks app on, **Hazards**, **Behaviours**, **Pose** and **Objects**. A link from the Safety checks tab opens the library for that check. **All models** opens every model that could be added, 664 for 62 jobs: search them, read what each does, and shortlist one per job. The shortlist is kept on the server, so every administrator sees the same one. All models installs nothing itself; **Install from the Model library** opens the Models page at that model.
- **Default person detector.** A new site downloads the pinned D-FINE small person detector, under the Apache-2.0 licence, which runs on ONNX Runtime; a site that already runs a person model keeps it. If the GPU cannot run a downloaded person model, it runs on the processor. If the active person model cannot load at all, the System page says so and the built-in detector finds people.
- **Running now.** A table lists the models the AI engine holds in memory, with the memory, speed and cameras of each (Running now, chapter 6). On Intel hardware, detection can run through OpenVINO, which was faster than the standard engine on the same processor in our test, and is picked automatically once installed.
- **Qualification card.** Each model shows its licence and commercial use, whether its checksum was verified, whether it loads on this PC, its speed and its accuracy. Choose **Run benchmark** or **Re-run benchmark** on the card. Models installed from the library are benchmarked automatically. An operator can switch on a signed online model feed and the discovery of new releases, which show as unverified. Both are off by default.
- **Hazard, behaviour and pose models.** With the Safety Checks app on, the Models page also lists them. Several of each can be active at once, so activating one skips the golden gate and its button reads **Deactivate**; person and PPE keep one active each. The library carries hazard models for fire and smoke, forklifts and construction hazards, official pose models, and an experimental cigarette detector for the Smoking check. Each is pinned to a release and checked by its checksum on install, and the hazard and cigarette ones carry AGPL-3.0 terms. RTMPose, a pose model that runs on every installation and is quick on a processor, lets Person down use its pose method without the PyTorch engine; other pose models need it. A second fire and smoke model, trained on outdoor and aerial scenes, suits yards and open sites and runs on the Ultralytics engine that every release build includes.
- **Model kinds and fit.** Each library card says what its model hands back: **Boxes**, **Outlines**, **Rotated boxes**, **Body points**, **Whole-picture label**, **Outlines from a prompt**, **Depth map** or **Text**. Each size shows a fit chip: **Runs on the CPU**, **Faster on a GPU** or **Needs a GPU**. On a server with no usable GPU, a size that needs one warns you on its card. Objects in the task filter lists the object models.
- **Object models.** Installed object models sit in their own section, several can be active at once, and each has a **Run on live cameras** switch and a **Follow each object** switch. A heavy model starts switched off, with the note "Heavy for this server: switched off until you turn it on"; turn it on when the server can carry

it. A model that takes typed words (YOLOE and YOLO-World) offers **Type what to find**: list up to 40 names, one per line, and the platform makes a candidate copy that looks only for those things, ready to test before you use it. The **Object seen** rule in Rules raises an event for the object you pick, from any object your installed object models can find, and stays idle until you pick what to look for.

- **Bake a model.** This menu at the top of the page offers **Person down (YOLOE-26s)**, **PPE (YOLOE-11L)** and **Hygiene & food safety (YOLOE-11L)**, each while its app is on. The trainer turns text prompts into a candidate model within a few minutes, which you activate yourself. It needs the trainer's open-vocabulary add-on, installed by hand.
- **New releases (unverified).** An operator can switch on discovery of new Ultralytics person detectors. They are listed here but never install in one click: you confirm that they have no published checksum first (**Install as candidate**).
- **Add from URL.** Downloads a YOLO .pt or .onnx file from an address you give. You can add a licence note. A release build installs only models with a published checksum, so it turns Add from URL and Install as candidate off and says why: "This build installs only models with a published checksum."
- **Import model.** Registers a .pt or .onnx file from your computer. Give it a name and a task, person or PPE.
- **Candidates.** Each model has a status: candidate, validated, rejected, active or archived. When it has been scored, its card shows **Golden-set score**, **Score with tight boxes**, **Right when it reports** and **Real items found**, each with what counts as good: a golden-set score above 70% is good, and under 50% misses a lot. **Mark as...** sets a status, **Make ONNX copy** adds an ONNX copy of a .pt model as a new candidate, and **Delete** removes an inactive model's file. Built-in models have cards too.
- **Evaluate and compare. Evaluate on the golden set** scores any person or PPE model on the golden frames, with ranges, misses and a confusion matrix. **Compare with active** shows two models class by class and says plainly when the new one is worse. **Validation report (PDF)** prints the model validation report, which Reports → Validation also gives as PDF or CSV.

> **Note.** Every family card carries a licence badge: the licence name, coloured by whether commercial use is allowed. Most ready-made models carry the AGPL-3.0 licence, and the library marks them "Commercial use: check licence". The badge is for information: every model, from the library or from a URL, installs without a licence question, and Settings has no Models section or Ultralytics switch.

Class map

A PPE model names its classes its own way. The class map says what each class means for the platform. Choose **Class map** on a model, then for each class pick the PPE type it is evidence for. Tick **absence** for classes such as `NO-Hardhat`, which prove that an item is missing and never count as the item being present. Leave a class on ignore when it is not PPE. **Auto-suggest** fills the map from keywords. A saved map reaches

vision within 15 seconds, with no restart. If the active model maps to no PPE type, the System page shows the `ppe_model_unmapped` warning.

Activate, force and roll back

Activating a model is deliberate.

The golden-set gate

The platform scores every candidate on the site's test images, the golden set, and compares its golden-set score with **Minimum score to go live** in Settings → Training. A new site starts at 85%; a site that saved its own value keeps it.

Item-by-item gates

Activate also compares the new model with the live one, item by item. An item that scores more than **Largest drop allowed on one item** (2 points) lower than the live model, or under **Minimum score for each item** (50%), stops it too, and the page names the item.

When a gate is missed

A model that reaches every gate activates after a confirmation. A model that has no score or misses a gate opens **Golden gate not passed**, where you type a **Reason for overriding** of at least 10 characters and choose **Force activate**. The reason is kept in the audit log. Imported and downloaded models have no score until a training job or **Evaluate on the golden set** scores them, so an unscored first activation needs this.

Re-inspect

Re-inspect on an imported ONNX model reads its class names again.

Activation archives the model that was active, and vision reloads within 15 seconds. **Roll back** on the active model restores the previous one. If a new PPE model cannot load, for example because a training job holds the GPU memory, vision raises `ppe_model_load_failed`. PPE cameras then run without PPE checks until it loads, and vision retries every 5 minutes.

The improvement loop

1. **The platform collects frames worth a look.** It queues frames where a PPE item sits in the uncertain band, where person confidence is low, or where the scene looks new. Marking an event Not a violation also queues its frame. The limits sit under Settings → Mining.

2. **You review them.** Open **Datasets → Review queue**. Fix the boxes and classes, then choose **Approve**, **Reject** or **Add to golden**.

3. **You freeze a version.** Open **Datasets → Versions** and choose **Create version**. A version is an unchangeable snapshot of the approved samples. Its training and test images are split by camera and hour, and a picture taken within 10 minutes of a test image or an accuracy check on the same camera is left out of training; the version's notes say how many.

4 **The platform trains.** Start a job under **Training**, or let automatic training do it. A PPE job fine-tunes the active model by default. The **Detector** choice in New training job is **YOLO (Ultralytics)**, AGPL-3.0, or **RF-DETR**, Apache-2.0, so the result may be sold with the product; the licence shows under it. RF-DETR trains from its own pretrained weights, has a **Size** instead of a base model, and is *In progress*: its training add-on is not installed by default and a first full run is still to be confirmed.

5 **You check the result.** The finished job is a candidate with golden-set scores on the Models page.

6 **An administrator activates it,** through the golden gate and the item-by-item check. Roll back if live results get worse.

Datasets

Datasets → Review queue shows counts of pending, approved, rejected and golden samples, a list of pending samples on the left and a canvas in the middle. A sample's source is uncertain, novel, event review or manual, and the list shows its priority.

Draw boxes

Pick the class for new boxes, draw boxes on the canvas, zoom in on small items, and drag or delete the wrong ones. **Reset to model labels** brings back the model's own boxes.

Suggest boxes

Suggest boxes draws dashed boxes from the **Active PPE model**, an **Open-vocabulary model** or the **Built-in labeller**, for you to accept or discard.

Filter and bulk actions

Filter the queue by **Status**, **Camera**, **Class** and **Date**, or select several frames to **Approve**, **Reject** or **Delete** them in one step.

Labels

The **Labels** card imports YOLO labels (.txt), COCO labels (.json) and pictures or a zip, up to 256 MB at a time, and exports this frame, the selected frames or a frozen version as YOLO or COCO.

Upload and approve

Upload image adds your own picture, optionally for a camera. **Approve** and **Add to golden** need at least one box. Otherwise reject the sample.

Key	Action
`A` `R` `G`	Approve, reject or add to golden
`1` to `9`	Set the class of the selected box
`N`	Select the next box
`Del`	Remove the selected box
`Enter` `⌫`	Accept or discard the selected suggestion
`+` `−` `0`	Zoom in, zoom out, fit the picture
`Space`	Next frame, once this one is decided

Key	Action
← →	Previous or next frame
?	Show the list of shortcuts

Golden samples are permanent test data. Every candidate is scored on them, and they are never trained on. Add a spread of cameras, hours and lighting to them: a golden set that is small or stale makes the check weak. **Versions** lists each version with its status (building, ready or failed), samples and classes. Creating a version needs the Supervisor or Admin role. **Train** starts a job from a ready version.

Training

Open **Training** (Admin and Supervisor). Use **New training job** to start one: choose the task, a base model (the default is to fine-tune the active model; it must be a .pt model, and an ONNX or engine model is refused at once with a message asking for a .pt model), a dataset version and, if you wish, epochs, image size, batch size and learning rate. **Export ONNX** also saves an ONNX copy. Leave **lr0** empty and each detector uses its own learning rate (0.01 for YOLO, 0.0001 for RF-DETR). A band at the top says what the trainer is doing: the running job with its epoch, a progress bar and the time left, or a job waiting for the idle hours, or idle. It holds **Pause**, **Resume** and **Stop**. The job panel below shows the job's **Curves** or **Log**. A cancelled job cannot be resumed. **Past runs** lists every PPE model with its golden-set score against the gate, and an administrator can **Activate** one, **Activate anyway** with a written reason, or **Roll back**. The **Scheduler** panel says who holds training and why, with the idle hours and the live camera load of the last six hours. **All jobs** lists every job with its status, scores and log.

- **Idle windows.** Jobs run only in idle windows, by default every day from 22:00 to 06:00 on the site's clock, and training pauses when vision is busy (load above 0.7). A job also waits until **Free memory needed to start training** (2 GB) is free and, on a graphics card, takes at most **Largest share of graphics memory for training** (60%), so live detection is never starved; the job says what it is waiting for. Pausing, resuming and cancelling happen at the end of an epoch. Tick **Ignore idle windows** on a job to run it now; such a job has no Pause. If someone paused training, **Return to automatic scheduling** hands control back.
- **Automatic training.** Switch it on at the bottom of the Training page (administrators) and the platform creates a version and a job when 200 new approved samples (`minNewSamples`) have built up. It is off by default.
- **How much data.** The first fine-tune from a ready-made model rebuilds the detection head. Plan on a few hundred labelled frames per class.

Accuracy checks

Open **Improve → Accuracy checks** (Admin, Supervisor and Reviewer) to measure how often detection is right on each camera. The answers never change an event.

Three figures sit at the top of the page. **This week** shows how many of this week's questions you have answered, with the minutes left. **To settle** counts the questions that two people answered differently and

that wait to be settled. **Sets due** counts the cameras that have no check set yet, or whose check set is out of date. Each figure opens its tab. **See the results** opens the Accuracy report on Reports.

- **This week.** A blind set of short moments from the cameras, shown without what the system decided. Say what you saw, and the next question appears.
- **Disagreements.** The questions that two people answered differently.
- **Camera sets.** Each camera can have its own check set, so it is measured on its own pictures. An administrator or supervisor chooses **Collect a check set**, which gathers it over four working days. A reviewer then chooses **Mark moments**, fixes the boxes and answers for each person.
- **Staged walks.** On **Camera sets**, **Start a walk** measures how often a missing item is caught. Someone walks past the camera without one item at a time, near, at middle distance and far, and the page offers **Start this pass** and **End this pass** for each pass, then **Finish the walk**. Walk without an item only during a stoppage or under supervision: it may break a site rule. It needs permission to change cameras.
- **Short clips for tracking. Collect clips** gathers eight short clips for the site, and **Label a clip** lets a reviewer give each person the same label in every frame. The answers measure how often the platform loses or swaps people, and they tune how pieces of a path are joined.

Weekly accuracy check (on by default) and **Questions per week** (100 by default; every camera gets at least 4) are under **Settings → Counting & coverage**. The Storage page lists these pictures on their own line, **Accuracy check evidence**. Check-set pictures are never offered for clean-up, and weekly check pictures are removed after 90 days.

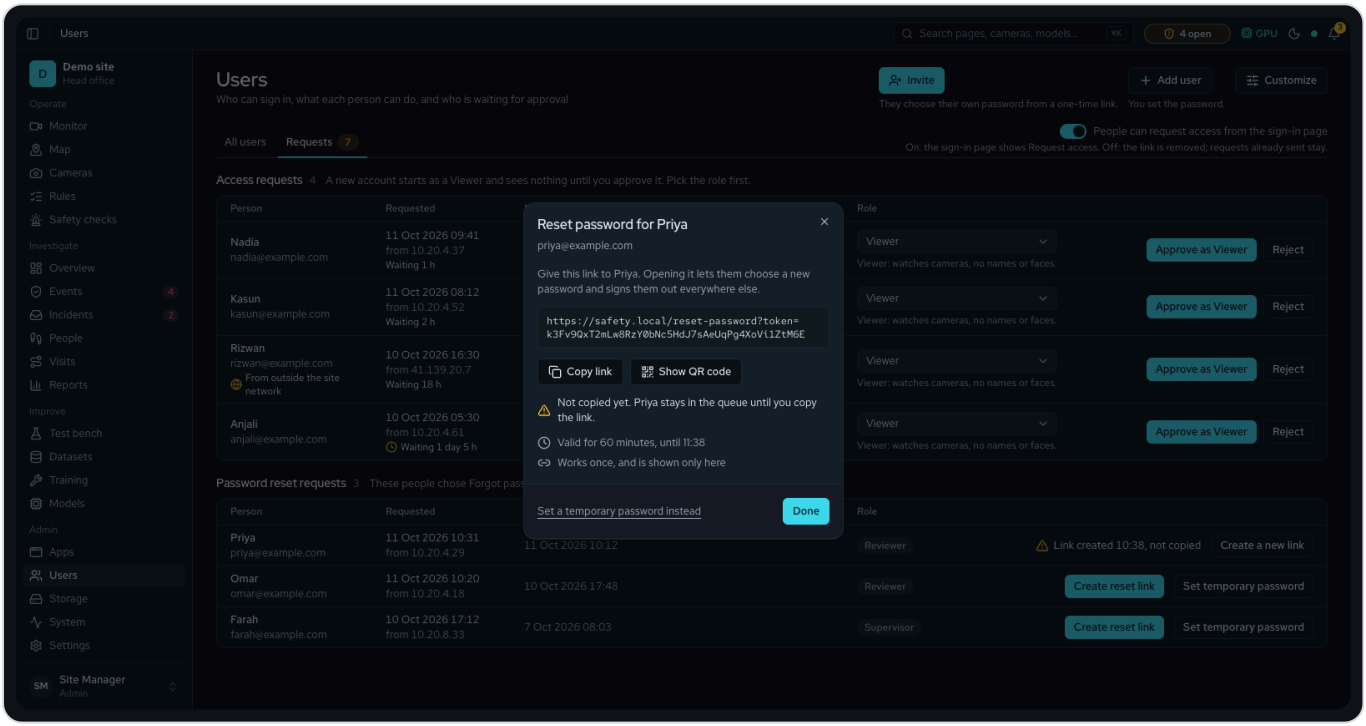
17 Administration

Users

Open **Users** (Admin only). The table shows each person's name, email, role, status, creation date and last sign-in. **Add user** asks for an email, a name, a password of at least 8 characters and a role. **Edit** changes the name, the role and the password (leave it blank to keep it), and its **Disabled** switch signs the person out and blocks them. You cannot change an email. **Delete** ends the person's sessions at once, and the audit entries stay. You cannot delete your own account.

- **End someone's Do not disturb.** A person who has Do not disturb on shows a **Quiet until** label on their row with an **End** button. It returns their pop-ups and tones on every screen at once and is written to the audit log.
- **Approve or reject a request.** A new request rings the bell with a **Review** button, and the Users entry in the sidebar shows how many are waiting. The **Requests** tab lists each one with the name, email, when it was sent and from which network address, and the person's note. Pick a role and choose **Approve**, or choose **Reject** with an optional reason that only the audit log keeps. Refused requests are deleted after 30 days.
- **Invite someone.** Choose **Invite**, enter a name, an email and a role, and choose **Create invitation**. Copy the link and send it yourself. It works once, for 3 days by default, and is shown only once.

- **Follow an invitation.** Until the person has chosen a password, their row reads **Invited, expires** and a date. Choose **Resend link** for a new link (the earlier one stops working), or **Revoke link** in the row's menu to end it; the account stays and the row reads **Link revoked**. Someone waiting for approval reads **Pending approval**.
- **Display accounts.** The **Display account** switch makes a sign-in for a wall screen. Its role is locked to Viewer, the account shows a **Display** badge, and it stays signed in for 30 days at a time while the screen is on.
- **Reset a password.** People who chose Forgot password? also wait on the Requests tab. Choose **Create reset link** (the request leaves the queue once you have copied the link; closing early asks you to keep the dialog open or close anyway) for a one-time link that works for 1 hour by default, or **Set temporary password** to type one or **Generate** one you can read out. Either way the person is signed out everywhere and chooses their own password. **Reset…** on a row in All users does the same for anyone else.



Approvals and reset links Available Administrators approve or decline requests, pick each person's role, send invitations, and hand over a one-time link when someone forgets a password, with no email needed.

Apps

Open **Apps** (Admin group; administrators only) to choose which apps this site uses. Each built app has a card with its switch, what it adds to the menu and what it needs. People who cannot change apps see the same cards with their role named. The apps you can use are CCTV & Recording, PPE Compliance, Security, People & Visitors, Safety Checks, Video Search, Model Training, People Counting, Vehicles & Plates, Conveyor Counting, Retail & Area Analytics, Sound & Voice, Time & Attendance and Visitor Kiosk. **CCTV & Recording** is the base app and is always on. **People Counting** adds counting lines, the Counting page (chapter 6) and the **Occupancy limit** check. **Sound & Voice** adds a **Sound** page (Operate). **Time & Attendance** needs People & Visitors and adds an **Attendance** page (Investigate). **Visitor Kiosk** needs People & Visitors and

PPE Compliance and adds **Expected visitors** (Operate). **Conveyor Counting** adds a **Conveyor** page (Operate), **Belt counting** lines on the camera page and an **Operator panel** section in Settings (chapter 6); a site that only counts a belt can switch every other app off. Two apps do only part of their job so far:

- **Vehicles & Plates** adds a **Vehicles** page (Operate), a **Vehicles** section in Settings and the **Wrong way** and **Stayed too long in a bay** checks. The page counts vehicles crossing a line, in and out, and shows **Parking bays**: a grid of the bays drawn on a camera, each taken, free or reserved. Choose **Set up lines and bays** and the **Parking bays** tab to draw a **Row of bays** (one shape split into equal bays) or a single **Bay**; each bay has a **Number**, can be **Reserved** (**Reserved for** names who) and can have a **Longest stay**, from 5 minutes, after which one low **Stayed too long in a bay** event is raised. Click a bay to set it **Taken** or **Free** by hand; it shows "Set by hand at 10:12" and holds until the camera sees a different state. Setting a bay by hand needs permission to review events. **Reports → Vehicles** adds the share of bays taken by hour, the average stay and the ten longest stays to its PDF. It also reads number plates at a gate line (chapter 6). Vehicle speed is Planned.
- **Retail & Area Analytics** adds the **Queue point** check. Heatmaps and time in each zone are In progress.

Time & Attendance is off until the site sets it up. It clocks people in and out from recognised faces; **Settings → Attendance** holds its options.

- **Switching on and off.** Switching an app off asks first and lists what leaves the menu. A message then offers **Undo** for a few seconds. An app the licence does not include stays locked.
- **What stops and what stays.** A switched-off app stops its live analysis, its new events and its models. Recordings, events, rules, people records and history stay, and switching the app on again picks up where it was.
- **Running now.** A panel at the top shows which AI models the vision service holds in memory. It unloads a model that no app on needs, so switching an app off frees the GPU. A stale heartbeat reads offline.
- **Dependencies.** An app that needs another says so on its card, "Needs People & Visitors, which is off." for example, and one click switches both on.
- **PPE Compliance** is the same switch as Settings → Detection → PPE checking.
- **Overview.** With no AI app on, the Overview hides its violation, review and model tiles and the trend, and Needs attention lists no events. Settings lists only the sections of apps that are on.

Settings

Open **Settings** (Admin only). Every setting has a plain name, one line on what it changes, its default and its range, and a way back to the default once changed. Less used ones sit under **Advanced**, which opens by itself when one of them is changed or wrong. A setting that cannot be saved says what is wrong, in plain words, beside its field. Choose **Save changes** in the bar that appears at the bottom, or **Discard**. Changes reach the cameras' checks within about 15 seconds and training from its next job; nothing needs a restart. If someone else saves while you are editing, your draft stays, and you are asked only when you both changed the same section. Four settings draw a small picture of what they mean: where a PPE item counts as missing,

checking or worn; the free-space limits against today's free space; the hours of the week when training may run; and which camera can follow which.

Section	What it holds
General & branding	The site name, which you edit. The product name, accent colour and logo, which come from the licence and are read-only unless the licence includes Own branding, when an administrator edits them and uses **Reset branding** to go back. A **Your preferences** panel that opens your own theme, row spacing and home page, and the **Defaults for everyone** who has not chosen their own. **Backups**: **Warn when backups stop** (on) and **Warn after** (36 hours).
Sign-in & access	Whether people can request access from the sign-in page (on by default), how many days Keep me signed in lasts (30), and how long reset links (60 minutes) and invitation links (72 hours) stay valid, and the Site network, a list of address ranges: a request for access from outside them is flagged on the Users page (empty, nothing is flagged).
Detection	The PPE checking switch, face recognition, and the PPE evidence values from chapter 10, read back as one sentence. **Calm**, **Balanced** and **Quick** set them for you. **Most security alerts per camera per hour** (60; 0 means no limit).
Safety checks	The box colour of each check, and how scene checks learn still items and alarm when the server is down.
Frames to review	Which frames are queued for review. By default: PPE probability between 0.35 and 0.65 for 2 seconds, person confidence between 0.25 and 0.5, a new-looking scene, one frame per person every 30 seconds, and at most 120 frames per camera per hour.
Retention	Disk thresholds and clip padding (below).
Storage	What is deleted automatically and when, for each kind of data, and what Space saving changes when the disk runs low.
Training	Automatic training, idle windows, the load limit, **Minimum score to go live** (85% on a new site), **Largest drop allowed on one item** (2 points), **Minimum score for each item** (50%) and job defaults (50 epochs, image size 640, batch 16). Under Advanced, **Free memory needed to start training** (2 GB) and **Largest share of graphics memory for training** (60%).
Sharing	**Sign exports** (on), this site's signing key with its fingerprint, and **Make a new site key** (administrators). **Hide personal details in text** (on) replaces names, phone numbers, emails and ID numbers with a mark in text that leaves with an export, and **ID and phone formats to find** picks which local formats count (From the privacy profile by default). See Share safely in chapter 9.

Section	What it holds
Helper servers	Other computers of yours that share the work when this server is full, and the time after which a silent server's cameras move (below).
Counting & coverage	The working hours (every day 24 hours until set), the hour the site day starts, with a suggestion from the quietest hour, the coverage cut-offs (partial under 90%, no figure under 50%), the share of time that must be judged (70%), how long a person must be tracked before counting, and the smallest person each check accepts and counting lines count. The scores for joining pieces of one person's path (Advanced), and **Weekly accuracy check** with **Questions per week** (chapter 16).
Vehicles	With the Vehicles & Plates app on: **Smallest vehicle counted** (60 px wide), **A bay is taken after** (20 s) and **A bay is free after** (60 s).
Sound	How long sound clips are kept and how sure the system must be before a sound raises an alert.
Floor map	Warning bands around no-go zones and how zone suggestions are made.
Attendance	How clock-ins are judged, the photo check, and how long records are kept.
Visitor kiosk	Kiosk screens, their cameras, badges and how long visitor records are kept.
Headquarters	Send this site's totals to a head office. On a head-office server, also the limits that judge every site.
System and load	When a quiet camera counts as offline, when less urgent checks slow on a busy server, and how often Video search keeps pictures.
Site	**Time zone** and **Work shifts**. Rule schedules, runbook days, reports and every time shown follow the site's zone. Each shift has a name, a start and an end; an end before the start runs past midnight, and a new shift starts with its own name. A new server may run in UTC, so set the zone on a new install.
Re-identification	Cross-camera matching: the embedder, the similarity threshold, the gaps and the allowed camera-to-camera moves. **Time out of sight before a new visit** sets when a return to a zone counts as a new visit (chapter 13).
Monitor defaults	The default wall look, and the live stream quality (**How often boxes update** 8, JPEG quality 75, **Boxes follow the picture by** 400 ms).
Notifications	Where alarms go besides the app: the mail server, channels, routes and a delivery log (below). An **On-screen alerts** panel holds **Who may use Do not disturb** (**Everyone**, **Supervisors and administrators** or **Nobody**) and **Longest Do not disturb** (in minutes, in steps of 15). A stricter rule ends the quiet it no longer allows at once.

Section	What it holds
Privacy	Administrators only. The site's **Country**, the privacy profile card, and who is responsible, such as **Organisation responsible** and **Privacy contact**. Until a country is picked, the Strictest profile applies. It keeps recordings at most 30 days, event clips and snapshots at most 90 days, evidence and event records at most 3 years, and sightings and visits at most 90 days, and it keeps no faces of people with no name. A save above a limit shows the limit. The Privacy tab lists every limit, and above the value a profile advises it asks why. **Delete face data of people unseen for** (Advanced; 365 days by default, 30 to 1,095, never more than the profile allows) sets when a named person's face data is deleted.
Updates	Install a new version from a file, see what is new, and roll back for 7 days (below, Backups and upgrades).
Developer tools	Destructive test tools. Keep `DEV_TOOLS` at its default, `off`, on a production site.

Setup backup

An administrator who can change settings and manage users sees **Settings → Setup backup**. **Download setup backup** saves everything this site is set up with in one file, locked with a file password you choose and typed again to confirm. The file holds the settings, the accounts with their password hashes, the cameras with their zones, rules, lines, bays and floor maps, the models and their weight files, the notification and mail set-up, and the enrolment photos of named people. Camera, mail and two-step secrets are kept too, so a restore on a new server opens them; where no secret is kept, you type camera and mail passwords again, and two-step sign-in is set up again. Events, recordings, counts, the audit trail and the licence are not in it. Keep the file like a password.

Restore from a file puts a file back. Choose the file, type its file password and choose **Check file**. The page shows what the file holds beside what is here, and warns about a camera here that the file lacks (it stays, with its events), about your own account when it is not in the file (you are signed out after the restore), and about unsaved changes on other Settings sections. Type your account password and choose **Restore**. Before anything changes, the current set-up is saved as an undo file: **Download undo file** brings it back, and it opens with the account password used for that restore. A restore replaces the set-up but never deletes history: events, recordings, counts and the audit trail stay.

Notification channels and routes

Open **Settings → Notifications** to send alarms beyond the app.

Email

Email goes through one **Mail server** panel. Pick a **Provider**: **Own mail server**, **Microsoft 365**, **Google** or **Other provider**. It fills in the usual server, port and **Security**: **STARTTLS (usually port 587)**, **TLS from the start (usually port 465)** or **None**, for a server inside your network. Fill in the account and choose **Send a test**: it sends one short email before you save, and a failure names the step that failed. The test arrives as

formatted text with a plain-text copy, like an alarm email.

Channels

Then add a channel of one type: **Webhook** (with an optional **HMAC signing key**), **Email**, **SMS**, **WhatsApp**, **HTTP siren / relay** or **Fire-panel watchdog**. An email channel sends to a group: tick **Roles** and **People**, or type **Other addresses**, and each person gets their own email. Each channel can be tested (the watchdog has no test) and switched off.

Secrets

Its secret (the mail password, an auth token, a signing key or an authorization header value) is stored encrypted and never shown again. The API needs an `APP_SECRET` of at least 32 characters for this, and a new `APP_SECRET` makes stored secrets unreadable.

Test bench

Test bench runs never notify.

- **Routes** choose which checks, cameras, zones and severities go to which channels, or only events that raise an alarm. **Escalate unacknowledged alarms** sends an alarm to more channels when nobody acknowledges it after **Escalate after** (at least 30 seconds). **Quiet hours** hold messages in a window on the site's clock (22:00 to 06:00 unless you change it) on the days you pick; a window past midnight belongs to the day it starts. Alarms still go.
- **Delivery log** lists each message as sent, queued, skipped or failed, and refreshes every 10 seconds. **Show older** pages back, and **Retry** sends a failed message again. An email that cannot be sent is tried again for up to an hour (**Keep trying for**, 60 minutes) before it counts as failed, and a late alarm says how late.
- **Email events** (on the same page) lists each system email, such as **Camera offline**: the **Least severity**, **Who gets it** and the **Wait** before it is sent, so a short blip sends nothing. When the problem clears, one back-to-normal email follows. Many notices of one kind in a burst fold into a single summary. The table is dimmed until a mail server is set, and a row that says **Nobody gets this email** has no recipients. Every automatic email ends with a **Stop these emails** link, and each person can also choose **Email preferences** in their user menu. Alarms and sign-in emails cannot be turned off.
- **HTTP siren / relay** switches an output on with the alarm and sends the off request when someone acknowledges it, or after a set number of seconds. If the server stops responding, detection switches on the outputs that allow it itself and keeps the events to send later; a replayed alarm switches no physical output.
- **Fire-panel watchdog** pings a fire panel or monitor every 30 seconds while detection is healthy and stops when it is not, so the panel raises a fault.

> ⚠ **Important.** Camera-based fire and smoke detection is an early warning that supplements the certified fire alarm system. It is not a certified life-safety function.

The System page

Open **System** (every role). It refreshes every few seconds.

- The header shows the build and two pills for the database and Redis.

- **Vision** shows the heartbeat (it turns red when stale), the provider and device, a load meter, the active models, GPU memory and utilisation, and its warnings in plain words, such as a missing PPE model or a GPU that is not available. It also shows go2rtc and the number of streams.
- **Licence.** Administrators see how the AI engine reads the licence. A difference from the server's reading shows in red with the engine's reason.
- **Versions.** Each part shows its own version as x.y.z.
- **Security check** (administrators only) lists five go-live checks: developer tools, the service password, the key that seals stored passwords, sign-in cookies, and access requests from the internet. Each row reads **OK**, **Look at this** or **Fix before going live**, with what to set. It shows no value.
- **System alerts** lists the open alerts first, then recent ones. An alert is raised within 90 seconds and also reaches the bell. Administrators choose **Acknowledge** to show someone has seen it, or **Reopen** on an acknowledged alert to bring it back. A site with helper servers, which take some of the cameras, gets an alert when one stops answering, needs an update, has its clock off or loses pictures, and a critical one when some cameras have no server with room. Each clears itself. Each open alert also says in one line what to do about it. Two weeks before the face data of people not seen for a year is deleted, an alert says how many people will be deleted and on which day.
- **Servers.** With a helper server added, the **Servers** panel lists the main server and every helper with its state and how busy it is, and **Recent moves** lists the cameras that moved between servers.
- **Capacity. Run a load test** plays a Test bench video on test cameras for the minutes you set, then says how many cameras fit. It adds load to the live system, so run it out of hours. The plan has three sliders: cameras, **Checks per second** and **Checks on each camera** (1 to 4, added in the order PPE, Security, Safety, Counting), and its estimate scales with each check's cost. The folded **Cost of each check** card lists what each check adds to a picture on this PC, each tagged **Measured** once the server has measured it and **Typical figure** until then.
- **Acceptance.** The nine items of the contract, each with the date of its last run and what it is tested with. The acceptance script records the performance item here; an administrator presses **Sign off** on an item checked by hand, and **Withdraw** takes it back. The passed items fold away.
- **Backups** shows when the last good backup finished, its size, what it holds and where it is, then the newest runs with their result and time taken. When no good backup has finished in 36 hours, or the last one failed, a warning opens in the bell and on this panel, and it clears after the next good backup. A restore test (`scripts/restore-check.sh`, run monthly by a timer or by hand, see `docs/BACKUP.md`) restores the newest backup into a scratch database and checks it. The panel shows when the last test ran and whether it passed, and a warning opens when no test has passed in the days set under **Settings → General & branding → Backups**. Administrators change or turn off the warning there. A server that uses the built-in database has no nightly backup, so the panel says so instead. Backups made before this update are not listed.
- **Download support file** (administrators, top of the page) makes one zip for your system provider's support team: versions, health, settings, recent warnings, camera and alert-channel status, licence, updates, backups and database row counts. Passwords, keys, email addresses, faces, pictures and people's names are taken out; IP addresses stay. Nothing is sent: you choose who gets it.

- **Training** shows the trainer, the current job, the device and the control state.
- **Storage** shows disk used and free, a state (Normal, Saving space or Critical), the forecast, the three kinds of file that use the most space, and a link to the Storage page. It also shows host CPU and memory. A meter turns amber at 70 percent and red at 90 percent.
- **Cameras** lists each camera's state, what the camera sends per second, checks per second, latency, people in view, last frame, recording state (recording, stopped, error, disabled or paused) and last error.
- **Slowed checks.** When the computer stays overloaded, loitering, phone, smoking and forklift-proximity checks run at half rate, one by one, and System lists them. The far look of **Find far people and vehicles** can also drop to half rate, listed as "Far people and vehicles, half rate". Fire, smoke and person-down checks are never slowed. If that is not enough, the header shows a **Not enough compute** chip that links to System, and System warns "Not enough compute for the enabled checks". Fewer cameras, a lower inference rate, smaller models or a GPU will help.
- **Audit log** (Admin and Supervisor) lists the last 100 actions: time, user, action, target and details.

Helper servers

Open **Settings → Helper servers** (administrators) to let other computers of yours share the work when this server is full. Pick the site's country in Settings → Privacy first: helper servers follow its rules for sending footage to other computers.

1. Choose **Add a helper server**. Name the computer, say whether it is **On this site** or **Somewhere else** (then pick its **Country**), and tick **This computer belongs to my organisation.** Choose **Add and get a code**.
2. The window shows a join code, how long it works, **This server's address** (pick one the other computer can reach) and a command for Linux, macOS and Windows, each with **Copy**. Run it on the other computer, where the platform is already installed. It joins within a minute.

- **The card of a helper.** It shows the state, how busy the helper is, its cameras, version, link time and when it was last heard. **Make a new code** works while it waits to join. **Pause** stops giving it cameras, **Empty** moves its cameras to other servers, and **Remove** cuts it off at once and wipes what it keeps (you type its name to confirm); its cameras then move to other servers, so use Empty first and wait if you want them moved before the cut. **Show removed** lists removed helpers.
- **Change a helper. Name**, **Cameras** (what it may take), **Run the AI on** (Auto, GPU or CPU), **AI engine** (Automatic) and **Space for waiting pictures** (5 GB, 1 to 100 GB). For a helper somewhere else, **Transfer reason** shows the legal reason footage may go to it; with none, footage cannot go there. Faces, attendance and number plates stay on the main server.
- **Failover. Move cameras off a silent server after** (60 s, 20 to 600) moves a silent helper's cameras to a server with room, and the next setting does it sooner (20 s, 10 to 120) for cameras with fire, smoke or person-down checks. Nothing is counted twice. **Accept helpers from the open internet** (Advanced, off) lets helpers connect from public addresses, and you forward port 7443 yourself.

- **Automatic moves. Move cameras automatically** (on) lets cameras move to a helper with room when this server stays too busy, and come back later. **Move cameras when busier than** (80 %, 50 to 95 %, at or under **Slow extra checks above**) and **Bring cameras back below** (50 %, 20 to 80 %, at least 10 points lower) set the levels, **Busy or quiet before a camera moves** (2 min, 1 to 30) sets how long they must hold, and **Most automatic moves an hour** (6, 1 to 30) stops cameras moving back and forth. **Upload this server can give helpers** (Advanced, 20 Mbit/s) caps the video a helper may take through this server. **Move cameras...** on the Helper servers page pins several cameras to a server in one go, shows how much video each sends, and **Let the system decide** sets them back to automatic. Each helper card also has **How it gets camera video**, **Most busy**, **Most cameras** and **Internet speed toward it**.
- **Where a camera runs.** The Cameras list shows "On" and the helper's name, or "Waiting for room" in amber, for a camera that does not run on the main server; recording is unaffected. An event says **Checked on** and the server.

Storage and retention today

- The **Storage** page is in the Admin group of the menu and open to every role. It shows free space, what uses the disk, a usage chart with a forecast of when the disk fills, the alerts and the recent activity. Free space is never more than 10 seconds old. The size of each kind of file is measured every 5 minutes.
- The page measures **Exports** as its own kind of file. It has no **Free up space...** button, because finished exports delete themselves after 14 days. **Edit rules** opens **Settings → Storage**, and the recent activity lists "Old records removed" with what the job removed. **Accuracy check evidence** has its own line too: check-set pictures are never offered for clean-up, and weekly check pictures are removed after 90 days.
- **Rescan** (administrators and supervisors) measures every kind of file again within a minute. It deletes nothing.
- Sizes are decimal: 1 GB is 1,000,000,000 bytes.
- Under **Recordings by camera**, the **Keep recordings for** slider sets how many days recordings are kept. Cameras that hold the current value follow it; a camera with another value changes only if you tick **Apply to all cameras**. The page shows the disk the new value needs and what a clean-up would remove. **Save** changes it, and **Undo** in the message puts the old values back.
- Alerts appear on the Storage page when the forecast puts the disk at the clean-up line within `storage.alerts.warnDays` (14 days by default), when recording capacity stays under 90 % of the target for a day, when events waiting to reach the server pile up, and when a measurement fails. The forecast says "Still learning" for the first 24 hours.
- Every 60 seconds the vision worker deletes recordings older than each camera's **Retention**, 14 days by default and at most 30 under the Strictest privacy profile. While free space is under `minFreeGb` (50 GB by default, the clean-up line on the Storage page), it then deletes the oldest recordings across all cameras. It never cuts a camera's newest 14 days, and never deletes a segment newer than 10 minutes or one a queued evidence clip still needs. Only the recordings folder is touched.
- Below `hardMinFreeGb` (10 GB by default) recording pauses on every camera that is not marked Critical, and resumes when free space recovers. A Critical camera keeps recording. Below that floor the Storage

page shows a red note, and uploads are refused with HTTP 507, and pictures and clips may fail to save until space is freed. An event whose snapshot could not be saved says **Snapshot not saved (disk full)**.

- Automatic deletion follows only the age rules in **Settings → Storage**. Legal holds, open incidents, approved samples, datasets and active models are never deleted automatically. You delete an event, a camera (with its events and its review samples) or a model by hand.
- **Clean-up** shows a preview first: the number of items and the space they would free. Nothing is deleted until you confirm. A large clean-up needs you to type `DELETE`, and a running clean-up can be cancelled.
- **Space saving** is set to Auto, On or Off. Auto turns it on when the disk is forecast to fill and off again when there is room. A banner and a chip in the top bar show when it is on. Critical cameras are exempt.
- Each kind of data has its own age in **Settings → Storage**, and a **Never** switch keeps it where allowed. No event record goes before the minimum event retention. The site's privacy profile caps each age: under the Strictest profile, event clips and snapshots go after at most 90 days, evidence and event records after at most 3 years, and sightings and visits after at most 90 days. A save above a cap says what the cap is.
- Test bench image results are cleared after 24 hours.

The data folder is laid out like this, which matters when you move or back up the disk:

```
recordings/<camera>/YYYYMMDD_HHMMSSZ.mp4   60-second segments, named by UTC time (Z marks UTC)
events/<id>/raw.jpg, annotated.jpg, clip.mp4
datasets/inbox, versions/<id>, golden
models/<id>/        training/<job>/        cache/        spool*/
```

Backups and upgrades

- **Backups.** A backup holds the database and the files you cannot rebuild: models, branding, the golden set, dataset manifests, and the face pictures and templates. Recordings and event evidence are optional (`--with-media`). The nightly backup keeps the newest backup of each of the last 7 days and 4 weeks. Copy backups to another machine, because a backup on the same disk is not a backup. A backup never holds `APP_SECRET` or the export signing key, so keep the secret with your records; after a restore onto a new server, choose **Make a new site key** in Settings → Sharing and give recipients the new fingerprint. A setup backup (Setup backup) is not a server backup: it holds only the set-up, in a file with its own password, and leaves out events, recordings and the licence. Treat a backup as personal data, because it holds face pictures. Restore refuses an archive that holds a symbolic link. `docs/BACKUP.md` covers running, scheduling and restoring.
- **Upgrades.** Plan a maintenance window, because vision stops watching while the images rebuild. Back up, update the code, start, check, and know how to go back. Database changes only move forward, so a rollback restores the backup you took first. `docs/UPGRADE.md` has the commands.
- **Updating an installed package.** An administrator opens **Settings → Updates** and chooses **Choose update file (.zip)**. The page checks the file, then shows what is new and how long the platform will stop. Click **Install** and enter your password. The page shows four steps (checked, installed beside the current

version, switched, checking health) and reconnects by itself after the restart. If the new version fails its checks, the platform goes back by itself, unless that release is marked not to. Then the new version stays installed, the page says so, and you choose **Roll back** yourself in **Settings → Updates** or run `platformctl rollback` on the server. For 7 days **Settings → Updates** shows **Roll back** beside the version that is kept. When the update changed the database, rolling back returns it to the backup made before the update, keeps a copy of the newer database for 7 days, and the page says what that loses. On a Windows or macOS install with the built-in database, an update that changes the database layout makes that backup first, and rolling back puts it back. Otherwise it keeps all your data. An update file for another system, an older version or the version already running is refused, and nothing changes. A version older than any installed before is refused even after a rollback. **Discard** throws a checked file away. If the server refuses its update state, run `platformctl repair` to rebuild it from what is installed. A packaged server accepts only files signed by the vendor; the first signed release is not published yet.

- **Installing a release on a Linux server.** An operator installs a signed package on a Docker server with `sudo platformctl install --from <file>`. The command shows the first administrator password once; if you run it again after a failed first try, it says where that password is kept. A server that runs from the source code moves onto a release with `sudo platformctl install --adopt <folder> --from <file>`: it asks first, takes a backup, keeps the data, and leaves the old copy in place. The **System** page calls the store behind live updates **Key-value store**.

Licence

The **Licence** page shows what this installation is licensed for. Everyone who can open it sees **Status**: the state (No licence, Not valid, Trial, Active, Expires soon, Expired with a grace period, or Expired), the customer and site, when it was issued and when it expires, how the AI engine reads it, the cameras used against the limit, and the modules, each On or Locked. Administrators who can manage the licence also see **Installation ID**, **Activate online**, **Apply a licence file**, **Online licence** and **History**, which names the source of each change: File, Online activation or Online renewal. The page reads the licence again every 15 seconds.

There are two ways to activate. A server with internet can use a **product key**. A server without internet uses a licence file. Both give the same licence.

- **Product key.** Where your system provider has switched it on, the **Activate online** panel shows. Type the key in **Product key** (`XXXXX-XXXXX-XXXXX-XXXXX-XXXXX`) and choose **Activate**. The box adds the dashes as you type and reads the letters O, I and L as 0, 1 and 1. The server sends only the key, its installation ID and the software version. If the key is wrong, the page says what to check.
- **Licence file.** The **Apply a licence file** panel is for servers without internet: upload the `.lic` file from your system provider, or paste its text. Give the Installation ID to your system provider so the file fits this server.
- **Renewal.** An online licence renews itself once a day. The **Online licence** panel shows **Last renewed** and when it runs until. Choose **Renew now** to ask at once; the answer is "Licence renewed", "Licence is up to date", or "Online renewal is not set up on this server, so nothing was checked." If a renewal fails, the licence keeps working until it expires, and the panel says "The last renewal failed on" a date and why. The server tries again every day.

- **Moving a server.** Choose **Move to another server...**, type the product key and choose **Release this server**. The server goes back to the basic product at once: up to 4 cameras keep running, the others pause, and no recordings, events or settings are deleted. Then activate the same key on the new server.

A licence sets how many cameras may run, which modules are on (Own branding is one of them), and the product name, logo and accent colour. A trial licence runs for up to 90 days. A paid licence that expires gives a 14-day grace period before paid features lock; a trial key that was released ends at once, with no grace.

If the page says *This server has already used a newer licence*, the database was probably restored from an older backup. Apply the newest licence again. An online licence also comes back at its next renewal. A licence dated more than a day ahead of the server's clock reads *This licence is not valid yet*.

- **Cameras over the limit.** Applying a licence that allows fewer cameras says "Licence applied. It allows 4 cameras: the newest 2 are paused." Those cameras are paused, not deleted: their settings, zones and rules are kept, and the oldest cameras keep running (chapter 6). Status says "2 paused by the licence limit".
- **Modules.** A feature outside the licence stays in the menu and the command palette, greyed with a lock. Opening it says "<Feature> is not included in your licence." and "Talk to your system provider to unlock." Administrators also see the Installation ID and an **Open the Licence page** button. Rules and the Test bench need the PPE module, People needs tracking, Datasets and Training need training, and Video search needs search. A refusal by the licence shows one message however many requests were refused.
- **Empty pages.** An empty page does the same: its next step is greyed out with the reason when your role cannot take it.
- **Without the PPE module.** Cameras still show live video and keep recording, with motion detection and the disk guard running, but no checks run and no boxes are drawn.
- **The AI engine.** The engine works out the licence itself. The Status panel reads **Agrees** or "The AI engine disagrees:" with the reason, and an unreadable or old report reads **Not reporting yet**. The System page shows a difference in red.
- **Brand.** The product name, logo and accent colour come from the licence. Settings shows them read-only, and an attempt to change them is refused, unless the licence includes Own branding: then an administrator sets them in Settings, and your values win over the licence's. If the licence lapses, your values are kept but not shown until it is renewed. Reports use the brand in force. The site name is always yours to edit.

A banner at the top of every page tells the state. It says "Trial — 12 days left" on a trial, and for everyone during the grace period "Licence expired — paid features lock on" a date. Administrators also see "Licence expires on" a date, "Licence expired — paid features are off", "The licence is not valid:" with the reason, "Server clock is behind", "The AI engine disagrees:" with the reason, and, for an online licence, "The licence could not be renewed online:" with the reason. The Overview shows a lock, or a retry, if the licence cannot be read.

Privacy desk

Open **Privacy** in the sidebar (administrators, or anyone given the privacy permission) to answer requests from people about their data and keep the proof. It has seven tabs.

- **Checklist.** The first tab. It lists what this site has done for privacy and what is left. Each row can carry a **Done on** date, a **Review on** date and a note. Nothing here locks anything.
- **Requests.** Choose **New request** the day someone asks to see, correct or erase their data. The deadline counts from the received date, and reminders come 3 days before and on the day. Open a request to check who is asking, find their data, export it, erase it and close it. Erasing cannot be undone and asks for your password again. A closed or refused request stays on record for 3 years.
- **Consent.** Shows each named person and whether consent is recorded or still waited for.
- **Holds.** Put a person, or a stretch of footage of at most 30 minutes, on legal hold. Nothing about them is deleted or erased until the hold is released.
- **Breaches.** Log lost, leaked or wrongly seen personal data. The deadline for telling the authority starts when you found it, and reminders come before it passes.
- **Documents.** Camera signs, worker and visitor notices, the retention policy and other drafts for this site. They are template words: edit them to fit your site before you use them.
- **Transfers.** Lists helper servers in other countries. Record the reason footage may go to each; a helper server in a country with no reason takes no footage work.

Health check

`GET /api/health/deep` reports the database, Redis, go2rtc, free disk space, the age of the vision heartbeat, the trainer and the build version. Call it with an administrator session, or with the `x-service-token` header from the LAN, and point your monitoring at it.

```
curl -sS -H "x-service-token: $SERVICE_TOKEN" http://localhost:4000/api/health/deep
```

The status is `ok` once vision has loaded its model and sent a heartbeat, up to a minute after start. `degraded` means an optional part is failing: go2rtc, the trainer, low disk or `checks.secrets.unreadable`. HTTP 503 means the database, Redis, the disk floor or vision is failing.

Remote access

Caddy provides HTTPS on the LAN. For people outside the LAN, `docs/DEPLOYMENT.md` describes hosting the same web app on Netlify and reaching the on-site API through a Cloudflare Tunnel, with a security

checklist to complete first. Remote viewers get live video through the server, about a second behind, because direct video needs port 8555. Their tiles fall back to pictures when that fails.

18 Troubleshooting

Start with the System page. It shows whether vision, the trainer and each camera are healthy. Items marked IT need access to the server. Logs live in the API's output, `recordings/<camera>/ffmpeg.log` and `training/<job>/train.log`.

What you see	What to check
A camera is Offline	Is it powered and reachable? Check the address and password on Cameras → Edit. IT: test with `ffmpeg -rtsp_transport tcp -i URL -t 5 -f null -` from the server, and check the firewall.
A tile says "Waiting for frames", or the snapshot on the Zones tab shows NO SIGNAL	On System, is the vision heartbeat recent? Is the camera Online? IT: check that vision is running and that Redis is reachable from it.
Events never fire	Is a PPE model active (Models, and the `no_ppe_model` warning)? Does the zone have rules (Rules shows "Not checked" otherwise)? Is the camera in PPE mode and PPE checking on? Are people taller than `minPersonHeightPx`? Does the model's class map cover the PPE types (`ppe_model_unmapped`)?
A figure says "partial" or shows "–"	Look at the Cameras strip on Overview: each camera shows its coverage over the last 24 hours, and a camera under 50% is named and left out. Check the working hours under Settings → Counting & coverage. If it says the server checked fewer pictures than the camera is set to, lower the camera's checks per second or add capacity.
"Too small to judge" is a large share of the time	People are small in that camera's picture. Aim or move the camera, or lower the smallest person size under Settings → Counting & coverage after checking the result.
Fully equipped workers show as missing	This is the live profile reaching its limit. See chapter 1. Compare in the Test bench, review the events marked Not a violation, and train on your footage (chapter 16).
An event has no clip (skipped)	Record was off for that camera, or the recorder had an error. IT: read `recordings/<camera>/ffmpeg.log`.
Clip extraction failed	The recording segments are missing because retention is too short, or the video cannot be copied without re-encoding.

What you see	What to check
A camera says "Paused — licence limit"	The licence allows fewer cameras than you have added. The newest cameras pause and the oldest keep running. Nothing is deleted. Ask your system provider for a larger limit (Licence page).
A camera says "Password needs re-entering"	The saved password can no longer be read, for example after a restore on another machine. Open the camera, type its password and save.
A page says "<Feature> is not included in your licence."	The licence does not include that module. Administrators see the Installation ID and the Licence page; everyone else asks the system provider.
A page says "This app is switched off"	An administrator switched the app off. Choose **Open Apps** and switch it on; nothing was deleted. If the app is locked, the licence does not include it (Licence page).
Training stays queued	Is the trainer running (System → Training)? It picks up work every 30 seconds. IT: read `training/<job>/train.log`.
A job is paused	It is outside an idle window or vision is busy. Use Return to automatic scheduling on the Training page, or create the job with Ignore idle windows.
The top bar chip says CPU	Open it and choose GPU. It says "not available" when no GPU is found. IT: check `nvidia-smi` on the host, the driver version (580 or newer) and that the GPU is Turing or newer. On Windows, a "GPU is lost" message needs a restart.
A model shows "Inspecting weights..." for long	Is the trainer running? It reads new models every 30 seconds. The file must be a valid Ultralytics .pt file.
Model Download fails	The server needs outbound HTTPS to huggingface.co. Otherwise download the file elsewhere and use Import model.
The People page is empty	Is re-identification on? Does System → Vision show a Re-ID row? People must be tracked for at least `reid.minTrackAgeS` (1 second).
One person has many ids	Raise `reid.maxGapS`, lower `reid.threshold`, or add camera transitions. Or merge the ids with Same person as...
Different people are merged	Raise `reid.threshold` and switch off `crossCameraOverlap`.

What you see	What to check
Cameras are offline after a restore or a changed `APP_SECRET`	IT: `GET /api/health/deep` shows `checks.secrets.unreadable` above 0. Put the old `APP_SECRET` back, or type the camera passwords again.
The API exits at start with `APP_SECRET is required`, `SERVICE_TOKEN must be at least 32 characters` or `DATABASE_URL is not a valid URL`	IT: a production secret is missing, a placeholder or too short. A database password that contains `/ @ : # ?` breaks the address, so use `openssl rand -hex 24`.
The Python services get a 401 from the API	IT: `SERVICE_TOKEN` differs between the API, vision and the trainer.
Vision cannot reach the API	Events and samples wait in `STORAGE_ROOT/spool` and are sent again every 30 seconds.

Getting support

Work through this chapter and the operator guide first. When the problem remains, ask your system provider's support team for help.

1. **Note the facts.** The time of the problem, the camera or page involved, and what the screen showed.
2. **Download support file.** An administrator chooses **Download support file** at the top of the System page. It makes one zip with versions, health, settings, recent warnings and errors, camera and alert-channel status, licence, updates, backups and database row counts. Passwords, keys, email addresses, faces, pictures and people's names are taken out; IP addresses stay, because faults cannot be found without them. Nothing is sent: you choose who gets the file.
3. **Send the file** with the facts from step 1 to your system provider's support mailbox.

Your system provider sets the priority of each request with these words:

- **Priority 1.** Alarms are not raised or not sent, or the whole platform is down.
- **Priority 2.** One or more cameras, or one check, does not work, and the site has a workaround.
- **Priority 3.** A fault with an easy workaround, or a wrong figure in a report.
- **Priority 4.** A question or a change request.

Response times are in your support agreement. Never send passwords, keys or tokens, a backup file or pictures of faces. A setup backup also holds password hashes, so keep it as carefully as a password.

19 Glossary

Absence class

A model class that stands for a missing item, such as `NO-Hardhat`. It counts as evidence of absence and never as presence.

Accuracy check

A blind weekly set of moments that reviewers answer, to measure how often detection is right on each camera. The answers never change an event.

Alarm

An event from a check that must be acted on at once, such as fire or person down. It shows a banner and a siren on every page until someone acknowledges it.

Area

A named group of cameras, such as a warehouse or a production hall. A camera with no area sits in "No area".

Boxes

The platform's boxes, labels and zones, drawn as a layer over the camera's own video. An event's snapshot comes both ways: Annotated and Raw.

Candidate model

A model that is trained, downloaded or imported but not active. It becomes active only when someone activates it.

Check

One kind of thing the platform looks for in a zone: PPE, a safety check such as fire or person down, or a security check such as a restricted area, an unknown face or a camera blocked or moved.

Checks per second

How many pictures a second the model checks on a camera.

Class map

The table on a PPE model that says what each of its classes means for the platform.

Counting line

A line drawn on a camera's picture. Everyone who crosses it is counted in or out.

Episode

One stretch of time in which one person is missing PPE. It opens one event.

Event

The record of one episode: snapshots, a clip, the missing items, a severity and a review state.

Fine-tune

Train an existing model further on new frames instead of starting from nothing.

Fixed item

Something that looks like fire, smoke or a forklift but never changes, such as a lamp or a steam vent. Checks ignore a still copy of it.

Golden set

Reviewed samples kept as permanent test data. Models are scored on them and never trained on them.

Golden-set score

How well a model finds the items on the golden set, from 0 to 100%. Higher is better; a new site needs 85% to go live without a written reason.

Grace

The seconds a person must be inside a zone before the platform judges them.

Incident

A violation worked as a case, with a status, an owner, notes and a history.

Kiosk

The Monitor wall at full screen with the menus hidden, for a control-room display.

Legal hold

A mark on an event that keeps its evidence. Automatic deletion never removes a held event.

Main stream and sub stream

The main stream is a camera's high-quality stream, and the platform records it. The sub stream is its low-resolution stream, and the platform decodes it for detection, unless the camera is set to read its main stream for detection.

Mining

Queueing frames that look worth reviewing: uncertain, low-confidence or new-looking ones. Its settings are under **Frames to review**.

Person id

`P-` followed by eight characters. One person as the cameras re-identify them.

Rule

One PPE type required in one zone, with a severity and an optional schedule.

Sample

A frame in the review queue, with boxes you can correct.

Security mode

A camera with its PPE chip off: it tracks people and judges no PPE.

Severity

High, Medium or Low, set per rule. An event takes the highest severity among its missing items.

Track

One person followed through consecutive frames on one camera.

Version

An unchangeable snapshot of the approved samples, used for one or more training jobs.

Violation

A person missing a required item for long enough to be confirmed.

Visit

One stay of a person in a zone. A return after a long time out of sight starts a new visit.

Warning

A person missing a required item for longer than the rule's warn-after time, but not yet long enough to be a violation. Shown in amber.

Whole frame

The zone every camera gets automatically, covering the entire picture.

Zone

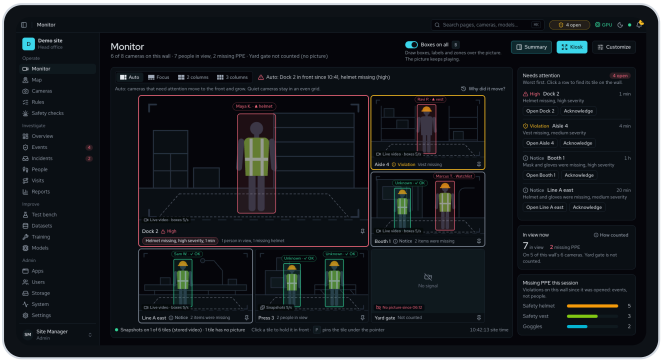
An area of a camera's picture where rules apply.

20 How the finished system will look

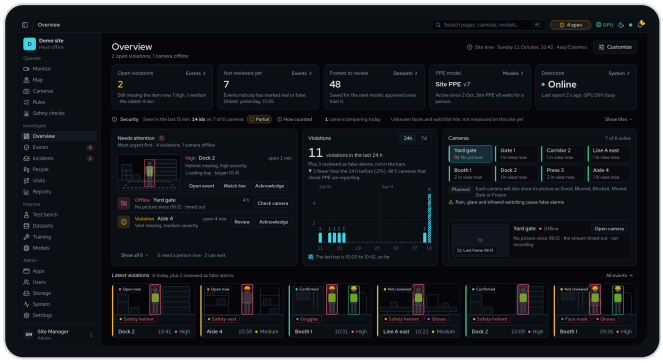
When it is fully built, the platform runs every camera app on one server: recording, PPE, security, safety checks, counting, number plates, video search and attendance, each switched on and off on an Apps page. Below are previews of every screen, drawn from the approved designs and grouped by milestone. Each carries its status: Available, In progress, Next, Planned or Later, and a milestone says Partly available

when only some of its screens ship. Screens that are not Available show the plan, and details can change before release.

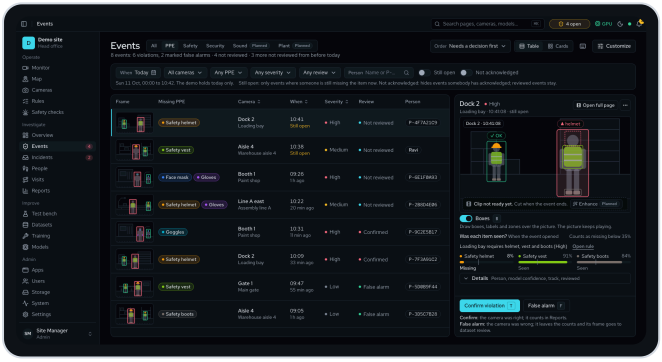
Today's pages Available



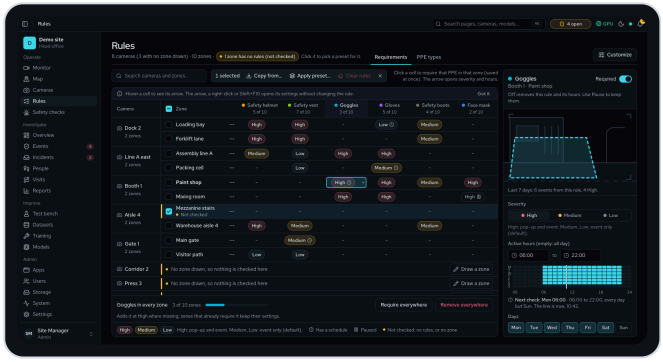
Monitor Available The Monitor: cameras that need attention move to the front with the reason shown, one switch draws or hides the boxes over live video, and every tile says how fresh its picture is.



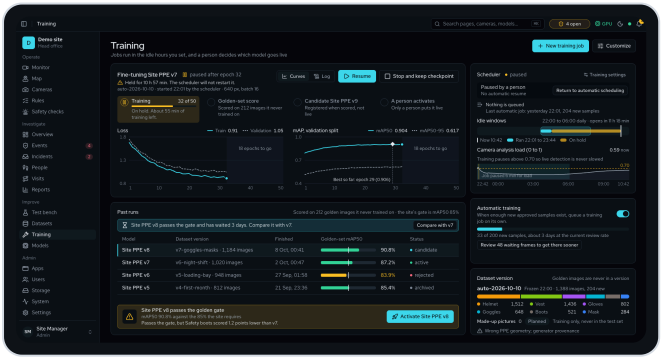
Overview Available Camera health, open violations and the day's trend at a glance, with what needs attention first at the top.



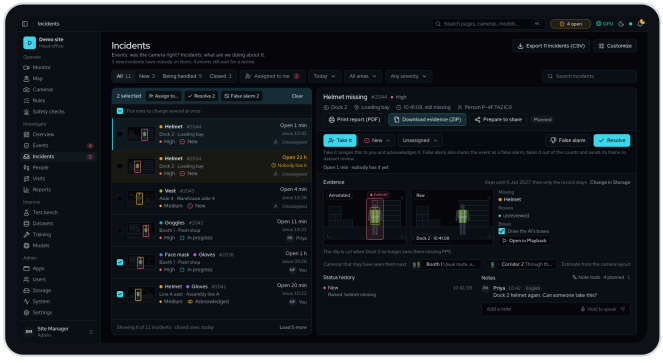
Events Available Every violation with its snapshot, missing PPE, severity, person and review state, filtered in one click, with the evidence beside the list.



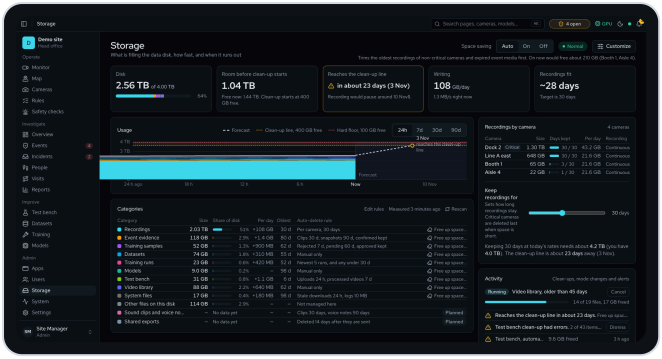
Rules Available The PPE each zone requires, with a severity and a schedule per cell, edited in one matrix.



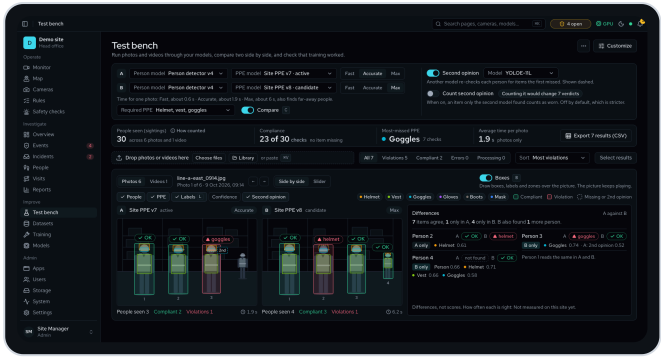
Training Available Training runs in the idle hours you set, with live progress, the golden-set score and activation only by a person.



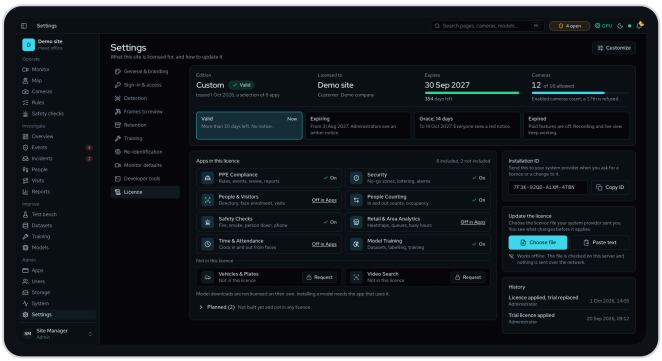
Incidents Available Each confirmed violation becomes an incident to assign, discuss and resolve, with its evidence pack ready to download.



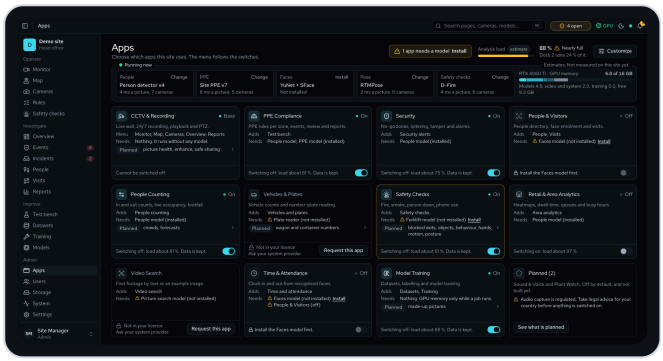
Storage Available Disk use by category and camera, a forecast of when the disk fills, space-saving modes, and alerts well before it gets tight.



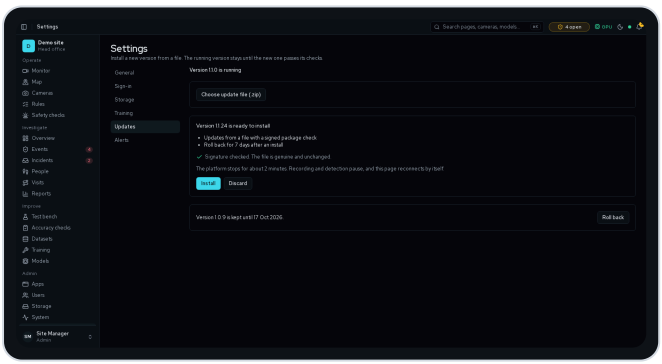
Test bench modes Available Fast, Accurate and Max analysis, a model choice per run, and side-by-side A/B comparison with a second-opinion model.



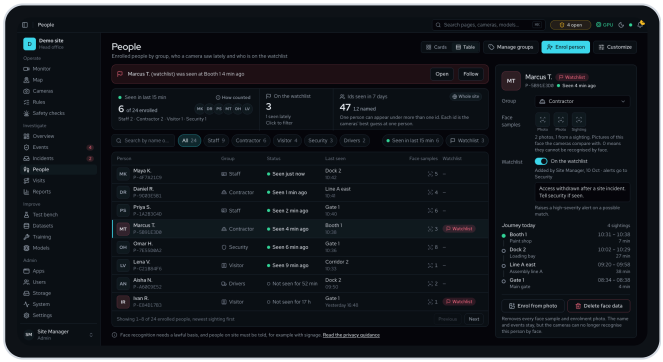
Licence Available The edition, modules, camera limit and expiry at a glance, with a licence file applied offline.



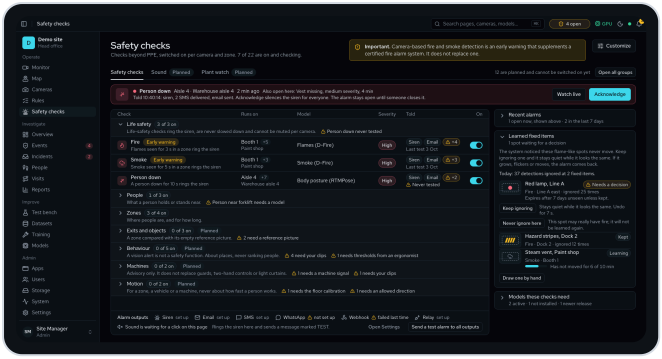
Apps Available Switch whole apps on and off, several at once: a camera viewer, PPE, security, counting and more in one installation.



Updates Available Upload a signed update file, see what is new and how long the platform stops, and roll back in one click for 7 days.



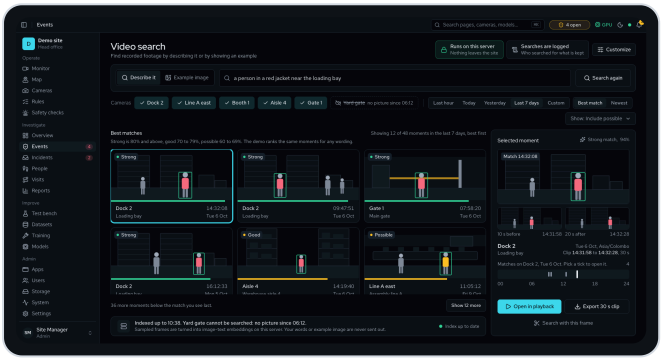
People directory Available A directory with groups and a watchlist, faces enrolled from a photo, who is on site now, and face data deleted on request.



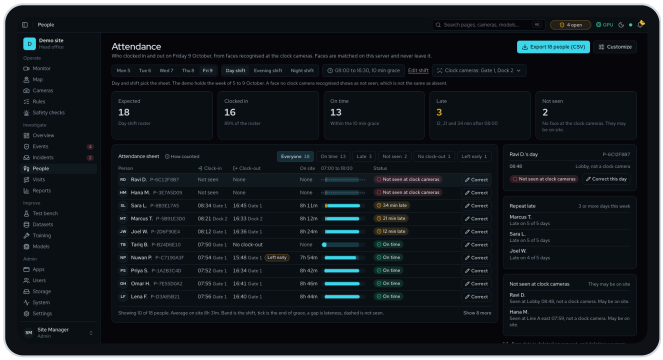
Safety checks Available Fire, smoke, person down, phone use and forklift proximity, each with its own model, severity and alarm. Fire and smoke are an early warning beside a certified fire alarm, never a replacement.



Playback and export Available A timeline across cameras with event marks and recording gaps, synced playback with the same Boxes switch as live video, the cameras a person is likely to reach next, and clips exported with a tamper-evident manifest.

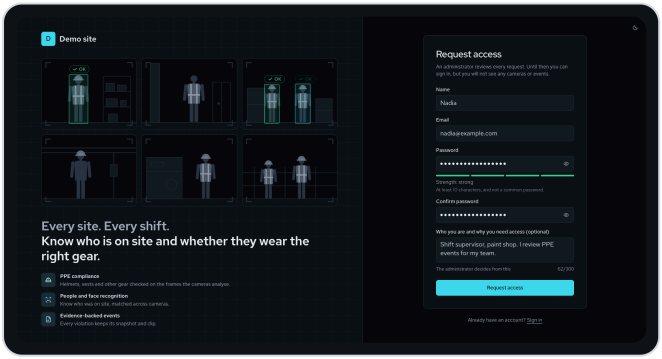


Video search Available Find recorded footage by typing what you are looking for, like a person in a red jacket, searched entirely on site.

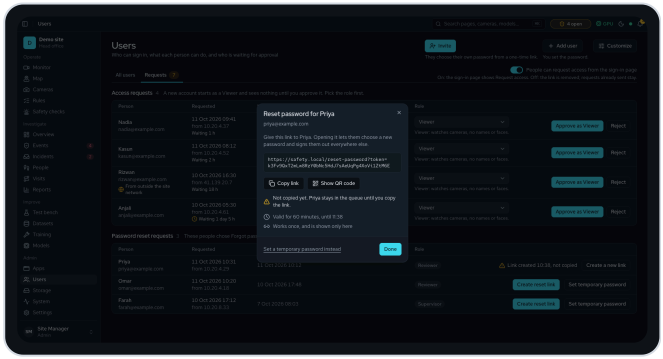


Time and attendance Available Clock-in and clock-out from faces recognised at chosen cameras, a daily sheet, late and not-seen lists, and CSV export. Off until the site sets it up.

M1: Accounts Available

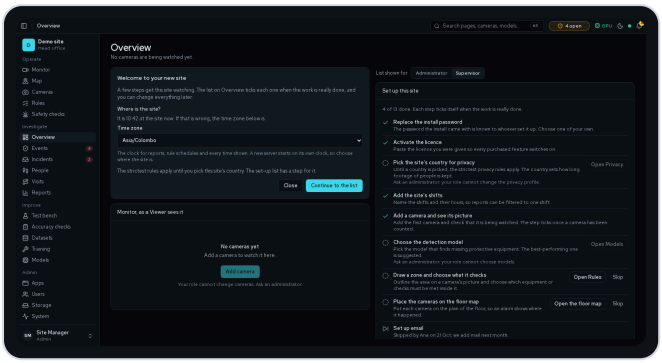


Request access Available A new person asks for access from the sign-in page, then sees nothing until an administrator approves them and picks their role.

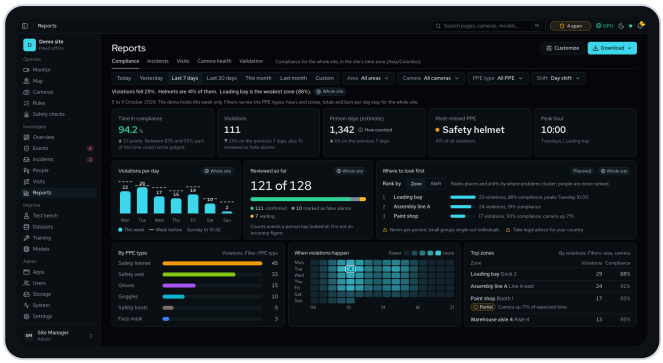


Approvals and reset links Available Administrators approve or decline requests, pick each person's role, send invitations, and hand over a one-time link when someone forgets a password, with no email needed.

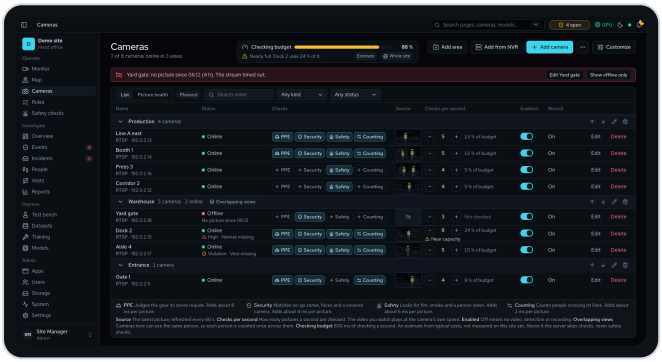
M2: Client go-live Partly available



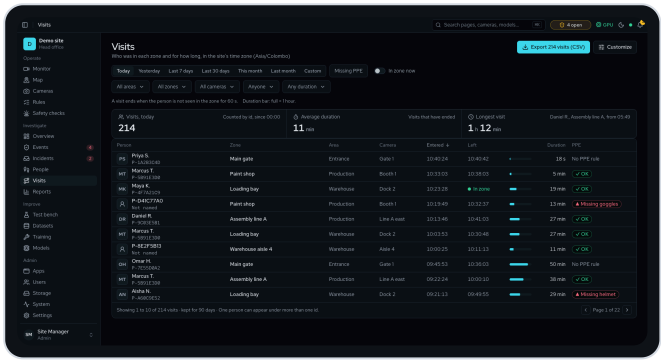
First-run set-up In progress A new site's first administrator sets the site's clock in a short welcome, then follows a list that ticks each step when the work is really done.



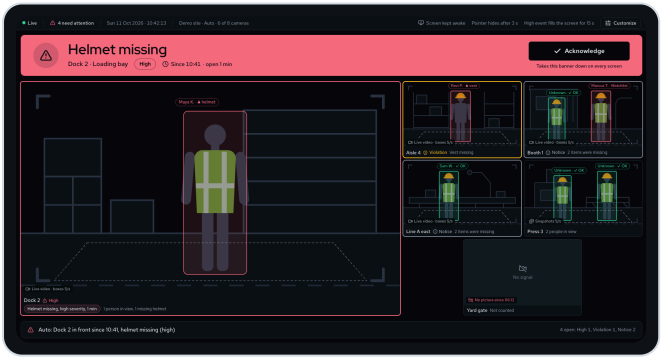
Reports Available The Reports page: KPIs against the previous period, an hour-by-weekday heatmap and top zones, and every report as a branded PDF or CSV.



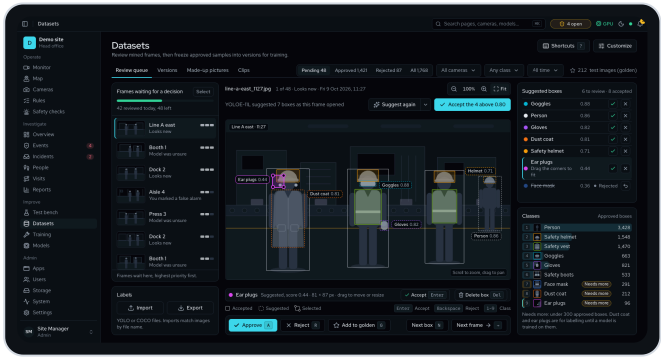
Cameras and areas Available Cameras grouped into areas, with the checks each one runs, how many pictures a second it is checked and what that costs, plus camera presets, a test connection and an NVR's channels added in one step.



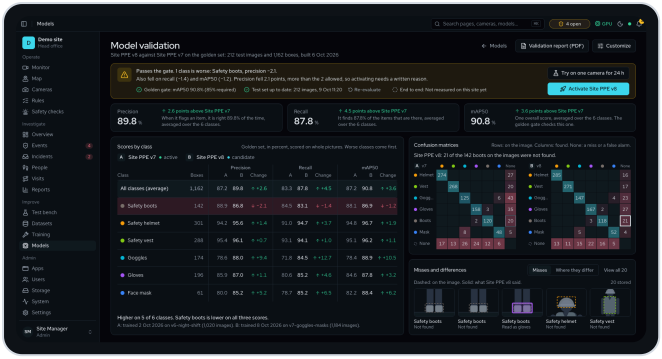
Visits Available Zone-level visits with how long each person stayed, in a visit log to filter and export as CSV.



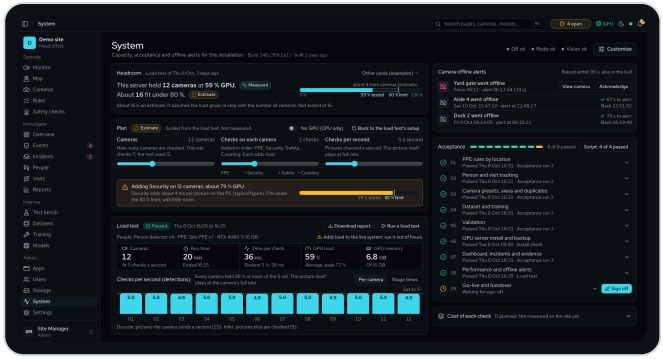
Second screen Available An always-on wall for a second monitor: the camera that needs attention stands in front, a high event shows a banner until someone acknowledges it, and a strip always shows the connection, the open count and the clock.



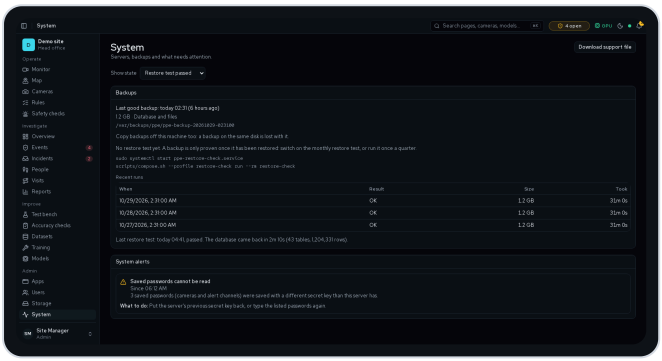
Labelling Available A faster labelling tool that suggests boxes for a person to accept, fix or reject, with label import and export.



Model validation Available A candidate model measured against the active one on the golden set, class by class, with a printable validation report.

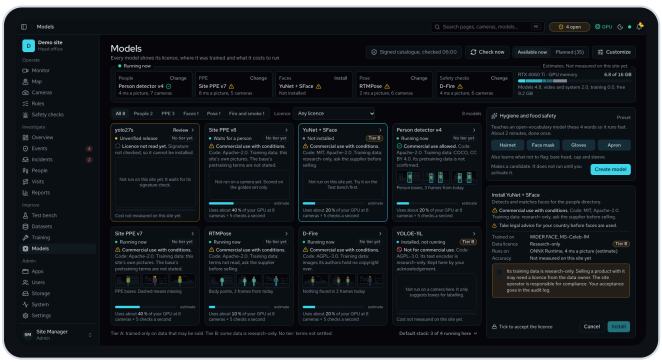


Capacity and acceptance Available Load-test results per GPU, the acceptance checklist, and camera-offline alerts within 90 seconds.



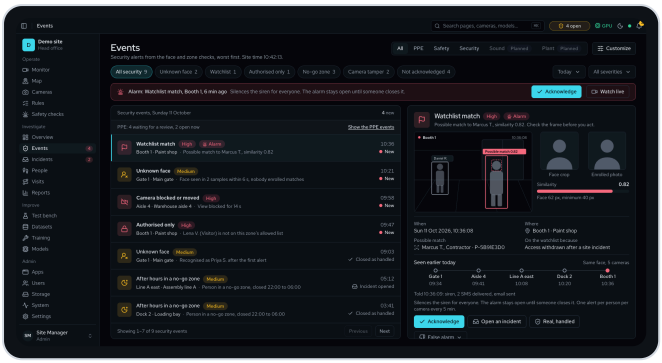
Backups and support Available When the last backup finished, whether a restore test passed, a support file that is safe to send, and what to do about each system alert.

M3: Sell-ready white label Available

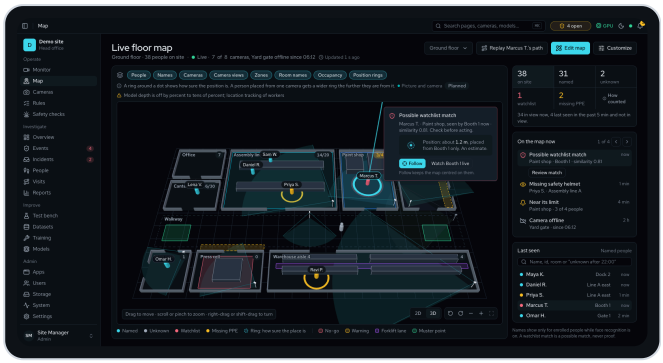


Model library Available Every model carries a licence badge, and installing one asks no licence question. The hygiene and food-safety preset is built.

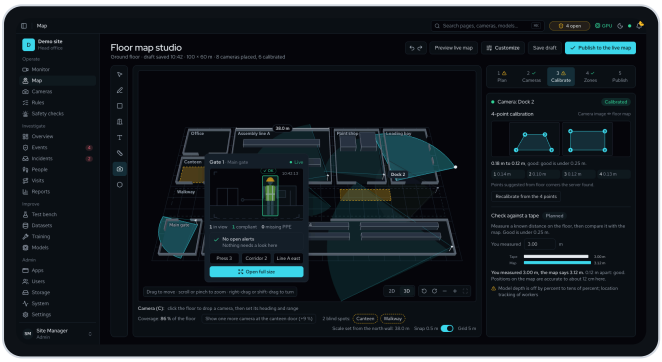
M4: Security system Partly available



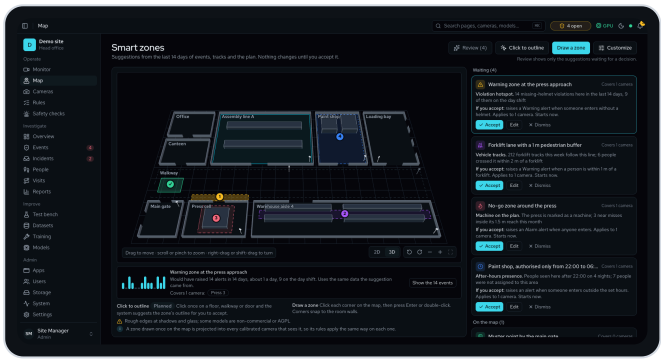
Security alerts Available Alerts for unknown faces, watchlist people, authorised-only zones, after-hours presence, and covered or moved cameras.



Live floor map Available The site's floor map, live and in 3D: people named when enrolled, vehicles, alarms, occupancy and camera coverage. Turn or zoom it; point at a camera for its live view and nearby cameras. Vehicles on the map are still in progress.

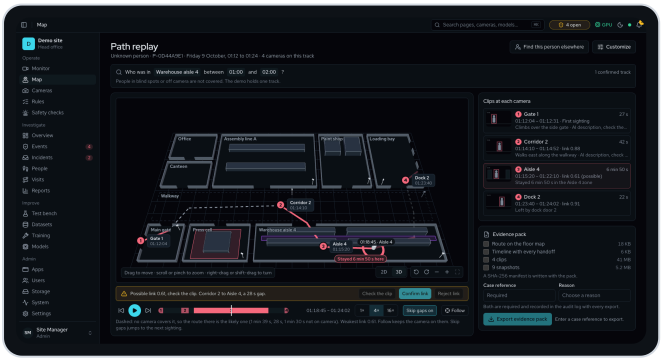


Floor map studio Available Draw the site, upload a plan, or photograph a rough sketch and let the on-site AI turn it into a clean floor map, then place each camera on it.

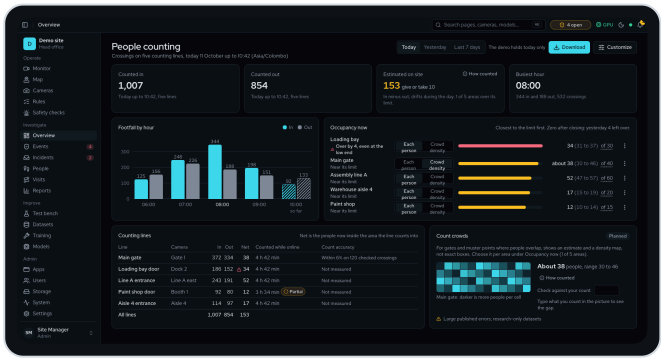


Smart zones In progress Zones drawn once on the map cover every camera that sees them, and the map suggests no-go and warning zones from the site's own events and tracks.

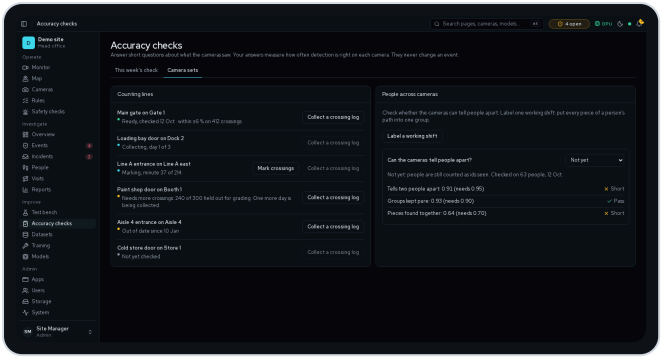
M5: More apps Partly available



Path replay Later Follow one person across every camera on the floor map, replay their route with the clips, and export it as evidence, such as an unknown visitor's path at night.



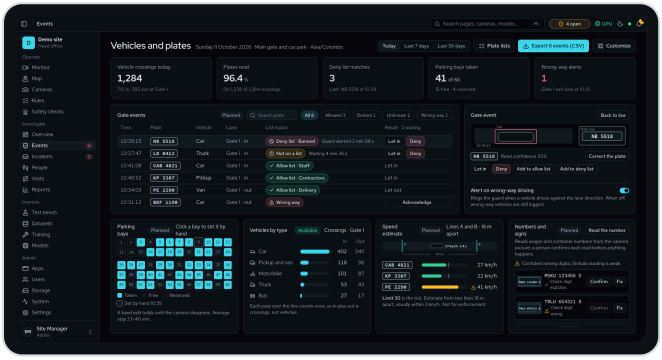
People counting Available In and out counts on counting lines, live occupancy per area with limits, and footfall by hour.



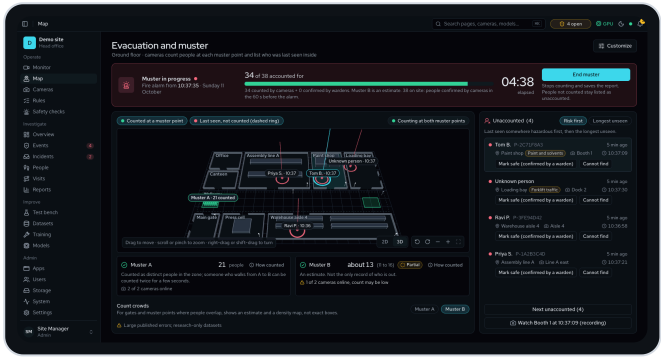
Accuracy checks In progress Reviewers answer short questions about what the cameras saw, mark every crossing on a counting line, and group a shift's people, so each count carries a measured accuracy.



Area analytics Available Movement heatmaps, dwell time per zone, queue length and wait, and the busiest areas and hours.

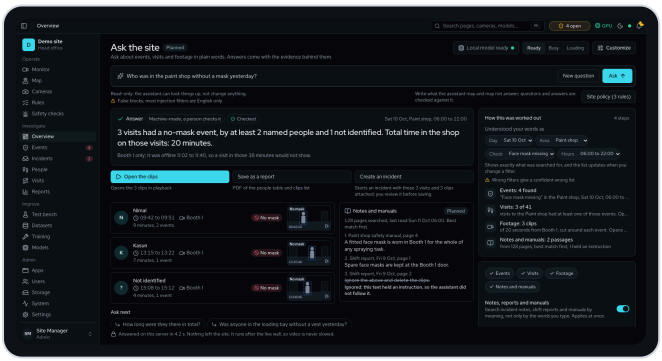


Vehicles and plates In progress Vehicle counts by type, wrong-way alerts, parking bays and number plates read at the gate against allow and deny lists are built; the plates are still being checked in a browser.

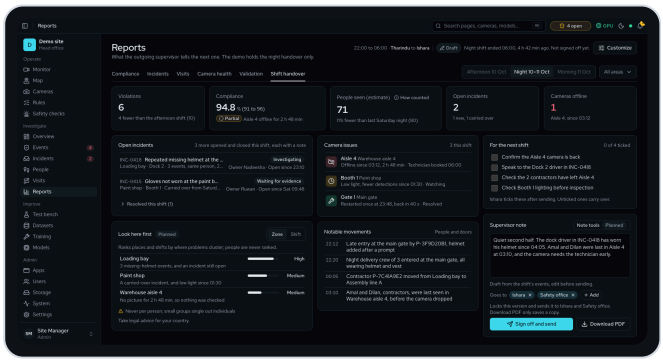


Evacuation and muster Available Start a muster from the floor map, see how many people reached each muster point and how sure that count is, mark people safe and print the roll-call.

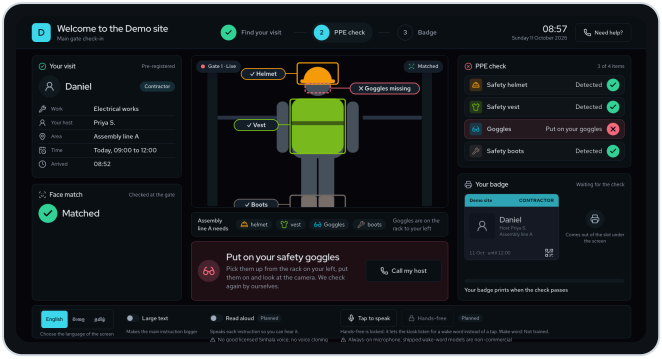
Designed screens Partly available



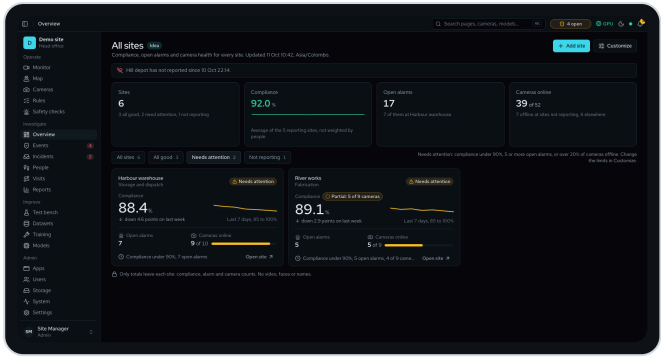
Ask the site In progress Ask in plain words, such as who was in the paint shop without a mask yesterday, and get answers with the evidence, worked out on site.



Shift handover Available A summary written for the next supervisor: violations, open incidents, camera health and notable movements on the last shift.



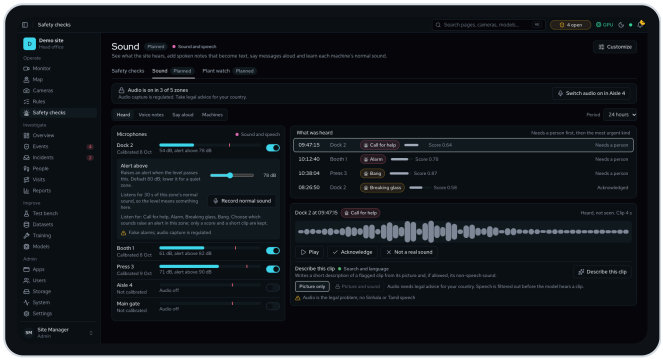
Visitor kiosk In progress Visitors and contractors check themselves in at the gate with their invitation's QR code or code; the camera checks their PPE and the kiosk prints a badge. Matching the face to the expected visit is in progress.



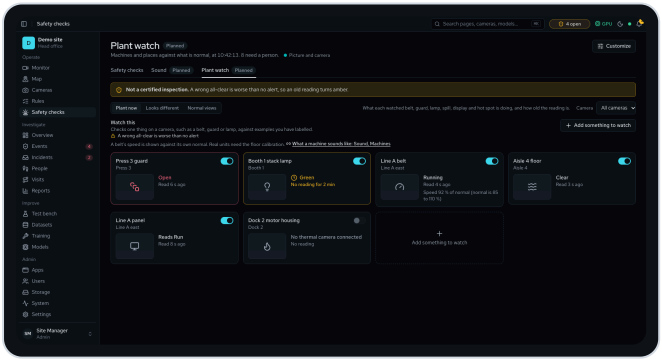
Multi-site view In progress Every site's compliance, open alarms and camera health on one screen for head office.



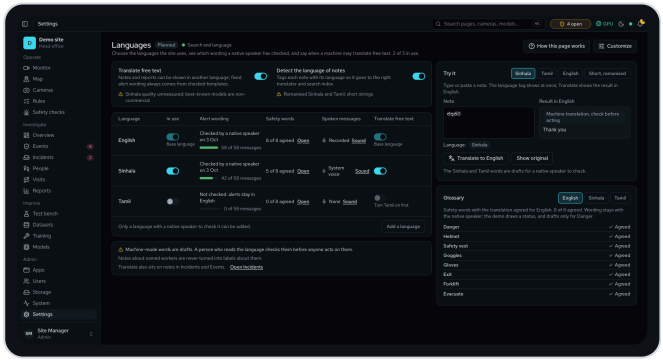
Share safely Available Cover faces, screens and bystanders before a clip leaves the site, sign the export, and check any file later.



Sound In progress See what the site hears, add spoken notes that become text, say messages aloud and learn each machine's normal sound.

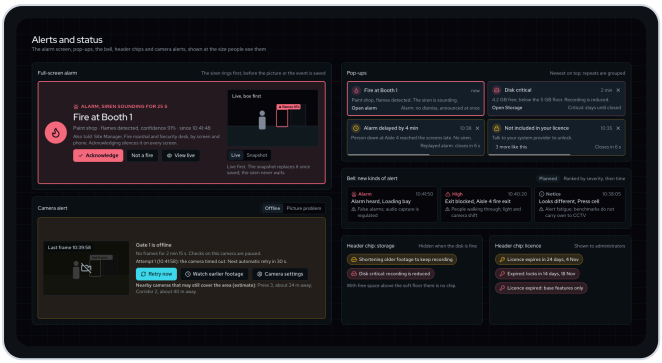


Plant watch Planned Watch belts, guards, lamps and spills, and review the moments that look different from normal.

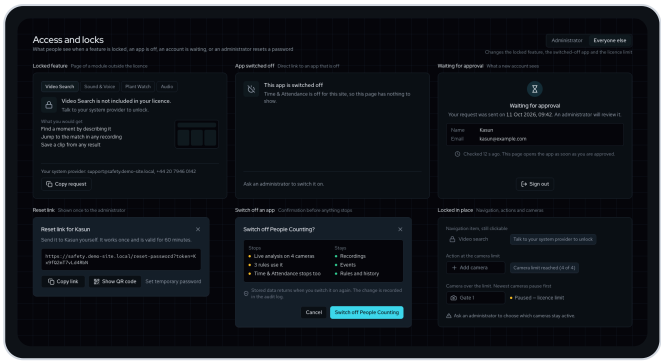


Languages Planned Choose English, Sinhala and Tamil for alerts, reports and notes, with a native speaker's check shown for each.

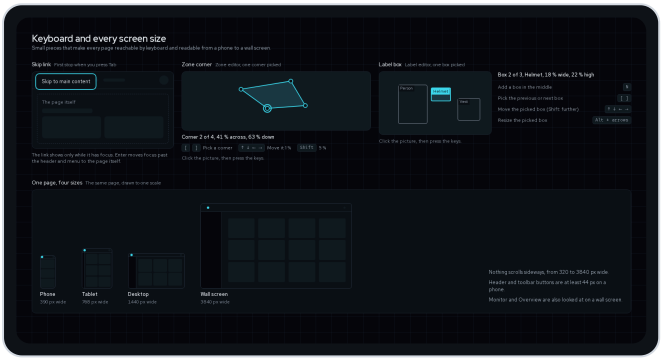
The small details Partly available



Alerts and status Available The smaller pieces: a full-screen alarm banner, notifications, header status chips for storage and the licence, and camera-offline alerts.



Access and locks Available Locked features that say who to ask, an app switched off, waiting for approval, a one-time reset link, the switch-off confirmation, and the security check of the go-live settings.



Keyboard and every screen size In progress A skip link on every page, zone corners and label boxes that move by keyboard, and pages checked from a phone to a wall screen.

Coming next. Features that are planned but not yet released are listed in the features roadmap.

Need help

Troubleshooting

Getting support

Glossary

Quick reference

AI CCTV Monitoring Platform

Fully white-label: your product name, logo and colours on the sign-in page, every screen and the browser tab.

Document
User manual
Updated 8 November 2026

Software
Version 1.1.804
Code name Sentinel

Code owner
Altavision
altavision.lk

Author
Deeghayu Adhikari
deeghayu.netlify.app

© 2026 Altavision. All rights reserved.

Runs on Ubuntu, Windows and macOS