

YUVIRED LLC

YuviRed Documentation

Manual

Installing, running and administering YuviRed, the AI Detection and Response platform.

Version 1.0 | Installer 1.0.322.1608 | Source commit 291d93d6

2026-11-19

Contents

Start here: YuviRed

- What YuviRed is
- Who it is for
- What it is not
- The three editions
- The first three things to do
- Where each guide is
- Getting help

Quick start: from download to your first report

- Before you start
- Step 1. Download and verify
- Step 2. Install
- Step 3. First-time setup
- Step 4. Activate
- Step 5. Let other computers reach the server
- Step 6. Add the first computer
- Step 7. Generate your first report
- What next

ADMINISTRATOR GUIDE

Installation guide

- Requirements
- Download and verify
- Interactive install
- Unattended install
- What the installer creates
- Upgrade in place
- Repair
- Collect support evidence
- Remove YuviRed
- The acceptance record
- Not verified in this guide

Setup guide

- **First-time setup**
- **Enroll multifactor authentication**
- **Add more administrators**
- **The Legal page, organization name and privacy contact**
- **First configuration decisions**
- **Let other computers reach the console (network mode)**
- **The browser certificate**
- **Not verified in this guide**

Licensing guide

- **How activation works**
- **Get an activation code (Lite)**
- **Activate with a code**
- **Activate with a license file (no internet on the server)**
- **Limits, and what counts**
- **Renewals and upgrades (Lite to Pro)**
- **Expiry**
- **Revocation**
- **The daily check-in, and what it sends**
- **A server moved to new hardware (server_changed)**
- **Every state of the license screen**
- **Recovery when locked**
- **What YuviRed receives from your server**
- **Not verified in this guide**

Endpoints: adding, checking, updating and removing computers

- **What an endpoint is**
- **Add a computer (Add Computer)**
- **Centrally managed deployment (group policy, Intune, other tools)**
- **Confirm an endpoint is connected**
- **What a check-in carries, and offline devices**
- **Update an agent**
- **Endpoint credential rotation**
- **Decommission and delete**
- **Remove an agent from a computer**

Troubleshooting

The YuviRed AI Prompt Monitor browser extension

What it does and does not do

How it reaches browsers in this release

Consent

Monitor Only, Warn and Block

Raw capture

Confirm the extension is connected

Update, and remove

Troubleshooting

The dashboard and the console tour

1. Before you start
2. The layout of the console
3. The Dashboard, widget by widget
4. How fresh is what you see
5. Reading the colors
6. Procedure: a five-minute morning check
7. Where to click next
8. Limits to keep in mind

AI Detection and Response (AIDR)

1. Monitoring versus enforcement, once and clearly
2. What AIDR discovers, and how
3. The AI Tools catalog and classification
4. Detections and alerts
5. Investigating
6. Policies (tuning detection)
7. Response actions that exist (Enforcement)
8. The blocklist and your own firewall (advisory)
9. Limits and honest notes

Agent Monitoring System (AMS)

1. Words used in the console
2. What is discovered, and how
3. The five AMS areas

4. Seats: 5 AI agents in Lite
5. Connectors for supported tools
6. Policy: Monitor only versus Enforce
7. Interventions
8. Messages and the Chat fallback
9. Approvals
10. Boundaries: WSL, virtual machines, containers and automated browsers
11. Alerts and findings in AMS
12. How long AMS keeps records
13. Permissions
14. Honest limits

Sensitive information in prompts (Prompt Monitor and Prompt Protection)

1. The three separate ideas
2. What is detected: rule families
3. The policy: observe, warn or block, per rule
4. What the person sees
5. Capture levels: what is kept about a finding
6. Reviewing what happened: the Prompt Monitoring page
7. Raw capture: enabling it, and the audit
8. Limits: what it covers and what it does not

Alerts, notifications, sign-in and the audit log

1. The alerts list and the alert detail
2. Notifications
3. Single sign-on (OIDC), if you want it
4. The audit history
5. Integrations that are off by default

Reports

1. Before you start
2. Make a report in four steps
3. Choices that shape a report
4. Reading a report
5. The same snapshot, a second report, and re-running
6. Who can see what

7. History, downloads and retention of report files
8. The Live overview tab
9. What a report cannot show
10. Common problems

Administering your YuviRed server

1. Updates
2. Back up (manual procedure)
3. Migration to a new server
4. Disaster recovery
5. Account recovery
6. Access management
7. Data handling
8. Privacy controls
9. Getting help

Uninstalling YuviRed

1. Remove the server (Apps and Features)
2. What is removed and what stays
3. Erase the data (the purge option)
4. Remove the agents from your other computers
5. What happens to your license

YuviRed on your work computer: a guide for the people who use it

- What runs on your computer
- What the agent reports
- What is not collected
- What the browser extension does
- Where to read the notice
- What you cannot change
- Questions and requests

REFERENCE

Console pages reference

- How to read this page
- The top bar and global controls
- First-run setup and sign-in screens

AIDR menu
AMS menu
Pages and features that do not exist
Related
Settings reference
How settings work
1. AI & Agent Monitoring
2. Prompt Protection
3. Detection & Response
4. Artifact Analysis
5. Notifications
6. Data & Retention
7. Security & Access
8. Integrations
9. System
Settings that interrupt monitoring, block activity, delete data, change access or affect licensing
Server-only and hidden settings
Related
Roles and permissions reference
Two layers: console role and system role
Permissions by role
What needs only a signed-in account
Multifactor authentication (MFA)
Passwords
Sign-in lockouts
Sessions
Related
Statuses and errors reference
1. License plan states
2. License lock messages (HTTP 402)
3. Activation messages (license screen)
4. Endpoint states

5. Browser extension and prompt protection status
6. Alert, incident and action statuses
7. First-run setup refusals
8. HTTP errors a person or an agent can meet
9. License banners, buttons and toasts
Related

Editions and limits reference
The three editions
What counts against each limit
What happens at a limit
The activation window (a new server)
Per-company limits on automatic Lite codes
At expiry
At revocation
When the server's hardware changes
Things these limits do not do
Related

Monitoring disclosures and permissions (reference for administrators)
What the agent needs
Browser and registry changes
What to give the people who are monitored
The acceptance record
Things people often miss

TROUBLESHOOTING

Troubleshooting
1. The server
2. Licensing
3. Computers (endpoints)
4. The browser extension
5. Notifications
6. Reports
7. Where the logs are

Collecting diagnostics

1. **Safe checks you can make first**
2. **Where the logs are**
3. **Run the support collector**
4. **What the ZIP holds**
5. **What is sensitive in it, and what to remove**
6. **What to write with it**

Start here: YuviRed

YuviRed shows you which AI tools the computers in your organization use, and helps you govern that use. This page tells you what the product is, what it is not, and where to go next.

What YuviRed is

YuviRed is AIDR, AI Detection and Response, for computers and browsers. It also covers AMS, AI agent Management and Security: it finds AI agents that run on your computers (for example coding assistants) and shows what they connect to.

You install one server on a Windows computer. You then add the Windows computers you want to watch. Each of them runs a small program called the YuviRed Agent, and (optionally) a browser extension for Chrome and Edge. You work in a web console that the server hosts. The console shows AI activity, alerts, reports, and the settings that decide what is recorded and for how long.

The console Dashboard of a lab install with two computers enrolled. The tiles at the top summarize what YuviRed has seen. The counts come from lab data.

Your monitoring data stays on your server. YuviRed LLC receives only what the license needs: see What YuviRed receives from your server.

Who it is for

- An IT administrator or security lead in a small or mid-size organization who must answer "who uses AI here, with what, and is it allowed?"
- Organizations that are entitled to monitor their computers and have told the people who use them. YuviRed is not for monitoring people who have not been told.

What it is not

- It does not block the network. It watches and reports; it is not a firewall. The only place it can stop anything is the browser extension, and only if an administrator later chooses a stronger Prompt Protection setting; a new installation starts in Monitor Only and nothing escalates by itself.

- It does not record keystrokes or take screenshots.
- It does not read your files or the pages people view. By default the extension keeps only a masked fragment of a prompt. Storing original prompt text is a separate setting that is off by default.
- Detection is probabilistic. YuviRed reports what it sees and how sure it is. Treat a finding as evidence to look at, not as proof.
- It does not support Firefox, macOS or Linux computers in this release. The server and the agents run on Windows. The extension supports Chrome and Edge.

The three editions

One installer serves all three. The license you activate decides the edition.

	Lite	Pro	Enterprise
Price	Free, never expires	US$1,500 per year	Custom
Computers	10, plus the server itself	50	As agreed
AI agents (AMS)	5	50	As agreed
How you get it	Request a code at https://yuvired.com/download	Request at https://yuvired.com/request; YuviRed approves it, you pay, then it is issued	Contact YuviRed
Trial or grace at the end	Not needed	No trial, no grace period	As agreed
Support	Email, reasonable efforts	Email, reasonable efforts	As agreed

Support is by email at support@yuvired.com, with no response-time promise for Lite and Pro.

One rule to know before you start: only the first account you create at setup (the owner) can activate the server, redeem a code or import a license file. Administrators you add later in the console cannot, and the console has no way to make a second owner. Create the owner account with an email address your organization keeps, and store its recovery codes.

The first three things to do

1. **Install the server.** Download `YuviRed-Lite.exe`, check its signature, and run it on a Windows computer that will stay on. Follow the Quick start, or the full Installation guide.
2. **Set up and activate.** Open the console on the server, create your administrator account, turn on multifactor authentication, and activate the server within 7 days with a code from the website. See Setup and Licensing.
3. **Add your first computer.** Use Add Computer on the Agent Deployment page, run the installer it gives you on that computer, and watch it appear under Endpoints.

Where each guide is

You want to	Open
Get running in about 10 minutes	quick-start.md
Check requirements, install silently, upgrade, uninstall	admin-guide/installation.md
Create administrators, set retention, notifications, the Legal page	admin-guide/setup.md
Activate, renew, move to Pro, recover a locked console	admin-guide/licensing.md
Read what YuviRed collects and your terms	Console, Legal page; or `https://<your server>:8443/static/legal/`
Read this documentation without internet	Console, **Help** page (an offline copy of the essential pages, served by your server at `https://<your server>:8443/static/docs/`); the full set is at https://yuvired.com/docs

The installer's last window also shows where the documentation is. The Help page is in the admin part of the navigation, so Analyst and View Only accounts do not see it.

The legal documents are the End User License Agreement, the Privacy Policy, the Monitoring and Data Notice, "What YuviRed collects and where it goes", and the Third-Party Software Notices. Read the Monitoring and Data Notice before you add the first computer.

Getting help

Email support@yuvired.com. Put "Privacy request", "Security report" or "Legal notice" in the subject when it is one of those. When you ask about a license, include the Installation ID shown on the license screen. It is not a secret. Website: https://yuvired.com.

Quick start: from download to your first report

This guide takes a new customer from the download to a first report. Plan for about 10 minutes of your own time, plus a few minutes of waiting while the installer and the first computer run.

Before you start

You need:

- A Windows computer for the server, signed in as a local administrator, that is on whenever you want monitoring. Windows Server 2022 or 2025, or Windows 10 or 11 Pro, 64-bit. Tested configuration: Windows Server 2025 Standard with 4 logical processors, 4 GB of memory and a 64 GB disk; use at least that. (YuviRed has not published a formal minimum; see the Installation guide for what the lab measured.)
- A second Windows computer (Windows 10, 11 or Server) to add as the first monitored computer.
- A work email address (a free email address is refused for a Lite download) and a company name.
- A modern browser (Chrome or Edge) on the server.
- Internet access from the server to https://activate.yuvired.com on port 443, once to activate and once a day afterwards. With no internet on the server, use the license file route in step 4.
- Permission from your organization to monitor the computers you add.

Example data in this guide: company "Northwind Dental", administrator a.rivera, server address 192.0.2.10. Use your own.

Step 1. Download and verify

1. Go to https://yuvired.com/download and enter your work email address and company name. You receive the email "Your YuviRed verification code" from YuviRed.
2. Enter the 6-digit code from that email on the page within 30 minutes. The page then gives you the download and your activation code; the email "Your YuviRed Lite activation code" repeats both. Keep the activation code: you use it in step 4. It works for 7 days and activates one server.

3. Download `YuviRed-Lite.exe` and copy it to the server.
4. Check the signature. In File Explorer, right-click the file, choose **Properties**, open the **Digital Signatures** tab, select the signature, choose **Details**. It must say "This digital signature is OK" and the signer must be **YuviRed LLC**. Or, in PowerShell, in the folder:

```
Get-AuthenticodeSignature .\YuviRed-Lite.exe | Format-List Status,
SignerCertificate
signtool verify /pa .\YuviRed-Lite.exe
```

5. Check the SHA-256 against the value shown on the download page (written here as `<published SHA-256>`; it is never in this guide):

```
(Get-FileHash .\YuviRed-Lite.exe -Algorithm SHA256).Hash -eq '<published
SHA-256>'
```

You should see: Status `Valid`, signer YuviRed LLC, and `True` from the hash check. If any of these is wrong, do not run the file: delete it and download it again.

Step 2. Install

1. Double-click `YuviRed-Lite.exe`. The window is titled **YuviRed Lite Installer** with the heading **License Agreement**.
2. Read the agreement. Tick **I have read and agree to the YuviRed License Agreement**. The **Install** button stays off until you do. Choose **Install**.
3. At the Windows prompt (User Account Control), choose **Yes**. It names the publisher YuviRed LLC.
4. The window **Installing YuviRed Lite** shows a progress bar and a checklist: Preparing installation, Installing YuviRed, Configuring database, Starting services, Securing installation, Verifying YuviRed. Wait. It takes a few minutes.
5. The window **Installation Complete** says "YuviRed Lite installed successfully", with **Service:** Running and **Status:** YuviRed Ready, and a note that starts "Documentation: https://yuvired.com/docs." Choose **Open YuviRed**.

You should see: the Installation Complete window. If instead you see **Installation Failed**, it shows a **Support code** and a **Log file** path, with a **Close** button: see Installation.

The installer's last page. The service shows Running and the status shows YuviRed Ready; the Finish and Open YuviRed buttons are at the bottom. The Windows evaluation watermark is part of the lab capture.

Step 3. First-time setup

1. Your browser opens https://127.0.0.1:8443 and warns that the connection is not private. This is expected: the server made its own certificate. Choose **Advanced**, then **Continue to 127.0.0.1 (unsafe)** (Chrome: **Proceed to 127.0.0.1 (unsafe)**). The address is this server itself, so the warning is safe to pass here.
2. The page **Welcome to YuviRed** asks you to **Create your administrator account**. Fill in **Full Name** (a.rivera), **Email**, **Password** and **Confirm Password**. The password needs at least 12 characters using at least 3 of: lowercase letters, capital letters, numbers and symbols. Check the email address before you continue: this account becomes the owner, the only account that can manage the license, and the console cannot change an account's email later.
3. Tick the box that starts "I accept the YuviRed End User License Agreement (version 1.0) and I have read the Monitoring and Data Notice" (both are links). Choose **Create Account**.
4. Because the owner role requires multifactor authentication (MFA; the screen calls it two-factor authentication), the next step asks you to scan a QR code with an authenticator app on your phone (any app that makes 6-digit time-based codes). Type the app's 6-digit code into **Code from your app** and choose **Enable & Sign In**. You have 2 minutes from choosing Create Account.
5. YuviRed shows your 10 recovery codes exactly once. Choose **Download .txt** or **Copy**, store them somewhere you can reach without your phone, tick "I have saved these recovery codes somewhere I can reach without this device. They cannot be shown again." and choose **Continue**.

You should see: the console dashboard, with an amber banner: "This YuviRed server is not activated. You have 6 days left to activate it. After that, the console is unavailable until it is activated." and the button **Activate now**.

The first-run setup form. Fill in the four fields, tick the acceptance box, then choose Create Account.

Step 4. Activate

Do this signed in as the owner account you just created; no other account can activate.

1. In the amber banner, choose **Activate now**. (Or open **Configuration**, find the **License** panel, choose **Activate or import license**.) The screen **Activate YuviRed** opens.
2. Under **Activation code**, type the code from step 1 and choose **Activate**.
3. **No internet on the server?** On a computer with internet, open the license file link in your YuviRed email (https://activate.yuvired.com/license-file), enter the activation code and the **Installation ID** from this screen, and choose **Download license file**. Back on the server choose **Upload license file** and pick it.

You should see: the toast **License activated**, and the License panel showing **Active**, **YuviRed Lite**, **Expires** Never and "YuviRed Lite is free and never expires." If a message appears instead, see Licensing.

The License panel after activation. Notice the plan, the Active state, and how many of your 10 computers and 5 AI agents are in use. Serial and installation ID are covered.

Step 5. Let other computers reach the server

A new server starts in loopback mode: it listens on 127.0.0.1:8443 only, so only a browser on the server itself reaches the console, and no other computer can enroll. Before you add a computer, switch network mode on.

1. On the server, open PowerShell as administrator and run (use your server's DNS name or fixed IP address; a DNS name survives an address change):

   ```
   $py = "C:\Program Files\YuviRed\runtime\python.exe"; $yr = "C:\Program Files\YuviRed\runtime\yuvired_service.py"
   & $py $yr network --name 192.0.2.10
   ```

2. Keep the window open. The command writes `config\network.conf`, adds the inbound Windows Firewall rule **YuviRed Server (TCP 8443)** for the Domain and Private profiles only, and restarts YuviRed: it prints "Restarting YuviRed. Keep this window open: this usually takes a minute or two and can take up to 7 minutes." The server's own agent is skipped in network mode.

3. Read the output. If it says "This server's network is set to Public, so Windows Firewall blocks port 8443 and no other computer can connect. YuviRed does not open the Public profile. ...", follow the instructions it prints (set only the trusted network to Private with `Set-NetConnectionProfile`), because computers cannot connect until you do.
4. Check it: `& $py $yr network --status` shows the bind, the firewall profile and rule, the server state and a line "certificate SHA-256: ..." (`& $py $yr network --off` returns to loopback and removes the rule.)

You should see: the command ends with exit code 0 and `--status` shows the server bound to the network with the rule present. From now on the console address is https://192.0.2.10:8443 (or https://<name>:8443), and other computers' browsers can reach it.

Step 6. Add the first computer

1. Reload the console, at https://192.0.2.10:8443 if you like (the browser warns about the certificate once more for the new address). Open **Agent Deployment**. The **Add Computer** panel shows the server address and "Computers: 0 of 10". Leave "Computers this download may enroll" at 1 and choose **Download installer**. Your browser saves `YuviRed-Add-Computer.zip`.
2. Copy the zip to the other computer. Treat it as a secret: it works for 24 hours and is a key to your server. Open `README.txt` inside it and follow its four steps in PowerShell as administrator (make the protected folder `C:\ProgramData\YuviRed-Install`, unzip into it, run `Install-YuviRed-Agent.cmd`). Do not run the MSI by itself and do not run from Downloads.
3. Wait until the installer finishes (up to 90 seconds for the server's answer).

You should see: "Configured. This computer is enrolled and reports to https://..." and "Done. Check that this computer is listed under Endpoints on the server." In the console, open **Endpoints**: the computer is listed within a few minutes, and Add Computer shows "Computers: 1 of 10". When you are done, delete the zip everywhere.

The Add Computer page. The Download installer button is at the top; check the computer count box before you download.

The Endpoints page with enrolled computers online. Both are lab computers.

Step 7. Generate your first report

1. Open **Reports**. On the **Reports** tab choose **Executive Summary** (about 2 to 4 pages).
2. Pick a period that covers today, then choose **Generate Executive Summary (PDF)**.
3. The report appears under **Report history**, with its number, period and files. Download it.

You should see: a report in Report history with status Done. A brand-new computer has little data yet: the report improves as the computer runs for a few hours or days. The **Live overview** tab shows the current picture at any time.

The report history. A finished report shows status Done, with buttons to download it as a PDF.

What next

- Finish setup: more administrators, retention and notifications, the organization name on the employee notice: Setup.
- Tell the people on the monitored computers. The notice is on the console's **Legal** page.
- Know your license rules and what to do before day 7: Licensing.
- Without internet, the console's **Help** page carries an offline copy of the essential pages of this documentation.

Installation guide

This guide covers everything about putting the YuviRed server on a Windows computer: requirements, download and verification, the interactive and unattended installs, what the installer creates, upgrade, repair and removal, and the record of your acceptance of the legal terms. For the shortest path, use the Quick start.

Product and versions documented: YuviRed 1.0, installer 1.0.322.1608, app 1.0.0, endpoint agent 1.8.2, browser extension 0.5.5.

Requirements

Item	Requirement
Server operating system	Windows Server 2022 or 2025, or Windows 10 (version 1809 or later) or 11 Pro; 64-bit. The installer itself refuses anything else with "YuviRed Lite requires 64-bit Windows 10 (1809) or later." It needs .NET Framework 4.7.2 or later, which every supported Windows ships
Monitored computers	Windows 10 and 11, Windows Server
Browsers (extension)	Chrome and Edge. Firefox is not supported
Permissions	You must be a local administrator on the server. The installer asks once for elevation
CPU, memory, disk	Tested configuration (the lab server running this release): Windows Server 2025 Standard, 4 logical processors, 4 GB of memory, 64 GB disk. Recommend at least that. After a fresh install with one enrolled computer, the YuviRed service and its bundled database used about 490 MB of memory, the program files about 251 MB and the data folder about 71 MB; Windows plus YuviRed used about 10 GB of the disk. The data folder grows with the number of computers and the retention settings. YuviRed has not published a formal minimum
Server stays on	Computers report to it continuously. Run it on a computer that is on and not asleep
Not a domain controller	The `network` command warns when the computer is a domain controller and refuses unless you pass `--allow-domain-controller`. Use a dedicated server

Network:

Direction	Port	Purpose
Inbound to the server, from your monitored computers and from administrators' browsers	TCP 8443	The console and the endpoint check-in. Not open until you turn on network mode
Outbound from the server	TCP 443 to https://activate.yuvired.com	Activation (once) and the daily check-in. No internet: use a license file
On the server only (loopback)	TCP 5440	The bundled PostgreSQL database. It listens on 127.0.0.1 only and is never reachable from the network

Network mode. A new server starts in loopback mode: it binds 127.0.0.1:8443, only a browser on the server itself reaches the console, and no other computer can enroll. To let others in, run on the server, in PowerShell as administrator:

```
$py = "C:\Program Files\YuviRed\runtime\python.exe"; $yr = "C:\Program
Files\YuviRed\runtime\yuvired_service.py"
& $py $yr network --name server.northwind.example     # a DNS name, or a fixed
IP such as 192.0.2.10
& $py $yr network --status                            # shows the bind,
firewall profile and rule
& $py $yr network --off                               # back to loopback, rule
removed
```

`network --name` writes `config\network.conf`, adds an inbound Windows Firewall rule named **YuviRed Server (TCP 8443)** for the Domain and Private profiles only, and restarts the service. Keep the window open: it prints "Restarting YuviRed. Keep this window open: this usually takes a minute or two and can take up to 7 minutes." and ends with exit code 0 when the server is ready again. The server's own agent is skipped in network mode. If the server's network is classed Public, the command tells you so: "This server's network is set to Public, so Windows Firewall blocks port 8443 and no other computer can connect. YuviRed does not open the Public profile. ..." and prints the `Get-NetConnectionProfile` and `Set-NetConnectionProfile ... -NetworkCategory Private` steps to move only the trusted network

to Private. The rule allows any remote address: to restrict which computers may connect, use a firewall outside the server. Prefer a DNS name: agents trust this server's certificate, and if you enroll by IP address and the address later changes, you must run `network` with the new address and make a fresh Add Computer download on each computer. After the command, the console address is https://<name>:8443. `network --status` also prints "certificate SHA-256: ..." (the value the Add Computer panel shows, for comparing with what a browser presents) and warns when the running server applies another setting than the saved one or when the certificate does not yet carry the configured names.

First-time account creation is accepted only from a browser on the server itself, even in network mode.

Download and verify

Purpose: be sure the file is the one YuviRed LLC published.

1. On https://yuvired.com/download enter your work email address and company name and ask for the verification code. A free email address is refused with the message "Use your work email address. Free email addresses cannot download YuviRed Lite. If your business has no work email address, contact us."
2. The email "Your YuviRed verification code" has a 6-digit code. Enter it on the page within 30 minutes. A wrong or old code says "That code is not right, or it has expired. Ask for a new one."
3. The page then gives you the download and your activation code. The email "Your YuviRed Lite activation code" repeats both: "The code works for 7 days, until <date> UTC, and activates one server." and "This link works until <date> UTC." for the download.
4. Limits per company for automatic Lite codes: 2 servers, 2 verified people, 3 codes per day. A third person or server, or a fourth code in a day, shows a message that ends "the YuviRed team will reply by email", with a form where you tell YuviRed why you need more.

5. Verify the signature. Windows Explorer: right-click `YuviRed-Lite.exe`, **Properties**, **Digital Signatures** tab, select the signature, **Details**: "This digital signature is OK", signer **YuviRed LLC**. PowerShell:

```
Get-AuthenticodeSignature .\YuviRed-Lite.exe | Format-List Status,
StatusMessage, SignerCertificate
signtool verify /pa /v .\YuviRed-Lite.exe        # needs the Windows SDK;
exit code 0 means verified
```

Status must be `Valid` and the signer subject must name YuviRed LLC.

6. Verify the SHA-256. The value is shown on the download page; this guide writes it as `<published SHA-256>` and never contains it:

```
(Get-FileHash .\YuviRed-Lite.exe -Algorithm SHA256).Hash -eq '<published
SHA-256>'
```

It must print `True`.

How to confirm: both checks pass. Common problems: the status is `NotSigned` or `HashMismatch`, or the hash differs: delete the file and download again; if it repeats, email support@yuvired.com. A SmartScreen warning that names an "Unknown publisher" means the file is not signed: do not run it. Rollback: delete the file.

Interactive install

Purpose: install the server with the installer window. Prerequisites: the verified file, local administrator rights, no other YuviRed installation of a newer version on this computer.

1. Double-click `YuviRed-Lite.exe`. The red window **YuviRed Lite Installer** opens on **License Agreement**, with the agreement text in a scrolling box.
2. Read it, then tick **I have read and agree to the YuviRed License Agreement**. **Install** is enabled only by that checkbox. **Cancel** closes the installer without changing anything.
3. Choose **Install**. At the Windows prompt (User Account Control) choose **Yes**. The prompt names the publisher; for a signed build it is YuviRed LLC.
4. **Installing YuviRed Lite** shows a progress bar and the live checklist: Preparing installation, Installing YuviRed, Configuring database, Starting services, Securing installation, Verifying YuviRed. Each step is ticked only when it really happens. Do not close the window; closing is refused while the install runs. Windows may first install the Microsoft Visual C++ runtime if a recent one is missing; that can ask for a restart.

5. **Installation Complete** appears only after YuviRed has proven itself ready twice: the installer package waits until the server is ready and stays ready for 15 seconds, and then the window checks both services, the database and web ports, and the health check. It shows a green check, "YuviRed Lite installed successfully", **Service:** Running and **Status:** YuviRed Ready, the note "Documentation: https://yuvired.com/docs. On this server after setup: https://localhost:8443/static/docs/ (Help in YuviRed). Legal documents: the Legal page in YuviRed.", and the buttons **Finish** and **Open YuviRed**. If the Visual C++ runtime needed a restart, the page adds "Windows needs to restart to finish installing a Windows component; YuviRed Lite will also be available after the restart."
6. Choose **Open YuviRed**. Your browser opens https://127.0.0.1:8443/. See Setup.

How to confirm: in PowerShell, `Get-Service YuviRed, YuviRedPostgres, YuviRedAgent` shows Running; the console loads.

Common problems:

- **Installation Failed**, "The installation could not be completed. Nothing from this attempt is left running; you can run the installer again." The page shows a **Support code** and a **Log file** path and a **Close** button. Run the installer again. If it fails again, run the support collector (below) before you reboot or uninstall, and email the ZIP to support@yuvired.com.
- "YuviRed was installed but is not answering on this computer (...). The services and their logs were kept for diagnosis.": the package finished but the window's own final check failed. Collect support evidence and email it.
- "A newer version of YuviRed Lite is already installed on this computer. To install this version, remove the newer one first."
- "An earlier YuviRed installation that this installer cannot upgrade was found on this computer. Remove it first, then run this installer again." (services named like YuviRed's exist that no installation record explains).
- "The license agreement could not be loaded from this installer package, so it cannot be accepted. Please download YuviRed Lite again.": the downloaded file is damaged; verify and download again.
- The browser warns that the connection is not private: expected, see Setup.

Rollback: a failed install is rolled back by the installer. To remove a finished install, see "Remove YuviRed".

The first installer page. Install stays greyed out until you tick the agreement box. Captured in the YuviRed lab; the Windows evaluation watermark is not part of your screen.

The progress page. The installer works through the checklist; wait until it finishes. Captured in the YuviRed lab; the Windows evaluation watermark is not part of your screen.

The completion page. Service Running and Status YuviRed Ready tell you the install worked. Captured in the YuviRed lab; the Windows evaluation watermark is not part of your screen.

Unattended install

Purpose: install without a window, for a deployment tool or a script. An unattended install records that an authorized administrator accepts the YuviRed End User License Agreement (version 1.0) and the Monitoring and Data Notice, so you must say so on the command line. Run it in PowerShell as administrator:

```
.\YuviRed-Lite.exe /quiet ACCEPT_EULA=1 /log C:\Temp\yuvired-install.log
$LASTEXITCODE
```

`/passive ACCEPT_EULA=1` is the same with a progress window and no questions.

Exit code	Meaning
0	Installed
3010	Installed, a Windows restart is needed to finish (a standard Windows code, usually from the Visual C++ runtime)
5101	Refused: `ACCEPT_EULA=1` was missing. Nothing was changed and no service was installed. The log says "YuviRed setup: unattended installation requires ACCEPT_EULA=1, which records that an authorized administrator accepts the YuviRed End User License Agreement (eula 1.0) and the Monitoring and Data Notice"
1602	Cancelled before it finished
other	A standard Windows Installer or installer-engine code; send the log to support

The parameter name is matched without regard to case; the value must be exactly `1`. Upgrade behavior: an unattended upgrade of a computer that already has the acceptance record (below) does not need `ACCEPT_EULA=1`. Removal (`/uninstall /quiet`) never needs it. Without `/log`, the installer writes its log to the `%TEMP%` folder of the account that ran it, in a file named like `YuviRed_Lite_*.log`. The account recorded in the acceptance record is the one that ran the installer (for a deployment tool running as the system account it reads `<computer or WORKGROUP>\SYSTEM`).

How to confirm: exit code 0, the three services Running, and `legal-acceptance.json` exists with `"mode": "silent"` (lab-verified on this release: a run without the parameter exited 5101 with no service and no record; with it, exit 0, services Running, record written with mode silent and installer version 1.0.322.1608).

Common problems: exit 5101 (add the parameter); running without administrator rights (the install fails at elevation: start PowerShell as administrator).

What the installer creates

Folders:

Location	Holds
`C:\Program Files\YuviRed\`	The program: `runtime\` (embedded Python), `pg\` (the bundled PostgreSQL 16), `versions\<version>\` (the application), `trust\` (public keys and addresses, no secrets), `support\` (the support collector), `browser\` (the browser native-messaging host), `agent-code\` (the server's own agent, copied at each start)
`C:\ProgramData\YuviRed\`	Your data, created once and kept across upgrades and removals: `identity\`, `config\` (settings, the TLS certificate and key, generated secrets), `pg\` (the database), `logs\`, `status\`, `agent\` (the server's own agent)
`C:\ProgramData\SyntinelRed\logs\`	The log of the server's own agent, `syntinel-agent.log` (the product's earlier name; it will change in a later release)

Services:

Service	Runs as	Notes
YuviRedPostgres	NT SERVICE\YuviRedPostgres, automatic	The database, on 127.0.0.1 port 5440
YuviRed	NT SERVICE\YuviRed, automatic	The console and the API, on port 8443 (loopback until network mode)
YuviRedAgent	LocalSystem, automatic (delayed)	YuviRed Endpoint Agent for the server itself; skipped in network mode

Other items: the Microsoft Visual C++ 2015-2022 x64 runtime if missing (kept at removal); the bundle's copy under `C:\ProgramData\Package Cache` (Settings, Apps uses it to uninstall); registry keys for the browser native-messaging host under Chrome, Edge and Mozilla; the Application event log sources `YuviRed` and `YuviRedAgent`; and, once you run `network --name`, the firewall rule **YuviRed Server (TCP 8443)**.

The certificate. First start creates a self-signed certificate (ECDSA P-256, valid 825 days) for the computer's name, `localhost` and 127.0.0.1, plus the names you give to `network`. Its key stays in `C:\ProgramData\YuviRed\config` and is protected by permissions. It is not installed into any Windows trust store, which is why browsers warn once. Agents trust exactly this certificate through the enrollment file they receive.

There is no administrator password to look up: you create the account at first run. Four secrets (a signing key, an encryption key, a token pepper and the database password) are generated once and never shown.

Upgrade in place

Purpose: move to a newer version, keeping your data. Prerequisites: the new `YuviRed-Lite.exe`, verified as above. YuviRed never installs an update by itself: after the daily check-in, a banner says "YuviRed <version> is available. This server runs <current version>. Nothing is installed automatically." with **Download** and **Dismiss**.

1. Run the new `YuviRed-Lite.exe` over the existing install (interactive, or `/quiet`; an unattended upgrade of a computer that holds the acceptance record needs no `ACCEPT_EULA=1`).
2. The installer detects the existing installation before anything changes. An older or same-version installation is upgraded in place: the previous version is removed inside the new install and `C:\ProgramData\YuviRed` is kept (accounts, computers, policy, network settings, certificate, installation ID and license). The interactive window shows the same

License Agreement page first; the checklist's "Installing YuviRed" step covers the replacement.
3. A newer installation stops the installer on the first screen with "A newer version of YuviRed Lite is already installed on this computer. To install this version, remove the newer one first."

How to confirm: the License panel and the console footer show the new version; Endpoints still lists your computers; the Legal page still shows the original installer acceptance (an upgrade never rewrites it). If the upgrade cannot finish, the installer shows "The upgrade could not be completed. The previous YuviRed installation and your data were left in place; you can run the installer again."

Not verified in this documentation pass: an in-place upgrade between two released versions on a customer computer (the detection and the wording are in the installer; the lab ran upgrades between pre-release builds, and this release is the first production installer). Take a manual copy of `C:\ProgramData\YuviRed` (stop the YuviRed services first) before any upgrade, because Lite ships no backup command.

Repair

Settings, Apps offers no Repair for YuviRed: the package disables the Windows Installer repair and modify entries. To fix a damaged installation, run the same or a newer `YuviRed-Lite.exe` again; the installer treats a same-version installation as an upgrade in place and keeps your data. This is not lab-verified; if it does not help, collect evidence and email support@yuvired.com. (Not verified: whether Settings, Apps shows a "Modify" or "Change" entry for the YuviRed Lite package itself; if it does, use the installer file instead.)

Collect support evidence

If an install fails, do not reboot, uninstall or clean up first. Run, as administrator, `C:\Program Files\YuviRed\support\collect-install-evidence.ps1` (or the copy beside the EXE). It makes one redacted ZIP on your Desktop. Look through it, then email it to support@yuvired.com.

Remove YuviRed

Open **Settings**, **Apps**, choose **YuviRed Lite**, **Uninstall**; or run `YuviRed-Lite.exe /uninstall /quiet` (no window, no questions). An interactive removal shows the window **Removing YuviRed Lite** with the steps Preparing removal, Stopping services, Removing YuviRed, and

ends on **Removal Complete**, "YuviRed Lite was removed from this computer", with "Your data, settings and logs were kept in C:\ProgramData\YuviRed. Installing YuviRed Lite again on this computer uses them." and "Computers you added keep their agent. Remove it on each computer with Uninstall-YuviRed-Agent.cmd, or from Settings, Apps." The removal does not delete your data and has no option to. Deleting `C:\ProgramData\YuviRed` yourself makes the next install a new server with a new Installation ID that must be activated again. The full procedure, including the agents, is in Uninstall.

The acceptance record

Every install records who accepted the legal terms. File: `C:\ProgramData\YuviRed\status\legal-acceptance.json`. Its fields: `document_id` ("eula"), `document_version` ("1.0"), `accepted_at` (UTC), `mode` ("interactive" or "silent"), `account` (the Windows account that ran the installer), `authority` ("deploying administrator"), `installer_version` and `recorded_by` ("installer"), plus a `history` list. A record of the same document version is never overwritten: an upgrade or a repeated install keeps the original acceptance. If a later release ships a new version of the agreement, the new acceptance becomes the current record and the earlier one moves into `history`. Example from the lab (silent install of this release): mode silent, account `WORKGROUP\SYSTEM`, installer version 1.0.322.1608.

To see it in the console: open **Legal**, panel **Acceptance records**: "Installer acceptance: version 1.0, accepted <time> by <account> (silent)" or "(interactive)", and "Console acceptance: version 1.0, accepted <time> by <email>". The console line is written when the owner accepts at setup; it also appears in the audit log as `legal.accept`.

Not verified in this guide

- A formal minimum for CPU, memory and disk: YuviRed has published none. The figures above are the one tested configuration and what it measured; a larger fleet or longer retention needs more disk.
- Repair behavior, whether Settings, Apps shows a Change entry, and an upgrade between two released versions on a customer computer.
- The exact wording of the yuvired.com download page (this guide follows the emails and service messages in the code).
- Exit code 3010 from an unattended install: expected from the Windows Installer engine when the Visual C++ runtime schedules a restart, not reproduced in the lab.

Setup guide

This guide covers the first sign-in and the decisions you make right after: the administrator account, MFA (multifactor authentication), more administrators, the Legal page, retention, notifications, and the browser certificate. Install first: Installation.

First-time setup

Purpose: create the one owner account for this server. Prerequisites: the installer finished; you are on the server itself (the form is accepted only from a browser on the server, even later in network mode).

1. Choose **Open YuviRed** on the installer's last window, or open https://127.0.0.1:8443 on the server. Visiting `/` or `/login` before an account exists sends you to `/setup`.
2. The browser warns that the connection is not private, because the server made its own certificate. Choose **Advanced**, then **Continue to 127.0.0.1 (unsafe)** (Chrome: **Proceed to 127.0.0.1 (unsafe)**). You can pass it here because 127.0.0.1 is the server you are sitting at.
3. **Welcome to YuviRed**, "Create your administrator account". Fill in **Full Name** (a.rivera), **Email** (a.rivera@northwind.example), **Password** and **Confirm Password**. The hint under the form ends "This account becomes the owner of this YuviRed installation." Check the email address: the console has no field to change an account's email later, and this one account is the only one that can ever manage the license (see "Add more administrators").
4. Tick the checkbox "I accept the YuviRed End User License Agreement (version 1.0) and I have read the Monitoring and Data Notice". Both names are links that open the documents in a new tab.
5. Choose **Create Account**.

What you should see: the MFA step (below), then the dashboard. The account is the owner. The setup page cannot be used again: once an owner exists `/setup` redirects to `/login`, also after a reinstall over the preserved data folder.

Password rules: at least 12 characters (at most 72 bytes), at least 3 of 4 kinds (lowercase, capital letters, numbers, symbols), at least 6 different characters, no common words such as "password" or "welcome" (also in disguised forms like P@ssw0rd), and not your email name or a 4-letter part of your name. Later changes may not reuse your recent passwords

(Configuration, Security & Access, "Prevent password reuse", default last 5). Passwords never expire on a schedule.

Common problems and what the screen says:

Message	Fix
Please enter your full name. / Please enter your email address.	Fill the field
Passwords do not match.	Retype both
Please accept the End User License Agreement and confirm you have read the Monitoring and Data Notice.	Tick the checkbox
Password must be at least 12 characters (or another rule)	Choose a stronger password
Setup is only available from the computer YuviRed is installed on	Open the form in a browser on the server itself
Setup has already been completed	An owner exists; sign in at /login
Cannot reach YuviRed. If it was just installed, give it a moment and try again.	Wait a minute and retry

Rollback: none. The owner account cannot be replaced through this form again, its email cannot be changed in the console, and the administrators you add later cannot take over the license. If you typed a wrong email, email support@yuvired.com before you activate.

The Welcome form that creates the first administrator account.

Enroll multifactor authentication

The owner role requires MFA (the screen calls it two-factor authentication: "Your role requires two-factor authentication. Scan this QR with an authenticator app ..."), so the first sign-in goes straight to it. You have 2 minutes from choosing **Create Account** (or **Sign In**); if it ends, sign in again to get a new QR code and delete the old YuviRed entry from your authenticator app first.

1. On your phone, open an authenticator app (any app that makes 6-digit time-based codes).

2. Scan the QR code on the screen, or type the **Secret** shown under it.
3. Type the app's 6-digit code into **Code from your app** and choose **Enable & Sign In**.
4. YuviRed shows your 10 recovery codes (format XXXXX-XXXXX) exactly once: "Save these recovery codes now: they are shown exactly once." Choose **Download .txt** or **Copy**, store them away from your phone, tick "I have saved these recovery codes somewhere I can reach without this device. They cannot be shown again.", and choose **Continue**. Each code works one time.

What you should see: the dashboard. At every later sign-in you enter your email and password (**Sign In**), then the **Verification code** (the app's 6-digit code, or one recovery code), then **Verify**. That step also has a 2-minute window.

Recovery (what exists in this release):

- Lost phone, still have a recovery code: sign in, enter the recovery code at **Verification code**. Then open your account menu, **Recovery Codes**: **Regenerate recovery codes** gives a new set (the old set stops working), and **Disable MFA** turns MFA off for your account. Both ask for a current code from your authenticator app, not a recovery code, so with the phone gone you can only keep signing in with the remaining recovery codes until a second administrator resets you (next point). The page shows how many codes are left.
- Lost phone and no codes: another administrator (a user with the Admin role) opens **Team / Users** and uses the shield button on your row ("Disable MFA for a lost phone with no recovery codes"); it is audited, and you enroll a new phone at your next sign-in.
- The only administrator, phone and codes both gone: the console has no self-service reset and nothing in the product restores access. Email support@yuvired.com from the owner's address; expect this to take time. Add a second administrator on day one and keep the codes safe so you never need this.

Add more administrators

Purpose: avoid a single point of failure. Prerequisite: signed in as an administrator.

1. Open **Team / Users** and choose **Add User**.
2. Fill **Full Name**, **Email**, **Password** (the same rules), **Role** (Analyst, View Only or Admin) and **Department**, then choose **Create User**.
3. Give the person the password through a safe channel and ask them to sign in. Anyone with the Admin role must enroll MFA at first sign-in (Configuration, Security & Access, "Require multi-factor sign-in for"; viewer accounts cannot be required).

What you should see: the toast "User created" and the person in the list.

What an added Admin can and cannot do. An Admin created here holds every permission except license management: activating, redeeming a code, importing a license file and downloading a license request belong to the owner account created at setup alone. Another user who tries sees "Only a YuviRed administrator can activate this server. Sign in with the administrator account created when YuviRed was installed." The console has no control to make a second owner or to move the owner role, and no field to change an account's email. Plan for it: keep the owner's credentials and recovery codes where your organization can reach them, and use the added Admin accounts for daily work.

The Legal page, organization name and privacy contact

The **Legal** page (under the admin navigation) has three panels:

- **Legal documents**: the End User License Agreement, Privacy Policy, Monitoring and Data Notice, "What YuviRed collects and where it goes", and Third-Party Software Notices, each with version and effective date. They are also served at https://<server>:8443/static/legal/.
- **Acceptance records**: the installer's and the console's acceptance (see Installation).
- **Employee notice preview**: the Monitoring and Data Notice with your organization name and privacy contact filled in, and **Copy link** to give people the notice address.

To set the names: open **Configuration**, section **System**, and fill **Organization Name** ("Used in notification subjects and reports", for example Northwind Dental) and **Privacy contact** ("Shown to the people your endpoints monitor, on the notice page", for example privacy@northwind.example). Save, then reload the Legal page: the preview now shows them. If you leave the fields empty the notice says "your organization" and "your organization's administrator". Tell the people whose computers you monitor where the notice is before you add their computers.

First configuration decisions

Open **Configuration**. Sections you should look at on day one:

Data & Retention (how long data is kept). **Automatic Data Cleanup** runs a daily cleanup. Fields: **Raw telemetry** (days; removed only after being summarized), **Daily summaries** (months), **Closed alerts** (days, 0 = keep forever), **Audit log** (days, 0 = keep forever; a window permanently deletes older audit records), **AI agent signals**, **AI prompt/output**

events, **Decided adoption requests**, and **Original prompt samples** (4 hours, 24 hours recommended, or 3 days; applies only if you turn on raw capture, which is off by default). Choose what your policy and the law where you operate allow, and write it in your own notice.

Notifications (who hears about findings). **Notify me about** (Critical only, High and above recommended, Medium and above, Everything); **Email alerts** with **Send alerts to**, **Mail server**, **Mail server port**, **Mail username**, **Mail password**, **Encrypt mail connection (TLS)** and **Send from address**; and **Slack / Microsoft Teams alerts**. Nothing is sent until you set a destination.

Prompt Protection. A new install starts in **Monitor Only (first-install default)**: detections are recorded, but no warning popup, confirmation or block appears. Nothing escalates by itself. Change it only deliberately, after you have read what users will experience.

Security & Access (MFA roles, password reuse) and **System** (the two names above).

What you should see after Save: the status line at the top of Configuration confirms the change.

Let other computers reach the console (network mode)

Before **Add Computer** can work, and before anyone else can open the console, turn on network mode. The server starts in loopback mode: it binds 127.0.0.1:8443, only a browser on the server itself reaches the console, and no other computer can enroll. In PowerShell as administrator on the server:

```
$py = "C:\Program Files\YuviRed\runtime\python.exe"; $yr = "C:\Program
Files\YuviRed\runtime\yuvired_service.py"
& $py $yr network --name server.northwind.example     # the server's DNS name
or fixed IP, e.g. 192.0.2.10
& $py $yr network --status                            # shows the bind and the
firewall rule
& $py $yr network --off                               # return to loopback
```

It writes `config\network.conf`, adds the inbound rule **YuviRed Server (TCP 8443)** for the Domain and Private firewall profiles only, restarts the service ("Restarting YuviRed. Keep this window open: this usually takes a minute or two and can take up to 7 minutes."), and the server's own agent is skipped. A server whose network is classed Public gets the warning "This server's network is set to Public, so Windows Firewall blocks port 8443 and no other

computer can connect. YuviRed does not open the Public profile. ..." with the `Set-NetConnectionProfile` steps to move only the trusted network to Private. What you should see: exit code 0, `--status` showing the rule present and the line "certificate SHA-256: ...", and the console reachable at https://<name>:8443 from another computer (with the same one-time certificate warning). Until it is on, the Add Computer panel says "Network mode is off, so other computers cannot reach this server." and prints this command. If the server's name or address changes later, rerun `network --name` and make a fresh Add Computer download on each computer; a DNS name makes this unnecessary.

Then add computers: **Agent Deployment**, **Add Computer**, **Download installer**; see the Quick start step 6.

The browser certificate

The server uses a self-signed certificate, so every browser that opens the console shows a warning the first time. The warning is expected. Pass it only for an address you know is your server. In the Add Computer panel the **Certificate SHA-256** is shown: compare it with the certificate the browser shows, or with `network --status` on the server; if they differ, stop.

Replacing it with a certificate from your own authority: the product has no console setting, installer option or documented procedure for this in this release, so this guide gives none. The certificate files are `tls-cert.pem` and `tls-key.pem` in `C:\ProgramData\YuviRed\config`, and agents trust exactly the certificate they were enrolled with. Replacing it strands every enrolled agent until each gets a fresh Add Computer download, and the browser extension on the server pins it too. Do not replace it by hand without asking support@yuvired.com. (Not verified: whether a hand-placed certificate that covers the configured names is kept at the next start.)

Not verified in this guide

- Whether a hand-placed certificate survives a restart (see above).
- The exact status line wording after saving Configuration.
- The recovery paths above were read from code, not run in the lab: the self-service **Disable MFA** and **Regenerate recovery codes** prompts ask for an authenticator code, and the Team / Users reset needs a second Admin.
- The exact names of the Configuration status line after Save and of the toast after Add User ("User created") were read from code.

Licensing guide

This guide explains how a YuviRed server is licensed: the 7-day activation window, activation codes, the license file route, the limits and what counts against them, renewals and upgrades, expiry, revocation, the daily check-in, a server moved to new hardware, every state of the license screen, and how to recover when the console is locked.

Editions: Lite is free and never expires (10 computers plus the server itself, and 5 AMS agents; AMS is Agent Monitoring System). Pro is US$1,500 per year (50 computers, 50 AMS agents), with no trial and no grace period at expiry. Enterprise is custom. One installer serves all editions; the license decides which one runs.

How activation works

A new server works for 7 days with Lite limits so you can try it. Within those 7 days you activate it once. After day 7 an unactivated console shows only the license screen until you activate. Computers already enrolled keep reporting, and the browser extension only monitors (it never blocks) until the server is activated. New computers and new AI agents are refused while the server is locked.

While the 7 days run, every signed-in user sees an amber banner: "This YuviRed server is not activated. You have 6 days left to activate it. After that, the console is unavailable until it is activated." with **Activate now**. The count is whole days left, rounded down, so it reads 6 right after install, then "1 day left", then "less than 1 day left". In **Configuration**, the **License** panel shows **Not activated** and the days left.

Only the owner account created at first-time setup can activate, redeem a code or import a license file; an Admin added later in the console cannot (see Setup).

A server activated once never gets the 7 days again. A server that never held a license (for example one upgraded from a pilot build) gets its 7 days from the first start of this version; get a code the same way as a new server.

Get an activation code (Lite)

1. On https://yuvired.com/download enter your work email address and company name, ask for the verification code, and enter the 6-digit code from the email within 30 minutes.

2. The page gives you the activation code, and the email "Your YuviRed Lite activation code" repeats it: "The code works for 7 days, until <date> UTC, and activates one server." It is valid once, for one server. Ask for a new one if it expires.

Limits per company for automatic Lite codes: 2 servers, 2 verified people, and 3 codes per day. Beyond that you see one of: "Two people at your company already use YuviRed Lite. To add another person, tell us why you need it and the YuviRed team will reply by email."; "Your company already has 2 YuviRed Lite servers. To add another server, contact YuviRed support and tell us why you need it."; or "Your company has reached today's limit for YuviRed Lite downloads. Tell us what you need and the YuviRed team will reply by email." Each has a form for your reason.

Activate with a code

Purpose: turn the server into a licensed YuviRed Lite (or Pro or Enterprise, with that edition's code). Prerequisites: signed in as the owner account created at first-time setup; the code; the server can reach https://activate.yuvired.com on TCP port 443.

1. Choose **Activate now** on the banner, or open **Configuration**, find the **License** panel and choose **Activate or import license**. The screen **Activate YuviRed** opens and says "You have 6 days left to activate this server. After that, the console is unavailable until the server is activated."
2. Note the **Installation ID** (with **Copy**; the toast says "Installation ID copied"). It identifies this server. Support may ask for it. It is not a secret.
3. Under **Activation code** type the code and choose **Activate**. Under the field: "Activating sends YuviRed this server's installation ID, computer name, Windows version and YuviRed version."

What you should see: the toast **License activated**; the License panel shows the badge **Active**, the title **YuviRed Lite**, **Licensed to** (your company), **Plan** YuviRed Lite, **Expires** Never, **Serial** filled in, **Installation ID**, **Last check-in** Not yet (until the first daily check-in), the counters **Computers used** and **AI agents used**, and the line "YuviRed Lite is free and never expires."

Only the owner account can activate. Anyone else sees "Only a YuviRed administrator can activate this server. Sign in with the administrator account created when YuviRed was installed." If you leave the code empty: "Enter the activation code."

Common problems (the text under **Activation code**, then "Need help? Email support@yuvired.com and include the installation ID."):

The screen says	What it means and what to do
That activation code is not valid, was already used on another server, or has expired.	A code works once, for one server, 7 days. Check the typing; ask for a new code
YuviRed could not be reached. Check this server's internet access and try again, or use a license file.	The server cannot reach activate.yuvired.com on port 443, or a proxy replaces certificates. Fix that or use the license file
Activation is temporarily unavailable. Try again later, or use a license file.	YuviRed's signing service is down. The code is not used up; try later
Your company already has 2 YuviRed Lite servers. To add another server, contact YuviRed support and tell us why you need it.	Use the **contact YuviRed** link at the bottom and give your reason
Too many activation attempts. Try again in an hour.	Wait an hour
Online activation is not set up on this server. Use a license file instead.	This build has no activation address; use the license file
YuviRed sent an unexpected answer. Try again in a few minutes. / YuviRed could not activate this server right now. Try again in a few minutes.	A fault on YuviRed's side or on the way; retry

Rollback: none needed; a refused code changes nothing.

The License panel. It shows the plan, the state and how many computers and AI agents you are using. Serial and installation ID are covered.

Activate with a license file (no internet on the server)

Prerequisites: your activation code, the server's Installation ID, and any computer with internet access.

1. On the server's activation screen, note the **Installation ID**.
2. On the internet-connected computer, open the license file link in your YuviRed email. The page **Get a YuviRed license file** says "Use this page for a YuviRed server without

internet access. You need the activation code from your YuviRed email and the installation ID shown on the server's activation page."

3. Enter the **Activation code** and **Installation ID**, then choose **Download license file** ("This sends YuviRed the activation code and the installation ID."). The page is at https://activate.yuvired.com/license-file.
4. Carry the file to the server. On the activation screen, under **License file** ("For a server without internet access, or when YuviRed sends you a license file. ..."), choose **Upload license file** and pick it. Accepted types: .txt, .lic and .yuvired-licence. A file that is too large is refused ("That file is too large to be a YuviRed license file."). The file is plain text: a few comment lines for you (plan, company, expiry, serial) and one license token; do not edit it.

What you should see: the same **License activated** result. If the console is still locked afterwards: "The license was stored, but the console is still locked. Email support@yuvired.com."

If the license-file page refuses: it shows the reason, "Need help? Email support@yuvired.com and include the installation ID.", and for a server-allowance refusal a form ("Tell us why you need another server and the YuviRed team will reply by email.") with **Your work email**, **Company**, **Why you need another server** and **Send to the YuviRed team**. If the service is down: "The YuviRed license service is temporarily unavailable. Try again later."

When YuviRed makes a license by hand (Pro, Enterprise, or special cases), choose **Download license request** on the same screen or in the Configuration License panel. It saves `yuvired-license-request.txt` with the Installation ID and the YuviRed version; email it to support@yuvired.com. YuviRed returns a license file bound to that Installation ID.

Limits, and what counts

- **Computers.** Lite 10, Pro 50, plus the server itself, whose own agent takes no seat in any edition. A computer takes a seat when it first connects after its Add Computer installer runs, not when you download the installer. A download does not reserve seats. If no seat is free, the server refuses the computer, even with a valid download, and the download is not used up.
- **AI agents (AMS).** Lite 5, Pro 50. These are seated AMS agents.
- A computer counts while it is enrolled, including a quarantined one. The console frees a seat when you remove the computer under **Endpoints** (there is no separate "decommission" control; a removal deletes that computer's history). Removing a computer that is reinstalled later takes a seat again.

- The License panel shows **Computers used** and **AI agents used**, and warns "Every computer seat is in use. A new computer is refused until one is free. To free one, open Endpoints and remove a computer you no longer use. Removing it deletes that computer's history."
- Existing computers are never cut off by a limit: they keep reporting.

Renewals and upgrades (Lite to Pro)

Lite never expires and needs no renewal. Pro and Enterprise are requested at https://yuvired.com/request (or by email to support@yuvired.com). YuviRed approves the request, you pay, and only then is the license issued as a code or a license file. There is no trial. Renewal is manual: request it the same way.

To move a server from Lite to Pro, activate the Pro code or license file on the same server, as above. A new license with a different serial supersedes the old one cleanly; the old record and its audit entry stay. Your computers, data and settings stay. A server that has had a paid license never takes a free Lite license again: the refusal reads "this server has held a Pro or Enterprise license, so it does not take a free Lite license. Contact YuviRed to renew the paid license".

Reminders: 30, 14 and 7 days before a Pro or Enterprise license ends, a banner says "Your YuviRed Pro license ends in <N> days (<date>). When it ends, the console is unavailable until a new license is activated. Enrolled computers keep reporting, and YuviRed only monitors until then. To renew, email support@yuvired.com." with **License**.

Expiry

A Pro or Enterprise license that ends locks the console: there is no grace period and no fallback to Lite. The screen titled **Your YuviRed license has expired** says "This YuviRed license has expired. The console is unavailable until a new license is activated. Enrolled computers keep reporting, and YuviRed only monitors until then." It shows the plan and end date, and the link **Request or renew a Pro or Enterprise license** (https://yuvired.com/request), "or email support@yuvired.com." Activate the new code or file on this same screen.

Revocation

- **By YuviRed.** The server learns of it at its daily check-in. The screen **This YuviRed license was revoked** says "YuviRed revoked the license for this server. The console is unavailable until a new license is activated. ... If you think this is a mistake, email support@yuvired.com." Email with the Installation ID. A new code or file (a new serial) unlocks it; the revoked license can never be used again.
- **On this server.** The Windows server's License panel has no revoke control in this release (the **Revoke Active License** button belongs to the Linux appliance's panel and is hidden here), and the server refuses to revoke a Lite license at all ("the free YuviRed Lite allowance cannot be revoked"). If a Pro or Enterprise license is ever revoked locally through the programming interface, the screen reads "This YuviRed license was revoked on this server. The console is unavailable until a new license is activated. ..." and a new license unlocks it.

The daily check-in, and what it sends

After activation (not during the 7 days), the server contacts https://activate.yuvired.com about once a day. It sends exactly six facts: the format name (`yuvired-checkin/1`), the Installation ID, the YuviRed version, the license serial, the number of computers, and the number of AI agents. It never sends computer names, user names, addresses, prompts or anything a computer did. When you activate, it sends the activation code, the Installation ID, the computer name, Windows version and YuviRed version. YuviRed also sees the internet address your server connects from and keeps a keyed hash of it in hourly counters to limit attempts; those rows are deleted after an hour. Your monitoring data never goes to YuviRed.

If YuviRed cannot be reached, nothing changes: the server keeps working and **Last check-in** in the License panel keeps its last value (**Not yet** until the first success). The reply can revoke this server's license or name a newer version, shown as a cyan banner: "YuviRed <version> is available. This server runs <current version>. Nothing is installed automatically." with **Download** and **Dismiss** (hides it for that version in this browser).

A server moved to new hardware (server_changed)

The server binds its license to its hardware identity. If the virtual or physical hardware changes (for example a VM restored or imported as a new VM, some hypervisor upgrades, or a copy), YuviRed treats it as a different server and locks the console. The screen **This server needs a new license** says "This server's hardware changed, so YuviRed treats it as a copy of

another server. The console is unavailable until a new license is activated for this server. Enrolled computers keep reporting, and YuviRed only monitors until then. If you moved or restored this server, email support@yuvired.com with the installation ID shown on this screen." The server now has a new Installation ID.

1. If it is a copy of another YuviRed server, shut the copy down. Do not activate it.
2. If you moved or restored it, email support@yuvired.com with the new Installation ID.
3. Activate the new code or license file YuviRed sends. The old license file is refused, and moving the server back does not bring it back.

Snapshot restores, host moves, and changes to network card, MAC address, CPU or memory do not change the identity. The **Add Computer** panel may also show "This server now runs on different virtual hardware. If it was moved or restored, confirm below. If it is a copy of another server, do not confirm: stop this copy. Computers already enrolled keep reporting either way." with the button **This server was moved or restored**. On a licensed server that button does not unlock anything: it answers "This server needs a new license because its hardware changed. Activate a new license on the license screen." The new license is the only way through.

If the screen says "YuviRed could not finish checking this server's hardware. Restart the YuviRed service, then try again.", restart the service (`Restart-Service YuviRed` in PowerShell as administrator) and open the license screen again.

Every state of the license screen

While the server is locked, every page shows only the license screen; its title names the state.

Title	Message	What to do
Activate YuviRed	This YuviRed server is not activated. The 7 days to activate it have ended, so the console is unavailable until you activate it. Enrolled computers keep reporting, and YuviRed only monitors until then.	Activate by code or file
Your YuviRed license has expired	This YuviRed license has expired. The console is unavailable until a new license is activated. ...	Renew at https://yuvired.com/request
This YuviRed license was revoked	YuviRed revoked the license for this server. ... (or "revoked on this server")	See Revocation
This YuviRed license could not be verified	This YuviRed license could not be verified. The console is unavailable until a valid license is activated.	Activate a new code or file; if it fails, email support with the Installation ID
This server needs a new license	This server's hardware changed ...	See above
YuviRed license	Opened from Configuration while the license is in force: "Activate this server, or import a new license file."	Close it, or import a new file

Badges in the License panel: Active, Ends soon, Expired, Revoked, Not valid, Not activated (amber during the 7 days, red after), Server changed; the title beside it reads Not activated, YuviRed Lite, YuviRed Pro, YuviRed Enterprise or Needs a new license. One more case: a Lite license that is "not valid" does not lock the console; the panel says "This license is not valid on this server, so new computers are refused. Activate a new license, or email support@yuvired.com with this server's installation ID."

Recovery when locked

1. Read the title and message on the license screen; find it in the table above.
2. Note the **Installation ID** (**Copy**). Nothing else is needed to ask for help.
3. Activate a valid code or license file on the same screen, signed in as the owner account; **Sign out** is there if you are on the wrong account. Sign-in still works while locked;

every other page answers with the license screen.
4. If you cannot activate, email support@yuvired.com with the Installation ID and the exact message.
5. Meanwhile, enrolled computers keep reporting and their data is kept. It is visible again once a new license is activated. A browser block an administrator applied earlier stays until rolled back, and no new blocking action is delivered while locked.

What YuviRed receives from your server

Only license facts, described above. Your monitoring data, computer names, users and prompts stay on your server. The full statement is the Privacy Policy on the Legal page.

Not verified in this guide

- The Pro and Enterprise purchase screens on yuvired.com and the exact wording of the download page (this guide follows the process owner-approved in the facts sheet and the messages and emails in the activation service, not the website pages).
- The expired, revoked, not-verified and server-changed screens were read from code; the lab ran the activation window, a code activation and a license-file activation on this release, not those states.

Endpoints: adding, checking, updating and removing computers

For the person who administers a YuviRed server. Describes YuviRed 1.0 with endpoint agent 1.8.2 and browser extension 0.5.5. Windows only: endpoints are Windows 10 and 11 and Windows Server.

What an endpoint is

An endpoint is a Windows computer that runs the YuviRed Endpoint Agent (the Windows service named `SyntinelRedAgent`, display name "YuviRed Endpoint Agent"; the service name, the state folder `C:\ProgramData\SyntinelRed` and the native messaging host `com.syntinelred.native_host` carry the product's earlier name; it will change in a later release). The agent reports to your own YuviRed server, never to YuviRed LLC. Each enrolled computer is one row on the **Endpoints** page. Lite allows 10 computers (the server itself does not use a seat) and 5 AI agents under Agent Monitoring System (AMS); other editions are in the licensing guide.

The agent collects metadata only. Eight kinds of telemetry are sent: `process`, `network`, `browser_ai`, `extension`, `software`, `local_ai`, `dns` and `agent_signal`. The exact contents are in the data notice (`https://<server>:8443/static/legal/data-notice.html`). Give people the notice before you install anything: see monitoring-disclosures.md.

Add a computer (Add Computer)

Purpose. Put the agent on another Windows computer and enroll it with your server.

Prerequisites.

1. **The server must be switched to network mode.** A new server answers only on the computer it is installed on (127.0.0.1). No other computer can enroll until you do this once, on the server, in PowerShell opened as administrator:

```
$py = "C:\Program Files\YuviRed\runtime\python.exe"
$yr = "C:\Program Files\YuviRed\runtime\yuvired_service.py"
& $py $yr network --name yuvired.northwind.example
```

Use a DNS name that every computer can resolve, or a fixed IP address. Prefer the DNS name. Every Add Computer download points agents at this name. **Do not change it later**: agents trust the server's certificate for that name and lose the server if the name or address changes (see "Troubleshooting"). The command restarts YuviRed and waits until it is ready; allow up to 7 minutes. `& $py $yr network --status` shows the binding, the firewall profile and whether the firewall rule exists. `& $py $yr network --off` goes back to the local-only setting. The command opens one firewall rule: TCP port 8443, Windows network profiles Domain and Private only. A network Windows has classed as **Public** is never opened. The server's own agent is skipped in network mode and holds no seat.

2. You are signed in to the console with an account that can manage endpoints (administrator).
3. A free computer seat: the page shows "Computers: 3 of 10".
4. On the target computer: a local administrator account and a network path to `https://yuvired.northwind.example:8443`.
5. The people who use that computer have been told (notice, see the disclosures reference). Running the installer confirms that your organization may monitor the computer and has told its users.

Steps.

1. In the console, open **Agent Deployment** (left menu).
2. Find the **Add Computer** panel ("Download an installer for another Windows computer on this network"). It shows the server address, "Computers: x of 10", "AI agents: x of 5" and the certificate SHA-256 fingerprint.
3. Set **Computers this download may enroll** (1 to the number of free seats).
4. Select **Download installer**. Your browser saves `YuviRed-Add-Computer.zip`.
5. The download is valid for **24 hours**. The zip holds `YuviRed-Agent.msi`, `enrol.yuvired` (server address, the server's certificate and a short-lived token), `Install-YuviRed-Agent.cmd`, `Uninstall-YuviRed-Agent.cmd` and `README.txt`. Treat it as a secret: `enrol.yuvired` is a key to your server. Copy it to the target computer over a channel you trust.
6. On the target computer, open PowerShell **as administrator**. Do the remaining steps in that window, not in File Explorer. (Do not use right-click, Run as administrator on the file: Explorer may offer to "give you access" to the folder, which makes the installer refuse it.)

7. Make the install folder, which only administrators can change. The command is in the zip's `README.txt`:

```
$d='C:\ProgramData\YuviRed-Install'; if (Test-Path $d) { 'That folder
already exists. Delete it, then run this again.' } else { $s=New-Object
Security.AccessControl.DirectorySecurity;
$s.SetSecurityDescriptorSddlForm('O:BAD:P(A;OICI;FA;;;SY)
(A;OICI;FA;;;BA)'); [void][IO.Directory]::CreateDirectory($d,$s); 'Ready:
' + $d }
```

(The version in `README.txt` also reads the folder back to prove it was made safely. Copy it from there.) It must print `Ready: C:\ProgramData\YuviRed-Install`.

8. Unzip into it: `Expand-Archive -Path "$HOME\Downloads\YuviRed-Add-Computer.zip" -DestinationPath C:\ProgramData\YuviRed-Install` (change the path if the zip is elsewhere).
9. Run the installer: `& C:\ProgramData\YuviRed-Install\Install-YuviRed-Agent.cmd`

What the installer does. It refuses to run unless it is elevated and unless the folder and every file in it can be changed only by SYSTEM, Administrators or TrustedInstaller (it stops with "Nothing was installed" otherwise). It then installs the agent silently from the MSI into `C:\Program Files\YuviRed Agent`, configures it from `enrol.yuvired` (the token goes into one protected file, never onto a command line), deletes `enrol.yuvired`, installs the Windows service `SyntinelRedAgent` (delayed automatic start, restart on failure), registers the native messaging host for Chrome, Edge and Firefox, writes no browser policy (the extension's store ids are empty in this release, so it prints "Browser extension: not in the Chrome and Edge stores yet. Add it by hand as the install guide shows."), and waits up to 90 seconds for the server to accept the computer. The monitoring notice is installed as `C:\Program Files\YuviRed Agent\NOTICE.txt`. The agent keeps its state in `C:\ProgramData\SyntinelRed` (readable by SYSTEM and Administrators only) and its log in `C:\ProgramData\SyntinelRed\logs\syntinel-agent.log`.

What you should see. The window prints "Installing or updating the YuviRed agent. This can take a few minutes." and ends with "Done. Check that this computer is listed under Endpoints on the server.", the monitoring notice location and the four browser steps (see extensions.md), then "Press any key". If Windows asks for a restart it says so. If it ends with "This computer is NOT enrolled yet", read the line above it (see "Troubleshooting").

How to confirm. See the next section.

Afterwards. Delete `YuviRed-Add-Computer.zip` from the Downloads folder and every other copy once all the computers it was for are added. To stop a download early: open **Deployment Campaigns**, open the campaign named "Add Computer" followed by the date

and time (UTC) the download was made, and choose **Revoke**. Computers it already added keep working. A download does not reserve seats: a computer takes a seat when it first connects, and the server refuses it if none is free.

Common problems. "The folder is not safe": delete the folder and make it again with the command in step 7. Code 1618: another installation is running, wait and run again. Not elevated: open PowerShell as administrator. Log: `%TEMP%\YuviRed-Agent-Install.log`.

Rollback. Run `& C:\ProgramData\YuviRed-Install\Uninstall-YuviRed-Agent.cmd` (see "Remove an agent"), then remove the computer under **Endpoints** to free its seat.

The Add Computer panel. Set the number of computers the download may enroll, then choose Download installer. Captured in the YuviRed lab; your screen shows your own server's address.

Centrally managed deployment (group policy, Intune, other tools)

What is supported: the Add Computer download, run as administrator on each computer, as above.

What is not supported, plainly:

- **There is no MSI property for the enrollment token.** The token travels only inside `enrol.yuvired`, never on a command line, so running `YuviRed-Agent.msi` alone (including through a Group Policy software package or an Intune line-of-business app) installs the agent without connecting it to your server. YuviRed warns against it.
- A custom install folder (`INSTALLDIR`) and `msiexec /fa` repair are not tested.
- No group policy template, Intune package or silent-switch script for enrollment is shipped. The download expires after 24 hours and enrolls at most the number of computers you chose (up to 10).
- The scripts and the MSI are not Authenticode-signed; a computer that allows only signed code can block the agent.

If your management tool can run `Install-YuviRed-Agent.cmd` elevated from a folder only administrators can change, it may work, but YuviRed has not tested that route. Test on one computer first.

Confirm an endpoint is connected

Steps.

1. Open **Endpoints**. The page title reads "Enrolled Endpoints" with the total.
2. Find the computer by name. The columns are Hostname, User, OS, Agent, IP, Risk, Last Seen, Status.
3. **Agent** should read v1.8.2. **Status** shows a dot and one of three words.

Status	Meaning
Online	the last check-in was 180 seconds ago or less
Stale	the last check-in was between 180 and 600 seconds ago
Offline	more than 600 seconds ago, or never

The agent checks in every 60 seconds, so a new computer normally shows Online within a few minutes. The installer waits up to 90 seconds for acceptance. The filter above the table (All status, Online, Stale, Offline) and the sort "last seen" help in a larger list. The refresh button on a row asks the agent to re-read its settings at its next check-in (about 60 seconds).

On the computer. In PowerShell: `Get-Service SyntinelRedAgent` should show Running. The log should contain "Registration OK" and then "Check-in OK" lines.

What a check-in carries, and offline devices

Every 60 seconds the agent runs its collectors and sends one batch (at most 2,000 items): the eight telemetry kinds, the computer name, the signed-in user name (domain removed), one IP address, the Windows version and the agent version, plus a short health report about the agent. At start and then every tenth cycle (about every ten minutes), and when you press the row's refresh button, it also pulls its settings.

The server can answer with settings, a refresh flag, a request to upload the agent log tail, approved response actions and AMS commands. There is no remote command execution.

Offline computers. If the computer is off or cannot reach the server, the console shows it Stale, then Offline, and keeps its last data. **Nothing is queued while it is away: a check-in that cannot be sent is dropped and that cycle's telemetry is lost.** After more than three failures in a row the agent slows down, to at most one try every 300 seconds, and keeps

trying. When the server is reachable again it resumes by itself. (The AMS part keeps up to 5,000 events on disk and sends them when it can.)

An offline computer still holds its seat. If the server is not yet activated after the 7 day window, or a paid license has ended, the console locks, but enrolled computers keep reporting and YuviRed only monitors.

Update an agent

A Lite agent is installed from the MSI, and **MSI-installed agents do not update themselves**, whatever the setting Automatic Agent Updates says. To move a computer to a newer agent (and a newer extension):

1. Install the new YuviRed server build (the signed in-place upgrade). Your computers, policy and settings are kept.
2. On the server, make a new Add Computer download (a new zip carries the newer MSI).
3. On each computer, repeat steps 6 to 9 of "Add a computer". The MSI installs over the old version and the computer keeps its identity and uses no second seat. The new download also replaces the extension files; see extensions.md.

Upgrading the server does not update the agent running on the server itself.

Endpoint credential rotation

What it is. Each agent authenticates with its own secret credential. Rotation replaces it without reinstalling. The server issues a new credential, the agent saves it and proves it works, and the old one stops being accepted (both work together for a 30 minute overlap).

When to use it. After you suspect a credential was exposed, when a computer is handed to another person, or on a routine schedule you set yourself. It needs agent 1.7.20 or newer; agent 1.8.2 qualifies.

Steps.

1. Open **Agent Deployment** and scroll to **Endpoint Credential Rotation** (administrators only).
2. Choose **Rotate Credentials for All Endpoints**, or tick computers in the list and choose **Rotate Credentials for Selected Endpoints**.
3. In the dialog choose a **Window** (default 7 days; 1, 3, 14 or 30 days are available) and select **Confirm Rotation**.

What you should see. A row appears under **Rotation Operations**. Online computers rotate on their next check-in. Offline computers stay visibly pending, with a reason, and finish when they reconnect, until the window ends. A completed row means the new credential was verified and the old one rejected.

Common problems. A computer pending for "waiting for agent 1.7.20 or newer" runs an older agent. A pending row on an online computer reads "waiting for the endpoint's next check-in (about once a minute)"; on an offline computer "endpoint is offline; the rotation starts automatically when it reconnects". "The set of eligible endpoints changed since you opened this dialog": the count was refreshed, confirm again. A failed row has **Retry**.

Rollback. There is none; rotation is one way. If a computer fails to complete after Retry, remove it and add it again.

Decommission and delete

The console has **one** control: the red **Remove endpoint** button on an Endpoints row (administrators). Decommission exists only in the server's interface for programs (`POST /api/v1/endpoints/{id}/decommission` , and `/reactivate`) and has no button in this release.

	Decommission (API only)	**Remove endpoint (console)**
Seat	freed	freed, if the computer used one
History, alerts, audit trail	kept	telemetry and alerts deleted permanently
Agent on the computer	refused with HTTP 410, backs off (at most one try every 300 seconds) and keeps its identity	its check-ins are refused; while the Add Computer download it was installed with is still valid (24 hours from the download) the agent can enroll again, on its next recovery attempt or restart, as a new row that takes a seat; after that it stops reporting for good
Undo	reactivate (takes a seat again)	cannot be undone

Remove endpoint shows counts first. Telemetry and alerts are deleted. AI content events are kept but unlinked; AI agent records are archived; open investigations stay open but

become unlinked. Select **Remove endpoint** in the dialog to confirm. Uninstalling the agent tells the server nothing: the row and the seat stay until you remove them. **Uninstall the agent on the computer first, then remove the row**, so the computer cannot enroll itself again.

Remove an agent from a computer

1. Preferred: run `& C:\ProgramData\YuviRed-Install\Uninstall-YuviRed-Agent.cmd` from an administrator PowerShell (keep the folder from the install). It prints "The YuviRed agent has been removed from this computer." and deletes the download files.
2. Or: **Settings, Apps** (Add or Remove Programs), choose **YuviRed Agent**, then **Uninstall**. This does not report what its cleanup could not remove.
3. Remove the computer under **Endpoints** to free the seat.

The removal deletes the service, the program files, the browser registration and `C:\ProgramData\SyntinelRed` (state and logs). If the cleanup did not finish the script lists what is left (exit code 3). A browser extension loaded by hand must be removed in `edge://extensions` or `chrome://extensions`. Install and removal logs stay in `%TEMP%` (`YuviRed-Agent-Install.log`, `YuviRed-Agent-Uninstall.log`).

Troubleshooting

Symptom	Likely cause and fix
The agent cannot reach the server (installer ends "NOT enrolled", exit 7, log "Cannot reach backend")	Check three things. (1) The server is still in local-only mode: on the server run `network --status`; if it says loopback, run `network --name <DNS name or fixed IP>`. (2) The server's network is on the **Public** Windows profile: the firewall rule covers Domain and Private only; set the network to Private or Domain. (3) The name does not resolve from the computer: test `Resolve-DnsName yuvired.northwind.example` and `Test-NetConnection yuvired.northwind.example -Port 8443`.
Service not running	`Get-Service SyntinelRedAgent`; then `Restart-Service SyntinelRedAgent` and read the log. The installer prints this command after a failed enrollment. It works for 24 hours after the download was made.
401 or 403 in the agent log	The server refused a credential. After a registration refusal the log gives the reason, for example "deployment token expired" or "enrollment cap reached" (make a new download), or "license seat limit reached" (remove a computer you no longer use, then restart the service). A 403 on a check-in means the credential no longer matches: rotation or a removed row. Reinstall from a new download.
Token expired	The download is 24 hours old, was revoked, or has added all its computers. Make a new download.
410 in the log	The endpoint was decommissioned through the API. Reactivate it, or reinstall.
402 in the log	Only a demo license produces it (expired, revoked or invalid). A Lite or Pro license lock does not stop check-ins; it locks the console and refuses new enrollments.
Time skew: "TLS certificate time check failed: check this computer's clock"	The computer's clock is wrong. Fix the time and restart the service.
TLS: "TLS certificate not trusted"	The agent trusts only the server's own certificate from `enrol.yuvired`. The certificate or address changed after the download. Make a new download and install it over the old one.

Symptom	Likely cause and fix
Seats show full	Open **Endpoints** and remove a computer you no longer use.

The YuviRed AI Prompt Monitor browser extension

For the person who administers a YuviRed server. Describes extension 0.5.5 with agent 1.8.2. Chrome and Edge (Chromium 121 or newer) on Windows. **Firefox is not supported.**

What it does and does not do

On 15 AI sites the extension checks, inside the browser and before anything is sent, what a person types, pastes, uploads or receives, for sensitive information such as passwords, keys, personal data and card numbers. It then applies your policy: observe silently, warn, or block. The sites are chatgpt.com, chat.openai.com, claude.ai, gemini.google.com, copilot.microsoft.com, perplexity.ai, m365.cloud.microsoft, copilot.cloud.microsoft, chat.deepseek.com, grok.com, chat.mistral.ai, meta.ai, poe.com, aistudio.google.com and huggingface.co/chat. (The extension's manifest lists 19 address patterns because some sites have a www form and HuggingChat has three.) Microsoft 365 Copilot, DeepSeek, Grok, Le Chat, Meta AI, Poe, Google AI Studio and HuggingChat have not been checked on the live sites; Warn and Block are not proven there until you test them.

It reports to **your** server only: the kind of finding, a masked fragment (at most the first 4 and last 4 characters; for a password only the beginning), the action taken, the site, the browser, and attachment counts. By default no prompt or reply text is kept. It does not read other websites, history, cookies or the clipboard, and it sends nothing to YuviRed LLC. It cannot see desktop applications or other browsers, and it is a control for people sending something by mistake, not against the computer's administrator.

How it reaches browsers in this release

Nothing installs the extension into Chrome or Edge automatically in YuviRed 1.0. The Add Computer installer puts the extension's files in `C:\Program Files\YuviRed Agent\extension` and prints the steps below; a person with local administrator rights then loads it into each browser profile by hand. That is the route today, and the only one this guide can tell you works.

Two policy-based routes exist in the product but neither is active on a YuviRed Lite server:

- *Store route (planned).* When the Chrome Web Store and Microsoft Edge Add-ons listings are approved, their ids ship in a later build (the file `setup\store-extension-ids.json` beside the agent; both ids are empty in 1.0). From then on the agent installer itself writes the Chrome and Edge force-install policy, the native messaging allowlist, and the extension's server address and capture level ("follow the console setting"). It never writes consent. Until the ids exist the installer writes no browser policy at all and prints "Browser extension: not in the Chrome and Edge stores yet. Add it by hand as the install guide shows." A store-hosted force-install works on any Windows computer that applies the policy, joined to a domain or not.
- *Self-hosted route (the console setting).* **Configuration, System, Remote Browser Extension Wiring** (off by default) lets Windows agents write a force-install policy for a packed extension that the server signs and serves itself. It needs a signing key on the server. **The YuviRed Lite server removes that key from its environment at start**, so on Lite the setting is present but inactive: every agent is told "disabled" whether the switch is on or off. Where it does apply (other editions), Chrome force-installs an extension from outside its store only on a computer joined to an Active Directory domain, joined to Microsoft Entra ID or enrolled in Chrome Enterprise Core, and Edge only on a domain-joined computer. Note that this route, where active, also sets the extension's consent by policy (see "Consent").

What the employee sees. Loaded by hand: nothing is added until you or they load it; after loading, the extension appears in `chrome://extensions` or `edge://extensions` with Developer mode on, and the person can remove it. Force-installed (store route, later): "Installed by your organization" and no Remove button.

Load the extension by hand

Purpose. Put the extension in a browser profile on an added computer. **Prerequisites.** The agent is installed and the computer is shown Online under **Endpoints**. A local administrator signs in to the computer (Developer mode and Load unpacked need no special right, but the folder is under Program Files and the steps are the administrator's). Each Windows user and each browser profile needs the extension loaded separately.

1. In Chrome open `chrome://extensions`, or in Edge open `edge://extensions`, and turn on the **Developer mode** toggle on that page. (Without it the **Load unpacked** button is not shown.)
2. Choose **Load unpacked** and select `C:\Program Files\YuviRed Agent\extension`. Use that folder, not a copy: a new Add Computer download replaces it, and deleting a copy silently removes the extension.

3. The extension's options page, titled **YuviRed AI Prompt Monitor**, opens by itself the first time (it opens whenever the extension is loaded without consent). If it does not, open the extension's **Details** and choose **Extension options**.
4. On that page leave **Backend URL** and **Deployment token** empty: with the agent on the computer the extension finds the server through the agent's native messaging host. Tick **I consent to local inspection of prompts and attachments on the monitored sites above** and choose **Save**. Nothing is recorded or reported until that is saved.
5. Restart the browser. If it offers to turn off Developer mode, keep it on.

What you should see. After Save the status next to the button reads "Saved. Ruleset loaded (n rules)." and the headline under **Connection** turns to "Connected to YuviRed Lite". Under **Endpoint ID** a green line reads "Auto-detected from local agent: <id> (used instead of the field above)". The disclosure card reads "Reports go to your organization's YuviRed server at https://<your server>:8443". **Common problems.** See Troubleshooting. **Rollback.** Remove it in the browser's extensions page.

The extension options page after Save. Notice the Connected heading and the policy line. Captured in the YuviRed lab; your screen shows your own server's address.

Consent

Nothing is recorded or reported until consent is given. (The page hooks are in place as soon as the extension is loaded, so text passes through the extension's memory on the monitored sites, but nothing is kept or sent.)

- **On a browser loaded by hand** (every browser in this release) the person ticks the consent box on the options page and saves. They can untick it later. The disclosure card on the same page reads "Policy consent value: not set by policy".
- **Consent by policy on managed browsers.** The extension accepts an organization-level `consented` value from managed browser policy (the extension's managed schema). When it is true, monitoring is active without the checkbox, the options page does not open by itself, the person is not asked and cannot withdraw it; the options page shows "Policy consent value: true (consent given by policy)". **The Add Computer installer never writes this value.** Two things in the product do, neither of them on a Lite server: the console's Remote Browser Extension Wiring route described above, where it is active, and the appliance's own extension installer script. If you set the value yourself

in your own browser policy (the `3rdparty\extensions\<id>\policy` key of Chrome or Edge), do it only after you have told people (the notice, see monitoring-disclosures.md).

What a person can change on the options page, and what they cannot. They can edit Backend URL, Deployment token, Endpoint ID, Monitored sites, Capture level (it can only lower what the console allows), and consent. With the agent on the computer, the endpoint id and the server the agent reports take precedence over the typed values. A site removed from Monitored sites is neither checked nor recorded, and capture level Disabled stops reports while warnings and blocks still run; **the console is not told about either**. A managed browser policy that sets those fields overrides the page. They cannot change your policy, the ruleset or the policy preset, and they cannot raise capture above what the server allows.

Monitor Only, Warn and Block

The policy is set in **Configuration, Prompt Protection** by an administrator. Presets: Monitor Only, Light Warnings, Warnings Only, Balanced, Strong (marked "recommended" in the console), Strict, Maximum, and Custom. **A new installation starts in Monitor Only.** Full tables are in the Prompt Protection guide.

Action	What the person sees on the page
Monitor Only (observe)	**Nothing.** The finding is recorded and an alert opens.
Warn	A dialog "Potentially sensitive information detected" and "Are you sure you want to send this information to {service}?" with **Go back and edit** (the default) and **Send anyway**. "Send anyway" is recorded. Escape or closing never sends.
Block	A dialog "Prompt blocked: sensitive information detected" with the items to remove and **Go back and edit** or **Close**. There is no override.
Draft notice	While typing, a small notice with a close (x) button: "Sensitive information detected in your draft: ... Your draft has not been sent." It never sends, and the finding is recorded.
Inspection unavailable	"Prompt not sent: inspection unavailable" when your policy holds what could not be checked.

Each dialog ends with a small line "YuviRed · ref <id>" that matches the event in **Prompt Monitoring**.

Raw capture

Off by default. When on, the first 20,000 characters of a prompt or reply that matched are stored on your server, encrypted, for 24 hours by default. It needs **all** of these: in **Configuration, Detection & Response**, Inspection Depth set to Full content inspection and **Keep Original Prompt Samples** on (an administrator with the configuration write permission), an HTTPS connection, and no per-endpoint exclusion. Reading a stored sample needs the separate "view raw content" permission, and every view is audited. Tell people first: the notice must say so.

Confirm the extension is connected

1. In the browser, open the extension's options. Under **Connection** read the headline: **Connected to YuviRed Lite** (green) means the last policy sync worked. "Waiting for the first policy sync", "Not synchronized" or "Not connected. Consent and save to start" mean it is not.
2. The rows show Product ("YuviRed Lite (local agent)"), Server, Transport ("local agent (native proxy)" on Lite), Policy (preset name and revision), Ruleset and Last sync.
3. In the console open **Prompt Monitoring** and find the **Browser protection** card. Each browser row reads **Protected**, **Monitoring only**, **Stale** or **Protection unavailable**, with the computer, browser, "ext" and the extension version, "policy applied", "policy pending" or "policy stale", and "last check-in". An enrolled computer no browser has reported from shows "No browser reporting yet". "update needed" marks an older extension. Before any browser has reported, the card reads "No browser extension has checked in yet" and its **Set up browser protection** steps are open. A row appears within a few minutes of consent (the extension reports its status every 5 minutes).

Update, and remove

Update. Make a new Add Computer download and run it on the computer (see endpoints.md). It replaces the files in `C:\Program Files\YuviRed Agent\extension`. Then restart the browser, or choose Reload on the extension. Nothing updates it silently; the console marks a browser "update needed" until it reports the new version.

Remove. Loaded by hand: choose Remove in `chrome://extensions` or `edge://extensions`. Removing the agent deletes the extension folder, so a hand-loaded copy stops working; remove the leftover entry in the browser (the Uninstall script prints this reminder). Policy-installed (store route, later): delete the force-install entry, then restart the browser.

Troubleshooting

Symptom	Cause and fix
Not in the browser after the agent install	The installer only puts the files on disk in this release. Load it by hand.
Headline "Not connected. Consent and save to start"	Consent not saved. Tick the box and choose Save.
Headline "Not synchronized", Last sync "failed: the server rejected the request" or "failed: no credential yet"	The agent is not on this computer, has not enrolled yet, or the computer is no longer enrolled. Check Endpoints.
Under Endpoint ID: "Local agent not detected. Using the Endpoint ID entered above." and Transport "direct HTTPS"	The native messaging host is not found. The extension was loaded before the agent, or the agent was removed: reinstall the agent, then Reload the extension.
Browser log: "Access to the specified native messaging host is forbidden"	The extension id is not allowed by the host manifest (a copy of the extension with a changed manifest), or a baseline policy sets `NativeMessagingBlocklist=*` without `NativeMessagingAllowlist`. Add `com.syntinelred.native_host` (the host keeps the product's earlier name; it will change in a later release).
Policy blocks it	Your own browser policy (for example `ExtensionInstallBlocklist=*`) blocks it; add its id to `ExtensionInstallAllowlist`. The id is fixed by the extension's manifest and is the same on every computer.
Disclosure card: "No server configured: nothing is reported" and Server "not set"	Leave Backend URL empty on a computer with the agent. Without the agent, enter the server's HTTPS address (the browser must then trust the server's certificate, which a Lite server's self-signed certificate is not by default).
Both a hand-loaded copy and a store copy	Remove the hand-loaded copy; both would check every prompt.
Row stays "Stale"	The browser is closed, the computer is off, consent was removed or its credential failed. The console cannot tell which.

Use **Copy status for support** on the options page for a support message. It holds counts and versions, no prompts.

The dashboard and the console tour

This chapter tours the YuviRed console: the sidebar, the top bar and the Dashboard page. It explains every number on the Dashboard, how fresh it is, how to read the red, amber and blue states, and where to click next. The example company in this guide is "Northwind Dental". All names, addresses and numbers are made up.

Examples use computer names such as NW-FRONTDESK-01, the user a.rivera and the address 192.0.2.10.

1. Before you start

- You need a signed-in console account. Role matters: an Admin sees every menu item, an Analyst sees most, and a View Only user sees the monitoring pages but not the Enforcement page or the Admin and Deploy sections. A page that your role cannot open is hidden, not greyed out.
- The console is reached at the address of your YuviRed server, for example https://192.0.2.10:8443. Your browser may warn about the server's certificate if the installation uses its own; that is covered in the installation guide.
- YuviRed Lite is the edition described here. It covers 10 computers plus the server itself, and 5 AI agents in AMS (Agent Monitoring System). Pro and Enterprise show the same pages with higher limits.

The Dashboard of a lab install with two computers enrolled. The counts come from lab data.

2. The layout of the console

2.1 The sidebar (left, always dark)

At the top is the YuviRed logo and a two-button switch:

- **AIDR** (AI Detection and Response): the monitoring and governance pages for AI use in browsers, programs, networks and prompts. Everything in this chapter and the AIDR chapter lives here.
- **AMS** (Agent Monitoring System): the pages for AI agents such as Claude Code. Covered in the AMS chapter.

Click either button to swap the list below it. The sidebar remembers whether you collapsed it (the **Collapse** control at the bottom turns it into an icon-only rail; hover an icon to see its name).

The AIDR list is grouped into sections:

Section	Pages	What they are for
Monitor	Dashboard, AI Tools, Alerts, Prompt Monitoring, Network, Endpoints, AI Map, User Behavior, Investigations, Artifact Analysis	Seeing what AI is in use and what needs attention. Artifact Analysis only appears for accounts with the artifact permission, and the feature is off by default.
Govern	Policies, Enforcement, Blocklist, Reports, Deployment Campaigns	Deciding what is allowed, taking response actions, producing reports. Enforcement is hidden from View Only users. Blocklist is Admin only.
Deploy (Admin)	Agent Deployment, Adoption Requests	Adding computers. See the deployment guide.
Admin (Admin)	Configuration, Legal, Help, Team / Users, Debug / Reset, System Logs	Settings, legal notices, support information, accounts, diagnostics.

The AMS list holds five areas: Overview, Agents, Governance, Alerts & Actions, Coverage & Connections, plus a Reports shortcut under "Related".

Two small badges draw the eye:

- On **Alerts**: the number of open alerts. If any open alert is critical, the badge changes to "N crit" (for example "2 crit") in red. When there are no open alerts, the badge is hidden.
- On the AMS list: an amber badge on **Agents** counts agents that need a decision, and an amber badge on **Alerts & Actions** counts approvals waiting for a person.

At the bottom of the sidebar:

- A live-updates line ("Connecting live updates..." and then the live state) and a count such as "4 endpoints online · 1 stale".
- Your name and role. Click it for the account menu: **My Account**, **Change Password**, **Authentication / MFA** (multifactor authentication), **Recovery Codes**, **Active Sessions** (shown as "soon" and disabled: this build does not track individual sessions) and **Sign Out**.

- A small build line that tells you which build the console is running. Quote it when you contact support.

2.2 The top bar

From left to right:

- The page title.
- A privacy badge. **Metadata mode** (blue) means YuviRed keeps activity metadata only for AI prompts and responses. **Full inspection** (red; with "(raw capture ON)" when original prompt samples are being stored) means content can be kept. Red here is deliberate: content is being stored, and you should be sure that is what you intended. Hover the badge for the full explanation.
- Your role badge and a clock.
- **Filter**: opens a search bar for the page you are on (search text, status, risk or severity). It filters the table on the current page only.
- **Export**: downloads the current page as a CSV file.
- The moon icon: switches between the light and dark theme. Light is the default the first time you sign in; your choice is remembered in this browser.
- **Refresh**: reloads all data now instead of waiting for the next automatic refresh.

2.3 Banners under the top bar

Up to three banners can appear above the page. They are not errors:

- A **license banner** during the 7-day activation window of a new server, and 30, 14 and 7 days before a paid license ends. The **License** link next to it opens the license screen.
- An **update banner** when the daily check-in learns that a newer YuviRed exists. It is a notice only: nothing installs by itself. **Dismiss** hides it for that version.
- A **critical alert banner** (red) naming the most recent open critical alert. **Investigate** opens the Alerts page. The small x dismisses it.

---

3. The Dashboard, widget by widget

Open **Dashboard** from the **Monitor** section. A small "Data freshness" label at the top right of the page explains timing; see section 4.

3.1 The five number cards

Each card is clickable and takes you to the page that explains the number.

Card	What the number counts	Sub-line under it	Click goes to
Endpoints	Computers enrolled in YuviRed.	"N online, N stale, N offline"	Endpoints
AI Tools	Distinct AI tools seen on your computers (programs, browser AI sites, extensions, local services).	"Detected"	AI Tools
Shadow AI	Tools classified Unauthorized plus tools still Under Review.	"N unauthorized · N review"	AI Tools, filtered to the shadow list
Open Alerts	Alerts that are still open (not resolved, not closed as false positive).	"N critical"	Alerts, filtered to open
Data Egress	Bytes sent to AI services. In this release it always reads **"Not measured"**.	"Byte counts not collected"	Network

How to read them:

- **Online, stale, offline.** A computer checks in about every 60 seconds. It reads **online** when its last check-in was within 3 minutes (180 seconds), **stale** from 3 to 10 minutes, and **offline** beyond 10 minutes. A stale computer is a soft warning (a sleeping laptop, a network blip). Offline means YuviRed has heard nothing for a while and cannot tell you what is happening on it.
- **Shadow AI** is "AI use nobody has approved". It is the sum of two different things, which the sub-line separates: tools you have classified **Unauthorized** (a decision you made) and tools **Under Review** (a backlog nobody has decided on yet).
- **Data Egress** reads "Not measured" in this release. The endpoint program does not measure how many bytes each connection sends, so YuviRed will not show a zero that could be mistaken for "nothing left the building". Treat the card as a placeholder; use the Network page to see which AI destinations are reached.

Colors on the cards:

- **Shadow AI** is red when at least one tool is Unauthorized, amber when the count is made only of tools under review, and neutral at zero. A value of "Unknown" is never red.
- **Open Alerts** is red whenever it is above zero.

3.2 Risk Distribution

A row of five tiles: **Critical**, **High**, **Medium**, **Low**, **Clear**. Each shows how many computers sit in that risk tier. A computer's tier is derived from the alerts and classifications that apply to it. Click a tile to open the Endpoints page filtered to that tier ("Clear" opens the whole list). The colors are red, orange, amber, green and grey.

3.3 Active Policies

The first six security policies, each with a switch. The switch turns a detector or rule on or off for everyone; only Admins can change it, and everyone else sees it dimmed. A small amber "no input" tag means the rule has nothing to check in this release (the data egress rule is the example, because byte counts are not measured). **Manage** opens the Policies page.

3.4 AI Traffic

A chart of AI connections per hour for the last 24 hours. If there was no AI traffic in the last 24 hours, the chart falls back to all time and says so in its sub-heading ("all time (no traffic in last 24h)"). It never claims 24 hours when it is showing something else.

3.5 Top AI Destinations

The six AI services your computers contact most, by number of connections. **View all** opens the Network page, where you can filter by 24 hours, 7 days or 30 days and download a destination list for your own firewall.

3.6 Recent Alerts

The eight most recently active alerts, newest first, each with a colored bar for severity, the title, the computer, a short detail and the time. An alert that keeps recurring moves to the top because the list sorts by last activity, not by when the alert was first created. Click a row to open the alert. If the alert list cannot be loaded, the panel says "Couldn't load alerts" with a **Retry** button. It never shows a reassuring "No Alerts" when it simply failed to ask.

3.7 Active Endpoints

The eight computers with the highest risk score, with the signed-in user, the operating system and the risk word. The dot on the left shows liveness: green online, amber stale, grey offline, and an orange warning dot for an online computer at high or critical risk. Click a row to open the computer's detail.

3.8 Detected AI Tools

A table of up to eight tools: **Tool**, **Vendor**, **Category**, **Status** (the classification), **Endpoints** (how many computers), **Severity** and **Last Seen**. Two chips at the right of the heading summarize: "N total" and, in red when above zero, "N shadow".

---

4. How fresh is what you see

The **Data freshness** label at the top of the page has a tooltip with the exact timings. In plain words:

- Each computer checks in every 60 seconds.
- The console re-reads all Dashboard data every 30 seconds. **Refresh** in the top bar does it now.
- New and repeating alerts are also pushed to the page within about 2 seconds of YuviRed recording them.
- Worst case from an event on a computer to the Dashboard: about 90 seconds for programs, network connections, installed software and local AI services (one 60-second check-in plus one 30-second refresh).
- A new visit to an AI website in a browser takes longer, about 6.5 minutes at the default settings, because the computer scans browser history every 5 minutes. The scan interval can be tuned in Configuration under AI & Agent Monitoring, "Browser AI Scan Frequency" (every 2, 5 or 15 minutes).

The Dashboard has no time filter of its own. Its cards show the current state, and its traffic chart covers the last 24 hours. The pages that have time filters are:

Page	Time choices
Alerts	All Time, Last 24 hours, Last 7 days, Last 30 days (by last activity)
AI Map	Last 5 minutes, Last 30 minutes, Last hour, Last 24 hours, Last 7 days (default), Last 30 days, All time, Custom range
Network	Last 24h, Last 7 days, Last 30 days
Prompt Monitoring	Last 1d, 7d, 30d, 90d
Reports	A start and end date in a time zone you choose

If a page looks out of date, check the sidebar line first. If it says live updates are not connected, use **Refresh**.

5. Reading the colors

The console uses color to mean one thing everywhere:

Color	Meaning	Examples
Red	Danger, a violation, something to act on now.	Critical severity; Unauthorized tools; a non-zero Open Alerts count; a blocked prompt; "Full inspection (raw capture ON)".
Orange	High severity, triage today.	High alerts; High risk tier.
Amber	Unknown, under review, a warning or a partial risk.	Under Review tools; Medium severity; a stale computer; a pending policy; a warning dialog shown to a person.
Blue, cyan, purple	Normal activity and informational states.	Monitored tools; Metadata mode; an observed event with no finding.
Green	Healthy or approved.	Online computers; Approved tools; Low severity; Protected browser.
Grey	Nothing to report, or not measured.	Clear risk tier; offline computers; "Not measured".

Two cautions:

- **Amber and "Unknown" are not green.** A card that says "Unknown" or "Not known" means YuviRed could not read the value. It is never a silent zero.
- **A zero is only as good as the reporting behind it.** If computers are offline, a zero for alerts says YuviRed has not heard about anything, not that nothing happened. Look at the Endpoints card first.

6. Procedure: a five-minute morning check

Purpose. Know in five minutes whether anything needs your attention.

Prerequisites. A signed-in account. Computers already enrolled.

Steps.

1. Open **Dashboard**. Read the sidebar line and the **Endpoints** card. You should see most computers online. If many are offline, stop and read the Endpoints page first (section 7 below): an alert count is unreliable when computers are silent.
2. Look at **Open Alerts**. If it is red, click the card. You should land on the Alerts page showing open alerts. Sort with **Risk first** to see the worst first.
3. Look at **Shadow AI**. If it shows red "N unauthorized", click it to see which tools. If it is amber (review only), decide whether to classify those tools now or later (see the AIDR chapter).
4. Scan **Recent Alerts** for anything marked critical (red bar). Click one to read its Why section.
5. Check the top-bar badge. It should say **Metadata mode** unless you deliberately enabled content capture.

What you should see. A page with no red items, or a short list of alerts you can open directly.

How to confirm. Click **Refresh**. The numbers should not change much. If they jump, something is active right now: stay on the page and watch Recent Alerts for a minute.

Common problems.

- *Cards show "...".* The page is still loading. Wait a few seconds, then use **Refresh**.
- *A card shows "Unknown".* The server could not answer for that number. Click **Refresh**. If it persists, open **System Logs** (Admin) and look for errors from the dashboard.
- *"Couldn't load alerts" in Recent Alerts.* Click **Retry**. Nothing is hidden; the panel refuses to show a false all-clear.
- *A computer you expect is not in Active Endpoints.* That panel shows only the eight highest-risk computers. Open **Endpoints** for the full list.

7. Where to click next

You want to	Go to	Then
See why a number is red	The card's linked page	Read the filtered list
Understand one alert	Alerts, click the row	Read Who, What, Where, When, Why, Action
Decide whether an AI tool is allowed	AI Tools	Use the **Classify** drop-down on the row
See every computer and when it last checked in	Endpoints	Use the status filter: Online, Stale, Offline
See what a person typed that was flagged	Prompt Monitoring	Open an event
See which AI services are reached	Network	Pick 24h, 7 days or 30 days
See who uses AI and how	AI Map, User Behavior	Use the time range and the Activity or Detections mode
Hand leadership a document	Reports	See the Reports chapter
Check AI agents	Switch the sidebar to AMS	Start at Overview

The Endpoints page in brief

Endpoints ("Enrolled Endpoints") lists every enrolled computer. Filters: a search box (host, user or IP), operating system, status (Online, Stale, Offline), risk (Critical to Clear), agent version, and a sort (name, last seen, risk). Click a row to open the computer's detail with its tools, alerts and telemetry. A second panel, **Extension Changes**, lists installs, removals, enables, disables, version and permission changes of AI browser extensions (metadata only, not what the extension does).

Any signed-in console user can read endpoint identity and raw telemetry, as stated in the YuviRed data notice. Give console accounts only to people who should see that.

The AI Map and User Behavior in brief

- **AI Map** answers "who is using AI, which apps, and which providers are involved". Two modes: **Activity** (what is happening) and **Detections** (what was flagged). The filter bar has a search box, the time range, a grouping (none, risk, operating system, status) and a status filter (All, Approved, Monitored, Under review, Flagged / Alert). The bookmark button says "Saved views are not implemented yet"; that is accurate.
- **User Behavior** shows **Highest-Risk Users**, **Behavior Patterns** and **Top AI Users by Activity** (User, Hosts, AI Tools, AI Events, Dependency, Risk, Flag). Users are attributed from the account that owned the process plus network correlation. There is no directory lookup, so a shared or service account appears as reported (for example SYSTEM).

The AI Map. Computers appear as actors on the left, AI applications in the middle and providers on the right.

---

8. Limits to keep in mind

- The Dashboard shows what YuviRed's endpoint program reported. Computers without the program, and AI use on unmanaged devices, are invisible.
- Byte counts are not collected, so Data Egress reads "Not measured".
- The Active Endpoints, Recent Alerts and Detected AI Tools panels each show at most eight rows.
- Red does not mean something was blocked. YuviRed's AIDR side watches and alerts; it changes anything only in the narrow cases described in the AIDR chapter (Enforcement: a managed-browser policy on named Windows computers), the Sensitive Data chapter (a warning or block in the browser before a prompt is sent) and the AMS chapter (Enforce mode ending an AI agent program). It never blocks network traffic.
- Times in the console follow your computer's clock and time zone; the clock in the top bar shows your local time.

AI Detection and Response (AIDR)

AIDR (AI Detection and Response) is the part of YuviRed that finds AI use on your Windows computers, sorts it into "allowed" and "not allowed", raises alerts when something breaks your rules, and gives you a few narrow ways to respond. This chapter explains what is discovered and how, how the AI Tools catalog classifies it, how alerts are named and deduplicated, how to investigate, which response actions really exist, and which do not.

The example company is "Northwind Dental". All names, addresses and numbers are made up (computers such as NW-FRONTDESK-01, the user a.rivera, the address 192.0.2.10).

---

1. Monitoring versus enforcement, once and clearly

Read this section first. Everything else depends on it.

- **Monitoring** means YuviRed sees something and records it. It may raise an alert. Nothing on the computer changes. The person using the computer is not told. Almost everything in AIDR is monitoring.
- **Enforcement** means YuviRed changes what the computer allows. YuviRed has three, and only three, enforcement mechanisms in Lite:
    1. **Browser policy actions** on the **Enforcement** page: "Block browser access" to a website and "Disable extension", carried out by the YuviRed program on **Windows Chrome and Edge** by writing a managed-browser policy. Off until a person asks for one and a different Administrator approves it in the console (section 7.3 explains the one-Administrator case), and reversible. Section 7.
    2. **Prompt Protection** in the browser extension: warn or block a prompt before it is sent to a supported AI site. Off until an Administrator chooses a stronger policy than the default **Monitor Only**. See the Sensitive Data chapter.
    3. **AMS Enforce mode**: end an AI agent program when a rule says Block. See the AMS chapter. That is a different area of the console with its own switch.
- **Alert-only labels.** A tool classified **Blocked** or **Denied** is a label and an alert level. It does **not** stop the tool from running or the website from loading. The console says "alert only" next to these labels on purpose.
- **No network blocking.** YuviRed does not block network traffic, change your firewall, sinkhole DNS, intercept encrypted traffic or install certificates. To block at the network, you give your own firewall, proxy or DNS filter the **blocklist** YuviRed can export (section 8). That list is advisory: YuviRed publishes it, your device applies it.

If a screen or an email says "Alerted (alert-only, no enforcement)", that is the truth for that alert: you were told, nothing was stopped.

2. What AIDR discovers, and how

The YuviRed program on each Windows computer (the endpoint agent, version 1.8.2 in this release) collects **metadata** and reports to your server about every 60 seconds. It reports these kinds of facts:

What is discovered	How it is found	Where you see it
Browser AI use: visits to AI websites such as chatgpt.com	The computer reads the browser history locally every 5 minutes (default) and sends only counts and times per AI site, not the full history	AI Tools, Alerts, AI Map
Installed AI tools: AI programs and software installed	A scan of installed software and running programs against the AI catalog	AI Tools, Endpoints
Local AI services: a model server such as Ollama listening on the computer itself	Checking ports and service names that local AI tools use	AI Tools, Alerts (Local / self-hosted model)
AI providers by network activity: connections to AI services	Every program's web connections (ports 443 and 80) are matched against known AI provider addresses. The agent sees which program connected to which destination, not the content	Network, AI Map, Dashboard
AI browser extensions	Extensions in Chrome and Edge are read from the profile folders: ID, name, version, permissions	Endpoints (Extension Changes), AI Tools, Alerts
DNS cache	Windows only. The computer's resolver cache is read for AI service names	Network, AI Map
AI agent signals	Redacted evidence of AI agents (see the AMS chapter)	AMS

Honest limits of discovery:

- The agent reports **every** program's web connections on ports 443 and 80, not only AI ones, and the server keeps them as raw telemetry. Any signed-in console user can read endpoint identity and raw telemetry. This is stated in the data notice you accepted at installation.
- Browser history is read **on the computer**. Only AI-site aggregates are sent.
- It cannot see the **content** of what people type. Content is covered only by the browser extension and Prompt Protection (Sensitive Data chapter).
- A computer with no agent, a browser profile the agent cannot read, a phone, and a personal laptop are not seen.
- Detection depends on the catalog. A new AI product that is not in the catalog appears as an unknown or not at all until the catalog is updated, or until you add it (section 3).
- Firefox is not supported by the browser extension; Chrome and Edge are.

The Network page. Each row is an AI destination seen from your computers, with its provider and classification.

---

3. The AI Tools catalog and classification

3.1 What the page shows

Open **AI Tools** ("AI Tool Inventory"). Each row is one tool observed on your computers:

Column	Meaning
Tool	The name. A lock icon means an administrator or analyst set the classification by hand (a manual override).
Vendor	Who makes it, or "Unknown".
Category	For example a chat assistant, a coding assistant, a local model, an extension.
Effective Classification	The decision currently in force (see below).
Severity / Risk	How serious this tool is considered, as a word and a score.
Endpoints	How many computers it was seen on.
Last Seen	When it was last observed.
Classify	A drop-down to change the classification (Analyst and Admin).
Actions	Admin only: remove an override, block browser access, delete, merge.

Click a row to see which computers and users used the tool.

3.2 The six classifications

Every tool has exactly one classification. YuviRed ships six. Your terms may be "sanctioned", "unsanctioned" and "unknown"; here is how they line up:

Classification	In your words	Alert it raises	Default severity	Counts as Shadow AI
Approved	Sanctioned	None (silent)	Info	No
Monitored	Sanctioned, but watched	None (silent, recorded)	Low	No
Under Review	Unknown	Medium alert	Medium	**Yes**
Unauthorized	Unsanctioned	High alert	High	**Yes**
Denied	Unsanctioned, policy label deny	High alert	High	No (it is a decision you already made)
Blocked	Unsanctioned, policy label block	High alert	Critical	No

Remember section 1: **Denied** and **Blocked** are labels. They raise the alert and mark the tool, but they do not stop it from running. The console shows "(alert only)" beside them.

A tool never seen before and not in the catalog starts as **Under Review**. That is on purpose: YuviRed does not mark unknown AI use as known without a source.

3.3 Procedure: classify one tool

Purpose. Tell YuviRed whether a tool is allowed, so alerts match your rules.

Prerequisites. Analyst or Admin role.

Steps.

1. Open **AI Tools**.
2. Find the tool (use the **Filter** button in the top bar to search by name or vendor).
3. In the **Classify** column, open the drop-down on that row and choose a classification, for example **Approved** for a tool your practice has licensed.
4. You should see a green message such as "Classified" and the row's classification change. If the tool had open alerts that the new classification no longer raises, they are folded away; if the new classification raises alerts, they appear within one check-in cycle.

How to confirm. Reload the page (**Refresh**). The row shows the new classification and, if you changed it by hand, a lock icon.

Common problems.

- *The drop-down is missing.* A View Only account cannot classify.
- *The classification changed back.* An agent check-in or a catalog update can replace a catalog-set value. A lock icon (manual override) protects your decision from both. Use the button "Use catalog default" (Admin) to remove the lock deliberately.

3.4 Procedure: classify many tools at once (the vendor bundle)

Purpose. Approve or review every tool from one vendor in one step.

Steps.

1. On **AI Tools**, type the vendor name in the **Filter** search so only that vendor's tools remain.
2. Tick the box in the table header to select every visible row (or tick rows one by one).
3. A blue bar appears: "N tools selected". Choose a classification in **Classify as**, then press **Apply**.
4. Confirm the dialog. It tells you how many manually locked rows will be replaced.

What you should see. One toast, and every selected row updated. The change is one audited record listing each tool's before and after values.

3.5 The catalog (Admin)

Manage Catalog opens the catalog manager: the definitions YuviRed matches against (names, vendors, domains, process names, extension IDs, URL patterns) with their default classification, severity and action. Admins can add or edit entries and see a freshness report: entries not reviewed within the review cadence (Configuration, AI & Agent Monitoring, "AI Catalog Review Reminder": monthly, quarterly or twice a year) are marked overdue so detection content does not go stale silently. Editing the catalog changes what YuviRed recognizes; it is a separate act from classifying a tool you already see.

3.6 Other tool actions (Admin)

- **Merge** folds a duplicate into another tool, for example two spellings of the same product.
- **Delete** removes the tool record. If the tool is still being detected, it reappears on the next check-in. Alerts and telemetry keep their references.
- **Use catalog default** removes a manual override.

4. Detections and alerts

4.1 What raises an alert

An alert (also called a finding) is raised when what YuviRed sees breaks a rule. Each alert has a **category** from a fixed list, so alerts mean the same thing every time:

Category	In plain words
Observed AI usage	A known AI service was used. Usually Info or Low.
Unreviewed AI tool	An AI tool nobody has decided on yet (Shadow AI).
AI policy violation	An Unauthorized, Denied or Blocked tool or destination was used.
Local / self-hosted model	A local AI service is running and could handle sensitive data without oversight.
AI browser extension	An extension with AI capabilities was found.
Potential AI data egress	A large transfer to an AI destination (dormant in this release, byte counts are not measured).
Prompt-level categories	Sensitive data in AI content, prompt injection, PII, and others, raised by the browser extension. See the Sensitive Data chapter.
Rogue agent activity	A governed AI agent acted outside its approved boundary (Critical). Observed only.

Each category also has a **tactic** (such as Shadow AI, Policy Violation, Data Exposure) used for sorting.

4.2 Severity

Severity	How to treat it
Critical	Act now. A blocklisted destination, a policy violation on a server, or a governed AI agent outside its boundary.
High	A policy violation. Triage today. An Unauthorized or Blocked tool running, or a flagged network destination.
Medium	Classify it. A tool under review, a browser visit to a flagged tool, a local model, or an unknown extension.
Low / Info	Awareness. Known, sanctioned AI activity.

A single browser page-load of a flagged tool counts one level lower than the tool actually running, because a visit is weaker evidence. Computers whose operating system name contains "Server" raise policy-violation alerts one step higher (up to Critical).

4.3 Deduplication: one alert, not a hundred

YuviRed folds repeats. An alert is identified by the **computer plus a stable subject** (the AI tool, the extension, the destination, or the prompt rule and site), not by its title. While an alert is open, acknowledged or investigating, repeats of the same thing on the same computer are **folded into it**: the "Last seen" time moves and the "Days seen" count rises by one per distinct day. Several events on one day count once.

- If you **resolve** an alert and the same thing happens again, YuviRed creates a **new** alert marked "Recurred after resolve", pointing at the earlier one.
- Reclassifying a tool does not create a duplicate, because the title is not part of the identity.
- Findings from scoped rules set to **Suppress** appear as **suppressed**: visible under the Alerts status filter "suppressed", never counted as open, and never notified.

4.4 Reading the Alerts page

Open **Alerts** ("Alert Management"). Filters across the top:

- **Sort**: Recent activity, or Risk first.
- **All Severity**, **All Status** (open, acknowledged, investigating, resolved, false_positive, suppressed), **All Classifications**, **All Categories**, **All Tools**.
- **All Time** (Last 24 hours, Last 7 days, Last 30 days), by last activity.

- **Search alerts...**, **All Endpoints** and **All Users**.

New and repeating alerts appear on the page within about 2 seconds.

4.5 Procedure: triage one alert

Purpose. Understand an alert, decide, and leave a record.

Prerequisites. Analyst or Admin role to change anything (View Only can read).

Steps.

1. Click an alert. A panel opens with sections **Who** (the signed-in user; a service account may show SYSTEM), **Owner** (assign it to a team member, or press the person icon to take it yourself), **What** (tool, action, evidence), **Where** (computer, operating system, address, agent version), **When** (first seen, last seen, days seen), **Why** (severity, classification, category, reason, and an expandable explanation of what each severity means), **Action** (the disposition) and **Triage**.
2. Read **Evidence** and **Reason**. Evidence is metadata (domains, process names, masked samples), never a prompt.
3. Under **Triage**, press **Acknowledge** to say you are looking at it. Its status changes to acknowledged.
4. Set **Status** (open, acknowledged, investigating, resolved, false_positive), write **Analyst Notes**, optionally enter a **Ticket reference** such as JIRA:SEC-1234 and press **Save**.
5. If it needs a deeper look, press **Investigate** (section 5). Use **Send** to push the alert to a configured channel (Alerts and Integrations chapter).

What you should see. A "saved" message. The alert's status tag updates; the sidebar badge count changes.

How to confirm. Reopen the alert. Your note and status are there. The SLA line ("SLA due") shows the response target for that severity from Configuration, Detection & Response, "Alert response-time targets" (a tracking goal only, nothing happens automatically when it passes).

Common problems.

- *"This alert changed since you opened it."* Someone else edited it. Close it, reopen it and decide again.
- *You resolved it and it came back.* That is correct: the activity happened again. See "Recurred after resolve" and consider a classification change or a scoped rule instead.
- *A Prompt Protection alert lists several users.* Those alerts group events from every user of the computer; each event row records its own user.

5. Investigating

5.1 The pages that help

Page	Use it to
Alerts	Start from a signal and read its evidence.
Endpoints	Start from a computer: its tools, alerts, extensions, last check-in.
User Behavior	Start from a person (as the account that ran the program).
AI Map	See actors, apps and providers connected; switch between Activity and Detections.
Network	See which AI destinations were reached and export a destination list.
Prompt Monitoring	See flagged prompts and outputs and what happened to each.
Investigations	Keep one case file for several related alerts (below).

5.2 The Investigations page

Investigations lists "Active Incidents". Press the **+** button to create one by hand (title and severity), or press **Investigate** inside an alert to open one that already holds that alert and moves it to investigating. An investigation has a status, an assigned owner, a checklist, linked alerts, and an **Export evidence** button that downloads what is attached. A button jumps to the Blocklist to add an entry. Deleting an investigation is a trash-can button, and is final.

5.3 Procedure: follow one user across computers

1. Open **User Behavior** and find a.rivera in **Top AI Users by Activity**.
2. Note the **Hosts** and **AI Tools** columns. A user on several hosts with many tools may be worth a look.
3. Open **Alerts**, set **All Users** to a.rivera, and read the open ones.
4. If several alerts tell one story, open one and press **Investigate**, then link the others in the investigation.

What you should see. One investigation listing the alerts and your notes.

Common problems. Attribution is by the account that owned the process. A shared Windows account, or a service account, appears as that account. YuviRed has no directory lookup, so it cannot tell you which human sat at a shared login.

6. Policies (tuning detection)

Policies ("Security Policies") lists rules that tune detection. An Administrator can create:

- **Approved AI Tool**, **User Exception**, **Endpoint Exception** and **General Tuning** rules. A rule limits itself by AI tools, provider domains, surfaces, detection categories, users, endpoints and severity, has an expiry (never, 24 hours, 7 days, 30 days or custom) and needs a **Business justification**. The **Action** drop-down offers **Alert**, **Log only (recorded as Suppressed, not raised)** and **Alert-only (block requested, not enforced)**. In the page's own words, a rule either suppresses (records matching findings as monitor-only, still visible under the Alerts status filter "suppressed", never hidden) or escalates (its **Minimum severity** is a floor that can raise a finding but never lower one).
- **Detector Toggle (legacy)** rules, which switch auxiliary detectors on or off.

Critical findings, blocklist hits and browser-extension prompt findings can never be suppressed. All policies are alert-only: the page's banner says "Policies are alert-only: no traffic, process, or access is blocked", and the action "Alert-only (block requested, not enforced)" is a label. Nothing on this page writes to a computer; the only AIDR control that does is the Enforcement page (section 7).

7. Response actions that exist (Enforcement)

7.1 What it is

Enforcement (Analyst and Admin; hidden from View Only) is where an Administrator asks YuviRed to change browser behavior on named Windows computers. Two actions exist:

- **Block browser access (managed-browser URL policy)**: adds a website to the Chrome and Edge `URLBlocklist` policy on the chosen computers. People who open that site in a managed Chrome or Edge see the organization's block page.
- **Disable extension (managed-browser policy: Chrome/Edge, Windows)**: adds an extension ID to the `ExtensionInstallBlocklist` policy so that extension is blocked from being installed or enabled.

7.2 Before you use it: what it changes

> **This writes to the computer.** Enforcement adds a value to the browser's machine-wide policy in the Windows registry on each named computer. The effect is **append-only and reversible**: YuviRed never deletes extension files or edits browser profiles, and the rollback removes only the value it added. It applies on that computer's next check-in, and the browser usually needs a restart or a policy refresh to show it. It applies only to **Chrome and Edge** on **Windows**. Another policy source (a group policy already managing the same list) can interact with it. Other browsers, other devices and other programs are not affected.

7.3 Trial run versus real

Mode	What happens
Dry run first (recommended)	A preview. YuviRed shows which computers would be affected and changes **nothing**. Status shows "Preview ready".
Enforce now	A real change, but it **still needs approval** from an Administrator before it is applied. Until approved nothing is enforced.

Enforcement actions are **never automatic**: every real action is requested by a person, approved by an Administrator, and reversible. The approver must be a **different** person from the one who requested the action or who changed its target, scope or mode while it was pending; the server refuses the approval otherwise ("Separation of duties: the approver must differ from the requester").

A Lite installation with one Administrator. With a single Administrator account, the console can create dry runs and request enforce actions, but it **cannot approve them**: the **Approve** button sends the request as that same person and the server refuses it. The server does have a recorded exception for exactly this case (when only one active account can approve, the request may carry a "single administrator self-approval" flag with a reason, and the audit entry records it), but **this release's console offers no button for it**; it can only be sent through the YuviRed web interface (API) by a technical administrator or with support's help. The practical route is to create a second Administrator account under **Team / Users** and have that person approve. The moment a second approver exists, the exception switches off.

A fleet-wide Enforce is refused: real enforcement requires **Specific endpoints** (named computers). A dry run may cover all endpoints.

7.4 Procedure: block one website in the browsers of two computers

Purpose. Stop managed Chrome and Edge on NW-FRONTDESK-01 and NW-FRONTDESK-02 from opening an AI site your policy forbids.

Prerequisites.

- You are an Administrator or Analyst to request, and a different Administrator to approve (see 7.3 if your installation has only one Administrator).
- The two computers are online (check-in within 3 minutes) and run the YuviRed agent.
- You have decided that blocking the site is right. This stops people reaching it in those browsers; it does not stop other programs or browsers.

Steps.

1. Open **Enforcement** and press **New Action**.
2. **Action**: leave "Block browser access (managed-browser URL policy)". **Domain**: type `example-ai.test` (use the real domain). **Browser**: leave "All supported" or choose Edge or Chrome.
3. **Scope**: choose **Specific endpoints** and type `NW-FRONTDESK-01, NW-FRONTDESK-02`.
4. **Mode**: leave **Dry run first (recommended)**.
5. **Business justification**: write why, for example "Not approved by the practice owner". Required.
6. If the domain is on the protected list (large shared services such as microsoft.com or google.com), a red warning says "Blocking this could break enterprise services." You must tick "I accept the risk (required to proceed)". Think before ticking: blocking a shared domain can break Office, sign-in or updates.
7. Press **Create dry run**. You should see a new row with status "Preview ready" and a list of computers that would be affected.
8. Review it. When satisfied, use the row's workflow button **Enforce**. A confirmation says "It still needs approval before it applies; nothing is enforced yet." Confirm. A new enforce action appears, pending. (If you choose **Enforce now** in the New Action dialog instead, the button reads **Create enforce action**; the result is the same pending action, still unapplied until approved.)
9. A **different** Administrator opens **Enforcement**, opens that row and presses **Approve** (with a decision note), or **Reject**.

What you should see. After approval the action becomes active. On each named computer's next check-in (about a minute) it is applied and the result is reported back. The row shows "Enforced" and the computers that confirmed the change.

How to confirm. On a named computer, open the browser, go to `edge://policy` or `chrome://policy`, and press "Reload policies". Under the URLBlocklist policy you should see the domain. Browsing to the site shows the block page. In the console, the row's per-computer results say changed or not changed.

How to undo. Open the row and press **Unblock / Rollback**. Confirm. YuviRed removes only the value it added, on the computer's next check-in. For an extension action the button reads **Re-enable / Rollback**.

Common problems.

- *"Enforce actions require specific endpoints, fleet-wide enforce is not allowed."* Name the computers.
- *You cannot approve.* You are the requester or editor. Ask a different Administrator. If you are the only Administrator, see 7.3: create a second Administrator account.
- *A computer reports "not changed".* Possible reasons: it is not Windows, the agent has no administrator rights for the registry write, or the value was already present. The self-guarded design writes nothing when it cannot do it safely.
- *Nothing happens.* The computer is offline or stale. The action waits for its next check-in.

7.5 Other tabs and buttons

Tabs: All, Active, Dry Run, Enforced, Export Feed, Rolled Back, Failed, Disabled, Deleted/Archived. Row buttons include Approve, Reject, Cancel, Enforce, Retry, Place on export feed, Remove from export feed, Unblock / Rollback, Re-enable / Rollback. **Delete** archives a record: the record leaves the list, the audit history stays. **Delete permanently** removes the record and cannot be undone.

7.6 Honest limits

- Browser policy blocking is a Windows Chrome and Edge control. It has been run for real in lab testing, a limited number of times; it is not a general web filter and not a substitute for a firewall.
- A person who uses another browser, a phone, or a personal computer is not affected.
- YuviRed does not end processes, block applications or quarantine anything through AIDR. (AMS can end AI agent programs; see the AMS chapter.)

8. The blocklist and your own firewall (advisory)

The **Blocklist** page (Admin) holds manual entries (a domain, IP address, process or application, with a reason). YuviRed combines them with the domains of tools classified Denied, Unauthorized or Blocked into a **destination list**. You can:

- **Export** it from **Network**: **Plain domain list** (one domain per line, for Squid and generic proxies), **Palo Alto EDL**, **pfSense URL-table alias** or **Cisco Umbrella**.
- Publish it as a **Blocklist Feed URL** (Network page): a token-protected address your firewall polls on its own schedule. **Generate** a token, **Rotate** it, or **Revoke** it.

The page says it plainly: **"Discovery only: YuviRed does not block."** Your firewall, proxy or DNS filter does the blocking, if you configure it to use the list. A manual Blocklist entry matching a tool also raises a **Critical** alert when it is reached.

Procedure: add a manual entry. Open **Blocklist**, press **Add Entry**, choose **Type** (Domain, IP Address, Process, Application), type the **Entry** and a **Reason**, save. It appears in the next export and in the feed your firewall polls. To stop serving it, delete the entry.

9. Limits and honest notes

- AIDR is mostly monitoring. Do not describe a Blocked or Denied label to your staff as "the tool is blocked".
- Detection depends on the agent being installed, running and reporting. Check the Endpoints page.
- Nothing here inspects prompt content; that is the browser extension's job (Sensitive Data chapter).
- Classification is your decision. YuviRed's catalog gives a starting point, never a verdict.
- The catalog is not exhaustive. Use **Manage Catalog** to add what you find.
- Times are shown in your browser's local time.

Agent Monitoring System (AMS)

AMS (Agent Monitoring System) is the part of YuviRed that finds the AI agents running on your Windows computers, shows what is known about each one, applies your rules to them, and lets you step in: send a message, approve or deny one action, pause, hold or stop. An "AI agent" here is a program that acts on its own on a person's behalf, for example Claude Code writing and running commands. It is not a person and not a chat window in a browser (browser chat is covered by AIDR and Prompt Protection).

Everything below describes YuviRed Lite 1.0. Where a control has limits, the limits are in the same place. The example company is "Northwind Dental"; names, computers and numbers are made up.

---

1. Words used in the console

- **AI agent**: one kind of AI agent on one computer, for example "Claude Code on NW-DEV-01".
- **Instance**: one running copy of an agent (a process tree). Restarting an agent starts a new instance. YuviRed tells instances apart by process ID plus start time, so a reused process ID is never mistaken for the old one.
- **Session**: one conversation or task inside an agent. Sessions are visible only where a connector is turned on for that agent's tool.
- **Collector** (also "sensor"): YuviRed's own program on each computer. It is not an AI agent and AMS never stops it.
- **Connector**: a small YuviRed hook that runs inside an agent tool and asks the collector before each tool call. It is what makes messages, approvals, tool rules and task cancel possible.
- **Intervention**: a deliberate administrator action: message, approve or deny, cancel task, pause, hold or stop.
- **Monitor only / Enforce**: the policy mode. Section 6.

The AMS Overview. The tiles count running agents and the items that need a decision.

2. What is discovered, and how

2.1 Which AI agent tools

The catalog recognizes these by name (display names as the console shows them): Claude Code, Claude Desktop, Cursor, Windsurf, VS Code (AI agent), OpenAI Codex CLI, Gemini CLI, Aider, Open Interpreter, Goose, Continue CLI, local model hosts (Ollama, LM Studio, llama.cpp, Jan, GPT4All), MCP servers and MCP configurations (MCP is the Model Context Protocol that lets agents use tools), the AI browsers Perplexity Comet and ChatGPT Atlas, and automated browsers (browser automation). Anything else that behaves like an AI client appears as **Unrecognized AI client**.

2.2 How: three detection methods

Every agent shows its detection method:

- **Catalog**: a known tool recognized by its executable, install path or command line. A program is never matched on part of its name alone, so `claudette.exe` is not Claude Code.
- **Behavior**: an unknown program (not a browser, not YuviRed) that connects to an AI provider's API, to a local model port (11434 or 1234), or to a company AI gateway you list in Governance. It is shown as "Unrecognized AI client" with its evidence. It is a candidate for review, not proof of a rogue agent.
- **Connector**: an agent that registered itself through the YuviRed connector kit for custom agents (the SDK, software development kit).

2.3 Process signals

The collector checks running programs every 2 seconds and network connections every half second, and sends redacted command lines (secrets removed), the Windows account that owns the process, a file fingerprint (SHA-256) of the program, the parent and child programs, and results of stop and hold actions. A new or stopped agent reaches the console within a few seconds.

2.4 Configuration signals

About once a minute, at each check-in, the agent also looks for AI agent settings files in every user profile on the computer. It lists folder names two levels deep (at most 500 folders per profile; links and junctions are skipped) and checks the fixed locations of the tools it knows (Claude Desktop, Claude Code, VS Code, Cline, Cursor, Windsurf, Codex, Gemini CLI and Continue). It opens only files named `.mcp.json`, `mcp.json`, `claude_desktop_config.json`,

`.claude.json` or `config.toml`, plus those known settings files. From them it sends: the file's path, MCP server names, the transport, commands with secrets removed, the **names** (never the values) of environment variables and top-level settings, the model or provider, and a SHA-256 of the file. Of a credentials file it sends only its path, that it exists, and a fingerprint of its name, size and modification time, **never its contents**.

A computer that has a settings file but no running program shows the agent as "configured", never "active". A running program is judged on the time the **server** received the evidence, not the computer's own clock.

2.5 What AMS cannot see

Also listed under Coverage & Connections:

- Agents on computers without the YuviRed collector.
- Agents inside WSL (Windows Subsystem for Linux), containers or virtual machines, unless a collector runs inside them (section 10).
- Work an agent hands off to a provider's cloud after the local program finishes.
- A renamed program with no known marker that never talks to an AI endpoint the collector recognizes.

---

3. The five AMS areas

Switch the sidebar to **AMS**. Each area has a one-line purpose at the top and a **Glossary** button.

Area	Purpose
Overview	What is happening and what needs attention. Six measures: Running agents now, Agents needing a decision, Pending approvals, Computers not reporting, Policy pending, Coverage gaps. Click a measure to open the matching list. A value shows "Not known" when monitoring is not working, never a reassuring zero. Below: **AMS health**.
Agents (heading "AI Workforce")	Find an agent and open its workspace. Filter by state, computer, runtime, user, control status, severity and **Runs in**. Select several agents to stop them together.
Governance	The rules agents must follow and whether each computer has applied them.
Alerts & Actions	Findings, the approvals inbox, and every intervention with its verified result.
Coverage & Connections	What AMS can see and control on each computer, which connectors are on, the connection test, boundaries, and known limits.

"Running agents now" means confirmed running within the last 90 seconds.

Reports stay in the Reports area (Executive Summary and Auditor's Report). AMS data, interventions, approvals and messages appear in both.

3.1 The agent workspace

Open an agent from the roster. The header names the computer, runtime, user, session and instance, with first and last seen. A computer that stopped reporting shows "computer not reporting since ..." and never "Stopped" or "Healthy". Sections: **Overview** (what it is, where it runs, which rule applies and what is still unknown), **Activity**, **Chat**, **Controls** (each control shows its own readiness: available, or unavailable with the reason), **History** (commands, events, messages, approvals and audit entries).

3.2 The record kept for each agent

Owner, human manager, department, business purpose, assigned task and compliance status are declared by your team. The computer reports what it observes beside them. The **approval state** is one of Unreviewed, Under review, Approved, Restricted, Unauthorized, Referred to AIDR. "Needs attention" means unowned, unreviewed, under review or referred. Ownership can be transferred but not cleared. Changing a record needs the `ams.manage` permission.

4. Seats: 5 AI agents in Lite

Lite includes **5 AMS agents** (Pro 50). A seat is used by an AI agent that has **execution evidence**: a running process or network activity. A settings file alone does not take a seat ("Registered (configuration only, uses no seat)"). Seats go to the agents with the oldest execution evidence first. Seated agents never lose their seat automatically.

What happens when all 5 seats are used. The audit of the pre-release build found that new agents were invisible and their evidence discarded. **The release behavior is different, and this is what you will see in 1.0:**

- A sixth agent is **kept and shown** with all its evidence, labeled **"Detected, not registered (seat limit reached)"**, and the Agents header counts it ("N more detected, not registered (plan limit 5)").
- It is **not registered for governance**: it cannot be approved, owned or referred to AIDR while it has no seat. The console explains this when you try.
- **Restricting it never needs a seat.** You can still mark it Restricted, Unauthorized or Under review.
- A running instance with no registry record yet (still linking, or over the limit) appears under "Running, not yet in the registry", and you can still **Stop**, **Pause** or **Resume** it there.
- It takes a seat automatically on the next correlation cycle (about 2 minutes) after a seat frees.

To free a seat, **archive** an agent you no longer need (Agents, the agent's controls). Archiving hides it and frees its seat; if it runs again it returns as the same record with its earlier decision and takes a seat again if one is free. Pro raises the limit to 50 through a new license.

The final-installer checks of this behavior were not run in the lab loop for the release build (the audit recorded them as outstanding). The text above comes from the release code; confirm it on your own installation by using the **Detected, not registered** label as your check.

5. Connectors for supported tools

5.1 Which tools

Tool	Where YuviRed writes its hook
Claude Code	`C:\Program Files\ClaudeCode\managed-settings.json`
Codex CLI	`%ProgramData%\OpenAI\Codex\requirements.toml` (a delimited block)
Gemini CLI	`C:\ProgramData\gemini-cli\settings.json`
Cursor	`C:\ProgramData\Cursor\hooks.json`
GitHub Copilot CLI	`C:\ProgramData\GitHub\Copilot\policy.d\yuvired.json` (its own file)
Windsurf or Devin	`C:\ProgramData\Devin\hooks.json`

Every connector is **off by default** and is turned on per tool under **Governance**, Connectors.

5.2 What turning one on does (read before you do it)

> **This changes another program's settings on every enrolled computer.** When you set a connector to On and save, the collector on each computer adds YuviRed's hook to that tool's official managed settings. It keeps one copy of the original file and up to five rolling backups, adds or replaces only YuviRed's own entries, and on turning the connector Off removes only those entries. It refuses to touch a file it cannot parse, a path with a link in it, or a file not owned by SYSTEM, Administrators or TrustedInstaller. A JSON settings file is rewritten (content kept, formatting not). From then on, that tool asks YuviRed before each tool call.

5.3 What the hook does

The hook is one small script. Before each tool call it reads the tool's event, asks the collector on the same computer (a local address, 127.0.0.1 port 47811, never the internet), and prints the tool's own documented answer: no objection, deny with a reason, or wait for an approval. It never reads the agent's transcript or files. If the collector cannot be reached, it follows a small read-only copy of your policy (section 6.3).

5.4 What is sent and stored

- **Tool-call previews.** For a tool call that needs approval, was denied, or raised an alert, the hook sends a **redacted preview of at most 2 KiB** of the tool input, plus a fingerprint of the full input. Calls that a rule simply allows are only counted per session, not recorded.
- **Replies.** An agent's reply to a message you sent is kept, redacted, **at most 8 KiB**. A reply is sent **only when an operator message is waiting for one**. YuviRed does not read the agent's other conversation.
- **Storage.** AMS connector previews are opt-in (a connector is off until you turn it on) and are stored **unencrypted** in the YuviRed database on your server. Treat database access accordingly. Control events are kept 180 days, messages 30 days (section 12).
- The endpoint agent itself never reads prompts or documents. Content capture in the browser is a separate, off-by-default feature (Sensitive Data chapter).

5.5 What was verified

The AMS lab audit (coverage matrix, rows P01-C, P03, P04, P06-C and P22) records:

- **Claude Code** with the connector, in a live session: deny (13 of 13), approval bound to one call (12 of 12), task cancel (13 of 13) and an operator message that reached the model with its reply visible in the console (12 of 12).
- Agents built with the **YuviRed SDK**: message, approval, cancel and pause verified.
- **Codex CLI** and **Gemini CLI**: detected, monitored and governed; the hook passes the connection test; a live session through the hook was **not** tested (no provider login in the lab).
- **Cursor**: detected, monitored, governed and stopped in the lab; a live session through its hook (approvals, deny) is not recorded.
- **GitHub Copilot CLI** and **Windsurf**: the hook is written by the same code, but the audit has no row for them. Treat their behavior as unverified.

For every tool except Claude Code and the SDK, treat the connector as "installable and testable with **Run connection test**, not proven in a live session".

5.6 Procedure: turn on the Claude Code connector and test it

Purpose. Let YuviRed ask before Claude Code tool calls and allow messages and approvals.

Prerequisites. `ams.policy` permission (Administrator). Claude Code installed on at least one computer that reports to AMS. A person who agrees that Claude Code settings on that computer will be edited (section 5.2).

Steps.

1. Open **AMS**, then **Governance**. Find **Connectors**.
2. Set **Claude Code** to **On**.
3. Press **Preview the change**. Read which computers and agents are affected and any warnings.
4. Press **Save policy**. A dialog asks why; type a reason (at least 8 characters) and confirm.
5. Open **Coverage & Connections**. For the computer, find the Claude Code connector row. Wait for it to read **Installed**. Press **Run connection test** and read each step.
6. Start a Claude Code session on the computer. The row should change to "installed with a live session seen".

What you should see. A policy version number increases; each computer shows "v# applied" in Governance once it has pulled the new version (until then it keeps its previous one).

How to confirm. The connection test shows every step passing, including a hook run against the collector with a test event. A heartbeat alone is never taken as proof.

How to undo. Set Claude Code to **Off** and save. The collector removes only YuviRed's entries.

Common problems. An error state with a reason (for example the settings file does not parse, or the folder has a link in it): fix the file or folder, then wait. A failed write backs off 60 seconds, doubling to an hour.

The AMS Governance page. The mode and a rule for each AI agent runtime are set here.

6. Policy: Monitor only versus Enforce

6.1 The mode

Mode	What happens
Monitor only (the default on a new installation)	AMS records and alerts. Nothing is stopped by policy. A rule that would block shows "would have been stopped". An Ask rule is recorded as "would ask" and the call proceeds.
Enforce	The rules apply. **A matching Block rule ends the AI agent's process tree (children first) on every computer that has applied the policy, without asking anyone.** An Ask rule makes a tool call wait for approval.

> **Stop before you switch to Enforce.** Enforce can end programs that people are using for work. A person can lose unsaved work in an agent session. Start in Monitor only, read the "would have been stopped" events for a week, fix your rules, preview the change, and only then switch. Your own manual actions (stop, hold, pause, message) work in **both** modes and never change the policy or its mode.

6.2 The rules

Standard view

- **Mode**: Monitor only (default) or Enforce.
- **Rules by AI client**: for each known tool, **Allow** (no action), **Alert** (an alert when it runs) or **Block** (stopped under Enforce; alert plus "would have been stopped" under Monitor only). An agent marked **Unauthorized** in approvals is treated as Block; **Restricted** as Alert.
- **Unrecognized AI clients (behavior detection)**: what to do with unknown programs that behave like AI clients. Default Alert.
- **Company AI gateways**: addresses (host:port, one per line) whose traffic does not count as an unrecognized AI client.
- **Connectors**: Section 5.

Advanced

- **Scoped rules** for the whole organization, one runtime, one computer, one user or one agent. The most specific matching rule wins; at the same level Block beats Ask beats

Alert beats Allow. Tool rules (Allow, Ask, Deny, optionally only when the tool input contains given text) need a connector. Limits: maximum run time, maximum running copies per computer, tool calls per minute.

- **Where it runs**: limit a scoped rule to agents running on the computer itself, in a WSL distribution, a container, a virtual machine or an automated browser. Ticking none means anywhere.
- **Exceptions**: allow something a rule would stop, for a scope, with a reason and an expiry. Expired exceptions stop applying by themselves.
- An agent marked Unauthorized always wins over an allow rule or an exception.
- **Approval wait**: 10 to 90 seconds, 90 by default. The console refuses a value outside that range.

Every save runs a **preview** first: which computers, agents and running instances are affected, conflicts between rules, and warnings (for example an Ask rule for a tool with no connector, which cannot stop the action before it happens). Each save is a new version. A computer that cannot read a new policy keeps its last valid one.

6.3 When YuviRed cannot check a call

Under Enforce, a tool call from a tool with a connector is checked by the collector. When that check cannot happen:

- **Hook tools** (collector unreachable): a cached Deny or Ask rule that matches refuses the call. Any other call runs, unless the rule "A tool call YuviRed cannot check" is set to Block (default: Let it run).
- **SDK agents**: the call runs and the SDK logs a warning, unless the developer set `on_unreachable="deny"`.
- A tool that ends its hook while an approval is still pending runs the call without the approval. YuviRed sets each tool's hook timeout to the approval wait plus 30 seconds, except Windsurf, which documents no timeout setting.
- Under Monitor only nothing is refused, except an SDK call whose developer chose deny.

6.4 Procedure: preview and switch to Enforce

Purpose. Move from watching to stopping, on purpose.

Prerequisites. `ams.policy`. You have run in Monitor only long enough to trust your rules, and you have told the people who use AI agents.

Steps.

1. **AMS**, **Governance**. Check each rule under **Rules by AI client**. For tools you allow, set **Allow**. For tools you forbid, set **Block**.
2. Change **Mode** to **Enforce**.
3. Press **Preview the change**. Read the affected computers, running instances, conflicts and warnings.
4. Press **Save policy**. The dialog repeats: "Mode is Enforce: every computer that applies this version stops agents matching a Block rule, without asking." Type your reason and confirm.

What you should see. The policy version increases. "Applied on each computer" shows "Pending" for each computer until it reports the version applied, then "v# applied".

How to confirm. Start a blocked tool on a test computer. Within seconds the tool should end. Under Alerts & Actions, Interventions shows the result "Stopped (verified by the computer)".

How to undo. Set **Mode** back to **Monitor only** and save.

Common problems. "Someone else saved meanwhile": the save is checked against the version you loaded; nothing was overwritten. Reload and redo it. A computer that is offline keeps enforcing its old version.

7. Interventions

Control	What happens	Needs
Message agent	Sends a message (section 8).	Connector, `ams.message`
Approve / Deny	Decides one waiting tool call.	Connector, `ams.approve`
Cancel task	Cancels the agent's current task through its connector. The program keeps running.	Connector, `ams.stop`
Pause / Resume	Suspends the agent's whole process tree, then resumes it. Resume is refused while a rule blocks or holds the agent.	Collector, `ams.stop`
Hold / Release	Stops the agent and keeps stopping it if it restarts, until you release it. Needs a known program identity.	Collector, `ams.stop`
Stop now	Ends the agent's process tree (children first) and checks they are gone.	Collector, `ams.stop`
Stop selected	Stops the agents you selected; you confirm the exact count; one result per agent.	Collector, `ams.bulk`
Stop container, Stop WSL distro, Close automated browser	Takes down a whole boundary (section 10).	Collector that can reach it

> **Before you press Stop now, Hold, Pause or Stop selected.** These end or freeze programs on a person's computer. **Stop** ends the agent and everything it started, which can lose unsaved work. **Pause** freezes the programs until you press Resume. **Hold** is stronger: it keeps ending the agent every time it starts, until you release it. None of them can undo work the agent already handed to a provider's cloud. YuviRed never stops its own collector, YuviRed's server or Windows system processes.

Every intervention asks for a reason, is audited, and shows **Requested** until the computer reports back ("the computer has not confirmed yet"). Only a computer-verified result reads as done. For Stop the answers are: "Stopped (verified by the computer)", "Partly stopped" (with what survived), "Not found", "Identity mismatch" (the process was replaced; nothing was

touched), "Refused (protected)", "Expired" or "Failed". A command not delivered within 5 minutes **expires and is never carried out later**.

The confirmation for Stop does not list the processes it will end; the stop ends the whole descendant tree. The Hold command appears in the console at the next 30-second refresh.

7.1 Procedure: stop one agent

Purpose. End an agent that should not be running.

Prerequisites. `ams.stop`. The computer is reporting.

Steps.

1. **AMS**, **Agents**. Find the agent (filter by computer or state).
2. Open it, go to **Controls**, and read the readiness of **Stop now**. If it is unavailable, the reason is written there (computer not reporting, missing permission, and so on).
3. Press **Stop now**. Read the confirmation, type a reason, confirm.

What you should see. The command reads **Requested**, then within seconds "Stopped (verified by the computer)". Under **Alerts & Actions** the intervention appears with the same result.

How to confirm. The agent's state changes and "Running agents now" on the Overview drops by one.

Common problems. "Identity mismatch": the program was replaced after you looked, so nothing was touched; reopen and try again. "Partly stopped": a child survived; the result names it. "Expired": the computer did not pick it up within 5 minutes; check the computer is reporting.

---

8. Messages and the Chat fallback

The **Chat** section sends a message to the selected agent session through its connector. Before you send, the console names the recipient session. Every message shows its delivery state: queued, sent, at the computer, accepted by the agent's inbox (not confirmed), handed to the agent, then the agent's reply, or failed, expired or unknown. A message is **never resent automatically**.

Which agents can receive a message. Message is offered only for runtimes whose hook actually delivers it: **Claude Code** and agents built with the YuviRed SDK. For Codex CLI, Gemini CLI, Cursor, GitHub Copilot and Windsurf, the hook does not deliver messages, so

Message is shown as unavailable with that reason. A message is information to the agent: it grants no permission, and typing "stop" in chat stops nothing. Use **Stop now**.

The Chat fallback note. If an agent has no connector session, or the workspace cannot load, Chat shows "Messaging: Not available" with the reason, and a summary derived from telemetry marked "not agent speech". This note is a statement about what YuviRed can do for that agent right now, not a fault.

---

9. Approvals

When an agent with a connector attempts a tool call that an **Ask** rule covers (under Enforce), the call waits and an approval appears in **Alerts & Actions** and in the agent's **Controls**: the exact tool, a redacted preview of its input, the rule, and when it expires.

Procedure. Open **Alerts & Actions**, find the pending approval (the amber badge on the menu counts them), read the preview, and press **Approve** or **Deny**, with a reason. Needs `ams.approve`.

- An approval is valid for **that one call**. The same call later needs a new approval; a changed input needs a new approval.
- If nobody decides before it expires, the call is **denied**.
- The decision reaches the computer on its next sync.

---

10. Boundaries: WSL, virtual machines, containers and automated browsers

An agent does not always run directly on the computer. It can run in a WSL distribution, a container, a virtual machine, or be driven by an automated browser (a script controlling Chrome or Edge).

What Lite can see.

- The Windows computer with its collector, and every agent that runs on it directly.
- That a WSL distribution, container, virtual machine or automated browser **exists and is running** on that computer. **Coverage & Connections**, Boundaries lists each one with its name, runtime, whether it is running, whether a collector is inside (present, absent or unknown), how many agents run in it, and when it was last seen. The agent list shows "Runs in WSL distro Ubuntu", "Runs in container web-1", "Runs in virtual machine ..." or "Runs in automated browser chrome". A **Runs in** filter narrows the list.

- A **Windows virtual machine** that you enroll as its own computer (Add Computer). It then links to its host by the machine's hardware identity and is seen as a full computer, shown as "Inside VM X on computer Y".
- An automated browser is attributed to the agent that drives it, seen from outside.

What Lite cannot see.

- **Agents inside a WSL distribution or a container.** Lite offers no Linux collector. The console says so on a red row with the next step: the Windows computer is monitored by its own collector, agents inside the distribution or container are not seen, and the whole distribution or container can be stopped from the host.
- **Agents inside a virtual machine that is not Windows.** Lite offers no installer for other guest systems.
- Work a local agent hands off to a provider's cloud.

A boundary with no collector is a **coverage gap**. These rows are never hidden. A computer that cannot read the Windows event feed is shown in **Detection feed health**: "Reduced visibility" means short-lived programs or connections can be missed.

Stopping a boundary is a separate control. It is named for what it does: **Stop container**, **Stop WSL distro**, **Close automated browser**. It is never what Stop on an agent does.

> **Before you stop a boundary.** Stopping a boundary takes down **everything inside it**, including programs that are not AI agents, and can lose work. The confirmation names the boundary, the computer, how many agents are inside and that other programs there also stop, and asks for a reason. Stopping a virtual machine is never offered. A stopped boundary is **not kept stopped**: it can be started again.

If a boundary button is grey, it says why: the collector cannot reach into it, the computer is not reporting, the kind is not supported, you lack `ams.stop` (or `ams.bulk` for several at once), or it is not running. After you confirm, each boundary reads **Requested** or **Refused** with a reason; nothing is shown as stopped until the computer reports back.

---

11. Alerts and findings in AMS

Alerts & Actions lists AMS findings with their evidence, the rule that caused them and links to the agent and its commands. You can **acknowledge** or **assign** a finding, and **resolve** it as *verified outcome* (only possible when a computer-confirmed result exists), *no action needed*

or *false positive*, each with a reason. AMS can also send a finding to AIDR as an investigation (Referred to AIDR). Two optional detections ship off or report-only:

- **Behavioral Detection** (Configuration, AI & Agent Monitoring): learns an agent's normal behavior and flags large departures. **Report only in this version**: it blocks nothing. A "Detection mode" choice "Active response" is shown as not available.
- **Rogue Agent Activity Detection**: scores governed agents against their approved boundary and raises a **Critical** alert when one steps outside it. Ships disabled. Observed only: it blocks nothing. It evaluates agents that have been through a governance decision, so a brand-new unreviewed agent is not evaluated until someone reviews it.

---

12. How long AMS keeps records

Set under Configuration (the retention settings). 0 keeps forever. AMS applies the first three itself about once an hour, whether or not Automatic Data Cleanup is on.

Records	Kept for
Control events from computers; decided or expired approvals (a pending approval is never deleted)	180 days
Messages with AI agents, both directions	30 days
AI agent instances that ended, with their ended connector sessions (running instances are never deleted)	90 days after they ended
Agent lifecycle changes and MCP server set changes (governance decisions are always kept)	365 days
Risk-score history (each agent's newest score is always kept)	180 days

AMS audit entries (who changed what, with before and after values) are part of the audit log, which is kept until purged.

13. Permissions

Permission	Allows
`ams.read`	view everything in AMS
`ams.message`	send messages to agents
`ams.approve`	approve or deny tool calls
`ams.policy`	change rules and connectors, run connection tests
`ams.stop`	stop, hold, pause, resume, cancel task, stop a boundary
`ams.bulk`	act on several selected agents or boundaries at once
`ams.manage`	edit agent records (owner, purpose, approval state)

Administrators hold all of them; analysts and View Only accounts hold `ams.read`. The server checks every action; a greyed button names the permission it needs.

14. Honest limits

- Windows computers only. Lite has no Linux or macOS collector.
- AMS stops programs and freezes them. It cannot undo what an agent already did, and it cannot stop work an agent handed to a provider's cloud.
- A message reaches only Claude Code and SDK agents. For the other five connector tools the console shows Message as unavailable, and a message that cannot be delivered is marked Failed or Expired and never resent.
- Live approvals and messages were proven in the lab for Claude Code and the SDK only (section 5.5).
- Agents inside WSL, containers and non-Windows virtual machines are not seen; only the boundary can be stopped.
- The workspace does not refresh by itself; press Refresh to see a new state.
- The Overview coverage-gap count does not include boundary gaps until you have opened Coverage & Connections once.
- A seat limit leaves extra agents unregistered (section 4); restricting them still works.
- "Healthy" in AMS means recent evidence, not "safe".

Sensitive information in prompts (Prompt Monitor and Prompt Protection)

The YuviRed browser extension, called the AI Prompt Monitor, looks at what a person types, pastes or attaches on supported AI websites **in the browser, before it is sent**. It can stay silent, warn the person and ask for a decision, or block the submission. This chapter explains what it detects, how you set the policy, what people see, how you review what happened, how raw content capture works and is audited, and what the extension does not do.

The example company is "Northwind Dental". Names, computers and numbers are made up.

---

1. The three separate ideas

- **Detection** finds sensitive information and says how sure it is. A checksum-valid payment card is high confidence. A bare label such as "password:" is medium. A sentence that only sounds clinical is low.
- **Policy** is yours. It decides what happens to each kind of finding and to content that could not be inspected. There is no hidden minimum: **Monitor Only** really shows nothing to the person typing.
- **Enforcement** is what actually happened to the submission. It is recorded truthfully in Prompt Monitoring: observed, warning cancelled, allowed after approval, blocked, allowed with incomplete inspection, or inspection unavailable.

Alerts follow findings, not outcomes. Every finding opens an alert, or joins an open alert for the same finding, at the rule's fixed severity. That is true under Monitor Only, and also when the person chose "Go back and edit" or the prompt was blocked and nothing was sent.

> **Prompt Protection is a control for users, not against the owner of the computer.** A person who administers the Windows computer can uninstall or change the software. On Lite the Administrator is the account created at first run; there is no separate "policy administrator" role.

---

2. What is detected: rule families

Detection rules are grouped into **categories** that you see in the policy editor, and into four **classes** that the presets act on.

2.1 Categories

Category (as named in the policy)	Examples
Passwords and authentication credentials	A labelled password ("password: ...", `DB_PASSWORD=...`, `-p...` on a command line), even a weak one
API keys, access tokens, session credentials and OAuth secrets	Provider keys, `Authorization` headers, live session tokens, connection strings with passwords
Private keys and recovery material	Private keys with key material, recovery codes
Government identifiers and sensitive personal records	Validated Social Security, National Insurance, SIN and Aadhaar numbers
Payment-card, banking and financial information	Luhn-valid cards with card context, IBANs, checksum-valid routing and account pairs
Employee, customer and health records	Personal-data tables, record dumps, a clinical detail tied to a patient identifier
Contact details and other lower-risk personal information	Phone numbers, email addresses, labelled dates of birth
Internal infrastructure and confidential business information	CONFIDENTIAL or internal-only markers, internal hostnames, roster schemas
Company-defined terms and internal domains	Words and hostnames that you list yourself
Medical, mental-health and care records; Children, students and education records; Special-category personal data; Biometric, genetic and identity-verification material; PINs, one-time codes	More specialized families, each with its own detectors

Category (as named in the policy)	Examples
and security answers; Legal, privileged and investigative material; Government and defence handling markings; Access-bearing links and invitations; Location, movement, travel and vehicle details	
Behavioural categories	Prompt-injection and similar behaviours: **always recorded, never blocked**

The exact rule-by-rule tables, confidence values and masking behavior are in the generated policy reference that ships with the product (`docs/PROMPT_PROTECTION_POLICY_REFERENCE.md` in the source tree, produced from the code). Detection is broad, not perfect.

2.2 The four classes that presets act on

Class	What it holds
Credentials	API keys, tokens, private keys, session secrets, any labelled password. An access key ID alone is borderline; it becomes a credential when its secret key appears within about 320 characters.
High-risk personal and financial	Validated government numbers, IBANs, validated cards with card context, name plus date of birth plus identifier, a clinical detail tied to a patient identifier.
Confidential business	Personal-data tables, confidentiality markers, roster schemas, your own company terms and internal domains.
Borderline	Contact details, labelled dates of birth, passport or driving license numbers with a label, internal hostnames, weak "secret =" assignments, expired tokens, a card number with no card context, a private key header without the key.

3. The policy: observe, warn or block, per rule

Open **Configuration**, then **Prompt Protection**. Only Administrators can change the policy; others see it read-only (the server refuses writes from anyone else, whatever the page shows).

3.1 Presets

A new installation starts in **Monitor Only**. A stronger preset is never applied by itself.

Preset	Credentials	High-risk personal and financial	Confidential business	Borderline	Cannot be inspected
Monitor Only	observe	observe	observe	observe	allow, record
Light Warnings	warn	warn	observe	observe	allow, record
Warnings Only	warn	warn	warn	warn	warn
Balanced	**block**	warn	warn	warn	warn
Strong (recommended)	**block**	**block**	warn	warn	block
Strict	**block**	**block**	**block**	warn	block
Maximum	**block**	**block**	**block**	**block**	block

The three actions:

- **Observe**: recorded silently; the person sees nothing.
- **Warn**: the person must approve explicitly before the submission is sent.
- **Block**: the submission is stopped before it leaves the browser, with no override in the page.

A prompt that is only too long to check warns at every preset except Monitor Only; no preset blocks it for its length alone.

3.2 Sensitivity settings (Advanced, "Custom")

Any change turns the policy into **Custom**. You can set:

- **Classes**: each of the four columns above to observe, warn or block.
- **Categories**: each administrator category to *inherit* or an explicit action. A category setting overrides its class.
- **Rule overrides**: a single detector's action.

- **Minimum confidence**: the lowest confidence that may interrupt a person. Below it, a finding is recorded but never interrupts. Every preset uses **medium**; Maximum uses **low**.
- **Options**: *contact: bulk only* (a single labelled phone number or date of birth is observe-only; lists still act) and *card requires context* (a valid number with no card wording is observe-only).
- **Uninspectable content**: six separate answers (allow, warn or block) for an unreadable attachment, an oversized attachment or form, encrypted or password-protected content, a scanner timeout or failure, a policy that could not be synchronized, and a prompt too long to check fully. Uninspected content is **never recorded as "safe"**.
- **Floors**: minimums that scopes and exceptions can never go below (for example credentials: block everywhere).
- **Scopes**: different settings per AI website, matched on the **exact** hostname. A look-alike domain never inherits a scope. Priorities must be unique; the first match wins. Lite has no device groups, departments or directory.
- **Exceptions**: scoped, accountable relaxations with an owner, a reason, what they cover, an action they lower to (observe or warn) and an expiry of at most 365 days, or a deliberate "permanent" choice. A value-exception cannot name a password, key or identifier.
- **Company terms and internal domains**: your own words, phrases and hostnames to be detected as confidential business material (literal text, no regular expressions).

> **The supported AI sites can read your company terms.** The check runs inside the AI site's own page, so that page can read the terms, internal domains, value-exception values and organization entries you add. Do not list anything that must stay hidden from those sites.

3.3 Procedure: choose a preset and apply it

Purpose. Move from watching to warning, or warning to blocking, on purpose.

Prerequisites. Administrator. The extension is installed on at least one managed browser and shows **Protected** on the Browser protection card.

Steps.

1. Go to **Configuration**, open **Prompt Protection**.
2. Read the active preset and revision.

3. Choose a preset card (for example **Balanced**). This creates an unsaved **draft**. The page shows "Unsaved draft".
4. Press **Test this policy...** to try it (section 3.4).
5. Press **Save / Apply...**. A summary lists the impact. Anything that becomes **more permissive** is listed at the top. Read it, then confirm.

What you should see. A new numbered **revision** becomes active. Prompt Monitoring shows the new revision. The Browser protection card shows each browser as **applied**, **pending** or **stale**.

How to confirm. A successful save is not proof that every browser applied it. Browsers pull the policy about once a minute. On a test computer, send a harmless made-up test prompt such as `password: Examplepass-123` to a supported AI site and see whether the dialog in section 4 appears.

How to undo. Open **Revision history**, find the earlier revision, press **Restore**. Restoring creates a **new** revision; numbers never roll back. If another Administrator saved in the meantime, the restore is refused: reload and try again.

Common problems.

- *Browser shows Stale or Protection unavailable.* It has not reported in 10 minutes, has no consent, or could not pull the policy.
- *Monitoring only.* Shown for a healthy browser under Monitor Only, with the reason "active policy is Monitor Only: warnings and blocking disabled". That is correct, not a fault.
- *A very large policy is refused* (over 700 KiB as delivered). Remove entries.

3.4 The policy preview in the console (64 KiB limit)

Test this policy... opens a dialog that runs your **draft** against built-in synthetic samples, or text you paste, using the same decision engine the browser uses.

1. Press **Test this policy...**.
2. Pick a built-in sample (**Ordinary question**, **Credential**, **Personal record**, **Internal business**, **Contact details**, **Uninspectable attachment**, **Too long to check fully**), or type or paste your own text in the box. Choose the destination site.
3. Run it. The result table shows, for the active policy and for your draft: the category, confidence, current action, proposed action and the reason chain.

The pasted sample is limited to **65,536 characters (64 KiB)**; longer text is refused. Use **made-up test data only**. The sample is inspected on your YuviRed server only. It is never

sent to an AI provider and is not kept after the test.

3.5 Export, import and revisions

Export saves the policy as a JSON file (no credentials in it). **Import...** loads one, checks every rule, value, destination and expiry, and refuses a document with any problem (listing them) rather than dropping parts of it. **Duplicate as Custom...** saves a named copy. Policy changes cancel pending warning approvals: a relaxed policy never sends a draft or a cancelled prompt by itself, a person must press Send again.

The Prompt Protection section. The active preset is marked; the buttons at the bottom open the policy test.

The Test this policy dialog. It compares what the current policy does with what a draft would do, and saves nothing.

4. What the person sees

Situation	What appears
Observe	Nothing. The finding is recorded.
Warn	A dialog: "Potentially sensitive information detected", the categories and masked evidence, and "Are you sure you want to send this information to {service}?" with two buttons: **Go back and edit** (the default) and **Send anyway**. Pressing Escape, closing the dialog or clicking elsewhere never sends. One approval releases exactly one submission; editing, changing attachments or destination, or a policy change needs a fresh decision.
Block	"Prompt blocked: sensitive information detected", what must be removed, and **Go back and edit** or **Close**. There is no override in the page.
Draft notice	A small dismissible notice while typing, when the draft would be warned about or blocked. It never sends anything. It is silent under Monitor Only. The finding is recorded when the notice appears, even if the person then edits the text out.
Inspection unavailable	When the policy says to hold content that could not be inspected.
A ref	Warning and block dialogs end with a small "ref" and an id. The same id is the **Correlation id (dialog ref)** in the event's Guard decision panel and the **Submission id** in the Auditor's Report, so a ref a person reports leads straight to that submission.

Evidence shown to the person is masked: at most four leading and four trailing characters of a provider token, and **no characters** of a password or free-form secret.

Where the person can reduce protection. On the extension's options page a person can remove a site from "Monitored sites" or set "Capture level" to "Disabled". A removed site is neither checked nor recorded; at Disabled the browser still warns and blocks but reports no findings. The console is not told about either setting, so that browser keeps reading Protected. A managed browser policy that sets these fields overrides the options page. YuviRed Lite writes no browser policy for this.

5. Capture levels: what is kept about a finding

Set under **Configuration**, **Detection & Response**:

Setting	Choices
Prompt Content Capture (the most content YuviRed may keep)	**Off: no capture at all**; **Activity only: no content** (metadata only); **Redacted excerpts (Recommended)**; **Full content**
Inspection Depth	**Activity metadata only (Recommended)**; **Full content inspection**
Keep Original Prompt Samples	A switch. Requires Full content inspection.

What each level means:

- **Metadata only**: records that AI was used, where and by whom, with the rule and confidence, never any text.
- **Redacted excerpts (the default)**: a short masked excerpt of each finding. Keys and tokens keep at most their first and last 4 characters; identity and card numbers at most their last 2; passwords none of their own characters. The extension stores this masked fragment by default.
- **Full content**: the full text of a matched prompt or response may be stored. See section 7.

The top-bar badge shows **Metadata mode** (blue) or **Full inspection** (red, with "(raw capture ON)" when samples are being stored). Endpoints can capture less than the console setting, never more.

---

6. Reviewing what happened: the Prompt Monitoring page

Open **Prompt Monitoring** (Monitor section). The top row holds:

- The **Browser protection** card: per browser, **Protected**, **Stale** (no report in the last 10 minutes), **Protection unavailable** (no consent, or the last policy pull failed) or **Monitoring only** (a healthy browser under Monitor Only), each with its reason, with whether its policy is applied, pending or stale, and onboarding steps while no browser is protected.
- Category cards (Prompt injection / jailbreak, Sensitive data in AI content, PII in AI content, Malware generation, Phishing / social engineering, Exploit development, Data exfiltration intent, Policy bypass, Sensitive company data, Suspicious code generation,

Unauthorized automation, AI agent misuse), each with a count of prompts and outputs and an **Alerts** button.

Below is **Suspicious content events**. Filters: **All types** (Prompt, Output), **All categories**, **All severities**, **All outcomes** (for example "Sensitive data blocked", "Sensitive data warning (any)", "warning pending", "warning cancelled", "warning approved (sent anyway)", "draft flagged", "Allowed by company policy; inspection incomplete", "Allowed by exception", "Inspection unavailable", "Enforcement unavailable", "Observed only"), **Provider (origin)**, **User**, and a window of **Last 1d, 7d, 30d or 90d**.

Columns: ID (a display id such as SR-EVT-000456, click to copy), When, Type, Severity, Outcome, Category, Rule, Confidence, Endpoint and user, Provider, and the **Matched excerpt** (a lock icon means full content was captured). At most 200 events are listed with "Showing N of M events".

Click a row for the detail: rule, provider, computer, user, browser, time, the **Guard decision** panel (what the policy intended, what actually happened, when the person answered a warning, the policy revision, scope or exception that applied, what could not be inspected), the redacted excerpt, the linked alert and **Open alert & triage**.

Outcome wording, in plain words

Outcome	What it means
Observed only	Recorded; the active policy does not warn or block for this finding.
Sensitive data warning, pending	The person was warned and has not decided; the submission is held in the browser.
Warning cancelled	The person chose Go back and edit. Nothing was sent.
Warning approved (sent anyway)	The person confirmed; that text was released once. This records permission to send, not delivery.
Sensitive data blocked	The extension cancelled the submission before it left the browser.
Draft flagged	The person saw an inline notice while typing. Nothing was sent at that point.
Allowed by company policy; inspection incomplete	Part could not be inspected and the policy allowed it. Never labelled safe.
Inspection unavailable	Inspection could not complete and the policy holds such submissions.
Enforcement unavailable	The enforcing part could not act; the submission may have left the browser.

Exports

An Analyst or Administrator sees **Export CSV** on the page; it exports the events matching the filters, and every exported row carries its alert id. The **Reports** area also holds a **Prompt protection** section (Reports chapter) that counts prompts and findings in a period. Reports never include prompt text.

Procedure: review a blocked prompt

1. **Prompt Monitoring**. Set **All outcomes** to **Sensitive data blocked** and a window.
2. Click a row. Read **Rule**, **Provider**, **Endpoint**, **User** and the **Guard decision** panel.
3. Press **Open alert & triage** to record a note and status.

What you should see. Outcome "Sensitive data blocked: blocked before transmission" and an alert for the same finding.

Common problems. *No events at all.* The extension may not be installed or consented to; check the Browser protection card. *An event with Full content level but no content.* The extension reported without content, or the sample was purged by the retention window.

The Prompt Monitoring page. In this lab capture no browser had reported yet, so the status says so.

---

7. Raw capture: enabling it, and the audit

Raw capture is off by default. It exists for investigations that genuinely need the original text. Enabling it stores prompt content on your server. Read this whole section first.

> **What turning it on means.** For matched events, up to the **first 20,000 characters** of the prompt (or AI response) are stored on your YuviRed server, **encrypted**, and erased automatically after a short window (**24 hours** by default; choices are 4 hours, 24 hours or 3 days). Only Administrators can view them, and **every view is written to the audit log** (`dlp.content_event_viewed`). Passwords, keys and tokens that YuviRed detects are shown as [REDACTED]. Content that people typed is personal data and may be confidential; tell your staff before you enable it, and keep the window short.

7.1 Procedure: enable raw capture

Purpose. Keep original text of matched prompts for a limited time.

Prerequisites. Administrator. A reason you can state. Connections to the server over encrypted HTTPS (capture over plain HTTP is allowed only in a lab by a separate, clearly labelled setting).

Steps.

1. **Configuration**, **Detection & Response**.
2. Set **Inspection Depth** to **Full content inspection**.
3. Set **Prompt Content Capture** to **Full content**. Confirm the dialog "Capture full prompt content? Complete prompt and output text will be stored for matched events instead of redacted excerpts."
4. Turn on **Keep Original Prompt Samples**.

5. Under **Data & Retention**, set **Original prompt samples** to the shortest window you can live with.
6. Optional (Advanced): set **Prompt samples: endpoint default**, and tag lists **Always keep samples from endpoints tagged** and **Never keep samples from endpoints tagged** (if both match, "never" wins).
7. Save.

What you should see. The top-bar badge turns red: **Full inspection (raw capture ON)**.

How to confirm. Trigger a test finding with made-up data. In **Prompt Monitoring** open the event: the lock icon is present and an Administrator sees **View captured prompt (audited)**.

How to undo. Turn off **Keep Original Prompt Samples**. Samples already stored are erased when their window expires. Set **Prompt Content Capture** back to **Redacted excerpts (Recommended)** and **Inspection Depth** to **Activity metadata only**.

7.2 How to view a captured prompt and where it is logged

1. Open the event in **Prompt Monitoring**.
2. Press **View captured prompt (audited)**. The text appears below.
3. The view is recorded in the audit log with who viewed it and when (Alerts and Integrations chapter, "The audit log").

An Analyst sees only that content exists, not the content.

Common problems. *No content for an event although capture is on.* The extension's own capture level may be lower, the content may be older than the window, or the prompt did not match a rule.

---

8. Limits: what it covers and what it does not

Sites. The extension is installed by default for **15 hostnames** that make up **13 AI services**: ChatGPT (`chatgpt.com`, `chat.openai.com`), Claude (`claude.ai`), Gemini (`gemini.google.com`), Microsoft Copilot consumer (`copilot.microsoft.com`), Perplexity (`perplexity.ai`), Microsoft 365 Copilot for work accounts (`m365.cloud.microsoft`, `copilot.cloud.microsoft`), DeepSeek (`chat.deepseek.com`), Grok (`grok.com`), Mistral Le Chat (`chat.mistral.ai`), Meta AI (`meta.ai`), Poe (`poe.com`), Google AI Studio (`aistudio.google.com`) and HuggingChat (`huggingface.co`, chat pages only).

- **Supported**: ChatGPT, Claude, Gemini, Microsoft Copilot (consumer) and Perplexity. Every site uses the same two checks (the typing box and the page's own requests), and

these five are the ones the extension's end-to-end test harness exercises against reproductions of each site's request shape (ChatGPT, Gemini and Copilot shapes among them; the 2026-10-27 run passed all 181 scenarios). The lab rule for this release was "no live AI site, no real account", so **no run against the real sites is recorded** for any of the thirteen; test on your own computers before relying on a block.

- **Added, not yet verified on the live site**: the other eight. The extension is injected and a reproduction of each site's request shape is inspected and blocked as expected on a local test, but nobody has run it on the real site. Treat them as **best effort**.
- Frames from another host (for example Copilot Chat inside a page from a different host) and Grok on x.com are **not** inspected. Google AI Studio sends what you type in the background before you press Run, so the dialog can appear while you type.

Browsers. Chrome and Edge only. Firefox is not supported. **In this release nothing installs the extension into Chrome or Edge automatically.** The console setting "Remote Browser Extension Wiring" exists but is inert on a Lite server (the Lite service does not serve the signed extension, so every agent is told it is disabled), and the store route waits for the store IDs. The working route today: a person with local administrator rights on the computer opens `chrome://extensions` or `edge://extensions`, turns on **Developer mode**, chooses **Load unpacked** and picks `C:\Program Files\YuviRed Agent\extension`. After store approval, an update will carry the IDs and the agent's installer will write the browser force-install policy so the extension installs from the store (planned, not in this release); see the deployment guide. Managed browsers carry consent by policy; the options page shows it.

Attachments that are read. Text files, `.docx`, `.xlsx`, `.pptx`, OpenDocument files, RTF, `.eml`, and ZIP, TAR and GZIP archives.

Attachments that are not read (they follow your "unreadable attachment" answer): PDF, images, older Office files (`.doc`, `.xls`, `.ppt`), Outlook `.msg`, and archives other than ZIP, TAR and GZIP.

Not inspected at all.

- Desktop AI applications and any AI site not on the list.
- Requests a page issues from a background worker thread (they are caught at the typing box, not at the network boundary).
- Text inside images or scanned documents.
- Clipboard history and keystrokes (the extension does not log either).
- Encrypted traffic: it never installs certificates or intercepts TLS (Transport Layer Security).
- Anything on a computer with no extension, a disabled extension, or a removed site.

Length. A prompt over 262,144 characters may be only partly checked; over 2,097,152 characters, about the first and last 1,040,384 characters are inspected and **the middle is not**. YuviRed raises a medium alert "Prompt only partly checked" for these at every preset, even Monitor Only, so a warning that replaced a block does not go unnoticed. Review those alerts.

Older extensions. Versions older than 0.5.0 receive only a legacy blocking setting and cannot show warnings. Update them.

Not against a hostile site. The check runs in the AI site's own page. It protects people from sending something by mistake through a site that behaves normally. It is not built to hold against an AI site whose own code works against it.

When something goes wrong. On the computer, open the extension's options page (Extensions, YuviRed AI Prompt Monitor, Details, Extension options) and press **Copy status for support**. It holds the extension version, browser, consent, policy revision, ruleset version, how many sites are monitored, whether the last policy sync worked, the last error code and counts of waiting or discarded detections, never content. Send it to support@yuvired.com with the computer, browser, site and time. Do not include the prompt or file.

Alerts, notifications, sign-in and the audit log

This chapter covers working with alerts day to day, getting alerts out of the console (email, Slack, Microsoft Teams, a webhook), single sign-on, the audit log, and the integrations that are off until you turn them on. The AIDR chapter explains where alerts come from and how to triage one; this chapter is about the list, the lifecycle, and everything that leaves the console.

The example company is "Northwind Dental". Names and addresses are made up (a.rivera, the mail server `smtp.example.com`, the address 192.0.2.10).

---

1. The alerts list and the alert detail

1.1 The list

Open **Alerts** ("Alert Management"). Each row is one deduplicated alert: title, severity, status, classification, category, computer, user, last activity and days seen. The sidebar badge counts open alerts and shows "N crit" when any open alert is critical.

Filters (all combine): **Sort** (Recent activity, Risk first), **All Severity**, **All Status**, **All Classifications**, **All Categories**, **All Tools**, **All Time** (24 hours, 7 days, 30 days, by last activity), **Search alerts...**, **All Endpoints**, **All Users**. Selecting a point on the AI Map or a Dashboard card can open the list pre-filtered; changing a filter clears that pre-filter so the two never silently combine.

New and repeating alerts push live onto the page within about 2 seconds. The page also refreshes every 30 seconds, and **Refresh** reloads now.

1.2 Severities

Severity	Meaning
Critical	Act now. A blocklisted destination, a policy violation on a server, a governed AI agent outside its approved boundary.
High	A policy violation. Triage today.
Medium	Classify it. A tool under review, a local model, an unknown extension.
Low / Info	Awareness. Known, sanctioned AI activity.

The severity is fixed by the rule that raised the alert, not by who looks at it. A Prompt Protection alert keeps its rule's severity whatever the person decided on the warning.

1.3 Statuses

Status	Use it when
open	New. Nobody has acted.
acknowledged	Someone is aware (press **Acknowledge**).
investigating	Under investigation. Pressing **Investigate** sets this and opens a case file.
resolved	Handled.
false_positive	Not a real problem.
suppressed	A visibility record from a scoped rule set to Suppress. Not counted as open and never notified.

Open, acknowledged and investigating are the "active" statuses: repeats fold into an active alert. After resolved or false_positive, a repeat creates a **new** alert marked "Recurred after resolve".

1.4 Deduplication

One alert stands for "this thing on this computer", not for each event. The identity is the computer plus a stable subject (the AI tool, the extension, the destination, or for prompt findings the rule's name and the AI site), never the alert title. While an alert is active, repeats update **Last seen** and raise **Days seen** (one per distinct day). Reclassifying a tool does not create a duplicate. A Prompt Protection alert lists the events folded into it, newest first, with each event's own user, result and excerpt.

1.5 The detail

Click an alert. The detail holds **Who**, **Owner**, **What**, **Where**, **When**, **Why**, **Action** and **Triage**, as described in the AIDR chapter. Footer buttons: **Close**, a channel drop-down with **Send**, **Investigate** and **Save**. Closing with unsaved changes asks first.

1.6 Procedure: close an alert with a note

Purpose. Leave a clear record of why an alert is closed.

Prerequisites. Analyst or Administrator.

Steps.

1. Open **Alerts**, click the alert.
2. Read **Evidence** and **Reason**.
3. Under **Triage**, choose **Status**: **resolved** (it was real and is dealt with) or **false_positive** (it was not a problem).
4. Write what you found or did in **Analyst Notes**. Optionally put your ticket's id in **Ticket reference** (for example `TICKET-1234`).
5. Press **Save**.

What you should see. A confirmation toast, the status tag updated, and the sidebar count lower by one.

How to confirm. Reopen the alert: your note and status remain. The change is written to the audit log as `alert.status_update`.

Common problems.

- *"This alert changed since you opened it."* Someone else edited it; reopen it.
- *It came back.* If the same activity happens again after you resolved it, a new alert appears marked "Recurred after resolve". For a false positive that will keep happening, create a scoped rule set to **Suppress** (AIDR chapter, Policies) instead of closing alerts one at a time. Critical findings and prompt findings can never be suppressed.
- *A Rogue agent activity alert.* The detail also offers **Mark as expected behaviour**, which records the behavior as normal for that agent and closes the alert as a false positive. That is stronger than resolving: it stops those facts from counting again for that agent. It needs a written justification.

1.7 Assigning and the response target

Use the **Assigned to** drop-down on the alert (or the person icon to assign it to yourself). The **SLA due** line comes from **Configuration**, **Detection & Response**, "Alert response-time targets" (Critical 4 hours, High 24 hours, Medium 1 week recommended; Low none). It is a tracking goal only: nothing happens automatically when it passes, except the alert shows "SLA breached". The target is set when the alert is first created and does not reset when repeats fold in.

2. Notifications

2.1 What triggers one

Notifications are sent for **new** alerts at the moment a computer's check-in creates them. Folded repeats do not notify. Three controls decide whether a new alert notifies:

1. **Severity threshold**: "Notify me about" (default **High and above (Recommended)**; choices: Critical only, High and above, Medium and above, Everything).
2. **Repeat suppression**: "Don't repeat the same alert for" (1 hour, **6 hours (Recommended)**, 12 hours, 24 hours). A repeat of the same finding on the same computer inside this window is not re-notified.
3. **Hourly cap**: "Max notifications per hour" (10, **20 (Recommended)**, 50, 100). When the cap is hit, YuviRed sends one "Notification rate cap reached" notice instead of silently dropping the rest.

Suppressed alerts never notify. A channel works only when it is both **enabled** and **configured** (a mail server and recipients, or a URL). Anything else is silent. These counters are kept in memory and reset when the YuviRed server restarts.

Daily summary email. If you turn on **Daily summary email**, medium and low findings are held and sent once a day after the hour you choose under **Send after** (shown in UTC, default 07:00 UTC). **High and critical findings always notify immediately** and are never held for the summary.

2.2 Where to set it up

Open **Configuration** and the **Notifications** section ("Which events reach your team, and where they are delivered"). Save with **Save Changes**. Below the fields a test panel lists each channel with a **configured** or **not configured** tag and a **Test** button that sends a small test message.

Field (as labelled)	What to enter
Email alerts	Switch for the email channel.
Send alerts to	One or more addresses separated by commas, for example `security@example.com, owner@example.com`.
Mail server	The host name of **your own** mail server (SMTP, the standard mail-sending protocol), for example `smtp.example.com`. YuviRed does not provide a mail server.
Mail server port	Default 587.
Mail username	The account on your mail server, if it needs one.
Mail password	That account's password. Stored encrypted. The page shows only "not set" or "configured", never the value.
Encrypt mail connection (TLS)	On by default. YuviRed starts an encrypted session (STARTTLS) before logging in. Leave on unless your server cannot do it.
Send from address	The "From" address, default `noreply@yuvired.local`. Use an address your mail server will accept.
Slack / Microsoft Teams alerts	Switch for the chat channel.
Slack or Teams webhook URL	An incoming-webhook URL created in Slack or Teams. Stored encrypted.
Webhook feed	Switch for the generic JSON webhook.
Webhook URL	The address that should receive each alert as JSON (a SIEM, a ticketing or automation tool). Stored encrypted. Must start with http or https.

Mail is sent through the mail server you specify, with a 5-second connection limit. A failure is logged and shown in the Test result; it never blocks monitoring.

2.3 Procedure: set up email notifications

Purpose. Get a message when a high or critical alert is raised.

Prerequisites. Administrator. A mail server that accepts connections from the YuviRed server, an account on it, and a recipient mailbox. YuviRed has to reach the mail server on its port: check your firewall.

Steps.

1. **Configuration**, **Notifications**.
2. Turn on **Email alerts**. Enter **Send alerts to**, **Mail server**, **Mail server port**, **Mail username**, **Mail password**, leave **Encrypt mail connection (TLS)** on, and enter **Send from address**.
3. Leave **Notify me about** at **High and above**.
4. Press **Save Changes**.
5. In the test panel, find **Email** (it should read **configured**) and press **Test**.

What you should see. "Test Email sent" and a message in the mailbox within a minute, with a subject like "[YuviRed] HIGH: Unreviewed AI tool detected".

How to confirm. Open the message. It should follow the layout in section 2.4.

How to undo. Turn off **Email alerts**, or clear **Mail server**.

Common problems.

- *"smtp.host not configured"* or *"notifications.email_to not configured"*: fill in the missing field and save.
- *"SMTP send failed: ..."*: the text after the colon is the mail server's or the network's answer. Typical causes: wrong port, wrong password, the server requires an app password, the YuviRed server cannot reach the mail server, or TLS does not match your server.
- *Test works but real alerts do not arrive.* Check **Notify me about**, the repeat window and the hourly cap, and that the alert is not suppressed.

2.4 What each channel sends

Every message carries one fixed sentence: **"YuviRed is alert-only. No traffic, process, or access was blocked."** The sentence is there so nobody reads a notification as proof that something was stopped.

Email (plain text). Subject: `[YuviRed] SEVERITY: alert title`. Body, in order: severity, priority and risk score; the title; the category and tactic; the disposition ("Alerted (alert-only, no enforcement)"); then five labelled blocks, **WHO** (computer id, host name, signed-in user), **WHAT** (the AI tool, its classification, its policy label, and evidence metadata such as domains or process names), **WHERE** (operating system, the computer's IP address, agent version), **WHEN** (first seen, last seen, accessed time and days seen), **WHY** (the rule's description and where the classification came from); then the alert-only sentence.

Slack or Teams (a short text message). First line: `YuviRed SEVERITY (NEW): title on hostname`. Second line: the tool name and its classification. Third line: the alert-only sentence. Slack and Teams incoming webhooks both accept this simple form.

Webhook (JSON, sent as an HTTP POST). Fields:

Field	Content
`source`, `kind`	"yuvired" and "ai_detection_alert"
`alert_id`, `display_id`	The alert's id and its human id (such as SR-ALRT-000123) for ticket titles
`external_ref`	Your ticket reference, if you entered one on the alert
`recurred_from`	The earlier alert's identity and ticket reference, if this alert recurred after being resolved, so your tool can route it to the existing ticket
`severity`, `priority`, `risk_score`	How serious it is
`title`, `category`, `category_label`, `tactic`	What it is
`disposition`	"Alerted (alert-only, no enforcement)", with the policy label when a tool is labelled deny or block
`who`	Computer id, host name, signed-in user
`what`	Tool name, classification, policy label ("NOT enforced") and evidence metadata
`where`	Operating system, IP address, agent version
`when`	First seen, last seen, accessed time, days seen
`why`	The rule's description and where the classification came from
`enforcement_note`	The alert-only sentence

What is not in any notification. Prompt text and raw prompt samples are never included: Prompt Protection alerts carry the masked evidence only. Notifications do include the computer's host name, the signed-in user's name and the computer's IP address. Think before you send them to an outside chat service. Webhook and chat URLs are secrets: anyone with the URL can post to that channel.

Manual send. On an alert, pick a channel next to **Send** and press it to push that single alert to the channel now, regardless of the threshold. A channel that is not configured appears as " (not configured)" and is disabled.

The Notifications section of Configuration. Email, chat and the webhook feed are set up here; the webhook address shown is a made-up lab example.

---

3. Single sign-on (OIDC), if you want it

OIDC (OpenID Connect) lets people sign in to the console through your organization's identity provider (Microsoft Entra ID, Google Workspace, Okta and similar). It is **off** until you turn it on. Local accounts keep working.

Prerequisites. Administrator. An application registered with your identity provider that gives you a client id and client secret. The provider's redirect (callback) address set to your YuviRed console address. YuviRed's own **public URL** (https) that people use to reach the console.

Settings. Configuration, **Integrations**, "Single sign-on (OIDC)": **Single Sign-On** (switch), **Login button label** (for example "Sign in with Northwind"), **Identity provider URL**, **Client ID**, **Client secret** (stored encrypted), **This appliance's public URL** (required). It stays inactive until the provider URL, client details and public URL are all set. Advanced: **Auto-create accounts on first SSO sign-in** (new people get a **read-only** account; has no effect until **Approved email domains for auto-created accounts** is set), **Extra trusted identity-provider hosts**, **Requested sign-in strength (OIDC acr_values)** and **Require provider MFA for roles that require MFA**.

> **MFA and SSO.** People who sign in through SSO do not go through YuviRed's own multifactor stage: your identity provider is the authority. Turning off "Require provider MFA for roles that require MFA" silently exempts SSO users from any MFA requirement you set on YuviRed roles. Leave it on.

Steps.

1. Register the application with your provider; note the client id and secret.
2. In **Integrations**, fill in the fields above and turn on **Single Sign-On**. Save.
3. Sign out. The sign-in box now has **Sign in with SSO** (or your label) under an "or" divider.

4. Sign in with a test person. You should return to YuviRed signed in.

How to confirm. The audit log records `auth.sso_login_success` (and, for a new person, `auth.sso_user_provisioned`).

Messages a person may see and what they mean. "No account exists for your SSO identity" (automatic creation is off and no local account matches: create or link one); "Your email domain is not approved for automatic access"; "Your identity provider email address is not verified" (the provider did not say the email is verified); "This account is already linked to a different SSO identity"; "An administrator must link this privileged account to SSO before you can sign in this way" (an Administrator account needs linking first); "Your role requires multi-factor authentication. Your identity provider did not report an MFA method"; "This account is deactivated"; "SSO sign-in expired or was already used, try again"; and generic provider errors, for which you should retry once and then read the audit log for `auth.sso_login_failed`.

Limits. SSO has been exercised against test identity providers; this guide has no record of a live sign-in test against a named provider on Lite. Test with your own provider before relying on it, and keep one local Administrator account for emergencies.

---

4. The audit history

4.1 What is logged

Every sensitive action is written to the audit log with who did it, when, the address they came from, what it acted on, and a detail record. Examples (the action names you will see in exports):

Area	Actions
Sign-in	`auth.login_success`, `auth.login_failed`, `auth.login_locked_out`, `auth.logout`, `auth.change_password`, `auth.mfa_challenge`, `auth.sso_login_success`, `auth.sso_login_failed`
Settings	`config.update` (with before and after values), `config.update_rejected`
Alerts and tools	`alert.status_update`, `alert.send`, `alert.expected_behavior`, `classification.*`, `ai_tool_catalog.*`, `blocklist.add`, `blocklist.remove`
Enforcement	response action requests, approvals, rollbacks
Prompt Protection	`gateway.prompt_blocked`, `gateway.prompt_warned_allowed`, `gateway.prompt_guard_unavailable`, policy revisions
Raw content	`dlp.content_event_viewed`, `dlp.raw_sample_viewed`, `aidr.events_exported`
AMS	`ams.agent.stop_requested`, `ams.agent.hold`, `ams.agent.message_sent`, `ams.approval.decided`, `ams.control_policy.update`, `ams.connector.test_requested`, finding changes
Computers	`agent.enrolled`, `endpoint.delete`, `endpoint.decommission`, `credential_rotation.*`
License and legal	`license.activate`, `legal.accept`
Reports	generating, downloading, cancelling and deleting a report (report number, formats and sizes only)
The audit log itself	`audit.export` (pulling the audit trail is on the audit trail) and `audit.retention_purge`

4.2 Where to see it

There is no Audit page in the console in this release. You can reach audit records in these ways:

- **In context.** Each AMS agent has an **Audit history** fold-out in its workspace (needs the audit read permission). The Reports area includes an **Administration and accountability** section (Auditor's Report) covering setting and policy changes and the account list.
- **By export.** The export exists **only as a web interface (API) route; the console has no button for it.** An account with the `audit.read` permission (Administrators and View

Only accounts have it; Security Analysts do not) can call `GET https://<server>:8443/api/v1/audit/export?format=csv` (or `format=json`, or `format=cef`) with the `Authorization: Bearer <token>` header of a signed-in session. The three formats are **CSV** (spreadsheets), **JSON** (everything, including the detail record) and **CEF** (Common Event Format, the form most security information and event management systems read). Optional query parameters `actor_id`, `action` and `target_type` filter it; an export holds at most 50,000 rows, and each export is itself recorded as `audit.export`. Typing the address into a browser is not enough, because the console signs in with a bearer token rather than a cookie; this is a task for a technical administrator with a tool such as curl or PowerShell. How to obtain a session token is not described in this guide.

- **Support.** Contact support@yuvired.com if you need an export and cannot make the call yourself.

4.3 How long it is kept

The audit log is **kept until purged**. The setting **Audit log** under **Configuration**, **Data & Retention** defaults to 0, meaning keep forever. Setting a number of days permanently deletes older records. Uninstalling YuviRed keeps `%ProgramData%\YuviRed` (including the database) unless the purge option is chosen.

4.4 Limits

- Records are in the same database as everything else. There is no tamper-proof chain: an administrator with database access could alter them. Export regularly to an append-only location if your auditors need that.
- Audit entries for prompt events say that someone viewed content, never what it said.

5. Integrations that are off by default

Integration	Default	What turning it on does
Email, Slack or Teams, Webhook	Inert until a server or URL is set	Sends alert details (host name, user name, IP address, tool) to those destinations (section 2).
Daily summary email	Off	Sends one daily summary of medium and low findings.
Single Sign-On	Off	Sends sign-in requests to your identity provider (section 3).
Blocklist Feed URL	No token	Lets your firewall fetch the AI-destination list with a token (AIDR chapter).
Artifact Analysis	Off	Lets analysts have uploaded files examined by external providers (MetaDefender Cloud, Hive). With the policy on **Upload disabled** nothing leaves YuviRed, not even a hash; **Hash only** may send file hashes; the upload policies may send the files. The console states what each policy permits before it is saved.
Raw prompt capture	Off	Stores prompt text encrypted for a short time (Sensitive Data chapter).
Connectors (AMS)	All off	Edits other tools' settings to add YuviRed's hook (AMS chapter).
Remote Browser Extension Wiring	Off, and inert on Lite	On a Lite server the switch has no effect in this release: the server does not serve the signed extension, so every agent is told the wiring is disabled whatever the switch says. Nothing installs the extension automatically; see the Sensitive Data chapter, section 8, for the route that works today.
Daily license check-in	On after activation	Sends six facts: format, installation id, version, serial, computer count and AI agent count. No telemetry, no names.

YuviRed has no connection to a SIEM, ticketing tool or directory beyond the webhook, the export files and the single sign-on described here. The product does not send telemetry to YuviRed.

Reports

YuviRed produces two reports, each downloadable on its own, plus a live on-screen overview:

- **Executive Summary**: a concise leadership overview of activity, risks, protection and recommended actions, for executives and security leads. Usually 2 to 4 pages as a PDF. The first page stands on its own.
- **Auditor's Report**: detailed findings, evidence, configuration history and supporting records for auditors and technical reviewers.
- **Live overview**: the older on-screen summary with CSV, HTML and print options (section 8).

Both reports describe **what YuviRed recorded**. Neither is a certification, and neither says that your organization complies with any law or standard. The example company is "Northwind Dental"; all names and numbers in examples are made up.

---

1. Before you start

- **Where.** Choose **Reports** in the left menu (AIDR or AMS list) and stay on the **Reports** tab, the first tab. The **Live overview** tab is the second.
- **Permission.** You need `reports.read` (every built-in role has it). Some sections need more (section 6). A section your role may not see is still listed and marked "Not included: your role cannot view this", so a report never hides that something was left out. Ask an Administrator to generate it if you need those sections.
- **Data.** A report can only describe computers that run the YuviRed agent and reported in the period.

---

2. Make a report in four steps

Procedure: make an Executive Summary for November

Purpose. Give the practice owner a one-page-first overview of the month.

Prerequisites. `reports.read`. At least a few days of data.

Steps.

1. **Choose the report type.** In **New report**, under **Report type**, pick **Executive Summary** ("About 2 to 4 pages"). Optionally type a **Report title (optional)**, for

example `Northwind Dental, November 2026`. Leave it blank for an automatic title.
2. **Choose the period.** Set **Start date** `2026-11-01` and **End date** `2026-11-30`, or press a quick range: **Last 7 days**, **Last 30 days**, **Last 90 days**, **This month**, **Last month**, **This year**. Set **Time zone** (it starts as your computer's; type to search, for example `America/New_York`). A line under it says "That is 30 days." The report covers the whole first day and the whole last day in that time zone and handles daylight saving changes.
3. **Choose the file format(s)** under "Executive Summary file formats", for example **PDF**. Each format you tick makes its own file.
4. Leave **Compare with the previous period of the same length** ticked. The report then also shows what changed since the previous 30 days and says plainly when the two periods cannot fairly be compared.
5. Optional: open **Optional filters** and **Privacy options** (section 3).
6. Press **Preview**. Read it: it names the report type, shows the exact period, how many computers are included, which areas are included or left out and why, rough record counts and known gaps.
7. Press the button that says what it will make, **Generate Executive Summary (PDF)**.

What you should see. A toast "Report RPT-000012 started" and the page scrolls to **Report history**, where the new row shows its progress. You can leave the page while it runs.

How to confirm. When the status reads complete, press **Download** next to each file. Open the PDF: the first page lists the period, the time zone and the key indicators.

Common problems.

- *"The period is longer than 400 days."* Choose a shorter period.
- *"You already have reports being generated."* You can run three at once and the server runs ten in total; wait for one to finish.
- *The report failed.* The row shows the reason. An internal error shows a short reference; send it to support@yuvired.com. If the server restarts while a report runs, it is marked failed; generate it again.

---

3. Choices that shape a report

3.1 Report type and sections

- **Executive Summary** always covers every area you are allowed to see.
- **Auditor's Report** has **All sections** ticked by default. Untick it to choose some of these sections:

Section	What it covers
Monitoring coverage	Which computers reported, how well, and the hour-by-hour visibility of each
AI inventory and usage	AI tools, websites, extensions and services seen and their classifications
AI agent activity	AMS: agents found, approvals, interventions and messages
Prompt protection	Prompt findings and outcomes from the browser extension
Policies and settings	Active policies and recorded setting changes
Alerts and response	Alerts, their handling, and response actions
Administration and accountability	Who did what: accounts and audit entries
Evidence quality and limitations	How complete the data is and what the report cannot say

Detail level (Auditor's Report): **Summary** (up to 25 rows per table in the PDF and Word files), **Standard** (200) or **Complete** (every row; can make a long report). The data files (CSV, Excel, JSON and the ZIP package) always hold every record.

3.2 Filters

Optional filters limit the report to some computers, groups, users, AI tools, AI agents, severities, policies, types of sensitive information, or event types (for example "AI websites visited in a browser", "AI programs running", "Network connections to AI services", "AI service name lookups (Windows)", "AI browser extensions", "AI services running on the computer", "AI software installed"). The tick-box **Include lab and test computers** is off by default: lab and test computers are left out. A counter next to "Optional filters" shows how many are active.

Users filter. Prompt protection records name their own user and are filtered by it. Other records do not name a user, so they are filtered by the computer's **last signed-in user**. A report for one person can therefore include things another person did on a shared computer.

3.3 Privacy options

Option	Default
Include user names	On
Include IP addresses	Off
Include masked samples of sensitive information (partly hidden values such as `4111********1111`)	Off
Include analyst notes	Off

Passwords, keys, tokens and prompt text are **never** included, whatever you tick. Notes and masked samples can still be sensitive; leave them off for reports that leave the practice.

3.4 Date and time handling

You pick whole days in a time zone; the report covers the first day from 00:00 and the last day to 23:59 in that zone. Times in the report carry the zone and its offset, for example `2026-11-03 14:05:00 -05:00`. The maximum period is 400 days. The comparison period is the period of the same length just before.

3.5 Formats

Format	Executive Summary	Auditor's Report
PDF	Branded leadership briefing: key indicators, takeaways, coverage, AI activity, sensitive information, protection in effect, findings, response, comparison, recommended actions, limitations	Complete branded report: contents, sections, findings with evidence references, appendices, page numbers
Word (.docx)	Editable copy of the PDF	Editable copy of the PDF
Excel (.xlsx)	Summary sheet, comparison, recommendations, limitations and the summary tables	One sheet per dataset, plus summary, findings, evidence index, data dictionary, limitations and glossary
CSV files (.zip)	One CSV per summary table, plus `report_metadata.json`	One CSV per dataset, including the evidence index and data dictionary
JSON	Structured summary	The whole report as structured data
ZIP package	The summary documents and data, with `manifest.json` and `SHA256SUMS.txt`	Everything for the auditor: documents, datasets, evidence index, manifest and checksums

Each report type has its own list of formats; switching type never carries the other type's choices over. Every format made from one report comes from the same snapshot of data, with the same totals. One format gives one file; the ZIP package is optional.

---

4. Reading a report

4.1 Section status

Each section says which of these applies: **Evidence available**; **No matching activity observed** (YuviRed was watching and saw nothing); **Evidence missing or not collected**; **Feature disabled**; **Not supported in this edition**; **Outside YuviRed's scope**; **Not**

included: your role cannot view this; **Partial evidence**; **Not selected for this report**. **Missing evidence is never shown as zero.** A section says "No matching activity observed" only when the computers in scope reported during the period and the records for that time are still kept. If records may have been deleted under retention, or no computer reported, it says "Evidence missing or not collected" instead.

4.2 Executive Summary numbers

Each key indicator says how it is counted and what it is out of. Example (made up): "Prompts with sensitive information: 7, counted once each even when a prompt matched several rules; out of 412 prompts to supported AI sites." Rules of reading:

- "Prompts with sensitive information" counts each prompt once. The Auditor's Report shows both **records** (one per matched rule) and **submissions** (distinct prompts).
- A warning the person accepted is shown as "sent after a warning", never as prevented.
- **Monitor Only** means nothing was warned or blocked by design; the summary says so rather than implying protection.
- A block is counted only when the browser extension reported that it stopped the prompt. A response action counts as having taken effect only when it ran in Enforce mode and succeeded; a dry run changes nothing and is not counted as an effect.
- Proposed owners in the recommended actions are suggestions, not assignments.

4.3 Findings and evidence references

Findings (F-001, F-002, ...) are listed most serious first. Each answers: What happened? Why does it matter? What evidence supports it? What did YuviRed do? What remains unresolved? What should you do next? "Confirmed" means the records show it directly; "suspected" means the records match a pattern (for example an AI agent's behavior). **Evidence references** (such as E-ALR-00001) point to one record in the data files; **Appendix D, the Evidence index**, lists each one with its dataset, record id, when it happened and when YuviRed received it.

4.4 Appendices (Auditor's Report)

A Methodology (how counts are made, duplicates and delayed records), B Glossary, C Limitations, D Evidence index, E Data dictionary.

4.5 Monitoring, warning and blocking are kept apart

The report never merges "seen", "warned" and "blocked". A person who pressed **Send anyway** is counted as sent after a warning. Alerts are counted as raised, not as outcomes.

5. The same snapshot, a second report, and re-running

Every report keeps the snapshot of data it was made from, for as long as its files are kept. In the history list:

- An Executive Summary row has **Auditor's Report from this snapshot**, and an Auditor's Report row has **Executive Summary from this snapshot**. They use exactly the same data, so every number and every finding reference (F-001…) matches. The new report gets its own number and says "Same snapshot as RPT-...".
- **Re-run as new snapshot** reads the data again as it is now, so its numbers can differ if records arrived late. It gets a new number and the same choices.

Procedure: make the auditor's companion to a summary

1. In **Report history**, find the Executive Summary row (for example RPT-000012).
2. Press **Auditor's Report from this snapshot**, choose the formats and confirm.

What you should see. A new row, RPT-000013, "Same snapshot as RPT-000012". Its totals match RPT-000012.

6. Who can see what

Section	Needs
Monitoring coverage, AI inventory and usage	`endpoints.read`
AI agent activity	`ams.read`
Prompt protection, Alerts and response	`alerts.read`
Policies and settings	`policies.read`
Administration and accountability	`audit.read` (the account list also needs `users.read`)
Evidence quality and limitations	`reports.read`

Inside a section, some tables need a further permission: *Setting and policy changes* need `audit.read`, *Response actions* need `response.read`, *Enrollment tokens* need `endpoints.manage`, *Background job health* needs `config.read`. A table you may not view is empty and listed under Limitations.

7. History, downloads and retention of report files

Report history lists every report you made: number (RPT-000012), title, type, period, who made it, when, status, size, and whether its files are still stored. Buttons: **Refresh**, **Download** (per file), **Cancel** (while running), **Delete files**, **Re-run as new snapshot**, and the two "from this snapshot" buttons.

- **Retention of report files.** Files are kept for **30 days** by default (setting `retention.report_file_days`), and the total is capped at **2 GB** (`retention.report_storage_mb`; the oldest files go first). The history entry stays after the files are deleted. Download anything you need to keep. The ZIP package is the form to archive.
- **Delete files** removes the stored files and keeps the history entry. It cannot be undone, but you can re-run the report (the numbers may differ).
- Administrators (`reports.manage`) can tick **Show reports by everyone**. Everyone else sees only their own.
- Generating, downloading, cancelling and deleting reports are recorded in the audit log (report number, formats and sizes only, never the report contents).
- Making a report never changes any monitoring or protection setting.

Checking that a file has not changed

The audit package lists a SHA-256 checksum for every file. On Windows: `certutil -hashfile report.pdf SHA256`. Compare the result with the line for that file in `SHA256SUMS.txt`. Checksums show whether a file changed since the package was made. On their own they do not prove who made it, that it is complete, or a chain of custody.

8. The Live overview tab

Live overview is the older on-screen summary built from live data. A drop-down chooses **Executive summary** or **Administrator report**. Buttons: **Refresh**, **CSV**, **HTML**, **Print / PDF**. It describes the situation now, not a period, and does not keep a snapshot. Use the **Reports** tab for anything you need to file or hand to someone else.

9. What a report cannot show

- Computers without the YuviRed agent.
- The text people type to AI (never stored in a report), and what AI agents are told.

- Encrypted network traffic.
- Activity before YuviRed was installed, and records deleted under retention.
- Hour-by-hour visibility of each computer is recorded from this version on; earlier hours are reported as unknown.

---

10. Common problems

You see	Why	What to do
A section reads "Evidence missing or not collected"	No computer reported, or the records were purged under retention	Check the Endpoints page and the retention settings
A section reads "Not included: your role cannot view this"	Your role lacks the permission in section 6	Ask an Administrator to generate it
A number differs between two reports of the same period	One was re-run later and records arrived late	Use "from this snapshot" to get matching reports
Files missing from the history row	Retention deleted them (30 days default, 2 GB cap)	Re-run; consider longer retention
The Word file looks different from the PDF	Same content, editable layout	Use the PDF as the reference

The new report form: pick the type, the period and the file format, then Preview or Generate.

The report history after generating a report. The row shows status Done and the download buttons.

Administering your YuviRed server

This guide is for the person who looks after a YuviRed server after it is installed: updates, copies of the data, moving to a new server, getting back into an account, who can do what, what is kept and for how long, privacy settings, and how to ask YuviRed for help.

It assumes the server is installed and activated. If it is not, start with the Installation guide and Licensing.

Example names used here are made up: the company "Northwind Dental", the server address `192.0.2.10`, the administrator `a.rivera`.

What is on the server. YuviRed keeps its programs in `C:\Program Files\YuviRed` and everything it learns in `C:\ProgramData\YuviRed`. Three Windows services run it:

Service name	Shown in Services as	What it is
`YuviRedPostgres`	YuviRed PostgreSQL	The built-in database. Listens on this computer only (port 5440). Runs as `NT SERVICE\YuviRedPostgres`.
`YuviRed`	YuviRed	The web console and the server. Listens on port 8443. Depends on `YuviRedPostgres`. Runs as `NT SERVICE\YuviRed`.
`YuviRedAgent`	YuviRed Endpoint Agent	The server's own endpoint agent. Depends on `YuviRed`. Runs as the local system account.

`C:\ProgramData\YuviRed` holds:

Folder or file	Holds
`pg`	The database (`pg\data`), including its own logs in `pg\data\log` .
`config`	The four generated secrets (`yuvired.env`), the database password file (`pg.pwfile`), the HTTPS certificate and private key (`tls-cert.pem` , `tls-key.pem`), the network setting (`network.conf`).
`identity`	The server's identity record (`appliance.json`), its private key sealed for this computer (`appliance_private.dpapi` , see section 2.2), and markers.
`agent`	The server's own agent: its settings, state and enrollment credential.
`status`	Status files that hold no secrets: `boot-status.json` , `platform-binding.json` , `network-status.json` , `legal-acceptance.json` (the license acceptance record).
`logs`	The server log `syntinel-red.log` and the agent start logs.
`backups` , `scratch` , `staging`	Created empty. Nothing in this release writes a backup into `backups` (section 2.1).
`update.journal`	The update engine's journal (not used by this release's installer updates).

Windows allows only administrators, the system and the two YuviRed service accounts to read the folder. The server's own agent writes its log elsewhere, in `C:\ProgramData\SyntinelRed\logs` (the folder name is the product's earlier name; it will change in a later release).

The installation ID is created at the first start and kept in `identity\appliance.json` and in the database. It is not calculated from the hardware; the hardware identity (the computer's firmware serial numbers) is recorded next to it, and a change of hardware makes the server give itself a **new** installation ID (section 3.1).

---

1. Updates

1.1 How a new version reaches you

YuviRed never installs an update by itself. There are two ways you learn that a newer version exists:

- **The console notice.** Once a day an activated server checks in with YuviRed (see what YuviRed receives from your server). If a newer version exists, a banner says "YuviRed `<version>` is available. This server runs `<current version>`. Nothing is installed automatically." with a **Download** link when YuviRed has published one, and **Dismiss**. Dismiss hides the banner for that version in that browser only.
- **The website.** New installers are published at https://yuvired.com. A server that cannot reach YuviRed never shows the banner, so check the website now and then.

Support and retirement policy (from the facts sheet):

- Email support, reasonable efforts, for Lite and Pro. There is no response-time promise. Enterprise is as agreed.
- Security fixes are made for the current version, and for 90 days for the version it replaced.
- YuviRed announces the retirement of a version at least 60 days ahead, on yuvired.com and in the console. (The console notice described above is the only update message found in the product code; a separate retirement message in the console was not verified.)

In practice: plan to move to the current version within 90 days of a new one appearing.

1.2 Install an update (in place)

Purpose. Replace the programs with a newer version and keep your accounts, computers, policy, settings, certificate, installation ID and license.

Before you start.

- You are signed in to Windows on the server as an administrator.
- You have the new `YuviRed-Lite.exe` from https://yuvired.com (or the link in the banner). Check its SHA-256 fingerprint against the one published with the download, and check the signature: right-click the file, **Properties**, **Digital Signatures**, and the signer reads **YuviRed LLC**. The Windows prompt that asks to allow the installer also names the publisher **YuviRed LLC**. A file whose prompt says **Unknown publisher** is not ours: do not run it. The installer, its removal program and the agent installer inside an Add Computer download are all signed the same way.
- Pick a quiet time. The console is unavailable while the update runs, and the server does not collect while the service is stopped. Computers keep trying, and what a computer could not send in that time is lost (the agent has no offline queue for most telemetry).
- Take a copy of the data first (section 2). **YuviRed takes no copy for you before an update.**

- Read the notes that come with the new version.

Steps.

1. Close any open YuviRed installer window.
2. Double-click the new `YuviRed-Lite.exe`.
3. Accept the license page and click **Install**, then accept the Windows administrator prompt.

   You should see: the installer window with the checklist (Preparing installation, Installing YuviRed, Configuring database, Starting services, Securing installation, Verifying YuviRed).

   The installer progress page, as it looks during an install or upgrade. Captured in the YuviRed lab; the Windows evaluation watermark is not part of your screen.

4. Wait for **Installation Complete**, then click **Open YuviRed**.
5. Sign in with your existing administrator account.

An unattended update uses `YuviRed-Lite.exe /quiet`. An update of a server that already has its acceptance record (`status\legal-acceptance.json`) does not need `ACCEPT_EULA=1`; the installer logs "upgrade of an installation with a recorded legal acceptance". A record of the same license version is kept, never overwritten.

How an update runs. The installer removes the previous version's registration first, then puts the new files in place, prepares the data folder, starts the services and waits up to five minutes until the new version reports **YuviRed Ready** and answers on port 8443. The whole update is one Windows Installer transaction: if any of those steps fails, the files and services of the previous version come back and the window says "The upgrade could not be completed. The previous YuviRed installation and your data were left in place; you can run the installer again."

What an update keeps. Everything under `C:\ProgramData\YuviRed`: accounts and roles, enrolled computers, alerts and history, policy, settings, the TLS certificate and key, network settings, the installation ID and the license. Computers keep their identity and keep reporting.

What an update does not do.

- It does not update the agents on your other computers. An agent installed from the Add Computer download never updates itself. A newer agent reaches a computer with a new

Add Computer download (Agent Deployment page) run on that computer. A new extension reaches a computer the same way.

- It does not ask before it changes the database. The new version changes the database layout when the service first starts (the boot step `migrate`, phase **Schema Current**), inside the same installer run.
- It does not take a copy of your data first.

How to confirm.

1. The console footer shows the new build number.
2. On the Windows server, open Services: `YuviRedPostgres`, `YuviRed` and `YuviRedAgent` all read **Running**. The update stops `YuviRedAgent` with the server; the new version starts it again itself, trying for up to three minutes and writing each try to `C:\ProgramData\YuviRed\logs\agent-start.log`. If it still reads **Stopped** after that, run `net start YuviRedAgent`.
3. In the console, **Configuration**, **License** still shows your plan as **Active**.
4. On **Endpoints**, computers show **Online** within a few minutes.

Common problems.

Problem	What to do
The installer says a newer version is already installed	You ran an older installer. Never run an earlier installer over a newer one. Use the newer file.
The installer fails and rolls back ("The upgrade could not be completed...")	The files and services return to the old version, but the database is not rolled back: if the new version had already changed its layout, the old version may not start ("Can't locate revision" in the log). Do not start trying again repeatedly. Run the support collector (Collecting diagnostics) before you reboot or retry, and send it to support@yuvired.com. Running the new installer again is usually the way forward, because the database then matches it.
The console asks you to activate after an update from an early build that had no license	A server upgraded from a pilot build that held no license has 7 days to activate. Your YuviRed contact sends a code. See Licensing.

Rollback. There is no supported way to go back to an older version. A newer version changes the database layout, and an older version cannot read it: its service does not start. The way

back is to restore the copy you took before the update (section 2.4) onto a server that has the older version installed. This path has not been tested by YuviRed (see section 2.1).

---

2. Back up (manual procedure)

2.1 What YuviRed offers, plainly

YuviRed Lite has no built-in backup or restore command. There is no backup button, no scheduled backup and no restore wizard. Nothing in the installer or in the console writes a copy for you. The `backups` folder under `C:\ProgramData\YuviRed` is created empty and stays empty. (The program contains code that can write a backup set, a database dump plus the secrets file, but no command, button or schedule runs it in this release: the server command `yuvired_service.py` has no `backup` command, and the update engine that would call it is not used by the installer-based updates of this release.)

What follows is a manual procedure that a Windows administrator can follow. **YuviRed has not run this procedure end to end.** It follows from how the product stores its state. Test a restore on a spare computer before you rely on it, and do it again after each major update.

If your server is a virtual machine, a copy of the whole machine made by your hypervisor is a second option. YuviRed has not tested restoring a server that way either, and a restore that gives the machine different virtual hardware makes YuviRed treat it as a different server (see section 3).

2.2 What must be in the copy

Copy the whole folder `C:\ProgramData\YuviRed`. The parts that matter:

Folder or file	Holds	If you lose it
`pg`	The database: all computers, alerts, accounts, policy, history.	All history and accounts.
`config\yuvired.env`	The four generated secrets: the token signing key, the encryption key, the token pepper and the database password.	The database cannot be opened by the server. **Encrypted data in the database (MFA secrets, saved mail and webhook secrets, original prompt samples) becomes unreadable for good.**
`config\pg.pwfile`	The database password the database itself uses.	The database service cannot start with its data.
`config\tls-cert.pem`, `config\tls-key.pem`	The server's HTTPS certificate and its private key. Enrolled computers trust exactly this certificate.	Enrolled computers stop trusting the server. Each needs a new Add Computer download.
`config\network.conf`	The address or name other computers use to reach the server.	You must set network mode again.
`identity`	The server's identity record (installation ID, public key, hardware binding) and its private key, sealed for this computer. (The license itself, and a copy of the installation ID, are in the database, `pg`.)	The server cannot start until the identity is dealt with (see below).
`agent`	The server agent's own settings, state and enrollment credential.	The server's own agent cannot report until it is enrolled again.
`status`	The license acceptance record and status files.	The acceptance record of the install is lost (the server still runs).

Keep the keys with the data. The encryption key lives in the secrets folder (`config`). A copy of the database without the matching `config` folder is not usable: the encrypted parts cannot be read. The folder holds secrets and the server's HTTPS private key **in plain files**.

Store the copy only on encrypted storage that only administrators can read, never in a shared folder, a ticket or an email.

How the identity key is sealed, and what it means for a copy. The private key in `identity` is not stored in the clear. On a new install the service tries the computer's TPM (the Trusted Platform Module) and, when the service account cannot use it, falls back to Windows Data Protection for the machine (DPAPI, machine scope, file `appliance_private.dpapi`). The code records which one it used in `appliance.json` (`key_protection`: `tpm2-pcp` or `dpapi-machine`); on the real installs YuviRed has run, the TPM path was refused to the service account and `dpapi-machine` was recorded. What that means:

- A key sealed with DPAPI opens only on the **same Windows installation** (Windows keeps the master key with the operating system). It does not open on another computer, in a different Windows installation on the same hardware, or in a virtual machine restored as a new machine. A key sealed with the TPM opens only on that TPM.
- At every start the service loads the key before it serves. If the key cannot be opened, the start stops at the step `provision_identity` with phase **Error** and the message "the appliance private key could not be loaded" in `status\boot-status.json`. The console does not open. There is no command that repairs this: the way forward is a fresh identity (a fresh install, section 3.2).
- So a copy of `C:\ProgramData\YuviRed` restores fully only on the same Windows installation (section 2.4). On other hardware, restore `pg` and `config` but not `identity` (section 3.3).

2.3 Make a copy

Purpose. A consistent copy of the server's data, taken while the database is stopped.

Before you start. An administrator account on the server. Somewhere to put the copy that is not the same disk: encrypted, access-limited storage, for example `E:\YuviRed-backup` on an encrypted drive. Enough free space for the size of `C:\ProgramData\YuviRed` (check it in Windows Explorer, Properties).

Steps.

1. Choose a quiet time. Computers cannot report while the service is stopped.
2. Open **Windows PowerShell** as administrator: right-click Start, then **Terminal (Admin)** or **Windows PowerShell (Admin)**.
3. Stop the three services, in this order:

```
net stop YuviRedAgent
net stop YuviRed
net stop YuviRedPostgres
```

You should see: "The ... service was stopped successfully." for each. (`net stop YuviRed` may ask to stop `YuviRedAgent` with it if you skipped the first line; answer Y.) Stopping the database while it is running is what makes the copy consistent. Do not copy the `pg` folder while `YuviRedPostgres` runs.

4. Copy the folder with its permissions and owners:

```
robocopy "C:\ProgramData\YuviRed" "E:\YuviRed-backup\2026-11-19" /E
/COPYALL /B /R:1 /W:1
```

You should see: a summary table at the end with **Failed** `0` for both Dirs and Files. A copy that reports failures is not a backup.

5. Start the services in the reverse order:

```
net start YuviRedPostgres
net start YuviRed
net start YuviRedAgent
```

6. Write down, next to the copy: the date, the YuviRed version and build shown in the console footer, and the installation ID (**Configuration**, **License**, or the license screen).

How to confirm.

- Open the console and sign in. **Endpoints** shows computers **Online** within a few minutes.
- The copy folder contains `pg`, `config`, `identity` and `agent`, and `config` contains `yuvired.env`.

Common problems.

Problem	What to do
`net stop` says the service is not started	Fine. Continue.
`robocopy` reports "Access is denied" or failed files	You are not running as administrator, or an antivirus holds files. Run PowerShell as administrator and try again.
The services do not start afterward	Run `net start YuviRedPostgres` first and wait ten seconds, then `net start YuviRed`. See Troubleshooting, "The YuviRed service does not start".

Rollback. The procedure only reads the data folder. If a step fails, start the services again.

How often. The decision is yours. The server's data changes all day. Take a copy before every update, and on a schedule that matches how much history you can afford to lose. Keep at least one copy off the server.

2.4 Restore on the same server

Purpose. Return this server's data to the state of a copy.

Before you start.

- The same YuviRed version as the copy is installed on this server. A copy made by an older version can be opened by a newer one (it changes the database layout when it starts), but a copy made by a newer version cannot be opened by an older one.
- The same Windows computer, the same Windows installation and the same hardware: the identity key opens only there (section 2.2). Restoring on different hardware, or after Windows was reinstalled, is section 3.
- You know that everything collected after the copy was made is lost, and that computers whose credentials were rotated after the copy may need attention (see the note below).

Steps.

1. Open PowerShell as administrator.
2. Stop the services (the same three commands as in 2.3).
3. Move the current data aside so you can go back:

```
mkdir C:\YuviRed-before-restore
move "C:\ProgramData\YuviRed\pg" C:\YuviRed-before-restore\
move "C:\ProgramData\YuviRed\config" C:\YuviRed-before-restore\
move "C:\ProgramData\YuviRed\identity" C:\YuviRed-before-restore\
move "C:\ProgramData\YuviRed\agent" C:\YuviRed-before-restore\
```

4. Copy the four folders back, keeping permissions and owners:

```
robocopy "E:\YuviRed-backup\2026-11-19\pg" "C:\ProgramData\YuviRed\pg" /E
/COPYALL /B /R:1 /W:1
robocopy "E:\YuviRed-backup\2026-11-19\config"
"C:\ProgramData\YuviRed\config" /E /COPYALL /B /R:1 /W:1
robocopy "E:\YuviRed-backup\2026-11-19\identity"
"C:\ProgramData\YuviRed\identity" /E /COPYALL /B /R:1 /W:1
robocopy "E:\YuviRed-backup\2026-11-19\agent"
"C:\ProgramData\YuviRed\agent" /E /COPYALL /B /R:1 /W:1
```

   You should see: **Failed** `0` each time. The permissions matter: the services refuse to run if these folders can be changed by ordinary users.

5. Start the services in order: `YuviRedPostgres`, `YuviRed`, `YuviRedAgent`.

How to confirm. Sign in with an account that existed at the time of the copy. **Endpoints** lists the computers of that time. **Configuration**, **License** shows your plan **Active**. Computers report again within a few minutes.

Common problems.

Problem	What to do
The `YuviRed` service stops at once	Read `C:\ProgramData\YuviRed\status\boot-status.json`: it names the boot phase and, for an error, the step that failed. Collect diagnostics and write to support.
The error names the step `provision_identity`, "the appliance private key could not be loaded"	The identity key does not open on this Windows installation (section 2.2). Roll back as below, or go to section 3.3.
"Can't locate revision" in the log	The copy was made by a newer version than the one installed. Install the newer version first.
Computers show **Offline** after the restore	Computers whose credential changed after the copy keep the previous credential for 30 days and recover on their own. Beyond that, add the computer again with a new Add Computer download.
The console says "This server needs a new license"	The hardware identity changed. See section 3.

Rollback. Stop the services, remove the four restored folders, move the four folders back from `C:\YuviRed-before-restore`, start the services.

---

3. Migration to a new server

3.1 What you must know first

- **A new computer is a new server to YuviRed.** The server records the computer's hardware identity (firmware serial numbers and product identifier) next to its installation ID at the first start and checks it at every start. When it finds different hardware (a different computer, a virtual machine restored or imported as a new machine, a motherboard change, some hypervisor upgrades), a server that held a license gives itself a **new installation ID** and locks the console with **This server needs a new license** until a license for the new ID is activated. The license is bound to the old ID. The original license does not work again on its own, even if you move the data back: moving back counts as another change. A fresh install on the new computer simply starts with its own new ID and the 7-day activation window.

- **Free servers are counted per company.** Lite codes from the website allow 2 servers per company, counted by installation ID, over all time. A replacement server uses one of the two. Write to support@yuvired.com **before** you migrate, say that the new server replaces the old, and give both installation IDs. Support can arrange a license for the new server.
- **Enrolled computers trust one certificate.** Each agent trusts exactly the certificate the old server made, and was enrolled using the server's address or name. A new server makes a new certificate. Use a DNS name for the server, not an IP address, if you can: then moving the name to the new server needs only a DNS change.
- **The identity key is tied to the Windows installation.** The server's private identity key is sealed with Windows Data Protection for this machine (or the TPM, section 2.2). Copied to a different computer it does not open, and the service stops at start with "the appliance private key could not be loaded". That is why a move keeps the **new** server's own `identity` folder (section 3.3). YuviRed has not tested moving a data copy to different hardware end to end.

3.2 The reliable path: install fresh and add computers again

This is the path the product's own limits describe: "A lost server is replaced by a new one, and each computer is added to it again with a new Add Computer download."

1. Email support@yuvired.com as described above and get a code or license file for the new server.
2. Install YuviRed on the new server (Installation guide) and create the administrator account.
3. Activate the new server (Licensing).
4. Recreate your settings by hand: users and roles, notification channels, policy, retention, organization name and privacy contact. Use the old console for reference.
5. In the new console, **Agent Deployment**, **Add Computer**: download the package and run it on each computer you want to keep. A computer that is already enrolled keeps its identity.
6. Keep the old server (switched on but not collecting) until its history is no longer needed, or export what you need first. Then uninstall it (Uninstall).

What you lose: alerts, history and accounts of the old server do not appear on the new one.

3.3 Moving the data too (not tested by YuviRed)

If you must keep the history, you can try to restore a copy (2.3) on the new server. Expect these steps, and test them on a spare computer first:

1. Install the **same version** on the new server. Do not activate.
2. Stop the three services.
3. Restore `pg` and `config` as in 2.4 (not `identity`: leave the new server's own identity, because the old key does not open here). `config` carries the encryption key, which the database needs.
4. Start the services.
5. If the console opens, expect a license screen. The installation ID the console shows is read from the database, so a copied database may show the **old** ID while the new server's identity file holds a new one; what the server does with the old license in that state has not been tested by YuviRed. Do not rely on the old license: send support both installation IDs and what you did, and activate what support sends.
6. Set the server's address with `network --name <dns name>` (Installation guide, network mode), then check **Agent Deployment**. Computers that enrolled by IP address need a new Add Computer download, because the new server's certificate differs.

If step 3 or 4 fails, the fresh-install path in 3.2 is your fallback. Keep the old server and the copy until the new one has run for a few days.

---

4. Disaster recovery

Purpose. Know in advance what you would need if the server computer were lost.

What to keep, off the server and offline or in encrypted storage:

1. A recent copy of `C:\ProgramData\YuviRed` made as in 2.3.
2. The installer you used (`YuviRed-Lite.exe`) of the version that made the copy, with its file fingerprint. You can download the same version again only while YuviRed still offers it.
3. The installation ID and the license details: the plan, and the email with the activation code or the license file. (The ID is not a secret. The activation code is single use and is of no use after activation, but keep the email as a record.)
4. A list of your computers and the server's address or DNS name, so you can add computers again.
5. A list of your administrator accounts, and the MFA recovery codes of each, kept in your password manager.
6. The notification settings you typed (mail server, webhook addresses), kept in your own records. The saved mail password and webhook secrets are in the database, encrypted, and are only useful with the matching `config` folder.

How long a rebuild takes. YuviRed has not measured a rebuild. The installer waits up to five minutes for the server to become ready. What decides the total time is how fast you can get a new Windows computer, a license for it, and how many computers you must add again by hand. Plan for hours, not minutes, and for more if you have many computers.

What the product does while you rebuild. Computers keep running their agent. They cannot send to a server that is not there, and the data of those hours is lost. They keep trying.

Steps in a disaster.

1. Get a Windows computer that meets the requirements (Installation guide).
2. Write to support@yuvired.com with the old and new installation IDs (the old one from your records).
3. Follow 3.2 (fresh path). If you have a copy and want the history back, try 3.3 afterwards on a second computer, not on the one that must serve.

Rollback. Not applicable. Keep the old computer's disk if you can, in case support or your own administrators need something from it.

---

5. Account recovery

5.1 Reset another user's password

Purpose. An administrator gives someone a new password.

Prerequisites. You are signed in as an account that has the **Team / Users** page and permission to manage users (the **Admin** role or **super admin**).

Steps.

1. In the left menu, under **Admin**, click **Team / Users**.
2. In the person's row click the key icon (hover text **Reset password**).

   You should see: a dialog **Reset Password** with **New Password** and **Confirm Password**.
3. Type a password that follows the policy (at least 12 characters; at least three of lowercase, uppercase, digit and symbol; not containing the person's name or email; not a common word; not one of their last passwords), type it again, and click **Reset Password**.

4. Give the person the password by a safe route (not by email).

You should see: a green message "Password reset".

What happens.

- The password is **temporary**. The person must choose their own at their next sign-in.
- It works for **72 hours**. If they have not signed in and changed it within 72 hours, it stops working, and the console answers "Invalid credentials", exactly as for a wrong password. Reset it again.
- Every session the person has open is ended at once.
- The change is written to the audit log.

How to confirm. The person signs in, sees the **Change Password** dialog ("Password change required"), sets a new one, and reaches the console.

Common problems.

Problem	What to do
"Password must be at least 12 characters" or "Password is too predictable"	Choose a stronger password.
The person says the temporary password is refused	More than 72 hours have passed, or the lockout (5.4) applies. Reset again.

5.2 Reset another user's multifactor authentication

Purpose. A person lost their authenticator phone and their recovery codes.

1. In the left menu, under **Admin**, click **Team / Users**.
2. In the person's row click the crossed-shield icon. It shows only when the **MFA** column reads **On**, and never on your own row.
3. Confirm the question.

You should see: "MFA disabled" and the **MFA** column reads **Off**. If their role requires MFA they enroll again at their next sign-in. The action is written to the audit log. If they still have a recovery code, they can use it at the sign-in code step without your help.

5.3 When the only administrator is locked out

Read this before it happens. **The product has no email reset, no "forgot password" link and no command that resets an administrator's password or turns off its MFA.** What exists:

Situation	What the product offers
You forgot your password and another administrator exists	The other administrator resets it (5.1).
You lost your phone and still have a recovery code	Use the recovery code in place of the 6-digit code. Each works once.
You lost your phone and recovery codes and another administrator exists	The other administrator resets your MFA (5.2).
Too many wrong attempts	Wait. After 8 wrong attempts for the same address and account within 5 minutes, further attempts are refused ("Too many failed login attempts. Please wait and try again.") until the window passes. A second, account-wide cool-down (30 wrong attempts in an hour from any addresses) never applies to a person using the console on the server itself at `https://127.0.0.1:8443`. The Lite server has no Redis, so it keeps these counters in memory: restarting the YuviRed service clears them.
You are the only administrator and you lost the password, or the phone and the recovery codes	**Nothing in the product recovers it.** There is no sole-administrator recovery: no reset command, no email reset, no recovery link. First-run setup does not reopen: it is closed for good once an administrator exists, even a deactivated one. Contact support@yuvired.com with your installation ID (shown on the license screen or in **Configuration**, **License**) and tell us what happened. Expect to be asked to prove that you control the computer. What support can then do is decided case by case and is not a product feature.

Prevent it.

1. Create a **second** administrator account for a different person, on day one (**Team / Users**, **Add User**, role **Admin**).
2. Remember that only the first account (**super admin**) can activate or revoke a license. An account made with **Admin** cannot. Keep that first account's password and recovery codes safe.
3. Store every administrator's MFA recovery codes in a password manager.

4. Use a copy of `C:\ProgramData\YuviRed` (section 2) as a last resort: a restored copy has the accounts of its day.

---

6. Access management

6.1 Roles

The **Team / Users** page offers three roles:

Role in the console	Can
Admin	Everything except activate or revoke a license: users, settings, endpoints, policy, deployment, notifications.
Analyst	Triage alerts, work on endpoints, policies and blocklist, read reports and settings. Cannot manage users or change most settings. Can ask for an enforcement action but cannot run it.
View Only	Read alerts, policies, endpoints, reports and the audit log. Cannot change anything.

The first account made at setup is a **super admin**, which is **Admin** plus the license permission. Changing a person's role replaces all their roles with the new one. You cannot change your own role, and you cannot demote or deactivate the last active administrator.

Important. Whatever the role, **any person who can sign in can read endpoint identity (computer names, user names) and raw telemetry**, and the per-user behavior summary. Give accounts only to people who should see all of it. Reading original prompt text needs its own permission and every view is audited.

6.2 Add a user

1. **Team / Users**, **Add User**.
2. Fill in **Full Name**, **Email**, **Password** (a temporary one that follows the policy), **Role** and **Department**, then **Create User**.

The Add User dialog on Team / Users. Choose the role carefully; it decides what the person can change.

You should see: the person appears in the table. Give them the password by a safe route.

6.3 Remove or deactivate a user

- **Deactivate** (keeps the account and its history): the edit icon in the person's row, set **Status**, **Save**. The person's next page load shows **Session Expired** and they cannot sign in.
- **Delete** (the bin icon): their past actions stay in the audit log by email. This cannot be undone.

You cannot delete or deactivate yourself, nor the last active administrator.

6.4 Multifactor authentication

- By default, accounts with the **Admin** role must use MFA. You can also require it for **Analyst**: **Configuration**, section **Security & Access**, row **Require multi-factor sign-in for**, tick the role, **Save Changes**.
- A person turns MFA on with the shield icon at the bottom of the sidebar (**My security**): **Enroll an authenticator app**, scan the code, type the 6-digit code,** **Enable MFA**, and save the 10 recovery codes (shown once).
- The phone and the server must have the right time. A code is only valid for the current 30 seconds.

6.5 Sessions

To end every session of a person without changing their password, an administrator can use `POST /api/v1/users/<id>/revoke-sessions` (no console button). Deactivating the account also ends sessions.

---

7. Data handling

Read the product's own statement first: **What YuviRed collects and where it goes** (the data notice), shown in the console under **Legal** and served at `https://<server>:8443/static/legal/data-notice.html`. This section is the administrator's side of it.

7.1 What is kept, and for how long

The figures below are the defaults (from the data notice and `docs/RETENTION.md`).

Data	Default	You can change it
Raw telemetry (all kinds)	30 days	Yes (**Raw telemetry**)
Daily summaries	13 months	Yes (**Daily summaries**)
AI agent signals	120 days	Yes (**AI agent signals**, 0 keeps forever)
Closed alerts not tied to an investigation	365 days	Yes (**Closed alerts**)
Prompt Monitoring event rows	90 days	Yes (**AI prompt/output events**)
Original prompt samples (only if you turned them on)	24 hours	Yes (**Original prompt samples**)
Decided adoption requests	30 days	Yes
AI agent (AMS) messages, events, ended instances	30, 180 and 90 days	Not in the console
Report files	30 days	Yes (setting `retention.report_file_days`)
Audit log	**Kept until you purge it**	Yes (**Audit log**, days; 0 means keep forever)
Endpoints	Until an administrator removes them	n/a
Server log files	5 files of 10 MB each, rotating	No

Open alerts, and closed alerts linked to an investigation, are never deleted by age. The cleanup job runs 5 minutes after the service starts and then once every 24 hours (not at a fixed time of day). The settings behind the table are named `retention.raw_days`, `retention.rollup_months`, `retention.agent_signal_days`, `retention.alert_days`, `retention.ai_content_event_days`, `retention.adoption_request_days`, `retention.audit_log_days`, `retention.report_file_days` (with `retention.report_storage_mb`, 2048), and `retention.ams_message_days`,

`retention.ams_event_days`, `retention.ams_instance_days`; the audit log lists changes to them.

7.2 Change a retention window

1. Left menu, **Admin**, **Configuration**.
2. Click the **Data & Retention** section chip.
3. Change a value. **Raw telemetry** and **Daily summaries** must be at least 1. Day-based windows accept 0 for "keep forever".
4. Click **Save Changes**. Shortening a window asks you to confirm.

You should see: "Saved 1 setting(s)". To see the effect first, scroll to **Telemetry Retention** and click **Run now (dry run)**: a green message says how many rows would be removed and nothing is deleted. **Run now** deletes (it asks to confirm and cannot be undone).

Shorter retention does not remove anything copied elsewhere, such as a copy of the data folder.

7.3 The audit log

The audit log is kept forever by default. Setting **Audit log** to a number of days turns purging on. Only do that when you copy the log elsewhere. The log can be exported through the API route `GET /api/v1/audit/export?format=csv|json|cef` (permission to read the audit log; there is no export button for it in the console, and this was not verified from the Lite console).

7.4 Delete one computer's data

1. Left menu, **Endpoints**. In the computer's row click the red bin icon (**Remove endpoint**; you need the administrator role).
2. The dialog **Remove endpoint "<name>"?** shows the real counts of what will be deleted and what stays. Read it.
3. Click **Remove endpoint**.

What removing an endpoint deletes: the computer record and everything the database keeps under it: its raw telemetry and **all its alerts** (whatever their status, even an open alert under investigation; the dialog shows both counts), plus its other per-computer rows (extension changes, rotation records and the like, which the database removes with the computer record).

What it does not delete:

- Prompt Monitoring event rows (they keep a copy of the computer name and user name). The dialog counts them as preserved.
- Records of AI agents (AMS): they are archived and detached (counted in the dialog).
- Investigations (they stay, unlinked; counted in the dialog).
- Audit log entries that name the computer (the removal itself is written to the audit log as `endpoint.delete`).
- Reports already generated, and any copy of the data folder.

The agent on that computer is **not** told and is **not** removed. It keeps running, its next check-in is refused, and it does not enroll again by itself (its Add Computer download is dead after 24 hours). Uninstall the agent from the computer (section 4 of Uninstall). "Decommission" keeps the data and frees the license seat.

There is no feature to find or erase everything about one person, or to export everything about one person. You assemble such an answer by hand, and you can delete by computer.

7.5 Exports

- **Reports** page: generate a report and download it (files kept 30 days).
- Tables: the **Export** button on a page exports the page's table to CSV. Prompt Monitoring has an **Export CSV** button for people with the permission.
- Audit log: through the API (7.3).

Exported files are copies and are outside YuviRed's retention. Handle them like the data.

---

8. Privacy controls

Control	Where	Default
Original prompt text ("raw capture", Mode 3)	**Configuration**, **Inspection Depth**, **Keep Original Prompt Samples**, **Prompt Content Capture**	Off. Needs all of: Inspection Depth set to full, raw capture on, a secure connection. Stored encrypted, 24 hours, every view audited.
Masked fragments	**Prompt Content Capture**	On: the first 4 and last 4 characters of a matched value
Per-computer exclusion	**Never keep samples from endpoints tagged**	Empty
Connectors for AI tools (AMS)	**AMS**, **Coverage & Connections**	Off. Previews up to 2 KiB and replies up to 8 KiB are stored **unencrypted**.
Extension force-install from your server	**Configuration**, **Remote Browser Extension Wiring**	Off, and without effect in this release: the Lite server does not serve a signed extension package, so the setting answers "disabled" to every computer. Installing from the browser stores automatically is planned.
Browser consent	Extension options page	Each person ticks a box, unless your administrator records consent by policy
Who sees data	**Team / Users**	See 6.1

The Legal page. Left menu, **Admin**, **Legal** shows the legal documents that ship with the product, the **Acceptance records** (who accepted the license terms at install), and an **Employee notice preview** filled from **Configuration**: **Organization Name** and **Privacy contact**. Set both before you give the notice to people. The notice text is a template: you are responsible for telling the people you monitor, and for your own legal basis. The Legal documents are also at `https://<server>:8443/static/legal/`.

Full-disk encryption. Most data (computer names, telemetry, alerts, masked samples, AMS text) is stored unencrypted in the database. Use full-disk encryption on the server and on every copy.

9. Getting help

Email support@yuvired.com. Write a clear subject for a problem, and use the exact subject "Privacy request", "Security report" or "Legal notice" for those. Postal address: YuviRed LLC, 4030 Wake Forest Rd Ste 349, Raleigh, NC 27609, United States. No response time is promised for Lite or Pro.

What to send

1. The YuviRed version and build number (console footer) and the installation ID.
2. What you did, what you expected and what happened, with the time (and time zone).
3. A screenshot of the message, with names you do not want to share removed.
4. For an install or service problem, the support collector ZIP (see Collecting diagnostics).

What the support collector is. A script installed with the server, at `C:\Program Files\YuviRed\support\collect-install-evidence.ps1`. It only reads. It writes one ZIP on the Desktop of the administrator who ran it. It leaves out the secrets file, the database password file, the HTTPS private key, the identity folder, the database files and the agent's state and settings files, and blanks known secret formats in the text it includes. It is installed on the server only; to use it on another computer, copy the file there.

What to check and remove before sending. Open the ZIP and read it. It contains computer names, Windows and product versions, file lists with names and sizes, service state and the text of the server logs. It can name your server and your administrator's Windows account. Remove anything you do not want to share. The redaction covers known formats only.

Do not send the `config` folder, `yuvired.env`, a copy of the database, an MFA code or recovery code, a password, or an activation code that is still unused.

Security reports. If you think you found a security problem, use the subject "Security report" and give the facts, not the data.

Uninstalling YuviRed

This page covers removing the YuviRed server from a computer, what stays behind, the option that erases the data, removing the agents from your other computers, and what happens to the license.

Read this first. Uninstalling the server does **not** delete your data, and it does **not** tell your other computers. Both are on purpose, and both need steps from you. **Best order: remove the agents from your computers first (section 4), then the server.** A computer whose server is gone keeps its agent, keeps trying to reach the server, and keeps applying the last policy its browser extension cached.

---

1. Remove the server (Apps and Features)

Purpose. Take the YuviRed programs and services off this computer. The data stays.

Prerequisites.

- You are an administrator on the server.
- If you may want the history later, take a copy of `C:\ProgramData\YuviRed` first (Administration, section 2). After a purge (section 3) nothing can be restored.
- You know that computers you added keep their agent running and will keep trying to reach this server (section 4). Their check-in slows down to one try every 5 minutes after repeated failures.

Steps.

1. Open Windows **Settings**, **Apps**, **Installed apps** (on Windows 10: **Apps and Features**).
2. Find **YuviRed Lite**, click the three dots (or select it), and choose **Uninstall**, then confirm.
3. Accept the Windows administrator prompt.

   You should see: the YuviRed window titled **Removing YuviRed Lite**, then **Removal Complete** with "YuviRed Lite was removed from this computer", "Your data, settings and logs were kept in C:\ProgramData\YuviRed. Installing YuviRed Lite again on this computer uses them." and "Computers you added keep their agent. Remove it on each computer

with Uninstall-YuviRed-Agent.cmd, or from Settings, Apps." A removal that needs a restart adds "Windows needs to restart to finish removing YuviRed Lite."
4. Restart the computer if the window asks you to.

To remove it without a window: `YuviRed-Lite.exe /uninstall /quiet` (the `ACCEPT_EULA` rule applies to unattended installs; whether a quiet removal is refused without it was not verified).

How to confirm.

1. In **Services**, there is no `YuviRed`, `YuviRedPostgres` or `YuviRedAgent`.
2. `C:\Program Files\YuviRed` is gone.
3. `https://127.0.0.1:8443` no longer answers.

Common problems.

Problem	What to do
"Removal Failed" or a message about restarting	Restart, then run the uninstall again. Before you restart, run the support collector if you want support to look at it (Collecting diagnostics): a restart removes the state that explains the failure.
A service is "marked for deletion"	Restart the computer. Windows removes it then.
The Apps list shows nothing	Do not delete folders by hand. Contact support@yuvired.com.

Rollback. Install the same version (or a newer one) again with `YuviRed-Lite.exe`. If `C:\ProgramData\YuviRed` was kept, the new install finds the identity, secrets and database and continues: accounts, computers and the installation ID are the same, and the first-run setup page does not open again.

---

2. What is removed and what stays

Removed by a normal uninstall (in this order: firewall rule, stop `YuviRed`, stop `YuviRedPostgres`, unregister the database service, delete the `YuviRed` service, remove the AI tool hooks, remove the version folders, the server agent's code copy, its launcher and the staging folder; the `YuviRedAgent` service and the browser registrations are removed first of all)

- The three services: `YuviRed`, `YuviRedPostgres` and `YuviRedAgent`.
- The programs in `C:\Program Files\YuviRed`, including the version folders of the application, the server agent's code copy (`agent-code`) and launcher, and the support collector.
- The Windows Firewall rule `YuviRed Server (TCP 8443)`, if you had turned on network mode.
- The browser connector registrations (the registry entries that let the extension talk to the agent) and the Event Viewer sources `YuviRed` and `YuviRedAgent`.
- The hooks YuviRed added to other AI tools' managed settings (Claude Code, Codex, Gemini CLI, Cursor, Copilot CLI, Windsurf), where a connector had been turned on.

Nothing is sent to YuviRed or to your computers: the removal does not notify anyone.

Stays after a normal uninstall

What	Where	Why it matters
The data folder	`C:\ProgramData\YuviRed`	The database (computers, alerts, accounts, history), the generated secrets and encryption key, the HTTPS certificate and private key, the server's identity and installation ID, and the logs. **All of it is readable by anyone who can open the folder as an administrator, and most of it is not encrypted.**
The server agent's log	`C:\ProgramData\SyntinelRed\logs`	The folder name is the product's earlier name; it will change in a later release.
Windows' Visual C++ runtime	Windows	Other programs may use it.
Windows service profiles	`C:\Windows\ServiceProfiles`	Windows creates a profile for each service account and may leave it (not verified). It holds no YuviRed data.
The browser extension	In each browser	Remove it in `chrome://extensions` or `edge://extensions`.
The agents on other computers	Every computer you added	See section 4.
Copies you made	Anywhere	YuviRed knows nothing about them.

The installer never deletes `C:\ProgramData\YuviRed`. Logs are kept even after a purge (section 3).

---

3. Erase the data (the purge option)

Purpose. Remove the server's data as well as its programs: when you retire a server, return a computer, or must meet a deletion duty.

What it is. A command-line option, `uninstall --purge`, that you run **by hand** from an administrator PowerShell. The installer never uses it, and the Apps page has no checkbox for it. It runs the whole removal plan of section 2 (it stops and deletes the three services, removes the firewall rule, the version folders and the browser registrations) and then erases the data.

What it erases.

- The identity files (`appliance.json`, the private key file `appliance_private.dpapi`, `.key` or `.tpm`, `entitlement_lease.json` and the `.provisioned-ok` marker) are **overwritten with random bytes and then deleted**.
- The folders `identity`, `config` (generated secrets, encryption key, certificate and key), `pg` (the database), `backups`, `scratch` and `status` are **removed**. Those are removed, not overwritten: on a solid-state drive deletion does not guarantee the old data cannot be recovered. For that, use full-disk encryption on the server from the start, and erase the whole drive when you retire it.

What it does not erase.

- `C:\ProgramData\YuviRed\logs`. Logs are kept.
- The server agent's folder `C:\ProgramData\YuviRed\agent` (its settings, state and **its enrollment credential**, a secret), the file `C:\ProgramData\YuviRed\update.journal`, and the folder `C:\ProgramData\YuviRed` itself.
- `C:\ProgramData\SyntinelRed` (the server agent's log; the earlier name, see section 2).
- **Any backup, copy of the data folder, export or report file you or anyone else made, anywhere.** The purge cannot see them.
- Anything the agents already stored on your other computers.

If you need those too, delete them yourself afterwards (carefully: `agent` and `logs` first, then the folders) and empty the Recycle Bin.

Prerequisites.

- An administrator PowerShell.
- You have decided you do not need the data. **This cannot be undone.**
- The server is installed. The command is part of the installed programs, and the normal uninstall removes them, so the purge has to be run **before** the normal uninstall in section 1. **YuviRed has verified this command only from its code and its tests, not by running it after a normal uninstall or on a real customer server.** If you purge first, check afterwards that Installed apps still lists YuviRed Lite and remove it as in section 1.

Steps.

1. Open PowerShell as administrator.
2. Run the purge command (nothing is on PATH; both files are named in full):

   ```
   & "C:\Program Files\YuviRed\runtime\python.exe" "C:\Program
   Files\YuviRed\runtime\yuvired_service.py" uninstall --purge
   ```

   You should see: a line "browser host removed: ..." and then one line `uninstall steps:` listing the steps that ran, in this order: `remove_firewall_rule, stop_yuvired, stop_postgres, unregister_postgres, delete_yuvired_service, remove_ai_tool_connectors, remove_versions, remove_agent_code, remove_agent_launcher, remove_staging, erase_identity, remove_data, remove_bytecode_caches`. A step that failed in part carries "(failed: ...)" with the error name, for example `remove_data (failed: PermissionError)`; the later steps still run.
3. Remove the program as in section 1 if you have not. (After the purge the services are already gone; Installed apps still lists YuviRed Lite until you remove it.)
4. Delete any folder named above that you want gone.

How to confirm. `C:\ProgramData\YuviRed` holds at most `agent`, `logs` and `update.journal`. The folders `pg`, `config`, `identity` and `status` are gone.

Common problems.

Problem	What to do
A step reads "(failed: PermissionError)"	A file was still open (a service still stopping, an antivirus scan). Check in Services that `YuviRedAgent`, `YuviRed` and `YuviRedPostgres` are stopped or gone, then run the command again.
A folder is still there afterwards	Restart the computer and run the command again, or delete the folder by hand as administrator and empty the Recycle Bin.

Rollback. None. After a purge the server's history, accounts and encryption key are gone. A new install on the same computer starts from the beginning with a **new installation ID** and a new 7-day activation window, and needs a new activation (section 5).

---

4. Remove the agents from your other computers

Uninstalling the server **does not notify the computers**. Their agents keep running, keep trying to reach the server (and writing failures to their log), and keep their copy of the server's address and trust. Remove them yourself, ideally **before** you remove the server:

Purpose. Stop monitoring on a computer that you no longer watch.

On these computers the agent's Windows service is named `SyntinelRedAgent` (shown in Services as **YuviRed Endpoint Agent**) and its data and logs live in `C:\ProgramData\SyntinelRed` (the earlier name, see section 2). Only the server's own agent is named `YuviRedAgent`.

Steps, per computer.

1. Sign in to the computer as an administrator.
2. Remove the agent in one of two ways:
   - Open **Settings**, **Apps**, **Installed apps**, find **YuviRed Agent**, and **Uninstall**. Or:
   - If you still have the folder you unzipped from **Add Computer**, run `Uninstall-YuviRed-Agent.cmd` in it as administrator. It removes the agent, checks what is left, and deletes the download's own files from that folder: `enrol.yuvired` (the key to your server), `YuviRed-Agent.msi`, `Install-YuviRed-Agent.cmd` and `README.txt`.
3. Remove the browser extension in `chrome://extensions` or `edge://extensions` (**Remove**), if it was loaded there.
4. Delete any copy of the Add Computer zip and its `enrol.yuvired` file from that computer and from shared folders.

You should see: the agent is gone from Installed apps, there is no `SyntinelRedAgent` service, and `C:\ProgramData\SyntinelRed` is gone.

If the server is still running, also remove the computer from the console first: **Endpoints**, the red bin icon (**Remove endpoint**), which frees the seat and deletes that computer's data (Administration, section 7.4).

Common problems.

Problem	What to do
The Add Computer folder is gone	Use **Installed apps**.
Uninstall is refused	Run it as administrator.
Settings says the uninstall worked, but something is left	An uninstall from Settings does not report what its cleanup could not remove. Check that the `SyntinelRedAgent` service and `C:\ProgramData\SyntinelRed` are gone. `Uninstall-YuviRed-Agent.cmd` checks for you and names what is left.

Rollback. Add the computer again with a new **Add Computer** download from a server.

---

5. What happens to your license

- The license is **bound to the installation ID** of the server it was activated on. The installation ID is kept in `C:\ProgramData\YuviRed\identity\appliance.json` and in the database. If you kept the data folder and reinstall on the **same computer and the same Windows installation**, the server keeps its ID and stays activated.
- If you purged, or deleted the data folder, a new install has a **new installation ID** and is not activated. It gets a new 7-day activation window.
- **A license is not moved by you.** To use your license, or your company's Lite allowance, on another server, contact **support@yuvired.com** and give the old and new installation IDs. For Lite, a company may have 2 servers, counted by installation ID over time: a replacement server uses one. Pro and Enterprise licenses are handled by YuviRed.
- Uninstalling does not tell YuviRed. The server stops sending its daily check-in. YuviRed sees a server that is not reporting. Nothing is revoked or refunded by uninstalling.

To confirm the old server is gone from your records, write to support@yuvired.com. There is no screen on which you remove a server from your company's allowance.

YuviRed on your work computer: a guide for the people who use it

Your organization put YuviRed software on this computer to see how AI tools are used at work and to keep sensitive information from being sent to AI services by mistake. This page explains what is on your computer and what it does. It does not replace the notice your organization gave you.

What runs on your computer

1. **The YuviRed Endpoint Agent.** A background Windows service named "YuviRed Endpoint Agent". It starts with Windows, runs as the system, and reports about once a minute to **your organization's own YuviRed server**. It does not report to YuviRed LLC, the company that makes the software.
2. **The YuviRed AI Prompt Monitor** (only if your organization loaded it into your browser; in this version it is added by hand, browser profile by browser profile, not installed automatically). A Chrome or Edge extension that works on AI chat sites such as ChatGPT, Claude, Gemini and Copilot.

What the agent reports

- Your computer name, your Windows user name, one IP address and the Windows version.
- Which AI programs are running (name, location, memory and processor use).
- Web connections of programs to web ports, with the address, its name and the program. This covers every program, not only AI. The amount of data is not recorded.
- Visits to AI websites in Chrome, Edge, Brave and Firefox: a count and the last time, for the last 24 hours. The agent reads a temporary copy of your browser history on the computer to find these. Other sites, page addresses and page titles are not sent.
- AI-related browser extensions, installed AI software, AI services running on the computer, and AI entries in the Windows name cache.
- Settings and names of AI tools such as coding assistants (secrets are removed; environment variable names, never values).

What is not collected

Keystrokes, screenshots, the clipboard, passwords, the contents of your documents, the contents of web pages, and chat transcripts. The agent does not read the text of your prompts or the AI's replies.

What the browser extension does

On 15 AI sites (ChatGPT, Claude, Gemini, Microsoft Copilot, Perplexity, Microsoft 365 Copilot, DeepSeek, Grok, Le Chat, Meta AI, Poe, Google AI Studio and HuggingChat) it checks what you type, paste or attach, and what the AI replies, for sensitive information (passwords, keys, personal data, card numbers). The check of what you send runs **before it is sent**, and it also looks at a draft while you type. The check runs in your browser. Your organization sees the kind of finding, a masked fragment (at most the first 4 and last 4 characters; for a password only the beginning), what happened, the site, and your browser. It does not send the text. Only if your organization has switched on full capture, which the notice must tell you, is the first part of a flagged prompt or reply stored on its server, encrypted, for 24 hours.

In "monitor only" mode you see nothing on the page, but findings are still recorded. Nothing is recorded or reported until consent is given. In this version that means you tick the box **I consent to local inspection of prompts and attachments on the monitored sites above** on the extension's options page and choose **Save**; the page opens by itself when the extension is loaded and consent has not been given. If your organization ever sets consent through browser policy instead, the same page shows "Policy consent value: true (consent given by policy)" and you are not asked.

If a warning appears

A box titled **Potentially sensitive information detected** says what kind of information was found and asks "Are you sure you want to send this information to <the AI service>?".

- Choose **Go back and edit** (the default) to take the information out. Nothing is sent.
- Choose **Send anyway** only if you are allowed to share it. Your choice is recorded. Closing the box or pressing Escape does not send.

If a block appears

A box titled **Prompt blocked: sensitive information detected** lists what to remove. Remove it and send again. You cannot override a block. A small "ref" code at the bottom of a warning or block lets your administrator find the event if you ask for help.

If a draft notice appears

While you type, a small notice may say "Sensitive information detected in your draft" and that your draft has not been sent. It has a close (x) button. Nothing is sent by the notice itself; it tells you what will happen when you press Send.

The options page

Open it from the browser's extensions page, Details, Extension options. It shows what is inspected, where reports go, and a Connection status ("Connected to YuviRed Lite" when it works). A **Copy status for support** button gives your administrator version and status numbers, with no prompt text.

Where to read the notice

- On this computer: `C:\Program Files\YuviRed Agent\NOTICE.txt`.
- On your organization's server: `https://<your server>:8443/static/legal/notice.html`. Ask your administrator for the address. The data notice (`data-notice.html`) gives more detail.

What you cannot change

You cannot stop, pause or remove the agent without administrator rights, and you cannot change your organization's policy or the rules. A computer whose agent stops reporting shows as Stale and then Offline on your organization's console, and a browser whose extension stops reporting shows as Stale there. The extension's options page lets you narrow the list of monitored sites or lower its capture level; your organization's console is not told when you do. If the browser extension is ever installed by browser policy (not in this version) it has no Remove button and cannot be turned off.

The agent's Windows service is named `SyntinelRedAgent` and its files are under `C:\ProgramData\SyntinelRed` (the product's earlier name; it will change in a later release).

Questions and requests

Contact your organization's privacy contact (shown on its notice page, or ask your administrator). You can ask what is held about you, to correct or delete it, or to object, as the

law where you work allows. YuviRed LLC does not hold your monitoring data; its address for software questions is support@yuvired.com.

Console pages reference

Every page of the YuviRed console in menu order: what it is for, its main controls, and who can see it.

This reference covers YuviRed 1.0 (app 1.0.0, installer 1.0.322.1608, endpoint agent 1.8.2, browser extension 0.5.5) for Windows. The console is one web page served by the YuviRed server, normally at `https://<server>:8443/`. Names below are exactly as they appear on screen.

How to read this page

- **Who can see it** has two parts. The left menu hides some items by role (the console decides from the account's role: Admin, Analyst or View Only). Separately, the server checks a permission on each request the page makes. The pages below say which permission the page's data needs. A page can open for a role and still show an error or an empty list if the server refuses the data. See roles-and-permissions.md for the permission table.
- "Admin" in this page means an account whose role is **Admin** (the first account created at setup, and any account you create with the **Admin** role). "Analyst" and "View Only" are the other two roles in the console.
- AIDR (AI Detection and Response) is the first product area. AMS (Agent Monitoring System) is the second. A switch at the top of the left menu (**AIDR** and **AMS**) changes which menu is shown. Both share the same top bar, sign-in and license.
- The guide chapters named in "Guide" are in `docs/customer/admin-guide/`.
- **Present but inactive on YuviRed Lite.** The console is one program for every edition, so a few controls that belong to the Docker appliance edition appear on a Lite server and change nothing there. They are named in the page that shows them: the **Deployment Tokens** panel on Agent Deployment, and the settings **Automatic Agent Updates**, **Remote Browser Extension Wiring**, **Extension distribution mode**, **Chrome Web Store extension ID** and **Allow extension wiring over http agents** on Configuration (see settings.md, section 9). The panel **Deploy Endpoint Agent** is hidden on Lite outright.

The top bar and global controls

These appear on every page after sign-in.

Control	What it does
Page title	Shows the current page's title (for example "Alert Management"). The browser tab title mirrors it, prefixed "YuviRed".
Role badge	Shows **Admin**, **Analyst** or **View Only**.
Clock	The current time.
Filter	Opens a filter bar for the current table: a search box ("Search the current page..."), a status drop-down ("All Status"), a risk drop-down ("All Risk / Severity": Critical, High, Medium, Low) and **Clear**. Hidden on Alerts, which has its own controls.
Export	Exports the current page to CSV.
Moon icon	Switches between light and dark theme. The choice is kept in the browser.
Refresh	Reloads the current page's data now.
License banner	Appears when the server is inside its 7-day activation window ("This YuviRed server is not activated. You have N days left to activate it. After that, the console is unavailable until it is activated.", with **Activate now**; the count reads "1 day" or "less than 1 day" at the end), or 30, 14 and 7 days before a paid license ends (with **License**). See statuses-and-errors.md.
Update banner	Appears when the daily license check-in names a newer YuviRed version: "YuviRed <version> is available. This server runs <version>. Nothing is installed automatically." It has **Download** (when a link was given) and **Dismiss**. Dismissal is kept per browser.
Critical alert banner	Shows the most recent open critical alert, with **Investigate** (opens Alerts) and a dismiss cross.
Live status (bottom of left menu)	**Connecting live updates...**, **Monitoring...** (live), **Reconnecting live updates...**, or **Live updates paused; refreshing every 30s**.
Account menu (click your name, bottom left)	**My Account**, **Change Password**, **Authentication / MFA**, **Recovery Codes**, **Active Sessions** (greyed out, marked "soon": this build does not track individual sessions), **Sign Out**.
Collapse	Collapses the left menu to icons. The choice is kept in the browser.

Guide: administration.md.

First-run setup and sign-in screens

Shown before the console, not in the menu.

Screen	Purpose and controls
Welcome to YuviRed / Create your administrator account	Shown once, on the server itself, while no administrator exists. Fields: **Full Name**, **Email**, **Password**, **Confirm Password**; a required box "I accept the YuviRed End User License Agreement (version 1.0) and I have read the Monitoring and Data Notice"; button **Create Account**. The account becomes the owner (the `super_admin` role).
Sign in	**Email**, **Password**, **Sign In**, and **Sign in with SSO** when single sign-on is set up.
Verification code	**Verification code** (6-digit authenticator code or a recovery code), **Verify**, **Back**. The step lasts 2 minutes.
Forced MFA enrolment	QR code and **Secret**, **Code from your app**, **Enable & Sign In**. Then recovery codes with **Download .txt**, **Copy** and **Continue**. Shown for a role that must use MFA (multifactor authentication).
Session Expired	Blocking dialog with **Log In Again**.
YuviRed license	Activation screen: **Installation ID** (**Copy**), **Activation code** (**Activate**), **License file** (**Upload license file**, **Download license request**), **Sign out**. Opens by itself when the server is locked.

Guide: setup.md and licensing.md.

AIDR menu

Monitor section

Dashboard (page title "AI Threat Intelligence")

- **Purpose:** the first page after sign-in: a summary of endpoints, AI tools, open alerts and AI network activity.
- **Main controls:** five cards (**Endpoints**, **AI Tools**, **Shadow AI**, **Open Alerts**, **Data Egress**) that open the related page; panels **Risk Distribution**, **Active Policies** (**Manage**), **AI Traffic**, **Top AI Destinations** (**View all**), **Recent Alerts**, **Active**

Endpoints and a **Detected AI Tools** table (Tool, Vendor, Category, Status, Endpoints, Severity, Last Seen).
- **Who can see it:** every signed-in role. The data route needs only a signed-in account.
- **Guide:** aidr.md.

AI Tools (page title "AI Tool Inventory")

- **Purpose:** every AI tool seen on the monitored computers, each with its effective classification. The catalog of known tools and their defaults is managed from here.
- **Main controls:** the inventory table; **Manage Catalog** (Admin only); a classification can be changed by roles that hold the `policies.write` permission (Analyst and Admin).
- **Who can see it:** every signed-in role (list needs only a signed-in account). **Manage Catalog** is Admin only.
- **Guide:** aidr.md.

Alerts (page title "Alert Management")

- **Purpose:** findings from the policy engine and the prompt guard, with status, severity and assignment. Click an alert to investigate it.
- **Main controls:** sort (**Recent activity**, **Risk first**); filters for severity, status (open, acknowledged, investigating, resolved, false_positive, suppressed), classification, category, tool, time (**Last 24 hours**, **Last 7 days**, **Last 30 days**), endpoint and user; **Search alerts...**. The menu item shows a badge with the open count.
- **Who can see it:** shown to every role in the menu; the data needs `alerts.read` (Admin, Analyst, View Only; not the Developer role). Changing a status or adding a note needs `alerts.write` (Analyst and Admin).
- **Guide:** alerts-and-integrations.md.

Prompt Monitoring

- **Purpose:** prompt-level events from the browser extension (the sensitive-data guard): what was observed, warned about or blocked, with a redacted excerpt. Also shows protection status per browser and the active Prompt Protection policy.
- **Main controls:** an **Alerts** link, filters, **Export CSV**, the table "Suspicious content events" (ID, When, Type, Severity, Outcome, Category, Rule, Conf., Endpoint / User, Provider, Excerpt (redacted)); a browser protection card (see the Protection status list in statuses-and-errors.md).
- **Who can see it:** the data needs `alerts.read`. **Export CSV** and original prompt samples need `alerts.write` and `dlp.view_raw_content` (Admin roles only for the latter).

- **Guide:** sensitive-data.md and extensions.md.

Network (page title "AI Egress & Destinations")

- **Purpose:** AI domains your computers resolve and reach, with volume. Gives you a list to hand to a firewall, proxy or DNS device. YuviRed does not block network traffic itself.
- **Main controls:** panel **AI Destinations** with **Refresh** and an export menu (the denylist for your device, discovery only); panel **Blocklist Feed URL** with **Generate**, **Rotate** and **Revoke** for a token-protected URL your device can poll.
- **Who can see it:** summary needs only a signed-in account. Exporting the denylist needs `blocklist.manage` (Analyst and Admin). Generating, rotating or revoking the feed token needs `config.write` (Admin).
- **Guide:** aidr.md.

Endpoints (page title "Endpoint Monitor")

- **Purpose:** every enrolled computer: identity, last check-in, liveness, risk and agent version. Also lists browser-extension lifecycle changes.
- **Main controls:** table "Enrolled Endpoints" with filters (search, OS, status Online/Stale/Offline, risk, agent version) and sort; per-row detail, refresh-now and **Remove endpoint** (with a confirmation that states what is deleted); panel **Extension Changes** with a change-type filter (installed, removed, enabled, disabled, version change, permission change, manifest change).
- **Who can see it:** the list, endpoint detail and raw telemetry need only a signed-in account (see roles-and-permissions.md). Refresh-now and log collection need `endpoints.operate`. Removing, decommissioning and reactivating need `endpoints.manage` (Admin).
- **Guide:** endpoints.md.

AI Map (page title "AI Map")

- **Purpose:** a map of who is using AI, which applications, and which providers are involved.
- **Main controls: Activity** and **Detections** mode buttons; search ("Search actors, apps, providers, domains..."); time range (Last 5 minutes to All time, Custom range); **Refresh**; paging arrows. A "Save view" button is present but not implemented (it shows "Saved views are not implemented yet").
- **Who can see it:** every signed-in role (needs only a signed-in account).
- **Guide:** aidr.md.

User Behavior (page title "User Behavior Analytics")

- **Purpose:** per-user AI usage and risk, attributed from process telemetry.
- **Main controls: Refresh**; panels **Highest-Risk Users**, **Behavior Patterns**, **Top AI Users by Activity**.
- **Who can see it:** every signed-in role (needs only a signed-in account).
- **Guide:** aidr.md.

Investigations

- **Purpose:** incidents: group related alerts and track them from triage to resolution.
- **Main controls: New Incident**; list "Active Incidents"; a detail pane with a status selector (open, investigating, contained, resolved).
- **Who can see it:** reading needs only a signed-in account. Creating, changing and deleting incidents needs `alerts.write` (Analyst and Admin).
- **Guide:** alerts-and-integrations.md.

Artifact Analysis

- **Purpose:** look up or submit one file (an "artifact") to an outside analysis provider. The feature is off by default (`artifact.enabled` is off), and with the default policy nothing leaves the server.
- **Main controls:** panel **Analyze an Artifact** (**Browse**, **Analyze Artifact**, **Choose another file**); panel **Recent Analyses**. Settings live under Configuration, Artifact Analysis.
- **Who can see it:** the menu item shows only to accounts holding `artifact.read` (Admin, Analyst, View Only). Starting an analysis needs `artifact.analyze` (Analyst and Admin). Deleting history and testing providers need `artifact.manage` (Admin).
- **Guide:** aidr.md.

Govern section

Policies (page title "Security Policies")

- **Purpose:** detection rules that decide what becomes an alert.
- **Main controls:** the rule list with an enable/disable toggle; **New Policy** (Admin only).
- **Who can see it:** every signed-in role can list. Creating, editing and deleting need `policies.write` (Analyst and Admin); the **New Policy** button is shown to Admin.
- **Guide:** aidr.md.

Enforcement (page title "Enforcement: Response Actions")

- **Purpose:** controlled response actions that follow a fixed life cycle: request, dry-run, approve, run, roll back. An Analyst may request; only an Admin may approve, run or roll back.
- **Main controls: New Action**; the action list with a status for each action (see statuses-and-errors.md).
- **Who can see it:** hidden from View Only. Requesting and viewing need `response.read` (Analyst and Admin); approve, reject, run, recall, roll back, disable, archive and delete need `response.enforce` (Admin only).
- **Guide:** alerts-and-integrations.md.

Blocklist (page title "Blocklist Manager")

- **Purpose:** domains, addresses and processes you want treated as blocked in detection.
- **Main controls: Add Entry**; a remove control per entry.
- **Who can see it:** Admin only in the menu. Changes need `blocklist.manage` (Analyst and Admin at the server).
- **Guide:** aidr.md.

Reports

- **Purpose:** download an Executive Summary (about 2 to 4 pages, PDF) or an Auditor's Report (PDF plus a ZIP package), and see a live overview. It is also reachable from the AMS menu under Related.
- **Main controls:** tabs **Reports** and **Live overview**; the report builder; **Report history** with **Refresh** and a **Show reports by everyone** box (Admin only); on Live overview: **Refresh**, **CSV**, **HTML**, **Print / PDF**.
- **Who can see it:** needs `reports.read` (all five built-in roles). Seeing, downloading and deleting reports made by other people needs `reports.manage` (Admin).
- **Guide:** reports.md.

Deployment Campaigns

- **Purpose:** tracks the roll-out of agents from package delivery to `fully_ready`. It tracks; it does not push software. On YuviRed Lite every Add Computer download appears here as a campaign named "Add Computer" with the date and time (UTC); **Revoke** stops that download, and computers it already added keep reporting.
- **Main controls:** status filter (All statuses, Draft, Active, Paused, Completed, Revoked, Expired), **Refresh**, **New Campaign** (Admin only), lifecycle actions on a campaign.

- **Who can see it:** `deployment.read` (Admin, Analyst, View Only). Creating or changing a campaign needs `deployment.manage` (Admin).
- **Guide:** endpoints.md.

Deploy section (Admin only)

Agent Deployment

- **Purpose:** put the endpoint agent on more computers.
- **Main controls on YuviRed Lite:** panel **Add Computer** ("Download an installer for another Windows computer on this network"), which shows the server address, **Computers: N of 10** and **AI agents: N of 5**, the server certificate's SHA-256 fingerprint to compare, the field **Computers this download may enroll** and the button **Download installer** (the download is valid for 24 hours); panel **Deployment Tokens** (**New**); **Agent Monitoring Scope**; **Recent Check-ins**; **Endpoint Credential Rotation** with **Refresh**, **Select all shown**, **Clear** and **Rotation Operations**. The panel **Deploy Endpoint Agent** (script commands, **Test Connection**, **Interactive**, **Unattended (RMM / Intune / SCCM)**) is hidden on Lite, because those commands do not produce a working agent on Lite.
- **Present but inactive on Lite:** the **Deployment Tokens** panel still shows. A token made with **New** is meant for the install commands of the hidden panel, so on Lite nothing in the console uses it; each Add Computer download carries its own short-lived enrollment credential instead (it appears under Deployment Campaigns as "Add Computer" with the date and time).
- **Who can see it:** Admin only in the menu. Add Computer, tokens and rotation need `endpoints.manage`. The button **This server was moved or restored** (shown only after a hardware change) needs `license.manage`, so only the first account can use it.
- **Guide:** endpoints.md.

Adoption Requests (page title "Endpoint Adoption Requests")

- **Purpose:** review computers asking to take over an existing endpoint's identity (for example after a reinstall). Approve or reject each one.
- **Main controls: Refresh**, and **Approve** or **Reject** per request. Request statuses: pending, approved, rejected, expired.
- **Who can see it:** Admin only. Needs `endpoints.manage`.
- **Guide:** endpoints.md.

Admin section (Admin only)

Configuration (page title "Platform Configuration")

- **Purpose:** every setting an administrator may change, plus the license panel. Described in full in settings.md.
- **Main controls:** panel **License** (subtitle "Plan, limits and license for this server"; buttons **Activate or import license**, **Download license request**, **Reload**; the plan badge, **Licensed to**, **Plan**, **Expires**, **Serial**, **Installation ID**, **Last check-in**, and the counters **Computers used** and **AI agents used**); panel **Platform Configuration** (**Discard Changes**, **Save Changes**, a chip per section, **Expand all**, **Collapse all**, **Advanced**, **Diagnostics**); **Notifications: Test Send**; **Classifications** (**Add Classification**, **Reload**); **Telemetry Retention** (**Reload**, **Run now (dry run)**, **Run now**); **Daily Digest**. If the panel cannot load every detail it says "Could not load all license details. Choose Reload."
- **Present but inactive on Lite:** in the System section, **Automatic Agent Updates** and **Remote Browser Extension Wiring** (and the advanced extension distribution rows) can be switched but change nothing on a Lite server. See settings.md, section 9.
- **Who can see it:** Admin only in the menu. Reading settings needs `config.read`; saving needs `config.write`; the license actions need `license.manage` (only the `super_admin` role, which is the first account).
- **Guide:** administration.md and licensing.md.

Legal (page title "Legal Documents")

- **Purpose:** the documents that ship with this product (EULA, privacy notice, notice, data notice, third-party notices), the acceptance record, and a preview of the employee notice.
- **Main controls:** panels **Legal documents** (one row per document, with version and effective date), **Acceptance records**, **Employee notice preview** (your organization name and privacy contact are filled in from Configuration), **Copy link**.
- **Who can see it:** Admin only in the menu. The acceptance record route needs `config.read`.
- **Guide:** setup.md.

Help (page title "Help and Support")

- **Purpose:** links to the documentation and the support contact, with the two facts support asks for.

- **Main controls:** one panel, **Help** ("Documentation and support"), with seven items in this order: **Documentation (offline copy on this server)** (opens the copy served by this server at `/static/docs/`, no sign-in needed), **Online documentation** (https://yuvired.com/docs), **Quick start**, **Troubleshooting**, **Collect diagnostics for support** (the offline pages for each), **Legal documents** (opens the Legal page), and **Contact support: support@yuvired.com (include your installation id and version)**, followed by the line **Installation id: ... Version: ...** read from this server (it says "unavailable" when the server does not answer; the links still work).
- **Who can see it:** Admin only in the menu. The installation id and version come from the license status route, which needs only a signed-in account.
- **Guide:** troubleshooting/index.md.

Team / Users (page title "Team & User Management")

- **Purpose:** create users and set their role, department and active state; reset passwords; turn off another user's MFA; sign a user out everywhere.
- **Main controls:** panel **Team Management** with **Add User** (Full Name, Email, Password, Role: Analyst, View Only, Admin, Department; **Create User**), a table (Name, Email, Role, Department, Last Login, Status, MFA) with edit, reset-password, turn-off-MFA and delete icons. You cannot change your own role or deactivate your own account.
- **Who can see it:** Admin only. Needs `users.read` and `users.manage`.
- **Guide:** administration.md.

Debug / Reset

- **Purpose:** live counts, the last 5 telemetry payloads, and a **Data Reset** that clears test data.
- **Main controls:** panel **Platform Debug Status** ("Admin only, live counts and last telemetry") with **Refresh** and **Last 5 Telemetry Payloads Received**; panel **Data Reset** ("Clear test data. Users and tokens are preserved") with the text "Select what to clear, then click Reset. This cannot be undone.", tick boxes **Telemetry records**, **Alerts**, **Endpoints (tick Telemetry records and Alerts too)**, **AI Tool catalogue**, and **Reset Selected**. Treat it as destructive.
- **Who can see it:** Admin only. Needs `debug.access`.
- **Guide:** administration.md.

System Logs

- **Purpose:** diagnose problems.

- **Main controls:** panel **Feature Status** ("Admin only, what's working right now") with **Refresh**; panel **Something Went Wrong?** (what to send your YuviRed contact; shown on Lite only); **Endpoint Agent Logs** with an endpoint drop-down, **Collect logs** and **Refresh** (logs are collected only when you ask; bundles are kept 7 days, at most 3 per endpoint; needs agent 1.4.9 or later); **Live Logs** with a level filter, **Autoscroll**, **Pause**, **Clear**.
- **Who can see it:** Admin only. Needs `debug.access`; collecting agent logs needs `endpoints.operate`.
- **Guide:** administration.md.

AMS menu

AMS pages are reached with the **AMS** switch. Each page opens with a one-line purpose. All AMS read pages need `ams.read` (Admin, Analyst, View Only). Acts on a live AI agent need admin-tier permissions: `ams.manage`, `ams.message`, `ams.approve`, `ams.policy`, `ams.stop`, `ams.bulk`. Guide for the whole group: ams.md.

Menu item	One-line purpose (as shown)	Main controls and notes
Overview	"What is happening and what needs your attention"	Fleet-level summary of AI agents.
Agents	"Find an agent and inspect or control it"	**Agents** roster button, **Search agents...**, a roster filter (attention, unowned, paused, or an approval state), an All Agents card and a per-agent card with tabs. Badge (amber) counts agents awaiting a governance decision. Approval states: Unreviewed, Under review, Approved, Restricted, Unauthorized, Referred to AIDR.
Governance	"The rules agents must follow, and whether they are active"	The control policy and its state.
Alerts & Actions	"Findings, interventions and their verified results"	Badge (amber) counts pending approvals.
Coverage & Connections	"What AMS can see and control, and how to close gaps"	Coverage gaps and connectors.
Reports (under Related)	Opens the same Reports area as AIDR	See Reports above.

Pages and features that do not exist

- **Active Sessions** in the account menu is present but disabled: this build does not track individual sessions.
- **Save view** on the AI Map is not implemented.
- There is no **Decommission** control in the Agent Deployment page; **Remove endpoint** on Endpoints is what frees a computer seat (and it deletes that computer's history).
- There is no Backup page. Lite ships no backup command.

- The console cannot change the log level or any setting marked server-only in settings.md.
- There is no console setting for the session length (8 hours) or for the sign-in lockout rules; see roles-and-permissions.md.
- Nothing in this release installs the browser extension into Chrome or Edge by itself. The Prompt Monitoring page's setup steps on Lite describe the manual route (Load unpacked from `C:\Program Files\YuviRed Agent\extension`). Automatic installation from the browser store is planned for a later update.

Related

- roles-and-permissions.md
- settings.md
- statuses-and-errors.md
- editions-and-limits.md

Settings reference

Every setting the YuviRed console exposes under Admin, Configuration, grouped the way the console groups them, with defaults, allowed values, who may change them and what to watch for.

Covers YuviRed 1.0 for Windows (app 1.0.0, agent 1.8.2, extension 0.5.5). Defaults come from the server's settings table (`SETTINGS_SCHEMA`); labels and choices come from the console. Where the two differ, the console's wording is used for the label and the server's value is the default.

How settings work

- **Where:** Admin, then **Configuration**, panel **Platform Configuration**. Nine sections (below), then a collapsed **Advanced settings: uncommon, granular controls** box and a **Diagnostics** box (read only). Chips at the top jump to a section. Rows marked "Advanced" in this page sit inside the collapsed advanced box of their section.
- **Saving:** change one or more rows, then click **Save Changes**. Each changed setting is saved on its own and every save is written to the audit log. A green toast says `Saved N setting(s)` . If the server refuses one value you see `Save failed - <key>: <reason>` ; the other values still saved. **Discard Changes** reloads the saved values. Leaving the page with unsaved edits asks first.
- **Confirmation boxes:** for a change that reduces detection, privacy or retention, or changes fleet behavior, the console asks you to confirm before saving. These are listed as "Warning" below.
- **Who may change settings:** reading a setting needs the `config.read` permission. Changing one needs `config.write` . In the built-in roles, `config.write` belongs to the two Admin roles only (the first account, and any account created with the **Admin** role). Analysts and View Only accounts can read most settings through the server but the Configuration page is Admin only in the console menu. Settings under `artifact.` also need `artifact.manage` (Admin) to read or write.
- **When a change takes effect:** the server keeps no copy of a setting between requests. Each feature reads its setting from the database when it next needs it, so no setting in this page asks for a restart of the YuviRed service. The practical delay depends on the feature: a sign-in rule applies at the next sign-in; a notification rule at the next finding; the retention job at its next run; a background analysis at its next cycle (the "frequency" settings below); an agent setting at the agent's next configuration pull. Where a delay is

longer than "next use" it is stated in the row. Treat "no restart needed" as a code observation: it was not tested for each row individually.

- **Secrets:** password, webhook and key fields show as masked. They are overwritten only if you type a replacement. Accounts with `config.read` but not `config.write` see only "configured" or "not set".
- **Server-only settings:** some values cannot be changed from the console. They are listed at the end under "Server-only and hidden settings".
- **Units:** "UTC" is Coordinated Universal Time. "MB" is megabytes.

---

1. AI & Agent Monitoring

"How YuviRed observes AI agents and learns their normal behavior." Mostly AMS (Agent Monitoring System) behavior analysis. Findings are report only; nothing is blocked.

Label	Default	Allowed	What it does	Notes
Behavioral Detection	On	On, Off	Builds a behavior profile per AI agent and reports when an agent departs from its own pattern or from similar agents.	Warning on Off: evaluating AI-ag anomalies." Exis Key `ams.behavic`
Detection Sensitivity	Balanced (Recommended)	Low: fewer findings; Balanced (Recommended); High: more findings	How strong and persistent a deviation must be before it becomes a finding. Sets two keys together (`modified_z_threshold` 5, 3.5 or 2.5 and `hysteresis_k` 3, 2 or 1).	Shows "Custom: if the two keys w
Learning Period	Standard (Recommended)	Fast: learn quickly; Standard (Recommended); Conservative: learn longer	How many observations are needed before an agent is judged against its own history (4, 8 or 16) and how much history a profile keeps (25, 50 or 100).	While an agent i compared to sim
Finding History	Standard (Recommended): last 500	Compact: last 100; Standard: last 500; Extended: last 2000	How many recent behavioral findings are kept in a rolling list.	Older findings ag
Agent Activity Tracking	On	On, Off	Links endpoint telemetry to the AI agents that produced it.	Warning on Off: longer linked to Detection has nc Also decides whi against the AMS and-limits.md). K `ams.correlation`

Label	Default	Allowed	What it does	Notes
Rogue Agent Activity Detection	Off	On, Off	Scores each governed AI agent against its approved boundary (unapproved destinations, tool expansion, self-configuration, credential reach, restricted-but-still-active, blocklisted destinations) and raises a Critical alert when one steps outside it.	Warning on On: ' alerts within abo `agent_boundary.`
Browser AI Scan Frequency	Standard: every 5 minutes (300 s)	Frequent: every 2 minutes; Standard: every 5 minutes; Reduced: every 15 minutes	How often endpoint agents re-scan browser history for visits to AI services.	The server accep seconds. Applied next configuratio costs slightly mo Key `agent.history_s`
AI Catalog Review Reminder	Quarterly (Recommended): 90 days	Monthly (30); Quarterly (90); Twice a year (180)	AI tool definitions not reviewed for longer than this are marked overdue on the catalog freshness report.	Server accepts 7 `catalog.review_`

Advanced rows in this section:

Label	Default	Allowed	What it does
Analysis frequency	Every 5 minutes (300 s)	Every 2 minutes; every 5 minutes; every 15 minutes; every hour	How often the behavioral engine runs a full pass.
Profile adaptation speed	Standard (0.3)	Slow: long memory (0.1); Standard (0.3); Fast: favors recent behavior (0.6)	How quickly an agent's learned normal follows its recent behavior. Fast can normalize a slow drift.
Similar agents needed for group comparison	3	whole number, 1 to 1000	Group comparison is skipped when fewer agents of the same type exist.
First-time behavior reporting	Report all new behaviors (Recommended) (0)	Report all (0); Ignore minor new behaviors (0.25); Significant new behaviors only (0.5)	Whether a first-time connection or capability is reported on its own.
Recent-activity window	24 hours (Recommended)	6 hours; 24 hours; 3 days; 7 days	Look-back window for counting recent activity.
Detection mode	Report only (Recommended)	Report only; "Active response: not available in this version" (disabled)	This version always reports without acting. Active response does not exist yet; choosing it would change nothing, and the console disables it.
Activity tracking frequency	Every 2 minutes (Recommended)	Every minute; every 2 minutes; every 5 minutes; every 15 minutes	How often new telemetry is linked to AI agents.

2. Prompt Protection

"What the browser extension observes, warns about or blocks before a prompt or attachment leaves for an AI service, using a preset or your own policy." This section is not a list of simple rows. It holds one versioned policy document with its own draft, impact summary, **Save** and apply steps. Saving the policy needs `config.write`; viewing it needs `alerts.read`. The full procedure is in sensitive-data.md.

- **Presets** (card per preset, plus a comparison table): **Monitor Only**, **Light Warnings**, **Warnings Only**, **Balanced**, **Strong**, **Strict**, **Maximum**, and **Custom** (a preset plus your own changes). A new installation starts in **Monitor Only** ("No popups, banners or blocks. Findings are still recorded in Prompt Monitoring."). A preset marked "Recommended default: review and apply" means no administrator has reviewed it yet.
- **Advanced policy** (collapsed): content classes, categories, rule overrides, **Minimum confidence for an intervention** (Low, Medium, High; Medium for every preset except Maximum, which uses Low), **Contact details: act on bulk lists only**, **Payment cards: require card context**, handling of content that cannot be inspected (Allow and record, Warn, Block), floors (administrator minimums that may only raise an action), scopes (per destination, endpoint or user), exceptions (may only lower an action; each has an owner, reason and expiry), **Confidential terms** (up to 200, 2 to 64 characters each), **Internal domains** (up to 200), and organization formats and namespaces.
- **Actions per class:** Observe, Warn or Block.
- **Warning:** "Block" stops a prompt or attachment from leaving the browser. The check runs inside the supported AI sites' pages, so those sites' own code can read your confidential terms and internal domains. Do not enter anything that must stay hidden from them. While the server is locked (not activated after 7 days, or a license that ended) the extension receives Monitor Only and nothing is blocked; blocks applied earlier by a response action stay until rolled back.
- The older setting "Prompt Blocking" (`enforcement.prompt_blocking.default_mode`) and its rules are history once a Prompt Protection revision exists; the server refuses changes to them with a 409 and points you to the policy. Extensions older than 0.5.0 receive a mode derived from the policy.

3. Detection & Response

"What gets inspected, what gets blocked, and how quickly alerts should be handled."

Label	Default	Allowed	What it does	Notes
Prompt Content Capture	Redacted excerpts (Recommended)	Off: no capture at all; Activity only: no content; Redacted excerpts (Recommended); Full content	The most content YuviRed may keep about an AI prompt or response event. Endpoints can capture less than this, never more.	Warning on Full content: complete prompt and outpu text is stored for matched events. Key `privacy.capture_level`.
Inspection Depth	Activity metadata only (Recommended)	Activity metadata only; Full content inspection	Metadata only: records that AI was used, where and by whom, not the prompt text. Full: prompt content becomes eligible for capture, encrypted at rest.	Warning on Full. The server accepts only `metadata` and `full`. Key `privacy.inspection_mode`.
Keep Original Prompt Samples	Off	On, Off	Stores encrypted copies of matched prompts for investigation. Needs Inspection Depth = Full.	Warning on On: copies are stored until the retention window ends. Viewing a decrypted sample needs `dlp.view_raw_content` (Adr roles) and is audited. Key `privacy.raw_capture_enabl`

Label	Default	Allowed	What it does	Notes
Critical alerts	4 hours	1 hour; 4 hours; 8 hours; 24 hours; 3 days; 1 week; No target (0)	Target response time for new alerts of that severity.	A tracking goal only, not an automated action. Key `sla.critical_hours`.
High alerts	24 hours	same list	same	Key `sla.high_hours`.
Medium alerts	1 week (168 hours)	same list	same	Key `sla.medium_hours`.
Low alerts	No target (0)	same list	same	Key `sla.low_hours`.

Advanced rows:

Label	Default	Allowed	What it does
Prompt samples: endpoint default	Capture samples by default	Capture samples by default (on); Skip samples by default (off)	Whether endpoints keep original prompt samples unless a tag rule says otherwise.
Always keep samples from endpoints tagged	empty	text (tags)	Endpoints with any of these tags always keep samples.
Never keep samples from endpoints tagged	empty	text (tags)	Endpoints with any of these tags never keep samples. If both lists match, "never" wins.
Allow sample capture without an encrypted connection (lab only)	Off	On, Off	Lets prompt-sample capture run over plain HTTP. Warning on On: prompt content could cross the network in cleartext. Leave off anywhere real.

4. Artifact Analysis

"Whether analysts may have uploaded artifacts examined by external providers, which providers are used, and how much of an artifact is permitted to leave this appliance." Off by default. Reading or changing any row needs `artifact.manage` (Admin) as well as `config.read` or `config.write`; without it the server answers HTTP 403 "Artifact Analysis configuration requires the artifact.manage permission." YuviRed ships no provider credential: the account, entitlements and terms are yours. The section opens with a numbered note, "Turn this on in order", and the master switch is the last row on screen, under the heading **Turn Artifact Analysis on**.

Label	Default	Allowed	What it does
Use MetaDefender Cloud	Off	On, Off	File and hash reputation by OPSWAT. Leave off until the key and the commercial-use attestation are saved. A **Test Connection** button under it sends the key and nothing else, and only works once Artifact Analysis is enabled.
MetaDefender Cloud API key	not set	secret	Your own key. Stored encrypted, never shown again.
Use Hive	Off	On, Off	AI-generated image detection. It has no hash lookup, so it is reached only when the policy permits an image to be submitted. Hive offers no connection test.
Hive image API key	not set	secret	Your own key; stored encrypted, never shown again.
Enable Artifact Analysis (last row)	Off	On, Off	Master switch. While off, no artifact route accepts work and no outbound provider call is made. Saving On shows a confirmation first that states exactly what the configuration then permits to leave the server.

Advanced Provider Settings (collapsed):

Label	Default	Allowed	What it does
Submission Policy	Hash only (Recommended)	Hash only (Recommended); Private upload allowed; Public upload allowed; Upload disabled: nothing leaves	The ceiling on what may be disclosed. Upload disabled: nothing leaves, not even hashes. Hash only: only supported hashes may be sent. Private upload allowed: file bytes only to a provider attested as private. Public upload allowed: file bytes to configured providers after an analyst authorizes. Warning when widening.
Ask the analyst to authorize each file submission	Off	On, Off	On: every analysis that would send a file waits for the analyst's confirmation. It never widens the policy ceiling.
Commercial-use attestation (each provider)	Off	On, Off	Your record that you may use the provider commercially. YuviRed records it and verifies nothing.
Private-processing attestation (each provider)	Off	On, Off	Your record that your agreement permits private processing.
Cached hash result freshness (MetaDefender)	24 hours	1 to 8760	How long a stored hash result is reused.
Cached file result freshness (MetaDefender)	168 hours	1 to 8760	How long a stored file result is reused.
Cached image result	720 hours	1 to 8760	How long a stored image result is reused.

Label	Default	Allowed	What it does
freshness (Hive)			
Largest accepted upload	20 MB	1 to 50	Larger uploads are refused. The 50 MB ceiling is fixed in the web server and cannot be raised here.
Unauthorized artifact lifetime	30 minutes	5 to 120	How long an uploaded, unauthorized artifact stays in encrypted staging.
Authorized artifact lifetime	2 hours	1 to 24	How long an authorized artifact may stay in staging while providers work.
Staging area capacity	200 MB	50 to 4096	New uploads are refused once staging holds more.
Analysis history retention	365 days	0 to 3650	How long analysis history is kept; 0 keeps it forever. Shorter values delete history.

5. Notifications

"Which events reach your team, and where they are delivered." A notification fires only for a NEW finding at or above the chosen severity, subject to the quiet window and the hourly cap. Use the panel **Notifications: Test Send** (save first) to check each channel.

Label	Default	Allowed	What it does
Notify me about	High and above (Recommended)	Critical only; High and above (Recommended); Medium and above; Everything	Minimum severity of a new finding that triggers a notification.
Email alerts	On	On, Off	Send email for findings. Needs a mail server and recipients.
Send alerts to	empty	text (addresses)	Recipients.
Mail server	empty	text	SMTP host name.
Mail server port	587	1 to 65535	SMTP port.
Mail username	empty	text	
Mail password	not set	secret	
Encrypt mail connection (TLS)	On	On, Off	Turning it off sends mail unencrypted.
Send from address	`noreply@yuvired.local`	text	Change it to an address your mail server accepts.
Slack / Microsoft Teams alerts	On	On, Off	Chat channel switch.
Slack or Teams webhook URL	not set	secret	An incoming-webhook URL from Slack or Microsoft Teams.

Label	Default	Allowed	What it does
Webhook feed	On	On, Off	Sends each notification as a JSON event to the URL below, for a SIEM (security information and event management system), SOAR or automation pipeline.
Webhook URL	not set	secret	
Daily summary email	Off	On, Off	Once-a-day summary of medium and low findings. High and critical findings always notify at once.
Send after	07:00 UTC (Recommended)	00:00 to 23:00 UTC	Hour for the daily summary, in UTC.
Don't repeat the same alert for	6 hours (Recommended)	1 hour; 6 hours; 12 hours; 24 hours	A repeat of the same finding on the same endpoint inside this window is not sent again.
Max notifications per hour	20 (Recommended)	10; 20; 50; 100	Hourly cap.
System-health alert repeat interval (Advanced)	24 hours (Recommended)	6; 12; 24; 48 hours	How long a repeat page about the same degraded background job is held back. Recovery notices always send.

Channel switches are independent of whether an address or URL is set, so you can pause a channel without clearing its credentials. The default of On for chat and webhook does nothing until a URL is saved.

6. Data & Retention

"How long telemetry, findings, and audit records are kept." The daily cleanup deletes data permanently.

Label	Default	Allowed	What it does
Automatic Data Cleanup	On	On, Off	Runs the daily cleanup that enforces the windows below. Manual runs stay available. Warning on Off: old data accumulates without limit.
Raw telemetry	30 days	whole days, 1 or more	Raw telemetry older than this is removed, but only for days already rolled up into daily summaries.
Daily summaries	13 months	whole months, 1 or more	Daily summary data older than this is removed.
Closed alerts	365 days	days, 0 = keep forever	Removes only closed alerts (resolved, false positive, suppressed) that are not part of an incident.
Audit log	0 (keep forever)	days, 0 = keep forever	The record of who changed what. Setting a window permanently deletes older audit records.
AI agent signals	120 days	days, 0 = keep forever	Raw AI-agent detection evidence rows. Only rows already processed are removed.
AI prompt/output events	90 days	days, 0 = keep forever	Prompt and output events from the browser extension.
Decided adoption requests	30 days	days, 0 = keep forever	Pending requests are never removed.
Original prompt samples	24 hours (Recommended)	4 hours; 24 hours; 3 days	How long an encrypted prompt sample is kept before its content is erased.

Warning: lowering any window asks "Shorten ... retention? Records older than N ... will be permanently deleted at the next cleanup run". Moving from "keep forever" (0) to a number deletes everything older at the next run. There is no undo and Lite ships no backup command.

Other retention rows with no plain-language control (shown as raw rows under Advanced, Other settings): `retention.report_file_days` (30, generated report files are deleted after this many days; history entries stay; 1 to 3650), `retention.report_storage_mb` (2048; oldest stored files go first when exceeded; 100 to 51200), `retention.ams_message_days` (30), `retention.ams_event_days` (180), `retention.ams_instance_days` (90) (0 = keep forever), `ams.activity_retention_days` (365), `ams.risk_snapshot_retention_days` (180), and `agent_boundary.retention_days` (0 = keep forever). These raw rows delete data at their next run like the rows above, but the console shows **no confirmation box** when you shorten them.

Panel **Telemetry Retention** on the same page: **Run now (dry run)** shows what would be deleted; **Run now** deletes it. Alerts, evidence and the audit log are never purged by the telemetry job; they follow their own windows above.

---

7. Security & Access

"Connection security and sign-in requirements."

Label	Default	Allowed	What it does	Notes
Agent Connection Security	Allow unencrypted connections	Allow unencrypted connections (`off`); Warn: log unencrypted connections (`warn`); Require encrypted connections (`enforce`)	How strictly agents and browser extensions must use TLS (Transport Layer Security) when reporting in.	Warning on Require: agents and extensions connecting without TLS are refused (HTTP 426) until reconfigured. Key `security.tls_enforcement`. Read live on every request.
Require multi-factor sign-in for	Admin	check boxes **admin** and **analyst**	Users with these roles must enroll an authenticator app at sign-in. View Only accounts can never be required.	Key `mfa.required_roles`. A user in a required role without MFA is sent to forced enrolment at next sign-in. There is no setting for the session length (8 hours) or the sign-in lockout rules; they are fixed (see roles-and-permissions.md).
Prevent password reuse	Last 5 passwords (Recommended)	Off; last 3; last 5; last 10; last 24	Refuses a new password equal to any of the last N.	The server accepts 0 to 24.
Encryption Key Rotation Reminder	Annually (Recommended)	Quarterly (90); Twice a year (180); Annually (365); Every 2 years (730)	When the last recorded rotation is older than this, System Status marks it overdue and an	A reminder only. Rotation is never automatic. The server accepts 30 to 730 days.

Label	Default	Allowed	What it does	Notes
			operational alert is sent.	
Require per-device gateway credentials (Advanced)	Off	On, Off	AI-gateway traffic is accepted only from devices holding their own short-lived credential; the shared deployment token is refused.	Warning on On: devices without a credential lose AI-gateway ingest until they check in and receive one. Turning it off again is the immediate rollback.
Gateway credential lifetime (Advanced)	8 hours (Recommended)	1 hour; 8 hours; 24 hours	Lifetime of a device credential.	
Minimum agent version (Advanced)	empty (no minimum)	a version such as `1.8.2` (the box's example text still reads "e.g. 1.7.4")	Agents older than this are refused (HTTP 426) when they report in, until updated.	Warning when set.

8. Integrations

"Single sign-on and external systems." Single sign-on uses OpenID Connect (OIDC). It stays inactive until the provider URL, client details and this server's public URL are all set.

Label	Default	Allowed	What it does
Single Sign-On	Off	On, Off	Adds a **Sign in with SSO** button.
Login button label	empty	text	Text on the button.
Identity provider URL	empty	URL	Your identity provider's address.
Client ID	empty	text	
Client secret	not set	secret	
This appliance's public URL	empty	URL	The address users reach this console at. Required: SSO stays disabled until set.
Auto-create accounts on first SSO sign-in (Advanced)	Off	On, Off	New users get a read-only (View Only) account. Has no effect until approved email domains are set.
Approved email domains for auto-created accounts (Advanced)	empty	text	
Extra trusted identity-provider hosts (Advanced)	empty	text	Only for providers that serve sign-in endpoints from other host names.
Requested sign-in strength (OIDC acr_values) (Advanced)	empty	text	Optional authentication-strength request.
Require provider MFA for roles that require MFA (Advanced)	On	On, Off	Warning on Off: SSO users bypass this platform's MFA requirements entirely.

The SIEM or webhook feed is configured under Notifications. The firewall blocklist feed token lives on the Network page. Setting up SSO against a real identity provider has not been verified in this release's documentation work; treat it as unverified until you test it with your provider.

---

9. System

"Platform identity and agent update behavior."

Present but inactive on YuviRed Lite. Every row below except the first two belongs to the Docker appliance edition. On a Lite server they can be switched and saved, and the audit log records the change, but they change nothing today:

- **Automatic Agent Updates** acts only on agents installed with an install command. Every agent on a Lite server is installed from the Add Computer package, so the switch has no agent to act on. To update an agent on Lite, run a new Add Computer download on that computer.
- **Remote Browser Extension Wiring** and the three advanced extension rows act only when this server serves the signed extension or a store listing exists. The Lite service does not serve the signed extension (the Windows service removes the signing key), so the setting answers "disabled" to every agent, and the Chrome Web Store ID is empty until a listing is approved. In this release nothing installs the extension into Chrome or Edge automatically; the working route is a manual "Load unpacked" (see extensions.md). Automatic installation from the store is planned for a later update.
- The two **Staged rollout** rows pace Automatic Agent Updates, so on Lite they have nothing to pace.

Label	Default	Allowed	What it does	Notes
Organization Name	`YuviRed`	text	Used in notification subjects and reports.	
Privacy contact	empty	an email address or an `https://` URL, up to 200 characters	Shown to the people your endpoints monitor, on the notice page.	The server refuses anything else.
Automatic Agent Updates	Off	On, Off	Agents installed with an install command update themselves, paced by the staged rollout settings.	Inactive on Lite (above). Warning on On: "Agents installed from a package, including every computer added with Add Computer on YuviRed Lite, are not changed."
Remote Browser Extension Wiring	Off	On, Off	Windows agents (1.7.12 or later) keep Chrome and Edge wired for the self-hosted extension: force-install policy, managed configuration, native messaging.	Inactive on Lite (above). On the appliance edition it writes machine policy as SYSTEM on every computer and is inert unless the server serves the

Label	Default	Allowed	What it does	Notes
				signed extension.
Extension distribution mode (Advanced)	Self-hosted: appliance serves the signed .crx (managed / air-gapped)	Self-hosted; Chrome Web Store: unlisted/private (required for unmanaged endpoints)	How Chrome and Edge force-install the extension.	Inactive on Lite (above). Store mode requires agents to check in over HTTPS.
Chrome Web Store extension ID (Advanced)	empty	32 letters a to p	Used in store mode; while empty, store-mode wiring stays inert.	Empty in this release.
Allow extension wiring over http agents (Advanced)	Off	On, Off	Permits wiring on computers whose agent checks in over plain HTTP.	Inactive on Lite (above). A network attacker can tamper with the instruction on plain HTTP regardless of this switch. Trusted networks only.
Staged rollout: target agent version (Advanced)	empty	a version	Inactive unless it matches the version bundled with the server.	Inactive on Lite (above).
Staged rollout: current stage (Advanced)	Canary: smallest group first (Recommended)	Canary; Early: wider group; Broad: whole fleet	Which cohort is unlocked for the target version.	Advance canary, early, broad by hand.

Settings that interrupt monitoring, block activity, delete data, change access or affect licensing

"Confirmation" says whether the console asks you to confirm before saving.

Setting	Effect	Class	Confirmation
Agent Connection Security = Require encrypted connections	Refuses agents and extensions that connect without TLS (HTTP 426).	Interrupts monitoring	Yes
Minimum agent version	Refuses older agents (HTTP 426) until updated.	Interrupts monitoring	Yes
Require per-device gateway credentials (Advanced)	Devices without a credential lose AI-gateway ingest.	Interrupts monitoring	Yes
Behavioral Detection Off; Agent Activity Tracking Off	Stops behavior analysis and telemetry linking; Agent Activity Tracking Off also stops new AI agents being seated against the AMS limit.	Interrupts monitoring	Yes
Detection mode (Advanced, AI & Agent Monitoring)	Only "Report only" works; the other choice is disabled.	Nothing	No
Prompt Protection presets Balanced, Strong, Strict, Maximum; any Block action; "Block" for content that cannot be inspected	Blocks prompts or attachments in the browser.	Blocks activity	The policy editor shows an impact summary before apply
Rogue Agent Activity Detection On	Raises Critical alerts (report only).	Alerts	Yes
Automatic Data Cleanup Off	Old data accumulates	Storage	Yes

Setting	Effect	Class	Confirmation
	without limit.		
Raw telemetry; Daily summaries; Closed alerts; Audit log; AI agent signals; AI prompt/output events; Decided adoption requests; Original prompt samples	Shortening any window permanently deletes older records at the next cleanup run.	Deletes data	Yes, when shortened
Analysis history retention (Artifact Analysis, Advanced)	Shorter values delete analysis history.	Deletes data	No
Raw rows `retention.report_file_days`, `retention.report_storage_mb`, `retention.ams_message_days`, `retention.ams_event_days`, `retention.ams_instance_days`, `ams.activity_retention_days`, `ams.risk_snapshot_retention_days`, `agent_boundary.retention_days`	Delete older records at their next run.	Deletes data	No
Debug / Reset page (not a setting)	Clears selected data, cannot be undone.	Deletes data	Yes (on the page)
Require multi-factor sign-in for; Prevent password reuse; Single Sign-On and its rows	Change who can sign in and how.	Changes access	No
Require provider MFA for roles that require MFA = Off	SSO users bypass the MFA requirement.	Changes access	Yes
Auto-create accounts on first SSO sign-in	Creates View Only accounts for unknown people from the approved domains.	Changes access	No

Setting	Effect	Class	Confirmation
Keep Original Prompt Samples On; Prompt Content Capture = Full content; Inspection Depth = Full	Store prompt content (encrypted).	Changes privacy	Yes
Allow sample capture without an encrypted connection (lab only)	Prompt content could cross the network in cleartext.	Changes privacy	Yes
Enable Artifact Analysis; Submission Policy widened	Can send hashes or files outside your network.	Changes privacy	Yes
Automatic Agent Updates; Remote Browser Extension Wiring and the extension rows; Staged rollout	Fleet-wide behavior on the appliance edition. Inactive on Lite (section 9).	Fleet-wide	Automatic Agent Updates: Yes; the others: No
License state	Not a setting: the server refuses any write to license state (HTTP 400, recorded in the audit log). Activate through the license screen (see licensing.md).	Affects licensing	Not applicable
Session length; sign-in lockout	Not settings: fixed at 8 hours, and 8 failures per 5 minutes per address plus 30 per hour per account.	Access	Not applicable

Server-only and hidden settings

These cannot be changed from the console.

- **Written by the system:** the **Diagnostics** box shows read-only values, in this order: **Behavioral analysis: last completed**, **last started**, **last pass**; **Learned behavior profiles**; **Recent behavioral findings (rolling buffer)**; **Activity tracking: last completed**, **last started**, **rows scanned last cycle**, **progress marker**; **Observed-activity progress marker**; **Rogue Agent Activity: last completed**, **last started**, **last cycle stats**; **Daily summary: last sent (UTC date)**; **Encryption key: last rotation recorded**; **Prompt-blocking policy version**; **Prompt Blocking (legacy): stored value** and **Legacy blocking rules: stored value** (both superseded by Prompt Protection); **Prompt Protection: active policy revision**; **Active license ID (managed via the License panel)**; **EULA accepted: version, time and by** (recorded at first-run setup; the server refuses edits); **Gateway credential revocation counter**.
- **License state** (`license.active_license_id`): the server refuses writes from every role with HTTP 400 ("<key> is license state, not a configurable value. Activate or revoke a license through the License panel...") and records the attempt.
- **EULA acceptance record** (`legal.*`): refused with HTTP 400; written once at setup.
- **Encryption key rotation stamp** (`security.encryption_key_rotated_at`): written only by the rotation script, refused otherwise.
- **Hidden because no feature reads them:** `ai.*` (provider addresses, default provider, MCP servers), `proxy.*`, `branding.logo_url`, `branding.accent_color`, `branding.theme_default`, `update.channel`, `update.auto_check`, `update.last_checked_at`, `logging.level`, `logging.json_format`, `retention.telemetry_days`, `ams.behavioral.peer_prior_strength`. In particular the log level is not changeable from the console, and the console has no proxy setting.
- **Other keys with no plain-language row** (shown as raw key rows under Advanced, Other settings): the internal progress and error markers `ams.correlation_*` (last error, retry, quarantine, resume, stats), the Rogue Agent tuning values `agent_boundary.*` (interval, window, thresholds, family weights, retention, agents per tick), and `gateway.allow_token_ingest_for_credentialed_endpoints`. Do not change them unless YuviRed support asks you to.
- **Environment and installer settings** (network mode, ports, certificate, the service account and similar) are set on the server itself (for example the `network` command that the Add Computer panel shows) or by the installer, not in the console. See installation.md.

- A test-only Artifact Analysis provider has five settings that exist only in an isolated test stack; they are absent on a real install.

Related

- console-pages.md
- roles-and-permissions.md
- editions-and-limits.md
- statuses-and-errors.md

Roles and permissions reference

The built-in roles, what each can and cannot do, which pages and API routes need only a signed-in account, and the sign-in, MFA and session rules the server enforces.

Covers YuviRed 1.0 for Windows. Everything here is read from the server's own tables and code; the source of each fact is in the comments.

Two layers: console role and system role

The **Team / Users** page lets an administrator pick one of three roles for an account: **Admin**, **Analyst** or **View Only**. Behind those three, the server keeps five **system roles**, each a list of permissions. The console pick maps to a system role like this:

Console role (Role column)	System role given	Description from the server
Admin	`administrator`	"Full operational access. Cannot manage licensing."
Analyst	`security_analyst`	"Day-to-day triage: alerts, policies, endpoints, reports."
View Only	`read_only`	"View-only access across the platform."
(not in the console)	`super_admin`	"Full access to every capability, including license and role management."
(not in the console)	`developer`	"API and integration access: AI providers, prompt templates, MCP servers."

Three consequences to know:

1. **The first account is the owner (`super_admin`).** First-run setup creates it, and it is the only account that holds the `license.manage` permission. An account you create with the **Admin** role is `administrator` and **cannot activate a license, import a license file or redeem an activation code**: the buttons show, but the request fails with `Missing required permission: license.manage`. Sign in with the first account to manage the license. (The license screen says: "Only a YuviRed administrator can activate this server. Sign in with the administrator account created when YuviRed was installed.")

2. **Changing the role of an existing user replaces all of that user's role assignments** with the new one.
3. `super_admin` **and** `developer` **and custom roles cannot be assigned from the console.** They are set only through the API route `POST /api/v1/rbac/assign`, which needs `roles.manage`. The five system roles cannot be deleted. A user with no role assignment at all is treated as the system role that matches the user's console role.

The console hides some menu items by console role (Admin only: Blocklist, Agent Deployment, Adoption Requests, Configuration, Legal, Help, Team / Users, Debug / Reset, System Logs; Analyst and Admin: Enforcement; Artifact Analysis shows only to accounts holding `artifact.read`). The hiding is a convenience. The server checks the permission on every request, and that check is what counts.

Permissions by role

38 permissions exist (counted from the server's permission catalogue; the "What it allows" column says where the console checks each one, which is more specific than the server's one-line description). `Y` means the role holds it. SA = `super_admin`, AD = `administrator` (console **Admin**), AN = `security_analyst` (console **Analyst**), RO = `read_only` (console **View Only**), DV = `developer`.

Permission	What it allows	SA	AD	AN	RO	DV
`users.read`	View user accounts	Y	Y	-	-	-
`users.manage`	Create, edit, deactivate, delete users; reset passwords; turn off another user's MFA; end sessions	Y	Y	-	-	-
`roles.manage`	Create and assign roles	Y	Y	-	-	-
`policies.read`	View policies	Y	Y	Y	Y	-
`policies.write`	Create, edit, enable or disable policies; classify AI tools	Y	Y	Y	-	-
`alerts.read`	View alerts, Prompt Monitoring events and Prompt Protection policy	Y	Y	Y	Y	-
`alerts.write`	Change alert status, add notes, create and change incidents, export prompt events	Y	Y	Y	-	-
`endpoints.read`	Read-only endpoint extras (credential rotation status, rollout status, server identity)	Y	Y	Y	Y	-
`endpoints.manage`	Remove, decommission and reactivate endpoints; deployment tokens; Add Computer; adoption requests; credential rotation	Y	Y	-	-	-
`endpoints.operate`	Refresh a computer now; collect agent logs	Y	Y	Y	-	-
`blocklist.manage`	Manage the blocklist; export the denylist	Y	Y	Y	-	-
`config.read`	View platform configuration	Y	Y	Y	Y	Y
`config.write`	Change configuration, including secrets; save the Prompt Protection policy; classifications; blocklist feed token	Y	Y	-	-	-
`ai_providers.read`	View AI provider records	Y	Y	-	Y	Y

Permission	What it allows	SA	AD	AN	RO	DV
`ai_providers.write`	Add or edit AI provider records and keys	Y	Y	-	-	Y
`prompt_templates.manage`	Manage prompt templates	Y	Y	-	-	Y
`audit.read`	View and export the audit log	Y	Y	-	Y	-
`license.manage`	View and change license information: activate, import, redeem a code, confirm a moved server	Y	-	-	-	-
`reports.read`	Generate and view your own reports	Y	Y	Y	Y	Y
`reports.manage`	See, download and delete reports other people made	Y	Y	-	-	-
`debug.access`	Debug / Reset, System Logs, system status, data reset	Y	Y	-	-	-
`tools.manage`	Delete and merge AI tools; manage the AI Tool Catalog	Y	Y	-	-	-
`dlp.view_raw_content`	View decrypted original prompt samples (every view is audited)	Y	Y	-	-	-
`response.read`	Request and view enforcement (response) actions	Y	Y	Y	-	-
`response.enforce`	Approve, reject, run, recall, roll back, disable, archive, delete actions	Y	Y	-	-	-
`ams.read`	View the AMS agent registry, governance records and risk	Y	Y	Y	Y	-
`ams.manage`	Edit AMS governance (owner, purpose, approval, pause); refer agents to AIDR	Y	Y	-	-	-
`ams.message`	Send messages to AI agent sessions	Y	Y	-	-	-
`ams.approve`	Approve or deny one pending AI agent tool call	Y	Y	-	-	-

Permission	What it allows	SA	AD	AN	RO	DV
`ams.policy`	Change the AMS control policy and connectors	Y	Y	-	-	-
`ams.stop`	Stop, hold, release, pause, resume and cancel AI agent work	Y	Y	-	-	-
`ams.bulk`	Act on several AI agents at once (each act also needs its own permission)	Y	Y	-	-	-
`deployment.read`	View deployment campaigns	Y	Y	Y	Y	-
`deployment.manage`	Create or change campaigns; mint bootstrap credentials	Y	Y	-	-	-
`metrics.scrape_mint`	Mint, rotate and revoke credentials for the unattended metrics feed	Y	Y	-	-	-
`artifact.read`	View Artifact Analysis results and provider status	Y	Y	Y	Y	-
`artifact.analyze`	Start, authorize and cancel an artifact analysis	Y	Y	Y	-	-
`artifact.manage`	Delete artifact history; test providers; read or change Artifact Analysis settings	Y	Y	-	-	-

In plain words

- **Admin (** `administrator` **)** holds every permission except `license.manage`.
- **Analyst (** `security_analyst` **)** triages: it reads and changes alerts, policies, the blocklist and incidents, can refresh a computer and collect its logs, can request (but never approve or run) an enforcement action, can run artifact analyses and read reports. It cannot change configuration, manage users, remove endpoints, view the audit log, see original prompt samples or act on AI agents beyond reading them.
- **View Only (** `read_only` **)** reads: policies, alerts, endpoints, configuration values, AI provider records, reports, the audit log, AMS, deployment campaigns and artifact results. It changes nothing.
- **Developer** holds no security-data permissions (no alerts, endpoints, policies or AMS). It reads configuration and reports and manages AI provider records and prompt templates.

What needs only a signed-in account

Some pages and routes check that you are signed in but do not check any permission. **Any signed-in user, of any role (including View Only and Developer), can read these.**

Plain statement: **any signed-in user can read endpoint identity and raw telemetry.** That means the list of endpoints (host name, user, IP address, operating system, agent version, risk), one endpoint's detail, and its raw telemetry rows (processes, network connections, browser AI visits, extensions, software, local AI, DNS names, AI agent signals), up to 1000 rows per request. Treat every account, including View Only, as having that access. Create accounts accordingly.

Routes that need only a signed-in account:

Area	**Routes (** `/api/v1/...` **)**
Endpoints	`GET /endpoints`, `GET /endpoints/{id}`, `GET /endpoints/{id}/telemetry`, `GET /endpoints/{id}/extension-changes`, `GET /extension-changes`
AI tools and catalog	`GET /tools`, `GET /tools/{id}/usage`, `GET /tool-catalog`, `GET /tool-catalog/freshness`, `GET /tool-catalog/{id}`
Classification, policy, blocklist	`GET /classifications`, `GET /classifications/{id}`, `GET /policies`, `GET /blocklist`
Dashboards	`GET /dashboard/stats`, `GET /topology`, `GET /reports/summary`, `GET /reports/executive`, `GET /behavior/summary`, `GET /network/summary`, `GET /network/provider-ip-ranges`
Incidents	`GET /incidents`, `GET /incidents/{id}`
License	`GET /license`, `GET /license/capacity`, `GET /license/activation`
Notifications	`GET /notifications/status`
Roles	`GET /rbac/roles`, `GET /rbac/permissions`, `GET /rbac/users/{id}/permissions`, `GET /rbac/users/{id}/roles`
Your own account	`POST /auth/logout`, `GET /auth/me`, `POST /auth/change-password`, `GET /auth/mfa/status`, `POST /auth/mfa/disable`, `POST /auth/mfa/recovery-codes/regenerate`

Everything else in the console routes needs a named permission (above). The agent and extension routes (register, check-in, configuration, gateway, extension download) are not for

people: they use a deployment token, an endpoint credential or a gateway credential instead of an account. A few routes need nothing (the health and version routes, sign-in and setup).

Multifactor authentication (MFA)

- **What it is:** a second sign-in step with a 6-digit code from an authenticator app (time-based one-time password), or a one-time recovery code.
- **Who must use it:** the roles listed in the setting **Require multi-factor sign-in for** (Configuration, Security & Access). The default is **Admin**. The console offers **Admin** and **Analyst**; **View Only** can never be required, whatever the setting says. A user in a required role who has not enrolled is sent straight to enrolment at sign-in, including the first account at the end of first-run setup. Console role is what counts: `admin` covers both `super_admin` and `administrator`.
- **Everyone else:** any user may enrol or disable MFA for their own account from the account menu (**Authentication / MFA**, **Recovery Codes**). Turning it off needs a current code. Regenerating recovery codes needs a current code and replaces the old set.
- **Recovery codes:** 10 codes of 10 characters, shown once, each works once. A new set replaces the old.
- **Time limit:** the step between password and code (and forced enrolment) lasts **2 minutes**. If it ends the console says: "This sign-in step has ended. It lasts 2 minutes, and a newer sign-in or a restart of YuviRed also ends it. Sign in again."
- **Wrong codes:** a wrong code counts as a failed sign-in on the same budget as a wrong password. In addition, more than 20 wrong or repeated code attempts for one account in an hour trigger a back-off with the message `Too many failed codes for this account. Retry in Ns.` (HTTP 429). It slows guessing; it does not lock the account for good.
- **Lost phone and no recovery codes:** an administrator with `users.manage` turns that user's MFA off (the shield icon on **Team / Users**); the user enrols again at next sign-in.
- **Single sign-on:** SSO users skip the server's own code step. By default the identity provider must prove it performed MFA for roles that require it (setting **Require provider MFA for roles that require MFA**, default On). Turning that off lets SSO sign-ins bypass the requirement.

Passwords

- At least 12 and at most 72 bytes; at least 3 of 4 character classes (lowercase, uppercase, digit, symbol); at least 6 different characters; must not contain the user's own name or email name, or a commonly guessed word (including leet-speak forms

such as `P@ssw0rd`); must not equal the current password or any of the last N (setting **Prevent password reuse**, default 5, range 0 to 24). Existing passwords are never re-checked.
- The same rules apply at first-run setup, **Add User**, **Reset Password** and **Change Password**.
- No scheduled expiry. A password change is forced only when an administrator sets it (reset with a required change, or **require password change**). Until the user changes it, every call except changing the password is refused with HTTP 403 `password_change_required`.
- A password an administrator issues as temporary stops working **72 hours** after it was set. The sign-in is then refused exactly like a wrong password ("Invalid credentials"), and counts as a failed attempt.

Sign-in lockouts

Rule	Value	What the user sees
Per source address and account	8 failed attempts inside a rolling 5 minutes blocks that pair until the window drains.	HTTP 429 `Too many failed login attempts. Please wait and try again.`
Per account from any address	30 failed attempts inside a rolling hour. A time-based cool-down, never a permanent lock. A successful sign-in does not clear it. A sign-in from the server computer itself (loopback) is exempt, so someone at the server can always sign in.	Same 429 message.
Wrong password and unknown user	Same answer for both.	HTTP 401 `Invalid credentials`
Audit	Failed sign-ins are audited, capped at 120 rows a minute per source address.	

On YuviRed Lite there is no Redis, so these counters live in the service's memory. A restart of the YuviRed service clears them. Because the per-account ceiling counts failures from every address, a person on the network who keeps failing can hold an account out of network sign-in for up to an hour; sign in at the server itself to get around it.

Sessions

- **Length:** a sign-in issues a token valid for **480 minutes (8 hours)**. After that the console shows the dialog **Session Expired** ("Your session has ended and background updates have been paused. Log in again to resume monitoring.") with **Log In Again**.
- **Ended early** by: **Sign Out**, a password change (your other sessions end), an administrator's reset, deactivating or deleting the user, and an administrator's revoke-sessions.
- **Lite limit:** Lite has no Redis, so the list of ended sessions is kept in memory. **A restart of the YuviRed service forgets sign-outs, session endings after a password change and revoke-sessions**: such a session works again until its own 8 hours run out. Deactivating or deleting a user still ends their sessions after a restart, and a required password change survives a restart.
- **Active Sessions** in the account menu is disabled: this build does not list individual sessions.
- A session also stops working while the server is locked for licensing (HTTP 402; see statuses-and-errors.md). Signing in, the account routes (your own password and MFA), the license screen, first-run setup and rolling back a response action (so a browser block applied earlier can still be lifted) keep working.
- Neither the session length nor the lockout numbers can be changed from the console.

Related

- console-pages.md
- settings.md
- statuses-and-errors.md
- editions-and-limits.md

Statuses and errors reference

Every status badge, license plan state, license lock message, endpoint and extension status, and the error codes and messages a person can meet, each with what it means and what to do.

Covers YuviRed 1.0 for Windows (agent 1.8.2, extension 0.5.5). Messages are quoted as the product shows them, in the product's US spelling "license".

---

1. License plan states

The license panel (Configuration, License) and the license screen show the plan state of this server. The server computes it from the active license, its dates and this server's installation ID.

Plan state	Badge in the console	Meaning	What to do
`grace`	**Not activated** (amber); the panel line "You have N days left to activate this server. After that, the console is unavailable until the server is activated." and a banner: "This YuviRed server is not activated. You have N days left to activate it. After that, the console is unavailable until it is activated." (**Activate now**). The count reads "1 day" and	A new server inside its 7-day activation window, counted from its first start. The console works with Lite limits (10 computers, 5 AI agents).	Activate the server within the window (activation code or license file). The day count is rounded down, so the banner may say 6 days right after install.

Plan state	Badge in the console	Meaning	What to do
	then "less than 1 day".		
`not_activated`	**Not activated** (red)	The 7 days ended with no activation. The console answers HTTP 402 `license_activation_required` (see section 2).	Activate the server. Enrolled computers keep reporting; YuviRed only monitors until then.
`active`	**Active** (green)	A valid license is in force. For Lite: "YuviRed Lite is free and never expires." For Pro or Enterprise: "This license ends on <date>. When it ends, the console is unavailable until a new license is activated. There is no grace period."	Nothing.
`expiring`	**Ends soon** (amber)	A Pro or Enterprise license ends within 30 days. A banner appears at 30, 14 and 7 days: "Your YuviRed Pro license ends in N days (date)... To renew, email support@yuvired.com." (the **License** link opens the license screen).	Renew before the end date at https://yuvired.com/request. The console keeps working until the end date.
`expired`	**Expired** (red)	A Pro or Enterprise license has ended. Console is locked (`license_expired`). There is no grace after expiry. Lite has no expiry, so Lite never reaches this state.	Activate a new license.
`revoked`	**Revoked** (red)	The license was revoked. For Lite only YuviRed can revoke, at the daily check-in. The lock text says who revoked it.	Email support@yuvired.com if you think it is a mistake; otherwise activate a new license.

Plan state	Badge in the console	Meaning	What to do
`invalid`	**Not valid** (red)	The license cannot be verified on this server (for example it is bound to another server, or its signature fails). Paid plans lock the console. Lite does not lock but refuses new computers: "This license is not valid on this server, so new computers are refused. Activate a new license, or email support@yuvired.com with this server's installation ID."	Activate a license bound to this server's installation ID.
`server_changed`	**Server changed** (red); the panel heading reads "Needs a new license" and the license screen's title "This server needs a new license"	This server's hardware identity changed (a copy, a restore on new hardware, or a virtual hardware change), so YuviRed treats it as a copy of another server.	Email support@yuvired.com with the installation ID shown on the license screen. The old license does not work again by itself, even if you move back.
`legacy`	an amber badge showing the word `legacy` (the console has no	A license that is not a Lite, Pro or Enterprise plan (an older tier). It keeps its earlier behavior and does not lock the console.	Not expected on a new install.

Plan state	Badge in the console	Meaning	What to do
	plain-language label for it)		
`unlicensed`	an amber badge showing the word `unlicensed`	Reported only by a server that is not in license mode. Not expected on a YuviRed install.	None.

Next to the badge the panel shows a heading: "Not activated" (grace and not activated), "YuviRed Lite", "YuviRed Pro" or "YuviRed Enterprise" (the plan in force), "YuviRed Lite (team build)" (a test entitlement, not a signed license), or "Needs a new license". Other labels you may see: the **Computers used** and **AI agents used** counters show "N / limit"; a counter turns red when usage is above the limit. A line "Every computer seat is in use. A new computer is refused until one is free..." appears when no seat is free. If the panel cannot load every detail it says "Could not load all license details. Choose Reload."

---

2. License lock messages (HTTP 402)

While the server is locked every console request answers **HTTP 402** with a code and a message, and the console opens the license screen. Only sign-in and the account routes, the license routes, first-run setup and the roll-back of a response action keep working. **Enrolled computers keep reporting. While locked, the browser extension gets the Monitor Only policy, which never blocks, and no new blocking action reaches a computer. Blocks applied earlier stay until rolled back (the roll-back route is the one exception to the lock, so an administrator can lift them).**

Lock code	Message shown	Screen title	What to do
`license_activation_required`	"This YuviRed server is not activated. The 7 days to activate it have ended, so the console is unavailable until you activate it. Enrolled computers keep reporting, and YuviRed only monitors until then."	Activate YuviRed	Activate with a code or license file.
`license_server_changed`	"This server's hardware changed, so YuviRed treats it as a copy of another server. The console is unavailable until a new license is activated for this server. Enrolled computers keep reporting, and YuviRed only monitors until then. If you moved or restored this server, email support@yuvired.com with the installation ID shown on this screen."	This server needs a new license	Email support with the installation ID.
`license_expired`	"This YuviRed license has expired. The console is unavailable until a new license is activated. Enrolled	Your YuviRed license has expired	Renew at https://yuvired.com/request and activate.

Lock code	Message shown	Screen title	What to do
	computers keep reporting, and YuviRed only monitors until then."		
`license_revoked` (by YuviRed)	"YuviRed revoked the license for this server. The console is unavailable until a new license is activated. Enrolled computers keep reporting, and YuviRed only monitors until then."	This YuviRed license was revoked	"If you think this is a mistake, email support@yuvired.com."
`license_revoked` (on this server)	"This YuviRed license was revoked on this server. The console is unavailable until a new license is activated. Enrolled computers keep reporting, and YuviRed only monitors until then."	This YuviRed license was revoked	Activate a new license.
`license_invalid`	"This YuviRed license could not be verified. The console is unavailable until a valid license is activated."	This YuviRed license could not be verified	Activate a valid license.

Notes: the lock is decided by the server, not by the screen. A locked server's Add Computer shows no free seats. A revoked license cannot be activated again. The screen title during the 7-day window (opened from the banner) is also "Activate YuviRed".

3. Activation messages (license screen)

The license screen sends the activation code, this server's installation ID, computer name, Windows version and YuviRed version to https://activate.yuvired.com. A code is valid for 7 days and activates one server, once. The console shows the service's message and, for refusals that need a person, a "contact YuviRed" link. Every refusal also shows: "Need help? Email support@yuvired.com and include the installation ID."

Shown by this server

Code	Message	HTTP	What to do
`invalid_code` (empty)	"Enter the activation code."	400	Type the code.
`activation_unreachable`	"YuviRed could not be reached. Check this server's internet access and try again, or use a license file."	502	Check internet access and the firewall, or use the license file route.
`activation_not_configured`	"Online activation is not set up on this server. Use a license file instead."	409	Use the license file.
`activation_bad_response`	"YuviRed sent an unexpected answer. Try again in a few minutes."	502	Retry later.
`activation_refused`	"YuviRed did not accept this activation code."	400	Shown when the service refused without naming a reason. Check the code; request a new one.
(service fault)	"YuviRed could not activate this server right now. Try again in a few minutes."	502	Retry later.
`license_server_check_incomplete`	"YuviRed could not finish checking this server's hardware. Restart the YuviRed service, then try again."	409	Restart the service, then retry.
other license refusals	`license rejected: <reason>`	400	See the reasons below.
permission	"Missing required permission: license.manage"	403	Sign in with the first (owner) account.

Reasons after `license rejected:`

"empty license"; "unrecognized license format"; "malformed license encoding"; "malformed license payload"; "license signature verification failed"; "license is bound to a different appliance installation" (the license is for another server: request one for this installation ID); "this license has already expired"; "this YuviRed server takes a Lite, Pro or Enterprise license"; "a Lite license allows at most 10 computers"; "a Lite license allows at most 5 AI agents"; "a Lite license never expires, so it cannot carry an expiry date"; "a Pro or Enterprise license must have an expiry date"; "a plan license must be bound to one server (installation ID)"; "license serial ... has been revoked"; "this server has held a Pro or Enterprise license, so it does not take a free Lite license. Contact YuviRed to renew the paid license" (a server that was Pro or Enterprise cannot go back to free Lite).

License file errors

"the license file is empty"; "the license file is too large to be a YuviRed license" (limit 64 KB); "no YuviRed license was found in the file"; "the file holds more than one license". Console: "That file is too large to be a YuviRed license file."

Sent by the activation service

Code	Message	HTTP	What to do
`invalid_code`	"That activation code is not valid, was already used on another server, or has expired."	400	Check the code (valid 7 days, one server, one use). Request a new code.
`invalid_installation`	"The server sent no valid installation ID."	400	Retry; contact support if it repeats.
`server_allowance_used`	"Your company already has 2 YuviRed Lite servers. To add another server, contact YuviRed support and tell us why you need it."	403	Contact support. This is the per-company limit (see editions-and-limits.md).
`rate_limited`	"Too many activation attempts. Try again in an hour."	429	Wait an hour.
`signer_unavailable`	"Activation is temporarily unavailable. Try again later, or use a license file."	503	Retry or use a license file.
`too_large`, `bad_request`	"The request is too large..." / "The request must be a JSON object."	413 / 400	Not expected from the console; contact support.

When you ask for a Lite code on the website (https://yuvired.com), you may instead see: `free_email_domain` ("Use your work email address. Free email addresses cannot download YuviRed Lite..."), `domain_not_allowed`, `email_allowance_used` ("Two people at your company already use YuviRed Lite. To add another person, tell us why you need it..."), `daily_limit_reached` ("Your company has reached today's limit for YuviRed Lite downloads..."), `invalid_verification` ("That code is not right, or it has expired. Ask for a new one."), `rate_limited_other` ("Too many attempts. Try again in an hour."), `email_unavailable`. Each tells you to wait, use another address or contact support.

4. Endpoint states

Liveness (Endpoints page and Dashboard)

Derived by the server from the time of the computer's last check-in. Not stored.

State	Rule	Meaning	What to do
Online	last check-in within 180 seconds	Reporting now. The agent checks in about every 60 seconds.	Nothing.
Stale	181 to 600 seconds	A soft warning: a few check-ins missed.	Wait a few minutes.
Offline	more than 600 seconds, or never checked in	Not reporting. A rejected credential looks identical to a powered-off computer.	Check the computer is on and can reach the server; see endpoints.md.

Freshness of AI activity

Live means the AI interaction was observed within 24 hours; **Historical** means older. This is separate from liveness.

Lifecycle

Lifecycle	Holds a computer seat?	Meaning
`active`	Yes	Enrolled and managed.
`quarantined`	Yes	Temporarily isolated, still enrolled.
`decommissioned`	No	Retired by an administrator. The computer is refused (HTTP 410) and cannot re-register with its old identity.
`archived`	No	Retired. Same refusal.

The server's own agent never uses a seat. A Windows Subsystem for Linux (WSL) distribution linked to its host computer shares the host's seat; a virtual machine or container with its own collector counts as its own computer.

Adoption requests

`pending`, `approved`, `rejected`, `expired`. Approving or rejecting a request already decided gives HTTP 409 "Request is already <status>".

AI agent (AMS) seat state

`seated` (counts against the AI agent limit), `unseated` (known only from a configuration or key file, no running process or connection seen, not counted), `over_limit` (detected above the limit; stays visible with its evidence, not counted).

AMS approval state

Unreviewed, **Under review**, **Approved**, **Restricted**, **Unauthorized**, **Referred to AIDR**. Amber badges mean awaiting a human decision.

---

5. Browser extension and prompt protection status

On Prompt Monitoring, each browser that has checked in shows a protection badge, a policy state and a version.

Badge	Meaning	What to do
Protected (green)	Seen within 10 minutes, the user consented, the last policy pull worked, and the active policy is not Monitor Only.	Nothing.
Monitoring only (grey)	Same, but the active policy is Monitor Only: warnings and blocking are off. Prompts are still recorded.	Choose a stronger preset if you want warnings or blocks.
Stale (amber)	Not seen for more than 10 minutes. Only the absence of a report is known: "(no recent report; YuviRed cannot tell why)".	Check the browser is open and the computer reaches the server.
Protection unavailable (red)	The user has not consented, or the last policy pull failed (the reason is shown).	Fix the cause shown (consent, connection, certificate).

Policy state (shown as "policy applied", "policy pending", "policy stale" or "policy unknown")	Meaning
`applied`	The browser runs the active policy revision.
`pending`	The browser runs an older revision and will update.
`stale`	Not fresh, so what it runs cannot be confirmed.
`unknown`	No state was reported.

A computer whose agent reports but whose browser has not yet appeared shows the amber label "No browser reporting yet" with the agent's own online or offline state. The card's count line reads "N protected, N monitoring only, N stale, N unavailable".

An **update needed** tag means the extension is older than the version on this server. On Lite, run a new Add Computer download on that computer, then restart the browser or choose Reload on the extension at chrome://extensions or edge://extensions. Extension lifecycle changes (Installed, Removed, Enabled, Disabled, Version change, Permission change, Manifest change) appear on the Endpoints page.

6. Alert, incident and action statuses

Where	Values
Alert status	`open`, `acknowledged`, `investigating`, `resolved`, `false_positive`, `suppressed` (active ones: open, acknowledged, investigating)
Alert severity	`critical`, `high`, `medium`, `low`
Incident status	`open`, `investigating`, `contained`, `resolved`
Response (enforcement) action	`pending` (requested), `approved`, `running`, `succeeded`, `failed`, `rejected`, `cancelled`, `rolled_back`, `expired` (delivery deadline passed; nothing is delivered after it), `recalled` (an administrator stopped a running action)
Deployment campaign	`draft`, `active`, `paused`, `completed`, `revoked`, `expired`
Campaign delivery steps	package_delivered, agent_installed, service_running, enrolled, identity_established, native_host_registered, extension_installed, extension_healthy, policy_synced, first_checkin, fully_ready; exceptions pending_offline, failed
Artifact analysis	Pending, Running, Completed, Partial, Failed, Cancelled; provider status OK, Degraded, Down, Inactive

7. First-run setup refusals

Message (code)	HTTP	Meaning and action
"Full name is required"	400	Fill in the name.
"A valid email address is required"	400	Enter an address with an @ and a domain.
"Passwords do not match"	400	Retype both.
Password rule text (for example length 12, three character classes, common words, your own name)	400	Choose a stronger password; see roles-and-permissions.md.
"You must accept the End User License Agreement and read the Monitoring and Data Notice" (`eula_not_accepted`)	400	Tick the acceptance box on the setup screen. For an unattended install, pass `ACCEPT_EULA=1`; without it the installer exits with code 5101.
"Setup has already been completed"	409	An administrator already exists. Sign in.
"That email address is already registered"	409	Use another address or sign in.
"Setup is only available from the computer YuviRed is installed on"	403	Open the setup page in a browser on the server computer itself.

8. HTTP errors a person or an agent can meet

HTTP	Where	Message	Meaning and action
401	Console	"Not authenticated"; "Invalid or expired token"; "User not found"	Sign in again (a session lasts 8 hours).
401	Sign-in	"Invalid credentials"	Wrong password, unknown or inactive user, or an expired temporary password. The three look the same on purpose.
401	MFA step	"Invalid or expired pre-auth token"; "Invalid code"	The 2-minute step ended; sign in again. Or the code was wrong.
401	Agent	"Invalid deployment token"; "Deployment token expired"; "X-Deploy-Token required"; "Not a campaign bootstrap credential"	The token or Add Computer download is wrong or old (an Add Computer download is valid for 24 hours). Make a new download.
402	Console	a license lock (section 2)	Activate the server.
402	Agent (enrolled computer)	a license stop	The agent logs it, backs off to 300 seconds, keeps its identity and resumes when the condition clears.
403	Console	"Missing required permission: <code>"	Your role lacks that permission; ask an administrator.
403	Console	`password_change_required`: "A password change is required before this account can be used. Change it via /auth/change-password."	Change your password.
403	Configuration	"Artifact Analysis configuration requires the artifact.manage permission."	Only an Admin can read or change the

HTTP	Where	Message	Meaning and action
			Artifact Analysis rows.
400	Configuration	"<key> is license state, not a configurable value..."; "<key> is the recorded acceptance of the End User License Agreement..."; "security.tls_enforcement must be one of: off, warn, enforce..."; "branding.privacy_contact must be at most 200 characters." and similar "must be" messages	The value is refused; the message names the rule. The first two are recorded in the audit log.
404	Configuration	"Unknown setting: <key>"	The key does not exist on this server version.
403	New computer enrolling	"license seat limit reached"; "license activation required"; "license revoked"; "license expired"; "license invalid"; "license server changed"; "server identity changed: an administrator must confirm it"	The computer is refused. See editions-and-limits.md.
403	Agent	"This credential is not authorized for that endpoint"	The credential does not match; the agent retries recovery with the credentials it holds.
404	Agent	"Endpoint not registered"; "Endpoint is not linked to a deployment campaign"	The computer is not enrolled (for example it was removed), or its installer reports progress for a campaign it does not belong to. Reinstall with a new Add Computer download.
409	Add Computer	"Only N computer seats are free."; "Every computer seat is in use. To free one, open Endpoints and remove a computer you no longer use. Removing it deletes that computer's history."	Free a seat or ask for fewer computers.

HTTP	Where	Message	Meaning and action
409	Add Computer	"Network mode is off, so other computers cannot reach this server..."; "The agent installer is not part of this server build."; "This server's certificate could not be read."; "This server's hardware identity changed. Confirm it before adding computers."	Fix the cause stated.
409	Configuration	"<key> is superseded by the Prompt Protection policy..."	Edit the policy under Configuration, Prompt Protection.
409	MFA	"MFA is already enabled; disable it first to re-enrol"; "MFA is not enabled"; "No enrolment in progress..."	Follow the message.
409	Campaign	"Campaign is '<status>'; no further progress reporting accepted"; "Cannot <action> a campaign in status '<status>'"	The campaign is not in a state that allows that action.
410	Agent	"This endpoint has been decommissioned by an administrator."	Retired computer; it stays refused.
410	Agent	"Campaign no longer exists"	Make a new Add Computer download.
426	Agent or extension	"Agent version X is no longer supported. Minimum required: Y..."	Update the agent (run a new Add Computer download).
426	Agent or extension	"Plaintext HTTP is not accepted for agent/extension endpoints..."	Point the agent at the HTTPS address; this appears only when Agent Connection Security is set to Require.
429	Sign-in, MFA	"Too many failed login attempts. Please wait and try again."; "Too many failed codes for this account. Retry in Ns."	Wait (8 failures per 5 minutes per address and account; 30 per hour per account). Signing in at the server

HTTP	Where	Message	Meaning and action
			itself is exempt from the per-account hour limit.
429	Agent	"Too many registration attempts; slow down"; "Too many rotation requests; slow down"; "Too many campaign report attempts; slow down"	Wait and retry.
429	Activation	"Too many activation attempts. Try again in an hour."	Wait an hour.

9. License banners, buttons and toasts

Text	Meaning
"License activated"	The license was accepted.
"Installation ID copied" / "Copy failed: select the installation ID and copy it by hand."	Result of the **Copy** button.
"The license was stored, but the console is still locked. Email support@yuvired.com."	The license is stored but does not unlock this server (for example it is for another installation ID).
"Activating sends YuviRed this server's installation ID, computer name, Windows version and YuviRed version."	The four facts the activation sends. After activation the server checks in once a day with six facts: format, installation ID, version, serial, computer count and AI agent count. No computer names, user names or prompts.
Update banner	A newer YuviRed exists. Nothing installs itself.

Related

- editions-and-limits.md
- roles-and-permissions.md
- settings.md
- console-pages.md

Editions and limits reference

What Lite, Pro and Enterprise include, what counts against each limit, what happens at a limit, the per-company limits on automatic Lite codes, and what happens at expiry, revocation and the end of the activation window.

Covers YuviRed 1.0 for Windows. One installer serves every edition; the license decides the edition. Edition terms below are the owner-approved terms recorded in the documentation facts sheet; the limits and behaviors are checked against the server's code.

The three editions

	Lite	**Pro**	**Enterprise**
Price	Free	US$1,500 per year	Custom
Expiry	None: "YuviRed Lite is free and never expires."	Has an end date; no grace after it	Has an end date; terms as agreed
Computers (endpoints)	**10**, plus the server itself	**50**, plus the server itself	As stated in the license
AI agents under AMS	**5**	**50**	As stated in the license
How you get it	Request an activation code on https://yuvired.com with a verified work email address (or a license file)	Request at https://yuvired.com/request; YuviRed approves it, you pay, then the license is issued	Contact YuviRed
Trial	Not applicable	**No trial**	As agreed
Renewal	Not needed	Manual, at https://yuvired.com/request	Manual
Support	Email support@yuvired.com, reasonable efforts, no response-time promise	Same	As agreed

How the server enforces the plan: a plan license is a signed license whose tier names the plan. The server refuses, at activation and again at every check, a Lite license that allows more than 10 computers or 5 AI agents or that carries an expiry date, and a Pro or Enterprise license that has no expiry date. A license must be bound to one server (its installation ID), and must allow at least one computer and state an AI agent limit. A database edit cannot widen a plan.

What counts against each limit

Computers (the endpoint limit)

- A **computer** counts when it is enrolled and its lifecycle is `active` or `quarantined`.
- **The server's own agent does not count.** On Lite at most one endpoint, the one enrolled with the server's own token, is exempt. If the server is listed under Endpoints it does not use a computer seat. A Lite server therefore monitors 10 other computers plus itself.
- A **Windows Subsystem for Linux (WSL) distribution** that the server has linked to its host computer shares the host's seat and does not count again (the host must be active). A **virtual machine** or a **container** with its own agent counts as its own computer, even when linked. An unlinked inner collector, or one whose host is retired or deleted, also counts.
- A **decommissioned** or **archived** endpoint does not count, and it may not report either. **Removing an endpoint** (Endpoints page, **Remove endpoint**) is what frees a seat, and it **deletes that computer's history**. The console has no separate Decommission control.

AI agents (the AMS limit)

- An **AI agent** counts only when it holds a seat: it must be known to be running, meaning a running process or a network connection was seen. An agent known only from a configuration or key file is **unseated** and is not counted.
- An agent detected beyond the limit is marked **over_limit**: it stays visible with its evidence and does not count.
- The correlation job (setting **Agent Activity Tracking**) is the only thing that assigns seats. Turning it off stops new agents from being linked.

Where you see the numbers

- Configuration, **License** panel: **Computers used** and **AI agents used**, each "N / limit". A counter turns red above the limit.
- Agent Deployment, **Add Computer** panel: "Computers: N of 10" and "AI agents: N of 5" (limits follow the license), with the note "If the server itself is listed under Endpoints, it does not use a computer seat."
- Active license limits are read from the license in force; on a Lite server with no license yet (inside the 7-day window) the compiled Lite limits apply. While the server is locked the counters are not shown.

What happens at a limit

The limit is enforced when a NEW computer or AI agent tries to join. Nothing already enrolled is switched off.

Situation	What happens	What you can do
Computer limit reached (for example 10 of 10 on Lite)	The new computer is refused at enrolment with HTTP 403 and the plain message `license seat limit reached`. Add Computer shows "Every computer seat is in use. A new computer is refused until one is free. To free one, open Endpoints and remove a computer you no longer use. Removing it deletes that computer's history." A download that would exceed the free seats is refused: "Only N computer seats are free."	Remove a computer you no longer use, or move to Pro.
Computers already enrolled	Keep reporting. Nothing is stopped.	
AI agent limit reached (for example 5 of 5 on Lite)	A newly detected AI agent is marked **over_limit**: visible with evidence, not counted, not seated.	An archived agent does not count; whether the console lets you archive an agent is covered in ams.md (not verified here). Or move to Pro.
Server locked (not activated, expired, revoked, changed)	New computers are refused (`license activation required`, `license expired`, `license revoked`, `license server changed`). The Add Computer panel shows 0 free seats. Computers already enrolled keep reporting.	See statuses-and-errors.md section 2.
Lite license that is not valid on this server	New computers refused: "This license is not valid on this server, so new computers are refused." The console stays open.	Activate a license for this server.
A server that held a Pro or Enterprise license	Does not take a free Lite license: "this server has held a Pro or Enterprise license, so it does not take a free Lite license. Contact YuviRed to renew the paid license".	Renew the paid license.

Enrolment checks run under a lock, so two computers enrolling at the same moment cannot both take the last seat.

The activation window (a new server)

- A **new server** gets **7 days** from its first start in which the console works with the Lite limits (10 computers, 5 AI agents). A banner shows the days left.
- **After the 7 days** with no activation, the console stops serving (HTTP 402 `license_activation_required`). **Computers already enrolled keep reporting, and the software falls to monitoring only**: the browser extension is told to only observe, no new blocking action reaches a computer, and blocks applied earlier stay until they are rolled back. New computers are refused (except the server's own agent).
- **Ways to activate:** an **activation code** (valid 7 days from issue, activates one server, usable once) typed into the license screen, or a **license file** for a server without internet access, made at https://activate.yuvired.com/license-file from the code and the server's installation ID.
- **A fresh database starts the 7 days again** (the first start is recorded in the database), so a purge reinstall counts as a new server with a new installation ID. A server that was activated once never gets the window again on the same database. The server's own clock decides; a clock set back extends the days left.
- Pro and Enterprise licenses are issued by YuviRed after approval (and, for Pro, payment). A Pro or Enterprise code made by YuviRed is activated the same way.

Per-company limits on automatic Lite codes

YuviRed's activation service, not your server, applies these when someone at a company asks for an automatic Lite code on https://yuvired.com. The server only shows the service's answer.

Limit	Value	When exceeded
Servers per company	**2** Lite servers	`server_allowance_used`: "Your company already has 2 YuviRed Lite servers. To add another server, contact YuviRed support and tell us why you need it." Activating the same server again (same installation ID) does not use a second slot.
Verified people per company	**2** people	`email_allowance_used`: "Two people at your company already use YuviRed Lite. To add another person, tell us why you need it and the YuviRed team will reply by email."
Codes per day per company	**3**	`daily_limit_reached`: "Your company has reached today's limit for YuviRed Lite downloads. Tell us what you need and the YuviRed team will reply by email."
Free email addresses	Not accepted	`free_email_domain`: "Use your work email address..." If the business has no work address, contact support.

Other fixed numbers: an activation code is valid at most **7 days** and works for one server once; a website verification code lasts 30 minutes and allows 5 attempts; a company's codes are counted by the email domain. Rate limits (service defaults): 3 website requests per email address per hour, 20 per network address per hour, and overall ceilings ("Try again in an hour"). **The three limits above are the activation service's defaults** (its settings `ACTIVATION_MAX_SERVERS_PER_COMPANY`, `ACTIVATION_MAX_EMAILS_PER_COMPANY` and `ACTIVATION_MAX_AUTOMATIC_PER_DOMAIN_PER_DAY`); YuviRed can set them differently, including to zero, in which case the refusal reads "To add a YuviRed Lite server for your company, contact YuviRed support and tell us why you need it." A refusal always offers a contact route. Per-company limits apply to automatic Lite codes only; Pro and Enterprise licenses are issued by YuviRed by hand, and a server activated with a code YuviRed made by hand also counts as one of the company's servers.

At expiry

- **Lite never expires.** A Lite license cannot carry an expiry date, and the server refuses one that does.
- **Pro and Enterprise have an end date and no grace after it.** Banners appear 30, 14 and 7 days before: "Your YuviRed Pro license ends in N days (date). When it ends, the console is unavailable until a new license is activated. Enrolled computers keep reporting, and YuviRed only monitors until then. To renew, email support@yuvired.com." The plan state shows **Ends soon**.

- **On the end date** the server locks: HTTP 402 `license_expired`, the license screen opens, enrolled computers keep reporting, and monitoring only applies (extension observes, no new blocks, earlier blocks stay). New computers are refused ("license expired").
- **To unlock:** activate a new license (renew at https://yuvired.com/request). The renewal is manual; nothing renews itself.
- The server's own clock is trusted: moving it back extends the license, moving it forward ends it early.

At revocation

- **Lite:** a local administrator cannot revoke it (the server answers HTTP 409 "the free YuviRed Lite allowance cannot be revoked"), and the console shows no revoke button on a YuviRed server (the **Revoke Active License** button belongs to the Docker appliance edition's panel). Only YuviRed can revoke a plan license, in its reply to the daily check-in.
- **When YuviRed revokes** (the reply names this server's exact serial): the console locks with `license_revoked`: "YuviRed revoked the license for this server. The console is unavailable until a new license is activated. Enrolled computers keep reporting, and YuviRed only monitors until then." If you think it is a mistake, email support@yuvired.com. A revoked license can never be activated again.
- **A server that cannot check in** (for example no internet) is never revoked remotely, and a revocation or a "newer version" notice can take up to about 25 hours to appear after YuviRed sets it.
- **The daily check-in:** once every 24 hours an activated server sends six facts: format, installation ID, version, serial, computer count and AI agent count. It never sends computer names, user names, endpoint names, addresses or prompts. A failed check-in changes no limit and locks nothing. The last attempt and last success are shown on the license panel (**Last check-in**).
- **A newer version** named in the reply shows as a banner. Nothing installs itself.

When the server's hardware changes

A restore onto new hardware, a copied disk, a motherboard swap or a hypervisor change that alters the virtual firmware identity gives the server a new installation ID. Until a license bound to the new ID is activated the console is locked (`license_server_changed`) and the state is **Server changed**. Moving back does not undo it. Email support@yuvired.com with the

installation ID. A boot that cannot read the firmware once can look like a change and lock the server the same way.

Things these limits do not do

- No limit stops a computer that is already enrolled from reporting, except that a decommissioned or archived computer is refused.
- The console is the only thing a lock closes; detection data keeps arriving.
- There is no trial of Pro and no grace period after a Pro license ends.
- Per-company limits cannot be seen or changed from the console. They live in YuviRed's activation service.

Related

- statuses-and-errors.md
- console-pages.md
- settings.md
- roles-and-permissions.md

Monitoring disclosures and permissions (reference for administrators)

YuviRed 1.0, agent 1.8.2, extension 0.5.5. One place for what the software needs on a computer, what it changes, and what you must tell people. This is not legal advice. You, not YuviRed LLC, are responsible for notice, consent and works council duties where you operate.

What the agent needs

Item	Detail
Install	a **local administrator**, running `Install-YuviRed-Agent.cmd` from PowerShell as administrator, from a folder only administrators can change
Runs as	the Windows service `SyntinelRedAgent` ("YuviRed Endpoint Agent"), **SYSTEM**, delayed automatic start, restart on failure. The service name, the state folder and the native messaging host name below carry the product's earlier name; it will change in a later release
Files	`C:\Program Files\YuviRed Agent` (includes `NOTICE.txt`, the `extension` folder for Load unpacked, and `setup\store-extension-ids.json`, empty in this release)
State and log	`C:\ProgramData\SyntinelRed` (SYSTEM and Administrators only); log `logs\syntinel-agent.log`
Network	outbound HTTPS to your server on TCP 8443 only; it trusts only the server's own certificate from `enrol.yuvired` and changes no Windows trust store. It also asks the computer's DNS server for the name of every remote address it reports (a reverse lookup)
Local port	the Agent Monitoring System (AMS) bridge listens on 127.0.0.1 port 47811 whenever the agent runs as a service, connector or not; it accepts nothing from other computers
Never collected by the agent	keystrokes, clipboard, screenshots, page contents, prompt and reply text, document contents

Browser and registry changes

Change	When
`HKLM\SOFTWARE\Google\Chrome\NativeMessagingHosts\com.syntinelred.native_host`, the same under Microsoft Edge, and under `SOFTWARE\Mozilla\NativeMessagingHosts`	always, at i
Chrome and Edge policy: force-install entry, native messaging allowlist, extension `backendUrl` and `captureLevel` (follow server)	**not in this** writes thes later build, records wh `HKLM\SOFTW` and remove "Browser e Edge stores by hand
Chrome and Edge policy written by the agent at run time, including `consented` = 1	only where **Browser E** YuviRed Lit holds no ex written
Browser blocked-URL or extension-disable policy values	only when response a
Managed settings files of Claude Code, Codex, Gemini CLI, Cursor, Copilot CLI, Windsurf (hooks)	only when connector f

The server also writes one inbound firewall rule (TCP 8443, Domain and Private profiles) when network mode is on.

What to give the people who are monitored

1. The monitoring notice (`notice.html` , template `frontend/legal/notice.md`) and, for detail, the data notice (`data-notice.html`). The notice marks lines "[Your organization: ...]" that you must complete or delete: purpose, legal basis, retention you set, whether full capture or AMS connectors are on, who has console accounts, where the server is. Set **Organization Name** and **Privacy contact** in **Configuration, System** so the notice shows them.

2. Notice **before** installation, and again before you switch on raw capture or AMS connectors.
3. The privacy policy (`https://yuvired.com/privacy`) and terms (`https://yuvired.com/terms`) where your policy needs them.

Every server serves: `https://<server>:8443/static/legal/` with `manifest.json`, `eula.html`, `privacy.html`, `notice.html`, `data-notice.html`, `third-party.html`. The console **Legal** page under Admin lists them. The agent notice is `NOTICE.txt` in the YuviRed Agent folder, and each Add Computer `README.txt` repeats the duty.

The acceptance record

Installing the server records who accepted the terms at `%ProgramData%\YuviRed\status\legal-acceptance.json` (unattended install: `YuviRed-Lite.exe /quiet ACCEPT_EULA=1`; exit 5101 without the parameter). Running the agent installer is a separate act: it confirms that your organization may monitor the computer and has told its users. YuviRed does not record that confirmation per computer.

Things people often miss

- Every web connection of every program is reported, not only AI.
- Any person who can sign in to the console can read endpoint identity and raw telemetry, whatever the role.
- Masked fragments are stored unencrypted; raw capture (off by default) is encrypted and kept 24 hours.
- A person using a browser can narrow the extension's checks on its options page; the console is not told.
- In this release every extension is loaded by hand and consent is the person's own tick on the options page. Nothing in the product sets consent by policy on a Lite server; the agent's wiring route that would is inactive.
- A check-in that cannot be sent is lost: the agent keeps no offline queue, so a computer that was away for a day reports nothing for that day.
- Uninstalling keeps `%ProgramData%\YuviRed` on the server; the agent removal deletes its own data.

Notice templates: `frontend/legal/notice.md`, `frontend/legal/data-notice.md`, `frontend/legal/privacy.md`. Related guides: endpoints.md, extensions.md.

Troubleshooting

Find your symptom, then follow its entry. Each entry has the likely causes, the steps, and how to confirm the fix. If you are stuck, collect diagnostics (Collecting diagnostics) and write to support@yuvired.com.

Two agent names. The server's own agent is the Windows service `YuviRedAgent`. On every other computer the agent's service is `SyntinelRedAgent` (shown in Services as **YuviRed Endpoint Agent**), and on every computer, the server included, the agent writes its log under `C:\ProgramData\SyntinelRed` (the folder name is the product's earlier name; it will change in a later release).

Commands run in **Windows PowerShell as administrator** on the server. The server command is long because nothing is on the PATH; it is written out in full where needed:

```
& "C:\Program Files\YuviRed\runtime\python.exe" "C:\Program
Files\YuviRed\runtime\yuvired_service.py" <command>
```

Contents

1. The server: service will not start, console unreachable, certificate warnings, port in use, database failed, disk full, time skew, sign-in problems
2. Licensing: code refused, activation service unreachable, locked console, server changed
3. Computers (endpoints): not appearing, stale, log codes, token expired, rotation failure
4. The browser extension
5. Notifications
6. Reports
7. Where the logs are

---

1. The server

1.1 The YuviRed service does not start (or stops at once)

Symptom. In Services, `YuviRed` is **Stopped**, or starts and stops again. The console does not open. The installer, if just run, said the install failed.

Likely causes.

- The database service `YuviRedPostgres` is not running. `YuviRed` depends on it.
- The boot stopped at a step. The boot status file names it.
- The data folder is not safe: it, or a folder in it, was replaced by a link or is not owned by an administrator, or its permissions were changed.
- The database is newer than the program (after installing an older build over a newer one, or after an update that was rolled back once the new version had already changed the database).
- The identity key does not open: the data folder was copied from another computer or Windows installation. The boot status names the step `provision_identity` and "the appliance private key could not be loaded" (Administration, section 2.2).
- Port 8443 is taken (see 1.4) or the disk is full (see 1.6).
- The hardware identity changed (see 2.4). The console is locked, but the service itself still runs.

Steps.

1. Read the boot status:

   ```
   Get-Content C:\ProgramData\YuviRed\status\boot-status.json
   ```

   You should see: `"phase": "YuviRed Ready"` when healthy. Otherwise the file shows the phase it stopped in (`Preparing`, `Secrets Ready`, `Database Ready`, `Schema Current`, `Platform Bound`, `Identity Created`, `Registration Pending`, `Restricted`, `Clone Suspected (Restricted)`, `Error`) and, for `Error`, which step failed. The file holds no secrets.

2. Start the database first, wait ten seconds, then start the server:

   ```
   net start YuviRedPostgres
   net start YuviRed
   ```

3. Open Event Viewer, **Windows Logs**, **Application**, and look at sources `YuviRed` and `YuviRedAgent`. An error "YuviRed stopped at phase ..." names where it stopped.
4. Read the end of `C:\ProgramData\YuviRed\logs\syntinel-red.log` and the newest file in `C:\ProgramData\YuviRed\pg\data\log`.
5. If the log says "Can't locate revision", the installed program is older than the data. Install the newer version over it.
6. If the boot status or the log says a path "is a link" or "was not created by an administrator", something replaced a YuviRed folder. Do not delete it. Collect diagnostics

and write to support.
7. Run the support collector before you restart the computer or reinstall.

How to confirm. `boot-status.json` reads `YuviRed Ready`, all three services read **Running** after a minute, and `https://127.0.0.1:8443` opens.

Common problems. If the service starts but stops after a few seconds, read steps 3 and 4. Do not reinstall as a first step: a reinstall over a damaged install can leave a partial state.

1.2 The console does not open

Symptom. The browser says "This site can't be reached", "connection refused" or times out.

Likely causes and steps.

Cause	Check and fix
The service is not running	See 1.1.
You are on another computer, and the server is in its default mode	By default the server answers only on the computer itself (`https://127.0.0.1:8443`). To serve other computers, turn on network mode on the server: `... yuvired_service.py network --name <dns name or fixed IP>`. See the Installation guide.
Network mode is on but the firewall blocks	The rule `YuviRed Server (TCP 8443)` applies to the **Domain** and **Private** network profiles only. A network set to **Public** is not allowed. Run `... network --status` to see the profile and whether the rule is present. Set the connection's network profile to Private, or add a rule of your own.
The name does not resolve	From the other computer, `Resolve-DnsName yuvired.northwind.example`. Fix DNS, or use the address `192.0.2.10`.
A firewall or proxy between	Test with `Test-NetConnection 192.0.2.10 -Port 8443` from the other computer.
Wrong address	Use `https://`, and port `:8443`.

How to confirm. The sign-in page opens from the other computer. `Test-NetConnection ... -Port 8443` shows `TcpTestSucceeded : True`.

1.3 Browser certificate warnings

Symptom. "Your connection is not private", "NET::ERR_CERT_AUTHORITY_INVALID", or a name mismatch.

Likely causes.

- **Expected on first use.** The server makes its own certificate and no browser knows it. It is valid for 825 days and names the computer's name, `localhost`, `127.0.0.1` and the names and addresses you gave to network mode.
- **Name mismatch.** You opened the server by an address or name that is not in the certificate.
- **Wrong clock** on the server or on your computer (see 1.7): the certificate looks not yet valid or expired.

Steps.

1. For the first sign-in, proceed past the warning (**Advanced**, then continue). It is your own server.
2. For a name mismatch, open the server by a name or address in the certificate, or run `... network --name <the name you use>` so the certificate includes it. **A new certificate stops enrolled computers trusting the server** (they trust exactly the old one): they need a new **Add Computer** download. Use a DNS name from the start to avoid this.
3. To remove the warning, import the server's certificate (`C:\ProgramData\YuviRed\config\tls-cert.pem`, the public part only) into the trusted store of the computers' browsers, by your own means. Never copy `tls-key.pem`.

How to confirm. The address bar shows the certificate without a warning, or you chose to proceed once.

1.4 Port 8443 or 5440 is already in use

Symptom. The service stops at start; the log says it could not bind; or the installer's wait fails.

Steps.

1. Find the program that holds the port:

```
Get-NetTCPConnection -LocalPort 8443 -State Listen | Select LocalAddress,
OwningProcess
Get-Process -Id <OwningProcess>
```

Use `5440` for the database port.

2. If it is another program, stop it or move it to another port. YuviRed's ports are fixed. A way to change them is not documented (not verified).
3. If it is a leftover `pythonservice` or `postgres` of YuviRed, stop the `YuviRed` and `YuviRedPostgres` services, wait, and start them again.

How to confirm. `Get-NetTCPConnection -LocalPort 8443 -State Listen` shows the YuviRed service's process (`pythonservice`).

1.5 The database service failed to start

Symptom. `YuviRedPostgres` will not start, or the console shows server errors; the `YuviRed` log says it cannot connect to the database.

Likely causes. Disk full. The data folder or its permissions were changed. A restore that left `config\pg.pwfile` and `pg` from different times. Antivirus holding the files. A crash that left the database needing recovery (it usually recovers itself at start).

Steps.

1. Read the newest file in `C:\ProgramData\YuviRed\pg\data\log` (names are `postgresql-<date>.log`). The last lines say why.
2. Check free space (`Get-PSDrive C`). If it is nearly zero, see 1.6.
3. Check that nothing else (antivirus, backup software) is locking the folder: exclude `C:\ProgramData\YuviRed` from scanning and from file-level backup while the service runs.
4. Start it again and wait a minute: after a crash it replays its log, which can take time.
5. **Do not delete files in** `pg\data` , do not run a repair tool on it by hand. Collect diagnostics, then write to support.

How to confirm. `YuviRedPostgres` is **Running**, port 5440 listens, and `YuviRed` starts.

1.6 The disk is full

Symptom. Services stop; the database log says "No space left on device"; the console is slow or errors.

What fills the disk. The database (raw telemetry is kept 30 days by default, and the agent reports every web connection of every program), the logs (bounded), report files (capped at 2 GB), and your own copies of the data folder.

Steps.

1. Free space outside YuviRed first: other files, old copies, installers.
2. When there is room to start, start `YuviRedPostgres` then `YuviRed`.
3. In the console, **Configuration**, **Data & Retention**: lower **Raw telemetry** (for example from 30 to 14 days), then, under **Telemetry Retention**, click **Run now (dry run)** and then **Run now**. A database does not shrink the disk file at once: space is reused, and may only be returned later.
4. Move report files off with **Reports**, **Delete files**.
5. Add disk space for the long term.

How to confirm. Free space is above 15 percent, no new "No space" lines appear in the logs, and the console loads.

1.7 Time skew

Symptom. MFA codes are refused although typed correctly; the certificate looks "not yet valid"; a paid license seems to end early or late; sign-ins end sooner than expected.

Cause. The server's clock (or a phone's, or a computer's) is wrong. The server trusts its own clock: moving it back extends a paid license's end date, moving it forward ends it early. Authenticator codes change every 30 seconds.

Steps.

1. `w32tm /query /status` shows the last sync. Repair: `w32tm /resync`. In a domain the server gets time from the domain; elsewhere, set an internet time source in Windows **Date & time**.
2. On a virtual machine, turn on the host's time sync, or use Windows time, not both fighting.
3. Fix the phone's clock to automatic.

How to confirm. MFA codes are accepted, and `w32tm /query /status` shows a recent sync.

1.8 Sign-in problems

Symptom	Cause and fix
"Invalid credentials"	Wrong password. **A temporary password given by an administrator stops working after 72 hours and the console then also says "Invalid credentials".** Ask an administrator to reset it again.
"Too many failed login attempts. Please wait and try again."	8 wrong attempts for the same address and account within 5 minutes. Wait 5 minutes. Restarting the YuviRed service clears the counters.
"Session Expired" box	Your session ended (time limit, or an administrator reset your password or ended your sessions). Click **Log In Again**.
The sign-in code is refused	Check the phone's time (1.7). Use a recovery code, or ask another administrator to reset your MFA.
Nobody can sign in	See Administration, section 5.3, "When the only administrator is locked out".

---

2. Licensing

The license screen opens when the server is not activated after its 7 days, or when the license is not in force. Enrolled computers keep reporting and the extension only monitors, so staff are never blocked because of a license. New computers and new AI agents are refused until a license is active.

2.1 An activation code is refused

Symptom. After **Activate**, the screen shows a message, then "Need help? Email support@yuvired.com and include the installation ID."

Every refusal the service can give, with its code (the code is what support sees):

The screen says	Code	Meaning and what to do
"That activation code is not valid, was already used on another server, or has expired."	`invalid_code`	A code works once, for one server, for 7 days from when it was issued. Check you typed it exactly; get a new one at https://yuvired.com/download.
"Your company already has 2 YuviRed Lite servers. To add another server, contact YuviRed support and tell us why you need it."	`server_allowance_used`	The company limit is reached (counted by installation ID, over time). Use the **contact YuviRed** link, or write to support@yuvired.com. If you are replacing a server, say so and give both installation IDs.
"Too many activation attempts. Try again in an hour."	`rate_limited`	Wait an hour.
"Activation is temporarily unavailable. Try again later, or use a license file."	`signer_unavailable`	YuviRed's side. Try later, or use the license file route (2.2). The code is not used up.
"The server sent no valid installation ID."	`invalid_installation`	Restart the YuviRed service and open the license screen again. If it persists, write to support.
"The request is too large..." / "The request must be a JSON object."	`too_large`, `bad_request`	A fault on the way (a proxy altering the request). Try the license file route.
"Only a YuviRed administrator can activate this server. Sign in with the administrator account created when YuviRed was installed."	(403 from your server)	Only the first account (super admin) can activate. Sign in with it.
"That file is too large to be a YuviRed license file."	n/a	You chose the wrong file.
"The license was stored, but the console is still locked. Email support@yuvired.com."	n/a	Write to support with the installation ID.

The website (https://yuvired.com/download) can give these answers when you ask for a code: "Use your work email address. Free email addresses cannot download YuviRed Lite..." (`free_email_domain`), "Two people at your company already use YuviRed Lite..." (`email_allowance_used`), "Your company has reached today's limit..." (`daily_limit_reached`), "We cannot accept this email address..." (`domain_not_allowed`), "That code is not right, or it has expired. Ask for a new one." (the 6-digit email code, valid 30 minutes, 5 tries). The limits are 2 servers and 2 verified people per company, and 3 codes per day.

Confirm the fix. The screen says **License activated**, and **Configuration**, **License** shows **YuviRed Lite**, **Active**.

2.2 The activation service cannot be reached

Symptom. "YuviRed could not be reached. Check this server's internet access and try again, or use a license file." or "Online activation is not set up on this server. Use a license file instead."

Likely causes. The server cannot reach `https://activate.yuvired.com` on TCP port 443; a firewall or proxy; a proxy that replaces certificates; no internet; DNS.

Steps.

1. From the server: `Test-NetConnection activate.yuvired.com -Port 443`. Fix what blocks it.
2. Open `https://activate.yuvired.com/v1/health` in a browser on the server. A healthy service answers.
3. If the server must stay offline, use a **license file**: on any computer with internet, open https://activate.yuvired.com/license-file, enter the activation code and the installation ID, download the file, then on the server choose **Upload license file**.

Confirm the fix. License activated. After activation the server checks in once a day. If it cannot, it keeps working, and **Last check-in** stays at its last value.

2.3 The console is locked

The title tells you why.

Title	Meaning	What to do
Activate YuviRed	The 7 days to activate ended.	Activate (code or file). A server that was activated once never gets the 7 days again.
Your YuviRed license has expired (Pro, Enterprise)	The end date passed. No grace.	Renew at https://yuvired.com/request or write to support. Activate the new code or file.
This YuviRed license was revoked	YuviRed revoked it ("YuviRed revoked the license for this server...") or an administrator did ("...was revoked on this server...").	Write to support with the installation ID. A new code (a new serial) unlocks it. A revoked license cannot be used again.
This YuviRed license could not be verified	The stored license no longer verifies.	Activate a new code or file; if that fails, write to support.
This server needs a new license	See 2.4.	See 2.4.

"YuviRed could not finish checking this server's hardware. Restart the YuviRed service, then try again." means exactly that: restart the `YuviRed` service, then open the license screen again before you ask for anything.

2.4 "This server needs a new license" (server_changed)

Symptom. After a restore, a copy, a move, or a change of the virtual hardware, the console shows "This server's hardware changed, so YuviRed treats it as a copy of another server."

Cause. The server records the computer's hardware identity (firmware serial numbers) next to its installation ID and checks it at every start. If it changes (a virtual machine restored as a new machine, imported or copied, a motherboard change, some hypervisor upgrades), a server that held a license gives itself a **new installation ID** and locks (the audit log records `license.server_changed` with the old and new IDs). The license is bound to the old one. YuviRed cannot tell a legitimate move from a copy. Moving back does not undo it: it counts as another change.

Steps.

1. If this is a **copy** of another YuviRed server, shut the copy down.
2. If you moved or restored this server on purpose, copy the installation ID from the screen (**Copy**) and write to **support@yuvired.com**: say what you did, and give the old installation ID if you have it.
3. Activate the new code or license file support sends you, as in 2.1.
4. Meanwhile your computers keep reporting.

Confirm the fix. The screen reads **License activated** and the console opens.

---

3. Computers (endpoints)

3.1 A computer does not appear in Endpoints

Likely causes and steps.

Cause	Check and fix
The Add Computer download expired	Its token lasts 24 hours. The agent log says so ("The Add Computer download has expired. Make a new download on the server..."). Make a new download.
All seats are used	The agent log says "All computer seats on the server are in use. Uninstall the agent from a computer you no longer use and remove that computer under Endpoints on the server, then restart this agent. If the Add Computer download is more than 24 hours old, make a new one instead." Lite allows 10 computers plus the server. **At 10 of 10 the console refuses a new Add Computer download** ("Only 0 computer seats are free."), so free a seat first.
It waits for an administrator	"The server is waiting for an administrator to confirm it under Agent Deployment, Add Computer. Restart this agent after that." Confirm it there, or under **Adoption Requests**, then restart the agent.
The computer cannot reach the server	From it: `Test-NetConnection 192.0.2.10 -Port 8443`. Check the address in the download, DNS, the server's network mode and firewall (1.2).
The computer does not trust the certificate	Happens after the server's certificate changed (1.3). Make a new Add Computer download.
The license is not in force	New computers are refused while the license is expired, revoked, not valid or the server not activated. See 2.
The agent service is not running	On the computer, Services, **YuviRed Endpoint Agent** (service `SyntinelRedAgent`), or read the agent log.

How to confirm. The computer appears under **Endpoints** with **Online** within a few minutes, and the agent log shows a successful check-in.

3.2 A computer shows Stale or Offline

Status comes from the last check-in: **Online** up to 3 minutes, **Stale** up to 10 minutes, **Offline** beyond that (or never seen). The agent checks in every 60 seconds.

Steps.

1. Is the computer on, awake and on the network? A laptop that sleeps is offline.
2. Is the agent service running on it (`SyntinelRedAgent`, shown as **YuviRed Endpoint Agent**; on the server itself `YuviRedAgent`)? Start it.
3. Read its log (`C:\ProgramData\SyntinelRed\logs\syntinel-agent.log`): "Cannot reach backend at https://... (...)" means network or name, and the bracket gives the reason; a 401 or 403 means its credential is refused (3.3).
4. Is the server healthy (1.1)? If many computers are offline at once, look at the server first.
5. After repeated failures the agent slows down to one try every 5 minutes at most. Restart the service to force a try now.

How to confirm. The row turns **Online** within a few minutes.

3.3 What the numbers in the agent log mean (401, 402, 403, 410)

Code	Where you see it	Meaning	What to do
401	Agent log, or in a script or browser console	The server does not accept the credential presented, or no credential was sent.	The agent first tries every credential it holds. If it keeps failing: the computer was removed under Endpoints, its credential was rotated or replaced, or the server's database was restored from before. Add the computer again (3.5).
403	Agent log at registration or check-in	The server refuses this computer. The agent log gives the reason in its own words and a fixed hint. Reasons: license seat limit reached, token expired, enrollment cap reached, invalid token, server identity changed, a decommissioned endpoint, a license revoked, expired or not valid.	Follow the hint in the log. See 3.1.
402	The console, a script calling the API	The console is locked because the license is not in force. This is a **console** answer. Computers keep reporting; they are not stopped by it.	Fix the license (2.3). If you see 402 in an **agent** log, that is not an expected message: collect diagnostics (not verified further).
410	Agent log during enrollment, or at check-in	During enrollment, "Campaign no longer exists": the Add Computer download (a campaign) that this computer was installed with has been completed, revoked or deleted. At check-in, the computer was retired (decommissioned) on the server.	Make a new **Add Computer** download and run its installer; for a retired computer, reactivate it under **Endpoints** first.

Code	Where you see it	Meaning	What to do
426	Agent log	The agent is older than the server's minimum agent version, or the server's connection security does not allow its connection.	Update the agent with a new Add Computer download.

3.4 A token or download has expired

An **Add Computer** download carries a token that lasts **24 hours** and can enroll a limited number of computers. When it has expired or its limit is reached, the agent log says so. Make a new download: **Agent Deployment**, **Add Computer**, then run the new installer on the computer. A computer that is already enrolled keeps its identity and uses no second seat. Delete old copies of the zip.

3.5 Credential rotation failed

Where. Agent Deployment, **Endpoint Credential Rotation**. Each computer has a rotation state:

State	Meaning	Action
Pending	Waiting: the computer is offline, below agent 1.7.20, or has no credential of its own yet	Wait; agent 1.8.2 supports rotation.
In Progress	The new credential was delivered; waiting for the next check-in	Wait a few minutes.
Completed	Verified: the new credential works, the previous one is refused	None. If the row adds "awaiting first check-in" and ten minutes pass, check that computer.
Failed	A terminal problem	Click **Retry** (or **Retry all failed/expired**).
Expired	The window closed before the computer came back	**Retry**, with a fresh window.
Skipped	Already in another rotation	None.

Steps.

1. Open **View** on the operation and read the reason for the failed computer.
2. If it is offline, bring it back, then **Retry**.
3. If the computer cannot authenticate at all (the agent keeps its previous credential for 30 days after a completed rotation; beyond that, or after a snapshot revert or restored database, nothing can reach it): uninstall the agent from the computer and add it again with a new **Add Computer** download.
4. A computer with the extension may show a short gap in prompt monitoring after a rotation. It recovers within a few minutes by itself.

How to confirm. The row reads **Completed** with "new credential verified, previous credential rejected", and the computer stays **Online**.

---

4. The browser extension

The extension (AI Prompt Monitor 0.5.5) works in **Chrome and Edge**. Firefox is not supported. Open its **Options** page: it shows a **Connection** card with the headline, the Product, Server, Transport, Policy, Ruleset and Last sync. **Copy status for support** at the bottom copies a text with versions and counts only (no prompts, findings or passwords).

4.1 The extension is not installed on a computer

Causes. In this release nothing installs the extension into Chrome or Edge by itself. The console setting **Remote Browser Extension Wiring** is off by default and, on a YuviRed Lite server, has no effect: the server does not serve a signed extension package, so it answers "disabled" to every computer. Store installs are planned and are not available until the store listings exist. Firefox is not supported.

Steps.

1. On a computer added with Add Computer, a person with local administrator rights loads the extension from `C:\Program Files\YuviRed Agent\extension` with **Load unpacked** (`chrome://extensions` or `edge://extensions`, Developer mode on). Use that folder, not a copy: a copy never updates.
2. Each person ticks the consent box on the options page, unless your administrator records it by browser policy.

Confirm. The extension appears in the browser's list and the options page shows a **Connection** card.

4.2 "No server" or "Not connected"

Headline or row	Meaning	Fix
Not connected. Consent and save to start	Consent is not given or not saved	Tick "I consent to local inspection of prompts and attachments on the monitored sites above", **Save**.
Waiting for the first policy sync	Consent given; no policy pulled yet	Wait a minute; **Save** again.
Not synchronized, Last sync "failed: ..."	The last pull failed; the last known good policy stays active	Read the reason below.
Server: **not set**	No server address and no local agent found	Install the agent on the computer, or enter the server's HTTPS address in the options.

Reasons shown after "failed:": it could not reach the server (check the network, and that the YuviRed agent is still installed, because on Lite the agent carries the browser's requests; or, if Transport says direct HTTPS, that this browser trusts the server's certificate); the server rejected the request (the computer is no longer enrolled, or its credential changed); the server returned an error; or the answer could not be read.

Steps. Confirm the agent service (`SyntinelRedAgent`) runs on that computer and the computer is **Online** in the console (3.2). If the computer was removed from Endpoints or its credential rotated, add it again (3.5). Then **Save** on the options page.

Confirm. The headline reads "Connected to YuviRed Lite" and Last sync reads "OK".

4.3 Policy: blocked, or nothing happens

- **A person was blocked or warned.** Policy is set in the console under **Configuration**, **Prompt Protection**. A stronger policy than **Monitor Only** warns or blocks. Read the policy and, if needed, change it there. A first install starts in **Monitor Only**, which shows nothing on the page.
- **Nothing is blocked but you expected it.** The extension only monitors when the license is not in force (expired, revoked, not valid, or the server not activated); **no block is applied then**. Fix the license (2.3). The options page's Policy row shows what the browser applies. "Company policy not received: the answer your administrator set for

this case applies" means the policy could not be fetched, and the "if the policy cannot be fetched" answer is used.
- **A page's first send is held a moment.** Until a page learns its policy (up to 5 seconds), a send is held. This is expected.

Confirm. The Policy row on the options page shows the preset and revision you set.

---

5. Notifications

Notifications go out for **new** findings at or above **Notify me about**, subject to the quiet window and the hourly cap. The **Test** buttons ignore those brakes.

Where. Configuration, the **Notifications** chip, and the panel **Notifications, Test Send** below.

Symptom	Likely causes	Steps
Test says "Channel 'email' is not configured yet."	Host and recipient (email) or the URL (Slack, Teams, webhook) are empty or unsaved	Fill them in and click **Save Changes** first.
Email test says "failed: ..."	Wrong mail server, port, login or TLS setting; the server cannot reach the mail server; the mail server refuses	The reason is the mail library's own text. Check **Mail server**, **Mail server port** (587 is preset), **Mail username**, **Mail password**, **Encrypt mail connection (TLS)**, **Send from address**, **Send alerts to**. From the server: `Test-NetConnection mail.northwind.example -Port 587`. Check the mail server allows relay from this server.
Slack or Teams test fails	The webhook address was revoked, mistyped or pasted with spaces; outbound blocked	Create a new incoming webhook in Slack or Teams and paste it into **Slack or Teams webhook URL**. Check the server can reach the internet on 443 (or goes through your proxy).
Webhook test fails	The receiving system refuses, a certificate it does not trust, wrong address	Check the URL and the receiver's log.
Test works but real alerts never arrive	**Notify me about** is above the alert's severity; the quiet window; the hourly cap (**Max notifications per hour**); **Don't repeat the same alert for** hides repeats	Lower **Notify me about**, raise the cap, and check that the alert is new.
It worked and stopped	A password changed, a webhook was rotated, the mail server moved	Re-run **Test**.

How to confirm. The channel's tag reads `configured`, **Test** says `sent`, and the message arrives ("... this is a test notification from the Configuration page...").

---

6. Reports

Symptom	Likely causes	Steps
The report is empty or a section says "Not included: your role cannot view this"	You lack the permission for that area; or there was no activity in the period; or retention removed the data (raw telemetry is kept 30 days by default)	Ask an administrator to generate it, check the period, and check the retention settings (Administration, section 7).
Numbers look low	Computers were offline or not enrolled; the agent has no offline queue (a failed check-in is lost)	Check **Endpoints** for the period.
It takes a long time	A long period (the longest allowed is 400 days) and a large database; the server runs a limited number of reports at once (the exact limits were not verified here)	Choose a shorter period; wait for another report to finish.
"You already have reports being generated. Wait for one to finish."	Your own queue is full	Wait.
Status **failed**	The service restarted while it ran; or an internal error (a short reference is shown)	Generate again. If an internal error repeats, send the reference to support.
Download says the files were deleted	Report files are kept 30 days by default and total 2 GB	**Re-run as new snapshot**.

Confirm. The report's status is complete and the file downloads.

7. Where the logs are

What	Where
Server log	`C:\ProgramData\YuviRed\logs\syntinel-red.log` (10 MB, 5 old files)
Server agent start logs	`C:\ProgramData\YuviRed\logs\agent-start.log`, `local-enroll.log`
Database log	`C:\ProgramData\YuviRed\pg\data\log\postgresql-<date>.log`
Boot status	`C:\ProgramData\YuviRed\status\boot-status.json`
Agent log (every computer, including the server's own agent)	`C:\ProgramData\SyntinelRed\logs\syntinel-agent.log` (5 files of 5 MB; the earlier name, see the top of this page)
Installer logs	`%TEMP%` and `C:\Windows\Temp`, names start `YuviRed_Lite_`
In the console	**Admin**, **System Logs**
Windows Event Log on the server	Event Viewer, **Windows Logs**, **Application**, sources `YuviRed` and `YuviRedAgent` (the agent's own messages may also appear under `SyntinelRedAgent`)
Windows Event Log on another computer	**Application**, source `SyntinelRedAgent`

Windows Event Log entries.

- Source `YuviRed`, information: "The YuviRed service has started", and "The YuviRed service has stopped" (event ID 4100) when it stopped because it was asked to.
- Source `YuviRed`, error: "YuviRed stopped at phase '<phase>': <error or authorization>". The phase is one from 1.1. This is the event to quote to support.
- Source `YuviRedAgent` (server) or `SyntinelRedAgent` (other computers): start and stop of the agent service, and the agent's own warnings.
- **System** log, source **Service Control Manager**: a service failed to start, or ended unexpectedly.

The default agent log level records counts and errors, not the AI domains, program names or URLs it saw.

Collecting diagnostics

When something goes wrong and you write to support@yuvired.com, the fastest answer comes from a few facts and, for an install or service problem, one ZIP file from the support collector.

What YuviRed never does. The product has no crash reporter, no usage analytics and no support upload. Nothing leaves your computer unless you send it. The collector writes a file on the administrator's Desktop. You look through it and decide whether to send it.

---

1. Safe checks you can make first

These only read. They change nothing. Run them in **Windows PowerShell** as administrator on the server.

Check	Command or place	Healthy result
Services	`Get-Service YuviRed, YuviRedPostgres, YuviRedAgent`	All three `Running`. (On another computer the agent service is `SyntinelRedAgent`, shown as **YuviRed Endpoint Agent**: `Get-Service SyntinelRedAgent`. The service name carries the product's earlier name; it will change in a later release.)
Service state	`& "C:\Program Files\YuviRed\runtime\python.exe" "C:\Program Files\YuviRed\runtime\yuvired_service.py" status`	Prints `RUNNING` (or `STOPPED`, `START_PENDING`, `STOP_PENDING`). It does not print the boot phase: read the next row for that.
Boot status file	`Get-Content C:\ProgramData\YuviRed\status\boot-status.json`	`"phase": "YuviRed Ready"`. It carries no secrets. A phase of `Error` names the failed step.
Health	`Invoke-WebRequest https://127.0.0.1:8443/api/health -SkipCertificateCheck` (PowerShell 7)	Status code 200. In Windows PowerShell 5, open `https://127.0.0.1:8443/api/health` in a browser instead.
Ports	`Get-NetTCPConnection -LocalPort 8443, 5440 -State Listen`	Both listening
Network mode	`& "C:\Program Files\YuviRed\runtime\python.exe" "C:\Program Files\YuviRed\runtime\yuvired_service.py" network --status`	Shows the bind address, the firewall rule and whether it applies
Disk space	`Get-PSDrive C`	Free space is not close to zero
Time	`w32tm /query /status`	A recent "Last Successful Sync Time"

Write down what you see. A line from the boot status file, or the phase from `status`, tells support more than "it does not work".

2. Where the logs are

Log	Where	Holds
Server application log	`C:\ProgramData\YuviRed\logs\syntinel-red.log`	The server's own messages. Rotates: 10 MB, 5 old files.
Service start logs	`C:\ProgramData\YuviRed\logs\agent-start.log` and `local-enroll.log`	How the server's own agent started and enrolled.
Database log	`C:\ProgramData\YuviRed\pg\data\log\postgresql-<date>.log`	One file per day.
Boot status	`C:\ProgramData\YuviRed\status\boot-status.json`	The boot phase and any error, no secrets.
Windows Event Log	Event Viewer, **Windows Logs**, **Application**, sources `YuviRed` and `YuviRedAgent` on the server, `SyntinelRedAgent` on another computer; and **System**, source **Service Control Manager**	Start and stop of the service. An error "YuviRed stopped at phase ..." names the phase it stopped in.
Server agent log	`C:\ProgramData\SyntinelRed\logs\syntinel-agent.log`	The server's own agent (the folder carries the earlier name, see section 1).
Agent log on another computer	`C:\ProgramData\SyntinelRed\logs\syntinel-agent.log`	Check-ins, errors, refused registrations. Up to 5 files of 5 MB.
Installer logs	`%TEMP%` of the administrator who ran the installer, and `C:\Windows\Temp`; names start `YuviRed_Lite_`	What happened during an install.
In the console	**Admin**, **System Logs**	Live logs and the status of features, for administrators.

3. Run the support collector

3.1 What it is

`collect-install-evidence.ps1` is a script installed with the server. It **only reads**: it starts, stops, deletes and repairs nothing. It makes one ZIP file.

It is made for a failed install or a service that does not run, but its ZIP is also the best starting point for most server problems.

3.2 Before you run it

Do not restart the computer, uninstall, repair or clean anything first. A restart finishes pending service deletions, and cleaning deletes the registrations. Either destroys the state that explains the failure. Run the collector first.

3.3 How to run it

Prerequisites. An administrator account on the computer where the problem is.

Steps.

1. Open File Explorer at `C:\Program Files\YuviRed\support`. If the install failed so early that nothing was installed there, write to support@yuvired.com and ask for the script. The script is installed on the server only. For a problem on a computer you added with **Add Computer**, copy `collect-install-evidence.ps1` from the server to that computer (any folder, for example the Desktop) and run it there: its `10-agent.txt` part is made for that computer.

2. Right-click `collect-install-evidence.ps1`, **Run with PowerShell**. Or, in a PowerShell window:

   ```
   powershell -NoProfile -ExecutionPolicy Bypass -File "C:\Program Files\YuviRed\support\collect-install-evidence.ps1"
   ```

   If you are not an administrator, it asks for permission through the Windows prompt (UAC) and starts again.

3. Wait until it prints the path of the ZIP. By default it looks back 3 days in the Windows event logs. You can change it with `-Days 7`.

You should see: a message that names the file: `YuviRed-install-evidence-<computer>-<yyyyMMdd-HHmmss>.zip` on the Desktop of the administrator who ran it.

Options.

Option	Does
`-Days 7`	Looks back 7 days in the Windows event logs (default 3).
`-OutDir D:\Evidence`	Writes the ZIP there, not on the Desktop.
`-IncludeRawEventLogs`	Also exports the raw Windows event logs. **Do not use unless support asks.** They cover every program on the computer and cannot be redacted.

How to confirm. The ZIP exists, and opens, and `README.txt` inside lists what was collected and any file that was removed because it could not be made safe. If the script cannot make a file safe and cannot remove it, it makes no ZIP at all.

Common problems.

Problem	What to do
"running scripts is disabled"	Use the full command above, which allows this one script.
No ZIP on the Desktop	You may have run it as a different administrator: look at that account's Desktop. Or the script named a file it could not make safe and stopped.
The script runs for a long time	It reads Windows logs and all installer logs. Wait a few minutes.

Rollback. Nothing was changed. Delete the ZIP when it is no longer needed.

---

4. What the ZIP holds

Every text file is passed through a cleaner before it is packed. The ZIP holds:

File or folder	Contents
`00-summary.txt`	Windows version and build, language, last restart time, the Visual C++ runtime, antivirus and security software present, pending restarts, the installer file fingerprints
`01-services.txt`	State, start type, account and failure settings of the three services, and the "marked for deletion" flag
`02-processes-network.txt`	The YuviRed and database processes, what listens on ports 8443 and 5440, and the health check answer
`03-installed-products.txt`	Installed YuviRed entries in Apps and Features, Windows Installer and the installer cache
`04-program-files.txt`	The files under `C:\Program Files\YuviRed`, the current version file
`05-acls.txt`	Folder permissions (names only) of the program and data folders
`06-programdata.txt`	The names, sizes and times of the data folder, and copies of the status files, the logs, the database logs and the update journal
`07-logs\`	The installer logs that name YuviRed, from every user's temporary folder and `C:\Windows\Temp`
`08-msi-log-extract.txt`	The lines of those logs about services, custom actions and errors
`09-eventlogs\`	A text extract of the Service Control Manager, Windows Installer, YuviRed and crash events
`10-agent.txt`	On a computer added with Add Computer: the state of the agent service `SyntinelRedAgent`, versions, file names (never the agent's state folder or its settings file) and, under `programdata\agent-logs\`, the agent's logs. On the server itself this part reports that service as absent (the server's agent is `YuviRedAgent`); the agent's logs are still copied
`README.txt`	What was collected, and what was removed

What it deliberately leaves out: the secrets file `yuvired.env`, the database password file, the HTTPS private key, the identity folder (private key and license cache), the database files, and the agent's state and settings files.

What the cleaner does: known secret formats (passwords, keys, tokens, enrollment keys, key blocks, bearer headers) are replaced. Database logs keep only the time, severity and

message header. The application log loses its SQL and parameter blocks.

---

5. What is sensitive in it, and what to remove

The cleaner only knows **known formats**. It is a best effort, not a guarantee. **Read the ZIP before you send it.**

Things the ZIP does contain that you may want to keep private:

- The **computer name**, the Windows user names that ran the installer (they appear in file paths and logs), and your server's address or DNS name.
- The names of **other security software** found on the computer.
- **Server log lines** that can mention computer names, user names and IP addresses of your endpoints and the email address of a person who signed in.
- File and folder **names** under the YuviRed folders.
- The product versions and the license state.

What to do:

1. Unzip the file into a folder.
2. Search it for anything you do not want to share (a person's name, a customer name, an address) and edit or delete that line or file. Say in your email what you removed.
3. Zip the result again, or send the original if you removed nothing.
4. Send it to **support@yuvired.com**. Do not post it in a public place. If it is too large for email, say so, and support will tell you how to send it.

Never send: the `config` folder, `yuvired.env`, the database folder, a password, an MFA or recovery code, or an unused activation code.

---

6. What to write with it

1. Your server's YuviRed version and build (console footer) and the installation ID.
2. What you did, what you expected, what happened, and when (with the time zone).
3. What you already tried.
4. Whether any computers are affected, and which Windows versions.
