



Cyber Unified Defense Platform

Administrator Guide

Every module, best practices, common workflows, troubleshooting, and backup/restore

Version 5.47.9

Updated August 30, 2026

DADavisDesigns: Cyber Defenders

Introduction

This guide is the day-to-day reference for administering Cyber Unified Defense Platform after initial setup. If you have not yet installed the platform, start with the Getting Started Guide instead. This guide assumes the server is running, at least one agent is reporting, and you have logged in successfully.

Cyber Unified Defense Platform has been extensively developed and tested in the DADavisDesigns: Cyber Defenders lab across real, active network environments. That said, no platform is bug-free on day one in a customer's hands — every network is different, and we expect some adjustments as real-world usage surfaces edge cases our lab testing didn't. Your feedback during this trial period directly shapes those fixes and priorities. If something feels off, incomplete, or could work better, we want to hear about it.

Module Reference

Dashboard & Security Score

The home dashboard summarizes overall network health: device count, Security Score (A-F grade), and live module status tiles for scanning, DNS, traffic, and DHCP monitoring. The Security Score reflects open findings across every module — resolving alerts and approving legitimate devices raises it; new unresolved findings lower it.

Device Discovery & Baseline

Scheduled scans discover devices via ARP/ping sweep across every subnet the server has a network interface on — not just its primary network. If the server is given a second, VLAN-tagged network connection (for example, to reach a segmented IoT or guest VLAN), discovery automatically extends to that subnet as well, with no additional configuration required. New devices are listed under New Device Alerts.

Each device should be reviewed and either approved (added to the trusted baseline) or flagged for investigation. Once approved, a device's MAC, IP, and hostname become the baseline that DHCP Monitoring and Traffic Monitor compare future activity against.

DNS Monitoring

Logs DNS queries across the network and flags suspicious lookups (newly registered domains, known bad infrastructure, DNS tunneling patterns). Top Queried Domains and Recent Queries give quick visibility into what endpoints are resolving. A whitelist is available for known-noisy internal or vendor domains so they do not clutter the suspicious-query view.

Traffic Monitor

Shows outbound/inbound connections per device, with connection counts and data volume. Use Investigate IP on any external endpoint to pull GeoIP and reputation context, or Flag Anomaly to manually escalate a connection for review. GeoIP lookups are cached server-side for several hours, so repeated lookups of the same IP will not trip the external lookup service's rate limit.

About GeoIP Lookups — Third-Party Service, Not Part of Cyber Unified Defense Platform

IP geolocation data shown under Investigate IP is provided by ipwhois.io (ipwho.is), an independent, third-party service. It is not owned, operated, or warranted by DADavisDesigns: Cyber Defenders, and Cyber Unified Defense Platform has no affiliation with it beyond using its public API.

Cyber Unified Defense Platform uses the free tier by default: 1,000 requests/day, no API key required, standard location accuracy, and no uptime guarantee — suitable for evaluation and light day-to-day use, which is why lookups are cached for several hours to stay comfortably within that limit.

For production environments with heavier investigation volume, or where guaranteed uptime and higher accuracy matter, ipwhois.io offers paid plans directly (Cyber Unified Defense Platform does not resell or mark these up): Basic — \$12.99/month: 500,000 requests/month, enhanced accuracy, 99.9% uptime SLA. Business — \$39.99/month: 3,000,000 requests/month, adds security/threat data and bulk lookups. Unlimited — \$99.99/month: unlimited requests, priority routing, 99.99% uptime SLA.

Pricing and current terms: <https://ipwhois.io/pricing>. Confirm figures there before purchasing, as third-party pricing can change independently of this document.

EDR Console

Endpoint detection and response data reported by the agent, across three tabs:

- Processes — real-time process activity per endpoint.
- Internet Activity — outbound connections at the process level, with risk flagging. Well-known safe ports (80/443/53) are never flagged high-risk regardless of the underlying process signal, to reduce noise.
- Process/File Watchlist — configure specific process names or file hashes to alert on if seen anywhere on the network.

Email Security Monitor

Two tabs share the same underlying dataset:

- Behavioral Events — endpoint-side detections (suspicious mail-client-to-child-process activity such as Outlook spawning a script interpreter, SMTP flood, suspicious temp-file drops) plus server-side message analysis: Business Email Compromise patterns (payment redirection requests, gift-card fraud requests, executive impersonation with urgency/secrecy language), risky/double-extension and password-protected attachments, Reply-To domain redirection, generic-greeting and poor-grammar templates, direct requests for personal/financial information (SSN, card numbers, account credentials), authentication-failure spikes by sender domain, and mailbox auto-forward/rule tampering (see below).
- Phishing Analysis — per-message scoring across configured mail providers (Microsoft 365, Google Workspace, Exchange, IMAP, personal OAuth), with severity from Low through Critical based on sender authentication (SPF/DKIM/DMARC), lookalike-domain detection, and content indicators.

The Total Emails KPI on either tab opens one shared, paginated listing of everything received in the current window (including clean messages) — there is only one such listing, not a separate copy per tab.

Mailbox Auto-Forward / Rule Tampering Detection

A classic Business Email Compromise persistence technique: once an attacker gains access to a mailbox, they add a rule that silently forwards (or forwards and deletes/marks-as-read) mail matching finance or credential-related criteria, so the real owner never sees it. Cyber Unified Defense Platform checks each monitored mailbox's rules/filters every poll cycle and compares them against a per-mailbox baseline — only rules that

are genuinely new since the last check generate an event, so existing legitimate rules do not cause a flood of alerts on rollout.

- A new forwarding rule alone is flagged High severity.
- A new forwarding rule combined with auto-delete or mark-as-read (hiding the forwarded mail from the owner) is flagged Critical — this combination is the strongest signal of active account compromise.
- Supported for Microsoft 365 (Graph API) and Google Workspace (Gmail API) connections. Not available for Exchange/IMAP or personal-OAuth connections — mailbox rules are not readable over plain IMAP.

Setup Required — Additional Permission Needed Per Provider

This detection needs read access to mailbox rules, which is separate from the message-reading access already configured for phishing scanning. Nothing is captured or usable until the appropriate permission below is granted; until then, the platform silently skips this check for that mailbox and normal message scanning is unaffected.

Microsoft 365: in the Azure AD app registration used for the M365 connection, add the MailboxSettings.Read Application permission (API permissions → Add a permission → Microsoft Graph → Application permissions), then have a Global Administrator click Grant admin consent.

Google Workspace: in the Workspace Admin Console, under Security → API Controls → Domain-wide Delegation, edit the existing service account's authorized scopes and add <https://www.googleapis.com/auth/gmail.settings.basic> alongside the existing gmail.readonly scope.

File Scanner

On-demand lookups for files, URLs, IP addresses, and hashes against known-threat sources, plus scheduled scanning support for Linux endpoints via an async queue (results populate once the agent fulfills the hash request).

File Integrity Monitoring (FIM)

Tracks changes to configured files/directories on monitored endpoints and alerts on unexpected modification, with noise suppression (a short debounce window) to avoid alert floods from legitimate, rapid successive writes.

Hash Watchlist

A list of known-bad file hashes maintained under Settings. Any hash added here with “Create Alert” enabled generates a Critical alert if matched. Hash matching currently runs against the hourly System32 integrity scan; extending it to FIM and File Scanner results directly is on the roadmap.

Trust Rules & Blacklisting

Trust Rules let an administrator mark a hash, URL/domain, IP, or process as Blacklisted or Whitelisted from a matching alert. Whitelisting suppresses future alerts for that indicator and retroactively resolves any already-open alert it matches; Blacklisting forces future detections of it to Critical severity and prevents it from being auto-suppressed, so it stays visible for review.

Important — Blacklist Does Not Block Traffic

Blacklisting an indicator does not intercept, block, or otherwise prevent the underlying network connection, file execution, or DNS resolution. Blacklisting only escalates severity and keeps the alert open for manual review — it is a detection and workflow control, not a containment control.

For domains and URLs, real enforcement is available separately via DNS Sinkholing (see below). IP-level and hash-level enforcement remain on the Roadmap.

DNS Sinkholing

Blocks malicious domains at the DNS level, per-endpoint or fleet-wide, under Watchlist → DNS Sinkholing. Unlike Blacklisting above, Sinkholing is real enforcement: any domain added here resolves to loopback (127.0.0.1) on every endpoint it is applied to, so the connection never reaches its intended destination.

The Endpoint Status column shows how many enrolled endpoints have confirmed the block is actually applied. Enforcement requires the agent to have permission to modify the local hosts file (root on Linux, Administrator on Windows) — a domain showing fewer confirmations than total endpoint count may indicate a permissions issue on the missing endpoint(s) rather than a sync delay.

Sinkholing Works Silently — No Dedicated Alert on a Blocked Attempt

DNS Sinkholing does not currently generate a notification when a device actually attempts to reach a blocked domain — it enforces the block, but does not flag the attempt as an event of its own. If a user reports being unable to reach a site, or a domain in the EDR Console's Internet Activity log shows a result of 127.0.0.1, that confirms Sinkholing successfully blocked that connection — it is not an error. Resolving to 127.0.0.1 (or 0.0.0.0) is a standard industry convention for DNS-level blocking, used by tools such as Pi-hole and enterprise DNS security products, not something unique to this platform.

DLP (Data Loss Prevention)

Monitors for sensitive data leaving the network (e.g. unauthorized data-sharing patterns) and flags matching activity for review.

Patch Compliance

Tracks OS and software patch status across endpoints so missing updates are visible from one place instead of checking each machine individually.

Endpoint Agents & Agent Updates

The Endpoint Agents page lists every registered agent with its live version, online/offline status, and basic host stats (CPU, memory, disk, uptime). New agent builds are rolled out from the Agent Update Package panel at the top of this page rather than reinstalled by hand:

- Choose File and select the agent package downloaded from the Cyber Unified Defense Platform website, then click Import & Stage.
- Once staged, online agents pick up the new version automatically on their next check-in. Use Push Update to All to prompt every agent to check immediately instead of waiting.
- Each endpoint card shows its active version and whether it is current relative to the staged package, so a rollout can be confirmed without checking machines individually.

This is a single, cross-platform workflow — one Import & Stage covers both Windows and Linux endpoints. Windows and Linux agents check in and self-update independently in the background, so no separate action is needed per platform.

Alerts

The central queue for every finding the platform generates, regardless of source module. Each alert carries a severity, a status (open, acknowledged/investigating, resolved/false positive), and links back to the originating event where applicable.

Threat Monitor

A dedicated view for triaging and understanding alerts at a category level, organized into two tabs:

- Live Detection Log — a real-time feed of incoming alerts as they occur, for day-to-day monitoring.
- Classification — every alert automatically classified by type and mapped to the relevant MITRE ATT&CK technique, joined against Threat Intelligence and Email Security event data so classifications reflect real feed and event context rather than the alert's raw type field alone. Purely operational alert types (e.g. device online/offline events) are excluded, keeping this view focused on genuine security findings.

Selecting any alert opens a full-screen detail view with complete context and a Download JSON button, for exporting the raw event data to share with a team member or attach to a ticket.

Executive & Weekly Reports

Board-level narrative reports summarizing the week's security posture in plain language, with a resolution-rate donut chart and category breakdowns (Data Protection, Network Security, etc.) — intended for stakeholders who do not need the technical detail views.

Settings, Licensing & Users

Server configuration, SMTP/email notification setup, license activation and status, and administrator user management (including MFA enrollment and MFA reset) all live here.

Under Platform Controls → Users & Roles, the Reset MFA button lets an administrator clear another user's TOTP enrollment directly — see Resetting a User's MFA under Common Workflows below.

Best Practices

- Approve your device baseline in the first week. Traffic and DHCP monitoring are only as useful as the baseline they compare against — an unreviewed New Device Alerts queue means those modules are effectively blind.
- Enable MFA on every administrator account, not just the default admin. Additional accounts are created under Settings → Users.
- Review the DNS whitelist early. Internal services and common vendor domains generate recurring “suspicious” noise until whitelisted, which can bury genuine findings.
- Set a data retention policy deliberately rather than leaving defaults. Retention affects both storage growth and how far back investigations can look.
- Populate the Hash Watchlist with known-bad hashes relevant to your environment (threat intel feeds, prior incidents) rather than leaving it empty — it only helps once it has entries.

- Confirm SMTP/email notifications are configured and send a test message before relying on them — an alert that cannot email you is an alert you may not see in time.
- Treat the Weekly Executive Report as a stakeholder-facing artifact, not an internal working document — use the Alerts queue and module dashboards for day-to-day operational work.
- When adding domains to DNS Sinkholing, remember it is a silent control — pair it with periodic spot-checks of Internet Activity logs (looking for 127.0.0.1 resolutions) if you want visibility into how often a block is actually being hit.
- Know the Reset MFA path before you need it. A lost or replaced phone is common — being familiar with Platform Controls → Users & Roles → Reset MFA ahead of time avoids delay during an actual lockout.

Common Workflows

Investigating an Alert

- Open the alert from the Alerts queue (or from a “#123” link shown inline in a module’s event table).
- Review the detail panel for source module, severity, and any linked event context (e.g. the specific email, connection, or file change that triggered it).
- For network/traffic alerts, use Investigate IP to pull GeoIP and reputation context before deciding.
- Set the alert’s status — Acknowledged/Investigating while working it, then Resolved or False Positive once concluded.
- If it is a recurring false positive (e.g. a known-noisy internal service), add the relevant entry to the appropriate whitelist so it does not keep re-alerting.

Onboarding a New Endpoint

- Generate a one-command install link from Settings → License → License Management → Generate Distributable Install Link, choosing the target platform.
- Run the link on the target endpoint.
- Confirm it appears under Devices with an active status within a few minutes.
- Approve it in the device baseline once confirmed legitimate.
- Check EDR Console and Email Security (if applicable) begin showing telemetry from the new endpoint.

Generating a Report for Stakeholders

- Go to Reports → Weekly Executive Report.
- Select the reporting window if not using the current default week.
- Review the narrative summary and category breakdown for accuracy before sending — it reflects resolved/open counts at generation time.
- Export or email the report directly to the intended recipient list.

Resetting a User’s MFA (Lost or Replaced Device)

If a user loses or replaces the device they enrolled MFA on, an administrator can clear their TOTP enrollment so they can re-enroll a new authenticator at next login — without disabling the account or resetting their password.

1. Go to Settings → Platform Controls → Users & Roles.
2. Click Reset MFA (next to Open Account Management).
3. Select the affected user from the dropdown.
4. Click Reset MFA and confirm.
5. Inform the user their MFA has been cleared — on their next login, after entering their password, they will be prompted to enroll a new authenticator (scan a fresh QR code) before continuing.

Note

Resetting MFA does not disable the account, change the password, or end other active sessions. It only clears the stored TOTP secret and enrollment flag, so a new authenticator can be enrolled. If MFA is required platform-wide (Require MFA enabled under Users & Roles), the user cannot complete login until they finish re-enrollment.

Troubleshooting

Agent installed but not appearing on the dashboard

- Confirm the endpoint can reach the server's IP and port over the network — no firewall or network segmentation blocking the connection.
- Check that the agent process is actually running on the endpoint (Task Manager / ps aux), not just installed.
- Allow a few minutes — initial registration and first check-in are not always instantaneous.

Platform suddenly read-only / write actions failing

- Check Settings → License. A trial license that has fully expired puts the platform into read-only mode until renewed — the dashboard still displays data, but changes are blocked.
- An in-app banner and email alert should have preceded this by up to 7 days; if email alerts were not received, verify SMTP configuration under Settings.

Email notifications not arriving

- Send a test notification from Settings → Email Notifications and confirm SMTP shows “Connected & Authenticated.”
- Check spam/junk filtering on the receiving side, particularly for Gmail-hosted SMTP relays.
- Confirm the recipient address on file is current — a stale or typo'd address will fail silently from the platform's perspective.

GeoIP / Investigate IP returning errors

- A “Too Many Requests” response indicates the external lookup service's rate limit has been hit. Lookups are cached for several hours once successful, so this should self-resolve; heavy first-time investigation of many distinct IPs in a short window is the most common cause.

Unexpected or duplicate device entries after an agent reinstall

- This is expected if the same physical endpoint registers with a new agent ID. Merge or remove the stale duplicate entry from Devices once confirmed.

DNS Sinkholing shows fewer than total endpoints confirmed

- A domain stuck below full confirmation usually means one or more endpoints lack permission to write the local hosts file. On Linux, confirm the agent service is running as root (`ps aux | grep lg_agent` should show root in the owner column, not a regular user). On Windows, confirm the agent service is running with Administrator rights.

User locked out — lost or replaced MFA device

- Have an administrator go to Settings → Platform Controls → Users & Roles → Reset MFA, select the user, and confirm.
- The user can then log in with their password and will be prompted to enroll a new authenticator.
- See Resetting a User's MFA under Common Workflows for the full walkthrough.

Mailbox auto-forward detection never fires (M365 / Google)

- This almost always means the additional permission has not been granted yet — see Setup Required under Email Security Monitor → Mailbox Auto-Forward / Rule Tampering Detection.
- For Microsoft 365, confirm MailboxSettings.Read is listed under the app registration's API permissions with a green "Granted" status, not just "Not granted."
- For Google Workspace, confirm the gmail.settings.basic scope was saved under Domain-wide Delegation for the correct Client ID — a typo'd Client ID silently authorizes nothing.
- This check runs silently when the permission is missing — no alert or error is shown in the UI, only a skip message in the server console, so a missing permission looks identical to "no forwarding rules exist" from the dashboard alone.

Backup & Restore

Cyber Unified Defense Platform stores its data in a single SQLite database file, which makes backup and restore straightforward.

What to back up

- The database file (`lan_guardian.db`) under the server's ProgramData directory.
- The secret key file (`secret_key.txt`) in the same directory — required for existing sessions and encrypted values to remain valid.
- Server configuration and any customized settings, if maintained outside the database.

Recommended procedure

- Stop the Cyber Unified Defense Platform server and watchdog processes before copying the database file, to avoid backing up mid-write.
- Copy the database and secret key files to a dated backup folder on separate storage from the production server.
- Restart the server and watchdog processes.
- Periodically verify a backup by restoring it to a test environment rather than assuming the copy is valid.

Restoring

- Stop the server and watchdog processes on the target machine.

- Replace the database and secret key files with the backup copies.
- Restart the server and watchdog processes, then confirm login and data appear as expected.

Important

Always take a fresh backup before applying an upgrade or reinstalling the server, in addition to routine scheduled backups.