




NetscanXi

Version 10r1 - Admin & User Guide

Discover every device on your network, identify software and vulnerabilities, audit Active Directory and TLS, score risk, map compliance, and watch for change - and now look **inside your Docker hosts** with deep container & image vulnerability scanning - secured with user accounts, roles, MFA and multi-tenant data separation, all from one self-hosted dashboard.



Contents

1. Part I - Everyone

- 2. 1. What is NetscanXi?
- 3. 2. The Dashboard at a Glance
- 4. 3. Signing In & Your Account
- 5. 4. Account Security: Password & Two-Factor

6. Part II - Using NetscanXi (Operators & Admins)

- 7. 5. Running a Scan
- 8. 6. Scan Options & Profiles
- 9. 7. Reading Asset Details & Sorting/Filtering
- 10. 8. Vulnerabilities, CVEs & Mitigations
- 11. 9. Software, Compliance, Certificates & AD Tabs
- 12. 10. Docker Host Detection
- 13. 11. Change Detection & Alerts
- 14. 12. Exporting Reports (Full & Single-Host)
- 15. 13. Scheduled Scans

16. Part III - Administration

- 17. 14. Roles & Permissions
- 18. 15. User Administration
- 19. 16. Multi-Tenancy & Data Separation
- 20. 17. Activity & Audit Log
- 21. 18. Installation & First-Run Setup
- 22. 19. Configuration Reference
- 23. 20. Safety, Tips & Troubleshooting

24. Part IV - New in Version 9.1

- 25. 21. Interactive Dashboard Charts
- 26. 22. CVE Finding Lifecycle & Status
- 27. 23. Multi-Select Asset Table & Bulk Actions
- 28. 24. The Tenant Admin Role
- 29. 25. Global Search Bar
- 30. 26. Data-Retention Policies
- 31. 27. Patch-Aware Vulnerability Assessment
- 32. 28. Historical Trend Graphs
- 33. 29. Service-Desk Integrations
- 34. 30. Customisable Dashboard
- 35. 31. API Additions in 9.1

36. Part V - New in Version 10r1 (Docker)

- 37. 32. Deep Docker Scanning Overview
- 38. 33. Authenticated Docker API Inventory
- 39. 34. Containers & Images in the Host View
- 40. 35. Image Vulnerability Scanning (Trivy / Gype)
- 41. 36. CIS Docker Benchmark Audit
- 42. 37. The Docker Dashboard Panel
- 43. 38. Container Drift & Exports
- 44. 39. Step-by-Step: Running a Deep Docker Scan
- 45. 40. API & Configuration Reference (10r1)

Part I - Everyone

For all signed-in users, whatever their role.

1. What is NetscanXi?

NetscanXi is a self-hosted web application that scans your local network and tells you exactly what's on it - now with full account security and multi-team data separation.

It wraps the industry-standard `nmap` scanner in a clean dashboard and layers on risk scoring, vulnerability enrichment, Docker detection, regulatory compliance mapping, Active-Directory and TLS auditing, change tracking and alerting. Point it at your network and you get a sortable, searchable inventory of every live host - its software, open ports, OS, known vulnerabilities (with fixes) and an at-a-glance risk score.

Who it's for: home-lab owners, sysadmins, MSPs and security-minded teams who want continuous visibility of one or many networks without heavy enterprise tooling. If you can run `docker compose up`, you can run NetscanXi.

What's new in Version 9.1

-  **Passive scanning** - a one-click, listen-only mode that discovers assets *without sending any probe traffic*. It reads the ARP/neighbour cache and overhears self-announcements (mDNS, SSDP/UPnP) to harvest IP, MAC, vendor, hostname, advertised services and a low-confidence device type. No ports, service versions, OS accuracy, vulnerabilities, TLS/AD/Docker or credential findings are produced - those require an active scan.
-  **Unique 8-character asset IDs** - every discovered device is assigned a stable `[A-Z0-9]` ID *bound to its MAC address* and stored in NetscanXi. The same physical device keeps the same ID across scans, so newly found vulnerabilities are attributed to the asset rather than duplicated.
-  **Day & time scheduling** - schedule scans by ticking the days of the week and entering an hour and minute, in addition to the classic "every N minutes" interval. Each schedule can run an Active or Passive scan.
-  **Editable asset Type** - correct a mis-identified device type inline. The correction is bound to the asset's ID and re-applied on every future scan.

- 📄 **Scan-type on every report** - CSV, JSON and PDF exports now clearly state whether the data came from a **Passive** or **Active** scan, and include the asset ID per device.

What's new in version 8

- 🏠 **Soft multi-tenancy** - users, scans and schedules carry a *tenant* tag; people see only their tenant's data, while admins see everything.
- 📋 **Activity & audit log** - login/logout times plus scans, exports and admin actions, all timestamped with user, tenant and client IP.
- 📄 **Single-host report export** - export CSV/JSON/PDF for just one host.
- **Sort & filter everywhere** - Software, Certificates and Compliance tabs gain filters and sortable columns.
- **Live scanning IP** - the progress area shows the exact host nmap is working on right now.
- **Refreshed UI** - a clean text banner replaces the logo, and the top-right links become action buttons (Schedule, Security, User Admin, Activity, Logout).

Version 8 builds on accounts & roles (v4), TOTP MFA (v5), credential-exposure, software inventory, compliance and UI scheduling (v6), and Cyber-Essentials mapping, AD/DC auditing, PDF reports and TLS cert tracking (v7).

2. The Dashboard at a Glance

The main screen is organised top-to-bottom:

Area	What it shows
Banner	The NetscanXi Version 10r1 text banner and tagline.
Top bar	Your username + role tag, your  tenant (if scoped), scheduler status, and the action buttons:  theme ,  Schedule ,  Security ,  User Admin (admins),  Activity (admins),  Logout .
Controls bar	Target box, profile selector, Scan , Passive Scan , Stop buttons, and Export CSV / JSON / PDF .
Scan options	Checkboxes: Ports, Running services, OS detection, Docker hosts, Vulnerabilities, Credential exposure, AD/DC audit, SSL/TLS certs.
Progress	Live bar with start time, elapsed timer, current phase, and the IP currently being scanned .
Summary cards	Hosts, Vulnerabilities, Known-Exploited (KEV), High-risk assets, Docker hosts, Critical changes.
Compliance / Top-10 / Recent Changes	Regulatory posture, the highest-priority findings, and what changed since last scan.
Tabs	Assets · Software · Compliance · Certificates · Active Directory · Scan History.

3. Signing In & Your Account

If your administrator has enabled authentication, you'll land on a sign-in page. Enter your **username** and **password**. If MFA is set up, you'll then be asked for a 6-digit code (see §4).

- **Your role** appears as a tag next to your name (viewer / operator / admin). It controls what you can do - see §14.
- **Your tenant** is shown as a  badge if you're scoped to one team's data.
- **Logout** - use the  **Logout** button in the top bar. Your login and logout times are recorded in the activity log.

Open mode: if no accounts exist and no password is set, the dashboard is open and everything runs under a single `default` tenant. Admins should secure it before exposing it beyond localhost (§18).

4. Account Security: Password & Two-Factor

Click  **Security** in the top bar to open the **Account Security** panel, where you manage your own password and two-factor authentication.

Change your password (self-service)

In the **Change your password** section, enter your **current password**, then a **new password** (at least 8 characters) and confirm it. Click **Update password**. The change takes effect immediately and is recorded in the activity log. If you get the current password wrong, the update is refused.

Locked out and can't sign in at all? You can't reach this panel without logging in — ask an administrator to reset your password for you (§15).

Two-factor authentication (TOTP)

1. Open **Security**; a QR code and secret appear.
 2. Scan it with an authenticator app (Google Authenticator, Authy, 1Password...).
 3. Enter the current 6-digit code to confirm and enable MFA.
 4. **Save your backup codes** - shown once, each usable a single time if you lose your device. You can regenerate them later (this invalidates the old set).
- At next sign-in you'll enter your password, then your authenticator code. Lost your phone? Choose **"Use a backup code instead."**
 - You can **disable MFA** yourself unless an admin has *enforced* it for your account.

Admins can require MFA per user. If required but not yet set up, you'll be prompted to enrol at login.

Part II - Using NetscanXi

Day-to-day scanning and reporting. Requires the **operator** or **admin** role (viewers can read and export, but not start scans).

5. Running a Scan

1. The **target** box auto-fills with your detected network (e.g. `192.168.1.0/24`). Edit it for a range or single host (`192.168.1.10`).
2. Pick a **profile** (§6) and tick the **scan options** you want.
3. Click **Scan**. The progress bar shows live percentage, current phase, start time, elapsed time, and the **IP currently being scanned**.
4. Click **Stop** to cleanly terminate a running scan.
5. When finished, the dashboard and tabs populate automatically.

Faster scans: NetscanXi runs a quick host-discovery sweep first, then only deep-scans hosts that are actually online - far faster than probing every address in a range.

Multi-tenant note: a scan you start is tagged with *your* tenant. Scoped users only ever see and re-run scans within their own tenant; admins see all (§16).

Active vs Passive scanning (v9)

NetscanXi offers two acquisition modes:

	Active Scan	Passive Scan
How	Sends nmap probes (SYN, service/version, OS, NSE scripts) to targets.	Sends no probe traffic. Listens only: reads the ARP/neighbour cache and overhears mDNS / SSDP-UPnP self-announcements.
Options	Full set of profile + selectable options.	None - one button, fixed output.
Finds	IP, MAC, vendor, hostname, OS+accuracy, open ports, services/versions, software, vulnerabilities/CVEs, Docker, TLS, AD, compliance, risk score.	IP, MAC, vendor, hostname, advertised services, low-confidence device type & OS hint, first/last seen, asset ID.
Does <u>not</u> find	—	Open ports, service versions, OS accuracy, vulnerabilities/CVEs, TLS/AD/Docker/credential findings, risk score (shown as n/a).
Use when	You need depth: ports, vulns, compliance.	You need a zero-footprint inventory, or to discover quiet/sensitive devices without touching them.

To run a passive scan, click **Passive Scan** in the controls bar. Passive assets appear in the Assets table with a **Passive** badge and **—** in the Ports/Vulns columns. When an *active* scan later runs against the same device (matched by MAC), its full findings attach to the same asset ID.

No options for passive: by design the passive scan has no checkboxes - it returns exactly the asset attributes a passive observer can harvest, and nothing that would require sending packets to the target.

6. Scan Options & Profiles

Profiles

Profile	What it does	Speed
Quick	Top 100 ports, light service detection, no vuln scripts.	Seconds
Standard	Top 1000 ports + OS & service detection.	Minutes
Deep	Adds vulnerability NSE scripts for CVE detection.	Slower

Selectable options

Independent of the profile, tick or untick any of these before scanning:

- **Ports** - scan ports at all (untick for discovery-only).
- **Running services** - identify software/version behind each open port.
- **OS detection** - fingerprint the operating system.
- **Docker hosts** - probe Docker/container-runtime ports and the Docker API.
- **Vulnerabilities** - run nmap's vuln scripts (slower, finds CVEs).
- **Credential exposure** - *defensive* checks for cleartext logins, exposed services and weak TLS. Never brute-forces or captures credentials.

- **AD / DC audit** - identify Domain Controllers and report directory hardening posture (defensive enumeration only; no password attacks).
- **SSL/TLS certs** - inspect certificates for expiry and weak crypto.

How it works: your selections translate directly into the underlying nmap command, so you only spend time on the checks you care about.

7. Reading Asset Details & Sorting/Filtering

The asset table lists every live host. Click any **column header** to sort, use the **filter box** to search by IP, hostname, OS, port, service or vulnerability, and click any **row** to expand its full detail panel.

Column	Meaning
Asset ID	(v9) The device's stable 8-character ID, bound to its MAC address and reused across every scan.
IP / Hostname	Address and resolved name. Passively-found assets carry a Passive badge here.
Type	Guessed device type with a Docker badge when applicable. (v9) Click the  pencil to correct a mis-identified type - the correction sticks to the asset on future scans.
OS	Detected operating system with confidence %.
Software	Lead application detected.
Ports	Count of open ports.
Risk	0-100 risk score (low med high). Shows n/a for passive-only assets.
Vulns	Number of vulnerabilities found (– for passive assets).

Unique asset IDs & correcting device type (v9)

Each device gets a permanent **8-character alphanumeric asset ID** the first time it's seen. The ID is keyed on the device's **MAC address** (or its IP when no MAC is visible, e.g. routed hosts), so the same physical device is recognised on every subsequent scan and its vulnerabilities are never double-counted.

If the scanner mislabels a device's **Type** (say it calls a NAS a "web server"), click the  next to the type, enter the correct value (e.g. **nas**, **printer**, **camera**, **router**, **iot**) and save. The override is stored against the asset ID and automatically re-applied whenever that device is scanned again. Operators and admins can edit type; clearing the field reverts to the scanner's guess.

The expanded panel shows an overview grid, the full vulnerability list, detected software with CPE identifiers, a per-port breakdown, raw script output, a **Re-scan** button, and per-host **Export** buttons (§12).

Sorting & filtering everywhere (v8): the Software, Certificates and Compliance tabs now have their own filter boxes and clickable sortable column headers, and the Activity log is filterable too.

8. Vulnerabilities, CVEs & Mitigations

When vulnerability scanning is enabled, each finding shows:

- A **CVE code** linked to its full NVD entry.
- A **CVSS score**, colour-coded by severity.
- A **KEV** tag if it's in CISA's Known-Exploited catalog, and an **EXPLOIT** tag if a public exploit exists.
- A **⚡ Mitigation suggestion** - CISA's official remediation for KEV vulns, or tailored hardening advice otherwise.

The **Top 10 Vulnerabilities** panel ranks findings network-wide by Known-Exploited → exploit-available → CVSS → number of affected hosts. Each entry shows the CVE id (linked to NVD), CVSS score, KEV/EXPLOIT tags, the **full CVE description** and a **🔧 suggested mitigation** - the same detail now carried into every export.

Treat findings as leads, not gospel. nmap's vulnerability scripts are heuristic; confirm against the vendor advisory before acting.

9. Software, Compliance, Certificates & AD Tabs

Software inventory

A flat, filterable, sortable list of every application/version detected across the network, with its host, service and port.

Compliance

Each host is cross-referenced against eight frameworks - **PCI DSS, GDPR, ISO 27001, SOC 2, HIPAA, Cyber Essentials, NIST CSF** and **NCSC CAF** - with pass/fail per framework and a summary of at-risk hosts. Expand a row for the failing controls and remediation. Filter by host or status.

New in Release 3: the compliance engine now also maps each finding to **NIST CSF** (Cybersecurity Framework) outcome categories - an internationally recognised standard for private-sector enterprise buyers - and to the **NCSC Cyber Assessment Framework (CAF)** objectives/principles, supporting UK public-sector and operators-of-essential-services assurance requirements. Both appear automatically in the dashboard summary, the per-host grid, the expandable control detail and the PDF report.

Certificates

With **SSL/TLS certs** enabled, this tab lists discovered certificates with subject, issuer, expiry date, days-left and a status (ok expiring expired/weak). Sort by days-left to find what to renew first.

Active Directory

With **AD / DC audit** enabled, detected Domain Controllers are listed with domain and forest names and hardening findings (e.g. SMBv1 enabled, SMB signing not required, LDAP without LDAPS).

Defensive only. The credential, AD and TLS checks never brute-force, capture or guess credentials.

10. Docker Host Detection

NetscanXi flags hosts running Docker / a container runtime by combining open-port signals (Docker API, Swarm, etcd, Kubernetes), service banners, and the `docker-version` probe. Detected Docker hosts show a  **Docker** badge, a dashboard count, and a dedicated section with engine version, ports and indicators.

New in 10r1: this network-level detection is now the *entry point* to a much deeper view. With **Deep Docker scan** enabled, NetscanXi pulls the container/image inventory from the Docker API, scans images for CVEs and runs a CIS Docker audit. See **Part V** (§32-40) for the full workflow.

⚠ **Exposed Docker API:** an open Docker API port (2375/2376/4243) raises a warning - an exposed daemon can allow full host takeover. Restrict it immediately.

11. Change Detection & Alerts

Every scan is compared to the previous one *within the same tenant*. The **Recent Changes** panel highlights new/gone devices, ports opening/closing, service or version changes, new vulnerabilities (critical if Known-Exploited), and OS changes. Notable changes can be pushed as **alerts** to Slack/Discord/generic webhooks or email (§19).

12. Exporting Reports (Full & Single-Host)

Whole-scan export

From the controls bar:

- **Export CSV** - flat asset table (includes a Docker column, plus dedicated *CVE Descriptions* and *Mitigations* columns) for spreadsheets.
- **Export JSON** - full structured data for other tools (every vulnerability carries its description and mitigation).
- **Export PDF** - a branded report: executive summary, compliance overview, host inventory, TLS certs needing attention, Domain-Controller findings, and a *Vulnerability Detail* table listing each CVE with its description and suggested mitigation.

CVE descriptions & mitigations (v8): all three export formats now include, for every finding, the full CVE description and a recommended mitigation - CISA's required action for known-exploited CVEs where available, otherwise tailored service/product hardening advice.

Scan type & asset IDs on every report (v9): CSV, JSON and PDF exports clearly state whether the data is from a **Passive** or **Active** scan - the PDF prints it prominently under the header, the CSV carries a leading `# Scan type: ...` line plus a per-row *Scan Type* column, and the JSON `scan.scan_type` field records it. Every export also includes each device's **Asset ID** so findings can be tracked to the same physical device over time.

Single-host export (v8)

Expand any host and use its **Export: CSV / JSON / PDF** buttons to download a focused report for just that host - overview, open ports/services, software and vulnerabilities (each with its CVE description and mitigation). Ideal for handing one machine's findings to its owner.

Scoped & audited: exports respect your tenant, and every export is recorded in the activity log.

13. Scheduled Scans

Click  **Schedule** to manage recurring scans from the UI (operator+). For each schedule set a name, target (blank = auto /24), profile, **scan type** (Active or Passive), option toggles (credential exposure, AD/DC, SSL), and a cadence. Enable/disable, **run now**, or delete any schedule. Schedules you create are tagged with your tenant.

Day & time scheduling (v9)

You can now schedule by **day of week and clock time** instead of (or as well as) a plain interval:

- **Run on days** - tick any combination of **Mon–Sun**.
- **At time** - enter the **hour (0–23)** and **minute (0–59)** in local time.
- **or every N minutes** - the classic interval, used only when no days are ticked.

Example: tick **Mon, Wed, Fri** and set **02:30** to run a scan at 2:30 am every Monday, Wednesday and Friday. Choose **Passive** as the scan type for a zero-footprint nightly inventory. The schedule list shows the next run time and the Active/Passive badge. A global default interval can also be set via environment variables (§19).

Part III - Administration

For users with the **admin** role: accounts, tenants, the audit log, install and configuration.

14. Roles & Permissions

Role	Can do
admin	Everything: run/cancel scans, manage schedules, manage users & tenants, view the activity log, enforce MFA. Sees data across all tenants.
operator	Run, cancel and re-scan; manage schedules; export. Scoped to their own tenant.
viewer	Read-only: dashboard, all tabs, history and exports. Cannot start scans. Scoped to their own tenant.

Safety rail: the system refuses to delete or demote the **last remaining admin**, so you can never lock yourself out.

15. User Administration

Click  **User Admin** (admins only) to open the user-management modal.

- **Add a user** - username, password (≥ 8 chars), role, **tenant**, and an optional **Require MFA** toggle.

- **Edit a username** inline (Release 2) - type a new name in the username field and press Enter / click away. Names are checked for uniqueness; renaming your own admin account keeps your session signed in.
- **Set / reset a password** (Release 2) - click  **set password** on a user's row and enter a new password (≥ 8 chars). Use this to recover a user who has **locked themselves out**; the reset applies immediately. Admins set passwords without needing the old one.
- **Change role** inline from the dropdown.
- **Change tenant** inline in the tenant field (v8).
- **MFA controls** - force MFA on/off per user; **reset MFA** if someone loses their device (they re-enrol at next login). The list shows each user's MFA state and last-login time.
- **Delete a user** with the X button (never the last admin).

Full editable access, by design-limit: admins have editable control over every account's **username** and **password**. Passwords are stored as one-way salted hashes, so they can be *set/reset* but never *viewed* - not even by an admin. This is intentional and protects users even from a compromised admin session.

Bootstrapping: on a fresh install you can create the first admin from environment variables or the CLI (§18).

16. Multi-Tenancy & Data Separation

NetscanXi supports **soft multi-tenancy**: one installation can serve multiple teams/customers while keeping their scan data separated.

How it works

- Every **user, scan and schedule** carries a **tenant** tag (default: `default`).
- A scan started by a user is tagged with that user's tenant; change-detection compares only against prior scans in the same tenant.
- **Scoped users** (viewer/operator) see only their own tenant's scans, history, exports and running-scan status.
- **Admins are unscoped** - they see and report across all tenants, and can filter the activity log by tenant.

Setting up tenants

1. Decide on tenant names (e.g. `acme` , `contoso` , `home`). They're free-text tags - just be consistent.
2. In **User Admin**, assign each user a tenant when creating them, or edit it inline. The tenant field auto-suggests existing tenants.
3. Those users' scans and schedules are then automatically separated.

"Soft" separation: all data lives in one database with tenant-scoped queries - not separate databases. It cleanly separates day-to-day visibility but is not a hard cryptographic boundary; admins can always see everything.

Existing data: scans and users created before v8 are tagged `default` automatically on upgrade. Re-tag users as needed; historical scans keep their original tag.

17. Activity & Audit Log

Click  **Activity** (admins only) to open the audit log. It records, with a timestamp, the acting user, tenant and client IP:

Event	When it's logged
 login /  logout	Successful sign-in and sign-out (login/logout times).
 login failed	A bad username/password attempt.
 scan /  re-scan	A scan or single-host re-scan is queued.
 export /  host export	A whole-scan or single-host export.
 schedule run	A schedule is run on demand.
 user added /  tenant change	User-administration actions.
 username change /  admin password reset	An admin renames a user or sets/resets their password (Release 2).

🔑 password changed / ❌ password change failed

A user changes their own password via Account Security (Release 2).

- **Filter** by free text (user, action, detail, IP) or by **tenant** from the dropdown.
- **Refresh** to pull the latest events (most recent first).

Use it for: answering "who scanned what and when", spotting failed sign-ins, and demonstrating access accountability for compliance.

18. Installation & First-Run Setup

Option A - Docker (recommended)

```
cd netscan-xi/v8
docker compose up --build
```

Open `http://localhost:5000` (or `http://<host-ip>:5000` from the LAN). Data persists in `./data`.

Linux only for real scanning. The container uses host networking plus `NET_ADMIN` / `NET_RAW` so nmap can see your physical LAN and do SYN/OS detection. On Docker Desktop for Mac/Windows it runs in a VM and won't reach your real network - use the native install there.

Option B - Native install

```
sudo apt update && sudo apt install -y nmap python3 python3-venv
cd netscan-xi/v8
chmod +x run.sh
sudo ./run.sh # http://localhost:5000
sudo ./run.sh 127.0.0.1 8080 # custom host/port
```

Why sudo ? OS detection (`-O`) and SYN scanning (`-sS`) need root for raw socket access.

Create the first admin

On a fresh install, either set the bootstrap env vars before first start:

```
NETSCAN_ADMIN_USER=admin
NETSCAN_ADMIN_PASSWORD=change-me-please
```

...or create one from the CLI:

```
python3 -m app.server seed-admin --username admin
```

Once an admin exists, manage everyone else from 👤 **User Admin**. For PDF export, ensure `reportlab` is installed (it's in `requirements.txt`).

19. Configuration Reference

All configuration is via environment variables - everything has a safe default.

Variable	Default	Purpose
NETSCAN_ADMIN_USER	<i>unset</i>	Bootstrap first admin on a fresh install.
NETSCAN_ADMIN_PASSWORD	<i>unset</i>	Password for the bootstrap admin.
NETSCAN_SECRET	random	Flask session signing key - set this for stable logins.
NETSCAN_PASSWORD	<i>unset</i>	Legacy single-password fallback (until accounts exist).
NETSCAN_MFA_ISSUER	NetScan Xi	Issuer name shown in authenticator apps.
NETSCAN_USE_NVD	0	1 to enrich CVEs via the NVD API.
NETSCAN_SCHEDULE_MINUTES	0	Global recurring scan interval (0 = off).
NETSCAN_SCHEDULE_PROFILE	standard	Profile for the global scheduled scan.
NETSCAN_SCHEDULE_TARGET	auto	Target for the global scheduled scan.
NETSCAN_WEBHOOK_URL	<i>unset</i>	Slack/Discord/generic webhook for alerts.
NETSCAN_ALERT_MIN_SEV	warn	Minimum severity to alert on.
NETSCAN_SMTP_*	<i>unset</i>	SMTP email alerting (host/port/user/pass/from/to).
NETSCAN_DB	./data/netscan.db	SQLite database path.

UI-managed schedules vs. the global schedule: the env-var schedule is a single built-in job; the  **Schedule** modal lets you create as many tenant-tagged schedules as you like.

20. Safety, Tips & Troubleshooting

Only scan networks you own or are authorised to scan. Port and vulnerability scanning other people's networks may be illegal.

- **Secure the dashboard** - create accounts (or set `NETSCAN_PASSWORD`) and a stable `NETSCAN_SECRET` before exposing it beyond localhost; put it behind a reverse proxy with TLS for remote access.
- **Enforce MFA** for admins and operators from User Admin.
- **Empty results?** Run as root / with the right container capabilities; without them OS detection and SYN scans are skipped.
- **Deep scans feel slow?** That's the vuln scripts - use Quick/Standard for fast discovery and reserve Deep (and the AD/SSL/creds toggles) for when you need them.
- **Can't see another team's scans?** That's multi-tenancy working - only admins are unscoped (\$16).
- **PDF export returns an error?** Install `reportlab` (`pip install reportlab`).

- **Theme** - your light/dark choice is saved in the browser and applied instantly.

Part IV - New in Version 9.1

The ten enhancements shipped in the 9.1 release: interactive analytics, finding lifecycle, bulk actions, a new admin role, global search, retention controls, patch-aware vulnerability assessment, trend history, service-desk integrations and a customisable dashboard.

21. Interactive Dashboard Charts

The dashboard presents four **equal-sized chart tiles**: **OS Distribution**, **Vulnerability Severity**, **Network Risk & Compliance** and **Regulatory Compliance**. Together they give you an at-a-glance breakdown of what operating systems are on your network, how findings spread across Critical / High / Medium / Low severity, your overall risk and compliance posture, and pass/fail counts per regulatory framework.

Each tile is **customisable**. Use the **type selector** in a tile's header to switch it between a **pie chart**, **horizontal bar graph** or **vertical bar graph** (the Network Risk & Compliance tile additionally offers a **historical trend line**). **Drag the :: handle** in a tile header to reorder the charts, and **drag a tile's bottom-right corner** to resize it. Your chosen types, sizes and order are **saved in your browser**, and **Reset charts** restores the defaults.

The OS and severity charts are not just pictures - they are filters. **Click a pie slice or a legend entry** to instantly filter the Assets table to just that operating system, or click a severity bar to show only assets carrying findings at that level. Click the same element again to clear the filter. All charts are drawn as **inline SVG**, so they render crisply at any size and **work fully offline** - no external chart library or internet connection required.

Combine with the table: chart clicks set the same filter the filter box uses, so you can click a slice and then refine further by typing in the Assets filter.

22. CVE Finding Lifecycle & Status

Every vulnerability can now be tracked through a lifecycle instead of simply reappearing on each scan. Open a host's detail panel, find the CVE, and click **Set status** to mark it as **Open**, **Remediated**, **Accepted Risk** or **False Positive**. You can add an **optional note** (for example a change-ticket reference or the reason a risk was accepted).

A status can be applied **per asset** or **tenant-wide** (so the same CVE is treated consistently everywhere it appears in your tenant). Findings marked **Remediated** or **False Positive** are **de-emphasised** in the host view and **dropped from the active severity chart**, so your dashboard reflects work you've actually done rather than noise. Because status is keyed to the stable asset ID, the lifecycle decision **follows the asset across future scans** - re-scanning won't reset your triage.

Accepted Risk vs Remediated: Accepted-Risk findings remain visible (they're a deliberate, documented decision) while Remediated and False-Positive findings are filtered out of the active risk picture.

23. Multi-Select Asset Table & Bulk Actions

The Assets table gains a **checkbox column** and a **select-all** checkbox in the header, so you can act on several hosts at once. As soon as one or more rows are ticked, a **bulk actions bar** appears with:

- **Export selected (CSV / JSON)** - download a focused report covering only the hosts you picked.
- **Re-scan selected** - queue a fresh scan of just those hosts.

Selection **respects the current filter**: if you've filtered the table (by typing in the filter box or clicking a chart slice), **select-all** ticks only the rows currently shown, so bulk actions apply to exactly the subset you're looking at.

24. The Tenant Admin Role

Version 9.1 introduces a new RBAC role - **tenant_admin** - that sits between **operator** and **admin**. A tenant admin has **full admin powers within their own tenant only**: they can manage that tenant's users, schedules, scans, integrations, retention policy and activity log, without needing a global administrator.

Crucially, the boundary is strict. A tenant admin **can never act across tenants** and **cannot create global admins** - they can only manage roles up to tenant_admin within their own team. This lets you delegate day-to-day administration of a team to someone you trust with that team's data, while keeping cross-tenant control reserved for global admins.

Role hierarchy (9.1): **viewer** < **operator** < **tenant_admin** < **admin**. Each level includes everything below it; only the global **admin** is unscoped across all tenants.

25. Global Search Bar

A **global search bar** now lives in the top navigation. Start typing and it performs a **live search** across your latest scan, covering:

- **Assets** - by asset ID, IP, hostname, OS, vendor or MAC address.
- **CVEs** - by CVE id and by text in the description.
- **Software** - applications and versions detected on hosts.

Results appear in a drop-down as you type; **click any result to jump straight to that asset** and open its detail. It's the fastest way to answer "where is host X?" or "which machines are affected by this CVE?" without manually sorting the table. Search is scoped to the data you're allowed to see.

26. Data-Retention Policies

The new **Admin** panel lets you control how long scan history is kept, so the database doesn't grow forever. You can:

- **Delete scans older than N days.**
- **Keep only the M most-recent scans.**
- Use either rule, or both together.
- Tick **auto-apply** so the policy runs automatically after every scan.
- Click **Save & purge now** to apply the policy immediately.

The **asset registry** (your stable asset IDs and per-asset history) and the **audit log** are **always preserved** regardless of retention settings, so you never lose accountability or device continuity. **Tenant admins** can manage and purge **only their own tenant's** scans; global admins set policy across all tenants.

Purging is permanent. Deleted scans cannot be recovered - test your N-days / M-scans values before enabling auto-apply.

27. Patch-Aware Vulnerability Assessment

A new **Patch-aware vulns** scan toggle enables version-aware vulnerability assessment. When ticked it adds the `vulners` NSE script, and the scanner now captures **patch / build / package-release tokens** (for example `4ubuntu0.5`) into a **per-host patch-level inventory** rather than relying on the distribution version alone.

Findings are tagged "**version-confirmed**" when the exact detected release proves the host is vulnerable, versus "**heuristic**" when the match is only inferred. Mitigations now **cite the precise detected release**, so advice points at the actual package on the box. The result is markedly more accurate than distro-version-only assessment, with fewer false positives for already-patched systems.

28. Historical Trend Graphs

Set the **Network Risk & Compliance** chart tile to its **Trend (line)** view to plot **historical trends** over time: your overall **network risk score** and your **compliance posture (%)**, charted across your full scan history. This turns a point-in-time snapshot into a visible direction of travel - are you getting more or less secure? Switch the same tile to a pie or bar view whenever you just want the latest values.

A **snapshot is recorded automatically after every scan**, so the trend builds up on its own as you keep scanning. Use it to demonstrate progress after a remediation push, or to spot a slow drift in posture before it becomes a problem.

29. Service-Desk Integrations

NetscanXi can now raise tickets directly in your service desk. It ships native integrations for **Jira**, **ServiceNow** and **Zendesk**. Configure them in the new  **Integrations** admin panel: enter the **base URL**, **authentication** and the target **project / table / group**, then **test the connection** before saving.

Once configured, open any CVE finding in the host detail view and click **Create ticket** to file it in your chosen system - no copy-paste. Every ticket created is **recorded for audit** in the activity log, so there's a clear trail from finding to remediation task.

30. Customisable Dashboard

You can now tailor the dashboard to how you work. The main components - **Summary**, **Interactive Charts** and **Top Vulnerabilities** - can be rearranged. Click  **Customise** to enter customise mode, where you can **drag-reorder** components, **pin** one to the top, or **hide** ones you don't use.

Within the **Interactive Charts** block, each chart tile is independently customisable - change its **type** (pie / horizontal bars / vertical bars / trend), **resize** it by dragging its corner, and **reorder** the tiles via their drag

handles (see §21).

Click ↺ **Reset layout** to restore the component defaults, or ↺ **Reset charts** for the chart tiles, at any time. Your layout **persists per user**, so each person gets the dashboard arrangement they prefer the next time they sign in.

31. API Additions in 9.1

The 9.1 features are backed by new HTTP endpoints, useful if you automate NetscanXi or build on top of it. All respect your role and tenant scope (see §14 and §24).

Endpoint	Purpose
GET /api/dashboard	Summary data for the charts: <code>os_distribution</code> , <code>severity_distribution</code> and <code>lifecycle_counts</code> (§21, §22).
POST/GET /api/cve/lifecycle	Set and read finding status (Open / Remediated / Accepted Risk / False Positive) with optional note (§22).
GET /api/trends?days=	Historical risk-score and compliance snapshots over the requested window, e.g. <code>?days=90</code> (§28).
GET /api/search?q=	Live global search across assets, CVEs and software (§25).
GET/POST /api/settings/retention	Read and update the data-retention policy; updating can purge immediately (§26).
/api/integrations*	Configure and test Jira / ServiceNow / Zendesk service-desk integrations (§29).
/api/tickets	Create and list service-desk tickets raised from CVE findings (§29).
GET/POST /api/dashboard/layout	Read and save the per-user dashboard layout (§30).

Auth & scope: these endpoints use the same session/role checks as the rest of the app; tenant-scoped users only see and modify data within their own tenant.

Part V - New in Version 10r1 (Docker)

Deep Docker container scanning & vulnerability assessment - look inside the Docker hosts NetscanXi finds.

32. Deep Docker Scanning Overview

Version 10r1 extends NetscanXi's Docker *detection* (§10) into deep, per-container and per-image visibility. When the **Deep Docker scan** option is enabled (it is on by default), for every detected Docker host NetscanXi:

- connects to the **Docker Engine API** and pulls the container & image inventory;
- scans each image for **CVEs** using a local scanner (Trivy or Gype);
- runs a **CIS Docker benchmark** misconfiguration audit;
- surfaces **registry**, **SBOM** and **orchestration** (Swarm/K8s) findings.

Graceful by design. If the Docker API isn't reachable or no image scanner is installed, you still get the original network-level detection plus a clear note about what was unavailable. Nothing here runs an exploit or changes the target - the API access is read-only.

33. Authenticated Docker API Inventory

NetscanXi talks to the Docker Engine API over plain HTTP (2375), TLS (2376), with optional client certificate), the legacy 4243 , or a local unix socket. From it you get:

- **Engine** - version, API version, storage driver, rootless vs root, live-restore, insecure registries.
- **Containers** - name, image, state, published ports, and security flags: **privileged**, **docker.sock** mount, host network/PID, added capabilities, restart policy.
- **Images** - tags, digests and size.

mTLS to 2376: set `NETSCAN_DOCKER_TLS_CA` , `NETSCAN_DOCKER_TLS_CERT` and `NETSCAN_DOCKER_TLS_KEY` to present a client certificate to TLS-protected daemons.

34. Containers & Images in the Host View

Open any Docker host's detail panel. Below the existing Docker section you'll now see:

- an **API status** line (connected / offline, endpoint, TLS);
- the **engine** summary;
- a **Containers** list - each row shows the image and risk badges (privileged, docker.sock, host-net, host-pid, +caps);
- **Image Vulnerabilities** and the **CIS audit** (next sections).

Nothing was removed. All existing host views - ports, software, vulnerabilities, compliance, certificates, AD - are unchanged; the Docker detail is simply richer.

35. Image Vulnerability Scanning (Trivy / Grype)

Each discovered image is scanned for OS-package and language-dependency CVEs using a locally-installed scanner. The host view shows a **severity strip** (Critical/High/Medium/Low) and, per image, the CVE count and the top findings as chips. Container CVEs run through the same **CISA KEV** enrichment as host findings, so known-exploited issues are flagged (🔥) and sorted first.

Install a scanner: put **Trivy** (preferred) or **Grype** on the NetscanXi host's PATH. The bundled Docker image installs Trivy automatically. Without a scanner, NetscanXi still lists the images and tells you a scanner is needed for CVEs.

36. CIS Docker Benchmark Audit

NetscanXi audits each Docker host against high-signal CIS Docker controls and reports pass / warn / fail per check, including:

Check	Why it matters
Docker API exposure	A remote API grants full control of the host.
Daemon TLS	Plain-HTTP API has no authentication.
Privileged container	Full host capabilities - a top escape risk.
docker.sock mount	A container with the socket can control the host.
Host network / PID	Removes container isolation.
Insecure registry / :latest	Supply-chain & reproducibility risks.
Rootless / live-restore	Engine hardening posture.

37. The Docker Dashboard Panel

A new  **Docker & Containers** panel appears on the dashboard whenever the latest scan found Docker hosts (it auto-hides otherwise). It shows tiles for docker hosts, running containers, images, image CVEs, critical image CVEs, privileged containers, `docker.sock` mounts and CIS failures. Like every dashboard component you can **pin**, **hide** or **drag** it.

38. Container Drift & Exports

Change detection now reports container/image drift between scans: new/removed containers, new images, a container becoming **privileged** or mounting **docker.sock** (flagged critical), and new known-exploited container CVEs.

Exports carry the new data: **CSV** gains Containers / Image CVEs / Image KEV / CIS-fail columns, **JSON** includes the full Docker block, and the **per-host PDF** gains a Docker container/image/CIS section.

39. Step-by-Step: Running a Deep Docker Scan

- 1. (Optional) install **Trivy** or **Grype** on the NetscanXi host for image CVE scanning.
- 2. On the scan bar, leave **Docker hosts** and **Deep Docker scan** ticked (both default on).
- 3. Run an **Active** scan of your range. NetscanXi detects Docker hosts, then pulls inventory, scans images and runs the CIS audit.
- 4. Review the **Docker & Containers** dashboard panel, then open a Docker host to see containers, image CVEs and the CIS audit.
- 5. Triage: start with KEV + critical image CVEs, privileged containers and docker.sock mounts; export or raise tickets as needed.

Authorisation: only scan hosts you own or are authorised to assess. Use read-only Docker API access (and TLS client certs where possible).

40. API & Configuration Reference (10r1)

Endpoint	Purpose
GET /api/docker/capabilities	Reports whether the Docker API inventory and an image scanner (Trivy/Grype) are available.
GET /api/dashboard	Now includes a <code>summary.docker_stats</code> rollup and the deep Docker block on each host.

Scan option: `docker_deep` (the “Deep Docker scan” toggle) enables/disables API inventory + image CVE + CIS audit per scan.

Environment variable	Effect
NETSCAN_DOCKER_API=0	Disable authenticated API inventory.
NETSCAN_DOCKER_IMAGE_SCAN=0	Disable image CVE scanning.
NETSCAN_DOCKER_SCANNER	<code>trivy</code> / <code>grype</code> / <code>auto</code> .
NETSCAN_DOCKER_MAX_IMAGES	Cap images scanned per host (default 25).
NETSCAN_DOCKER_TLS_CA CERT KEY	Client materials for mTLS to 2376.