



OFFICIAL DOCUMENTATION

NetscanXi

Training & Reference Guide

Version 13

Includes Operator Training, Administrator Guidelines,
Troubleshooting, and FAQ.

Operator Training

Administrator Guidelines

Account Security & MFA

Scan Profiles

Compliance Mapping

 Docker Detection

 Multi-Tenancy

 Activity Log

 Passive Scanning

 Asset IDs

 Day/Time Schedules

 Deep Docker Scanning

 Image CVE Scanning

 CIS Docker Audit

 Patching & Remediation

Troubleshooting & FAQ

Welcome

Welcome to NetscanXi, your self-hosted network intelligence and vulnerability management dashboard.

This platform is designed to give you continuous visibility into your infrastructure, mapping software, open ports, Docker environments, and regulatory compliance posture. This guide covers daily operator workflows, administrator responsibilities, troubleshooting, and frequently asked questions.

New in Version 13: Patch and Remedy - apply **operating-system updates & upgrades** (apt / dnf / yum / zypper / apk, auto-detected) over SSH to a single asset, an asset group, or all assets - or update Docker images. NetscanXi **auto-creates and updates a remediation item** for every run, and uses **per-run credentials that are never stored**. Hands-on lessons are in **Part 6**.

New in Version 10r1: deep **Docker container scanning** - authenticated Engine API inventory (containers + images), per-image **CVE scanning** (Trivy/Grype) with KEV flagging, a **CIS Docker audit**, a Docker dashboard panel and container drift detection. Hands-on lessons are in **Part 5**.

New in Version 9.1: one-click **Passive Scanning** (zero-probe asset discovery), a stable **8-character Asset ID** bound to each device's MAC, **day & time scheduling** (tick days + set a clock time), an **editable device Type** that sticks across scans, and a **Passive/Active scan-type label** on every exported report.

Contents

1. Part 1 - Operator & Viewer Training

2. 1. Securing Your Account
3. 2. Running & Managing Scans
4. 3. Reading the Dashboard & Data Tabs
5. 4. Exporting Reports

6. Part 2 - Administrator Training

7. 5. User & Role Management
8. 6. Managing Multi-Tenancy (Data Separation)
9. 7. Monitoring the Activity Log

10. Part 3 - Troubleshooting & FAQ

11. 8. Troubleshooting Guide

12. Part 4 - Version 9.1 — What's New

13. 9. Interactive Dashboard Charts
14. 10. CVE Finding Lifecycle (Status Tracking)
15. 11. Multi-Select Asset Table & Bulk Actions
16. 12. The Tenant Admin Role
17. 13. Global Search Bar
18. 14. Data-Retention Policies (Admin Panel)

19. 15. Patch-Aware Vulnerability Assessment

20. 16. Historical Trend Graphs

21. 17. Service-Desk Integrations

22. 18. Customisable Dashboard

23. Part 5 - Version 10r1 — Docker Scanning

24. 19. Deep Docker Scan & the Host View

25. 20. Image CVEs & the CIS Audit

26. 21. Hands-On Lab: Your First Deep Docker Scan

27. Part 6 - Version 13 — Patch and Remedy

28. 22. The Patch and Remedy Console

29. 23. Credentials, Safety & the Audit Trail

30. 24. Hands-On Lab: Patch an Asset's OS

Part 1 - Operator & Viewer Training

Daily operations for users assigned the **Operator** (can scan and export) or **Viewer** (read-only) roles.

1. Securing Your Account

Before diving into the network data, ensure your access is secure.

- **Changing Your Password:** Click **Security** in the top bar. Enter your current password, type a new one (minimum 8 characters), and click **Update password**.
- **Enabling Two-Factor Authentication (TOTP):** In the Security panel, scan the displayed QR code with an authenticator app (like Authy or Google Authenticator). Enter the 6-digit code to confirm.

Tip: Always save your backup codes in a secure location. They are your way back in if you lose access to your authenticator device.

2. Running & Managing Scans

Operators can initiate scans to discover assets and identify risks.

- **1 Set the Target:** The target box will automatically fill with your detected network, but you can edit this to target a specific IP range or a single host.
- **2 Select a Profile:** Choose the depth of your scan based on your current diagnostic requirements.
- **3 Execute:** Click **Scan** for an active scan, or **Passive Scan** for a zero-probe listen-only sweep. The progress bar displays live percentage, elapsed time, and the exact IP currently being scanned.

Active vs Passive scanning (v9)

Choose the right acquisition mode for the job:

Mode	What it does	When to use
------	--------------	-------------

Active	Sends nmap probes. Returns ports, services, OS, software, vulnerabilities, Docker/TLS/AD findings, compliance and a risk score. Supports profiles and the selectable options.	You need depth and risk assessment.
Passive	Sends no probe traffic. Listens to the ARP cache and mDNS/SSDP announcements to harvest IP, MAC, vendor, hostname, advertised services and a low-confidence device type. No ports, services, OS accuracy, vulnerabilities, TLS/AD/Docker/credential findings, or risk score. No options - one button.	You want a quiet, zero-footprint inventory or must avoid touching sensitive devices.

Asset continuity: a device found passively shares the same **8-character Asset ID** it will use in active scans (matched by MAC), so when you later run an active scan its vulnerabilities attach to the same asset instead of being duplicated.

Correcting a mis-identified device Type (v9)

If the scanner labels a device's **Type** incorrectly, click the  pencil beside the type in the Assets table, enter the correct value (e.g. `printer`, `nas`, `camera`, `router`, `iot`) and save. The correction is bound to the device's Asset ID and re-applied automatically on every future scan. Clear the field to revert to the scanner's guess.

Scan Profiles

Scan Profile	Functionality Overview	Expected Speed
Quick	Scans the top 100 ports with light service detection. Skips vulnerability scripts.	Seconds
Standard	Scans the top 1000 ports and includes deep OS and service detection.	Minutes
Deep	Adds nmap vulnerability (NSE) scripts to detect specific CVEs.	Slower

3. Reading the Dashboard & Data Tabs

Once a scan completes, the dashboard populates with actionable intelligence.

- **Asset Details:** Click any row in the asset table to view an expanded panel showing software, open ports, and vulnerabilities.
- **Vulnerabilities:** Findings are ranked by severity, highlighting CVSS scores and tagging vulnerabilities found in CISA's Known-Exploited (KEV) catalog. Crucially, the system provides official CISA mitigations or tailored hardening advice for every finding.
- **Compliance:** NetscanXi maps your network against **eight** major regulatory frameworks. This includes pass/fail mappings for **GDPR, PCI DSS, ISO 27001, SOC 2, HIPAA, Cyber Essentials, NIST CSF**, and the **NCSC CAF**.
- **Docker Detection:** Hosts running container runtimes will display a Docker badge. The platform probes container ports and APIs to identify engine versions.

Critical: Exposed Docker APIs (ports 2375, 2376, 4243) represent a critical risk and should be restricted immediately.

4. Exporting Reports

You can export data for external audits, ticketing, or distribution to system owners.

- **Whole-Scan Export:** Use the top control bar to download the entire network inventory as a flat CSV, a structured JSON, or a branded PDF executive summary.
- **Single-Host Export:** Expand a specific host in the asset table and click the dedicated export buttons to generate a focused report (CSV/JSON/PDF) for just that machine.
- **Scan-type & Asset ID (v9):** every export states whether the data is from a **Passive** or **Active** scan (printed on the PDF, a leading comment line + column in CSV, and the `scan_type` field in JSON) and includes each device's Asset ID.

Part 2 - Administrator Training

Strictly for users assigned the **Admin** role, which provides unscoped visibility across all data and full control over system configuration.

Scheduling scans by day & time (v9)

Open  **Schedule** to manage recurring scans. As well as the classic “every N minutes” interval, you can now schedule by **day of week and clock time**:

- **1 Run on days:** tick any combination of Mon–Sun.
- **2 At time:** enter the hour (0–23) and minute (0–59) in local time.
- **3 Scan type:** pick **Active** or **Passive** for the scheduled run.
- **4** Leave the days unticked and set **every N minutes** instead for a plain interval.

Example: tick Mon/Wed/Fri at **02:30** with scan type **Passive** for a zero-footprint inventory three nights a week. The schedule list shows the next run time and an Active/Passive badge.

5. User & Role Management

Admins have full editable control over every account to prevent lockouts and ensure security.

- **Adding Users:** Navigate to **User Admin** to create accounts. You must assign a username, password, role, and a tenant tag.
- **Resetting Passwords:** If a user is locked out, admins can click **"set password"** on their row to immediately assign a new password without needing the old one.
- **Enforcing MFA:** Admins can force Multi-Factor Authentication on or off per user, or reset a user's MFA if they lose their authentication device (forcing them to re-enrol at their next login).

6. Managing Multi-Tenancy (Data Separation)

NetscanXi supports **soft multi-tenancy**, allowing one installation to serve multiple isolated teams or customers.

- **How it works:** Every user, scan, and schedule carries a text-based  **tenant tag**.
- **Visibility Boundaries:** Scoped users (Viewers and Operators) will only ever see scans, history, and exports that share their exact tenant tag.
- **Admin Visibility:** Admins bypass these restrictions and can see data across all tenants.

7. Monitoring the Activity Log

For compliance and accountability, the system maintains a timestamped audit log. Navigate to **Activity** to view successful/failed logins, scan executions, data exports, and administrative actions (like password resets).

Accountability: Every action logs the acting user, their tenant, and their client IP address.

Part 3 - Troubleshooting & FAQ

Common issues and their resolutions.

8. Troubleshooting Guide

Issue	Solution
Scans are returning empty results or missing data.	Ensure the underlying installation has root access. If NetscanXi does not have <code>NET_ADMIN</code> or <code>NET_RAW</code> capabilities, it will skip OS detection and SYN scans.
A passive scan found few or no devices.	Passive scanning only sees devices that have recently talked on the LAN (ARP cache) or that broadcast mDNS/SSDP announcements. Give it time, ensure the container shares the host network so it can read the neighbour cache and receive multicast, then re-run. For a complete inventory, run an Active scan.
A device's Type is wrong.	Click the  pencil next to the Type in the Assets table and enter the correct value. The override is bound to the device's Asset ID and persists across future scans.
A day/time schedule didn't fire.	Times are interpreted in the server's local timezone. Confirm at least one day is ticked and the hour/minute are set; if no days are ticked the schedule falls back to the "every N minutes" interval.

Part 4 - Version 9.1 — What's New

Ten new capabilities land in Version 9.1. Each lesson below is hands-on: open the named control and follow the steps. Unless noted, these features respect your tenant scope and role.

9. Interactive Dashboard Charts

The dashboard renders four equal-sized chart tiles — **OS Distribution**, **Vulnerability Severity**, **Network Risk & Compliance** and **Regulatory Compliance** — that you can re-style, resize and rearrange. The OS and severity charts also double as live filters for the asset table.

- **1 Find the charts:** after a scan completes, look at the **Interactive Charts** block near the top of the dashboard.
- **2 Choose a chart type:** use each tile's **type selector** to switch between a **pie chart**, **horizontal bars** or **vertical bars** (the Network Risk & Compliance tile also offers a **trend line**).
- **3 Resize a tile:** drag the **bottom-right corner** of any tile to make it bigger or smaller.
- **4 Move a tile:** drag the **:: handle** in the tile header to reorder the charts.
- **5 Filter by clicking:** click a pie **slice** (e.g. *Windows*) or a severity **bar** (e.g. *Critical*) to filter the asset table to only matching assets; click again or use **Clear filter** to reset.
- **6 Reset charts:** click  **Reset charts** to restore the default types, sizes and order.

Tip: your chart types, sizes and order are saved in your browser, so the dashboard looks the way you left it next time you sign in.

10. CVE Finding Lifecycle (Status Tracking)

You can now triage individual vulnerabilities and have that decision follow the asset across future scans, instead of re-reviewing the same finding every time.

- **1 Open host detail:** click an asset row, then expand a vulnerability under its **Vulnerabilities** list.
- **2 Set status:** click **Set status** and choose one of **Open** **Remediated** **Accepted Risk** **False Positive**.
- **3 Add a note (optional):** record why — e.g. a change ticket, a compensating control, or evidence it's a false positive.
- **4 Choose the scope:** apply the status to **this asset only**, or **tenant-wide** so the same CVE is treated consistently across every host in your tenant.

Persistence: a status is bound to the asset's Asset ID, so a finding you mark **Accepted Risk** stays that way when the host reappears in the next scan — until the CVE is no longer detected or you change the status.

11. Multi-Select Asset Table & Bulk Actions

The asset table gains a checkbox column so you can act on many hosts at once.

- **1 Select assets:** tick the checkbox at the start of each row, or use the **select-all** checkbox in the header to grab every row in the current (possibly filtered) view.
- **2 Export selected:** click **Export selected** and choose **CSV** or **JSON** to download just the chosen assets.
- **3 Re-scan selected:** click **Re-scan selected** to queue a fresh scan against only those hosts.

Tip: combine with chart filters — click the *Critical* bar, press select-all, then **Re-scan selected** to immediately re-validate every high-risk host.

12. The Tenant Admin Role

Version 9.1 adds a new RBAC role, **Tenant Admin**, sitting between Operator and the global Admin.

Role hierarchy: viewer < operator < tenant_admin < admin

- **What they can do:** a Tenant Admin has full administrative control **within their own tenant** — managing users, scans, schedules, retention, and integrations for that tenant.
- **What they cannot do:** they **cannot** see other tenants' data and **cannot** create global Admins.
- **Assigning it:** a global Admin opens **User Admin**, edits the user, and sets the role to `tenant_admin`.

Least privilege: prefer Tenant Admin over global Admin for per-customer or per-team administrators — it confines blast radius to a single tenant.

13. Global Search Bar

A search box now lives in the top navigation bar for fast lookups.

- **1 Type a query:** click the search bar in the top nav and enter an IP, hostname, CVE ID, or software name.
- **2 What it searches:** assets, CVEs, and software in the **latest scan**.
- **3 Jump to the asset:** click any result to open the associated asset's detail panel directly.

Scope: results obey your tenant boundary, so scoped users only match assets, CVEs, and software within their own tenant.

14. Data-Retention Policies (Admin Panel)

The new **Admin** panel lets administrators control how long scan history is kept.

- **1 Open the panel:** go to **Admin** and find the **Data Retention** section.
- **2 Set the rules:** choose **delete scans older than N days** and/or **keep the M most-recent** scans. You can use either or both.
- **3 Auto-apply (optional):** tick **auto-apply** to enforce the policy automatically after each scan.
- **4 Purge now:** click **Save & purge now** to save the policy and immediately remove out-of-policy scans.

Preserved data: purging removes only scan results. The **asset registry** (Asset IDs, Type overrides, finding statuses) and the **audit log** are always preserved.

15. Patch-Aware Vulnerability Assessment

A new scan toggle improves accuracy by confirming findings against captured patch levels.

- **1 Enable the toggle:** before running a scan, switch on **Patch-aware vulns**. This adds the `vulners` NSE script to the scan.
- **2 Run the scan:** NetscanXi captures patch and build tokens into a per-host **patch-level list** visible in host detail.
- **3 Read the tags:** each finding is labelled **version-confirmed** (matched against the detected build) or **heuristic** (banner/inference only).

Why it matters: version-confirmed findings dramatically cut false positives — prioritise remediation on those first.

16. Historical Trend Graphs

The dashboard plots how your posture changes over time.

- **1 Open the trend chart:** set the **Network Risk & Compliance** tile's type selector to **Trend (line)** to plot **network risk score** and **compliance posture** across your full scan history.
- **2 Switch views:** change the same tile to a **pie** or **bar** view to see the latest posture values instead of the time series.
- **3 Understand the data:** a snapshot is recorded automatically **after each scan**, so the line fills in as you scan over time.

Tip: use the trend line in management reviews to show risk trending down after a remediation push.

17. Service-Desk Integrations

Open tickets for findings directly in **Jira**, **ServiceNow**, or **Zendesk**.

- **1 Configure:** go to the  **Integrations** panel, pick your service desk, and enter the connection details (URL, project/queue, credentials or API token).
- **2 Test:** click **Test** to verify the connection before saving.
- **3 Create a ticket:** in host detail, expand a CVE finding and click **Create ticket** to open a pre-filled ticket in your service desk.

Workflow: pair this with the finding lifecycle — raise a ticket, then mark the CVE **Open** with the ticket reference in the note for full traceability.

18. Customisable Dashboard

Tailor the dashboard layout to the way you work.

- **1 Enter edit mode:** click  **Customise** in the dashboard toolbar.
- **2 Rearrange:** **drag** components to reorder them, **pin** the ones you always want on top, or **hide** components you don't need.

- **3 Exit:** leave Customise mode to lock in your layout; it persists for your account.
- **4 Start over:** click ↺ **Reset layout** to return to the default arrangement at any time.
- **5 Customise the charts too:** inside the **Interactive Charts** block you can change each chart's type, resize it and reorder it independently — see §9.

Per-user: your customised layout is saved to your account and doesn't affect what other users see.

Part 5 - Version 10r1: Docker Scanning

Hands-on training for deep Docker container scanning & vulnerability assessment.

19. Deep Docker Scan & the Host View

Version 10r1 looks *inside* the Docker hosts NetscanXi finds. With the **Deep Docker scan** option enabled (default on), NetscanXi connects to each host's Docker Engine API and pulls the container and image inventory.

- **1 Run a scan:** on the scan bar keep **Docker hosts** and **Deep Docker scan** ticked, then run an Active scan.
- **2 Open a Docker host:** its detail panel now shows API status, engine info and a **Containers** list with risk badges (privileged, docker.sock, host-net, host-pid, +caps).
- **3 Dashboard:** the  **Docker & Containers** panel summarises hosts, containers, images, image CVEs, privileged containers, docker.sock mounts and CIS failures.

Read-only: NetscanXi only reads from the Docker API - it never starts/stops containers or changes configuration.

20. Image CVEs & the CIS Audit

- **1 Image CVEs:** if Trivy or Gype is installed, each image is scanned for CVEs. The host view shows a severity strip and per-image CVE chips; **known-exploited (KEV)** CVEs are flagged  and sorted first.
- **2 CIS audit:** review the pass/warn/fail checks - exposed API, no TLS, privileged containers, docker.sock mounts, host network/PID, insecure registry, :latest tags. Fix the **fail** items first.
- **3 Exports & drift:** CSV/JSON/PDF now carry container & image data, and change detection flags new containers, new-privileged / docker.sock changes and new KEV container CVEs.

21. Hands-On Lab: Your First Deep Docker Scan

1. (Optional) install **Trivy** on the NetscanXi host so images get CVE results (the bundled Docker image already includes it).
2. Run an Active scan of a range that includes a Docker host.
3. Open the  **Docker & Containers** dashboard panel and note the counts.
4. Open a Docker host → review Containers, Image Vulnerabilities and the CIS audit.
5. Pick your top 3 actions: a KEV image CVE, a privileged container, and any docker.sock mount. Export a per-host PDF to share.

Authorisation: only scan hosts you are authorised to assess; use read-only Docker API access.

Part 6 - Version 13: Patch and Remedy

Version 13 lets NetscanXi **act** on what it finds. It applies **operating-system updates and upgrades** over SSH, and can also update Docker images. Patching is an **operator**-level action and is fully tenant-separated.

22. The Patch and Remedy Console

Open the **Patch and Remedy** tab. Choose a **Patch type** — **Operating system (SSH)** or **Docker images** — then a target:

- **Single asset** - patch one asset from the latest scan.
- **Asset group** - patch every asset in a tenant asset group.
- **All assets in scope** - patch every asset found (OS patching is not limited to Docker hosts).

Operating-system patching connects over SSH, auto-detects the package manager (`apt` , `dnf` , `yum` , `zypper` or `apk`) and runs an update + upgrade. Tick **Full distribution upgrade** for a heavier dist-upgrade, or **Simulate only** to run the package manager in dry-run mode so nothing is changed.

Docker patching instead pulls the latest image for each running container over the same authenticated Docker Engine API used for scanning, with an optional experimental **Recreate containers after pull**.

Preview first. Preview (dry run) lists the targeted assets and their detected OS (or images) with **no side effects** — the OS preview does not even contact SSH — and creates no remediation items.

23. Credentials, Safety & the Audit Trail

SSH and sudo credentials (for OS patching) and **registry credentials** (for Docker) are entered **per patch run**. They are used once and are **never written to the database, never logged, and never returned**. Password fields are cleared in the browser the moment the request is sent.

Every patch run is recorded two ways:

- **Remediation Tracking** - one item per targeted asset is auto-created (status *in progress*) and then updated to **resolved** or **blocked** with a result summary, assignee and timestamp.
- **Activity log** - an `os_patch` (or `docker_patch`) entry records who patched what, with counts only (e.g. `targets=3 ok=2 simulate=False`) — never any credential material.

Authorisation & change control: OS patching performs **privileged remote execution** and makes live changes to running systems. Patch only assets you are authorised to change, prefer the simulate/dry-run preview, and use a maintenance window for production hosts.

24. Hands-On Lab: Patch an Asset's OS

1. Run an Active scan of a range containing a Linux host you may change.
2. Open the **Patch and Remedy** tab; keep **Patch type = Operating system**.
3. Choose **Single asset** and pick your Linux host (it now appears whether or not it runs Docker).
4. Click **Preview (dry run)** and confirm the detected OS / package manager.
5. Enter the **SSH username/password** (and sudo password if different) — note they are not stored.
6. Tick **Simulate only** first, click **Patch now**, and review what *would* change.
7. Untick Simulate and run again to apply the updates; watch the per-asset log.
8. Open **Remediation Tracking** and confirm the item was created and updated; (admins) check the `os_patch` activity entry shows no credentials.