NETRIX™ On-Premise Installation Guide
Draft

Introduction

The NETRIX™ software is designed to host cloud-based applications developed by ADF. It should be deployed onto a powerful Ubuntu Linux server located on the same local network as the workstations hosting the ADF desktop applications.

This guide is designed to help IT administrators set up the complete platform (the ADF Desktop clients and the NETRIX™ on-premise server).
Setting up the ADF Desktop clients is required only once and can be automated by deploying the configuration scripts to the Windows workstations where the ADF Desktop applications are installed.

This guide is not intended for the end-users!

System Requirements

NETRIX™ is designed to run on the following computers:

Operating System	Minimum System Requirements
Linux Ubuntu	32GB of RAM, 4TB of free hard drive space

Recommended System Specifications
Ubuntu Linux Desktop 24.10 Intel i9 CPU 32GB of RAM 1TB PCIe NVMe SSD hard drive for Ubuntu 4TB PCIe NVMe SSD hard drive for data

Platform Components

NETRIX™ is a containerized application running on Docker with the following main containers:

- Frontend container is a web server of static files and the web app.
- Gateway container is a web server of backend API endpoints.
- Elasticsearch container is the search engine storing all the case data.
- Kibana container is the dashboard engine for the Audit Trail app.

- Zenko container is the object store for all the case files.
- Airflow containers manage the data processing.

Installation Instructions

If you have received a preconfigured NETRIX™ appliance you can jump to this section.

1. Prerequisite
    a. Make sure you have installed Ubuntu on the Linux server and make sure it is fully up-to-date. This is how ADF configures their appliances:
        i. Name: ADF-NETRIX
        ii. Computer name: netrix.adfsolutions.com
        iii. Username: adf-netrix
        iv. Password: adfnetrix$123
    b. Obtain the offline installation package from ADF. Note that this package is around 4GB.
    c. Download the most recent offline map file from this link. You can download a specific continent or the entire planet (2025-03-28-planet.mbtiles which is over 90 GB).
    d. Obtain a license file from ADF. To create the license file, the ADF support team will need the server UUID or your previous license file.
    e. Save all these files on a flash drive to connect to the Linux server.
2. Installing third party software. The scripts are available after in the installation folder/scripts.
    a. Docker (install-docker.sh)
    b. Portainer (install-portainer.sh)
    c. DNS dnsmasq (setup-dns.sh) - optional and only if a local DNS is required
    d. mDNS (install-avahi.sh) - optional to make the server discoverable at URL: https://netrix.adfsolutions.com.local
    e. SSH Server (install-ssh.sh) - optional
    f. Power button (install-shutdown.sh) - optional
3. Connect the flash drive to the Linux server.
4. Open a terminal from the flash drive.
5. Make sure to replace the text highlighted in orange with the actual value. Use the admin password when prompted.
6. Execute the installation command `sudo apt install ./NETRIX-VERSION.deb`.
7. During the installation process, you will be prompted to enter a few paths (see this chapter for more detail) and also point the installer to the license file and the offline map file.

Renewing your License

When receiving a new license file, it needs to be copied to **/opt/netrix/license** and replace the old license file. Once the new file has been copied, you need to restart NETRIX (see instructions here).

Setting up a Preconfigured Appliance

If you have received a preconfigured NETRIX™ appliance from ADF, follow these instructions to get started:

- Unbox the appliance.
- Connect a power cord, monitor, keyboard and mouse.
- Connect it with an Ethernet cable to your router.
- You can log in as user adf-netrix with password adfnetrix$123 (we recommend you change the password).
- The NETRIX™ application starts automatically in Docker when the appliance starts.
- Follow the instructions in the next section to connect to the network.

Connecting the Appliance to your Network

The NETRIX™ appliance needs to be connected to the same local area network with the workstations that will be accessing it.

For this connection to use HTTPS, a TLS certificate has been pre-installed on the appliance. In order for this certificate to be validated, a local DNS must have an address record (A) connecting netrix.adfsolutions.com to the appliance's IP address. You may have a router offering both DHCP and local DNS, but generally, most routers only offer DHCP. In order to solve this issue, it is possible to run a local DNS on the NETRIX™ appliance. Here is the procedure:

Configuring the Router with a Local DNS

- Configuring the Router:
    - Find out the IP address of the LAN port of the router (on small networks it is usually 192.168.0.1 and that is the value we will use in this example).
    - Reserve an IP address for the NETRIX™ appliance:
        - Find the DHCP server menu.
        - Add an entry in the Address Reservation
            - Enter the MAC address of the NETRIX™ appliance. Run the command `ip link show` from a Terminal to identify the MAC address.
            - Enter the IP address to reserve. It must be in the same subnet with the LAN port IP address, for example, enter 192.168.0.100.
            - There may be a button that will show the connected devices and automatically fill the MAC address and assigned IP address.
    - Configure the DNS address:
        - Find the DNS menu.
        - Enter the IP address of the NETRIX™ appliance.
- Configuring the local DNS:
    - Make sure your NETRIX™ appliance was rebooted after the DHCP configuration so it can obtain its reserved IP address.
    - On the NETRIX™ appliance, start a terminal.

- Go to the scripts folder of the installation package `cd /opt/netrix/scripts`.
- Execute the DNS setup script `bash ./setup-dns.sh` and verify that the IP address displayed is the one reserved in the DHCP.

Note that the workstations will only use that DNS after they have obtained their IP address from the newly configured router. A reboot of the workstations is likely necessary.

Configuring Each Workstation with a hosts File

If the local DNS configuration is not an option, you can configure each of the workstations with a local DNS entry so it can find the NETRIX™ appliance.
For Windows workstations, edit the file C:\Windows\System32\drivers\etc\hosts as admin. Add the following entry at the end of the hosts file and save your changes:

```
192.168.1.100 netrix.adfsolutions.com
```

Verifying the Installation

1. From one of the workstations, open a web browser and navigate to "https://netrix.adfsolutions.com"
2. The login page of the NETRIX™ app should open.
3. The admin account credentials are: admin/adfnetrix$123 (we recommend you change the password).

Checking NETRIX™ Health

The Portainer application is installed by default and running in Docker. This application is a container management platform that makes it possible to check the health of the NETRIX™ application.
To use Portainer, open a web browser on the NETRIX™ appliance or server and navigate to https://127.0.0.1:9443 and use admin, admin to log in.
From the Home page, click on the "Live connect" button and navigate to local>Containers to see the status of the various containers. Note that it is normal for some containers to not be running.

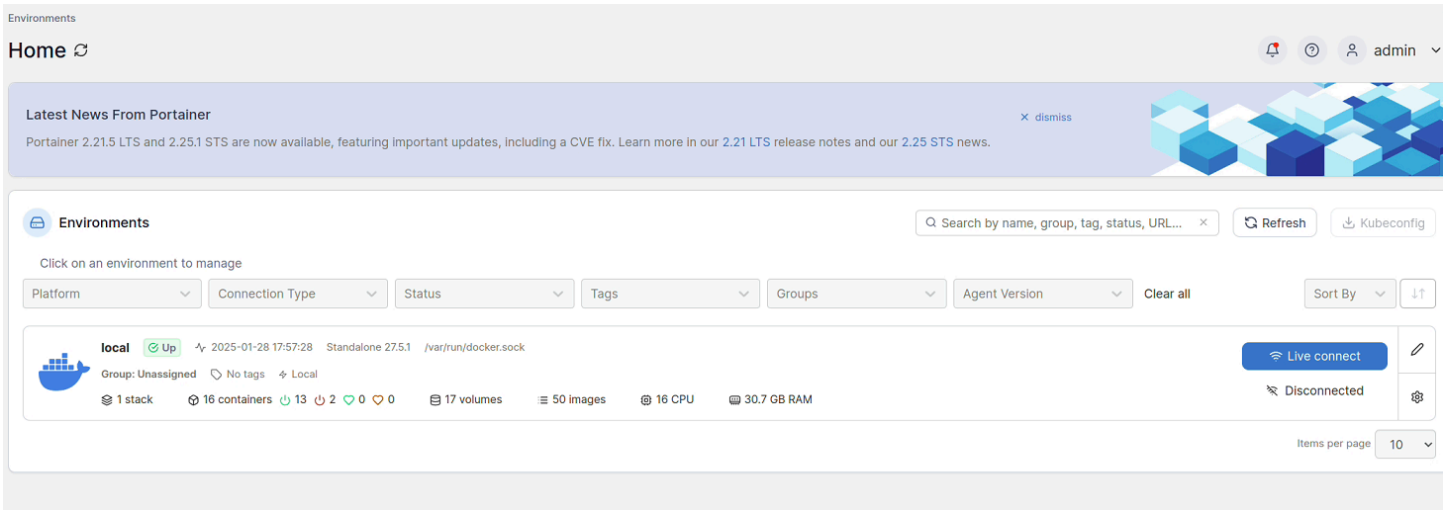

Upgrading NETRIX

In order to upgrade NETRIX to the latest version, follow these steps:

1. Make sure to backup the system data before upgrading.
    a. Refer to Configuring NETRIX™ Storage to locate the data to be backed up.
    b. It is recommended to stop NETRIX™ before backing up the data.
2. Download the latest version of NETRIX from https://www.adfsolutions.com/downloads-netrix.
3. Save the installer on a flash drive and connect it to the appliance.
4. Open a terminal and navigate to the flash drive.
5. Execute the installation command `sudo apt upgrade ./NETRIX-VERSION.deb`.
6. Select the Update option.

Configuring NETRIX™ Storage

NETRIX™ uses the following variables to define where to store its data.

Variable Name	What it does	Recommendation	Example
BLOB_PATH	This folder contains the original files, the previews, and the converted files.	This data can be stored on HDD.	/mnt/hdd1/netrix/blob
REPORTS_PATH	This folder contains the MS Word reports.	This data can be stored on HDD.	/mnt/hdd1/netrix/reports
TEMP_PATH	This folder is used to store the uploaded data containers temporarily.	This data should be stored on a fast SSD.	/mnt/ssd1/netrix/temp

ELASTIC_PATH	This folder contains the main database.	This data **must** be stored on a fast SSD.	/mnt/ssd1/netrix/elastic
POSTGRES_PATH	This folder contains a secondary database.	This data **must** be stored on a fast SSD.	/mnt/ssd1/netrix/postgres
MAPS_PATH	This folder contains the offline map data.	This data can be stored on HDD.	/mnt/hdd1/netrix/maps

These variables are set during the installation process and are saved in **/opt/netrix/docker/.env**.

Mounting Storage

Here are some instructions to make the mount points are persistent:

- Connect the additional storage device.
- Open a terminal.
- List the block devices `sudo lsblk -l`.
- Identify the drive in this list and note down the device name (we will use /dev/sdb1 as an example).
- Find the UUID of that storage device `sudo blkid /dev/sdb1` and note down the UUID.
- Create a mount point `sudo mkdir /mnt/hdd1`.
- Edit the /etc/fstab file, add this new line `UUID=<UUID> /mnt/hdd1 ext4 <options> 0 2` and save your changes.
- Mount the drive `sudo mount -a`.

Moving Storage Location

In order to move the data you have to:

- Stop NETRIX™.
- Create the new folder.
- Manually move the data from the old folder to the new one.
- Modify the path variables defined in *INSTALLATION_PATH*/docker/.env.
- Restart NETRIX™.

Sharing Log Files with Technical Support

In case you encounter an issue, the ADF Technical Support team will ask for the log files in order to troubleshoot the issue.

To download the log files from the web application:

- Log into NETRIX™
- Click on the Settings icon

- Select Download System Logs and the log timeframe between the last hour, 24 hours, or the last 7 days. Select the shortest timeframe that is the most likely to contain the error you are facing. Click DOWNLOAD when ready.
- The log archive will download to your workstation, ready to be shared with the ADF Technical Support Team.

Remoting into NETRIX™

It is possible to connect to a NETRIX terminal remotely by using SSH.
From a Windows terminal execute `ssh adf-netrix@netrix.adfsolutions.com`.

If you do not have a local DNS configured, then use the IP address of the appliance instead.

Stopping NETRIX™

If you need to stop the NETRIX™ application:

- Open a terminal.
- Stop the NETRIX™ containers `sudo systemctl stop netrix`.

Restarting NETRIX™

If you need to restart the NETRIX™ application:

- Open a terminal.
- Stop the NETRIX™ containers `sudo systemctl stop netrix`.
- Start the NETRIX™ containers `sudo systemctl start netrix`.

Shutting Down the NETRIX™ Appliance

If you need to shutdown the NETRIX™ appliance from a workstation:

- Open a terminal.
- Stop the NETRIX™ containers `sudo systemctl stop netrix`.
- Shutdown the appliance `sudo shutdown -h now`.

Repairing NETRIX™

If a critical file or container was mistakenly deleted, you can repair the NETRIX™ application by re-installing it:

- Open a terminal.
- Uninstall NETRIX™ `sudo apt reinstall adf-netrix`.

Deleting NETRIX™

This procedure permanently deletes all your data, so make sure you really want to remove everything!

If you need to delete the NETRIX™ application and all its data:

- Open a terminal.
- Uninstall NETRIX™ `sudo apt purge adf-netrix`.
- Remove any dependencies `sudo apt autoremove`.

Configuring the Desktop Applications

When using the Token Server app or the Audit Trail app, each ADF Desktop application has to be configured to properly connect to the NETRIX™ server.

To work with the Token Server App

On each workstation:

- Edit the configuration file C:\Users*NAME*\AppData\Local\ADF Solutions Inc*ADF_PRODUCT_NAME*\config.json.
- Add the “Token Server Options” with the properties below:
    - "url" is the IP address of NETRIX™ which is "https://netrix.adfsolutions.com".
    - The next two options are used in the Prepare Collection Key screen of the ADF desktop application.
        - Set the "ckLeaseDurationInDays" to a value between 1 and 10. This is the default value presented when borrowing an offline license token for a Collection Key or Authentication Key.
        - Set the "ckLeaseEditable" to true to let the user modify the lease duration.

```
{
  "Audit Trail": {
    ...
  },
  "Token Server Options": {
    "url": "https://netrix.adfsolutions.com",
    "ckLeaseDurationInDays": 1,
    "ckLeaseEditable": true
  },
  "Display Default Captures": {
    ...
  },
  ...
```

```
}
```

You can confirm that the ADF desktop application is properly connected to the token server in the Settings menu by making sure the License section shows the IPv4 address of the server. The license details displayed on screen are coming from the token server (except for the Hardware Identifier coming from the workstation).

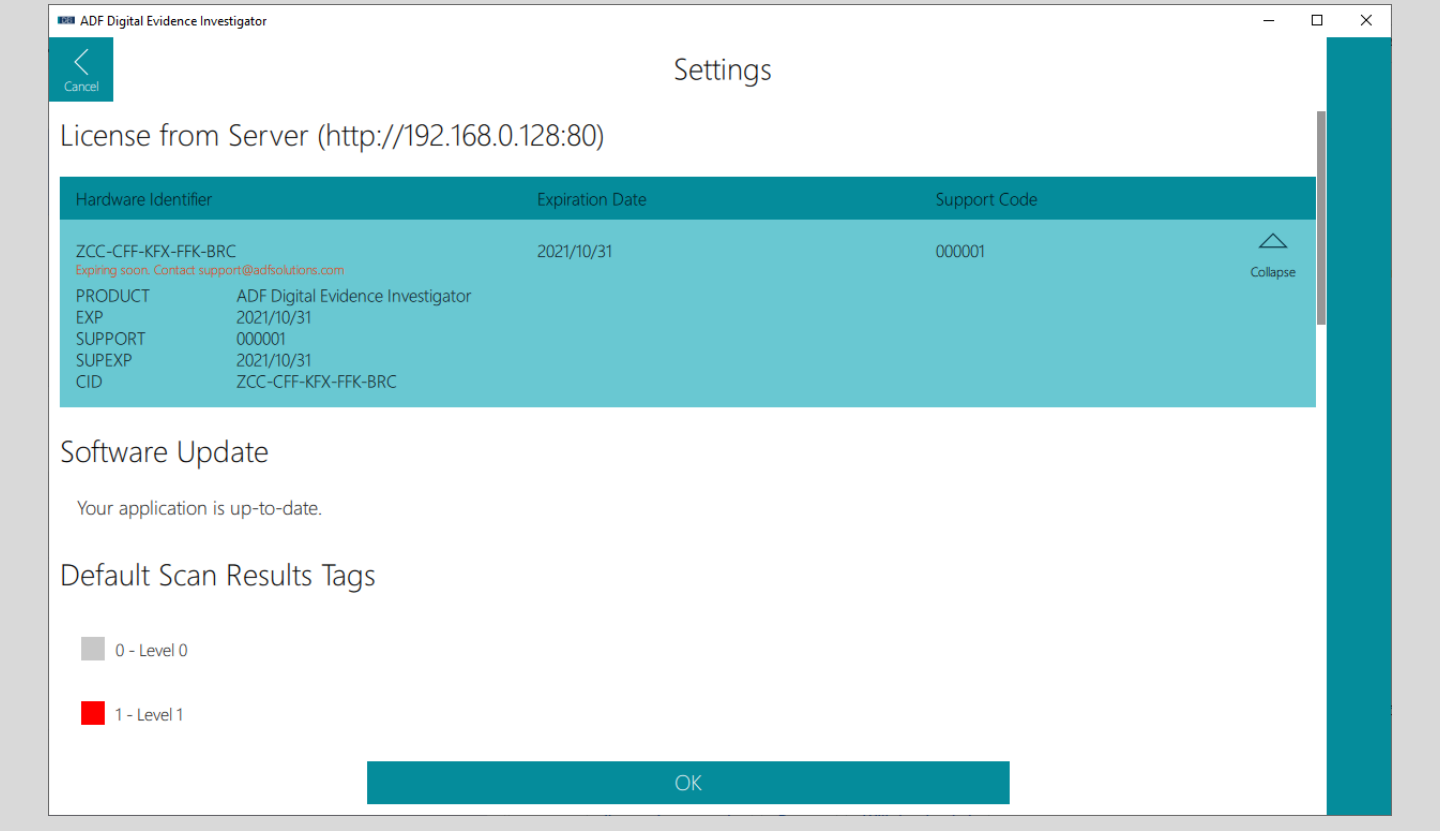

To work with the Audit Trail App

On each workstation:

- Edit the configuration file C:\Program Files\ADF Solutions Inc*ADF_PRODUCT_NAME*\filebeat\filebeat.yml.
- Edit the hosts property with the
- Add the three properties: hosts, username, password.
    - "hosts" is the IP address of NETRIX™ on port 9200 which is "https://netrix.adfsolutions.com:9200".

```
output.elasticsearch:
  hosts: ["https://netrix.adfsolutions.com:9200"]
```

Also, on each workstation:

- Edit the configuration file called config.json located at C:\Users*NAME*\AppData\Local\ADF Solutions Inc*ADF_PRODUCT_NAME*
- Set the Audit Trail enabled to true.

```
"Audit Trail": {
    "enabled": true,
    "max records": 100000
},
```

To see what data is being uploaded to the Audit Trail app, you can find the audit trail files in "C:\ProgramData\ADF Solutions Inc\v4\AuditTrail\audit_trail_date_time.n.log files".

Installing and Configuring the Uploader Service

This section details the installation and configuration of the ADF Uploader, a Windows desktop service designed to automatically upload scan results and attachments to NETRIX Case Review.

Uploader Overview

The uploader service monitors specified folders on a workstation for new data containers. When a new container is detected, the service automatically uploads the scan results and any associated attachments (such as images and PDFs) to a corresponding case in NETRIX. It's designed for easy deployment across multiple workstations, with an installer that can reuse a configuration file to ensure consistency.

Installation

The uploader is installed using an executable file (e.g., *ADF-Uploader-N.N.N-NNN.exe*). The installation process differs slightly depending on whether it's a first-time setup or a subsequent installation on another machine.

First-Time Installation

The initial installation is an interactive process that creates a configuration file which can be reused later.

1. Launch the installer executable. A setup wizard will appear.
2. **Screen 1: NETRIX Connection**
    - You will be prompted to enter the following:
        - **NETRIX URL**: The web address of your NETRIX instance (e.g., *https://netrix.adfsolutions.com*).
        - **Admin username and password**: Credentials for an admin account on the NETRIX server.
    - Click **Next**. The installer will validate the connection and credentials. If it fails, you may see an error like "NETRIX server not found at specified URL" or "Login or Password is incorrect".

3. **Screen 2: Upload Configuration**
    - Configure the paths and rules for the uploader:
        - **Case path root**: The base folder where your case folders are located. The uploader uses the subfolder names within this path as the case names in NETRIX Case Review. A browse button is available to help you select the folder.
        - **Relative scan result path**: The path within a case folder where the scan results are stored (e.g., *Reports\Standalone Viewer\ScanResults*).
        - **Relative attachment path (optional)**: The path within a case folder where attachments are located (e.g., *Notes*).
        - **Case status (optional)**: Select a default status from a dropdown list (fetched from the NETRIX server) to assign to a new case or when a data container is uploaded.
        - **Only upload folders created after (optional)**: Use the date picker to instruct the service to ignore any folders created before the selected date.
    - Click **Next**.
4. **Finalizing Installation**: The installer will authenticate to NETRIX and create a special **system user account** for the service's future interactions. It retrieves a long-lived authentication token and saves all the configured settings into a file named *uploader_config.json* in the same folder as the installer.

Subsequent Installations

For deploying the uploader on additional workstations with the same settings:

1. Copy the entire installer folder, including the *ADF-Uploader-N.N.N-NNN.exe* and the generated *uploader_config.json* file, to the new workstation.
2. Run the installer.
3. The installer will detect the *uploader_config.json* file, test the stored authentication token, and skip directly to the second configuration screen with all fields pre-populated from the file.
4. Review the settings and proceed with the installation as usual.

Configuration

The uploader's behavior is controlled by the *uploader_config.json* file. This file contains the connection details, authentication token, and folder monitoring rules.

Below is an example of the configuration structure:

```JSON
{
  "MonitoredFolder": {
    "CaseAbsolutePath": "C:/shared1/cases",
    "DataContainerRelativePath": "*/Reports/Standalone Viewer/*/ScanResults",
```

```
    "DataContainerNamePattern": "shared1/cases/.+/(.+)/Reports/Standalone
Viewer/*/ScanResults",
    "AttachmentRelativePath": "*/Notes",
  },
  "CaseStatus": "Case Ready for QA",
  "IgnorePriorToDate": "2024/01/01T00:00:00Z"
}
```

Configuration Options:

- **CaseAbsolutePath**: The absolute path to the directory containing case folders (e.g., *C:/shared1/cases*). The names of the subfolders here are used as case names in NETRIX Case Review.
- **DataContainerRelativePath**: The relative path from the case folder to the data container. Wildcards (*) can be used to replace a variable folder name.
- **DataContainerNamePattern (optional)**: A regular expression used to derive the data container's name from its path. If omitted, the name is taken from the data container itself. For example, in this folder hierarchy *C:\shared\cases\case123\exhibit2A\ScanResults\scanOfEx2A*, you should enter *shared/cases/.*/(.*)/ScanResults* for the data container to be named exhibit2A instead of scanOfEx2A. The parentheses indicate which part of the matching expression should be used
- **AttachmentRelativePath (optional)**: The relative path from the case folder to find associated attachments. The attachment will be accessible in the Attachments menu, in a subfolder named after the Data Container.
- **CaseStatus (optional)**: When a data container is uploaded, the corresponding case in NETRIX Case Review will be set to this status.
- **ignorePriorToDate (optional)**: The service will ignore any case folders with a creation date earlier than this timestamp.

Service Operation and Logging

Workflow

- The uploader runs as a Windows service, checking for new data containers every minute based on the *CaseAbsolutePath* and *DataContainerRelativePath* settings.
- When a new data container is found, the service determines if a new case needs to be created in NETRIX Case Review.
- It then uploads the scan result. If attachments are found in the attachmentRelativePath (limited to .jpg, .png, .webp, and .pdf files), they are also uploaded and placed in a subfolder named after the data container.
- If an upload fails, the service will retry after a set period.
- The service keeps track of successfully uploaded containers to prevent duplicates.

Logging

The service maintains detailed logs of its activities.

- **Log Location**: Log files are created in the service's installation directory.
- **Log Retention**: To save disk space, log files older than 30 days are automatically deleted.
- **Logged Activities**: The logs capture important events with as much detail as possible (timestamps, paths, filenames, etc.), including:
    - Detection of a new data container.
    - Creation of a new case.
    - Start, success, or failure of scan result uploads.
    - Start, success, or failure of individual attachment uploads.
    - Errors, such as failure to create a case or an invalid case status.