

Neuropacs™ Agent Usage Guide

Authored By: Kerrick Cavanaugh



Document Version History

Date Updated	Version	Notes
12/01/2025	v1.0.0	Initial release
03/02/2026	v1.0.1	Added "Quality Control" section to Settings.
03/12/2026	v1.0.2	Added Docker standalone installation method and Swagger page
04/01/2026	v1.0.3	Added Windows and virtualization support

Table of Contents

- Document Version History..... 2**
- Table of Contents..... 3**
- System Requirements..... 6**
 - Supported Environments..... 6
 - Supported Runtime Environments..... 6
 - Supported Virtualization Platforms..... 6
 - Supported Cloud Deployment Targets..... 6
 - Supported CPU Architecture..... 7
 - Software Requirements..... 7
 - Important Compatibility Notes..... 7
 - Hardware Requirements (Cloud Preprocessing)..... 7
 - Hardware Requirements (Local Preprocessing)..... 7
- Agent Application Releases..... 8**
- Installation..... 9**
 - Linux/macOS Docker Standalone Installation..... 9
 - Prerequisites..... 9
 - Installation..... 9
 - Linux/macOS Script-based Docker Installation..... 10
 - Prerequisites..... 10
 - Installation..... 10
 - Windows Installation..... 12
 - Prerequisites..... 12
 - Installation..... 12
- Application Management..... 15**
 - Linux/macOS Docker Standalone Installation..... 15
 - Linux/macOS Script-based Docker Installation..... 15
 - Windows Installation..... 16
- Data Collection..... 18**
 - Overview..... 18
 - Collected Metrics..... 18
 - Data Privacy and Security..... 19
- Usage..... 20**
 - Configuring the Agent..... 20
 - Configuring a PACS to Send Studies to the Neuropacs Agent..... 27
 - Navigating the UI..... 28
 - Navigation Menu..... 28
 - Viewing Studies..... 29

View Additional Information.....	30
Quality Control Reports.....	31
Group Assignment.....	33
Manual Uploads.....	34
Viewing Results.....	36
Group Management.....	38
Logs.....	41
Settings.....	42
Updates.....	42
API Documentation.....	43
On-Premise Quality Control (QC).....	45
DICOM-Based Studies.....	45
NIfTI-Based Studies.....	45
Security Overview.....	46
Key Security Features.....	46
Automated Credential Protection.....	46
Secure Storage & Encryption.....	46
Separation of Configuration & Credentials.....	46
Support.....	47

Neuropacs Agent Description

The Neuropacs Agent is an on-premise software service designed to securely receive, validate, and transmit medical imaging studies from a PACS to the Neuropacs Cloud platform. It supports automated DICOM routing, quality control, manual study uploads, and optional local preprocessing, while maintaining HIPAA-compliant handling and transport of all imaging data. The Neuropacs Agent also provides an intuitive, modern interface for managing studies and overseeing the Neuropacs workflow.

The agent can be configured to operate in one of two modes:

Cloud Processing

Ingested imaging studies are validated on-premises through an initial quality control check and then securely transmitted to the Neuropacs Cloud platform for full processing.

Local Processing

Imaging studies are validated and preprocessed on-premises, enabling stricter control over the data transmitted outside the host facility. Full imaging datasets remain within the local data center, while only derived non-PHI features are securely forwarded to the cloud for further analysis.

System Requirements

Supported Environments

The platform is supported in environments that can run Linux-based x86_64 / AMD64 containers either natively or within a supported Linux virtual machine. This provides flexible deployment across Linux, Windows, macOS, on-premise virtualization platforms, and certain cloud environments while maintaining a consistent Linux-based runtime.

Supported Runtime Environments

The following deployment environments are supported:

- **Native Linux hosts**
 - Modern Linux distributions with support for Docker Engine and Docker Compose
 - Ubuntu 20.04 or later is recommended
- **Windows hosts**
 - Windows 10/11 Pro, Enterprise, Education, and Server with Hyper-V/VMware virtualization capabilities
- **macOS hosts**
 - Modern macOS versions with Docker Desktop
 - macOS systems using VMware Fusion

Supported Virtualization Platforms

The platform may also be deployed inside a supported Linux virtual machine on the following platforms:

- **VMware**
 - VMware Workstation (Windows/Linux)
 - VMware Fusion (macOS)
 - VMware vSphere / ESXi
- **Hyper-V**
 - Windows Server Hyper-V
 - Windows 10/11 Pro, Enterprise, and Education with Hyper-V enabled

Supported Cloud Deployment Targets

Virtual machine images may be imported for deployment in supported cloud environments, including:

- AWS
- Microsoft Azure
- Google Cloud Platform (GCP)

Supported CPU Architecture

The platform currently supports:

- **x86_64 / AMD64**

Other CPU architectures may be compatible only if they provide a supported virtualization or emulation mechanism capable of running x86_64 Linux guest environments or containers.

Software Requirements

- Docker Engine / Docker Compose
- Docker Desktop where applicable
- Docker version 2+

Important Compatibility Notes

- The application runtime is Linux-based in all supported deployment models.
- Windows and macOS are supported as host environments only and do not run the application as native Windows or macOS workloads.
- Native Windows-only container deployments are not supported.
- Linux container execution must occur either:
 - directly on a supported Linux host, or
 - within a supported Linux virtual machine
- Custom prebuilt or converted virtual machine images may be provided for:
 - VMware-compatible environments
 - Hyper-V-compatible environments

Hardware Requirements (Cloud Preprocessing)

- CPU: 2+ vCPUs
- RAM: 4+ GB
- Storage: 50–200+ GB

Hardware Requirements (Local Preprocessing)

- Application
 - CPU: 2 vCPUs
 - RAM: 4 GB
- Processing
 - CPU: 6+ vCPUs / container
 - RAM: 8+ GB / container
- Storage: 50–200+ GB

Agent Application Releases

Release Date	Version	Notes
12/01/2025	v0.1.61	Initial production release

Installation

Note: To install the Neuropacs Agent, you must first obtain an API key and service user credentials. Contact Neuropacs support for more information.

Linux/MacOS Docker Standalone Installation

Prerequisites

Before installation, ensure the following are available:

- A supported **Ubuntu 20.04+ Linux host**
- Sufficient CPU, memory, and disk resources for the Docker containers
- Docker version 2+ installed and running
- Network settings and any site-specific configuration values required by your environment

Installation

1. Authenticate with GitHub Container Registry (required a GHCR token)

```
echo "<ghcr_token>" | [sudo] docker login ghcr.io -u neuropacman --password-stdin
```

2. Pull the Neuropacs Agent Docker image

```
[sudo] docker pull ghcr.io/neuropacs/neuropacs-agent:latest
```

3. Basic Usage

```
[sudo] docker run -it --name neuropacs-app \  
-p <UI_PORT>:3001 -p <DICOM_PORT>:104 \  
-v /var/run/docker.sock:/var/run/docker.sock \  
-v <path_or_volume_name>:/data \  
-v <path_or_volume_name>:/usr/src/app/data \  
ghcr.io/neuropacs/neuropacs-agent:latest
```

4. Options and Flags

Option	Description	Required	Default
<code>-v neuropacs-container-workspaces:/usr/src/app/processing</code>	Docker managed volume for local processing I/O. Must be passed exactly as defined in the Option column.	Yes (in local-processing mode)	N/A
<code>-v /absolute/path/to/certs:/data/certs:ro</code>	Docker bind mount for TLS certificates. Within this directory must include a cert.pem and key.pem file.	Yes (in tls-enabled mode)	N/A

5. Verify Installation

Confirm that:

- the required containers are running
- the application is reachable on the expected interface
- the environment-specific configuration has been applied successfully

Linux/MacOS Script-based Docker Installation

Prerequisites

Before installation, ensure the following are available:

- A supported **Ubuntu 20.04+ Linux host**
- Sufficient CPU, memory, and disk resources for the Docker containers
- Docker version 2+ installed and running
- Network settings and any site-specific configuration values required by your environment

Installation

1. Download the installer script

```
curl -O https://d13kmmwctl1jvu.cloudfront.net/agent/install/linux/install.sh &&
chmod +x install.sh
```

2. Basic Usage

```
./install.sh --institution-name NAME [OPTIONS]
```

3. Options and Flags

Option	Description	Required	Default
<code>--institution-name</code>	Name of host institution for registration	Yes	N/A
<code>--local</code>	Install from local source	No	No
<code>--version</code>	Application version (ex. v1.2.0)	No	LATEST
<code>--host-mount</code>	Path where application data will be stored	No	<install-dir>/volumes
<code>--use-volumes</code>	Use Docker-managed named volumes instead of bind mounts	No	No
<code>--local-processing</code>	Enable local processing and download Docker image	No	No
<code>--normalization-image IMAGE</code>	Custom normalization Docker image	No	neuropacs_normalization:latest
<code>--aidp-features-image IMAGE</code>	Custom AIDP features Docker image	No	neuropacs_aidp_features:latest
<code>--addl-features-image IMAGE</code>	Custom ADDL features Docker image	No	neuropacs_addl_features:latest
<code>--use-tls</code>	Enable HTTPS with custom TLS certificates	No	No
<code>--tls-cert PATH</code>	Absolute path to TLS certificate file	Yes (with <code>--use-tls</code>)	N/A
<code>--tls-key PATH</code>	Absolute path to TLS private key file	Yes (with <code>--use-tls</code>)	N/A
<code>--tls-ca PATH</code>	Absolute path to CA bundle file	No	N/A
<code>--install-dir</code>	Path where application will be installed	No	./neuropacs-agent
<code>--api-base-url</code>	URL of neuropacs server	No	Current production URL
<code>--help</code>	Display help information	No	N/A

4. Verify Installation

Confirm that:

- the required containers are running
- the application is reachable on the expected interface
- the environment-specific configuration has been applied successfully

Windows Installation

Prerequisites

Before installation, ensure the following are available:

- A supported **x86_64 / AMD64 Windows host**
- Sufficient CPU, memory, and disk resources for the virtual machine
- One of the following:
 - **VMware Workstation**
 - **Hyper-V**
 - **VMware vSphere/ESXi**
- The provided prebuilt virtual machine files
- Network settings and any site-specific configuration values required by your environment

Installation

1. Open or Import the Virtual Machine

Using VMware: Open your VMware environment and import or open the provided virtual machine files. Available formats include: **OVA**, **OFV**, and **VMDK**.

Using Hyper-V: Open Hyper-V Manager on the Windows host and import the provided **VHDX** file.

For local workstation deployments, use **VMware Workstation**. For managed server or enterprise virtualization environments, deploy the VM to **VMware vSphere / ESXi**.

2. Review Virtual Machine Settings

The provided virtual machine files are preconfigured with recommended resource allocation and networking configuration.

Before starting the VM, review and adjust the following as needed for your environment:

- **CPU allocation**
- **Memory allocation**
- **Disk location or datastore**
- **Network adapter / port group configuration**
- Any other VMware-specific settings required by your organization

Contact Neuropacs support for additional information on virtual machine configurations. Improper configurations can lead to unpredictable application behavior.

3. Start the Virtual Machine

Power on the virtual machine and confirm Linux guest operating system starts successfully.

4. Access the Guest System

Log into the guest operating system using the credentials and access method provided with the deployment package.

5. Apply Environment-Specific Configuration

Update any required configuration values for your environment, such as:

- hostnames
- network settings
- mounted storage paths
- application configuration values
- credentials or secrets, if applicable

To enable TLS, mount a path to your certificates and follow the installation procedure.

6. Start the Application

The virtual machine environment automatically provisions a **neuropacs** user account with the required permissions preconfigured. The user's home directory, **/home/neuropacs**, contains the following files:

1. **README.md** – comprehensive documentation covering application management and troubleshooting
2. **init.sh** – the initialization script used to begin the application setup process

From the neuropacs home directory, run:

```
./init.sh
```

Follow the on-screen prompts to configure the application. Once setup is complete, the application will start automatically.

7. Verify Installation

Confirm that:

- the virtual machine is running
- the required containers are running
- the application is reachable on the expected interface
- the environment-specific configuration has been applied successfully

Application Management

Linux/macOS Docker Standalone Installation

Management is performed via Docker commands, allowing administrators to start, stop, restart, inspect, and monitor the application using standard Docker and Docker Compose workflows.

Example usage:

```
docker start neuropacs-app # Start all services
docker stop neuropacs-app neuropacs-db # Stop all services
docker restart neuropacs-app # Restart services
docker ps -a \ # Display service status
  --filter name=neuropacs-app
  --filter name=neuropacs-db
docker logs -f neuropacs-app # View service logs
docker ps \ # Show service Docker images
  --filter name=neuropacs-app --filter name=neuropacs-db \
  --format 'table {{.Names}}\t{{.Image}}\t{{.Status}}'
docker rm -f neuropacs-app neuropacs-db && \ # Uninstall
  docker volume rm \
  neuropacs-data neuropacs-agent-data \
  neuropacs-db-data neuropacs-container-workspaces
docker run --rm neuropacs/neuropacs-agent help # Show help message
```

Linux/macOS Script-based Docker Installation

Upon successful installation, the Neuropacs Agent deploys a command-line management utility that enables streamlined administration of the application and its associated Docker services. This utility simplifies routine operational tasks such as starting, stopping, monitoring, and troubleshooting the Agent. The management utility is installed as the `neuropacs` executable. It can be run from the installation directory or from any location if added to the system `PATH`.

Example usage:

```

<install-dir>/neuropacs start                # Start all services
<install-dir>/neuropacs stop                 # Stop all services
<install-dir>/neuropacs restart              # Restart services
<install-dir>/neuropacs status               # Display service status
<install-dir>/neuropacs logs                 # View service logs
<install-dir>/neuropacs update [options]    # Update services
<install-dir>/neuropacs images              # Show service Docker images
<install-dir>/neuropacs uninstall           # Uninstall
<install-dir>/neuropacs help                # Show help message

```

Windows Installation

Management is performed via Docker commands, allowing administrators to start, stop, restart, inspect, and monitor the application using standard Docker and Docker Compose workflows.

```

docker start neuropacs-app                    # Start all services
docker stop neuropacs-app neuropacs-db       # Stop all services
docker restart neuropacs-app                 # Restart services
docker ps -a \                               # Display service status
  --filter name=neuropacs-app
  --filter name=neuropacs-db
docker logs -f neuropacs-app                 # View service logs
docker ps \                                  # Show service Docker images
  --filter name=neuropacs-app --filter name=neuropacs-db \
  --format 'table {{.Names}}\t{{.Image}}\t{{.Status}}'
docker rm -f neuropacs-app neuropacs-db && \  # Uninstall
  docker volume rm \
  neuropacs-data neuropacs-agent-data \

```

```
neuropacs-db-data neuropacs-container-workspaces
```

```
docker run --rm neuropacs/neuropacs-agent help
```

```
# Show help message
```

Data Collection

Overview

The Neuropacs Agent continuously collects operational telemetry to ensure the secure, reliable, and compliant operation of the system. These metrics are used exclusively for system health monitoring, performance optimization, fault detection, and regulatory compliance.

Collected Metrics

The following categories of operational metrics are collected by the Neuropacs Agent:

Runtime Metrics

Metric	Description
Agent ID	Logical Agent identifier (UUID)
Runtime Type	Runtime of the Agent (i.e., Docker, Singularity, Native)
Application Version	Running agent version

System Metrics

Metric	Description
CPU Core Count	Number of logical CPU cores available
CPU Load	Current system CPU utilization
Total Memory	Total physical memory available
Available Memory	Unused system memory available
Disk Free Space	Available disk storage on the working volume
Disk I/O Activity	Aggregate read/write activity indicators

Container and Service Health

Metric	Description
Container Status	Runtime state of each Neuropacs service container

Container Uptime	Duration each service has been continuously running
Service Availability	Health status of critical processing and API services

Processing and Workflow Metrics

Metric	Description
Order Status	Status of executed orders
Processing Queue	Local processing queue metrics (if applicable)

Data Privacy and Security

Metrics are strictly limited to infrastructure-level operational information and do not contain patient identifiers, clinical imaging data, diagnostic results, or host system information. All telemetry communications are encrypted in transit and authenticated.

Usage

Configuring the Agent

To configure the Neuropacs Agent to align with your custom workflow, access the integrated web portal available on port **3001**. When accessing the portal locally on the host system, navigate to:

```
http://localhost:3001/
```

If accessing the portal from another machine on the network, you must have network access to the host system and use its IP address, for example:

```
http://<host_ip_address>:3001/
```

If the installation was successful, you should now see the following screen:




Figure 1: Integrated UI Studies page

Navigate to the **Settings** tab within the user interface.

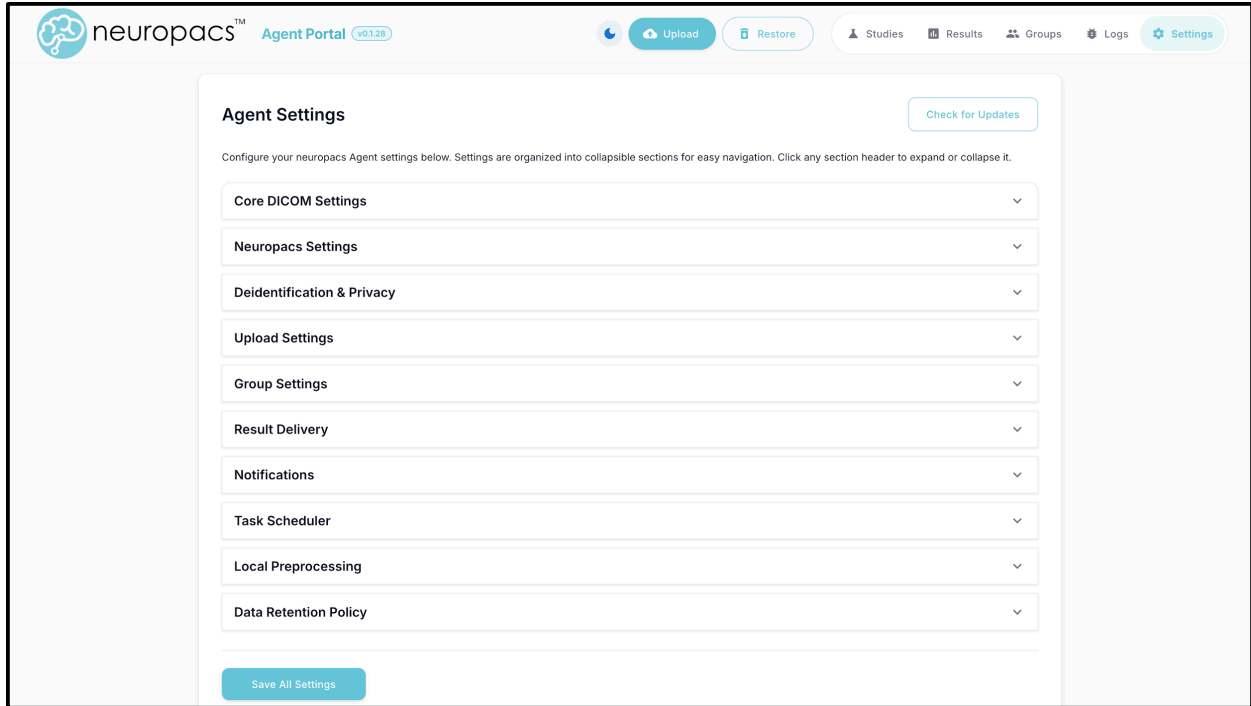


Figure 2: Settings page

Core DICOM Settings:

The Neuropacs Agent uses the configured DICOM settings (AE Title, Host, and Port) to securely receive imaging studies from PACS and other modalities. These values must match the corresponding routing configuration on the sending system.

- **Check for Updates:** Check if a new version of the agent is available.
- **AE Title:** A unique identifier used to register the Neuropacs Agent within the DICOM network for communication and routing.
 - **Recommended value: NEUROPACS_SCP**
- **DICOM Host:** The hostname or IP address where the Neuropacs Agent will receive DICOM connections. This value should reference the system on which the Agent is installed.
 - **Recommended value: localhost**
- **DICOM Port:** The TCP port on which the Neuropacs Agent listens for incoming DICOM connections. This port must be accessible to connected modalities or PACS systems.
 - **Recommended value: 104**

Neuropacs Settings:

The Neuropacs settings define how the Agent connects to the Neuropacs Cloud platform, including API access and authentication. These values must be configured correctly to enable secure data transmission and processing.

- **Neuropacs API URL:** The endpoint URL used by the Neuropacs Agent to communicate with the Neuropacs Cloud platform. This value must be reachable from the host system for successful data transfer and processing.
- **Service User:** The username of the service account provisioned for authentication to the Neuropacs Cloud. This account is used exclusively for system-to-system communication.
- **Service Password:** The password associated with the provisioned service account used to authenticate the Neuropacs Agent with the Neuropacs Cloud.
- **Select Product(s):** Specify one or more Neuropacs processing modules to be executed by the Agent for DICOM-routed workloads.

Deidentification and Privacy Settings:

The Neuropacs Agent applies HIPAA Safe Harbor–compliant de-identification procedures, removing protected health information (PHI) elements prior to transmission. Only non-identifiable data required for analysis is sent to the Neuropacs Cloud. The Agent also sends a periodic heartbeat message to the Neuropacs Cloud for monitoring of registered Agents. The content of these heartbeat messages can be configured to include only the desired information.

DEIDENTIFICATION SETTINGS ONLY AVAILABLE FOR CLOUD PROCESSING ENABLED SYSTEMS.

- **Date Handling Strategy:** Defines how date information is managed during de-identification. The Agent can either shift dates to preserve relative timelines while removing identifiable calendar references, or remove dates entirely for maximum privacy protection.
 - **Recommended value: Shift dates**
- **Shift Date Range:** Specifies the range of days used when shifting dates during de-identification. Dates are randomly adjusted within this range to preserve relative timelines while preventing re-identification.
- **Mask Zip Codes:** Controls whether ZIP codes are masked or generalized to protect patient identity.
 - **Recommended value: ON**
- **Send Accession Number to Neuropacs™ Cloud:** Controls whether the DICOM Accession Number is included in uploads to the Neuropacs Cloud. Useful for tracking, but may introduce identifying information.
- **Include System Metrics (CPU, Memory, Disk):** When enabled, system resource usage metrics are included in heartbeats for monitoring and troubleshooting.
- **Include Order Statistics:** When enabled, study/order counts (processed, failed, in-progress) are included in heartbeats.

- **Include Health Status:** When enabled, application and database health status are included in heartbeats. [RECOMMENDED]

Results Settings:

Manages visibility of completed studies in the Results page

- **Show External Orders in Results:** Displays completed orders from the neuropacs™ Cloud that were not initiated by this Agent but are associated with the same organization as the service user.

Upload Settings:

Controls how the Neuropacs Agent manages the transfer of imaging data to the neuropacs™ Cloud, including retry behavior and upload reliability. **Maximum upload size:** 5GB.

- **Maximum Upload Retries:** Specifies how many times the Agent will retry an image upload if it fails.

Quality Control Settings:

Controls what quality control checks are enforced by the Agent for incoming imaging. Disabling quality control checks can allow incomplete or non-compliant studies to continue processing. Use only when you intentionally accept this risk. Bypassed checks will be marked as bypassed in the QC report and treated as passing.

- **DICOM QC Checks:** Toggle QC checks for DICOM-based datasets.
- **NIfTI QC Checks:** Toggle QC checks for NIfTI-based datasets.

Group Settings:

The Neuropacs Agent supports grouping of studies into cohorts or project-defined collections, even when the studies originate from different subjects. These settings define how the Agent organizes, processes, and manages operations for studies that are associated with the same group.

- **Auto-Assign DICOM Uploads to Group:** All newly received DICOM studies will be automatically assigned to a specified group upon ingestion.
- **Automatic Report Sending:** Controls how group reports are configured to automatically send to a destination.
 - **Auto-Send Report When Group Completes:** When enabled, reports are automatically generated and sent once all orders within a group have reached a terminal state. Additional studies may be added to the group after completion, in which case a new report will be generated that includes the newly added studies.
 - **Send Group Reports on a Scheduled Interval:** Select how frequently group reports should be sent automatically, even when no new study completions have occurred. Multiple days may be selected for recurring

weekly delivery, or you may modify the cron expression directly to define a custom scheduling pattern.

- **Group Report Format:** Define the format of group reports, including the report type and filename to be generated and sent.
 - **Report Type:** Choose an available report type
 - **Prediction Indexes:** A CSV file containing the aggregated prediction indices generated by each selected product for every study within the group.
 - **Filename Template:** Define a filename template for group reports. The format supports the following placeholders: **{group_name}**, **{report_type}**, and **{date}**. Placeholders will be auto-populated prior to submission.
- **Group Report Delivery:** Specify the destination to which group reports will be delivered. **TIP:** Test your selected delivery method using the **“Test Delivery”** button.
 - **Delivery Method:** Supports email, REST API/webhook and None (manual download only).

Result Delivery Settings:

Controls how the Neuropacs Agent delivers and routes individual analysis result reports to external systems. **TIP:** Test your selected delivery method using the **“Test Notification”** button.

- **Delivery Protocol:** Controls the protocol used to route reports to external systems and the destination parameters. Supports email and REST API/webhook. **TIP:** Test your selected delivery method using the **“Test Delivery”** button. The following events are available:
 - **New Dataset Received**
 - **New Dataset Received**
 - **Quality Control Passed**
 - **Quality Control Failed**
 - **Upload Failed**
 - **Upload Succeeded**
 - **Report Delivered Successfully (Report Available)**
 - **Report Delivery Failed**
 - **Processing Failed (Failure Report Available)**
 - **Group Processing Completed**
 - **Group Report Sent Successfully**
 - **Group Report Send Failed**
 - **Local Processing Completed**
 - **Local Processing Failed**

Notification Settings:

Defines how and when the Neuropacs Agent delivers event-based notifications to the configured destination.

- **Notification Protocol:** Controls the protocol used to route notifications and the destination parameters.
 - **DICOM Encapsulated PDF:** Embeds the report within a DICOM object and transmits it using standard DICOM messaging.
 - **HL7 FHIR:** Encodes the report as a base64-encoded payload within an HL7 FHIR message for interoperability with clinical systems.
 - **REST API:** Sends the report as part of an HTTP POST or PUT request to a configured REST endpoint.
 - **None:** Stores reports locally on the Neuropacs Agent without transmitting them to any external system to be downloaded manually.

Scheduler Settings:

Controls when automated management tasks are executed by the Neuropacs Agent, including scheduled processing, cleanup, and maintenance activities. The frequency of these task executions determines how often the Agent communicates externally with the Neuropacs Cloud.

- **Available options:**
 - Every minute
 - Every 5 minutes
 - Hourly
 - Daily at midnight
 - Daily at 3:00 AM
 - Weekly on Sunday at 3:00 AM

Local Processing Settings:

ONLY AVAILABLE FOR LOCAL PROCESSING ENABLED SYSTEMS. Specifies how the Neuropacs Agent manages local preprocessing workloads, including when preprocessing is performed and how computational resources are utilized.

- **Enforce CPU and Memory Limits:** When enabled, local worker containers run with the configured CPU core and memory limits below. Disable to let containers run without explicit CPU/memory caps.
- **Container Configuration:** Determines how containers are configured for local preprocessing.
 - **Maximum Concurrent Containers:** Defines the maximum number of preprocessing containers that may run simultaneously.
 - **CPU Cores per Container:** Specifies the number of CPU cores allocated to each preprocessing container.
 - **Memory Limit per Container:** Specifies the amount of memory (in GB) allocated to each preprocessing container.

Data Retention Policy:

Defines how the Neuropacs Agent manages data retention policies, including the duration under which studies, logs, and related artifacts are stored or removed

- **Order Retention Period:** Specifies the number of days that order-related information is retained.
- **Log File Retention Period:** Specifies the number of days that log data is preserved.

Configuring a PACS to Send Studies to the Neuropacs Agent

To enable DICOM routing from your PACS to the Neuropacs Agent, configure a new DICOM destination (Storage SCP) within the PACS using the Agent's DICOM settings.

1. Gather Neuropacs Agent DICOM Settings

From the Agent configuration page, note the following values:

```
AE Title:      NEUROPACS_SCP
DICOM Host:    <agent_ip_address>
DICOM Port:    PORT
```

Example:

```
AE Title:      NEUROPACS_SCP
DICOM Host:    10.10.20.50
DICOM Port:    115
```

2. Create a New Destination in the PACS

In your PACS administration interface:

- a. Navigate to DICOM Nodes, Destinations, or Storage Targets (terminology varies by system).
- b. Select Add New Node or Add Destination.
- c. Enter the Neuropacs Agent settings:

```
AE Title: NEUROPACS_SCP
Hostname/IP: 10.10.20.50
Port: 115
```

3. Save the configuration.

4. Test the Connection

Most PACS systems provide a DICOM Echo / C-ECHO test.

Run:

- Ping or Verify
- Test Connection
- C-ECHO

Navigating the UI

Navigation Menu

This menu allows the user to:

- Navigate between pages in the interface
- Toggle dark mode
- Initiate a manual upload
- Refresh the page or enable automatic refresh




Figure 3: Navigation menu

Viewing Studies

Access the **Studies** page to view all imaging studies ingested by the Neuropacs Agent. Studies appear automatically once they are received and validated. Newly received studies will initially display **“No”** under the **Stable** column until processing completes and the study is marked as stable.

The **Actions** column provides the following options:

1. **Check Status** – Displays the current processing status (enabled after processing begins)
2. **View Quality Control Results** – Opens the QC details for the study
3. **Preview Report** – Preview PDF reports for completed/failed studies
4. **Cancel Order** – Cancels the study’s processing request
5. **View Study Information** – Shows metadata and study details
6. **Assign to Group** – Assign study to a group

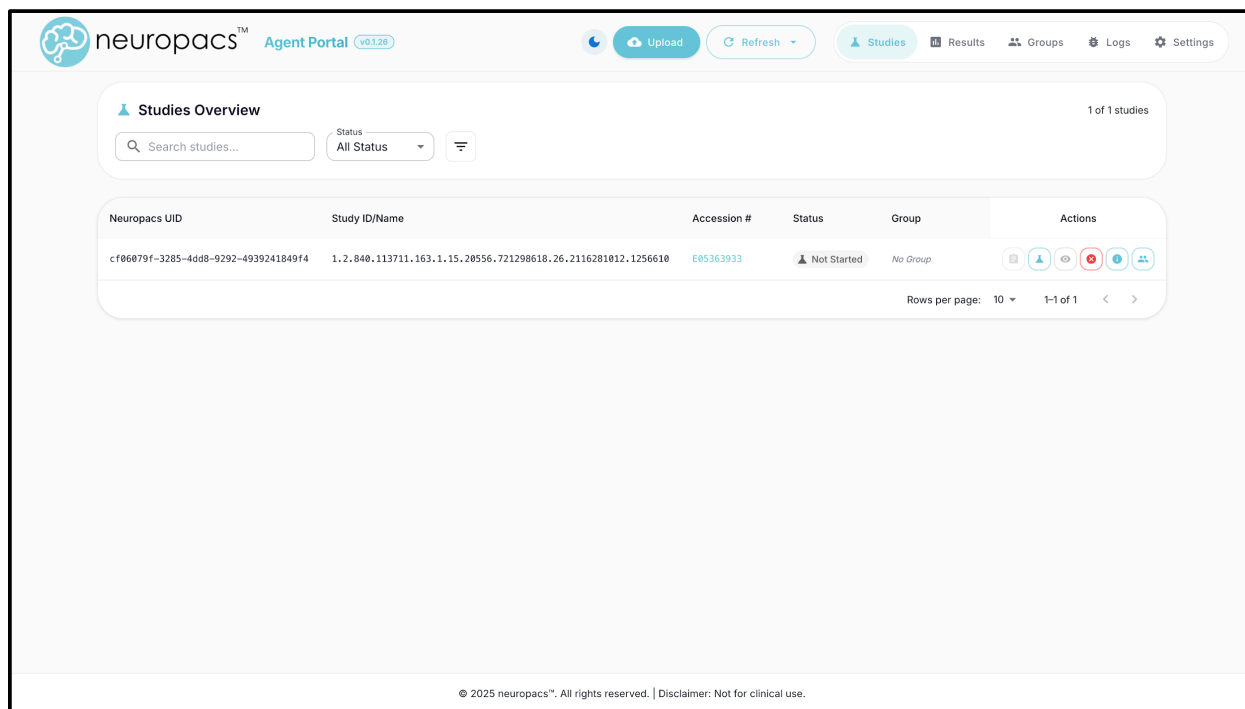


Figure 4: Unstable study on Studies page

View Additional Information

To view additional information for a study/order, click the **Clipboard** button in its corresponding row.

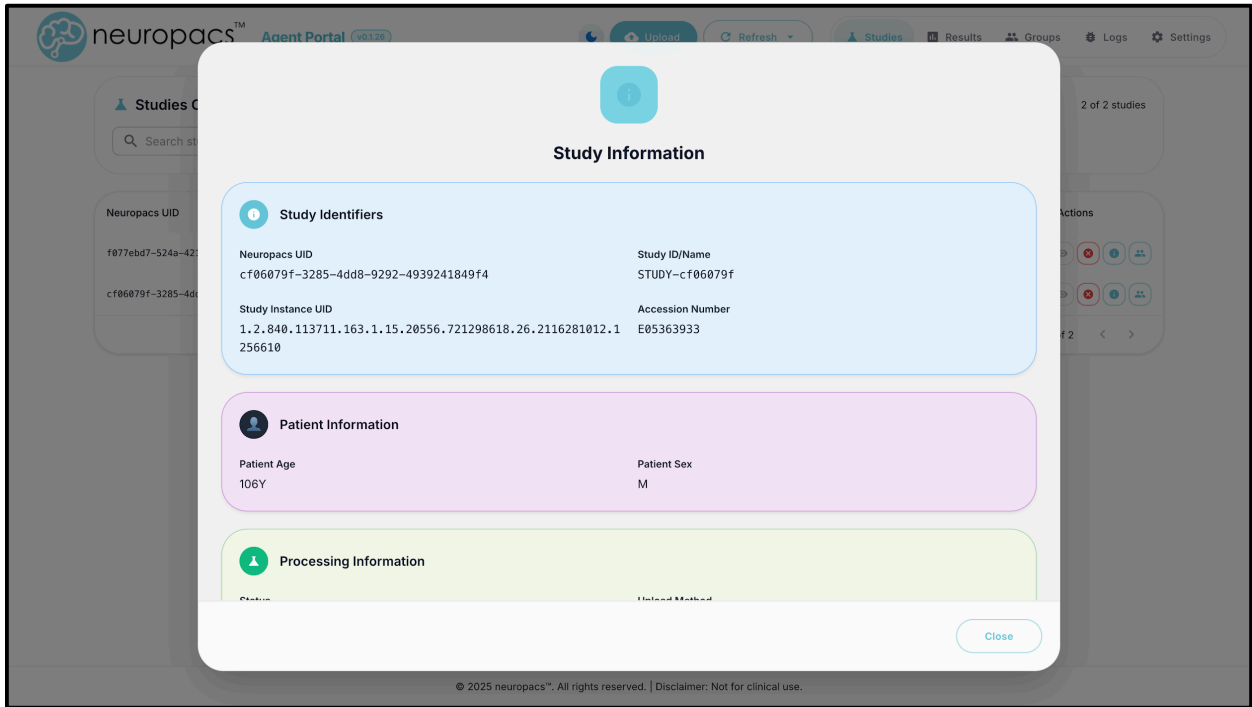


Figure 5: Study information popup

Quality Control Reports

After a study received through the DICOM SCP interface becomes stable, it is automatically evaluated by the Neuropacs quality control system to ensure that only appropriate datasets proceed to processing. To view the quality control results for a study, click the **Eye** button in its corresponding row.

The screenshot displays a 'Quality Control Report' window in the Neuropacs Agent Portal. The report is titled 'Quality Control Report' and 'Study Validation Analysis', generated on Dec 3, 2025, at 5:23 PM. It provides an overall score of 75% (3/4 checks passed) and indicates that the study has 1 quality issue(s) that require attention before processing. The report includes detailed test results for 'MRI Series Detection', which is marked as 'PASSED' with a timestamp of 17:23:05. A 'Close' button and an 'Override QC' button are visible at the bottom of the report window.

Figure 6: Quality control report 1/2

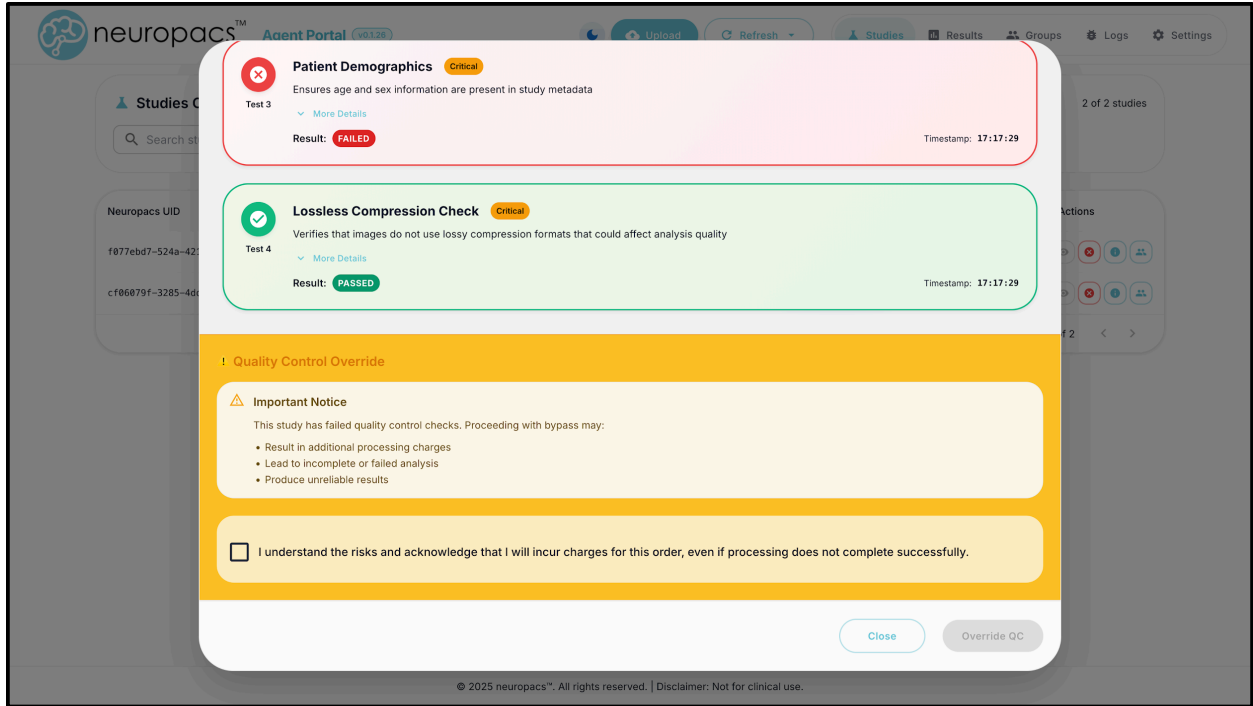


Figure 7: Quality control report 2/2

A failed quality control check may be bypassed when necessary; however, this should be done with caution.

Warning: Bypassing QC can result in additional processing charges, incomplete or failed analyses, or unreliable results.

Group Assignment

To manually assign a study to a group, select the **Group** button in the corresponding row. After assignment, a group tag will appear in that row under the *Group* column.

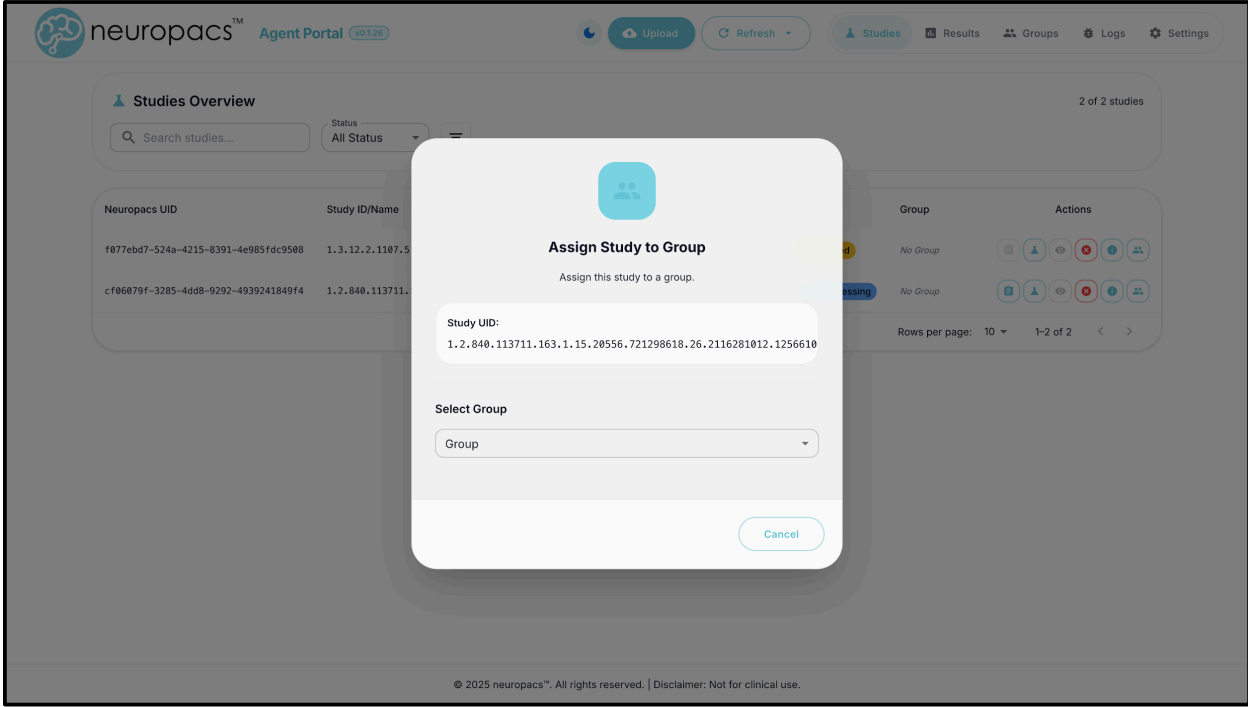


Figure 8: Group assignment dialog

Manual Uploads

To manually upload an order, click the **Upload** button in the navigation bar. Then follow the steps below to submit a new study:

1. **Study Name (Optional):** Provides a human-readable identifier for the study and is combined with a Neuropacs-generated UUID to form a unique reference. The name must be 5–50 characters and may include letters, numbers, spaces, slashes (/), colons (:), and hyphens (-).
2. **Accession Number (Optional):** Enter an accession number if available. This field supports downstream tracking and reconciliation with clinical systems.
3. **Patient Age (Optional):** Enter the patient’s biological age associated with the imaging study.
4. **Patient Sex (Optional):** Enter the patient’s biological sex associated with the imaging study.
5. **Assign to group (Optional):** Assign the study to a group. If selected, the product selection will be autopopulated with the products of that group.
6. **Select Product(s):** Select one or more products to be run on the uploaded study.
7. **Select Upload Type:** Choose the type of study data being uploaded. Available options include:
 - **Folder:** A folder containing DICOM files with optional JSON sidecar, or NIFTI/BVEC/BVAL files with required JSON sidecar.
 - **ZIP Archive:** A single ZIP archive containing DICOM files with optional JSON sidecars, or NIFTI/BVEC/BVAL files with required JSON sidecar
8. **Select Upload Button:** Start the upload process. A progress bar will appear that will show the progress of the upload. Do not navigate away or close the window during an upload as this may corrupt the upload process.

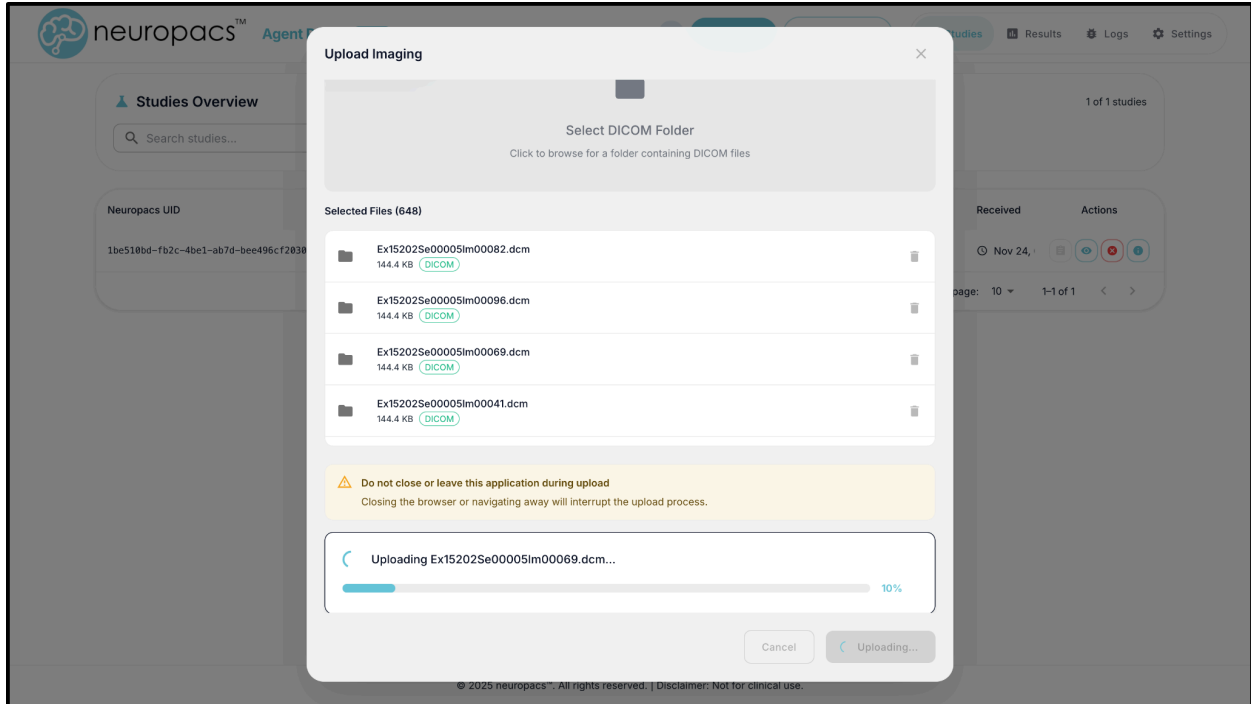


Figure 9: Manual upload indicator

Note: Patient biological age and sex must be provided through one of the following:

- this upload form,
- DICOM metadata, or
- a JSON sidecar when uploading via ZIP archive

Note: The Neuropacs Agent applies the following order of precedence when determining metadata values (highest to lowest):

1. Metadata entered in the upload form
2. JSON sidecar metadata
3. DICOM metadata (when available)

Note: For more information regarding upload specifications, visit our [Documentation Website](#).

Viewing Results

Access all available reports through the **Results** page. This page allows you to view reports directly in the browser, download them in multiple formats (JSON, PNG, PDF, TXT, XML), and manually re-send reports if delivery fails or a retransmission is required.




Figure 10: Integrated UI Results page

Manage results by selecting the **Gear** icon in the **Actions** tab.

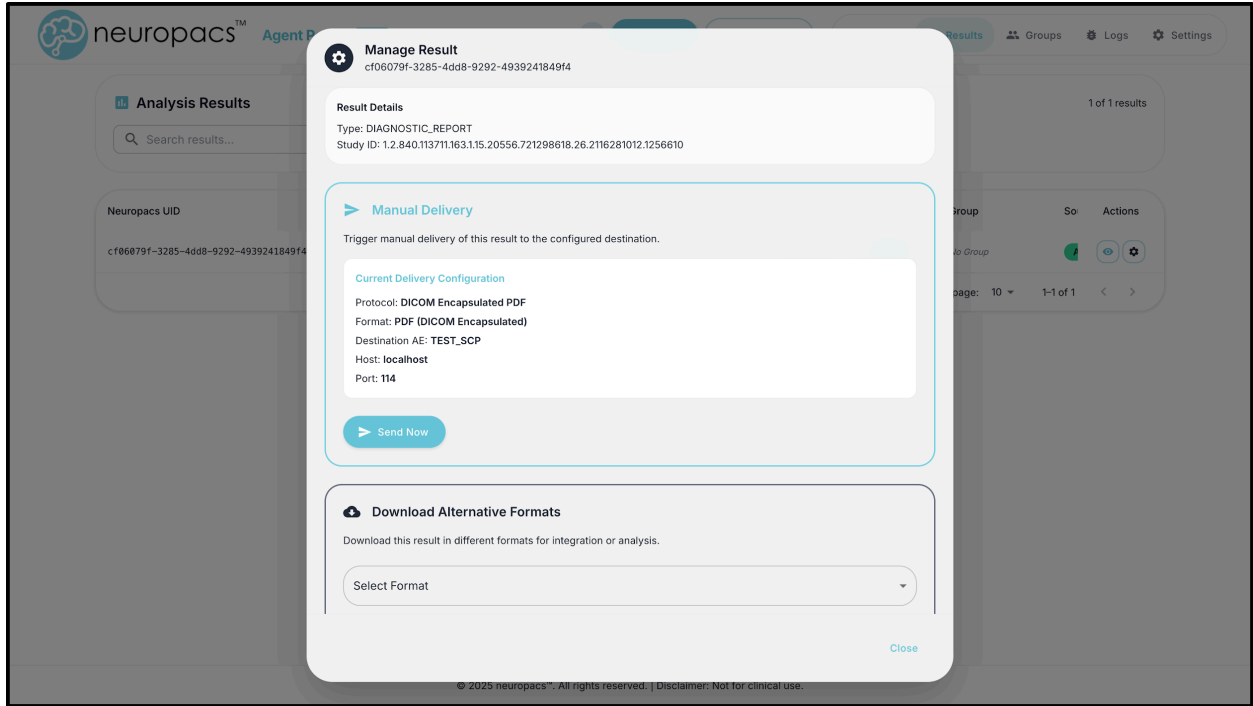


Figure 11: Manual result management popup

Manage results by selecting the **Eye** icon in the **Actions** tab.

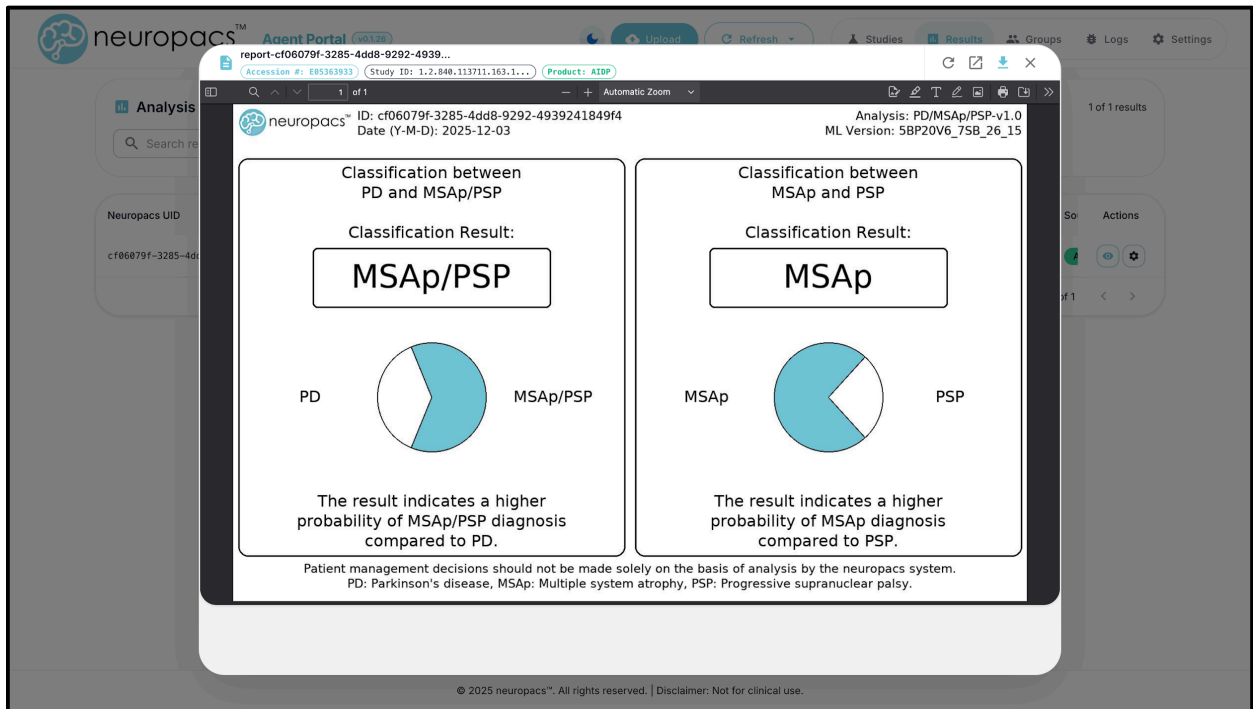


Figure 12: Integrated PDF report viewer

Group Management

Access the **Groups** page to create and manage groups and download group reports generated by the Neuropacs Agent. Configurations related to group reports can be managed in the **Settings** page.

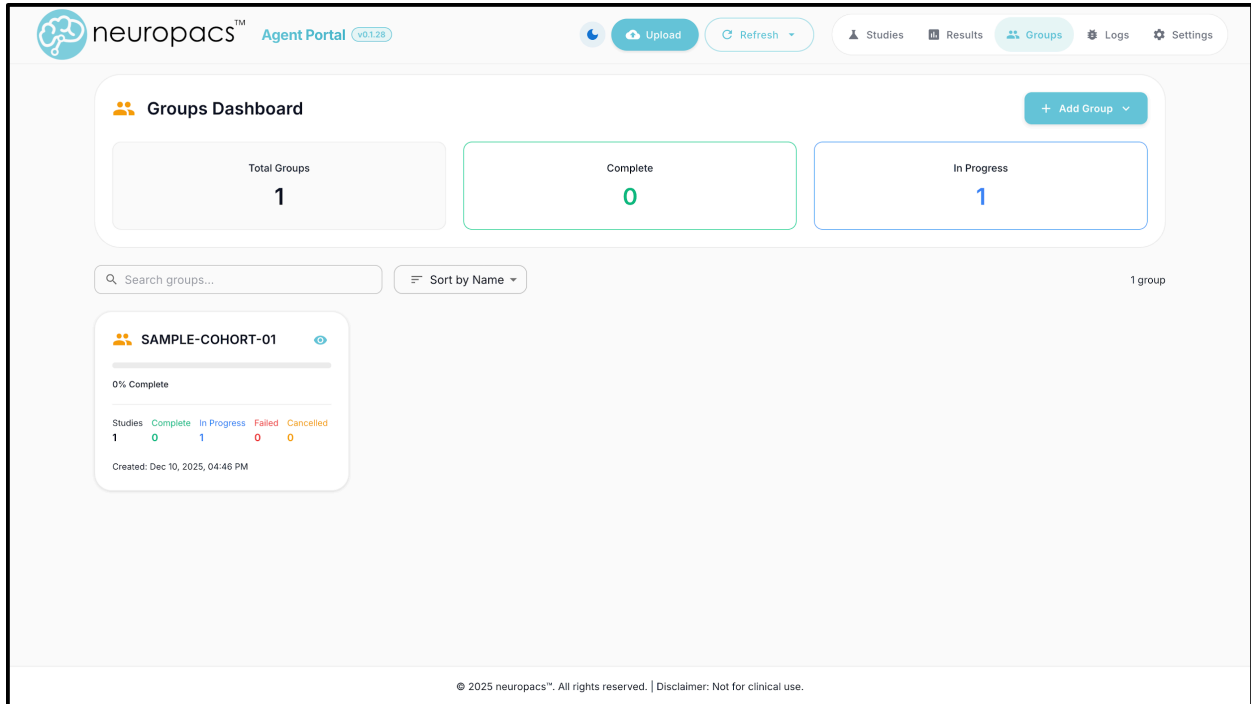


Figure 13: Groups page

Select the **Eye** icon for a group to view details and manually manage reports. Select **Send Report** to manually send the report to the configured destination. Select **Download Report** to instantly download the CSV report to disk.




Figure 14: Group details

Create groups by selecting the **“Create New Group”** button. **Note:** All studies associated with a group MUST use the same products.




Figure 15: Group creation form

Uploads a group of studies and auto-assign to a group by selecting the **“Upload New Group”** button. **Note:** A group of studies must be uploaded as a single folder containing a single ZIP archive for each study. Each ZIP must include either a DICOM study folder with an optional JSON sidecar, or a NIfTI (.nii/.nii.gz) with matching .bvec, .bval, and JSON sidecar. Select **“Download JSON sidecar template”** to download a template file.

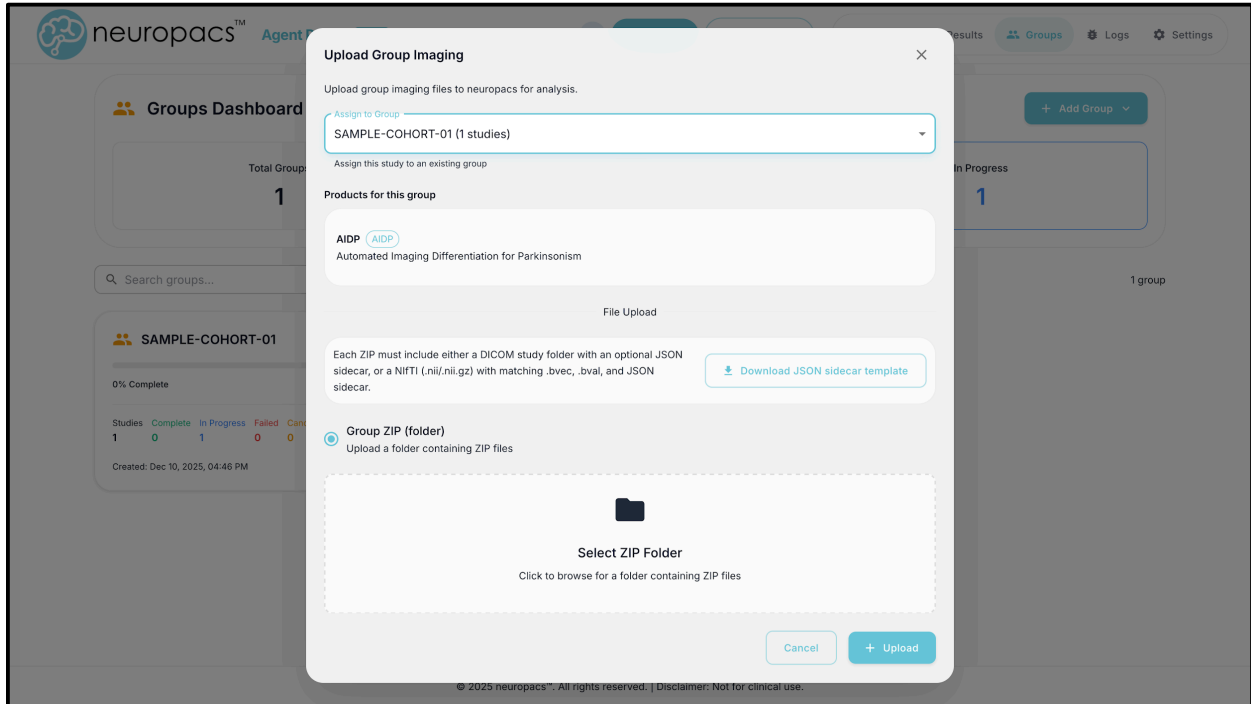


Figure 16: Group upload form

Logs

Access the **Logs** page to view, filter, search, stream, and download logs generated by the Neuropacs Agent. Selecting **File, CSV, or JSON** options will download the current log file in the corresponding format. Selecting the **ZIP** download option will download all available log files. Select **Stream logs** to enable live streaming of new logs.

Neuropacs Cloud logs can be provided upon request.

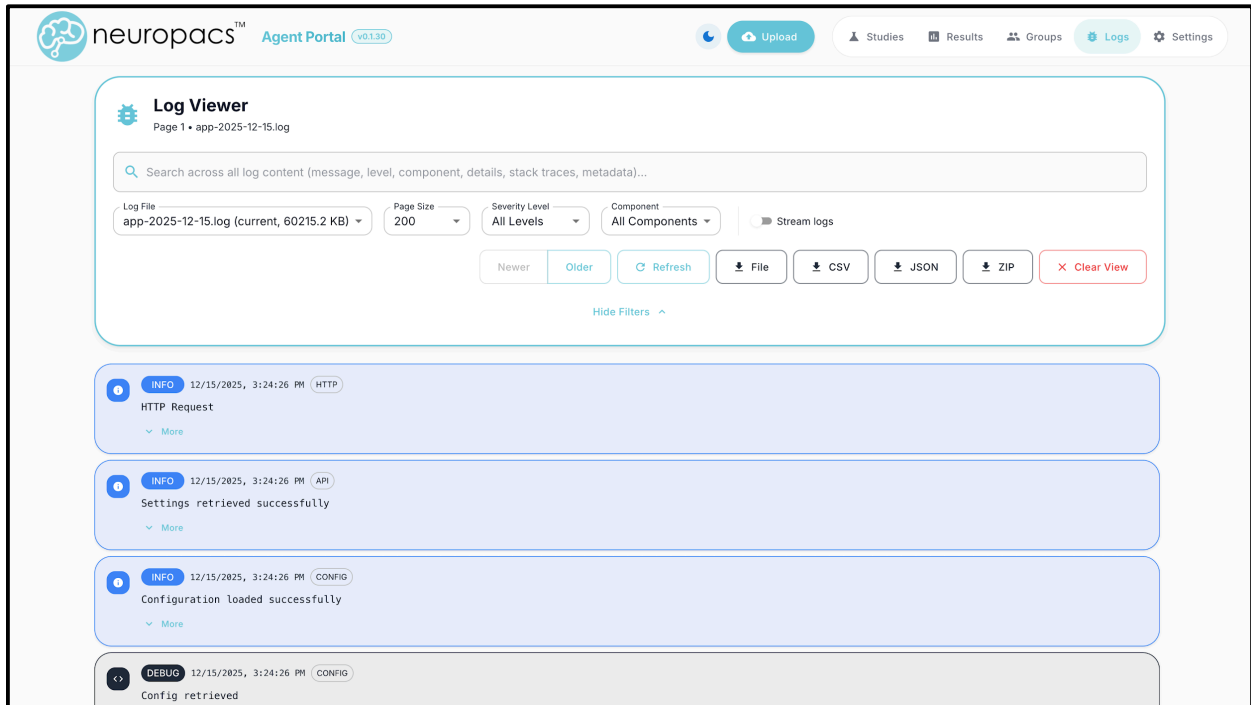


Figure 17: Logs page

Settings

Access the **Settings** page to view/modify Neuropacs Agent configurations.

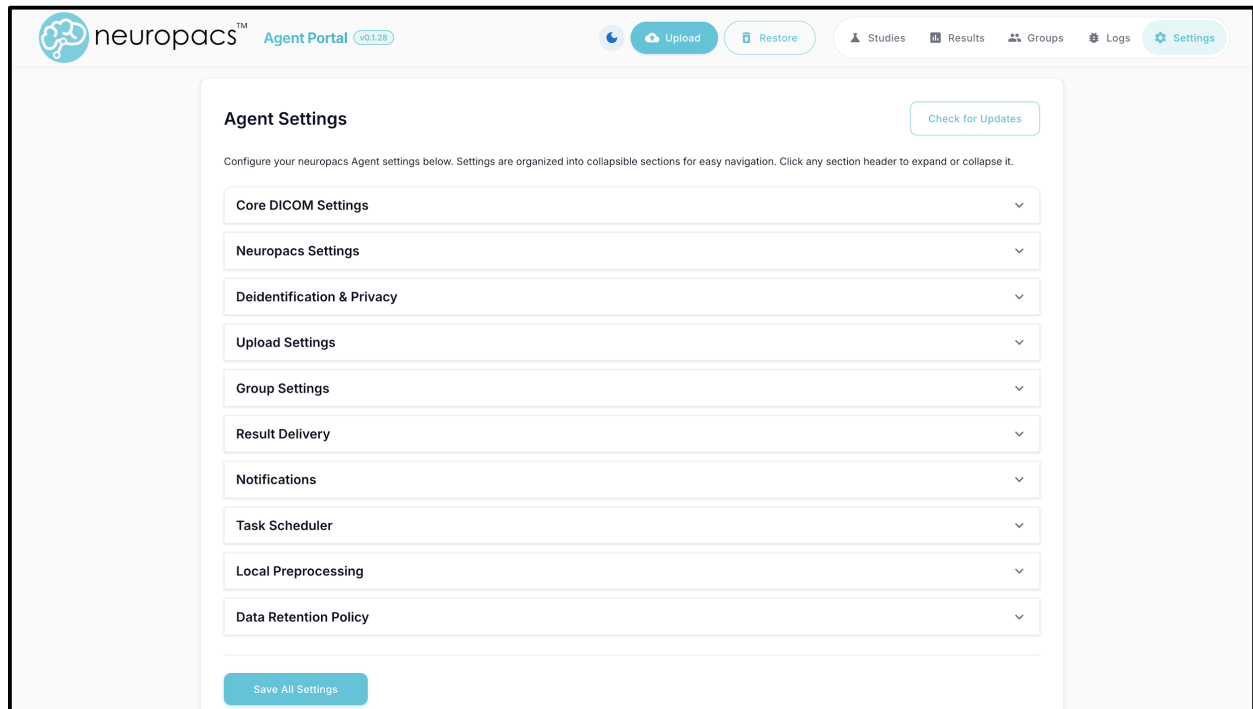


Figure 18: Settings page

Updates

Updates can be performed directly in the **Settings** page. To check if your installation is out-of-date, select **“Check for Updates”**. If your system application version does not match the latest release, an option will be available to update your system. This process takes ~5 minutes and retains all application data.

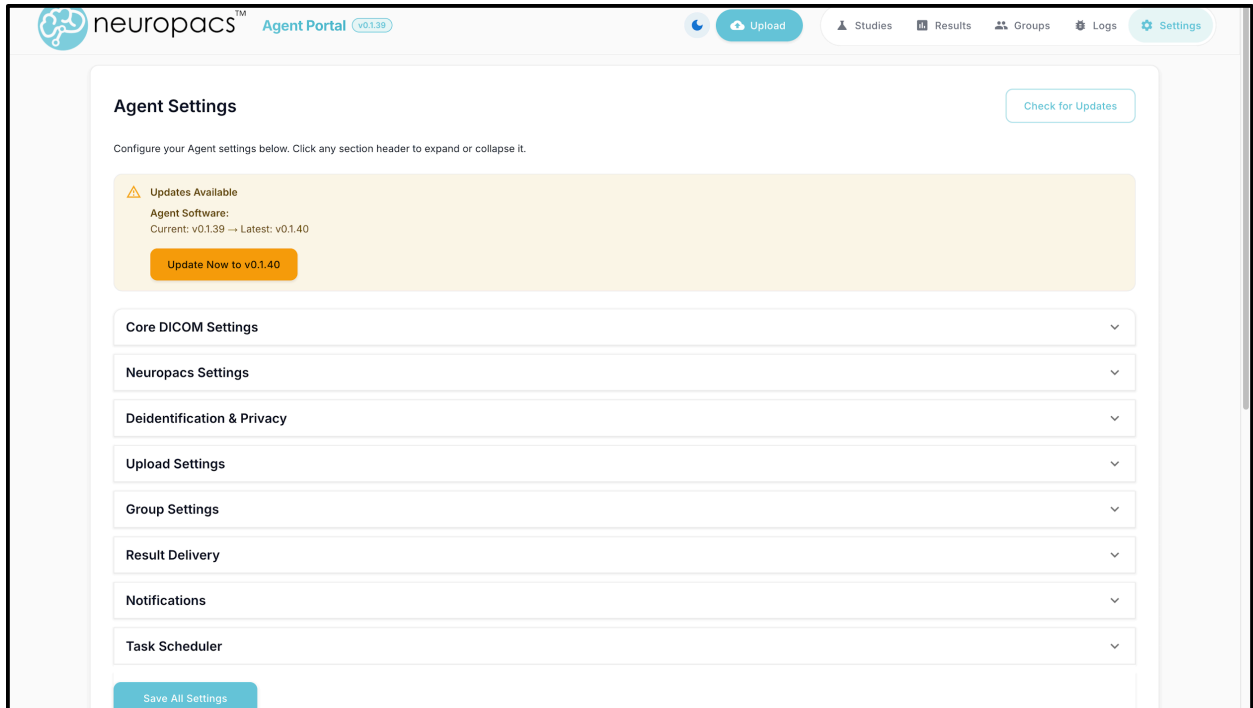


Figure 19: Update dialog

API Documentation

The Neuropacs Agent includes built-in interactive API documentation powered by Swagger UI. It provides a complete reference for all REST endpoints exposed by the agent, organized by functional area: health checks, study management, order lifecycle, result retrieval, batch operations, delivery, settings, updates, logs, and more.

Access: Navigate to /api-docs on your running agent (e.g. <https://localhost:3001/api-docs>).

Neuropacs Agent API 0.1.60 OAS 3.0

REST API for the Neuropacs Agent – manages DICOM studies, cloud order processing, result delivery, local processing jobs, batch management, and system administration.

Servers
/ - Local agent server

Health

Health checks and system status

- GET /health Health check (DB + DICOM)
- GET /api/health Simple health check
- GET /api/status Detailed system status

Studies

Study management and QC operations

- GET /api/studies List all studies
- GET /api/study-info/{neuropacsId} Get combined study information
- PUT /api/bypass-qc Bypass QC for a study
- PUT /api/force-stable Force a study to stable state

Results

Diagnostic and failure report retrieval

- GET /api/results Get all results
- GET /api/reports Get PDF report

Figure 20: API Swagger page

On-Premise Quality Control (QC)

The Neuropacs Agent performs on-premise quality control to validate the integrity and suitability of submitted imaging data prior to processing. Quality control checks vary based on the imaging data type and are designed to ensure compliance with Neuropacs analysis and clinical requirements.

DICOM-Based Studies

- **MRI Series Validation** — Confirms the study contains appropriate MRI imaging series.
- **DTI/DWI Detection** — Identifies the presence of Diffusion Tensor Imaging (DTI) or Diffusion Weighted Imaging (DWI) sequences.
- **Patient Demographics Validation** — Ensures required patient metadata, including age and sex, is present.
- **Instructions for Use (IFU) Compliance** — Verifies that study and series descriptions conform to Neuropacs Instructions for Use (IFU) requirements.
- **Lossless Compression Verification** — Confirms images are not encoded using lossy compression formats that may degrade analysis quality.

NIfTI-Based Studies

- **NIfTI File Detection** — Validates the presence of NIfTI-format imaging files.
- **DTI Required File Validation** — Ensures all required DTI files are present (.nii/.nii.gz, .bvec, .bval, and .json).
- **Metadata Sidecar Validation** — Confirms the JSON sidecar contains required patient demographic information.

Important: Bypassing QC checks may:

- Result in additional processing charges
- Lead to incomplete or failed analyses
- Produce unreliable or non-diagnostic results

Security Overview

The Neuropacs Agent is designed with security as a foundational principle. It uses modern, industry-standard methods to protect sensitive information such as passwords, service credentials, and system communications.

Key Security Features

Automated Credential Protection

- Database passwords are automatically generated during installation
- No manual password handling is required
- Sensitive credentials are never stored in plain text files
- Access to secrets is restricted to only the components that need them

This reduces the risk of accidental exposure and removes the burden of managing credentials manually.

Secure Storage & Encryption

- Credentials are stored using secure container-based secret management
- Sensitive data is encrypted at rest (when stored on disk)
- All network communications can be encrypted in transit using TLS (Transport Layer Security)
- Strict file permissions prevent unauthorized access

TLS ensures that data exchanged between systems — including medical images and reports — cannot be intercepted or read by unauthorized parties.

Separation of Configuration & Credentials

- System configuration settings are kept separate from passwords and sensitive data
- No credentials are stored in source code or version control
- Service credentials within the database are encrypted
- Secure credential rotation is supported without requiring system downtime

This architecture follows Zero Trust security principles, meaning every component must explicitly authenticate before accessing protected resources.

Support

For additional information regarding the Neuropacs Agent or this user guide, please contact our support team at **support@neuropacs.com**.