iDNA Applications Data Flow in Detail

This page follows the data that iDNA Applications reads from your Notes environment. It covers where the data is stored, how it is processed, who sees it, and what leaves your network, including to AI clients. The reading itself is described on iDNA Applications Data Collection and Access.

- Overview
- Storage on the Appliance
- Processing
- Who Sees What
- What Leaves Your Network
- AI Clients and the MCP Server
- Related Topics

Overview

1. **Read:** iDNA reads from your Domino servers with the collection Notes ID (see iDNA Applications Data Collection and Access).
2. **Stored** on the appliance, a Linux host in your network.
3. **Processed** every night on the appliance into the analyses.
4. **Shown** to the people you give an iDNA login.
5. **Retrieved by AI clients**, only if you enable the MCP server.

For reviews by a works council, a data-protection officer or a compliance function, panagenda provides two documents. The AI position paper covers the iDNA Applications Companion and the MCP server and is available from panagenda on request. The **note on user name pseudonymization** is available in English and German. This page summarizes both. Where a detail matters, the documents apply.

Storage on the Appliance

- iDNA runs as a set of containers on a Linux host that you provide and operate in your network. Access to the host is yours to control.
- iDNA stores the collected data in a PostgreSQL database on that host, with real names. It applies pseudonymization later, when it prepares the data for the analyses.
- Files on the host include the design exports, as compressed files, and the Notes runtime. The Notes runtime holds the collection ID file and its password, which iDNA needs to collect without anyone logged in.
- Each collection replaces snapshot data such as the directory, server documents and the catalog. Activity data such as session records builds up over time.
- Raw data that the nightly processing has already used is deleted after 30 days by default. An administrator can change the period, or switch the deletion off, in the Portal under [Data Retention and Cleanup Options]. The processed figures are kept.
- Backups, exports and copies of the host are yours to control. panagenda has no standing access to the appliance. You grant access for support or consulting.

Processing

iDNA computes the analyses every night on the appliance, so dashboards show the data up to the previous night. The nightly processing pseudonymizes user names as described in the note on user name pseudonymization. In short:

- **Pseudonymization is the default.** iDNA replaces user names in the per-user usage analyses with a pseudonym derived from the name. These analyses are who used which application and when, the per-user summaries, the license analysis, and the editor statistics of the content analysis. A person keeps the same pseudonym over time, so trends still work. The organization part of the Notes name stays readable.
- **It is pseudonymization, not anonymization.** iDNA has no function that turns a pseudonym back into a name, but anyone who holds the list of names and the method can rebuild the mapping.
- **Small departments and locations:** while pseudonymization is on, a department or location with fewer than 15 assigned users is shown under a placeholder. The license can set another threshold. The row and its figures stay, so totals still add up.
- **Shown with real names on purpose:** access-control entries, agent signers, group memberships, the people who last changed Profile Documents, who last accessed and who last wrote to each application, person documents from the Domino Directory, and MarvelClient data. These describe permissions, responsibilities and recent use. They are not intended for assessing anyone's work.
- **Switching it off** takes a new license. panagenda issues it only after your organization has confirmed in writing that it has taken the necessary internal steps, including with the works council. On a trial license, the usage analysis is always pseudonymized.

Who Sees What

- **Inbound connections:** a standard installation accepts browser connections on port 443 (HTTPS) and redirects port 80 to it. Behind port 443 are the Portal, the admin interface, the Metabase reports, DataMiner and, when enabled, the MCP server. Two more ports are optional and closed by default. VNC (TCP 5901) is open only while an administrator runs the VNC server on the appliance host for its graphical desktop. PostgreSQL (TCP 5432) is open only if the database port is published for direct database access.

- **Logins:** the Portal and the admin interface use local iDNA accounts or your LDAP directory, set up separately for each. Metabase has its own login. You decide who gets a login.
- **Roles:** Portal users have a role (for example Admin, Editor or User) that decides which pages they can use. Separate roles cover MCP and API access. The admin interface and the system pages are for administrators only.
- **Pseudonymization** applies to the prepared data, so every view built on it shows the pseudonyms. Per-user views in the Portal show pseudonyms while pseudonymization is on.
- **Outgoing connections for logins and email:** iDNA connects to your LDAP directory if you use it for logins. It connects to your mail server to send alerts and report subscriptions to the recipients you configure.
- **On a trial license**, some names and databases are masked or hidden. iDNA Applications Trial License explains which.

What Leaves Your Network

Connection	From to	What is sent	When
Software images	Appliance panagenda's download server (docker.panagenda.com)	Requests for the iDNA container images	Online installation and upgrades. The offline installer needs no connection.
Update check	The user's browser panagenda's file server (files.panagenda.com)	A request for a version file, without iDNA data	When the Portal login page is shown, unless the license has expired
Content analysis tokens	The user's browser panagenda's licensing server (licensing.panagenda.com)	The license key and the replica ID of the database to scan. Only the license key when the token status is shown.	When a user starts a scan of a replica set not yet authorized on the appliance, or opens the token status
Help links	The user's browser panagenda's documentation	A normal page request	When a user clicks a help link
AI clients	Appliance the user's AI assistant platform	The results of the tools the assistant calls	Only when the MCP server is enabled. See AI Clients and the MCP Server
Support bundles	An administrator panagenda support	The appliance's configuration and log files. Review the bundle before you send it.	Only when an administrator downloads a bundle and sends it

- The license is a file you upload, and iDNA checks it on the appliance. iDNA's own components send no usage statistics or crash reports to panagenda.
- **Metabase**, the reporting component, has its own settings for anonymous usage tracking and for checking for updates. Check them in Metabase's admin settings.

AI Clients and the MCP Server

The MCP server lets an AI assistant that your organization already uses retrieve analysis results from iDNA. The AI position paper describes it in full. In short:

- **Off by default.** Your organization requests it from panagenda as a license option, and your administrator switches it on. Your organization acknowledges the position paper when it opts in.
- **Restricted mode is the default:** a fixed set of tools, no ad-hoc queries, no per-user analysis tool and no access to the tables that record individual activity. The AI sees application metadata, aggregates and source code. Notes records some names as part of an application, for example access-control entries and agent signers. The server's database account can read them, but no current restricted-mode tool returns them.
- **Permissive mode** (ad-hoc queries and per-user analysis) must be chosen deliberately, as a separate license decision.
- **Access:** every user gets a personal access token, valid for 365 days by default. All tokens have the same data scope. The appliance records every tool call in an audit log.
- **Where results go:** tool results travel over HTTPS to the user's assistant platform and into the AI model's context. From there, storage, retention and use are governed by your organization's contract with the platform's vendor. In its default configuration the MCP server opens no connection to any external service.
- **Purpose:** the AI components are intended for application portfolio analysis. They are not intended for monitoring or evaluating individual employees.

Setup: Setting up the MCP Server and Setting up iDNA Applications Companion.

Related Topics

- iDNA Applications Data Collection and Access: what is read, the access it needs, the connections to Domino.
- iDNA Applications License Types and Options: license types and options, including pseudonymization and MCP.
- iDNA Applications Trial License: what a trial reads, analyzes and shows.
- Setting up the MCP Server: enabling the MCP server and issuing tokens.
- Content Age and Data Retention Analysis in iDNA Applications: what content age measures.
- SSL Certificate: using your own certificate for port 443.
- User Roles (Authority): the Portal roles.
- LDAP Settings: logins through your LDAP directory.
- Appliance Maintenance and Updates: software updates.

- Download Log Files: support bundles.
- Configuration: Usage by Organizational Units: department and location data.