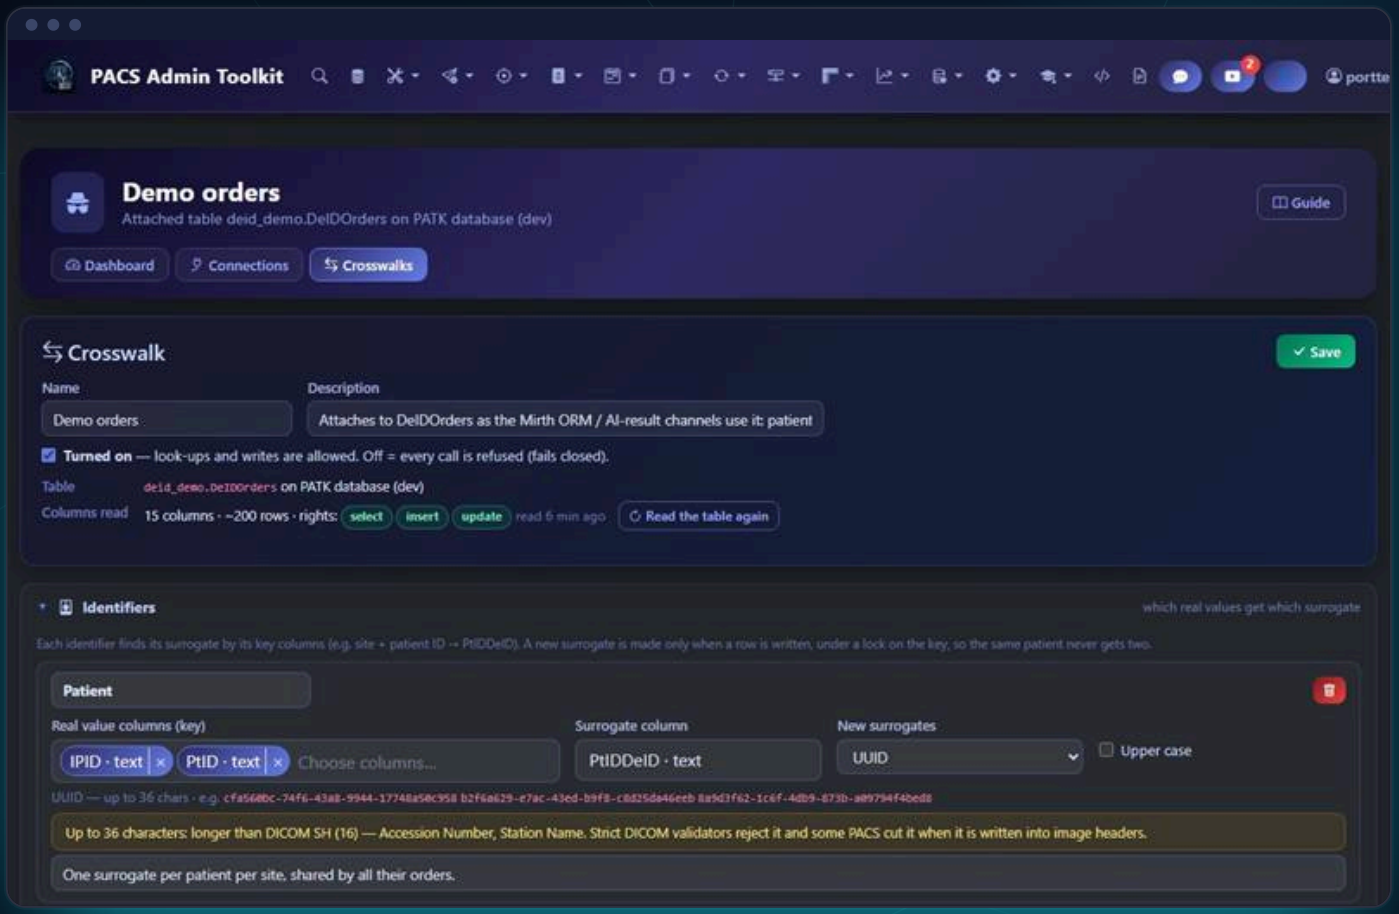
Platform & Integration | Leads with Trust

De-identification

Send studies and reports out under stable surrogates, and match the results back to the right patient.



Crosswalk tables give every patient, accession and physician one stable surrogate, created atomically so two messages for a new patient never get two. Results that come back are matched to the real values only when exactly one row fits, and every re-identification is audited.

One surrogate

per patient, even when messages arrive at the same moment

Fails closed

a down database or ambiguous match stops the message, never leaks it

Audited

every re-identification needs a permission and a reason

Replaces: **Interface-engine scripts and hand-built crosswalk tables**

License module `DEID`

THE CHALLENGE

Why it matters

Research partners, AI vendors and teaching files need imaging without patient identity, but the results they send back have to land on the right patient. That needs a crosswalk: a table pairing each real identifier with a surrogate.

Crosswalks are usually built inside interface-engine scripts. Two messages for a new patient can create two surrogates, a missing index slows every message, a lookup that matches two rows silently picks one, and nobody can say who looked up a real identity or why.

IN SHORT

Crosswalk tables give every patient, accession and physician one stable surrogate, created atomically so two messages for a new patient never get two. Results that come back are matched to the real values only when exactly one row fits, and every re-identification is audited.

CAPABILITIES

What it does

- **Crosswalks you can see**
  Attach an existing crosswalk table or let PATK create one; map which columns hold the real values and which hold the surrogates.
- **One surrogate per value**
  New surrogates are created under a lock on the real key (or by unique indexes on a PATK-managed table), so concurrent first messages always share one.
- **Surrogates that fit**
  UUIDs, keyed hashes, prefixed random patterns, counters or fake names, with warnings when a value is too long for the DICOM field it goes into.
- **Safe row writes**
  An upsert by the table's unique key, with a rule per column: overwrite, keep the first value, fill only if empty or never touch. Existing surrogates are never replaced, and a write that changes nothing leaves the row alone.
- **Re-identification that refuses to guess**
  A match names the de-identified columns that must find exactly one row; two rows means refused, not picked.
- **Index advice**
  Checks that every lookup has an index and suggests concurrent builds and redundant indexes to drop, with a script for a DBA.
- **Secure connections**
  Encrypted passwords, SSL, short timeouts and read-only connections for sites that only re-identify.
- **A bench to try it**
  De-identify a value, dry-run a row write or re-identify with a reason, before the crosswalk is turned on.

Presets for existing tables

Ready-made settings for common order and physician crosswalk tables, including physicians keyed by site so the same name at two sites stays two people.

De-identify inside the rules you already use

De-identify, Re-identify and Capture steps in the HL7 Field Editor and the DICOM Tag Editor: an order goes out with surrogates, the study's images get the same patient surrogate, and an AI result comes back with the real IDs.

Held, never leaked

A step that can't de-identify holds the message (an HL7 sender gets a negative ACK and resends; a DICOM image is refused), a PHI leak guard holds anything that still contains a value a step replaced, and a surrogate too long for a DICOM tag is refused instead of being cut.

Filter and trim what goes out

Drop messages that match (a NO READ report), blank a segment except the fields named, and delete report lines that contain any of a list of phrases.

Tests that don't reveal or write

Rule tests look surrogates up without creating any, and mask real values for users without the re-identify permission.

Profiles instead of hand-written rules

Tick what is replaced, emptied, shifted or removed (patient identity, demographics, order, providers, dates, free text, PS3.15 attributes), preview the result on a sample, and compile it into an HL7 Field Editor or DICOM Tag Editor template. Re-identify profiles restore results from a match of de-identified columns.

A field map for the crosswalk row

Each table column with its type and length next to the message value that fills it, tested with a rolled-back write.

Bring an existing crosswalk along

Load an export of any size from CSV or Excel; existing surrogates are kept as they are, conflicts are counted and left out, and a failed load changes nothing.

Whole files in one go

De-identify the columns of a spreadsheet for a research partner, or re-identify a returned file with a reason on record.

A review queue for early results

A result that arrives before its order reached the crosswalk is acknowledged, kept encrypted and retried until its row appears, then failed or passed after a hold time you set.

Look-ups with a paper trail

Find a surrogate, or re-identify one with a reason; every look-up is audited and counted.

People whose details change

When a last name or a patient ID changes, a key alias keeps the same surrogate for the new values; re-identify gives back the original, the newest or a preferred key, or stops for a person to decide.

Point and pick, not variable names

Every "which message value goes here" box suggests the likely HL7 field or DICOM attribute for each key and opens a searchable list of all of them, so nobody needs to know placeholder syntax.

Ready-made packs

Load a complete setup for an AI vendor, a research export, a teleradiology service or a teaching file, or bring your own exported pack. Crosswalks arrive switched off and profiles uncompiled, so nothing runs until it has been checked.

Compare before you switch

Run recorded interface-engine traffic through the new rules, read-only, and see field by field where the output would differ before a channel is moved.

Files that arrive on their own

Watch a folder, a share or an SFTP / FTPS server: each new file is de-identified or re-identified and delivered, the source moved aside, and no file taken twice.

A data API with limits

Let a research team or an AI vendor query just the columns, rows and look-ups an administrator allows, with an API key or token, daily limits and every call audited.

WORKFLOW

How it works

1

Connect

Add the database that holds the crosswalk, or use PATK's own.

2

Describe

Attach the table (or let PATK create one) and map identifiers, keys and surrogate columns.

3

Check

Build the suggested indexes and try look-ups, dry-run writes and re-identification with test values.

4

Add the steps

Put De-identify (and Capture) steps in the outbound HL7 and DICOM templates, and Re-identify in the template on the listener that receives results; each editor's Test tab shows the outcome as a dry run.

5

Turn on

Look-ups are cached; new surrogates are created atomically; every refusal is counted on the dashboard.

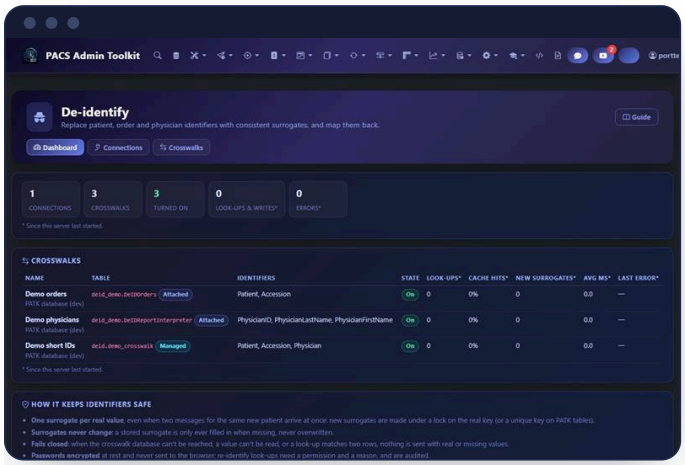
6

Match back

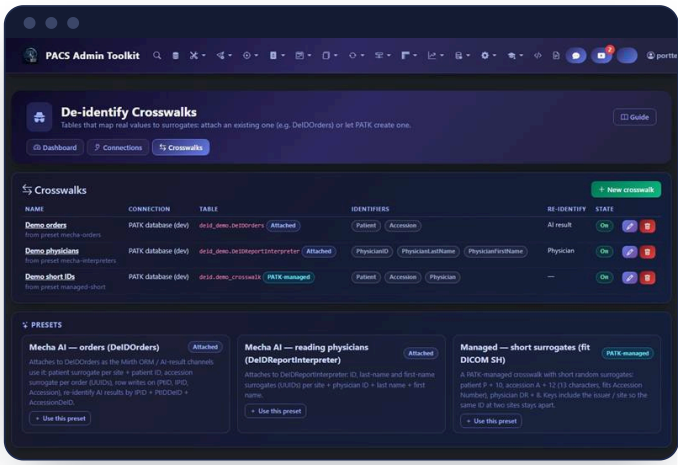
Results are re-identified only when exactly one row fits, with the reason in the audit log.

IN THE PRODUCT

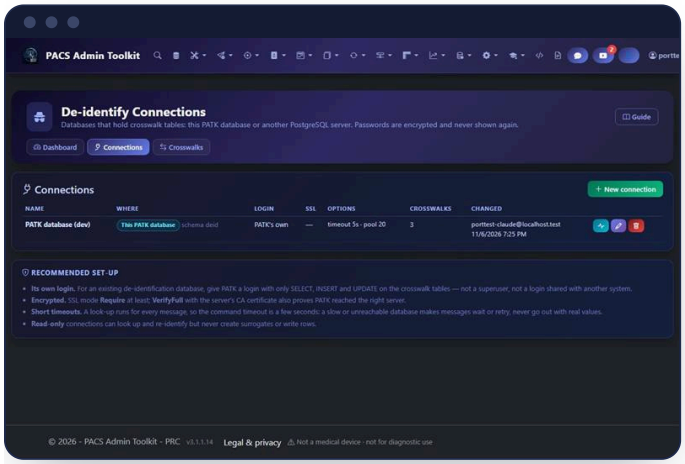
See it in action



The dashboard: every crosswalk with its look-ups, new surrogates, writes and errors



Attach an existing crosswalk table, start from a preset, or let PATK create one



Connections with encrypted passwords, SSL and short timeouts

USE CASES

Where it pays off

Studies to an AI vendor and results back

CHALLENGE	An imaging group sends studies to an AI vendor under surrogate patient IDs and accessions; results return with only the surrogates.
WITH PATK	The order crosswalk gives each patient and accession a surrogate on the first order; the AI-result match finds the one row by site, patient and accession surrogates and returns the real values.
OUTCOME	**Results land on the right study, and a result that matches two rows waits for a person instead of landing on the wrong patient.**

Taking over a crosswalk from an interface engine

CHALLENGE	A hospital network's crosswalk tables were filled by engine scripts for years; the network wants to retire the scripts without re-keying anyone.
WITH PATK	The existing tables are attached as they are; PATK reads the surrogates already there, fills new rows by the same keys and suggests the indexes the tables lacked.
OUTCOME	**Every existing patient keeps their surrogate, and the switch-over is a settings change rather than a data migration.**

A research export with short surrogates

CHALLENGE	A research team needs patient and accession surrogates short enough for DICOM fields that hold 16 characters.
WITH PATK	A PATK-managed crosswalk with prefixed counters (P, A, DR) and unique indexes on both directions.
OUTCOME	**Surrogates fit every field, can never collide, and stay the same for the same patient across exports.**

3T

Trust · Truth · Transform

Trust

Surrogates are created atomically, lookups fail closed and re-identification needs a permission and a reason.

Truth

The dashboard counts lookups, new surrogates, writes and errors per crosswalk; the audit log records every change and every re-identification.

Transform

Crosswalks that lived in engine scripts become settings that can be read, checked and tested in the browser.

VALUE MODEL

What it can be worth

An example imaging group sending studies to an AI vendor and matching results back.

Studies de-identified and matched back per year	40,000
Results needing manual matching without an exact crosswalk match	0.5% (200) at 15 minutes each
Interface-engine script maintenance avoided per year	40 hours
Blended staff cost per hour	$60
Duplicate surrogates found and fixed by hand per year without atomic creation	50 at 30 minutes each

$6,900

Estimated yearly value

Illustrative example. Your volumes, staffing and costs will differ; we will build this model with your own numbers.

DIFFERENTIATORS

Why it stands out

- **Concurrency handled**
  Surrogate creation is locked on the real key; it was tested with a hundred simultaneous first messages producing one surrogate.
- **Refuses rather than guesses**
  Ambiguous matches, missing indexes and unreachable databases are reported, not worked around.
- **Works with the tables you have**
  Existing crosswalks are attached without changing their structure.
- **Same platform**
  Connections, audit and permissions sit alongside routing, HL7 and DICOM tools.

SPECIFICATIONS

At a glance

Databases	PostgreSQL (another server or PATK's own, schema deid); SSL Require / VerifyCA / VerifyFull
Identifiers	Any number per crosswalk, keyed on one or more columns (site, patient ID, accession, names)
Surrogates	UUID, keyed hash (HMAC-SHA256), prefixed random pattern, counter, fake name
Controls	Upsert key, per-column write rules, re-identify matches, read-only connections, crosswalk on / off
Tools	Index advice with concurrent build / drop, try-it bench, dashboard counters, audit log
Rule steps	De-identify, Re-identify, Capture (HL7 Field Editor and DICOM Tag Editor); HL7 Drop message, Blank all fields except, contains any of; PHI leak guard
Profiles	Checklists for HL7 and DICOM, de-identify and re-identify, compiled into editor templates; capture field map; before / after preview
Packs	Five ready-made packs (AI vendor orders / reports / results, research DICOM export, generic AI vendor, teleradiology, teaching file); import and export
Parity check	Recorded interface-engine message exports compared field by field with what PATK would send; read-only; ignore list
Folder pick-up	Watched folders on the server, a share, or an SFTP / FTPS server; each new file de-identified or re-identified and delivered; never taken twice
Data API	Data sets outside systems query with an API key or token: chosen columns, required and fixed filters, rows per call and per day, allowed callers and networks; every call audited
Master Patient Index	With the MPI toolset: patient surrogates keyed by MPI ID (one surrogate per person, Exact and Confirmed links only) and re-identify to the person's ID for a chosen issuer

See De-identification with your own data

Ask for a guided demo or a trial license. PACS Admin Toolkit installs on your own Windows server in under an hour, and every module comes with step-by-step guides and in-app guided tours.

sales@3tconsults.com | support@3tconsults.com | +1 (612) 564-9194 | www.3tconsults.com

Not a medical device. PACS Admin Toolkit is not FDA-cleared or CE-marked and is not intended for primary diagnosis. It manages imaging and reporting data and workflows around the clinical systems that do.

Other product and company names may be trademarks of their respective owners. Screenshots show sample data. Value models are illustrative estimates, not guarantees.
© 2026 3T Consults and Coaching.