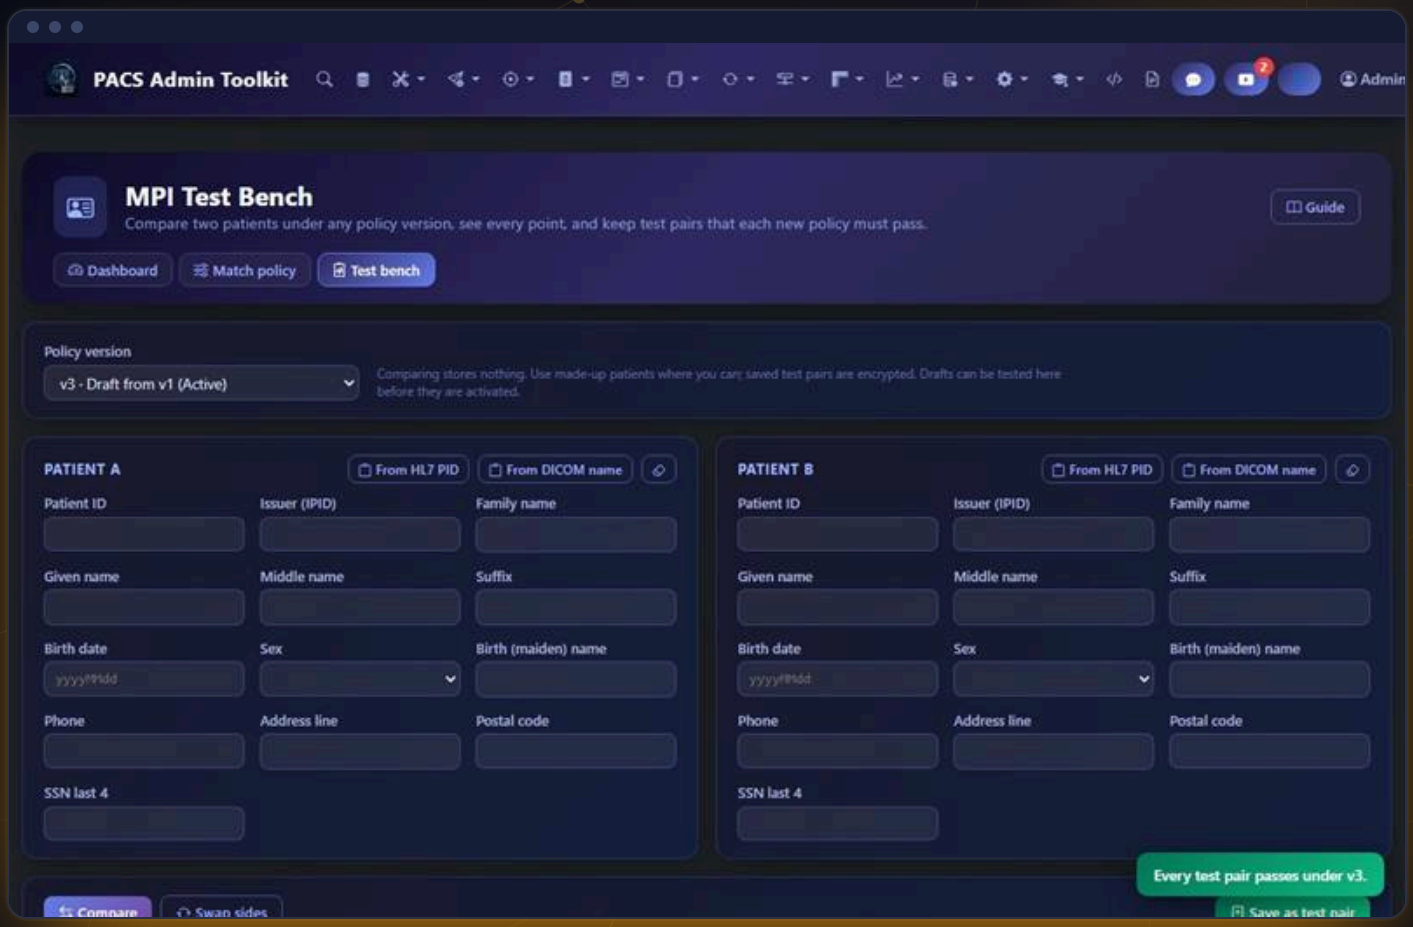
Platform & Integration | Leads with Truth

# Master Patient Index

Recognise one patient across facilities, IDs and systems, with the confidence and the reasons shown.




An index that collects the patients your DICOM, HL7, worklist and upload traffic already carries, and a matching engine that compares them by ID, issuer, names, birth date and more, sorts every comparison into a clear tier with a percent and plain-language reasons, and never links placeholders, twins or conflicting IDs on its own. Every rule lives in a versioned policy that is tested against known pairs before it goes live.

**100 %**
only when ID, issuer, names and birth date all agree

**Every point shown**
each field's outcome and score, not a black box

**Tested first**
a policy can't go live while a critical test pair fails

Replaces: **Spreadsheet look-ups, one-off matching scripts and guesswork across patient ID schemes**

License module `MPI`

THE CHALLENGE

## Why it matters

A patient seen at three facilities often has three patient IDs, a name spelled two ways and a birth date typed once with the day and month swapped. Prefetch, migrations and result routing all have to decide whether two records are the same person, and each tool decides differently.

Simple ID matching misses the same patient under another ID and links two people who share one. Hand-written matching scripts give a yes or no with no reason, can't be tested before a change, and leave nobody able to say why a link was made.

IN SHORT

An index that collects the patients your DICOM, HL7, worklist and upload traffic already carries, and a matching engine that compares them by ID, issuer, names, birth date and more, sorts every comparison into a clear tier with a percent and plain-language reasons, and never links placeholders, twins or conflicting IDs on its own. Every rule lives in a versioned policy that is tested against known pairs before it goes live.

CAPABILITIES

## What it does

- **Clear tiers**
  Exact, probable, possible, weak or no match, plus identity conflict (one ID on two people) and never link (a reviewer's decision), each with a percent and the reason in plain words.
- **Rules first, then points**
  Exact and probable matches are decided by rules; everything else is scored field by field: family and given name, birth date, sex, middle name, phone, address, postal code and SSN last 4.
- **Names compared the way people write them**
  Accents, apostrophes, titles and suffixes are set aside; sounds-alike and similar spelling are recognised; swapped names, birth names and nicknames count.
- **Dates compared the way people type them**
  Day and month swapped, one digit wrong or the same year only scores partly; placeholder and impossible dates don't count as agreement.
- **Guards against the classic mistakes**
  Placeholder patients (TRAUMA, UNKNOWN, DOE^JOHN), possible twins, IDs without an issuer, two IDs at one facility and records with too little to compare are capped and never linked on their own.
- **Issuer ID formats**
  Known patient ID patterns per issuer flag an ID that doesn't fit, so a mistyped ID can't make an exact match.
- **Versioned policies**
  Every threshold, weight and list lives in a numbered version; drafts are tuned and compared with the active one, and active versions never change.
- **Test pairs**
  Built-in synthetic pairs cover every tier and guard; your own pairs are stored encrypted with what the match should say. Critical pairs block any policy that breaks them.

### A test bench
Compare two patients under any version, paste an HL7 PID segment or a DICOM name, and see every point.

### Collects as it goes
DICOM listeners, HL7 listeners, the worklist and the uploader hand each patient to a background queue (observe only: nothing waits, nothing else changes). A back-fill adds the patients already stored.

### A golden view with its sources
Each person shows the values chosen from the most trusted source and where each came from, every source record with how it was linked, and the studies and orders seen for them.

### Records located, checked and pulled back
Each study stored or forwarded, each worklist order and each HL7 message becomes a record locator: where the record is and when it was last confirmed. Locators are re-checked on demand, when a person changes or on a schedule, within limits set per archive, and a study can be pulled back from an outside archive with a preview of what changes on the way. As the locator tables grow, check intervals are learned from how long records really stay put, each archive's check rate adapts to how it answers, and indexes are advised and built only in a quiet maintenance window.

### Review queues
Identity conflicts, probable links to confirm, possible matches, possible duplicates and placeholders wait in queues; each opens a side-by-side compare with every field coloured agree, partly, disagree or missing.

### Decisions with a reason
Confirm, link, merge, split or mark as different people; "different people" is a never-link the matcher respects from then on. Every decision can be undone, as a new recorded event.

### A patient's whole lifecycle
One timeline per patient across merged persons and records, exportable as CSV, with a tamper-evident hash chain that names the first entry changed, inserted or deleted.

### Second-person approval where you want it
Off by default; switch it on per decision at Global, facility or site level, with named approvers. A higher level is a floor the levels below can't remove.

### Retention the facility controls
History kept permanently by default; audit and requests for periods you set (never under a year), with federal and state reference periods and their sources beside the fields, legal holds and a permanent purge log.

### The MPI in your mappings
A Tag Editor or HL7 Field Editor step writes the MPI ID, the patient's ID for another issuer, all confirmed IDs or the golden demographics into images and messages. Only exact and reviewer-confirmed matches are written unless a step allows probable ones at a threshold, with a disclaimer, a reason and the name recorded; anything less is left, blanked or stopped.

### A rejected link stops at once
A reviewer's decision clears every cached answer, so no step uses a rejected link again.

### Exposure report
Every image and message a step changed is logged with its level; when a link is later rejected or split, the report lists what was sent using it, so the receiving systems can be corrected.

### Alerts
E-mail for new identity conflicts, decisions waiting for approval and review items over their SLA, with counts and links only. Config-impact alerts name the links your tools used that changed.

### Search like other systems do
Find a person by any mix of ID, issuer, name, birth date, phone, postal code or SSN last 4, with each candidate's level, warnings and points.

### Other systems can ask
HL7 v2 PDQ and PIX queries, FHIR Patient search, $match and PIXm, and a JSON API answer from the same matcher. Each answer says how sure it is; PIX answers carry exact and confirmed links only unless a caller is allowed probable ones with the disclaimer and a reason.

### Follows your registration system
When a sender merges, links or unlinks patients (HL7 ADT A40, A47, A24, A37 and the older merge events), the MPI sends it to review or, for senders you trust, merges or splits at once. Second-person approval still applies, and a never-link, a placeholder, an identity conflict or records that don't agree always go to a reviewer. A PIX identity feed listener takes registrations and merges straight from registration systems, and a merge log shows what happened to each one.

### Only known callers, only their domains
Each caller is an API key, API user or HL7 sender with its own interfaces, issuer domains, lowest level and daily limit; every query is logged without values and appears on the person's lifecycle.

### Priors under every ID
Prefetch workflows, the priors browser, the viewer's comparisons and HL7 fetch can ask the source for the patient's other confirmed IDs, keep priors the MPI links to the same person and leave out priors it knows belong to someone else. Off by default; a shadow mode reports what would change before anything does, and if the MPI can't answer the tool runs as before.

### One identity everywhere it travels
Migrations can rewrite old Patient IDs to the new system's IDs straight from the MPI, with a pre-flight list of every patient it can't map confidently. The worklist can carry each patient's other confirmed IDs, de-identification can give one surrogate per person, patient merges show what the MPI knows, and protections suggest the patient's other IDs. Only Exact and Confirmed links rewrite or key anything unless an administrator allows more, with a reason.

### Tune with evidence
Before a draft policy goes live, a dry run scores the matches already made under it and shows how many would newly link, stop linking or change tier. Reviewers' decisions suggest each weight, common names count for less, open review items are re-scored when a policy is activated, and any earlier version can be brought back.

**Audited changes**
Activating a policy needs a permission, a reason and passing tests; the change list is written to the audit log.

WORKFLOW

# How it works

1
**Start from calibrated defaults**
The first policy version is set from a survey of real patient data: how often names, birth dates and IDs agree.

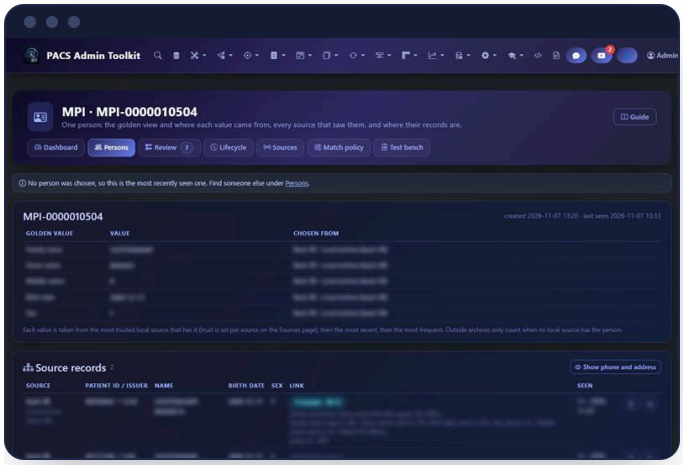
2
**Try it**
Compare known patients on the test bench and save the results you want every policy to keep.

3
**Tune a draft**
Change thresholds, weights, placeholders or ID formats and see the test pairs re-run, then dry run it on the matches already made.

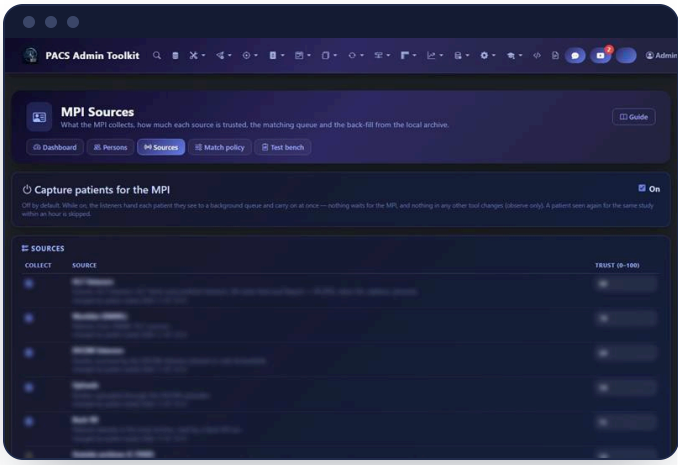
4
**Activate with a reason**
The old version is retired in the same step and the change is audited.

IN THE PRODUCT

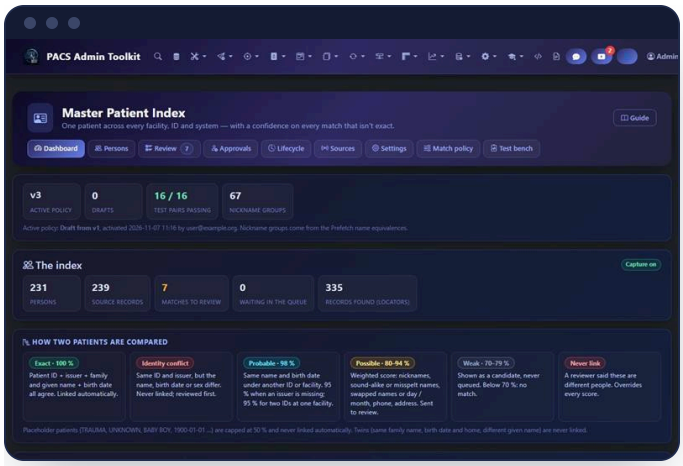
# See it in action



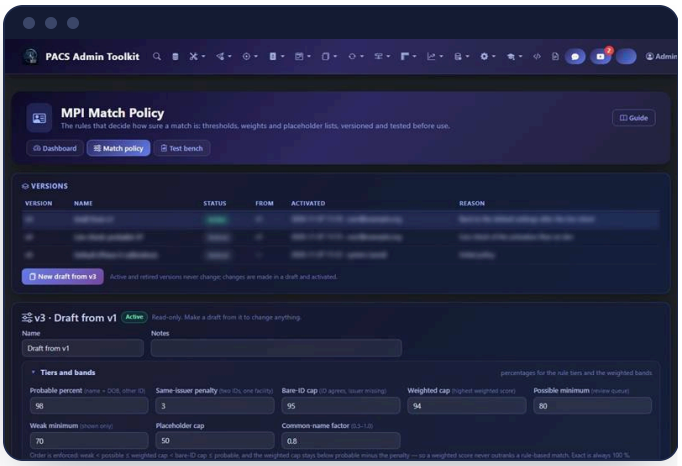
One person: the golden view and its sources, how each record was linked, and where their studies are



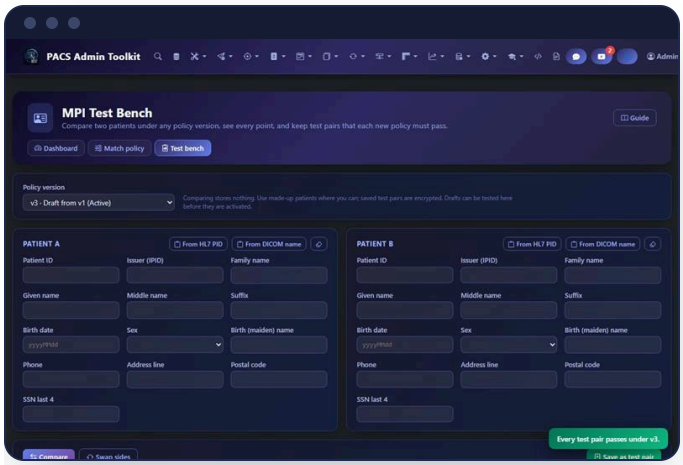
Capture switches and trust per source, the matching queue and the back-fill




The dashboard: the active policy, drafts waiting, test pairs passing and how the tiers work




Policy versions with thresholds, weights, placeholders and ID formats




The test bench: two patients compared, with the tier, warnings and every field's points

USE CASES

# Where it pays off

## Prefetch across facilities

| | |
|---|---|
| CHALLENGE | A radiology group reads for four hospitals, each with its own patient ID scheme; priors from the other hospitals are often missed. |
| WITH PATK | The matcher recognises the same name and birth date under another facility's ID as probable, with its percent and reasons. Prefetch, run in shadow mode first, then asks each archive for the patient under every confirmed ID. |
| OUTCOME | **Radiologists see the likely priors with an honest confidence, instead of only those that share an ID.** |

### Cleaning up before a migration

| | |
|---|---|
| CHALLENGE | An archive migration has to merge patients from two systems without linking two people by accident. |
| WITH PATK | Test pairs taken from the site's known problem cases are made critical, so no policy change can link them. The MPI crosswalk rewrites each image's old ID to the new system's ID as it migrates, and its pre-flight lists every patient that would be left unchanged and why. |
| OUTCOME | **The matching rules are agreed and tested, and the ID mapping reviewed, before any data moves.** |

### One ID on two people

| | |
|---|---|
| CHALLENGE | A registration error put two patients under one ID. |
| WITH PATK | The same ID and issuer with a different birth date is an identity conflict: never linked, shown first. |
| OUTCOME | **The error is caught instead of mixing two patients' studies.** |

3T

## Trust · Truth · Transform

### Trust

Only exact and probable matches (same name and birth date) link on their own, probable ones flagged for confirmation; placeholders, twins and conflicting IDs never do.

### Truth

Every comparison shows its tier, percent, warnings and each field's points; every policy change is audited with its reason.

### Transform

Matching rules that lived in scripts become a versioned, tested policy that can be read and tuned in the browser.

VALUE MODEL

## What it can be worth

An example radiology group reading for several facilities with separate patient ID schemes.

| | |
|---|---|
| Studies read per year | 150,000 |
| Studies whose priors sit under another facility's ID | 5% (7,500) |
| Staff time to find those priors by hand | 4 minutes each (500 hours) |
| Share of that search the matcher's candidates remove | 60% (300 hours) |
| Blended staff cost per hour | $60 |

**$18,000**

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

- **Explains itself**
  Every result shows why, field by field, in plain words.
- **Conservative by default**
  Only exact and probable matches link automatically, probable ones flagged for confirmation; everything else is a suggestion with its confidence.
- **Tested before it changes**
  Policies can't be activated while a critical test pair fails.
- **Same platform**
  Permissions, audit and licensing sit alongside prefetch, routing, HL7 and DICOM tools.

SPECIFICATIONS

# At a glance

| | |
|---|---|
| **Tiers** | Exact, Probable, Possible, Weak, No match, Identity conflict, Never link |
| **Fields** | Patient ID and issuer, family / given / middle name, suffix, birth name, birth date, sex, phone, address, postal code, SSN last 4 |
| **Name handling** | Accents and titles removed, suffixes kept apart, NYSIIS sounds-alike, Jaro–Winkler spelling, swapped names, nicknames |
| **Guards** | Placeholder names and dates, twin guard, ID without issuer, same issuer with different IDs, sparse data, issuer ID formats |
| **Policy** | Numbered versions (draft, active, retired), thresholds, weights, options, change list against the active version; dry run against existing matches with a percent chart; weight suggestions from review decisions; common-name table refreshed daily; open review items re-scored on activation; rollback to any earlier version |
| **Testing** | Built-in synthetic pairs, encrypted saved pairs with expected outcomes, critical pairs block activation, test bench |
| **Capture** | DICOM listeners (stored or forwarded), HL7 listeners, worklist, uploader, back-fill, outside archives (opt-in); master switch, per-source switches and trust, per-listener overrides |
| **Person page** | Golden view with provenance, source records with tier / percent / points, identifiers, matches to review, record locators, events |
| **Record locators** | Checks of this archive, the worklist and DICOM archives (study C-FIND) on demand, after person changes and on schedules; per-archive rate, concurrency, hours, night window, daily budget and circuit breaker; pause switches; Gone only after repeated misses over days; pull-back by C-MOVE through the DICOM queue to allowed destinations with a template preview; suppress, re-point, note and add by hand with a reason; check log and daily counts without patient values |
| **Locator health** | Learned check intervals per archive from the still-there curve (pinnable), adaptive check rate, answer times, table and index statistics, index advice with scripts and a guarded automatic builder (maintenance window, quiet-database checks, one partition at a time, never drops) |
| **Review** | Queues for conflicts, probable, possible, duplicates, placeholders, re-scored, deferred; side-by-side compare; confirm / link / merge / split / never link with a reason; undo |
| **Lifecycle** | Timeline across merged persons and records, CSV export, SHA-256 hash chain check |
| **Approval** | Optional second person per decision (confirm, reject, link / merge, split, reverse never-link, undo, policy activation, retention) at Global / facility / site, approver roles or users, self-approval off, expiry |
| **Retention** | Per kind, permanent or 1 year and up; federal and 50-state reference periods with citations; legal holds; oldest-first purge with a permanent log |
| **Editor steps** | Tag Editor and HL7 Field Editor MPI lookup: match on ID, issuer, name, birth date; write MPI ID, ID for an issuer, all IDs (CX repetitions), golden demographics, level; Exact + Confirmed by default, Probable per step with threshold, disclaimer and reason; leave, blank, stop or hold |

| | |
|---|---|
| **Usage** | Every real use logged with level, MPI ID and message; values encrypted; exposure report and CSV per rejected or split link; daily counts kept permanently |
| **Alerts** | E-mail: identity conflicts, approvals waiting, daily SLA digest, config impact; no patient values in e-mails |
| **Search** | In-app search by MPI ID, patient ID (issuer optional), name, birth date, sex, phone, postal code, SSN last 4; level, warnings and score breakdown per candidate |
| **Priors and fetch** | Prefetch workflows, priors fetch profiles, viewer comparisons and HL7 fetch: Off, Shadow (report only) or On; other IDs queried with their issuer, same-person priors kept, other-person priors left out, level badges in the viewer and priors browser; Exact + Confirmed, Probable 98 % opt-out per config; fails open; per-run report without patient values |
| **Query interfaces** | HL7 v2 PDQ ($QBP^{Q22}$ / $^{RSP}$K22) and PIX query ($QBP^{Q23}$ / $^{RSP}$K23) over MLLP; FHIR Patient search, $match and PIXm; JSON resolve / lookup; registered domains (namespace, OID, FHIR system); callers by API key, API user or MSH-3 / MSH-4 / address, with interfaces, domains, lowest level, limits; query log on each person's lifecycle |
| **HL7 merges and feeds** | ADT merges (A40, A39, A34, A30, A18, A36), ID changes (A47, A46), links (A24) and unlinks (A37) from HL7 listeners, worklist HL7 sources and the PIX feed; Off / Review / Apply per source type and per sender; never-link, placeholder, identity-conflict and agreement checks; second-person approval; PIX identity feed listener (IHE ITI-8, MLLP, ACK AA / AE / AR); merge log without patient values |
| **Migrations, worklist, De-ID** | Migration Tag Changes MPI crosswalk (Patient ID + issuer to the new issuer, pre-flight list, Probable only when allowed with a reason); worklist Other Patient IDs Sequence with optional MPI ID; De-ID surrogate keyed by MPI ID and re-identify to a chosen issuer; patient-merge advice with optional MPI merge; related-ID suggestions in Quick Protect |
| **Audit** | Policy activation with reason and change list; test pair changes; capture and source switch changes; interface, domain and caller changes; phone and address views; every decision, lifecycle view and export |

## See Master Patient Index 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.