



INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

DICOM De-Identification for Clinical Trials: Profile Guide

Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-10-20 · dicom de-identification · clinical trials · dicom ps3.15 annex e · attribute confidentiality profile · dicom anonymization · burned-in pixel data · hipaa safe harbor · imaging validation

Original article: <https://intuitionlabs.ai/articles/dicom-deidentification-clinical-trials-profile-validation>



In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes and Decision Points in Annex E
4. Pixel, Feature, UID, Date and Trial-Field Decisions
5. Implementation Considerations and Process Changes
6. Data Analysis and Evidence
7. Implications and Future Directions
8. Frequently Asked Questions (FAQs)
9. Conclusion

Executive Summary

DICOM de-identification for clinical-trial imaging is best treated as a protocol-specific selection problem: choose the PS3.15 Annex E options that match the recipient and the endpoints, then prove with tests that identity was removed and the quantitative signal survived. The edition banner on the current Annex E read **DICOM PS3.15 2026d** on 2026-10-20 (dicom.nema.org). Annex E defines one **Basic Application Level Confidentiality Profile** and **12 options**, recorded with CID 7050 codes **113100 to 113112** (dicom.nema.org).

Annex E is explicit about its limits: using the profiles does not guarantee removal of all identifying information (dicom.nema.org), and the Basic Profile leaves pixels alone unless a pixel option is added (dicom.nema.org). Empirical evidence agrees. The NCI Medical Imaging De-Identification (MIDI) project found that no automated method reached a 100% success rate (datacommons.cancer.gov), and its coordinators concluded that some human review remains necessary (datacommons.cancer.gov). In facial-recognition work, correct matching of head MRI to photographs rose from **83%** of 84 participants (pmc.ncbi.nlm.nih.gov) to **97%** of 157 (adni.loni.usc.edu).

The legal frame supplies no numeric pass mark. HHS states that there is no explicit numerical level of identification risk that meets the Expert Determination standard (www.hhs.gov), and that correctly applied methods leave a risk that is very small but not zero (www.hhs.gov). The IHE De-Identification Handbook likewise offers no universal threshold (profiles.ihe.net). This report therefore defines eight measurable test metrics and leaves every acceptance threshold to the protocol's privacy and imaging-science reviewers.

Utility is the second half of the decision. FDG-PET standardized uptake value (SUV) analysis needs patient weight (qibawiki.rsna.org), administered dose and injection time, and Annex E warns that date transformation must not disrupt the injection-to-acquisition comparison (dicom.nema.org). A repository that documents its own practice cautions that adding options raises the chance of residual PHI and demands a more rigorous process (wiki.cancerimagingarchive.net). The deliverables below are an option-by-use-case matrix (Table 1), a synthetic corpus specification (Table 2), an evidence summary (Table 3) and a release gate. A clean sample bounds risk only loosely: at 95% confidence, n clean items imply an upper bound near $3/n$ (editorial arithmetic).

Introduction and Background

Annex E of DICOM PS3.15 addresses the removal and replacement of attributes within a DICOM data set that may leak individually identifiable information (dicom.nema.org). It is a toolbox, not a process: the standard says use of the profiles does not replace a de-identification process (dicom.nema.org). The IHE De-Identification Handbook describes the process for removing individually identifiable information from healthcare data (profiles.ihe.net). This report sits between the two. It translates Annex E into a protocol-specific profile and options decision, plus a test pack that shows whether privacy controls damaged longitudinal or quantitative utility.

The audience is imaging core labs, trial data managers, privacy and security officers, platform engineers and validation teams. This reference is tool-neutral: it does not rank, compare or endorse software. It is an implementation reference, not legal advice, and it does not claim that any draft or voluntary standard is mandatory. Three labels are used throughout. **Requirement** means the cited source uses mandatory language. **Recommendation** means a cited source advises it. **Editorial choice** means this report's own suggestion, which no source requires.

Purpose, recipient and threat model

Every option decision starts with who receives the data and what they could join it to. Four recipient classes recur in trials:

- **Blinded central readers.** FDA guidance expects primary-endpoint image readers in randomized trials to be blinded to treatment assignment (www.hhs.gov), and in some diseases also to acquisition date (www.hhs.gov).
- **Sponsors and contract research organizations.** They usually need traceability to a subject and time point, which pushes toward keeping trial identifiers and sometimes UIDs.
- **Public or controlled-access repositories.** These face the widest set of unknown linkage sources, so conservative defaults dominate.
- **Algorithm developers.** They often need acquisition parameters held in private attributes, which pushes toward a safe-private retention option.

The threat model should name linkage sources, not only identifiers. One general study estimated that 99.98% of Americans could be correctly re-identified in any dataset using 15 demographic attributes (pmc.ncbi.nlm.nih.gov). That result is not imaging-specific, but it explains why dates, ages, sites and device details are assessed together.

The US legal frame in one paragraph

US HIPAA offers two de-identification methods. Expert Determination requires a qualified person to determine that the risk of identification by an anticipated recipient is very small (www.ecfr.gov). Safe Harbor removes listed identifiers, including full face photographic images and any comparable images (www.ecfr.gov), device identifiers and serial numbers (www.ecfr.gov), and all date elements except year for dates directly related to an individual (www.ecfr.gov). A catch-all covers any other unique identifying number, characteristic or code

(www.ecfr.gov), which HHS guidance applies to clinical trial record numbers (www.hhs.gov). Whether a DICOM release satisfies either method is a legal and statistical question for privacy counsel and a qualified expert; Annex E conformance alone does not answer it.

Key Changes and Decision Points in Annex E

From Supplement 142 to the current edition

The option structure originates in DICOM Supplement 142, prepared by the DICOM Standards Committee's Working Group 18 for clinical trials (www.dicomstandard.org). Its rationale was that it is impractical to leave decisions about what to retain or remove to individual sites or trials (www.dicomstandard.org). The IHE handbook's reference list annotates Annex E as covering clinical trials, double-blinding, traceability to original content and preserving data needed for the trial (profiles.ihe.net), and its DICOM example notes that many trial plans require these standard methods (profiles.ihe.net).

Two version cautions apply. First, the DICOM edition changes several times a year; the banner read 2026d at access, and earlier 2026 editions may differ, so the exact release used must be pinned in the protocol. Second, implementations evolve. One public repository documents the Basic Profile amended by options that include Clean Pixel Data, Clean Descriptors and Retain Longitudinal With Modified Dates (wiki.cancerimagingarchive.net), while the NCI MIDI pipeline started from the Basic Profile with the Clean Descriptors option (datacommons.cancer.gov) and the NCI EDNRN process recommends the Basic Application Confidentiality Profile (edrn.nci.nih.gov). Those are program-specific choices, not statements about what other trials should do, and repository documentation changes over time, so the option set must be read from the version in force on the date of use.

The Basic Profile and the action codes

The Basic Application Level Confidentiality Profile is intended for use in clinical trials and other scenarios (dicom.nema.org). It is extremely conservative: it removes all information related to patient identity and demographics, responsible parties, personnel and organizations, along with match-enabling data such as UIDs, dates and times (dicom.nema.org). Table E.1-1 assigns each attribute an action code, and reading the codes correctly is the foundation of an attribute-action audit:

- **D and Z** replace a value: D with a non-zero length dummy, Z with a zero length value or a VR-consistent dummy.
- **X and K** remove or keep: X removes the attribute and, for a sequence, all items; K keeps it unchanged, or cleaned for sequences. Together, D, Z, X and K cover the basic replace, remove and keep actions (dicom.nema.org).
- **C and U** clean and remap: C substitutes values of similar meaning known not to identify; U replaces a UID with a non-zero length UID that is internally consistent (dicom.nema.org).

Even a fully conforming run leaves residual duties, because the de-identifier bears responsibility for removing all identifying information, including from private attributes, retired attributes and free text (dicom.nema.org).

Option-by-option decision matrix

Table 1 below is the reference asset of this report. It lists the Basic Profile and the 12 options with a typical protocol trigger, the effect on retained, removed or replaced content, the primary test from the pack defined later (T1 to T10), and the residual risk to document. The triggers and risk notes are editorial choices drawn from the cited sections, not requirements.

Annex E element (CID 7050 code)	Typical protocol trigger (editorial)	Effect on content	Primary test	Residual risk to document
Basic Profile (113100)	Always the base	Removes or replaces identity, personnel, organization, UIDs, dates, private attributes; pixels untouched	T1, T2, T10	Free text, pixels and recognizable features remain
Clean Pixel Data (113101)	Any modality that can burn in text: ultrasound, secondary capture, screenshots	Changes stored pixel values; overlaying a graphic is not sufficient (www.dicomstandard.org)	T3	Operator review may be required (www.dicomstandard.org)
Clean Recognizable Visual Features (113102)	Head and neck CT or MRI; 3D reconstructions	Removes or distorts pixels so recognition is prevented	T4, T7	May render pixel data unusable for its purpose (www.dicomstandard.org)
Clean Graphics (113103)	Overlays, curves, annotations in group 60xx	Cleans overlay and graphic content outside Pixel Data	T3	Overlay groups other than 6000 (up to 601E) missed
Clean Structured Content (113104)	Structured reports, measurements, acquisition context	Cleans identifiers in content items	T8	Names inside nested text
Clean Descriptors (113105)	Descriptions and comments authored by technologists	Cleans operator free text; may keep technique text but drop diagnosis that could bias reads (www.dicomstandard.org)	T8	Over-removal of useful descriptions
Retain Longitudinal, Full Dates (113106)	Rarely; only with a legal pathway for dates	Keeps dates and times	T6	Dates are match-enabling data
Retain Longitudinal, Modified Dates (113107)	Response assessment, serial quantitative reads	Transforms dates but preserves gross longitudinal relationships (www.dicomstandard.org)	T6, T7	Shift key or mapping becomes sensitive
Retain Patient Characteristics (113108)	Dose, SUV or body-size normalization	Keeps age, sex, height, weight	T7	Quasi-identifier combinations
Retain Device Identity (113109)	Scanner-specific correction or qualification	Keeps device values, cleaning embedded titles	T1, T7	Serial numbers can point to a site
Retain UIDs (113110)	Audit trail back to originals is required (www.dicomstandard.org)	Keeps instance UIDs	T5	Enables matching to originals

Annex E element (CID 7050 code)	Typical protocol trigger (editorial)	Effect on content	Primary test	Residual risk to document
Retain Safe Private (113111)	Vendor acquisition parameters needed for analysis	Keeps only privates known to be safe	T2, T7	Opaque vendor blocks
Retain Institution Identity (113112)	Site-level analyses or site-specific QA	Keeps institution identity fields	T1	Site re-identification

The matrix has two practical readings. A protocol that needs only change over time at a single endpoint typically combines the Basic Profile with Modified Dates and a pixel option. A protocol with vendor-specific quantification adds Retain Safe Private and Retain Patient Characteristics, and accepts the extra review burden, since each added retain option widens the surface where PHI can persist (wiki.cancerimagingarchive.net). The IHE handbook observes that for DICOM the default profile plus selected options commonly suffices, but that private attributes need design-specific treatment (profiles.ihe.net).

Header attributes, private tags and the file preamble

Header work has three distinct surfaces, each with its own test.

State tags and preamble. A de-identifier adds CID 7050 codes for the profile and options to De-identification Method Code Sequence (0012,0064) (dicom.nema.org). Patient Identity Removed (0012,0062) asserts removal from both attributes and pixel data (dicom.nema.org), so a header-only run should not set it without a pixel decision.

Private attributes. Without the Retain Safe Private option, all private attributes are removed (dicom.nema.org). A removed block may keep or lose its Private Creator element (dicom.nema.org), so a test should not treat a leftover creator as a failure by itself. Bulk binary data in OB elements is particularly unsafe (dicom.nema.org), vendor blocks can mix text and binary components (arxiv.org), and some ultrasound vendors use large opaque binary private elements whose safety may be difficult to ascertain (arxiv.org). Blanket deletion has a cost: a common practice of deleting all private elements removes scientifically useful data as well as PHI (pmc.ncbi.nlm.nih.gov), while retaining safe private information greatly enhances utility and reuse (arxiv.org).

Free text. Operator-entered text may contain identities, which is why Clean Descriptors exists. The NCI evaluation team treated a partially de-identified person name as more serious than a partially de-identified place name (datacommons.cancer.gov). Header tests therefore need both residual-PHI and false-removal counts.

Pixel, Feature, UID, Date and Trial-Field Decisions

Burned-in pixels, overlays, graphics and recognizable features

The Basic Profile does not address information in the pixels (dicom.nema.org). With Clean Pixel Data, the stored pixel values must actually be changed (blacked out) (dicom.nema.org). If Burned In Annotation (0028,0301) is absent, the image may or may not contain burned-in annotation, so absence is not evidence of a clean image (dicom.nema.org).

Overlays are a second channel: sixteen separate overlay planes may be associated with one image (dicom.nema.org), so a review that opens only group 6000 is incomplete. Clean Graphics covers identifying information encoded in graphics, text annotations or overlays (dicom.nema.org).

Modality drives prevalence and method. Computed tomography does not normally contain burned-in annotation, whereas ultrasound often does (www.dicomstandard.org). The MIDI task group states that ultrasound images frequently, if not always, contain burned-in identification (arxiv.org), and recommends stratified review targeting higher-risk images such as ultrasound and objects with secondary capture or screenshot characteristics (arxiv.org). In the NCI evaluation, header information flagging a high risk of burned-in content was used to quarantine images for pixel review (datacommons.cancer.gov). Fully automatic removal still requires a human in the loop due to an excessively high failure rate (arxiv.org).

Triage can be efficient. One published validation on 415,182 images across ten modalities (pmc.ncbi.nlm.nih.gov) needed optical character recognition (OCR) for only 6050 images, or 1.5% of the dataset (pmc.ncbi.nlm.nih.gov). In the MIDI-B challenge, half of the submissions failed one pixel action (arxiv.org), and the authors noted only a very small number of burned-in images existed, so pixel scoring rested on few cases (arxiv.org).

Recognizable features are a different problem. The original supplement warns that cleaning them may render pixel data unusable for the collected purpose (www.dicomstandard.org). The risk is measured: face recognition software chose the correct MRI for 70 of 84 participants (83%) in 2019 (pmc.ncbi.nlm.nih.gov), the same group later matched 97% of 157 scans (adni.loni.usc.edu), and after popular defacing software 28% to 38% of scans still allowed matching (adni.loni.usc.edu). The first study's authors concluded that removing only metadata may be insufficient (pmc.ncbi.nlm.nih.gov). Defacing has analytic cost: a multi-algorithm study of 300 scans found success rates varied strongly between defacers (pmc.ncbi.nlm.nih.gov), with some outcome measures differing by as much as 5.4% (pmc.ncbi.nlm.nih.gov); another found small systematic biases that produced significantly different brain measures (pmc.ncbi.nlm.nih.gov); and an eight-algorithm comparison warned that rare catastrophic segmentation failures can occur (pmc.ncbi.nlm.nih.gov). Face is not the only biometric channel: a verification model on chest radiographs could reveal the same person ten and more years after the initial scan (pmc.ncbi.nlm.nih.gov). For head and neck imaging the decision is a measured trade between recognition risk and analysis bias, tested on the protocol's own algorithms.

UIDs, dates and longitudinal consistency

UIDs. A conforming de-identifier must always protect the SOP Instance UID and references to other instances (dicom.nema.org), with replacement values kept consistent across related instances. Retain UID is justified only when the risk of reaching the originals through UIDs is judged small (dicom.nema.org). One repository

builds replacement UUIDs from its own root plus a hashed value of the original UUID (wiki.cancerimagingarchive.net), which keeps replacement deterministic for a given input.

Dates. Two mutually exclusive options address longitudinal needs: Full Dates retains dates and times, and Modified Dates transforms them while preserving gross longitudinal relationships (dicom.nema.org). A consistent shift across separate runs requires a mapping or database (dicom.nema.org). Three constraints follow:

- **Intervals, not absolute dates, carry utility.** One public implementation normalizes dates to January 1, 1960 and offsets them relative to a registration date (wiki.cancerimagingarchive.net); a study two days after registration becomes 01/03/1960 (wiki.cancerimagingarchive.net). Retained private date tags must receive the same offset as the standard tags (wiki.cancerimagingarchive.net), and dates in other free-text tags are removed (wiki.cancerimagingarchive.net).
- **A zero shift is a defect.** A randomly chosen per-patient shift could be zero days, which would leave dates unchanged. Per-patient random shifts are a published approach (pubmed.ncbi.nlm.nih.gov); a shift must still be bounded away from zero for every subject.
- **Safe Harbor is stricter than Annex E.** HHS says elements of dates not permitted include day and month (www.hhs.gov), and dates tied to test measures are dates too (www.hhs.gov). A Modified Dates release therefore needs Expert Determination, or a limited data set with a data use agreement (a limited data set is PHI that excludes enumerated direct identifiers (www.ecfr.gov)). Counsel must decide which pathway applies.

Quantitative timing. Supplement text says time transformation must not disrupt computations (www.dicomstandard.org), and the task group adds that intervals between series time and radiopharmaceutical start must be preserved (arxiv.org). An older QIBA draft quantifies the sensitivity: an unhandled daylight-saving shift yields a 32% SUV error (qibawiki.rsna.org), and a 15-minute clock mismatch an avoidable calibration error of 10% (qibawiki.rsna.org).

Clinical-trial identification, consent and distribution fields

Subject identification in a trial generally requires a combination of sponsor, protocol, subject or reading identifier and site (dicom.nema.org). Clinical Trial Subject ID may, but need not, equal Patient ID (dicom.nema.org), so a pseudonymization plan must treat the two as separate attributes. For other trial identifiers, read the Table E.1-1 row for each attribute in the release used and record the action chosen in the specification; this report does not assume them.

Consent fields describe permission, not privacy status. Distribution Type is required if the distribution flag equals YES or WITHDRAWN (dicom.nema.org). The standard defines no semantics for what action is appropriate when consent is withdrawn (dicom.nema.org), so a withdrawal workflow is a protocol rule, not a DICOM behavior. Blinding adds a constraint: if a reader should not know the total number of time points per subject (www.hhs.gov), time-point labels, offsets and series counts need review as potential leaks even when no identifier remains.

Implementation Considerations and Process Changes

Utility-preservation tests

Start from the imaging charter, which FDA guidance says should describe image alteration procedures performed at the site, with removal of subject-identifying information as an example (www.hhs.gov). From the charter, build a **protocol-required feature list**: every attribute, pixel region, timing relation and identifier link that an endpoint needs.

The list can be short and still sensitive. QIBA's FDG-PET/CT profile requires patient weight in (0010,1030) (qibawiki.rsna.org), dose in (0018,1074) (qibawiki.rsna.org) and injection start in (0018,1078) (qibawiki.rsna.org). Its claims are a within-subject coefficient of variation of 10 to 12% for SUVmax (qibawiki.rsna.org) and a true-change threshold of a 39% or greater increase or 28% or greater decrease (qibawiki.rsna.org). RECIST 1.1 requires baselines never more than 4 weeks before treatment start (dctd.cancer.gov), the same method and technique across visits (dctd.cancer.gov), and progression defined as at least a 20% increase in the sum of diameters (dctd.cancer.gov), all of which depend on correct ordering and intervals of studies.

Device fields need a deliberate choice. Common practice keeps manufacturer and model but not serial number (arxiv.org). Inferring institution identity from device databases is described as a theoretical concern at present (arxiv.org), while HHS treats the production-identifier portion of a unique device identifier as not permitted in Safe Harbor data (www.hhs.gov).

The test pack is a set of ten tests:

- **T1 Header attribute-action audit.** Compare every attribute in a corpus inventory with the declared action.
- **T2 Private-block audit.** Record a disposition for every private block observed.
- **T3 Pixel and overlay review.** Review pixel frames selected by risk stratum and all 60xx overlay planes.
- **T4 Recognizable-feature check.** Render volumes and test with the face-matching method the threat model names.
- **T5 UID referential integrity.** Rebuild the reference graph and count broken links and collisions.
- **T6 Date order and interval.** Compare study order and intervals before and after.
- **T7 Protocol-feature retention.** Re-run endpoint computations on the corpus and compare.
- **T8 Descriptor and structured-content scan.** Scan free text for names, IDs and dates.
- **T9 Consent and distribution fields.** Verify presence, value domain and withdrawn-consent handling.
- **T10 State tags and preamble.** Verify Patient Identity Removed, method codes, preamble and group 0004.

Synthetic test corpus specification

A corpus with seeded identifiers and an answer key turns review into measurement. One published design builds an evaluation dataset of 1,693 CT, MRI, PET and digital X-ray images (www.ncbi.nlm.nih.gov) from 21 patients, 22 studies and 26 series (www.ncbi.nlm.nih.gov), pairs it with a curated de-identified version, and validates with a per-tag answer key and an evaluation script (www.ncbi.nlm.nih.gov). The later MIDI dataset

extends the design to 538 subjects, 605 studies and 708 series (arxiv.org), and NCI publishes answer keys, mapping files and a validation script (datacommons.cancer.gov). The task group advises special attention to private data elements in synthetic data (arxiv.org).

Table 2 below specifies a protocol-level corpus. Counts are design rules, not universal minimums.

Stratum	Specification (editorial)	Seeded content	Expected action	Tests
Header text	Every text VR present in the protocol's modalities	Names, IDs, dates, phone and address strings in free text and codes	Removed or cleaned per option	T1, T8
Private blocks	At least one block per vendor and model in scope, plus one unknown block	Text, numbers, UIDs, dates, opaque binary	Keep only blocks on a documented safe list	T2
Overlays and graphics	Overlay groups other than 6000 (up to 601E) populated	Bit-mapped text in several planes	Cleaned or removed	T3
Burned-in pixels	Each risk stratum, including ultrasound and screenshots	Text in header band, corner and mid-image	Pixels changed, not covered	T3
UID graph	Multi-series study with referenced objects and a re-run of the same study	Cross-references, shared frame of reference	Same input gives same replacement	T5
Longitudinal set	At least three time points per subject, including same-day pairs and a year boundary	Known intervals	Order and interval preserved	T6
Quantitative set	One case per endpoint algorithm	Weight, dose, times, vendor privates	Endpoint within the protocol's tolerance	T7
Trial fields	Consent states YES, NO, WITHDRAWN	Sponsor, protocol, subject, site, time point	Handled per the specification	T9, T10
Structured content	Reports with nested content items	Names inside text values	Cleaned	T8
Head series	Volumetric head scans if in scope	Real or phantom faces under approved handling	Recognition prevented or risk accepted	T4

The table is a template for the validation plan. Rows that do not apply, for example head series, should be removed with a written reason rather than left blank.

Residual-risk sampling and escalation

Eight metrics make results comparable across runs. Each is an editorial definition, not a DICOM or regulatory requirement.

- **Attribute-action coverage.** Attributes with a documented and verified action divided by attributes observed in the corpus inventory.

- **Private-tag review coverage.** Private blocks with a recorded disposition divided by private blocks observed.
- **Pixel and overlay review coverage.** Frames and overlay planes reviewed divided by those selected by the sampling plan.
- **UID-consistency defects.** Broken references plus replacement collisions, counted per run and per 1,000 UID references.
- **Date-order preservation.** Study pairs with order and interval preserved divided by pairs tested.
- **Protocol-required feature retention.** Required features retained with valid values divided by required features tested.
- **Residual-PHI findings.** Count by severity, with names weighted above places as the NCI team did (datacommons.cancer.gov).
- **False-removal findings.** Count of useful content removed or altered, such as the numeric strings altered by one heuristic (datacommons.cancer.gov).

Sampling must be stratified: review high-risk strata fully, sample the rest, and record the sampling frame. By the rule of three, if n independent items are reviewed with zero findings, the 95% upper confidence bound on the true finding rate is about $3/n$ (editorial arithmetic).

(Hypothetical Example) A team reviews 300 randomly chosen non-ultrasound images from a low-risk stratum and finds no burned-in text. The upper bound is about $3/300$, or 1.0%. That does not show the rate is zero, and it says nothing about the ultrasound stratum, which is reviewed exhaustively. If 5 findings had occurred, the team would stop the run, widen the sample, find the cause, fix the rule and repeat from T3.

Escalation rules are editorial: any residual-PHI finding in pixels or a name field stops release for that batch; a UID or date defect stops release for affected subjects; false-removal findings above the protocol's tolerance trigger option review. Two operational controls from the NCI evaluation are worth copying: verify that output file count equals input file count (datacommons.cancer.gov), and regression test the pipeline before deploying a new release of any component (datacommons.cancer.gov).

Evidence, approvals and the release gate

The release package should contain the specification (profile, options, CID 7050 codes, Annex E edition), the corpus and answer key, test results and metrics, the sampling plan and findings log, and sign-offs. The IHE handbook states that an implementation must be rigorously validated with test data before operational use (profiles.ihe.net) and recommends a small representative subset for design review (profiles.ihe.net). Public programs show the pattern: one repository inspects every image for burned-in PHI at curation (wiki.cancerimagingarchive.net). NCI's EDNRN process requires burned-in content to be removed from the pixels themselves (edrn.nci.nih.gov) and sites to validate de-identified images before upload (edrn.nci.nih.gov). FDA guidance asks for archiving that supports a clear audit trail with date and time stamps (www.hhs.gov).

A reasonable release gate has five checks (editorial):

- **Specification.** It matches the protocol and the legal pathway.
- **Coverage.** Every test in the pack ran on the pinned software version.
- **Thresholds.** Metrics meet limits the protocol set in advance.

- **Findings.** Open findings have documented dispositions.
- **Approvals.** The privacy officer, an imaging scientist and counsel have signed off.

Privacy counsel and imaging scientists should review the threat model and utility tests before any release; nothing here substitutes for their judgment.

Data Analysis and Evidence

Table 3 below collects the quantitative findings used in this report, each with its design implication. The values come from different designs (volunteers, archives, synthetic sets, challenges), so they size tests; they are not benchmarks to compare across rows.

Finding	Value	Design implication
Face matching on unaltered head MRI	83% of 84 participants (pmc.ncbi.nlm.nih.gov); 97% of 157 scans (adni.loni.usc.edu)	Treat head volumes as recognizable unless tested
Face matching after popular defacing	28% to 38% of scans still matched (adni.loni.usc.edu)	Defacing needs its own T4 test
Defacing effect on a downstream measure	Up to 5.4% difference (pmc.ncbi.nlm.nih.gov)	Run T7 on the protocol's algorithm
OCR workload after header triage	6050 images, 1.5% of 415,182 (pmc.ncbi.nlm.nih.gov)	Triage lowers cost, not the need for review
MIDI-B challenge scores	97.91% to 99.93% correct actions (arxiv.org)	High action-level scores can coexist with residual findings
NCI validation dataset size	23,921 images (datacommons.cancer.gov)	Public answer keys can seed a local corpus
FDG-PET SUVmax true-change threshold	39% increase or 28% decrease (qibawiki.rsna.org)	Date and weight errors can approach this scale
Unhandled daylight-saving shift (older 2009 QIBA draft)	32% SUV error (qibawiki.rsna.org)	T6 must include clock-change cases

Two patterns stand out. First, pixel privacy and analytic utility fail in different places: recognition risk is high for unaltered head volumes, yet defacing shifts measures by amounts that matter for small effects, so neither can be settled by policy alone. Second, aggregate accuracy hides critical residuals. Challenge scores near 98% to 100% are expressed over all actions, while the pixel actions that matter most were evaluated on very few images (arxiv.org). The MIDI authors state that unsupervised automated tools often fail to remove all identifiers, especially burned-in annotations and unstructured metadata, and recommend algorithmic processing with expert oversight (arxiv.org). Free text stays a persistent source of risk because many elements allow operator free text (arxiv.org).

Limits of the evidence are material. The NCI figures describe particular pipelines and datasets, not the whole field. Synthetic corpora omit patterns that real archives contain. Some measures come from small samples. Where no public benchmark exists for a scenario, none is invented here, and no private implementation

results, error rates or costs are implied.

Implications and Future Directions

Three trends affect how long a profile decision stays valid. First, recognition technology improves: matching rates in one research group rose from 83% to 97% between studies, so a defacing decision should carry a review date. Second, standards move: the DICOM edition changes, and the IHE De-Identification Handbook version in use should be pinned and rechecked before each release (profiles.ihe.net). A specification should name the Annex E edition, CID 7050 version and handbook version it used, and be rechecked immediately before each release. Third, regulatory interpretation is not frozen. HHS states that the standard does not mandate a particular risk-assessment method (www.hhs.gov), and the Safe Harbor text also requires that the covered entity have no actual knowledge that remaining information could identify an individual (www.ecfr.gov). Counsel should confirm how each applies.

For platform teams, the practical implication is to build the test pack into the pipeline. Make T1, T2, T5, T6 and T10 automatic regression checks, keep T3, T4 and T7 as sampled and reviewed, and store results with the specification. For validation teams, Table 1 gives a traceable link from each option to a test and a risk note. For core labs, the protocol-required feature list prevents privacy controls from silently removing endpoint inputs. A hash-based code can itself be an identifier: HHS guidance treats a code derived from a secure hash function without a secret key as an identifying element for Safe Harbor purposes (www.hhs.gov), which bears on any hashed UID or accession scheme. HHS also notes that identifiers must be removed regardless of their location in a record (www.hhs.gov), which supports scanning every text VR, not only listed attributes.

Open questions need local study, not assumption: modality-specific pixel behavior, vendor private-tag content that changes by software version, and the interaction between Retain Safe Private and quantitative tools.

Frequently Asked Questions (FAQs)

How should a team choose a DICOM de-identification profile for a trial? Start from the recipient and the protocol-required feature list, begin with the Basic Profile, and add only the options whose retained content an endpoint needs. Table 1 maps each option to a trigger and a test. Each added retain option widens the surface for residual PHI (wiki.cancerimagingarchive.net).

What is the PS3.15 attribute confidentiality profile? It is the Annex E mechanism: a Basic Application Level Confidentiality Profile plus options, each with action codes in Table E.1-1, intended for clinical trials among other uses.

How is imaging anonymized for a clinical trial? Combine header, private-tag, pixel, UID and date handling, record the profile and options in De-identification Method Code Sequence, and run the ten-test pack in this report.

How can DICOM de-identification be validated? Use a synthetic corpus with an answer key, compute the eight metrics, review high-risk strata fully, and escalate on residual findings. The IHE handbook says validation with test data must precede operational use (profiles.ihe.net). This report does not compare software; the same pack applies to any tool.

How should de-identification tools be compared? This report intentionally ranks no products. Run every candidate through the same synthetic corpus and the same eight metrics, with the same options configured, and record residual-PHI and false-removal findings per stratum.

Does HIPAA Safe Harbor work for DICOM? Safe Harbor lists face images, device serial numbers and date elements other than year (www.ecfr.gov). Retaining modified dates for longitudinal analysis therefore normally needs another pathway. This is not legal advice.

How is burned-in pixel data handled? Add Clean Pixel Data, which must change stored pixel values, review every overlay plane, and sample by modality risk, with ultrasound and screenshots in the highest-risk stratum (arxiv.org).

What does retain longitudinal temporal information do? Two exclusive options keep full dates or modified dates that preserve gross relationships; the modified option needs a consistent mapping across runs (dicom.nema.org).

Conclusion

DICOM Annex E gives trial teams a precise vocabulary for what a de-identified object retains, removes or replaces, but it does not decide for them. The decision is protocol-specific. It begins with a threat model that names the recipient and linkage sources, moves to an option set justified by endpoint needs, and ends with evidence: a synthetic corpus, defined metrics, stratified review and signed approvals. The Basic Profile is conservative by design, each retain option trades privacy for utility, and no profile removes contextual re-identification risk.

The tools in this report are meant to be reused. Table 1 links each option to a trigger, a test and a residual risk. Table 2 specifies a corpus that can be adapted to a protocol. The eight metrics make runs comparable without inventing universal thresholds. Because a sample with no findings cannot prove zero risk, release decisions should rest on documented coverage, explicit residual-risk statements and review by privacy counsel and imaging scientists. Edition numbers, regulator pages and implementation guides should be rechecked immediately before each release, since each can change independently of the protocol.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/chapter_E.html
2. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/sect_E.3.html
3. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/sect_E.2.html
4. <https://datacommons.cancer.gov/news/final-report-medical-imaging-de-identification-midi-project>
5. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7091256>
6. https://adni.loni.usc.edu/adni-publications/Schwarz_2021_Changing%20the%20face%20of%20neuroimaging.pdf
7. <https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html>
8. <https://profiles.ihe.net/ITI/DeId/process.html>

9. https://qibawiki.rsna.org/images/c/c7/QIBA_FDG-PET_Profile_v114.pdf
10. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/sect_E.3.6.html
11. <https://wiki.cancerimagingarchive.net/x/KACkAw>
12. <https://profiles.ihe.net/ITI/Deld/index.html>
13. <https://www.hhs.gov/guidance/sites/default/files/hhs-guidance-documents/FDA/Clinical-Trial-Imaging-Endpoint-Process-Standards-Guidance-for-Industry.pdf>
14. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6650473>
15. <https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-C/part-164/subpart-E/section-164.514>
16. <https://www.dicomstandard.org/News/ftsup/docs/sups/sup142.pdf>
17. <https://profiles.ihe.net/ITI/Deld/references.html>
18. <https://profiles.ihe.net/ITI/Deld/dicom-example.html>
19. <https://datacommons.cancer.gov/media/302>
20. <https://edrn.nci.nih.gov/data-and-resources/informatics/labcas-help/edrn-dicom-de-identification-process>
21. <https://profiles.ihe.net/ITI/Deld/ihe-use.html>
22. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_C.7.html
23. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/sect_E.3.10.html
24. https://dicom.nema.org/medical/dicom/current/output/chtml/part05/sect_7.8.html
25. <https://arxiv.org/pdf/2303.10473>
26. <https://pmc.ncbi.nlm.nih.gov/articles/PMC4450976>
27. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_C.7.6.html
28. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_C.9.2.html
29. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/sect_E.3.3.html
30. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11522224>
31. <https://arxiv.org/pdf/2507.23608>
32. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7943842>
33. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8249889>
34. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10246049>
35. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9434540>
36. https://dicom.nema.org/medical/dicom/current/output/chtml/part15/sect_E.3.9.html
37. <https://wiki.cancerimagingarchive.net/x/dQSwC>
38. <https://pubmed.ncbi.nlm.nih.gov/27013522>
39. https://qibawiki.rsna.org/index.php/Profile:_FDG-PET_Whole_Body
40. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_C.7.2.3.html
41. <https://dctd.cancer.gov/research/ctep-trials/for-sites/recist-guidelines-v11.pdf>
42. <https://www.hhs.gov/hipaa/for-professionals/faq/can-device-identifier-di-portion-unique-device-identifier-udi-be-part-limited-or-de-identified/index.html>
43. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8285420>

44. <https://arxiv.org/pdf/2508.01889v1>
45. <https://profiles.ihe.net/ITI/DeId/history.html>

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore [AI enablement](#), [pharma and biotech software development](#), [data engineering and BI](#), and [Veeva services](#). [Contact IntuitionLabs](#) to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: <https://intuitionlabs.ai>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.