JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Patient Record De-Identification and Privacy Controls Course for Health AI

23 – 27 November 2026
Geneva

Patient Record De-Identification and Privacy Controls Course for Health AI

Ref: 1414_52814 **Date:** 23 – 27 November 2026 **Location:** Geneva **Fees:** 7500 GBP

Patient Record De-Identification and Privacy Controls for Health AI is a 1-week training course for health information managers, privacy officers and clinical IT staff, producing a Health AI Data Protection Dossier for each model that touches patient records. Hospitals and insurers now feed clinical notes, images and claims into AI tools, so re-identification risks grow with each dataset. Jupiter Prime Training delivers Patient Record De-Identification and Privacy Controls at Core level.

Why Patient Record De-Identification and Privacy Controls Matter Now

Three groups shape the subject today. Data science teams request large extracts of patient records to train and test models. Suppliers offer hosted AI services that process records outside the organisation. Patients and ethics boards expect their information to stay private and to be used only for agreed purposes. A provider that de-identifies and controls data well can support research and AI services with confidence; one that does not risks re-identified patients, breaches and lost trust. The course uses the ISO/IEC 42001 AI management system standard and the NIST AI Risk Management Framework NIST AI RMF to place privacy controls within AI accountability.

Who Gains Most from Patient Record De-Identification and Privacy Controls

- Health information managers who release record extracts for analytics and model training
- Privacy officers who run impact assessments and handle patient requests
- Clinical IT staff who configure access, encryption and logging around AI services
- Research governance staff who approve secondary use of clinical data
- Procurement and contract staff who assess AI suppliers handling patient information

A hospital group or health insurer gains most when its records, privacy and IT staff attend together and agree one release process for AI datasets.

Patient Record De-Identification and Privacy Controls Capabilities

The outcomes centre on releasing useful data for AI while keeping patients unidentifiable.

- Apply suppression, generalisation and pseudonymisation to clinical extracts — data release decisions
- Test de-identified datasets for re-identification risk using quasi-identifier checks — ethics boards
- Specify role-based access, encryption and audit logging for model pipelines — clinical IT owners
- Assess AI suppliers against data handling and residency requirements — contract approvals
- Design a breach response path for leaked training data or model outputs — incident leads

- Communicate data use limits to clinicians and patients — public trust

Key Patient Record De-Identification and Privacy Controls Terms

- De-identification — removing or transforming details in a record so that a patient cannot reasonably be identified from it.
- Pseudonymisation — replacing direct identifiers with codes, while a separately held key could still link records back to patients.
- Quasi-identifier — an attribute such as birth date or postcode that can identify someone when combined with others.
- k-anonymity — a measure showing that each record shares its quasi-identifier values with at least k minus one others.
- Privacy impact assessment — a structured review of how a project collects, uses and protects personal information and what risks follow.
- Model inversion — an attack that tries to recover training records or sensitive attributes from a model's outputs.

Patient Record De-Identification and Privacy Controls Day by Day

Day 1: Patient Data in AI Projects

- Clinical Notes, Images and Claims Datasets — Where Identifiers Hide
- Primary Versus Secondary Use of Records — Purpose Limits for AI
- Hospital and Health Insurer Cases — Contrasting Data Release Pressures
- Data Flow Mapping for Model Training — Tracing Every Copy
- Patient Consent and Notice Options — What Patients Were Told

Day 2: De-Identification Techniques

- Suppression, Generalisation and Masking — Choosing Per Data Field
- Pseudonymisation and Key Custody — Who Can Relink Records
- Free-Text and Image Metadata Scrubbing — Hidden Names and Dates
- k-Anonymity and Quasi-Identifier Testing — Measuring Re-Identification Risk
- Synthetic Data Generation — When Real Records Are Not Needed

Day 3: Securing Model Pipelines

- Role-Based Access to Training Environments — Least Privilege for Analysts
- Encryption at Rest and in Transit — Protecting Extracts and Checkpoints
- Audit Logging and Query Monitoring — Detecting Unusual Extraction
- Model Inversion and Membership Inference Risks — Output Release Rules
- NIST AI RMF Map and Manage Functions — Privacy Risk Treatment

Day 4: Assessment, Suppliers and Breaches

- Privacy Impact Assessment for Health AI — Approve, Modify or Decline
- Supplier Due Diligence Questionnaire — Hosting, Retention and Subcontractors
- Data Processing Agreement Clauses — Limits on Supplier Model Training
- ISO/IEC 42001 Roles and Records — Accountability for Data Controls
- Breach Response Path for AI Datasets — Containment and Patient Notification

Day 5: Practising the Health AI Data Protection Dossier

- De-Identification Exercise on a Sample Extract — Applying Day Two Techniques
- Re-Identification Risk Test — Applying Quasi-Identifier Checks
- Pipeline Access and Logging Design — Applying Day Three Controls
- Supplier Assessment Role Play — Applying Due Diligence and Contract Clauses
- Health AI Data Protection Dossier — Final Version for Release Approval

Applying Patient Record De-Identification and Privacy Controls After the Course

- Map the data flows of one live AI project and record every copy of patient records it holds.
- Run a quasi-identifier test on the next dataset released for model training and document the result.

Patient Record De-Identification and Privacy Controls Questions FAQ

What are Patient Record De-Identification and Privacy Controls and why do they matter now?

They are the techniques and safeguards that keep patients unidentifiable when their records support AI. They matter now because AI projects request larger extracts and often send them to outside suppliers.

Which roles gain most from Patient Record De-Identification and Privacy Controls?

Health information managers, privacy officers, clinical IT staff and research governance staff gain most.

How does Patient Record De-Identification and Privacy Controls differ from a healthcare cybersecurity course?

It focuses on de-identifying and controlling patient records for AI use rather than defending networks against attacks. A cybersecurity course protects systems; this one decides what data may leave them.

What do you bring back from Patient Record De-Identification and Privacy Controls?

You bring back a tested de-identification method, a re-identification risk check, pipeline access rules, a supplier questionnaire and a breach response path.

Patient Record De-Identification and Privacy Controls: The Way Forward

The Health AI Data Protection Dossier first decides whether a requested dataset may be released for model training. Records managers, privacy officers and ethics boards use it before each release. After its first use, revisit the re-identification thresholds against the actual extract, since rare conditions and small hospitals often need stronger generalisation than expected.