



Global AI
Certification Council



GAICC HAIGP

HEALTHCARE AI GOVERNANCE PROFESSIONAL

Examination Content Outline
1st Edition | 2026 | v1.0

Global AI Certification Council
(GAICC)

Examination Content Outline

1st Edition | 2026 · GAICC-ECO-HAIGP-001

The Credential at a Glance

80 MCQs

120 minutes · Examination

70%

Pass mark · scenario-based

60 CPDs

Every 3 years · Renewal

Certification	GAICC Healthcare AI Governance Professional (GAICC-HAIGP)
Document code	GAICC-ECO-HAIGP-001
Version	1.0 · 23 October 2026
Level	Tier 3, Senior and Specialty (sector professional)
Examination	80 questions · 120 minutes · 70% pass
Training	32 contact hours, completed before the exam
Contact	certify@gaicc.org gaicc.org certifications.gaicc.org

Published by: Global AI Certification Council Limited (GAICC)

3900 Westerre Pkwy,
Richmond, VA 23233,
USA

Level 2, 697 Collins Street,
Melbourne VIC 3008,
Australia

Level 3, 21 Putney Way
Manukau, Auckland,
New Zealand 2104

© 2026 Global AI Certification Council. All rights reserved. GAICC, the GAICC logo, Healthcare AI Governance Professional, HAIGP, Certified Professional in AI Governance, CPAIG, and related marks are trademarks of the Global AI Certification Council.

Contents

The Credential at a Glance	2
1 About this Examination Content Outline	4
2 Purpose of the Certification	4
3 Who the Certification Is For	5
4 The Examination at a Glance	6
5 Domain Weighting	6
Domain I — Healthcare AI Landscape and Accountability	7
Domain II — Health AI Regulation and Medical Device Pathways	8
Domain III — Clinical Risk, Safety and Impact Assessment	9
Domain IV — Health Data, Privacy and Consent	10
Domain V — Health Equity, Ethics and Patient Transparency	11
Domain VI — Clinical Validation, Evaluation and Assurance	12
Domain VII — Post-Market Monitoring, Change Control and Incident Response	13
Domain VIII — Generative and Agentic AI, Supply and Scale	14
7 Reference Frameworks	15
8 Sample Questions	18
9 Eligibility	20
10 Maintaining the Certification	20
11 Frequently Asked Questions	21

1 About this Examination Content Outline

This Examination Content Outline (ECO) defines what the Healthcare AI Governance Professional (HAIGP) examination measures. It lists the eight domains, their weightings, and the tasks and knowledge a certified professional is expected to demonstrate. Every examination item is written to, and traceable to, one task code in this outline.

Use this outline to plan your preparation. Training providers use it to design courses, but no training provider sees live examination content, and every training route is treated equally.

How the content was defined

The domains and tasks were drafted from a role analysis of AI governance work in healthcare and life sciences, current regulation and standards, and employer needs across regulated enterprises and the technology and services firms that support them. The draft is subject to validation by the GAICC Subject Matter Expert Panel before the first live examination window.

2 Purpose of the Certification

AI now sits inside clinical care, coverage decisions, drug development and the administrative core of health organisations: early-warning models, imaging triage, ambient scribes, prior authorisation tools, patient messaging and care navigation agents. The harms are clinical and they scale. A drifting sepsis model misses deterioration, a scribe omits an allergy, a payer algorithm denies care on a prediction rather than the patient's circumstances, or a model performs worse for an under-served group. Accountability is split across manufacturers, deploying organisations, clinicians, payers and the services firms that build and run these systems, and the rules are layered across device law, AI law, privacy law and coverage regulation.

HAIGP holders lead and assure AI governance across that lifecycle. They classify use cases and determine regulatory status, build clinical safety cases, hazard logs and impact assessments, govern health data, consent and data sovereignty, test for bias and design patient transparency, specify validation and output audits, run post-market monitoring, change control and incident reporting, and set the autonomy limits, human checkpoints and contract terms that generative and agentic AI require. They do this from inside a health system, a payer, a medtech or pharma company, or a technology and services provider delivering AI to those clients.

The timing matters. Under the EU AI Act as amended in 2026, high-risk duties apply to Annex III health uses from 2 December 2027 and to AI medical devices and IVDs from 2 August 2028, alongside MDR and IVDR. The FDA has final guidance on predetermined change control plans and clinical decision support and has opened discussion of generative AI devices and foundation model dependencies. UK post-market surveillance rules require serious incident reports within 15 days, and sooner for deaths and public health threats, and US states are legislating on AI in care and coverage, with Colorado's rules applying from 1 January 2027. Organisations need practitioners who can turn these requirements into controls that work on the ward, in the claims queue and in the code.

3 Who the Certification Is For

The HAIGP is designed for professionals who govern, assure, deliver or advise on AI in healthcare and life sciences, including:

- Health system AI and digital governance leads
- Medtech regulatory affairs and quality professionals
- Pharma and life sciences AI leads
- Consultants and architects delivering health AI
- Clinical safety officers and clinical informatics leaders
- Payer utilisation management and compliance leaders
- Privacy, information governance and security officers
- Procurement, legal and risk professionals in health

Roles the certification supports

Head of Clinical AI Governance

Clinical Safety Officer (Digital and AI)

AI Regulatory Affairs Manager (SaMD)

Responsible AI Lead, Healthcare and Life Sciences

Health AI Assurance Consultant

What certified professionals can do

- 01 Classify health AI use cases and determine regulatory status across major device and AI regimes.
- 02 Plan evidence for AI medical devices, including PCCPs and dual MDR/IVDR and AI Act conformity.
- 03 Build clinical safety cases, hazard logs and impact assessments proportionate to clinical risk.
- 04 Govern health data use for AI under HIPAA, GDPR, EHDS and NZ privacy rules.
- 05 Test health AI for subgroup bias and design patient transparency that meets disclosure laws.
- 06 Specify validation and output audits for predictive models, ambient scribes and generative AI.
- 07 Run post-market monitoring, change control and incident reporting across jurisdictions.
- 08 Set autonomy limits, human checkpoints and contract terms for agentic and vendor-supplied AI.

4 The Examination at a Glance

Component	Details
Examination	80 questions (70 scored + 10 unscored pilot)
Duration	120 minutes
Pass mark	70% of scored items (scaled score of 350 on a 100 to 500 scale)
Format	Scenario-based multiple choice, single best answer from four options, no negative marking
Delivery	Online with AI-assisted proctoring through the GAICC Hub, available 24/7. A person makes any final integrity decision
Conditions	Closed book. No reference material, notes or AI tools
Training requirement	32 contact hours, completed before the exam
Validity and renewal	3 years · 60 CPD hours per cycle
Credential awarded	GAICC Healthcare AI Governance Professional (GAICC-HAIGP)

Cognitive levels

Most items present a realistic sector scenario and ask for the most appropriate action, judgement or control. Recall of clause numbers is not tested.

Cognitive level	Share of scored items	What it looks like
Understanding	25% (about 18 items)	Explain concepts, obligations and the purpose of controls in the sector context.
Application	45% (about 32 items)	Apply frameworks, regulations and controls to realistic sector scenarios.
Analysis and Evaluation	30% (about 21 items)	Weigh competing risks and obligations, judge evidence, and choose the most defensible course of action.

5 Domain Weighting

The 70 scored items are distributed across eight domains as follows. Item counts are approximate and may vary by one or two items between examination forms.

Dom.	Title	Weight	Items
I	Healthcare AI Landscape and Accountability	11%	8
II	Health AI Regulation and Medical Device Pathways	15%	11
III	Clinical Risk, Safety and Impact Assessment	15%	11
IV	Health Data, Privacy and Consent	12%	8
V	Health Equity, Ethics and Patient Transparency	10%	7
VI	Clinical Validation, Evaluation and Assurance	13%	9
VII	Post-Market Monitoring, Change Control and Incident Response	12%	8
VIII	Generative and Agentic AI, Supply and Scale	12%	8
6	Total scored items	100%	70

Maps clinical, operational, administrative and research AI to the people and organisations accountable for it. Covers governance bodies, the enterprise AI inventory and an ISO/IEC 42001 management system adapted to health organisations.

Code	Task: the certified professional can	Knowledge of
HAIGP-1.1	Classify healthcare AI use cases across clinical, operational, administrative and research settings by intended purpose, users, decision impact and patient exposure.	› Categories of clinical, operational, administrative and research AI with examples › Intended purpose as the anchor for regulatory and risk decisions › Differences between predictive, generative, multimodal and agentic AI systems › Patient exposure pathways: direct, clinician-mediated and back-office
HAIGP-1.2	Allocate accountability among developers, manufacturers, deploying health organisations, clinicians, payers and services providers using regulation, contract and professional duty.	› Provider, deployer, importer and distributor roles under the EU AI Act › Manufacturer and legal manufacturer duties under medical device regulation › Clinicians' non-delegable duty of care when acting on AI outputs › Payer accountability for AI-assisted coverage and utilisation decisions › Roles of services providers that build, integrate or operate AI for clients
HAIGP-1.3	Establish AI governance structures with clear decision rights for the AI committee, clinical safety officer, CMIO and CNIO, privacy, security and procurement functions.	› Terms of reference, membership and escalation routes for health AI committees › Clinical safety officer responsibilities for health IT and AI › Roles of chief medical and nursing information officers in clinical AI › Links to clinical governance, quality, risk, ethics and research committees
HAIGP-1.4	Build and maintain an enterprise AI inventory recording intended use, regulatory status, risk tier, owner, supplier, model version and monitoring status for every system.	› Minimum data fields for a health AI register › Discovery of shadow AI and AI features embedded in vendor products › Linking the inventory to procurement, change control and incident records › Using inventory data for board, regulator and external assessor reporting
HAIGP-1.5	Implement ISO/IEC 42001 in a health organisation, integrating it with ISO 13485, ISO/IEC 27001, clinical governance and the Joint Commission and CHAI RUAIH elements.	› ISO/IEC 42001 clauses, Annex A controls and statement of applicability › Integration points with ISO 13485 and ISO/IEC 27001 management systems › The seven RUAIH elements and how they map to ISO/IEC 42001 › Scoping a management system across hospitals, payers and suppliers

Health AI Regulation and Medical Device Pathways

Covers the qualification, classification and market pathways for AI-enabled medical devices across major jurisdictions, and the rules that apply to non-device and drug development AI. Includes dual conformity under MDR/IVDR and the AI Act and the risk of regulatory status drift.

Code	Task: the certified professional can	Knowledge of
HAIGP-2.1	Determine whether an AI function is a medical device or IVD by applying intended purpose, qualification criteria and decision support and wellness exclusions across jurisdictions.	› Software device qualification under FDA, EU, MHRA, TGA, HSA, MDA and CDSCO rules › FDA clinical decision support criteria and January 2026 final guidance › MDCG qualification and classification guidance, including MDR Rule 11 › Why time-critical and single-recommendation outputs change the analysis
HAIGP-2.2	Plan regulatory pathways and evidence for AI-enabled devices, including FDA 510(k), De Novo and PMA routes, PCCPs, and UK, EU and APAC registrations.	› FDA marketing submission content for AI-enabled device software functions › PCCP components: description of modifications, protocol and impact assessment › UK registration, MHRA AI Airlock and post-market surveillance expectations › TGA, HSA GL-04, MDA and CDSCO requirements for AI software
HAIGP-2.3	Prepare AI medical devices and IVDs for dual conformity under EU MDR/IVDR and AI Act Annex I, using one integrated technical documentation set.	› AI Act Annex I Section A and application from 2 August 2028 › MDCG 2025-6 on the interplay of MDR/IVDR and the AI Act › Notified body assessment of AI Act requirements within MDR/IVDR procedures › Mapping AI Act Articles 9 to 15 to general safety and performance requirements
HAIGP-2.4	Identify obligations for non-device health AI, including AI Act Annex III triage and eligibility uses, CMS utilisation rules and US state laws on AI in care and coverage.	› Annex III health uses: emergency triage, public health eligibility, insurance pricing › CMS Medicare Advantage medical necessity guidance and CMS-0057-F duties › California SB 1120, AB 3030 and AB 489; Texas SB 1188 and TRAIGA › Colorado SB 26-189 and HB 26-1139 duties for healthcare AI from 2027 › ONC HTI-1 decision support duties and the proposed HTI-5 changes
HAIGP-2.5	Apply regulatory expectations for AI in drug and biologic development, including the EMA-FDA guiding principles, the FDA credibility framework and GxP validation.	› The ten EMA-FDA Guiding Principles of Good AI Practice in Drug Development › FDA credibility assessment: context of use, model influence and decision consequence › GxP validation, GAMP 5 and draft EU GMP Annex 22 on AI › AI uses in trials, pharmacovigilance, manufacturing and regulatory submissions
HAIGP-2.6	Monitor regulatory status drift as functions, claims, users or integrations change, and trigger requalification before an unregistered device enters clinical use.	› Triggers that turn administrative tools into regulated medical devices › TGA digital scribe review findings and UK expectations for scribes › Feature creep from summarisation towards diagnosis, coding or recommendation › Horizon scanning: UK Commission recommendations, MDR simplification, new FDA papers



Clinical Risk, Safety and Impact Assessment

Applies device risk management, clinical safety case practice and AI impact and risk assessment to health AI. Covers automation bias and human factors, and the tiering of use cases so controls are proportionate to clinical consequence.

Code	Task: the certified professional can	Knowledge of
HAIGP-3.1	Apply ISO 14971 and ISO/TS 24971-2:2026 to identify, estimate, control and evaluate machine learning risks in medical devices across the lifecycle.	- ISO 14971 risk management process and benefit-risk determination - Machine learning hazards in ISO/TS 24971-2: data, bias, overfitting, drift - Scope limits of ISO/TS 24971-2, including the exclusion of LLM devices - Supporting guidance in ISO/TR 24971 and AAMI TIR34971
HAIGP-3.2	Develop and maintain clinical safety cases and hazard logs for health IT and AI under DCB0129 and DCB0160, including the manufacturer to deployer handover.	- DCB0129 manufacturer and DCB0160 deployer clinical risk management duties - Structure and content of a clinical safety case report and hazard log - Hazard identification techniques for AI: SWIFT, FMEA and bow-tie analysis - National review of DCB0129 and DCB0160 to address AI
HAIGP-3.3	Conduct AI impact assessments under ISO/IEC 42005 and AI risk assessments under ISO/IEC 23894 for non-device and organisation-wide AI uses.	- ISO/IEC 42005 impact assessment process, scope and documentation - ISO/IEC 23894 AI risk management aligned with ISO 31000 - Links between AI impact assessment, DPIA, AI Act Article 27 FRIA and equality assessment - Involving patients, clinicians and communities in impact assessment
HAIGP-3.4	Design human oversight controls that counter automation bias, alert fatigue and deskilling, using human factors and usability engineering evidence.	- Automation bias, complacency and alert fatigue in clinical decision support - IEC 62366-1 usability engineering and summative evaluation - Oversight design: presentation, confidence display, override and escalation - Deskilling risk and how clinician competence is maintained
HAIGP-3.5	Tier healthcare AI use cases by clinical consequence, autonomy and reversibility to set proportionate approval, validation and monitoring requirements.	- Tiering criteria: severity, autonomy, reversibility, population and time criticality - IMDRF significance and healthcare situation matrix for software devices - FDA autonomy versus consequence grid for generative AI devices - Proportionate controls per tier for approval, validation, oversight and monitoring

Health Data, Privacy and Consent

Governs the use of health information to build and run AI across US, EU, UK, NZ and APAC regimes. Covers secondary use, de-identification and synthetic data, consent for scribes and other capture technologies, Indigenous data sovereignty and cross-border flows.

Code	Task: the certified professional can	Knowledge of
HAIGP-4.1	Determine lawful bases and contractual safeguards for health information used in AI, including HIPAA business associate agreements, minimum necessary and GDPR special category conditions.	› HIPAA Privacy and Security Rules, BAAs and the minimum necessary standard › GDPR Article 6 lawful bases and Article 9 conditions for health data › Processor and sub-processor duties for AI vendors and services providers › Status of the HIPAA Security Rule update and current risk analysis duties
HAIGP-4.2	Apply NZ, UK, Australian and APAC health privacy rules, including the Health Information Privacy Code 2020, to AI collection, use and disclosure.	› Health Information Privacy Code 2020 rules, including rule 3A from May 2026 › Biometric Processing Privacy Code 2025 where voice or face data identify people › UK GDPR, Caldicott principles and the national data opt-out › Core duties under Australian, Singaporean and Malaysian privacy laws
HAIGP-4.3	Plan secondary use of health data for AI development under the European Health Data Space, research ethics and data access governance.	› EHDS timeline: application 2027, priority EHR categories and secondary use 2029 › Health Data Access Bodies, data permits and secure processing environments › Research ethics approval and consent waivers for AI development › Trusted research environments and data access committee decisions
HAIGP-4.4	Assess de-identification, re-identification and synthetic data approaches for AI datasets, including risks from free text, imaging, genomics and model memorisation.	› HIPAA Safe Harbor and Expert Determination de-identification methods › Re-identification risk from free text, imaging metadata and genomic data › Testing synthetic data for utility, fidelity and privacy › Model memorisation, membership inference and training data extraction
HAIGP-4.5	Design consent, notice and opt-out processes for AI that captures patient interactions, such as ambient scribes and voice or image analysis.	› Consent versus notice models for ambient scribe recordings › Retention and deletion of audio, transcripts and draft notes › Recording opt-outs and handling them in clinical workflow › Sensitive settings: mental health, paediatrics, interpreters and intimate examinations
HAIGP-4.6	Apply Indigenous and Māori data sovereignty principles and cross-border transfer controls to health AI data flows, hosting and model training.	› Te Mana Raraunga Māori data sovereignty principles and CARE principles › Iwi and community governance of data in population health models › Cross-border transfer mechanisms and health data residency requirements › Offshore hosting and access to health data by global delivery teams

Health Equity, Ethics and Patient Transparency

Ensures health AI performs fairly across populations and that patients and clinicians understand when and how AI shapes care and coverage. Covers subgroup testing, proxy discrimination, disclosure laws, ethics and explainability.

Code	Task: the certified professional can	Knowledge of
HAIGP-5.1	Evaluate subgroup performance and bias of health AI across ethnicity, sex, age, language, disability and socioeconomic groups, and set remediation thresholds.	- Subgroup metrics: calibration, sensitivity and false negative rate parity - Data representativeness and label bias in clinical datasets - Accent, dialect and language bias in speech and language tools - AI Act rules on processing special category data for bias detection
HAIGP-5.2	Identify and prevent proxy discrimination in payer and population health algorithms, applying ACA Section 1557 and individualised determination requirements.	- Cost used as a proxy for need and other proxy variable failures - ACA Section 1557 duties for patient care decision support tools - CMS position that medical necessity rests on individual circumstances - Physician review of adverse determinations under laws such as California SB 1120
HAIGP-5.3	Design patient transparency and disclosure for AI in care and communications that meets state disclosure laws, AI Act Article 50 and professional standards.	- California AB 3030 disclaimers and Texas patient disclosure duties - EU AI Act Article 50 transparency obligations applying from 2 August 2026 - Colorado SB 26-189 notice, explanation and human review rights - Plain-language, accessible and multilingual disclosure design
HAIGP-5.4	Apply ethical frameworks, including WHO guidance on large multi-modal models, to informed consent, autonomy, beneficence and justice in AI-supported care.	- WHO ethics and governance guidance for AI and large multi-modal models - Informed consent when AI materially influences diagnosis or treatment - Clinical ethics committee input on contested AI uses - Limits on AI-delivered therapy, such as Illinois law, and companion chatbots
HAIGP-5.5	Specify explainability suited to clinicians, patients and reviewers, including source attributes, uncertainty, limitations and reasons for adverse decisions.	- Global and local explanations and their limits in clinical use - Source attributes and transparency for predictive decision support - Specific denial reasons required under CMS-0057-F - Communicating uncertainty and known limitations to clinicians and patients

Sets the evidence needed before and after go-live, from Good Machine Learning Practice and external validation to structured audits of generative outputs. Covers vendor-independent assurance, health software standards and AI management system audits.

Code	Task: the certified professional can	Knowledge of
HAIGP-6.1	Apply Good Machine Learning Practice (IMDRF N88) to dataset design, reference standards, independence of training and test data, and model documentation.	› The ten GMLP principles and the evidence each implies › Reference standard selection and label quality in clinical AI › Train, tune and test independence and data leakage risks › IMDRF N93 draft framework for AI lifecycle management
HAIGP-6.2	Design external and prospective validation, silent-mode trials and local performance verification before go-live and after major change.	› Internal, external, temporal and geographic validation designs › Silent-mode and shadow deployment before clinical use › Reporting guidelines: TRIPOD+AI, CONSORT-AI and DECIDE-AI › Local acceptance testing and site-specific operating thresholds
HAIGP-6.3	Evaluate generative AI and ambient scribe outputs with structured audits for omission, hallucination, misattribution and clinically significant error, using defensible sampling plans.	› Error taxonomies for clinical documentation: omission, fabrication, misattribution › Sampling, stratification and clinician review protocols for output audits › High-risk content: allergies, medications, doses, laterality and negation › Limits of generic LLM benchmarks versus task-specific competency evaluation
HAIGP-6.4	Review vendor assurance evidence, model cards and health AI labels, and commission vendor-independent testing proportionate to risk.	› CHAI Applied Model Card content and the gaps to probe › Evidence pack contents: intended use, data, performance, limits, monitoring › Independent testing, red teaming and local benchmarking options › Signs that vendor evidence will not transfer to the local population
HAIGP-6.5	Verify health software lifecycle, security and AI management system conformity using IEC 62304, IEC 82304-1, IEC 81001-5-1 and ISO/IEC 42006 audit criteria.	› IEC 62304 Edition 2 changes for all health software and AI › IEC 82304-1 product requirements for health software › IEC 81001-5-1 security activities across the health software lifecycle › ISO/IEC 42006 requirements for bodies auditing AI management systems › Voluntary RUAIH organisational certification and its evidence expectations

Post-Market Monitoring, Change Control and Incident Response

Keeps deployed health AI safe as data, models and practice change. Covers real-world performance and drift monitoring, change control and significant-change decisions, vigilance and AI safety event reporting, device cybersecurity and decommissioning.

Code	Task: the certified professional can	Knowledge of
HAIGP-7.1	Design real-world performance and drift monitoring for deployed health AI, covering input, output, outcome and subgroup signals with defined action thresholds.	› Data, concept and label drift in clinical settings › Outcome lag, feedback loops and confounding by clinical intervention › Monitoring metrics, statistical process control and alert thresholds › FDA interest in real-world performance monitoring of AI devices
HAIGP-7.2	Operate change control for AI models using PCCPs and algorithm change protocols, deciding whether changes, including vendor foundation model swaps, are significant.	› Significant change criteria under FDA, EU MDR and UK guidance › PCCP scope limits and when a new submission is required › HSA and CDSCO change management and algorithm change protocols › Model swaps, prompt changes and retrieval updates as controlled changes
HAIGP-7.3	Manage vigilance and serious incident reporting for AI devices across jurisdictions, meeting FDA MDR, EU vigilance and UK post-market surveillance timelines.	› UK serious incident reporting: 15 days, 10 for deaths, 2 for public health threats › FDA Medical Device Reporting duties for manufacturers and user facilities › EU MDR vigilance, trend reporting and periodic safety update reports › AI Act serious incident duties and their interplay with MDR vigilance
HAIGP-7.4	Run AI safety event reporting, investigation and learning for clinical and non-device AI, including near misses, automation bias and documentation errors.	› AI safety event taxonomy and internal reporting channels › Investigating AI-involved incidents, including human factors contributors › Systems-based patient safety incident response and just culture › Sharing learning with vendors, regulators and peer organisations
HAIGP-7.5	Govern device cybersecurity and decommissioning of health AI, including FDA section 524B requirements, SBOMs, patching, data retention and safe withdrawal.	› FDA section 524B cyber device requirements and June 2025 guidance › SBOMs, vulnerability management and coordinated disclosure › AI-specific threats: data poisoning, model extraction and adversarial inputs › Decommissioning: clinical fallback, record retention and stakeholder notification

Governs generative and agentic AI in clinical, patient and payer workflows, and the supply chain behind them. Covers foundation model dependencies, procurement and contracts, services-provider obligations and scaling governance across a health system.

Code	Task: the certified professional can	Knowledge of
HAIGP-8.1	Govern ambient scribes and generative AI patient communications, setting clinician review, disclosure, scope limits and output audit requirements before scaled roll-out.	- NHS England ambient voice technology guidance requirements - Health NZ guidance on generative AI and endorsed scribe pilots - Clinician sign-off, edit tracking and accountability for AI drafts - Disclaimers and routes to a human for AI patient messages
HAIGP-8.2	Define autonomy limits, human checkpoints, tool permissions and audit trails for agentic AI in prior authorisation, orders, scheduling and care navigation.	- Levels of agent autonomy and action reversibility in health workflows - Approval gates, least-privilege tool scoping and kill switches - Agent logging, traceability and replay for audit and investigation - Risks of multi-agent and multi-model orchestration
HAIGP-8.3	Mitigate threats specific to generative and agentic AI, including prompt injection through EHR text, data leakage, over-retrieval and unsafe actions.	- Direct and indirect prompt injection via notes, messages and documents - PHI leakage and over-retrieval in retrieval-augmented systems - Red teaming and guardrail testing for clinical and payer agents - OWASP guidance for LLM and agentic applications
HAIGP-8.4	Assess foundation model dependencies and supply chain risk, applying the FDA generative AI discussion paper and AI Act general-purpose AI obligations.	- FDA discussion paper concepts, including a voluntary foundation model master file - AI Act general-purpose AI provider obligations applying since August 2025 - Model version pinning, deprecation and upstream change notice - Concentration and continuity risk from a single model supplier
HAIGP-8.5	Conduct procurement due diligence and negotiate AI contract clauses covering liability, audit rights, notice of model changes, data use, performance and exit.	- Due diligence, NHS DTAC and NICE evidence standards for health AI - Liability allocation and the revised EU Product Liability Directive - Clauses: audit rights, change notice, data use, incident support, exit - Service levels for performance, monitoring and model updates
HAIGP-8.6	Determine services-provider obligations when building, integrating or operating health AI, including provider status through substantial modification, and scale governance across a health system.	- AI Act Article 25: becoming a provider through modification or rebranding - Supplier duties to deployers: documentation, logging and oversight design - HIPAA, GxP and processor duties for global delivery teams - Scaling governance: federated committees, AI literacy and assurance metrics

7 Reference Frameworks

Candidates should be familiar with the intent, scope and practical consequences of the following laws, regulations, guidance and standards. Questions test the ability to apply them, not to recite them. Where a framework is proposed, delayed or being phased in, candidates are expected to understand its status and direction.

Framework, law or standard	Jurisdiction	Status (October 2026)	What a holder must be able to do
EU AI Act (Regulation (EU) 2024/1689), as amended by Regulation (EU) 2026/1744	EU	Annex III health uses apply 2 Dec 2027; Annex I devices 2 Aug 2028	Classify health AI, assign provider and deployer duties, meet transparency rules and plan dual conformity for AI medical devices.
EU Medical Device Regulation (EU) 2017/745 and IVD Regulation (EU) 2017/746, with MDCG 2025-6	EU	In force; Commission simplification proposal (Dec 2025) under negotiation	Qualify and classify AI software, prepare technical documentation and integrate AI Act requirements into one conformity file.
European Health Data Space Regulation (EU) 2025/327	EU	In force; applies Mar 2027; EHR systems and secondary use mainly from Mar 2029	Plan lawful secondary use of health data for AI training through data permits and secure processing environments.
EU Product Liability Directive (EU) 2024/2853	EU	Transposition due 9 December 2026; software and AI treated as products	Assess liability exposure for defective AI and allocate it in supplier and deployment contracts.
General Data Protection Regulation (EU) 2016/679 and UK GDPR	EU / UK	In force; health data is special category data	Establish lawful basis and Article 9 conditions, run DPIAs and set processor terms for AI vendors and services providers.
FDA guidance on AI-enabled device software functions and PCCPs	US	PCCP guidance final Dec 2024; lifecycle guidance draft Jan 2025	Plan marketing submission content and design predetermined change control plans for AI-enabled devices.
FDA Clinical Decision Support Software and General Wellness guidances	US	Revised final guidances issued 6 January 2026	Decide whether decision support or wellness software falls outside device oversight or under enforcement discretion.
FDA discussion paper on generative AI-enabled medical devices	US	Published 18 August 2026; comment period closed 19 October 2026	Anticipate expectations on autonomy and consequence, competency benchmarking, agentic behaviour and foundation model master files.
FDA premarket cybersecurity guidance (section 524B) and Quality Management System Regulation	US	Cybersecurity guidance final Jun 2025; QMSR effective 2 February 2026	Specify cyber device controls, SBOMs and quality system expectations aligned to ISO 13485 for AI devices.

7 Reference Frameworks (continued)

Framework, law or standard	Jurisdiction	Status (October 2026)	What a holder must be able to do
EMA-FDA Guiding Principles of Good AI Practice in Drug Development and FDA credibility framework	EU / US	Principles published 14 Jan 2026; FDA credibility guidance draft Jan 2025	Define context of use, assess model credibility and govern AI across trials, pharmacovigilance and manufacturing.
HIPAA Privacy, Security and Breach Notification Rules	US	In force; Security Rule overhaul delayed, now expected July 2027	Set business associate agreements, minimum necessary limits and risk analysis for AI vendors, scribes and agents.
CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F) and Medicare Advantage guidance	US	Timeframes and denial reasons from Jan 2026; APIs from Jan 2027	Govern AI in utilisation management so determinations rest on individual circumstances and state specific denial reasons.
ACA Section 1557 nondiscrimination rule	US	2024 rule covers patient care decision support tools; monitor enforcement position	Identify and mitigate discrimination risk in clinical and payer decision support tools.
US state health AI laws (California, Texas, Colorado, Illinois)	US states	Several in force; Colorado SB 26-189 and HB 26-1139 from 1 January 2027	Meet patient disclosure, physician review, utilisation review and automated decision notice and appeal duties.
ONC HTI-1 decision support intervention requirements	US	In force; proposed HTI-5 rule (Dec 2025) would remove source attribute requirements	Assess source attribute transparency and risk management for predictive decision support in certified health IT.
NIST AI Risk Management Framework 1.0 and Generative AI Profile (AI 600-1)	US	Voluntary; AI RMF revision in progress	Structure enterprise AI risk management and generative AI controls in a common vocabulary.
UK MHRA medical device framework, post-market surveillance regulations and AI Airlock	UK	PMS regulations in force June 2025; AI Airlock phase 3 runs 2026-2029	Plan UK registration, meet serious incident timelines and use sandbox learning for AI devices.
NHS DCB0129 and DCB0160 clinical risk management standards and ambient voice technology guidance	UK (England)	Mandatory in England; standards under national review to add AI provisions	Produce clinical safety cases and hazard logs and assure ambient scribes before and after deployment.

7 Reference Frameworks (continued)

Framework, law or standard	Jurisdiction	Status (October 2026)	What a holder must be able to do
APAC device regulators: TGA (Australia), HSA (Singapore), MDA (Malaysia), CDSCO (India)	APAC	TGA scribe review Oct 2026; HSA GL-04 Rev 4; CDSCO software guidance draft	Determine software device status, registration and change management for AI products across APAC markets.
NZ Health Information Privacy Code 2020, Health NZ AI guidance and Māori data sovereignty	New Zealand	Code in force; rule 3A applies from 1 May 2026; Medical Products Bill not introduced	Govern health information in AI, follow precautionary generative AI guidance and honour Māori data governance.
IMDRF Good Machine Learning Practice (N88) and AI Life Cycle Management framework (N93)	International	N88 final 2025; N93 draft April 2026	Apply internationally agreed principles for data, training, testing and lifecycle management of AI devices.
ISO 14971:2019, ISO/TR 24971:2020 and ISO/TS 24971-2:2026	International	Published; ISO/TS 24971-2 covers machine learning but excludes LLM devices	Run device risk management for machine learning hazards and justify benefit-risk decisions.
ISO/IEC 42001, 23894, 42005 and 42006	International	Published 2023-2025	Implement an AI management system, assess AI risk and impact, and prepare for management system audits.
IEC 62304, IEC 82304-1 and IEC 81001-5-1 health software standards	International	Edition 2 at final stage; publication expected late 2026	Verify software lifecycle, product safety and security processes for AI health software.
Joint Commission and CHAI Responsible Use of AI in Healthcare guidance and CHAI Applied Model Card	US / International	RUAIH guidance Sep 2025; voluntary organisational certification announced 2026	Benchmark organisational AI governance and request structured transparency from AI developers.
WHO guidance on ethics and governance of large multi-modal models	International	Published January 2024	Apply ethical principles and stakeholder duties to generative and multimodal AI in health.

Status reflects published information as at October 2026. GAICC reviews this list at least annually; examination items are not based on changes published less than six months before an examination window.

8 Sample Questions

These questions show the style and level of the examination. They are not drawn from the live item bank.

Sample question 1 HAIGP-2.6 Application

A hospital in England deployed an ambient scribe that the vendor supplies as a documentation tool, not registered as a medical device. The clinical safety case was approved on that basis. A vendor update now adds a feature that suggests differential diagnoses and proposes problem list codes in the draft note. Clinicians still review every note before sign-off. What should the AI governance lead do first?

- A. Enable the feature, because clinician review of every note before sign-off keeps the product outside device regulation
- B. Keep the feature off until the vendor confirms its regulatory status, and update the safety case and hazard log first
- C. Enable the feature for a pilot group and rely on the existing monthly scribe output audit to detect any problems
- D. Ask the vendor to add an on-screen notice that suggestions are not clinical advice, then enable the feature for all users

Answer: B. Suggesting diagnoses and codes changes the intended purpose and can make the product a medical device, so regulatory status must be requalified and the DCB0160 safety case updated before clinical use. Clinician review does not by itself take software outside device regulation, so A is the strongest distractor but wrong. A disclaimer or pilot does not resolve unregistered device use.

Sample question 2 HAIGP-7.2 Analysis and Evaluation

A manufacturer's FDA-authorized AI device triages adult chest radiographs. Its authorised PCCP covers retraining with data from additional sites, with pre-specified acceptance criteria and an update protocol. The commercial team wants to extend use to patients aged 6 to 17. A retrained model meets every PCCP acceptance criterion on a new paediatric test set. What should the regulatory lead advise?

- A. Implement the change under the authorised PCCP, because the retrained model meets every pre-specified acceptance criterion
- B. Implement the change and record the rationale in an internal change file, because the device's core function is unchanged
- C. Implement the change once each customer hospital has completed local validation on its own paediatric images
- D. Prepare a new marketing submission, because a new patient population is outside the authorised modifications

Answer: D. A PCCP only authorises the modifications it describes. Extending the intended population to children changes the indications for use and is outside a plan limited to retraining for adults, so a new submission is needed. Meeting acceptance criteria (A) matters only for changes within scope, and local validation (C) is good practice but no substitute for authorisation.

8 Sample Questions (continued)

Sample question 3 HAIGP-5.2 Analysis and Evaluation

A Medicare Advantage plan uses an AI model that predicts length of stay in skilled nursing facilities. Utilisation reviewers routinely issue denials of continued coverage once a member exceeds the predicted stay. A member's family complains that the denial ignored a documented post-surgical wound infection. As the plan's AI governance lead, which corrective action best addresses the compliance gap?

- A. Retrain the model on more recent claims data to improve the accuracy of its length of stay predictions for members
- B. Add a plain-language explanation of the model output and its confidence level to every denial letter sent to members
- C. Base each decision on the member's clinical circumstances and review past denials that relied on the prediction
- D. Raise the model's prediction threshold so that fewer members are flagged for review of continued skilled nursing coverage

Answer: C. CMS expects medical necessity determinations to rest on the individual member's circumstances, with an algorithmic prediction used only as an input, so the decision process and past denials must be fixed. Better denial explanations (B) support CMS-0057-F denial reason duties but leave the unlawful basis in place. Retraining or threshold changes do not address the process failure.

Sample question 4 HAIGP-8.3 Analysis and Evaluation

A health system pilots a care navigation agent that reads patient portal messages and referral letters, books and cancels appointments, and drafts referral orders for clinician signature. During red teaming, a referral letter containing hidden text instructs the agent to cancel all appointments in a clinic list, and the agent does so. Which control set most directly addresses the root cause while keeping the intended service?

- A. Treat document text as untrusted, limit tools to the requesting patient, and require approval for irreversible or multi-patient actions
- B. Add an instruction to the system prompt telling the agent to ignore any instructions it finds in referral letters, portal messages or attachments
- C. Fine-tune the underlying model on a large set of examples of injected instructions so that it learns to recognise and refuse them
- D. Switch the agent to read-only access for every patient record until the model vendor releases a version that resists injection

Answer: A. Indirect prompt injection cannot be reliably prevented inside the model, so the defence is architectural: untrusted inputs, least-privilege tool scope and human checkpoints on irreversible actions limit what any injection can do. Prompt instructions (B) and fine-tuning (C) reduce but do not remove the risk. Read-only access (D) stops the service and waits on a vendor fix that may not come.

Sample question 5 HAIGP-6.3 Analysis and Evaluation

A health system audits 200 randomly sampled ambient scribe notes each month, and the clinically significant error rate is within the agreed threshold. Clinicians at a community clinic report that medication omissions seem more frequent in consultations conducted through interpreters, which make up about 6% of encounters system-wide. What should the AI governance lead do next?

- A. Accept the result, because the overall error rate in a properly random monthly sample is within the approved threshold
- B. Commission a stratified audit that oversamples interpreter-mediated consultations, focused on medications and allergies
- C. Suspend the scribe across the whole health system until the vendor retrains and revalidates its speech recognition model
- D. Ask the vendor for its published accuracy figures for multilingual speech recognition and file them with the safety case

Answer: B. A random sample of 200 contains only about 12 interpreter encounters, too few to detect a subgroup problem, so a stratified audit focused on high-risk content is needed before deciding on restrictions. Accepting the overall rate (A) hides a possible equity and safety issue, system-wide suspension (C) is disproportionate without evidence, and vendor figures (D) do not show local performance.



9 Eligibility

Certification is open to anyone who completes the training requirement. There is no academic or years-of-experience requirement.

Requirement	Status	Detail
Training	Required	32 contact hours of HAIGP training, completed before the exam, by any of the three routes below. Hours are counted per certification and do not transfer between certifications.
Experience	Recommended, not required	At least five years in healthcare, life sciences, health IT or a related governance, risk, quality, regulatory, privacy or clinical role, with some exposure to AI or digital health, is recommended but not required. Experience is never used to refuse a candidate.
Prior credential	Recommended, not required	GAICC Certified Professional in AI Governance (CPAIG) or equivalent foundational AI governance knowledge.

Three equal training routes

Training route	What it involves
GAICC self-paced course	Structured eLearning on the GAICC platform with a certificate of completion.
GAICC Authorised Training Partner or Certified Instructor	Instructor-led, in person or live online, delivered to the 32-hour GAICC curriculum.
GAICC-recognised institution	A university, professional body or employer programme that GAICC has recognised as covering the Examination Content Outline.

No route gives an examination advantage. Certification decisions are made independently of training.

10 Maintaining the Certification

- The credential is valid for 3 years from the certification decision.
- Earn 60 CPD hours in each three-year cycle, split Technology 27, Humanity 12 and Strategy 21.
- Hours can be earned through healthcare AI governance work such as risk and impact assessment, assurance and audit, incident handling, policy and control design, teaching, writing, standards work and structured learning.
- Reaffirm the GAICC Code of Ethics, remain in good standing and pay the renewal fee.
- A three-month grace period applies after expiry. After the grace period, re-examination is required.
- Re-examination is also required where CPD evidence is insufficient, the credential has lapsed, or a major scheme change requires it.

Competency	What it covers	CPD hours
Technology	AI systems, data, models, tooling and technical controls.	27
Humanity	Ethics, human oversight, fairness, transparency and impact on people.	12
Strategy	Governance, policy, risk, assurance and organisational alignment.	21
Total per cycle	No elective category. Every hour maps to one competency.	60

The GAICC Candidate Handbook and the GAICC Recertification Procedure (GAICC-QMS-PRO-REC-001) are the authority on renewal.

11 Frequently Asked Questions

Do I need work experience?

No. Experience is recommended, not required, and is never used to refuse a candidate. The only eligibility requirement is the 32 contact hours of training.

Do I need the CPAIG before the HAIGP?

No. The CPAIG or equivalent foundational knowledge is recommended because the HAIGP builds on core AI governance concepts and applies them to healthcare and life sciences.

Does this cover medical devices only?

No. AI medical devices are a major part, but the credential gives equal weight to non-device AI: ambient scribes, patient messaging, payer utilisation management, population health models, operational tools and drug development AI. It also covers how a tool can drift into device regulation as its functions change.

I work for a technology vendor or IT services firm. Is this for me?

Yes. Much health AI is built, integrated or run by vendors and services providers. The course covers when a supplier becomes a provider under the EU AI Act, processor and business associate duties, evidence packs and model cards, change notice, and how to support a client's clinical safety case and ISO/IEC 42001 programme.

Do I need a clinical background?

No. The credential is designed for governance, regulatory, quality, privacy, technology and clinical professionals alike. Clinical concepts such as safety cases, automation bias and validation designs are taught from first principles, and mixed teams in the course reflect how health AI committees work in practice.

Which jurisdictions does the examination cover?

The examination tests principles that work across jurisdictions, using the US, EU and UK as primary reference regimes, with Australia, Singapore, Malaysia, India and New Zealand where they differ. Candidates are not expected to recall every local rule, but must recognise which regime applies and what it requires.

Is it relevant to payers and pharmaceutical companies?

Yes. Payer topics include AI in prior authorisation and medical necessity decisions, proxy discrimination and specific denial reasons. Life sciences topics include the EMA-FDA principles for AI in drug development, credibility assessment, GxP validation and AI in pharmacovigilance and manufacturing.

Where are the fees?

Examination, membership and renewal fees are published in the applicable GAICC Candidate Handbook and on the GAICC website. To maintain a single source of truth, fee amounts are not reproduced in this document.

certify@gaicc.org | support@gaicc.org | gaicc.org



Global AI
Certification Council

GAICC HAIGP

HEALTHCARE AI GOVERNANCE PROFESSIONAL

For senior practitioners who govern AI across healthcare and life sciences, from regulatory qualification and clinical safety to validation, monitoring and agentic workflows.

gaicc.org | certifications.gaicc.org | certify@gaicc.org