Supplementary Material

The following tables provide detailed supporting material, scoring rubrics, per-study enumerations, and extended comparative and risk analyses, that underpins the synthesis but is not essential to the main narrative. They are referenced from the main text as Tables S1–S18.

Table S1 Three-level ordinal rubric for thematic coverage assessment (Table 1)

Rating	Score	Decision criterion	Operational definition
✓ Covered	2	The dimension is a named focus or substantive contribution of the paper	The term (or close synonym) appears in the title, abstract, or dedicated section; the paper proposes, evaluates, or critiques the concept; at least one finding directly addresses it
! Partial	1	The dimension is mentioned but not systematically addressed	The concept appears in passing (for instance, in background, limitations, or future work) but is not investigated as a primary or secondary objective; no dedicated method, result, or recommendation is provided
✗ Absent	0	The dimension is not addressed	No substantive mention in title, abstract, or body; or only incidental use of a keyword without conceptual engagement

Two reviewers rated each dimension independently. Disagreements resolved by consensus. $\kappa = 0.79$ (95 % CI: 0.71–0.87), indicating substantial agreement.

Table S2 Clinical studies demonstrating personalisation technologies, classified by primary clinical domain. To ensure a consistent level of analysis and comparability, the clinical domain is assigned at the level of medical specialty; study-specific conditions, populations, and applications are reported separately in the Study Focus column, and the AI/clinical task type is reflected there (e.g. screening, monitoring, outcome prediction, treatment delivery).

Study Focus	Clinical Domain	Technology Type	Implementation Stage
Personalised lung cancer care (Welford et al. 2023)	Oncology	Digital database, patient portal	Active, not recruiting
Remote cardiovascular monitoring (Kennel et al. 2022)	Cardiology	ECG monitor, remote algorithm	Recruiting
Digital mental health service (Lattie et al. 2022)	Psychiatry & Mental Health	SMS messaging, coaching	Completed
Remote monitoring after transplant (Chang et al. 2021)	Transplant Medicine	Telehealth platform	Completed
Remote monitoring for elderly (Ahmed et al. 2023)	Geriatric Medicine	Connected devices, coordination	Not yet recruiting
Digital primary care (Xiong et al. 2023)	Primary Care	Digital clinic, app, website	Not yet recruiting
Paediatric sepsis screening (Gilholm et al. 2023)	Paediatrics	AI tool, EHR integration	Active, not recruiting
Telemedicine prenatal care (Iacoban et al. 2024)	Obstetrics	Remote monitoring, telehealth	Completed
Telemedicine for depression (Shih et al. 2023)	Psychiatry & Mental Health	Telemedicine platform	Recruiting
Predicting therapy outcome (Rost et al. 2023)	Psychiatry & Mental Health	AI models, dataset analysis	Active, not recruiting
Personalised vision/cognitive care (Muntean et al. 2023)	Neurology	Software platform, remote monitoring	Not yet recruiting
CBT training tool (Schaeuffele et al. 2024)	Psychiatry & Mental Health	NLP, ML, simulation	Completed
Personalised medication (AML) (Rodríguez-Maciñeiras et al. 2025)	Oncology	ML, genomics	Not yet recruiting
AI for skin disease (Wang et al. 2024)	Dermatology	AI, augmented reality, imaging	Recruiting
Digital prehabilitation (Blumenau Pedersen et al. 2023)	Cardiac Surgery	Mobile app, telehealth	Not yet recruiting
Digital treatments for substance use (Monarque et al. 2023)	Addiction Medicine	Mobile app, treatment modules	Active, not recruiting

Table S3 Analytical scoring rubric for expert consensus dimensions — Part A: 1–5 ordinal scale (Platform Suitability Score, Table 2). Anchor definitions agreed by two independent reviewers prior to scoring; residual disagreements (> 0 points) resolved by structured discussion. Intermediate scores (2, 4) assigned proportionally between anchors.

Dimension	Score 1 (lowest)	Score 3 (moderate)	Score 5 (highest)
Service breadth (healthcare-specific managed services)	No dedicated healthcare services; generic cloud infrastructure only	A small suite of healthcare-relevant managed services with limited clinical scope	Comprehensive portfolio of named healthcare services covering EHR, imaging, NLP, and IoT
FHIR/HL7 interoperability	No native FHIR or HL7 support; requires third-party middleware for all health data exchange	Partial FHIR support (e.g., via marketplace or limited API) or HL7v2 only	Native FHIR R4 server with SMART on FHIR, HL7v2, and DICOM natively supported
AI/ML clinical capability	General-purpose ML tooling only; no clinical AI models or health-domain NLP	General ML platform with some health-adjacent models (e.g., NLP) but no validated clinical AI	Validated clinical AI models (medical imaging, NLP extraction, clinical LLMs) with active healthcare deployment evidence
HIPAA/GDPR compliance coverage	No HIPAA BAA available; GDPR compliance unspecified or limited to select regions	HIPAA BAA available for core services; GDPR compliance in major EU regions with some service exclusions	HIPAA BAA across $\geq$100 services; full GDPR data residency controls with EU Sovereign Cloud option
EHR integration depth	No documented EHR connectors; integration requires fully custom development	API-based integration with one or two major EHR vendors; limited out-of-the-box connectors	Deep native connectors to leading EHR platforms (Epic, Cerner) plus open FHIR APIs for others
Genomics/omics support	No genomics or omics-specific tooling; standard object storage only	Partner-based genomics pipelines or limited WGS/WES support via marketplace	Dedicated managed genomics service supporting WGS, WES, multi-omics pipelines, and reference datasets
Vendor lock-in risk (inverted: 5 = lowest risk)	Pervasive proprietary APIs, formats, and services with no open-standard export mechanisms	Mix of proprietary and open-standard components; FHIR export available but not default	Open-standard first design; FHIR R4, open-source components, and multi-cloud portability by default

Scoring process: two reviewers (A.A., M.A.-R.) scored independently; disagreements resolved by discussion; T.J. available as arbiter. See Part B (Table S4) for 0–10 interval scale.

Table S4 Analytical scoring rubric for expert consensus dimensions — Part B: 0–10 interval scale used in Tables 3, S17, S15, 4, 7, and 8. Intermediate values interpolated proportionally between anchors. TRL scores follow the European Commission Horizon 2020 TRL scale and are not consensus-scored.

Dimension	Low: 0–2	Mid: 4–6	High: 8–10
Implementation Urgency (Table 3)	Capability is a future enhancement; clinical workflows function adequately without it	Capability gaps cause periodic audit failures or compliance friction but no immediate clinical risk	Absence directly impedes AI auditability, regulatory compliance, or patient safety in active deployments
Regulatory Penalty (Table 3)	No known regulatory sanction for absence; voluntary best practice only	Regulatory guidance recommends the capability; non-compliance may draw scrutiny but carries no automatic penalty	Absence constitutes a demonstrable breach of GDPR, HIPAA, or equivalent statutory instrument with enforceable penalties
Security Risk (Table S17)	Absence creates a theoretical vulnerability with very low exploitation probability given other controls	Absence creates a realistic attack vector or privilege escalation path that existing controls do not fully mitigate	Absence constitutes a critical unmitigated vulnerability; exploitation probability is high and impact is patient-safety-level
Workflow Impact (Table S17)	Over-restriction causes minor inconvenience; clinicians can accomplish tasks via documented workarounds	Over-restriction adds measurable steps to clinical workflows; documented workarounds increase error risk	Over-restriction blocks time-critical clinical access or forces unsafe workarounds (e.g., shared credentials)
Regulatory Maturity (Table S15)	No national regulatory framework; ad hoc oversight only	Framework in development or partially implemented; classification guidance exists but enforcement is inconsistent	Mature statutory framework with mandatory pre-market assessment, post-market surveillance, and enforcement precedent
Reimbursement Readiness (Table S15)	No reimbursement pathway; all costs borne by patient or institution	Ad hoc or pilot reimbursement available in selected jurisdictions; no national pathway	National reimbursement pathway established with defined eligibility criteria and outcomes-linked renewal
Regulatory Obligation (Table 4)	Voluntary guidance only; no legal instrument mandates the governance dimension	Sector-specific guidance with regulatory expectation but no statutory mandate	Explicit statutory mandate under enforceable legislation (e.g., EU AI Act, GDPR Art. 22, MDR 2017/745)
Governance Maturity (Table 4)	No governance standard exists; practices are entirely ad hoc	Emerging standards or frameworks proposed but not yet widely adopted or enforced	Mature, widely adopted governance standard with enforcement mechanisms and auditable compliance criteria
Integration Deficit Score — IDS (Table 7)	Dependency substantively addressed in $\geq$20 % of relevant corpus studies; integration mechanisms specified	Dependency acknowledged in corpus but implementation mechanisms are absent; conceptual references only	Zero corpus studies specify the integration mechanism; the dependency is structurally unaddressed across the entire evidence base
Outcome Strength (Table 8); scale 1–10	Conceptual or theoretical proposal only; no empirical evidence (score 1–2)	Observational, case-based, or survey evidence; no controlled trial (score 4–6)	RCT-validated or prospective cohort evidence with reported effect size and confidence interval (score 8–10)

Scoring process: two reviewers (A.A., M.A.-R.) scored independently; disagreements (>1 point) resolved by structured discussion; T.J. available as arbiter.

Table S5 Analytical Assessment of Predictive Modelling Approaches for Longitudinal Patient Trajectories in Hospital-of-the-Future Systems

Modelling approach	Core temporal assumption	Uncertainty method	Calibration requirement	Typical clinical application	Authors' TRL assessment
Cox proportional hazards	Constant covariate effect over time	Bootstrap CI	Hosmer–Lemeshow test	Readmission, mortality prediction	TRL 7–8 (deployed)
LSTM / RNN	Sequential dependency; no stationarity required	MC Dropout / ensemble	ECE across deciles	ICU deterioration, sepsis screening	TRL 6–7 (validated)
State-space (Kalman / HMM)	Latent Markov state evolution	Posterior distribution	State-conditional calibration	Wearable fusion, remote monitoring	TRL 5–6 (piloted)
Transformer (temporal attention)	Global temporal dependencies via self-attention	Conformal prediction	Reliability diagrams	EHR sequence modelling, polypharmacy	TRL 4–5 (emerging)
Bayesian hierarchical model	Partial pooling across patient subgroups	Full posterior	MCMC convergence + calibration	Personalised dosing, subgroup analytics	TRL 4–5 (emerging)

TRL = Technology Readiness Level. Authors' assessment based on deployment stage evidence in the 116-study corpus.

Table S6 Implementation Frameworks for Personalisation Technologies in Clinical Settings

Technology/Study Focus	Data Sources	Interoperability	Ethical Considerations	Usability/Patient Interface	Clinician Integration
Personalised lung cancer care	EHR, patient-reported data, imaging	Linked with portals and hospital databases	Consent management, data sensitivity	Web portal with structured inputs	Integrated into oncology workflows
Remote cardiovascular monitoring	ECG telemetry, procedural data	Cloud-linked with remote diagnostics	Risk model transparency, alert fatigue	Mobile app with real-time feedback	Cardiologist alert dashboard
Predicting therapy outcome	Text transcripts, engagement logs	Federated data learning	Bias mitigation, informed opt-in	Minimal interface; back-end inference	Not yet integrated
AI for skin disease	Dermatoscopic imaging, patient input	Partial EHR linkage	Algorithmic bias, informed consent	Augmented reality interface	Dermatology decision support
Personalised medication (AML)	Genomic sequencing, drug response	Not yet standardised	Genomic privacy, ethics	Clinician-led interface	High integration with haematology systems
Paediatric sepsis screening	Vital signs, EHR sensor inputs	Full EMR embedding	Parental consent, paediatric autonomy	Passive (backend)	Integrated into paediatric emergency alert systems
CBT training (Psychosis)	NLP on simulated speech	Offline simulation	Simulation ethics	Simulated patient dashboard	Clinician training
Vision/cognitive care (MCI)	Wearables, memory assessments	Edge-cloud synchronisation	Digital dignity, cognitive load	Simplified interface with alerts	Limited neurology EHR integration

Table S7 Comparative analysis of major cloud platforms for Hospital-of-the-Future implementations across seven critical dimensions. AWS services are used as illustrative examples in Sections 4.3–4.7; equivalent capabilities exist across all platforms listed.

Platform	Healthcare services	FHIR/HL7	AI/ML (clinical)	HIPAA/GDPR	EHR integration	Genomics	Lock-in risk
AWS	HealthLake, Comprehend Medical, HealthOmics, SageMaker, HealthScribe	Native FHIR R4; HL7v2 via HealthScribe	SageMaker AutoML, Bedrock LLMs, Comprehend Medical NLP	165+ HIPAA-eligible services; GDPR data residency (Kuo 2011b)	Open APIs; Epic, Cerner connectors	HealthOmics WGS/WES (Subramanian and Ramamoorthy 2024)	High (Opara-Martins et al. 2016)
Microsoft Azure	Health Data Services, DICOM Service, MedTech IoT, Azure OpenAI	Managed FHIR R4; SMART on FHIR (Mandel et al. 2016a)	Azure OpenAI GPT-4, Azure ML, Health Bot	HIPAA BAA; EU data boundary for GDPR	MS Teams, Epic connector	Cromwell on Azure	Moderate
Google Cloud	Cloud Healthcare API, Vertex AI, BigQuery, MedPaLM 2 / Med-Gemini	HL7v2, FHIR R4, DICOM natively (Rajpurkar et al. 2022)	Med-Gemini, Med-PaLM 2, AutoML Vision	HIPAA BAA; EU Sovereign Cloud	Cloud Healthcare API; Apigee HL7 routing	Google Life Sciences; AllOfUs	Moderate
IBM Cloud	Watson Health (Merative), IBM Cloud Satellite, Watson Assistant	HL7 FHIR via IBM FHIR Server (open-source) (Dash et al. 2019a)	Watson Oncology, Watson Clinical Trial Matching	HIPAA BAA; GDPR EU regions	Deep Cerner/MEDITECH heritage	IBM Genomics (partner-based)	Moderate
Oracle Health	Oracle Health EHR (Cerner), Clinical Digital Assistant, OCI	Native Cerner FHIR R4 APIs	Oracle Clinical AI, Generative AI Service	HIPAA BAA; GDPR EU regions	Tightest EHR-to-cloud path (native Cerner)	Health Data Intelligence	High (Cerner) / Low (greenfield)
Alibaba Cloud	Healthcare Brain, Medical IoT, EHOP	HL7 FHIR; limited Western documentation	PAI platform; medical imaging AI	HIPAA limited; strong China regulatory scope	Strong Asia-Pacific HIS systems	Alibaba Health omics	High
Huawei Cloud	EIHealth, HIS connector	HL7 FHIR via EIHealth	ModelArts ML; EIHealth drug discovery AI	GDPR EU regions; HIPAA limited	HIS/EMR connectors for Chinese hospitals	EIHealth gene sequencing	High
Snowflake	Healthcare & Life Sciences Data Cloud, Clinical 360	FHIR sharing via Marketplace	Cortex AI; partner ML integrations	HIPAA BAA; GDPR compliant	EHR data sharing via Data Clean Rooms	Genomics data marketplace	Low (multi-cloud)
Salesforce Health Cloud	Patient 360, Care Management, Health Timeline	FHIR R4; MuleSoft HL7	Einstein AI; NLP for clinical notes	HIPAA BAA; GDPR Shield	Deep CRM-to-EHR; Epic/Cerner connectors	Limited	Moderate
SAP Health	S/4HANA Healthcare, SAP Health Engagement	HL7 FHIR via SAP Integration Suite	SAP AI Core; predictive analytics	HIPAA BAA; GDPR compliant	ERP-to-EHR integration (finance/ops)	Limited	Moderate
Philips HealthSuite	HealthSuite Digital Platform, IntelliSpace Portal	FHIR R4 APIs; DICOM for imaging	AI diagnostic imaging; remote monitoring AI	HIPAA BAA; CE-marked medical device software	Deep imaging workflow (PACS, RIS)	Precision diagnosis pipelines	Low–Moderate

Sources: Kuo (2011a); Kruse et al. (2016); Dash et al. (2019b,a); Opara-Martins et al. (2016); Mandel et al. (2016a); Rajpurkar et al. (2022); Kuo (2011b)

Table S8 Edge deployment latency benchmarks and safety-critical response constraints for Hospital-of-the-Future AI applications. T_{req} = maximum permissible end-to-end latency; $\bar{\delta}_{\text{edge}}$ = measured median edge latency; $\bar{\delta}_{\text{cloud}}$ = measured median cloud latency; P_{fail} = probability of deadline violation.

Application	T_{req} **(ms)**	$\bar{\delta}_{\text{edge}}$ **(ms)**	$\bar{\delta}_{\text{cloud}}$ **(ms)**	$P_{\text{fail}}^{\text{edge}}$	$P_{\text{fail}}^{\text{cloud}}$	**Standard**
ICU EWS detection	500	18–45	320–890	<0.001	0.31–0.67	IEC 62443-3-3
Closed-loop insulin MPC	200	8–22	90–300	<0.001	0.12–0.48	ISO 14971
Surgical haptic feedback	100	7–14	45–120	<0.001	0.05–0.32	IEC 80601-2-77
Wearable arrhythmia alert	1,000	35–80	200–600	<0.001	<0.01	IEC 60601-2-47
Neonatal apnoea detection	300	12–30	150–500	<0.001	0.09–0.41	ISO 14971

Table S9 Human–AI task allocation framework for safety-critical edge AI deployments. Automation tier determines whether AI output is acted upon without human review (Tier A), with asynchronous review (Tier B), requires synchronous confirmation (Tier C), or informs human decision only (Tier D).

Decision type	**Example**	**AI accuracy**	**Risk**	**Tier**	**Max latency**
Alarm suppression	Non-actionable artefact removal	$F_1 \geq 0.91$	Low	A	50 ms
Deterioration alert	EWS threshold breach	Sens. ≥82 %	Medium	B	500 ms
Drug dose adjustment	Insulin pump correction	TIR ≥74 %	High	C	200 ms
Diagnosis	Sepsis classification	Spec. ≥91 %	High	C	2 min
Treatment initiation	Antibiotic prescribing	N/A	Critical	D	Human-only
Surgical actuation	Robot arm movement	N/A	Critical	D	Human-only

Table S10 Energy and carbon accounting framework for a 500-bed Hospital-of-the-Future AI deployment. Baseline figures derived from Equations 21–27. Mitigation savings are additive; combined application of all three levers reduces total annual carbon by 94.2 %.

Component	**Baseline (kWh/yr)**	**Baseline (kgCO_2e/yr)**	**Renewable saving**	**Federated saving**	**Edge-first saving**	**Optimised (kgCO_2e/yr)**
Edge nodes	1,261	294	−265	—	+206	235
Network	10,896	2,539	−2,285	−762	—	−508
Cloud compute	15,505	3,613	−3,252	—	−2,529	−2,168
Clinical devices	1,284	299	−269	—	—	30
Total	**28,946**	**6,744**	**−6,071**	**−762**	**−2,323**	**≈392**

Table S11 Risks of Depersonalisation in Digital Healthcare

Risk Area	Description	Consequences for Patients	Consequences for Clinicians	Potential Mitigation
Screen-mediated interactions	Clinicians focus on electronic records instead of patients	Patients feel unseen, loss of trust	Reduced attention to non-verbal cues	Adaptive interfaces, ambient documentation tools
Algorithmic decision-making	Automated recommendations without transparency	Reduced patient autonomy	Clinical deskilling, blind reliance	Explainable AI, shared decision dashboards
Remote consultations	Absence of physical presence	Reduced empathy, perceived distance	Loss of rapport, contextual awareness	Hybrid models with in-person follow-up
Administrative burden	Excessive data entry and navigation	Delays in care, fragmented communication	Burnout, less time for empathy	Workflow optimisation, voice-enabled documentation
Over-standardisation of care	Protocol-driven pathways override personalisation	One-size-fits-all care	Reduced professional judgement	Adaptive clinical pathways, personalised medicine
Data overload	Continuous data streams overwhelm interpretation	Confusion, anxiety over raw data	Alert fatigue, decision paralysis	Intelligent filtering, context-aware alerts
Privacy intrusions	Continuous monitoring feels invasive	Erosion of dignity, mistrust of systems	Ethical stress, liability concerns	Transparent consent, granular data control
Digital exclusion	Barriers for older adults or disadvantaged groups	Unequal access, health disparities	Frustration with inequitable outcomes	Inclusive design, training, alternative access modes
Reduced human contact	Increased reliance on automation in routine tasks	Sense of isolation, loss of comfort	Deskilling, detachment from patient care	Embedding face-to-face time, empathy training

Table S12 Analytical Assessment of Co-Design Evaluation Evidence in the Included Corpus: Structured Framework, Validated Instruments, and Comparative Evaluation

Study / context	**Co-design approach**	**Structured iterative cycle**	**Validated outcome instrument**	**Comparative usability evaluation**	**Authors' rating**
Foresman et al. (2025) — patient AI transparency	Focus groups; XAI interface co-specification	Partial (2 rounds)	None reported	None	Weak
Marko et al. (2025) — inclusive interface design	Participatory workshops with marginalised users	Partial (iterative)	Accessibility checklist (non-validated)	Pre/post within-group only	Moderate
Han et al. (2023) — patient preference elicitation	Structured preference elicitation protocol	Yes (full cycle)	Discrete choice instrument	No comparator arm	Moderate
Binzer et al. (2024) — health app trust	Secondary survey analysis	N/A (survey)	Trust in Health Technology scale	Public vs. commercial provider	Strong
Corpus-level gap (authors' assessment)	*Most co-design studies lack full iterative cycles, use non-validated instruments, and provide no comparative arm*				*Field gap*

Authors' rating: Strong = all three components present; Moderate = two; Weak = one or none.

Table S13 Analytical Assessment of Empathy-Driven Interface Design Evidence: Usability Heuristics, Cognitive Workload, and Comparative Evaluation in the Included Corpus

Study / interface type	Design claim	Heuristic evaluation reported	Cognitive workload instrument	Controlled comparative evaluation	Authors' evidence rating
Foresman et al. (2025) — XAI patient portal	Transparency improves trust	None	None	Pre/post qualitative only	Weak
Marko et al. (2025) — inclusive digital health	Accessibility reduces exclusion	Accessibility audit (partial)	None	None	Weak
Allande-Cussó et al. (2025) — digital workload	HCD reduces burnout	None	Burnout inventory (indirect)	Pre/post observational	Moderate
Han et al. (2023) — patient portal preferences	Co-design improves usability	None	None (preference only)	No comparator	Weak
Cho et al. (2022) — nursing CDSS usability	Heuristics surface design flaws	Nielsen's 10 heuristics (6 experts; severity 0.66–2.00)	None	Expert inspection (no user comparator)	Moderate
Kunz et al. (2025) — CKRT ICU CDSS prototype	HCD lowers workload	ISO 9241-11 walkthrough	SUS (median 87.5), UEQ, modified NASA-TLX	Simulated single-arm (n small); no A/B	Moderate
Corpus-level gap (authors)	*Instrument gap now partially addressed via other-methods studies (heuristic evaluation, SUS, and NASA-TLX reported); a controlled A/B interface comparison in a live clinical setting remains absent*				*Residual gap*

Authors' rating: Strong = heuristic + workload + comparator; Moderate = two present; Weak = one or none.

Table S14 Human-Centred Design Strategies in Future Hospitals

Domain	Humanisation Focus	Example Interventions	Expected Outcomes
Patient-facing tools	Empathy and communication	Co-designed patient portals, plain-language health reports	Improved trust, shared decision-making
Clinical interfaces	Usability and transparency	Adaptive EHR dashboards, predictive alerts with context	Reduced workload, safer decisions
Architecture & environment	Healing and dignity	Natural light, green spaces, noise reduction	Reduced stress, faster recovery
Inclusivity & accessibility	Equitable care	Multilingual apps, cognitive-friendly interfaces	Broader patient engagement, reduced exclusion
Staff wellbeing	Burnout prevention	Ergonomic spaces, digital literacy training	Greater resilience, job satisfaction

Table S15 Jurisdictional Readiness Matrix for Digital Therapeutics Governance

DTx governance dimension	Reg. maturity /10 (global)	Reimb. readiness /10 (AU/NZ)	Corpus studies	Coverage rate %	Governance gap /10
Regulatory classification (SaMD)	7	4	3	3	3.79
Clinical validation standards	5	3	0	0	5.00
Reimbursement pathway	3	1	0	0	6.60
Institutional governance / formulary	3	2	0	0	6.20
EHR interoperability (FHIR)	4	2	2	2	5.80

Scores 0–10 per rubric in Table S3. Reg. maturity: global best-practice frontier (EU/US/DE). Reimb. readiness: AU/NZ context specifically.
Governance Gap $= 0.40 \times (10 - \text{Reg. Maturity}) + 0.40 \times (10 - \text{Reimb. Readiness}) + 0.20 \times (1 - \text{Coverage rate})$. Higher score = more critical unaddressed gap.

Table S16 Quantitative risk register for Hospital-of-the-Future EHR migration. Probability $P \in [0, 1]$ and clinical impact $I \in [1, 10]$ are scored by structured expert consensus (adapted from ISO 14971 medical device risk management). Composite Risk Score $R = P \times I$. Post-countermeasure residual risk R' assumes full implementation of the specified countermeasure.

Failure mode	P	I	R	**Countermeasure**	R'
Terminology mapping error (medication dose)	0.18	9	1.62	Dual-coding validation + pharmacist sign-off	0.09
Patient identity mislink (MRN collision)	0.12	8	0.96	Probabilistic linkage (Jaro-Winkler $\geq$0.94)	0.04
Data loss during ETL	0.23	7	1.61	Row-count reconciliation + SHA-256 checksums	0.05
Downtime exceeding 4 h at cut-over	0.31	6	1.86	Blue-green deployment with auto-rollback	0.06
Semantic drift in longitudinal records	0.40	5	2.00	Semantic version control + annual audit	0.20

Table S17 Quantitative Risk-Priority Matrix for RBAC Dimensions in Hospital-of-the-Future Systems (authors' analytical scores, 0–10, per rubric in Table S4). Security Risk = estimated breach probability × impact if dimension is absent. Workflow Impact = clinician friction if over-restricted. Address Rate = proportion of 116-study corpus mentioning dimension. Weighted Priority = 0.45×Security + 0.25×Workflow + 0.30 × (1−Address rate).

RBAC dimension	**Security risk /10**	**Workflow impact /10**	**Corpus studies**	**Address rate %**	**Weighted priority /10**
Break-glass protocol	8	9	0	0	6.15
Privilege granularity (ABAC)	9	7	0	0	6.10
Clinical hierarchy alignment	7	8	2	2	5.44
Dynamic / zero-trust auth.	8	5	1	1	5.15
SIEM audit & monitoring	7	3	6	6	4.18

Security risk and workflow impact scored by structured author consensus. Address rate from full 116-study corpus extraction.

Table S18 Weight-sensitivity analysis of the Platform Suitability Score (PSS) ranking (Table 2). Each cell gives the platform's rank under the stated weighting scheme. Baseline = weights of Table 2; Equal = 1/7 per criterion; Sec+Int = HIPAA/GDPR and FHIR/HL7 weights doubled; Lock-in×3 = lock-in weight tripled; Clin-AI = AI/ML and genomics weights doubled (all schemes renormalised to sum to one). Spearman rank correlation ρ is computed against the baseline ranking.

Platform	**Baseline**	**Equal**	**Sec+Int**	**Lock-in×3**	**Clin-AI**
AWS	1	1	1	1	1
Google Cloud	2	2	2	2	2
Microsoft Azure	3	3	3	3	3
Oracle Health	4	4	4	7	5
IBM Cloud	5	7	5	6	6
Philips HealthSuite	6	6	7	5	4
Snowflake	7	5	6	4	7
Salesforce Health Cloud	8	8	8	8	8
SAP Health	9	9	9	9	11
Alibaba Cloud	10	10	10	10	9
Huawei Cloud	11	11	11	11	10
Spearman ρ vs baseline	1.00	0.96	0.99	0.91	0.95

The top three platforms (AWS, Google Cloud, Microsoft Azure) are invariant across all schemes, and AWS ranks first throughout. Mid-tier reordering is confined to the lock-in-stressed scenario, under which lock-in-favourable platforms (Snowflake, Philips HealthSuite) rise.

Table S19 Traceability from the synthesised corpus to the conceptual framework of analysis (Section 1.1). For each of the eleven analytical dimensions, the table records the evidence basis observed across the corpus, representative grounding studies (illustrative, not exhaustive), and the pillar in which the dimension is principally analysed. The same eleven dimensions are operationalised as the scoring rubric of the coverage analysis (Table 1).

Analytical dimension	**Evidence basis in the corpus**	**Representative grounding studies**	**Pillar**
Personalisation *(core)*	Tailoring of care to individual genomic, physiological, and behavioural profiles	(Godman et al. 2013; Pandey and Gupta 2024; Mani et al. 2025)	P
AI integration	Embedding of ML/DL models into clinical workflows	(Aburass et al. 2025; Choi et al. 2023)	P
Continuous monitoring	Longitudinal physiological sensing from wearable and bedside devices	(Dubey and Dubey 2022; Worlikar 2025)	P
Cloudisation *(core)*	Cloud-native aggregation, governance, and delivery of AI-processed data at scale	(Ahmed and Jayapandian 2024; Lakhan et al. 2023)	C
Interoperability	Standards-based clinical data exchange (HL7 FHIR)	(Mandel et al. 2016a; Pereira et al. 2025)	C
Data security	Protection of patient data under HIPAA, GDPR, and equivalent regimes	(Banerjee et al. 2025; Nowrozy et al. 2024)	C
AI governance & regulation	Accountability, audit, and oversight mechanisms specific to clinical AI	(Tazi et al. 2024; Seeram and Kanade 2024)	C
Sustainability	Energy and resource efficiency as an Industry 5.0 design objective	(Adel 2022; Alves et al. 2023)	C
Humanisation *(core)*	Human-centred design and governance keeping AI empathetic and accountable	(Villela and Ely 2022; Bienefeld et al. 2024)	H
Human–AI collaboration	Principled allocation of decision authority between clinicians and AI	(Bienefeld et al. 2024)	H
Empathy in design	Dignity, communication, and well-being embedded in interfaces and care environments	(Foresman et al. 2025; Villela and Ely 2022)	H