Adding Years of *Healthy Life*

Agentic AI as a Workforce

Building the conditions for a Managed Digital Workforce

An NHG Health strategic white paper for consultation and collaboration, supported by CHI HEAL

24th June 2026

Core provocation

What would need to be true for NHG Health to safely deploy, manage and scale a digital workforce over the next 12 months?

How to engage: *This is a consultation draft, not a final policy*. The aim is to build shared understanding and practical collaboration, so that AI is adopted responsibly, sustainably and in service of better care. NHG Health invites feedback from healthcare institutions, Institutes of Higher Learning, industry partners and the wider ecosystem to inform our strategic consideration on how we design, deploy and govern digital co-workers safely and effectively.

Comments may be submitted to CHI.HEAL@nhghealth.com.sg.

Contributions will inform further iteration of our strategy and a more detailed delivery roadmap which we intend to publish by end July 2026.

1 Mandalay Road, Singapore 308205
www.nhghealth.com.sg

Executive summary

Agentic AI marks a shift from AI that is prompted to AI that can take autonomous action. As models acquire memory, tool use, planning, retrieval, and multi-step reasoning capabilities, the strategic question for healthcare is no longer simply how to provide staff with better tools or how to deploy validated AI products. The deeper question is how to design, govern, supervise, and improve work when people and AI agents operate together inside real clinical, operational, research, and corporate workflows.

This white paper proposes that NHG Health should understand this next stage as the emergence of a new workforce capability: managed digital workers operating alongside people within redesigned workflows. This is the idea of AI as a Workforce. It does not replace AI as a Tool or AI embedded in Workflows; it extends them.

The proposition aligns closely with NHG Health's existing approach through HEAL: AI should be anchored in real health system priorities, built with trust and governance from the start, and moved beyond experimentation towards scalable impact. It also aligns with national direction. Singapore's refreshed AI posture is moving towards national AI missions, with healthcare explicitly identified as a priority sector. The opportunity for NHG Health is to translate this national ambition into a practical operating model for responsible AI deployment in health.

The central question this paper asks is: what would need to be true for NHG Health to safely deploy, manage and scale a digital workforce over the next 12 months?

The answer is not a single technology choice. It requires a linked set of capabilities: workflow redesign, human-AI collaboration, agent identity and access, sandbox and validation environments, orchestration and observability, lifecycle governance, cybersecurity by design, cost and model routing discipline, internal AI engineering, clinical AI leadership, and a platform model that allows development to happen close to the work while remaining governed through common standards. Additionally, we recognise that advances in AI raise important questions for our staff: about job security, about how roles will change, and about what the future of work looks like in healthcare. This paper addresses those questions directly, It offers is a clear strategic frame, an honest account of both opportunity and risk, and a set of explicit commitments on how NHG Health will manage this transition responsibly.

The paper therefore moves from concept to conditions. It is not intended to be a detailed delivery roadmap, but to establish the strategic frame and operational requirements that can inform the roadmap: what needs to exist, what decisions need to be made, and what capabilities NHG Health needs to build if agentic AI is to become a safe, useful, accountable, and scalable part of the healthcare workforce.

Thesis

AI as a Workforce in Healthcare is not primarily a technology strategy. It is a business and care transformation strategy, enabled by technical infrastructure, responsible governance, clinical leadership cybersecurity, workforce capability, and disciplined execution.

1. The emergence of a new workforce capability

AI capabilities have evolved rapidly from predictive models and generative assistants towards agents that can reason, remember, use tools, retrieve information, interact with systems, and undertake multi-step tasks. The change is not simply that models have become more capable. The change is that AI is beginning to move from support to participation: from systems that answer questions to systems that can act within workflows.

Healthcare organisations must ask how work itself should be organised when selected tasks can be undertaken by managed digital workers operating alongside people. The phrase 'AI as a Workforce' is intentionally practical. It is not intended to anthropomorphise agents or imply that agents are equivalent to people. It is a way of making visible the workforce management functions that become necessary when digital systems can take action.

The workflow becomes the strategic unit of change. Agentic AI will not create durable value if bolted onto old processes. It creates value when work is intentionally redesigned around the complementary strengths of people and AI.

Augmentation of the Workforce, Not Replacement.

We are aware that the language of 'AI as a Workforce' raises a legitimate concern: does this mean AI replaces our workforce? The answer is no. AI is deployed strictly to augment our human workforce, never to replace it. Our goal is to absorb administrative burdens so our staff can focus more on human interactions, empathy, and complex clinical judgements that algorithms cannot replicate.

2. Applying the Innovation Cycle to AI: Tool, Workflow, Workforce

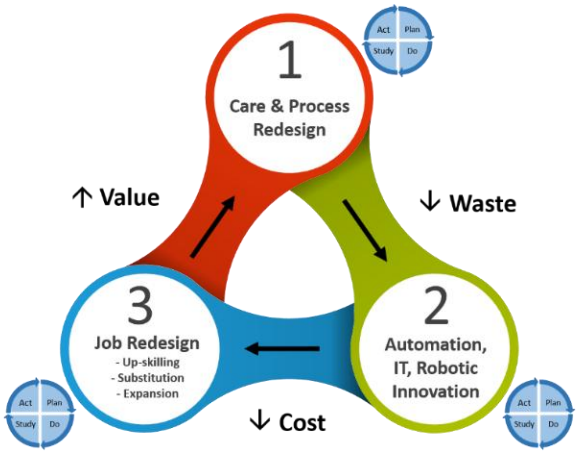

Centre for Healthcare Innovation's Innovation Cycle

NHG Health and CHI already have an established approach to innovation and transformation through the Innovation Cycle: care and process redesign, automation and technology innovation, and job redesign. Each

move enables and depends on the others; none, on its own, produces durable change. This paper extends that approach into the agentic AI era.

The cycle helps prevent a common mistake: assuming that technology implementation alone creates transformation. In healthcare, value emerges when technology is linked to redesigned care models, redesigned processes, and redesigned jobs. Agentic AI makes this even more important because it changes not only the toolset, but the distribution of work between people, teams, systems, and digital workers.

This also helps avoid another important trap: assuming that agentic AI should be the default answer to every workflow challenge. In many cases, conventional automation, robotic process automation (RPA), low-code tools or workflow redesign will be sufficient to improve efficiency, reliability and productivity. These capabilities are not made less important by agentic AI; they are often the foundation that makes agentic AI easier to apply safely. By making processes more explicit, structured and measurable, automation can create the conditions in which agents can later augment more complex, context-dependent or multi-step work.

CHI has a strong track record supporting teams to automate workflows, develop digital skills and grow citizen developer communities. The next stage extends that capability: building AI-fluent champions and workflow owners who can identify where standard automation is enough, where agentic augmentation adds value, and how both should be deployed within shared standards.

NHG Health's AI approach can therefore be described across three connected modes: AI as a Tool, AI embedded in Workflows, and AI as a Workforce. These modes coexist. The purpose of distinguishing them is not to choose one over the others, but to recognise that each has a different unit of value, different governance requirement, and different delivery model.

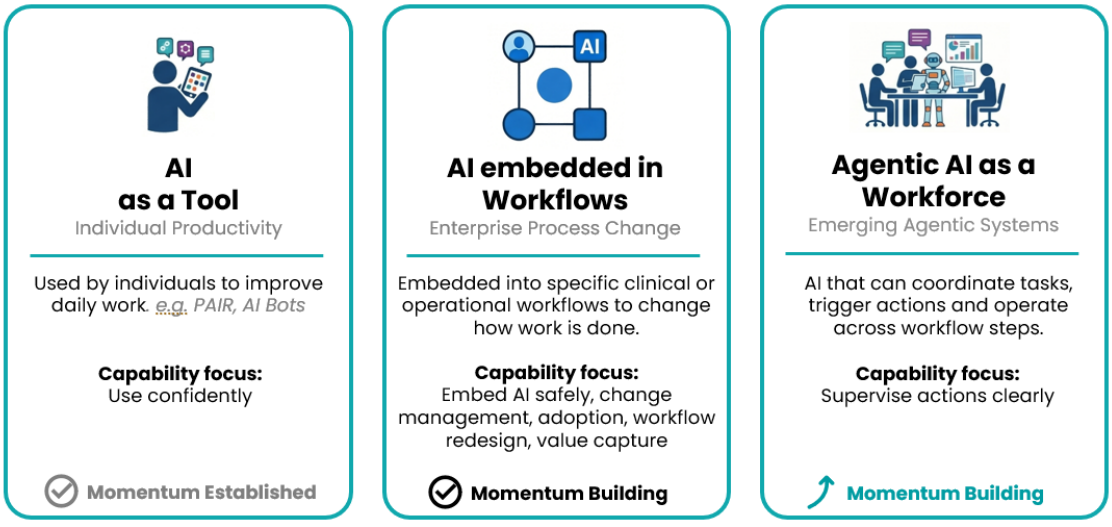

NHG has built momentum in AI as a tool and embedding within workflows; the next shift is AI as a workforce

Mode	Primary purpose	Unit of value	Governance challenge
AI as a Tool	Support individuals in daily work through generative AI, assistants, copilots and productivity tools.	Personal productivity and quality of work.	Responsible individual use, literacy, safe access, privacy and professional judgement.
AI embedded in Workflows	Deploy validated models or AI capabilities into specific clinical or operational workflows.	Workflow improvement, risk reduction, decision support or operational performance.	Validation, clinical accountability, monitoring, safety, integration and model lifecycle management.
AI as a Workforce	Deploy managed AI agents as digital co-workers within redesigned multi-step workflows.	Workflow capability: Redesigned around complementary strengths of people and AI to deliver deeper healthcare outcomes.	Agent identity, role definition, supervision, access control, observability, escalation, performance management and retirement.

AI as a Tool is the most familiar mode: a clinician using an ambient scribe, a manager using a generative assistant to draft a paper, an analyst using AI to summarise information. The governance challenge is responsible use at scale.

AI embedded in Workflows is where AI becomes a validated capability embedded in a defined clinical or operational workflow: imaging triage, risk prediction, prioritisation, sepsis alerts, or operational optimisation. The governance challenge is product-style assurance, post-deployment monitoring, and accountability for safe use.

AI as a Workforce is qualitatively different: agents undertake meaningful portions of multi-step work, while humans retain accountability but shift towards direction-setting, supervision, exception handling and judgement. We acknowledge that workforce upskilling may be required to enable this shift. Overall, this shift is closer to AI in the loop than simply human in the loop.

3. Why now: from experimentation to production

The current moment is defined by a shift from AI experimentation to AI production. Across recent strategy discussions, workshops and selected external learning attended by the authors, several themes have been consistent: AI capability is advancing rapidly; agentic systems are becoming more capable and practical.

One phrase captured the challenge well: some organisations have had 'more pilots than Singapore Airlines', but the task now is to land them. This is highly relevant to health systems, where proof-of-concept activity can easily proliferate without creating durable operational capability.

Another strong message was that the bottleneck is no longer intelligence; it is responsible deployment. The harder work is creating the organisational conditions to use agentic capabilities well: the data and context layer, workflow redesign, platform architecture, cybersecurity model, evaluation approach, staff capability, and governance needed for real-world deployment. Traditional annual funding models compound this challenge by throttling agile projects with rigid annual capital expenditure models.

Singapore's national direction reinforces this shift. The national AI strategy refresh has moved from broad ecosystem building towards mission-oriented deployment in priority sectors, including healthcare.

The message reflects NHG Health's direction that AI must be grounded in real system problems, responsibly deployed, and focused on public value over novelty or productivity alone; it highlights both the opportunity to develop a distinctive, mission-driven approach to agentic AI and the discipline required to ensure true transformation by recognizing that sandbox experiments are not production models and promising agents are not yet managed digital workers.

Strategic provocation

What would need to be true for NHG Health to safely deploy, manage and scale a digital workforce over the next 12 months?

4. AI as a Workforce: the strategic proposition

AI as a Workforce is the intentional deployment of AI agents as managed digital co-workers within redesigned processes, so that selected tasks can be undertaken more effectively by a combination of people and AI. This is not primarily a technology strategy. It is a business, care and workforce transformation strategy, enabled by technical infrastructure, governance and change capability.

The proposition has four implications. First, agents should be designed around workflows, not tasks in isolation. Second, they should be managed through a lifecycle, not released as one-off applications. Third, their value should be measured through outcomes, not novelty. Fourth, the organisation must build the platform and capability to manage agents as a portfolio, not simply approve isolated use cases.

5. Building a managed digital workforce: HR for agents

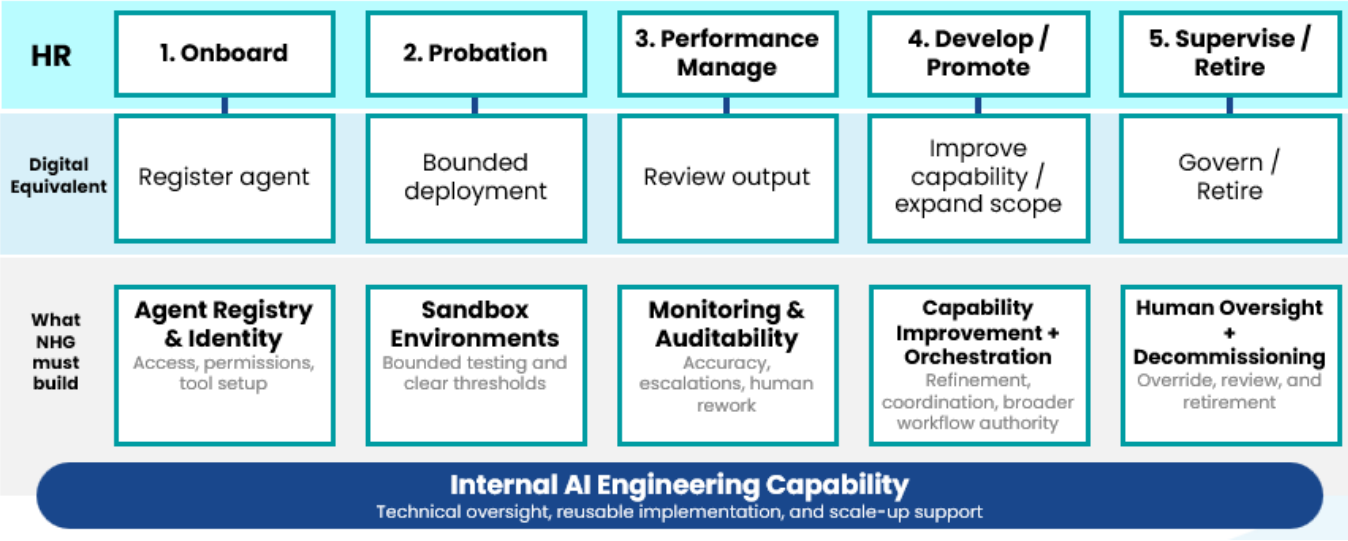

How to map HR functions to digital equivalents and platform requirements.

The 'HR for agents' metaphor remains useful because it translates abstract governance into familiar management functions. It should be used carefully: agents are not people, and the metaphor should not obscure human accountability. Its value is that it makes clear that every management function maps directly onto a technical, governance or operational requirement.

Workforce management function	Digital equivalent	Operational requirement
Onboarding	Register the agent, define role, assign identity, permissions, policies, memory boundaries and tool access.	Agent registry, identity management, role definition, risk classification and access approval.
Probation	Deploy in a bounded workflow with limited permissions, close monitoring and explicit success thresholds.	Sandbox, validation protocol, test cases, human review, release gates and evidence thresholds.
Performance Management	Monitor accuracy, completion, escalation, turnaround, rework, user feedback, safety events and cost.	Observability, logging, dashboards, evaluation layer and review cadence.
Development	Improve prompts, tools, context, retrieval, memory, model routing and workflow fit.	Capability improvement pipeline, version control, change control and performance comparison.
Promotion	Expand scope or autonomy only when there is evidence of safe and useful performance.	Staged deployment, autonomy levels, approval pathways and enhanced monitoring.
Supervision	Maintain human accountability, exception handling, override rights and audit review.	Named owner, escalation rules, audit trails, review boards and incident response.
Offboarding	Withdraw access, archive logs, retire versions and decommission agents no longer fit for purpose.	Lifecycle management, kill-switches, permission removal and records retention.

If an organisation has no agent registration, it has no reliable onboarding. If it has no sandbox and success thresholds, it has no probation. If it has no observability and metrics, it has no performance management. If it

cannot adjust permissions, context, prompts, tools or scope, it has no development pathway. If it cannot pause, rollback or retire an agent, it has no offboarding.

Core Principle
No anonymous agents in production. Every production agent should have an identity, a defined role, a permission set, a named human owner, and an observable record of what it has done.

6. AI-in-the-Loop, Job Redesign, and Workforce Transformation

Three key considerations when redesigning jobs. Gartner's research emphasizes that AI adoption succeeds when organizations prioritise workforce readiness, change management, and integration into real workflows and not just technology deployment.

The purpose of building a managed digital workforce is to expand the capacity and capability of the overall system. Deploying an agent creates the opportunity for capacity gain; redesigning the work is what captures it. The fuller benefit emerges when freed capacity is translated into redesigned roles, deeper human work, better care delivery, and more time for judgement, empathy, relationship-building and leadership.

In AI as a Tool and AI embedded in Workflows, the dominant pattern is often human-in-the-loop: a person drives the work and consults AI at controlled points. In AI as a Workforce, the relationship changes to AI-in-the-Loop. Agents may carry out meaningful parts of the workflow, while humans set direction, supervise, handle exceptions, approve consequential actions and intervene when judgement is needed.

This is why job redesign cannot be an afterthought, and is guided by three design principles

- **Design at task level, not role level.** Healthcare roles are bundles of tasks with different levels of complexity, risk, value and emotional content. Redesign should identify which tasks remain human, which are AI-assisted, and which can be delegated.

- **Design for the team, not the individual.** Agents will operate alongside people, other agents and systems. Redesign must address team configuration, accountability, handoffs, exception handling and trust.
- **Design to hedge against residue and deskilling risks.** When automation absorbs routine work, remaining human work may become more complex, less frequent or more judgement-heavy. Redesign must include deliberate skill maintenance and new human-agent teaming skills.

Workforce Transformations

NHG's stance is that some tasks will be automated, some roles will evolve, and the nature of work across the health system will change. History consistently shows that technological change creates new categories of work even as it transforms existing ones. The challenge is not whether change will happen, but whether NHG Health manages it responsibly. Our ongoing commitment is that it will, in the form of reskilling, reimagining workforce mobility pathways and guided by the principle of augmentation over replacement.

7. What would need to be true? Seven Conditions for operationalising AI as a Workforce

The next step is to identify the conditions that would need to be in place for NHG Health to safely deploy, manage and scale a digital workforce. These conditions should be applied proportionately as a risk-tiered approach allows NHG Health to move at pace without treating every use case as either trivial or prohibitive.

Seven conditions appear particularly important.

These conditions should be applied proportionately. Not every agent requires the same level of control. Low-risk agents that summarise, draft or retrieve information may need lighter controls; higher-risk agents that access sensitive data, update patient records, trigger actions, interact with patients or residents, or influence clinical, operational or financial decisions require the highest level of assurance, monitoring, human supervision and release control. A risk-tiered approach will help NHG Health move at pace without treating every use case as either trivial or prohibitive.

Condition	What it means for NHG Health
7.1 Clear Production ambition	Which priority workflows should demonstrate measurable improvement through safe AI-enabled redesign within the next 12 months? Our success may be defined through outcomes such as: • Number of priority workflows improved • Measurable operational, clinical or workforce value, • Safer, faster, more reliable workflow performance • Reduced waste, improved responsiveness, better staff experience, better resident/patient outcomes

7.2 Development close to workflows, enabled by common platforms	Agents should be designed close to the work but not built in isolation. Workflow proximity is essential because safe and useful agents depend on local context: how work is actually done, where variation occurs, what information is trusted, what exceptions matter, and where human judgment must remain central. However, proximity should sit on top of common architecture. NHG Health should avoid creating multiple disconnected local solutions. The preferred model is a matrixed one: workflow teams identify, shape and test opportunities close to users, while shared platform teams provide reusable components, standards, orchestration, security, observability, evaluation methods and lifecycle management. This means the strategic capability is not simply "citizen development" or "central build". It is a partnership model in which local teams bring workflow knowledge and ownership, while central teams provide safe speed, technical depth and common infrastructure.
7.4 People-first transition and workforce readiness	AI transformation is fundamentally a people transformation. Gartner's research emphasizes that AI adoption succeeds when organizations prioritise workforce readiness, change management, and integration into real workflows and not just technology deployment. NHG Health must invest accordingly in reskilling, role redesign, and deliberate protection against cognitive deskilling as routine tasks are automated.
7.5 A platform that enables safe speed	If every agent requires a bespoke path to validation, deployment and monitoring, NHG Health will not scale. HEAL should become the mechanism that allows teams to move faster safely through common sandbox, orchestration, observability, access, evaluation, cybersecurity and lifecycle capabilities.
7.6 Strong foundations of data, context and governance	Agents need context, not just models. They require reliable data, semantic clarity, workflow knowledge, organisational policy, role definitions, access rules, escalation criteria and outcome measures. Weak foundations produce unreliable agents, higher risk and limited value.
7.7 Production disciplines for resilience and scale	Production agents need observability, least-privilege access, failure design, named accountability, incident response, cost management and lifecycle control. Without these disciplines, early success can erode trust in the broader programme.

8. The platform and capability stack

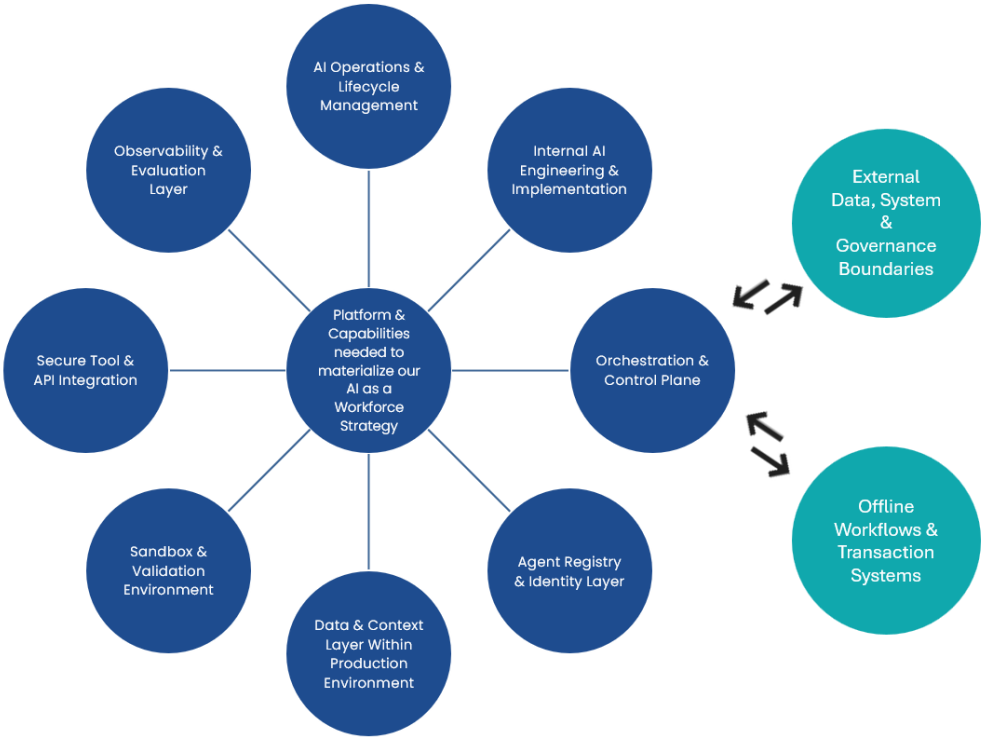

The platform and capability stack needed to materialise our AI as a Workforce Strategy

If AI as a Workforce is to move beyond concept, NHG Health needs a practical platform and capability stack to support the safe development, validation, deployment and management of digital workers. This is the technical and operational foundation that makes the workforce metaphor real rather than rhetorical.

The platform should be understood less as a single product and more as a set of linked capabilities. These capabilities should allow teams to identify suitable workflows, build and test agents safely, assign identity and permissions, connect approved tools and data sources, monitor behaviour, measure outcomes, manage cost, respond to incidents and improve or retire agents over time. It should also preserve model and vendor optionality as agentic AI is evolving too quickly to lock workflow redesign to any single vendor or architecture

Platform component	Purpose
Sandbox and validation environment	Safe spaces to test agent behaviours, workflows, policies, permissions, context and failure modes before production deployment. This is where onboarding and probation happen in practice.
Agent registry and identity layer	A record of all agents: roles, owners, versions, permissions, risk tier, tools, data access and deployment status. No anonymous agents in production.
Orchestration and control plane	How agents interact with tools, systems, humans, triggers, escalation paths and other agents

Data and context layer within production environment	Trusted data from internal and external sources, retrieval, semantic definitions, workflow knowledge, policy context, workflow knowledge, and memory boundaries that allow agents to act within organisational intent.
Secure tool and API integration	Approved pathways for agents to access enterprise systems, retrieve information, update records or trigger workflows within defined limits.
Observability and evaluation layer	Logging, monitoring, evaluation, drift detection, human rework tracking, escalation analysis, safety events, user feedback and performance dashboards.
AIOps and lifecycle management	Versioning, model routing, prompt management, memory management, guardrails, evaluation assets, rollback, cost monitoring, retirement and decommissioning.
Internal AI engineering and implementation	A multidisciplinary capability that absorbs learning, supports workflow teams, enforces standards, scaffolds solutions from validation to scale and prevents each use case from becoming a bespoke project.
Offline workflows and transaction systems	Operational systems, scheduling, case management, CRM-style systems, finance/admin platforms

9. Trust, cybersecurity and resilience by design

Our proposed HR for Agents concept establishes the broad guardrails for a managed digital workforce. This section establishes the details.

Agentic AI changes the cybersecurity model because the organisation is not only securing software that produces outputs. It is supervising non-human actors that may have credentials, memory, tool access, system edit access and the ability to act. This makes runtime governance as important as pre-deployment approval.

A helpful metaphor is that agents can behave like 'super-intelligent but gullible 8-year-olds': highly capable, but still vulnerable to manipulation, misleading context, unsafe instructions, or prompt injections.

Traditional application security remains necessary but is insufficient. AI introduces a broader attack surface: prompts, retrieved content, tool calls, APIs, vector stores, training and fine-tuning data, model endpoints, memory, inference, caching, agent-to-agent interaction and human feedback loops. The security model must therefore be designed across the full AI infrastructure end to end.

- **Identity and ownership**. Every production agent should have a unique identity, defined scope, version, risk tier and named human owner. The principle should be simple: no anonymous agents in production.
- **Least-privilege access.** Agents should only access the data, systems and tools needed for the task. Read-only should be the default unless write access is clearly justified.
- **Runtime observability.** Every action should be logged, queryable and auditable. If the organisation cannot see what the agent did, it cannot govern it.
- **Prompt and context security.** Prompt injection, malicious retrieved content, misleading documents and unsafe external inputs need to be treated as real attack vectors.
- **Adversarial testing and red-teaming**. Higher-risk agents should be deliberately tested before production release to identify unsafe tool use, data leakage, prompt injection susceptibility, permission escalation,

inappropriate refusal or escalation avoidance, memory manipulation, malicious inputs and behaviour outside intended scope.

- **Memory management.** Memory can be useful but also risky. It needs rules for retention, inspection, correction, reset, poisoning detection and sensitive data handling.
- **Failure design.** Agents should know when to stop, escalate or refuse to continue. The most dangerous agent is one that guesses confidently when it should pause.
- **Incident response.** AI incidents may involve harmful outputs, data leakage, unauthorised tool use, memory poisoning, cost runaway or unexpected agent-to-agent behaviour. These require a specific playbook.
- **Lifecycle control.** Agents need rollback, pause, permission withdrawal, decommissioning and retirement processes.

Trust by design

Trust is not a brand issue; trust is a design issue. For agentic AI, trust must be built into identity, architecture, access control, observability, failure design, incident response and lifecycle management.

10. Economics of a digital workforce

AI compute and token consumption should be treated as a new form of productive capacity. It is not equivalent to people, but it does create a resource that must be planned, budgeted, monitored and governed.

A use case that is impressive in a pilot may become unaffordable in production if it relies on expensive models for tasks that could be routed to cheaper models, if it loops unnecessarily, if it retrieves too much context, or if it performs actions that require excessive human review. Cost discipline is therefore not only a finance issue. It is part of scalability and operational resilience.

NHG Health should develop the ability to understand cost per useful outcome, not simply total AI spend. Relevant questions include: what does one resolved case cost in tokens? Which model should handle which part of the workflow? Who owns the AI budget? How often is usage reviewed? When is a frontier model justified? How do we prevent a successful pilot from becoming a hidden operational cost?

- **Cost per workflow and cost per outcome.** Measure spend against value, such as cost per resolved case, triage completed, document processed or escalation avoided.
- **Model routing strategy.** Use the right model for the task. Frontier models should be reserved for tasks where their additional capability is justified.
- **Token and compute budgets.** Set thresholds, alerts and review processes so scaling does not create uncontrolled spend.
- **Portfolio view.** Manage agents as a portfolio, including reuse, retirement, performance, cost and cumulative risk.
- **Monthly ownership.** Assign a clear owner for AI budget review and usage decisions, not only an IT line item.

11. Capability building: from literacy to deep deployment capability

In section 7 of this paper, we posited that, among others, citizen development capability and the adequate expertise and capacity of our people are key criteria to operationalize AI as a Workforce. This section expands on the building the capabilities of our Human workforce.

AI capability should not be treated only as a specialist programme or optional training offer. It will need to become part of NHG Health's core organisational competency, in the same way that operational improvement, clinical quality, digital literacy have become expected capabilities for modern healthcare delivery.

This has implications for roles, job design and leadership expectations. Future job descriptions, competency frameworks and development pathways should increasingly recognise AI fluency, responsible adoption, data-informed improvement and human-AI supervision as part of how work is led and improved.

Broad AI literacy is necessary but insufficient. If NHG Health is to move agentic AI from concept to routine practice, it needs deeper capability in evaluation, implementation, workflow integration, cybersecurity, data stewardship, product management and clinical leadership. This is already recognised in the CHI Clinical AI Fellowship programme hosted by NHG Health, which aims to build clinicians who can lead evaluation, implementation and workflow integration rather than simply use AI tools.

Recent discussions also reinforced the importance of people who can bridge technical capability and real-world work. The forward deployed engineer concept is useful here. The important feature is not the job title, but the capability: people who can work close to the workflow, understand user needs and operational constraints, and translate frontier capability into safe, useful deployment.

NHG Health will need a layered capability model embedded as a core organizational competency. AI capability should not be treated only as a specialist programme or optional training offer. Over time, it will need to become part of NHG Health's core organisational competency. At the broadest level, staff need AI literacy and confidence. Workflow teams need AI fluency: the ability to understand where AI can and cannot safely improve their own work. Clinical and operational leaders need deeper capability in evaluation, implementation, workflow integration and responsible adoption. Technical teams need specialist capability in engineering, cybersecurity, data, product management, architecture and lifecycle operations.

A further layer should build on CHI's existing automation and citizen developer community. The skills that have helped teams identify, redesign and automate workflows are highly relevant to agentic AI. A future community of AI champions and citizen developers could help surface suitable workflows, distinguish automation from agentic opportunities, support local adoption, and create a practical bridge between central platform standards and frontline service knowledge. This should be encouraged, but not left uncontrolled: locally developed ideas should be supported by common design patterns, security guardrails, review pathways and escalation routes.

This has implications for roles, job design and leadership expectations. Future job descriptions, competency frameworks and development pathways should increasingly recognise AI fluency, workflow redesign, responsible adoption, data-informed improvement and human-AI supervision as part of how work is led and improved. Not every staff member needs to become a technical expert, but many more staff will need to understand where AI can help, where it should not be used, how to supervise it safely, and how to redesign work so that human and digital capabilities are used well.

This shifts the capability question from "how do we train people on AI tools?" to "what skills, roles and competencies will NHG Health need in a world where AI is part of the workforce?"

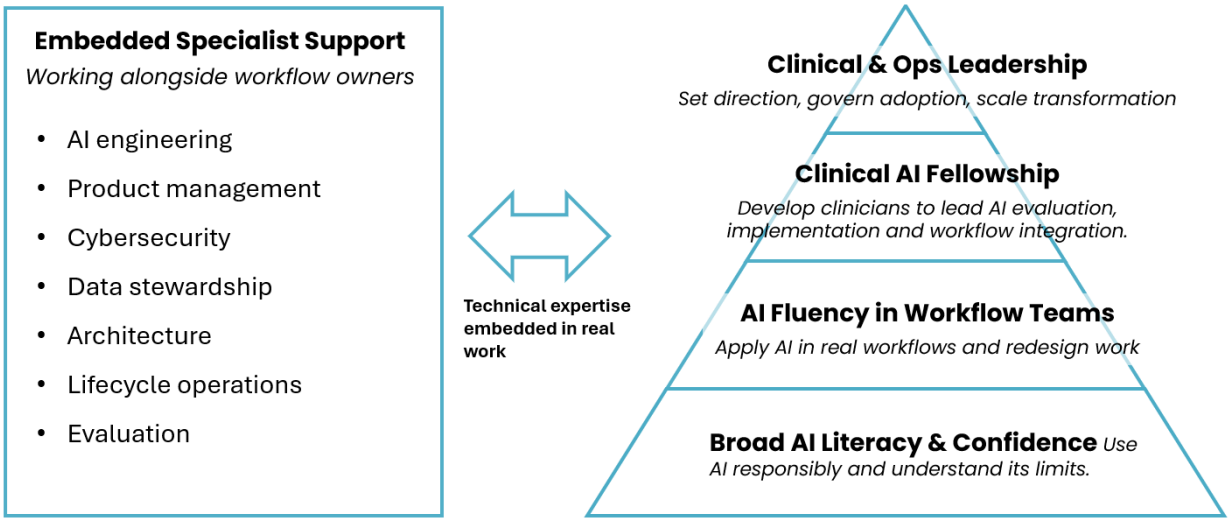

The Capability Building Model

12. Moving beyond academic concepts to production and outcomes

The value of this paper will not be in the concept alone. "AI as a Workforce" will remain an academic framing unless NHG Health can translate it into a small number of real, prioritised, production workflows that demonstrate measurable value. The next phase should therefore focus on selecting priority use cases, building the enabling platform, assigning accountable owners, deploying safely, and measuring outcomes.

The test is not whether the theory is compelling, but whether it helps us deliver better, safer and more sustainable work in practice

13. From white paper to delivery: the questions the roadmap must answer

This paper is not intended to be a detailed delivery roadmap. Its purpose is to frame the change and identify the conditions required. The next step should be to translate these conditions into a practical roadmap for HEAL and NHG Health over the coming year. The roadmap should answer a set of concrete questions. These questions are deliberately framed around operational readiness rather than technology alone.

Roadmap question	Why it matters
Production ambition and workflows	Which priority workflows should demonstrate measurable improvement through safe AI-enabled redesign within the next 12 months? Which workflows are sufficiently important, bounded, measurable and ready for agentic redesign, and which workflows should we not use agentic AI for? Number of priority workflows improved, Measurable operational, clinical, or workforce value Safer, faster, more reliable workflow performance Reduced waste, improved responsiveness, better staff experience, better resident/patient outcomes
Ownership	Who owns each workflow, each agent, each output and each risk?
Platform	What sandbox, orchestration, registry, observability, evaluation, deployment and model/vendor optionality capabilities are required?
Governance	What approval pathways, risk tiers, release gates, evaluation metrics, review cadence and incident processes are needed? How would governance keep up with agentic activity – in other words, how do we ensure that staff not end up spending more time checking agentic output vis a vis keeping their attention on the patient, and how would we augment with more technical personnel?
Cybersecurity	How will we identity, least privilege, prompt security, memory, tool access, logging, runtime monitoring, adversarial testing and decommissioning be handled? How would we balance the level of security and usability of agentic systems, and what tradeoffs between potential and realised benefits are we willing to accept?
Data and context	What data, knowledge, policy, semantic and workflow foundations must be strengthened?
Capability	What expertise and capacity must be built, hired, partnered or embedded close to the work, including champions, citizen developers and specialist AI deployment capability? How would we ensure that staff embrace agentic AI and not feel threatened, or even actively sabotage it?
Economics	How will we manage token budgets, model routing, cost per outcome and portfolio economics? How big would the initial investment look like, taking into consideration the sliding cost of memory and tokens across a long time horizon, and the capabilities of current end-user devices in our clinics and wards?
Learning system	How will each deployment generate reusable learning for the next?

14. Conclusion: building the next operating model for healthcare

AI as a Workforce is not simply a new label for agentic AI. It is a way of making visible the organisational challenge that agentic AI creates. If agents can act, then they must be managed. If they operate inside workflows, then those workflows must be redesigned. If they access systems and data, then identity, access, cybersecurity and observability must be built in. If they contribute to service delivery, then performance, accountability, cost and lifecycle management become core operating disciplines.

The opportunity for NHG Health is to use this moment not only to adopt agentic AI, but to build the institutional capability to deploy it responsibly, safely and at scale. That means bringing together the best of NHG Health's existing strengths: mission-oriented population health and care redesign, CHI's innovation cycle, HEAL's AI adoption platform, clinical leadership, evaluation capability, data and digital infrastructure, and a growing culture of responsible AI experimentation.

The aim is not digital substitution for its own sake. The aim is to redesign human-agent teams so that scarce human capacity is focused where it matters most: clinical judgement, relationships, empathy, leadership, complex problem solving, and accountable decision-making.

The immediate task is to move from concept to conditions: to define what would need to be true for NHG Health to safely deploy, manage and scale a digital workforce over the next 12 months. The more detailed roadmap can follow. But the direction is clear: build right, afford to scale, and keep agents safe and effective in production.