



PWYLLGOR DIGIDOL, DATA AC ARLOESI
DIGITAL, DATA AND INNOVATION COMMITTEE

DYDDIAD Y CYFARFOD: DATE OF MEETING:	21 July 2026
TEITL YR ADRODDIAD: TITLE OF REPORT:	Review of Organisational Artificial Intelligence (AI) Readiness
CYFARWYDDWR ARWEINIOL: LEAD DIRECTOR:	Huw Thomas, Executive Director of Finance
SWYDDOG ADRODD: REPORTING OFFICER:	Anthony Tracey, Director of Digital, Data and Technology Gareth Jenkins, Head of Data Science

Pwrpas yr Adroddiad (dewiswch fel yn addas)

Purpose of the Report (select as appropriate)

Ar Gyfer Penderfyniad/For Decision

DRODDIAD SCAA
SBAR REPORT

Sefyllfa / Situation

This report provides assurance to the Committee on Hywel Dda University Health Board's (HDdUHB) current level of organisational Artificial Intelligence (AI) readiness, together with recommendations, based on the independent assessment undertaken with CGI.

The Committee is asked to receive assurance on the progress made to date, the maturity of current arrangements, and the identified priorities required to enable safe, scalable and responsible AI adoption across the organisation.

Cefndir / Background

CGI undertook a comprehensive AI Readiness Assessment to evaluate HDdUHB's ability to adopt and scale AI across clinical, operational and corporate functions.

The assessment:

- Reviewed AI readiness across **six core pillars** (governance, people, data, models/assets, platforms, monitoring)
- Drew on:
 - Stakeholder interviews
 - Organisation-wide survey data
 - Existing documentation and live initiatives

The purpose was to:

- Understand the **current state of AI capability**
- Identify **risks and gaps**
- Define **practical next steps for safe and scalable adoption**

Asesiad / Assessment

Overall Position


The assessment concludes that HDdUHB has demonstrated emerging AI readiness, although this has not yet scaled across the organisation. Strong foundations are in place; however coordinated development and implementation are required to support enterprise-wide adoption. The evidence shows strong appetite for AI, active Copilot use and visible experimentation across administrative, clinical, operational and analytical areas. A separate workstream focussed on Shadow IT has identified use of several cloud-based AI tools. HDdUHB has important foundations in place to build on, including data quality activity, an information asset register, platform investment, Information Governance (IG) and an established data science function. The data science function and initial rollout of Copilot provide foundations at each end of the AI spectrum.

Appendix 1 provides the full report of the AI Assessment. It is intended to support the recommendations outlined within this paper by providing additional detail and assurance regarding the scope, implementation considerations, and anticipated outcomes of the proposal. The content within the appendix should be considered alongside the main report when reviewing the proposed next steps and decision points.

The opportunity is to establish a comprehensive approach to AI governance and oversight across the organisation, ensuring that all forms of AI are appropriately considered. Current AI adoption outside the Data Science function remains at an early-stage and has yet to be scaled, with activity largely focused on productivity tools and localised areas of experimentation. The next step is to align the data, platform, tool and governance landscape, so it is clear where different types of AI activity should occur, where assets should reside and how they transfer from idea to supported use at scale. The overall direction is to build on existing foundations, make AI activity visible to the right people, strengthen data quality for AI use, clarify approved tools, refine platform roles and identify priority use cases to create reusable delivery patterns.

Hywel Dda and the AI Spectrum

Hywel Dda has emerging activity across the AI spectrum, but custom generative AI remains the least clearly owned area. This will need clearer ownership, capability and governance as use cases develop.



CGI identified five key areas to prioritise that would safely integrate AI into the organisation. These areas take into account the parallel workstreams CGI are already working on with HDdUHB. Investment in governance, resources and user adoption will be required. Although these streams can be approached individually, logistically it would make more sense to group

them together as a program of work. Funding is potentially available from partners such as Microsoft to offset some costs.

Five Priorities for HDdUHB

The five priorities below provide a practical route from HDdUHB's current position of emerging AI readiness towards safe, scalable and responsible organisational adoption. They reflect the need to move beyond isolated experimentation by strengthening governance, improving visibility of AI activity, building staff confidence and capability, clarifying the data and platform foundations, and focusing early delivery on a small number of high-value use cases. Taken together, these priorities should be treated as a coordinated programme of work rather than separate initiatives, ensuring that AI adoption is aligned to organisational objectives, supported by appropriate controls, and able to deliver measurable benefits for patients, staff and services:

1. **AI Governance Foundations** - Implement AI principles, policy and the appropriate guardrails to enable responsible and scalable deployment.
2. **AI Enablement / Adoption** - Create a dedicated AI team responsible for building solutions, establishing staff training and fostering AI awareness.
3. **Data & Platform Strategy** - Clarify the data, content and platform foundations needed to support priority AI use cases.
4. **AI Tooling and Funding** - Optimise value from existing AI tooling, including Copilot, and leverage AI Vendor funding opportunities. Eradicate areas of shadow AI.
5. **Scaling Efficiently and Effectively** - Build priority AI use cases with reusable components (Radiology AI use cases in progress), to create repeatable and modular patterns.

Recommendations for HDdUHB

The recommendations below translate the findings of the AI Readiness Assessment into practical actions for HDdUHB. They are intended to strengthen organisational oversight, make current and future AI activity visible, clarify the safe and approved use of tools, and ensure that early priority use cases are selected and delivered in a way that creates reusable patterns for wider adoption. The recommendations also recognise that AI adoption must be supported by clear governance, robust data foundations, staff confidence and defined routes from idea, to pilot, to scaled implementation.

- **Make AI activity visible to the right people:** align the information asset register with data science tracking, Copilot, vendor AI and foundational AI use cases so the AI Governance Group has clear oversight.
- **Build on existing governance processes:** enhance Responsible AI principles, policy scope, approval routes and assurance expectations so AI can be managed consistently. This includes those identified for Copilot.
- **Strengthen data quality for AI use:** build on existing data quality work to make source, ownership, metadata, known limitations and suitability visible for priority AI use cases.
- **Prioritise a small number of use cases:** agree the highest-value opportunities, define success measures, and use them to tighten working patterns from idea to pilot and wider adoption decision.

- **Clarify approved tools and safe use:** make it clear which AI tools can be used, what they can be used for, what data boundaries apply and where staff should go for guidance.
- **Refine the platform and tool landscape:** agree what Databricks, Microsoft, Posit, Copilot, other vendor tools and existing data platforms are each used for, including boundaries and hand-offs.
- **Build reusable patterns from early work:** utilise priority use cases, including radiology where appropriate, to create repeatable approaches for data checks, evaluation, approvals, release, monitoring and feedback.

Implications and Risks

The assessment highlights that the principal risk for HDdUHB is not AI adoption itself, rather it is the potential for AI activity to develop in an uncoordinated manner without appropriate organisational oversight, clinical assurance, IG, data quality controls or approved mechanisms for scaling effective solutions. In the absence of a clear governance framework and visible pipeline of AI activity, there is a risk that local experimentation, supplier-led functionality and productivity tools are adopted inconsistently, creating variation in practice, uncertainty for staff, avoidable duplication and gaps in accountability. This could expose the organisation to data protection, cyber security, clinical safety, equality, legal and reputational risks, while also limiting the ability to realise the anticipated benefits of AI in improved productivity, decision support, patient flow, quality and service sustainability:

- Data protection and security risks (particularly where governance and classification are immature)
- Inconsistent or unsafe AI usage across departments
- Failure to realise value at scale due to fragmented approaches
- Reputational and clinical risk if AI is implemented without sufficient oversight and assurance

Conclusion and Next Steps

The assessment confirms that HDdUHB has a clear opportunity to build on its existing AI foundations, including the data science function, active Copilot use, emerging clinical and operational interest, and established IG and data quality arrangements. However, AI adoption has not yet been scaled across the organisation and there is a need to transition from isolated experimentation to a coordinated, well governed and benefits-led approach. This will require stronger oversight, clearer approved mechanisms for AI use, improved visibility of current activity, defined data and platform foundations, and sufficient capacity to support safe adoption. The proposed next steps are therefore intended to provide a structured route from current readiness to controlled implementation, ensuring that AI is adopted in a way that is safe, clinically assured, legally compliant, operationally useful and aligned to organisational priorities.

AI Readiness Assessment – Proposed Next Steps and Delivery Plan

The AI Readiness Assessment identified that HDdUHB has strong foundations and several emerging areas of AI activity but requires a more coordinated, governed and benefits-led approach to support safe scaling across the organisation. The recommended next phase is therefore focused on establishing governance, creating organisational visibility, enabling safe adoption and developing the supporting data and technology foundations:

- Establish an AI Governance Group, or equivalent governance route, with clear Terms of Reference, decision rights and reporting arrangements into the Digital, Data and Innovation Committee.

- Create a visible pipeline of current and proposed AI activity, bringing together data science work, Copilot adoption, vendor-enabled AI, shadow AI findings and priority clinical, operational and corporate use cases.
- Define and communicate approved AI tools, safe-use principles, data boundaries and escalation routes so staff understand what can be used, for what purpose and with what safeguards.
- Agree a small number of high-value priority use cases, including where appropriate radiology and productivity-focused opportunities, with clear success measures, clinical safety requirements, IG requirements and benefits realisation plans.
- Clarify the data and platform architecture needed to support AI, including the role of existing Microsoft, Databricks, Posit, vendor and data science environments, and the hand-offs between them.
- Develop an implementation plan that sets out the required capacity, governance products, training and awareness activity, delivery milestones and reporting cycle, with a progress update presented to the Committee.

Proposed Phasing

Phase 1 – Foundations (Months 1–3)

- Establish AI Governance Group.
- Approve Terms of Reference and reporting arrangements.
- Create AI activity register and pipeline.
- Publish interim AI guidance.
- Commence awareness and training programme.
- Define architecture principles and governance products.

Phase 2 – Prioritisation and Control (Months 4–6)

- Prioritise initial AI use cases.
- Complete architecture blueprint.
- Implement assurance and approval processes.
- Baseline benefits and success measures.
- Confirm funding and resource requirements.

Phase 3 – Delivery and Scaling (Months 7–12)

- Deliver priority use cases.
- Monitor benefits, safety and compliance.
- Refine governance arrangements based on experience.
- Produce AI Roadmap.
- Report progress and outcomes through DDIC.

Argymhelliad / Recommendation

The Committee is requested to:

- **NOTE** the findings of the AI Readiness Assessment, including that HDdUHB has emerging AI readiness, strong foundations and clear opportunities to expand the safe and responsible use of AI, but is not yet operating at scale.
- **RECEIVE ASSURANCE** from the progress already made through the data science function, Copilot adoption, existing information governance arrangements, data quality activity and wider organisational engagement in AI opportunities.
- **ENDORSE** the proposed five priority areas as the strategic framework for a coordinated programme of work to support safe, scalable and responsible AI adoption across HDdUHB, covering AI governance, enablement and adoption, data and platform strategy, tooling and funding, and the development of reusable delivery patterns.

- **APPROVE** the establishment of an AI Governance Group, or equivalent governance route, with clear Terms of Reference, decision rights, assurance requirements and reporting arrangements into the Digital, Data and Innovation Committee.
- **NOTE** that an implementation plan is being developed setting out governance products, approved tool guidance, data and platform requirements, capacity implications, training and awareness activity, milestones and a reporting cycle, with a progress update presented to a future meeting of the Committee.

Amcanion: (rhaid cwblhau) Objectives: (must be completed)	
Cyfeirnod Cylch Gorchwyl y Pwyllgor: Committee ToR Reference:	2.1.1 That the direction, development and delivery of the Digital Strategic Plan is to drive continuous improvement and support digitally enabled health care through a digitally enabled workforce to achieve the objectives of the Health Board's Annual Plan/Integrated Medium-Term Plan (IMTP). 3.1.19 Seek assurance on delivery against all Planning Goals aligned to the Committee, in accordance with the Board approved timescales, as set out in the Health Board's Annual Plan, considering, scrutinising and supporting and endorsing plans and programmes developed to deliver the Committee's objectives.
Cyfeirnod Cofrestr Risg Datix, Teitl a Sgôr Risg Cyfredol: Datix Risk Register Reference and Score:	Not Applicable
Parthau Ansawdd: Domains of Quality Quality and Engagement Act (sharepoint.com)	7. All apply
Galluogwyr Ansawdd: Enablers of Quality: Quality and Engagement Act (sharepoint.com)	6. All Apply
Amcanion Strategol y BIP: UHB Strategic Objectives:	All Strategic Objectives are applicable
Nodau Cynllunio 2026/27 Planning Goals 2026/27	All Planning Goals Apply

Hypergyswllt i Amcanion Llesiant HDdUHB 2026 Hyperlink to HDdUHB Well-being Objectives 2026	9. All HDdUHB Well-being Objectives apply
Model Cymdeithasol ar gyfer Iechyd a Llesiant : (wedi'i gwblhau ar gyfer newid strategol a/neu newidiadau sylweddol i wasanaethau) A Social Model for Health and Wellbeing : (<i>complete for strategic change and/or significant service changes</i>)	7. All Apply

Gwybodaeth Ychwanegol: Further Information:	
Ar sail tystiolaeth: Evidence Base:	Hywel Dda AI Readiness final report v1.0
Rhestr Termiau: Glossary of Terms:	Included within the report
Partïon / Pwyllgorau â ymgynhorwyd ymlaen llaw y Pwyllgor Digidol, Data ac Arloesi Parties / Committees consulted prior to Digital, Data and Innovation Committee:	Not Applicable


Effaith: (rhaid cwblhau) Impact: (must be completed)	
Ariannol / Gwerth am Arian: Financial / Service:	There are no immediate financial or capital commitments arising directly from this assurance paper. However, delivery of the proposed AI readiness priorities is likely to require further consideration of capacity, governance support, training, platform use, tooling and potential supplier or partner funding opportunities. Value for money will need to be assessed through prioritised use cases, clear benefits realisation plans and avoidance of duplicated or uncoordinated AI investment.
Ansawdd / Gofal Claf: Quality / Patient Care:	The report supports quality and patient care by setting out the governance, assurance and data foundations required for safe AI adoption. Potential benefits include improved productivity, decision support, patient flow, quality improvement and service sustainability. Any clinical AI use case will require appropriate clinical safety assessment, information governance review, evaluation, monitoring and clear accountability before wider implementation.

Gweithlu: Workforce:	The workforce impact is primarily related to capability, confidence and capacity. AI adoption will require staff awareness, training, clear safe-use guidance and access to approved tools. There may also be a need for dedicated capacity to support governance, evaluation, implementation and benefits realisation. The report does not propose workforce reductions; rather, it identifies opportunities to support staff productivity and reduce avoidable administrative burden where AI is safely applied.
Tegwch lechyd: Health Equity:	Potential health equity impacts have been considered at a strategic level. AI has the potential to improve access, consistency and service responsiveness, but may also introduce or amplify bias if data quality, population representation, evaluation and monitoring are not adequately addressed. Equity considerations should therefore be built into prioritisation, design, testing and ongoing monitoring of each AI use case. • Potential impacts on health equity have been considered: Yes, at strategic level. • An equity-focused health impact assessment has not been undertaken at this stage; this should be considered for individual AI use cases where there is potential impact on access, outcomes or experience.
Risg: Risk:	The key risks are uncoordinated AI adoption, inconsistent safe-use practice, unclear accountability, immature governance, data protection exposure, cyber security risk, clinical safety risk, equality risk and failure to realise benefits at scale. The proposed mitigations are to establish clear AI governance, create a visible AI activity pipeline, define approved tools and data boundaries, strengthen data quality controls, prioritise use cases, and report progress and assurance through the Digital, Data and Innovation Committee.
Cyfreithiol: Legal:	There are no immediate legal impacts arising from this assurance paper. Legal considerations will need to be addressed as part of individual AI use cases, particularly in relation to data protection, information governance, procurement, intellectual property, clinical safety, transparency, accountability and compliance with relevant legislation and national guidance.
Enw Da: Reputational:	There is potential reputational benefit if Hywel Dda adopts AI in a safe, transparent and benefits-led way, demonstrating responsible innovation and improved support for services. Conversely, ungoverned or poorly assured AI use could create reputational risk, particularly where it affects patient care, privacy, fairness, staff confidence or public trust. The governance and assurance

	approach proposed in the report is intended to mitigate this risk.
Gyfrinachedd: Privacy:	Privacy and confidentiality are central considerations for AI adoption. The report identifies the need for approved tools, clear data boundaries, visible oversight of AI activity, alignment with information asset management, and robust information governance review. Any use of personal, confidential or sensitive data in AI tools must be subject to appropriate assessment, controls, security safeguards and ongoing monitoring.
Cydraddoldeb: Equality:	Equality impacts have been considered at a strategic level. AI adoption has the potential to improve consistency and access, but may also create differential impacts if underlying data, model design, implementation or monitoring do not adequately account for protected characteristics and population diversity. Equality considerations should therefore be assessed for individual AI use cases as part of governance, approval and evaluation. • Equality Impact Assessment (EqIA) screening has not been undertaken for this assurance paper; it should be completed for individual AI use cases where implementation may affect access, experience, outcomes, staff groups or protected characteristics.

AI Readiness Assessment

Hywel Dda
June 2026



Contents

Page	Section
3	Executive Summary
6	Why and How we Arrived at These Findings
12	Approach
16	Conclusions and Roadmap
19	Responsible AI at Hywel Dda
28	AI Readiness: Technical Pillars
37	Microsoft Copilot Readiness
45	User Research
56	Acknowledgements
58	Appendix

Executive Summary

Introduction

CGI worked with Hywel Dda to assess its current AI readiness and identify practical next steps for safe and scalable adoption

The core question for this assessment was:

What capabilities does Hywel Dda already have to support AI development and use in clinician/patient focussed areas, and what needs to be done to further integrate AI into the organisation safely and with confidence?

- CGI identified five key areas to prioritise in order to safely integrate AI into the organisation.
- They take into account the parallel workstreams CGI are already working on with Hywel Dda.
- It is clear that investment in governance, resources and user adoption will be required.
- Although they can be approached individually, logistically it would make more sense to group them together as a program of work. N.B. Funding is potentially available from partners such as Microsoft to offset some costs.
- The Roadmap slide (18) identifies the items to work on now, in the near term and further down the line.

Five key priorities for Hywel Dda

AI Readiness Assessment

01

AI Governance foundations

Implement AI principles, policy and the appropriate guardrails to enable responsible and scalable deployment.

02

AI Enablement / Adoption

Create a dedicated AI team responsible for building solutions, establishing staff training and fostering AI awareness.

03

Data & platform strategy

Clarify the data, content and platform foundations needed to support priority AI use cases.

04

AI tooling/Funding

Optimise value from existing AI tooling, including Copilot, and leverage AI Vendor funding opportunities. Eradicate areas of shadow AI.

05

Scaling Efficiently & Effectively

Build priority AI use cases with reusable components*, to create repeatable and modular patterns.

* Radiology AI Use Cases in flight

Enterprise AI success depends on aligning business priorities with scalable technology foundations.

Why and How we Arrived at These Findings

Why AI Readiness is an important enabler for Hywel Dda

Many organisations are focusing on adopting AI without realising full value
Employees are adopting AI fast and bringing it to work if its not provided



88% of organisations reported regular AI use in 2025

McKinsey – State of AI, 2025

60% not yet realising material benefits

BCG – Widening AI Value Gap, 2025



43% increase in AI incidents reported globally since 2024

Stanford AI Index, 2026

Data privacy and patient confidentiality remain top worries, with up to 57% of healthcare leaders citing security concerns.

McKinsey Generative AI in Healthcare, 2026



78% of knowledge workers bringing own AI tools to work

Microsoft – Work Trend Index, 2024

78% of healthcare staff are entirely unaware of the HIPAA risks they introduce via shadow AI.

Healthcarediver – Shadow AI Report, 2026

Summary



CGI assessed Hywel Dda's **current AI readiness** across six pillars, drawing on stakeholder interviews, an organisation-wide survey and available internal documentation.



The evidence shows **strong appetite for AI**, active Copilot use and visible experimentation across administrative, clinical, operational and analytical areas. A separate workstream focussed on **Shadow IT** has identified use of several cloud-based AI tools



Hywel Dda has **important foundations in place to build on**, including data quality activity, an information asset register, platform investment, AI Governance and an established data science function.



The data science function and initial rollout of Copilot provide foundations at each end of the AI spectrum. The opportunity is to make sure **all types of AI** are covered, with data science overview and AI governance oversight.



Current AI use outside of Data Science remains **early-stage and not yet scaled**, with activity mainly focused on productivity support and pockets of experimentation.



The next step is to **align the data, platform, tool and governance landscape**, so it is clear where different types of AI activity should happen, where assets should reside and how they move from idea to supported use at scale.



The overall direction is to build on existing foundations, make AI activity visible to the right people, strengthen data quality for AI use, clarify approved tools, refine platform roles and identify priority use cases to create **reusable delivery patterns**.



CGI's engagement with Hywel Dda on AI-assisted radiology and data integration work provides **focused areas for further exploration** that could help connect readiness findings with practical delivery activity as AI adoption moves forward.

* Microsoft offer funding options related to AI/copilot readiness, governance, adoption and deployment.

Key Findings



AI readiness is emerging, not yet scaled: high awareness and appetite are clear, with current use mainly focused on productivity and administrative support, alongside pockets of clinical, operational and analytical experimentation.



Key elements of Responsible AI and AI Governance still emerging: requiring further policy, oversight, assurance, and compliance processes to manage AI risks effectively across Hywel Dda.



Need to connect technical foundations: Hywel Dda has data, analytics, data science, Copilot, platform and governance components, but not yet a consistent route for AI ideas, vendor tools and priority use cases to move into supported use at scale.



Clearer guidance needed on approved tools: survey and interview evidence show awareness and use of Copilot and other AI tools, while staff still need clearer guidance on available tools, permitted use, data boundaries and access.



People enablement will shape adoption: staff appetite is positive, but survey evidence shows strong demand for formal training, approved tools, written guidance and ongoing support.




Radiology is a focused opportunity: the use cases being explored with CGI could provide a practical route to test value, controls and delivery patterns before wider adoption decisions.

These key findings are an enabler to the key recommendations

Hywel Dda and the AI Spectrum

Hywel Dda has emerging activity across the AI spectrum, but custom generative AI remains the least clearly owned area. This will need clearer ownership, capability and governance as use cases develop.



Key Recommendations



Make AI activity visible to the right people: align the information asset register with data science tracking, Copilot, vendor AI and foundational AI use cases so the AI governance board has clear oversight.



Build on existing governance processes: enhance Responsible AI principles, policy scope, approval routes and assurance expectations so AI can be managed consistently. This includes those identified for Copilot.



Strengthen data quality for AI use: build on existing data quality work to make source, ownership, metadata, known limitations and suitability visible for priority AI use cases.



Prioritise a small number of use cases: agree the highest-value opportunities, define success measures, and use them to tighten working patterns from idea to pilot and wider adoption decision.



Clarify approved tools and safe use: make it clear which AI tools can be used, what they can be used for, what data boundaries apply and where staff should go for guidance.



Refine the platform and tool landscape: agree what Databricks, Microsoft, Posit, Copilot, other vendor tools and existing data platforms are each used for, including boundaries and hand-offs.



Build reusable patterns from early work: utilise priority use cases, including radiology where appropriate, to create repeatable approaches for data checks, evaluation, approvals, release, monitoring and feedback.

* Envisioning Workshop Funding is available from Microsoft

Approach

Assessment Scope and Approach


The assessment was structured across six key areas (AI readiness pillars) and drew on three main evidence sources:

- 21 stakeholder interviews across technical, governance, clinical, operational and people-focused roles
- A 16-question organisation-wide survey, issued to over 14,000 employees and completed by 433 respondents
- Available internal documents and artefacts relating to AI, governance, delivery and current initiatives

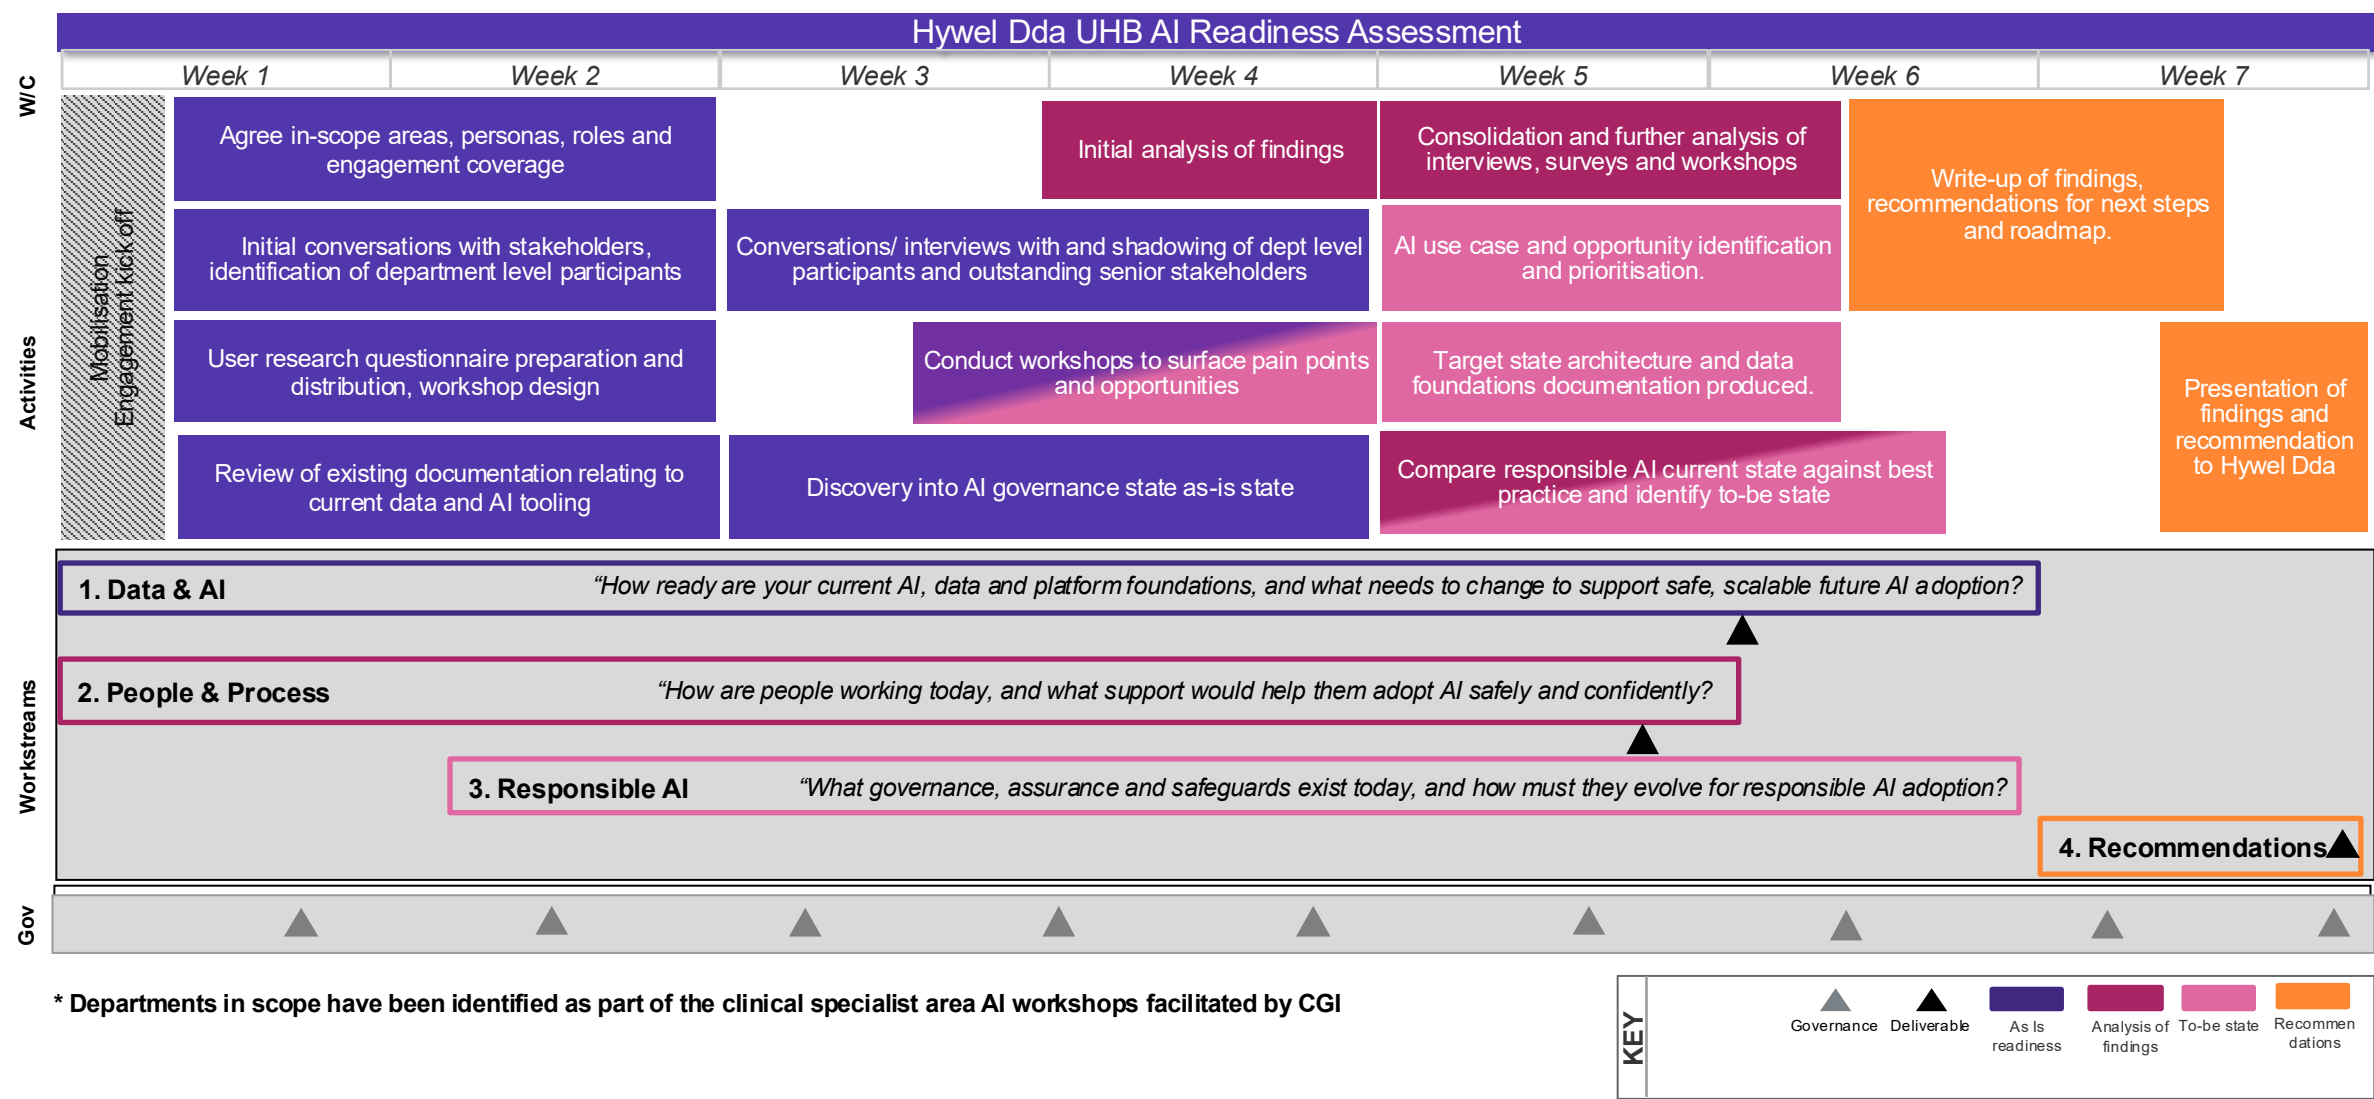
Note : Access to M365 was not available to enable deeper analysis of its current usage

AI Readiness Pillars

Six areas were used to frame Hywel Dda's current AI readiness and identify practical next steps



Plan on a Page: AI Readiness Assessment



Conclusions & Roadmap

Conclusions

AI is being effectively used in small specialist teams. To meet the broader User Appetite and gain authority-wide benefits more training, control and documentation is required around Data, Architecture, Tools and Process

We recommend the next step is to develop your AI Governance foundations further, including:

- ❖ establishing the AI Governance Board,
- ❖ creating AI principles,
- ❖ issuing interim AI guidance, followed by a dedicated AI policy and
- ❖ rolling out staff AI Awareness training.

From there, Hywel Dda can select an AI use case to explore further and move onto the wider AI Roadmap.

AI Roadmap

AI READINESS PILLARS	INITIAL	NEXT	BEYOND
DATA	Define priority AI use cases - Identify data sources required - Analyse data for suitability and readiness	Sandbox with route to live in mind - Curate data required for proving use cases - Map future data requirements for use cases	Operationalise data controls - Determine path for quality, access and lineage - Promote high quality, reusable data products
AI MODELS & ASSETS	Ensure asset register is future proof - Ensure fields for all potential information - Link to the models and assets in situ	Test register alongside development - Develop use cases as Proof of Concepts - Use POCs to refine register fields	Curate assets beyond POC - Apply release gates and model cards - Ensure promoted assets are registered
PLATFORM & PIPELINE	Clarify platform roles - Map the AI capabilities of all available platforms - Design environments for development and prod	Build delivery patterns - Determine path to live for POCs - Build CI/CD pipelines	Harden route to live - Package pipelines with CI/CD controls - Set rollback and support ownership
MONITORING & OPS	Define baseline and monitoring requirements - Forecast expected cost and usage - Agree AI-specific metrics and logging to capture	Build monitoring and logging - Build logging into the POCs as a first pass - Ensure monitoring is visible and actionable	Run live service controls - Track defined metrics - Monitor and act where necessary
FOUNDATIONS Responsible AI	Define Responsible AI at Hywel Dda - Agree AI principles and publish AI guidance - AI awareness training for staff	Responsible AI in action - Refine AI risk management and escalation - Review AI compliance	AI Assurance - Review audit processes - Plan for monitoring and reporting
PEOPLE & PROCESS	Plan communications and assign owners Define how people stay informed and involved	Human in the loop - Subject Matter Experts input into POCs - Keep the workforce informed	Embed into BAU - Communicate progress or change regularly - Ensure user feedback is considered

Responsible AI at Hywel Dda

Enterprise-level

Key headlines for Responsible AI and AI Governance

1

Priority: Essential foundations

Critical to get AI Principles, interim AI guidance and mandatory 'AI Awareness' training out to provide clarity for Hywel Dda staff.

2

Update existing processes

Review and update existing Governance processes to make sure they include AI risk identification and management.

3

Assess compliance

Look at AI considerations for Hywel Dda and any new requirements that might need to be included for suppliers of tools with AI features.

4

Develop AI assurance further

Refine structures and processes for monitoring, auditing and reporting on AI systems and organisational AI Governance processes.

AI is powerful but it
does need to be
used **thoughtfully.**



We need to consider:

**People, regulations,
global standards,
sustainability, and
AI risk management.**

With the right
governance, we can
innovate with confidence.



AI Governance Board

Effective AI Governance needs a **holistic overview** across ethical, legal, operational and societal aspects. Focusing on only one at a time can mean **vital opportunities** for innovation and risk management are missed. An AI Governance Board ensures this holistic overview happens.

Benefit: Robust, long-term structure that will provide a solid foundation as AI use grows within Hywel Dda, compared to alternative approaches.



AI Governance Board members:

- * Director of Digital (Anthony Tracey)
- * Head of Information Governance (Patrycja Duszynska)
- * Head of Digital Operations (Gavin Jones)
- * Other members (e.g. Clinical Safety team, Caldicott Guardian, Risk and Assurance team members)

Ways of working:

AI Governance Board meets regularly to own AI Governance and progress the tasks in the framework. Chaired by Director of Digital, unless delegated.



1) Make it clear: What does Responsible AI mean for Hywel Dda?

Current status	Next steps	Responsible
AI principles	- Existing values: https://hduhb.nhs.wales/healthcare/services-and-teams/our-values/ - Get AI Board input on what Hywel Dda's AI principles should be - Get Executive sponsorship - Who needs to be considered and/or who needs to test the principles? Operations – especially Cybersecurity, Information Governance, Clinical safety (CNO, CIO, clinical safety forums), Caldicott Guardian, Risk and Assurance team.	Director of Digital (Anthony Tracey, Executive Lead).
AI policy	- N/A - Define what's included under 'AI' at Hywel Dda - Issue interim best practice guide for staff as priority, and then write, approve and publish the new AI policy for staff. - Identify process for reviewing AI policy annually (e.g. take to relevant Board, go to consultation, then take to executive group). - Signpost to other resources for changing info (e.g. list of approved tools) - Refer to other related policies - both All-Wales and local policies.	Director of Digital (Anthony Tracey, Executive Lead).
Staff training	- Microsoft Copilot training exists (365 Centre of Excellence – not mandated) - Microsoft Copilot guidance exists - Compliance platform – have some training on AI - Establish mandatory 'AI Awareness' training for staff (e.g. 30 mins). - Link to the AI policy and Acceptable Use policy in the training - Consider tiers of learning - who only needs general 'AI Awareness' training and who needs additional, more detailed responsible use training (e.g. building AI systems, clinical use) - Clarity about what AI is, and broadly, traditional AI vs. generative AI. Use in business processes vs. clinician use.	Director of Digital (Anthony Tracey, Executive Lead).

2) Put it into action: Responsible AI at Hywel Dda

Current status	Next steps	Responsible
AI risk management	- Information Asset Register exists (review process is quite slow). - Clinical sign-off process for AI systems in clinical use – they map the risks. - Ethics Panel for general reviews - IT Digital Project Request Form - cannot install software without liaising with Information Gov. and Security. - Automated Decision Making question in DPIA, and data use and model use with AI.	- Update existing Information Asset Register and Governance processes to cover AI risks. Include AI impact assessment questions. - Clarify mandatory process for AI tool or feature approval – all tools must be approved before use (even ‘free’ online ones). - Check if any further AI-specific questions need to be added to Data Protection Impact Assessments (DPIA) and Security Assurance questionnaires. - Include process for updating AI systems, and phasing them out safely (especially clinical ones). - Establish future requirements for Responsible AI in the AI development lifecycle
Risk escalation	General incident response process - all incidents inc. IT incidents (Cyber investigation, Welsh Gov, and Comms team would coordinate)	- Examples of AI incidents are clearly outlined, staff are clear on how this works, and crisis communications process would kick in - Identify any contingencies for AI systems or tools (e.g. if one fails or goes down, what happens next?)
Compliance	- Plan to look at considerations - ISO 27001 is required for suppliers, as is ISO 9000.	Legal and Compliance

3) Show we're doing Responsible AI well: Checking it's working

Current status	Next steps	Responsible
Audits	• Currently dealing with pilots • Information Governance – do internal audits. • Have audits for other aspects of Governance.	• Look at whether any further AI points / processes need to be added to existing internal audits (i.e. align with AI risk management and impact assessments) • Bring in experts to audit AI systems regularly (e.g. for bias or explainability) and/or audit Governance processes. • Consider AI Governance external audit, aligned with ISO 42001 on AI Management Systems.
Monitoring	• Plan to monitor AI systems/tools and organisational AI Governance processes	• Head of Information Governance (Patrycja Duszynska)
Reporting	• Plan to monitor AI systems/tools and organisational AI Governance processes • Reports are currently regulations/legislation-driven	• Identify what you need to monitor (e.g. how many AI systems in use, how many people have taken training). • Monitor regulatory changes and enforcement. • What tools/systems/data do you need to be able to monitor AI use and/or AI Governance processes?
Reporting	• Plan to report AI systems/tools and organisational AI processes • Reports are currently regulations/legislation-driven	• Identify what you need to report on, who needs to be informed (internally and externally) and to what level of detail (e.g. credit usage and cost, sustainability) • Look at what Hywel Dda needs to be able to report (e.g. specific systems or dashboards).

AI Readiness: Technical Pillars

**Data,
Models and Assets,
Pipeline and Platform,
Monitoring and Ops**

Data Pillar

AI depends on data and content that are suitable, accessible and controlled for the intended use.

What this means

- Traditional ML needs structured data that is stable, trusted and well understood.
- Foundational AI, RAG and agents rely more heavily on documents, content, prompts and organisational knowledge.
- Vendor and commodity AI still need clear controls over local inputs, permissions, outputs and operational use.

Why it matters

- Data quality, ownership, sensitivity and context shape whether AI outputs can be trusted.
- Clear controls help AI systems operate safely, consistently and within agreed permissions.
- Data issues can create clinical, operational, privacy or reputational risk if they are carried into AI workflows.

Data Themes

Five themes used to assess whether Hywel Dda’s data and content can support useful, repeatable AI delivery.

Theme	Current status	Evidence signal	Next steps
Structured data for traditional ML and prediction	Structured data already supports reporting, analytics and some data science, but AI reuse depends on clearer metadata, quality and ownership.	SQL Server warehouse, Power BI reporting and data science activity provide a useful base, but the warehouse has grown organically and some datasets are scoped or duplicated.	Start with defining priority AI use cases. Record source, meaning, owner, refresh pattern, quality limits and historic recording issues for required datasets.
Unstructured content for generative AI	Documents, spreadsheets, SharePoint, Teams and shared folders are widely used, but priority knowledge sources are not yet mapped for AI use.	Key organisational content is spread across multiple document and collaboration stores.	Identify high-value knowledge domains. Record ownership, permissions, sensitivity, retention and update process before wider AI use.
Data quality, metadata and meaning	Data quality work exists, but AI-specific checks for meaning, completeness, history and suitability need to be made explicit.	Data quality worklists, SQL reports and patient-data quality activity already exist, but known issues are not always captured in a way that supports AI reuse.	Define a proportionate AI data check for early use cases. Capture known limitations and historic data issues.
Ownership, sensitivity and approval routes	Ownership, permitted use and approval routes are not consistently visible across systems and content stores.	Ownership and control visibility still depend on local knowledge in some areas.	Align data ownership and sensitivity information with existing asset, DPIA and information governance routes.
Operational and reusable data access	Databricks, Microsoft, Posit and the CGI / Hywel Dda integration-layer project all have roles to clarify for AI use.	Databricks is a strategic direction, the current warehouse supports reporting, Posit supports data science delivery, and the integration-layer work is a parallel route that needs clear positioning.	Assess platform roles and avoid duplicated routes. Define how priority AI use cases should access data and content.

Models and Assets Pillar

AI use becomes more repeatable when models and related assets are visible, tested, owned and controlled.

What this means

- AI assets include models, agents, RAG implementations, configured vendor tools and other AI-enabled capabilities.
- Artefacts are the specific items linked to an AI asset, such as prompts, model files, evaluation evidence and configuration records.
- Traditional ML, foundational AI and vendor AI need different validation and lifecycle controls.

Why it matters

- Hywel Dda needs to know what AI assets exist, what they do and where they are being used.
- Changes to prompts, models, data or configurations can change how an AI asset behaves.
- Lifecycle information helps teams move deliberately from idea to pilot, operational use and review.

AI Models And AI Assets Themes

Five themes used to assess whether Hywel Dda can manage AI models and assets through their lifecycle.

Theme	Current status	Evidence signal	Next steps
AI asset ownership and lifecycle control	AI activity spans data science, pilots, vendor tools and research-linked work. The lifecycle view needs to be consistent as possible across these routes.	The AI application form and related governance material provide a starting point for recording AI use cases and oversight needs. R&D routes, clinical safety routes and data science work already provide partial lifecycle controls.	Use the information asset register as the lifecycle view for material AI assets, from exploration to operational use.
AI asset visibility, registry and versioning	An information asset register exists, and the data science function also tracks related work.	Existing register activity, project tracking and data science work provide starting points, but artefacts such as prompts, notebooks, models, evaluation evidence and configuration records need consistent links.	Use the register as a current snapshot of AI assets and their linked artefacts. Align it with data science tracking of AI solutions rather than creating a separate competing inventory.
Use-case definition and value through the lifecycle	Demand exists for AI-enabled insight, prediction and support. The route from idea to operational use needs to be clearer.	AI ideas and pilots are emerging across predictive, vendor and clinical productivity use cases.	Apply a consistent use-case route covering problem, user, value, risk, data need, evaluation approach and intended operating model.
Evaluation and validation by AI type	Different AI patterns need different evaluation, including performance, explainability, grounding and user review.	Research and pilot activity uses expert feedback, but evaluation criteria are not yet consistent across AI types.	Define proportionate evaluation expectations for traditional ML, foundational AI, RAG, agents and vendor AI.
Guardrailed experimentation and vendor AI control	Staff and services are interested in AI tools and vendor products. The route from experimentation to managed use needs to stay clear.	Demand exists for clinical productivity tools and vendor AI, alongside emerging governance routes.	Apply consistent technical evaluation, asset visibility and review evidence to vendor and internally developed AI.

Pipeline And Platform Pillar

AI needs a clear technical route from exploration to tested, supported use

What this means

- Pipelines move data, code, configuration and artefacts through build, test, release and support.
- Platforms provide the environments, tools, integrations and controls that those pipelines run on.
- The route will differ for traditional ML, foundational AI, Copilot agents, vendor AI and BI-enabled use cases.

Why it matters

- AI delivery is harder to repeat if teams rely on local scripts, one-off environments or manual handovers.
- Platform choices should aim reduce duplication across Databricks, Microsoft, Posit and integration-layer work.
- Release and rollback routes help teams understand what changed, what it depended on and how to recover.

Pipeline And Platform Themes

Five themes used to assess whether Hywel Dda can move AI from experimentation to repeatable delivery. Visualisation of Current State Architecture & Target State Architecture can be found in the appendix.

Theme	Current status	Evidence signal	Next steps
Platform roles and delivery routes	Databricks, Microsoft, Posit, SQL Server and the CGI / Hywel Dda integration-layer project all have potential roles in AI delivery.	Databricks is a strategic direction, SQL Server and Power BI already support reporting, Posit supports data science outputs, and integration-layer work is in flight as a parallel project.	Assess platform roles against defined priority AI use cases. Agree where each platform should be used and where duplication should be avoided.
Development and runtime environments	Current capability includes SQL jobs, stored procedures, Python/Linux environments and BI workspaces. The production route for AI is less clear.	Data science delivery uses a defined technical environment, while some reporting and quality routines use warehouse SQL and Excel/Power BI patterns. Databricks progress depends on Azure/platform readiness.	Distinguish exploration, pilot and operational environments. Use priority use cases to define minimum environment expectations.
Reusable delivery patterns	Reusable patterns for ingestion, validation, masking, retrieval, prompts, evaluation and monitoring are not yet joined up as a standard approach.	Current work includes SQL reporting routines, data quality worklists, Posit-hosted analytical products, research validation, and emerging platform direction.	Build reusable patterns through early AI use cases, starting with data access, validation, permissions, evaluation and monitoring.
Deployment, release and rollback control	Source control, release management and environment separation appear uneven across current data, BI, data science and experimental AI routes.	Code and artefact management have been recognised during interviews as areas for improvement.	Define a lightweight release route for AI-related data products, models, agents and dashboards, including return-to-known-state expectations.
Artefact and dependency management	Prompts, retrieval indexes, datasets, notebooks, models, evaluations and configurations will need to be traceable as AI use grows.	Evidence points to many artefact types already emerging: notebooks, code, dashboards, reports, prompts, generated outputs, model work and data-sharing documentation.	Link key technical artefacts to the relevant information asset record so changes, dependencies and evaluation evidence remain traceable.

Monitoring And Operations Pillar

AI needs clear monitoring, support and review once it moves beyond experimentation

What this means

- Monitoring covers use, quality, value, cost, issues and ongoing performance.
- Traditional ML may need drift and accuracy checks; foundational AI may need output, grounding and user feedback checks.
- Operational support covers ownership, incident routes, defect handling and review cycles.

Why it matters

- AI can change over time as data, prompts, models, users and suppliers change.
- Benefits, usage and issues need to be visible after pilot, not only during approval.
- Existing monitoring patterns can be reused, but AI-specific signals need to be taken into account.

Monitoring And Operations Themes For AI Use

Five themes used to assess whether Hywel Dda can keep AI use measurable, supportable and under review.

Theme	Current status	Evidence signal	Next steps
Business validation and value tracking	AI ideas and pilots are emerging, but value measures are not yet consistently visible across the lifecycle.	Existing performance reporting uses dashboards, escalation summaries and review patterns.	Use existing reporting practice to define baseline, expected benefit and review points for priority AI use cases.
AI-specific monitoring by type	Current monitoring strengths are strongest around reporting and data quality, not yet AI outputs, model drift or grounding quality.	Data quality worklists, operational reports and performance dashboards provide useful monitoring patterns, but live AI monitoring evidence is limited.	Define proportionate monitoring expectations for traditional ML, foundational AI, RAG, agents and vendor AI. Include data quality, output quality and human feedback where relevant.
Cost, usage and adoption monitoring	Wider AI adoption will create new cost, usage and value questions, especially for licensed tools and platform services.	Copilot and other AI routes are visible, but benefits and usage evidence need to be connected to decisions.	Track usage, cost and adoption for priority AI tools and forecast before wider rollout or scale-up.
Operational support and ownership	Live AI use needs named business, technical and operational ownership, linked to the AI asset view.	Existing governance and data quality routes show ownership can be made visible where processes require it.	Link operational ownership, review cycle and support route to each material AI asset. Make support expectations explicit before pilot-to-live transition.
Technical issue and defect management	AI-related technical issues need to connect to existing service, data, platform and supplier routes.	Responsible AI work covers incident response; technical readiness needs clear hand-offs into those routes.	Map AI-related defects, data issues and supplier issues to existing reporting, triage and escalation routes.

Microsoft Copilot Readiness

Where we are, the risk that matters, and what we recommend

Copilot Readiness Summary

CGI advisory assessment across the dimensions of Copilot readiness. Overall position: **conditionally ready**. The platform and the appetite are in place; the controls are not yet, so scale should follow the governance, not precede it.

Readiness dimension	Status	Basis (verified figures)
Adoption and demand	ESTABLISHED	97% of paid users active in 90 days; median 144 prompts; strong staff-survey appetite
Foundational governance	DEVELOPING	Information Governance sign-off and a platform Data Protection Impact Assessment (DPIA) in place; no Copilot Acceptable Use Policy, DPIA, Application Form or approval gate
Information protection	ATTENTION REQUIRED	8 sensitivity labels defined but 0 applied; the Copilot Data Loss Prevention policy(DLP) is in test mode, not enforcing
Data access and oversharing	ATTENTION REQUIRED	External sharing on for 15,002 of 16,087(DHCW –Tenant Level) inactive sites (93%); Restricted Content Discovery(RCD) review begun, output pending
Identity and access	DEVELOPING	Conditional Access broadly enforced (124 policies); privileged access largely standing (tenant-level, DHCW)
Copilot Studio and agents	ATTENTION REQUIRED	105 agents, 99 without a recorded owner, 64 in unmanaged environments; no approval gate or catalogue
Monitoring and assurance	DEVELOPING	Audit and Copilot interaction monitoring available; assurance model still maturing

Established in place and operating. **Developing** partially in place, or evidence incomplete. **Attention required** not yet in place, or a material gap.

Reflects CGI's assessment confidence from the evidence reviewed, not a formal compliance benchmark. The basis column cites measured figures from the DHCW tenant reports.

Copilot licensing and adoption

Three distinct Copilot capabilities are in use, with different licensing and reach. Adoption of the paid assistant is high, and appetite to expand is strong.

Microsoft 365 Copilot Chat	Microsoft 365 Copilot	Copilot Studio
<i>Free, included tier</i>	<i>Paid, 120 licences</i>	<i>Agents, separate licence</i>
▪ Web-grounded chat available through Microsoft 365 licensing	▪ Productivity assistant in Word, Excel, Teams and Outlook ▪ 120 licences, individually assigned ▪ 97% active in 90 days; median 144 prompts	▪ Low-code agent-building platform ▪ 105 agents already built ▪ Licensed separately; seat count not in the reports

DHCW reporting identified approximately 120 Microsoft 365 Copilot licences, high levels of active usage, and an established Copilot Studio estate.
Sources: DHCW reports R-01, R-02 (23 May 2026) and R-13; the staff survey; the Digital Operations interview.

Copilot adoption is ahead of formal governance

EXISTING FOUNDATIONS

- M365 platform signed off by Information Governance
- Platform-level DPIA and Microsoft agreements in place
- Public consumer AI tools are restricted
- Copilot enabled following platform-level assurance
- Conditional Access broadly enforced within the DHCW tenant

GOVERNANCE PRIORITIES

- Copilot license application and assessment process
- Confirm whether a Copilot-specific DPIA is required
- AI risk classification approach and Acceptable Use Policy for Copilot
- Approval and review process for deployment of Agents
- Extend existing governance controls to cover Copilot Agents
- Review and agree Copilot DLP approach

The assessment found that governance foundations are already in place. The next step is extending these existing controls, processes and oversight arrangements to Copilot agents and future AI capabilities.

Some governance and security controls operate at DHCW tenant level and require joint ownership.

The risk that matters most

Copilot surfaces information user can already access. Where information classification is limited and sharing practices have not yet been fully reviewed, Content may be surfaced more widely than intended.

Permissions are inherited

Copilot creates no additional/new access. It reflects what a user can already open across SharePoint, Teams, Outlook and OneDrive.

Information Classification remains Limited

Eight sensitivity labels are defined but none are applied, so 0% of content is classified to guide what Copilot surfaces.

Oversharing risk requires review

External sharing is enabled on 15,002 of 16,087 inactive SharePoint sites (93%), increasing the importance of implementing a retention policy, information classification and access governance.

This is primarily an information governance and access management challenge rather than a technology issue.

Copilot Studio: Current Agents Landscape

Agent development is already occurring across the organisation. Copilot Studio is open to all licensed users, with no approval step and no catalogue. This is a Hywel Dda (HDD-scoped) view; agents owned by DHCW and other NHS Wales bodies are not included.

105



Agents already built

99



No recorded owner

64



Unmanaged environments

55



Production

GOVERNANCE STATE

- Available to licensed users
- No approval gate before publishing
- No formal catalogue or register identified
- Agents inherit the builder's access

AGENT DISTRIBUTION

- 30 Hywel Dda environments (HDD-scoped)
- 73 agents within single solution (Portal Booking)
- 55 production agents
- 45 in sandbox

DHCW Tenant report R-13 and Copilot Studio Agent inventory. Most Agents currently have no recorded owner, and many are deployed outside managed environments.

Copilot Studio: establishing governance

The aim is not to stop agent-building, but to put light, clear guardrails around it so it can continue safely.

CONTROL THE RISK

- **Approval gate** no agent reaches production without a defined approval and risk check.
- **Catalogue and ownership** a register of every agent, each with a named, accountable owner.
- **Lifecycle and review** test-and-release discipline; retire stale or ownerless agents on a schedule.
- **Control the Risk** assessment and classification before deployment

ENABLE SAFELY

- **Environment strategy** managed environments and clear rules on who can build where, agreed with DHCW.
- **Data guardrails and Access controls** Data guardrails and access controls aligned to agreed governance policies.
- **Monitoring ,Reporting and Operational oversight** to support ongoing governance and assurance .

Some controls operate at DHCW tenant level. Clarifying the Hywel Dda and DHCW responsibility model is a key first step .

A practical way forward

Move quickly, but safely

Stabilise & Govern

- Copilot AUP within the AI governance framework
- Formalise the License Application Form/official request process and DPIA
- Review SharePoint permissions and inactive sites/content that can be archived, begin labelling
- Copilot Studio approval gate and catalogue
- Clarify Hywel Dda and DHCW responsibilities for Copilot governance and controls

Enable & Scale safely

- Controlled adoption with clear criteria
- Establish ownership and visibility of AI solutions
- Support and oversight arrangements
- Monitoring and assurance
- Structured evaluation route for new tools

Optimise & Mature

- Benefits realisation and outcome measurement
- Scale approved, governed use cases
- AI operations, monitoring and support
- Periodic governance, risk and assurance reviews
- Continuous improvement

User Research

Purpose of the User Research

Working with CGI, Hywel Dda has undertaken an AI Readiness Assessment as part of its wider digital transformation programme to understand how AI can be adopted safely, effectively, and responsibly across the organisation.

This user research was conducted to establish a high-level organisational baseline of AI readiness from a user and stakeholder perspective, helping to identify:

-
- Current awareness and usage of AI tools
 - Confidence, capability, and digital readiness
 - Barriers to adoption and organisational challenges
 - Perceptions of trust, governance, and risk
 - Emerging opportunities and areas of value for AI adoption

The research aimed to complement the wider technical readiness assessment by exploring the people, operational, and organisational factors that may influence future AI adoption across Hywel Dda.

Research Methods & Observations

Participant outreach

We aimed to gather insight from a range of different stakeholders across multiple departments and roles.

The Survey was successful in capturing this diverse range, but 1:1 interview participation was lower, due to non-responsive contacts and busy schedules.

As a result, findings were largely dependent on the Survey results which may not represent the full range of perspectives across all departments within Hywel Dda.

Who provided the feedback

Survey results have shown that 50 respondents (17%) were from Managers. The remaining respondents are from varied roles and departments.

Some departments had a lower response rate, such as Diagnostics and Planned & Emergency Care. Encouraging participation from these departments could ensure a higher response rate in the future.

Survey responses

The digital Survey was live from the 29th April and closed on the 11th May. During this period, it received a total of 433 responses.

The Survey exceeding the expected 1% response rate and was closer to 3%. Maximising the number of responses is crucial for ensuring a representative and diverse sample.

Readiness Position: User & Organisational perspective

Strengths

- + High organisational awareness of AI
- + Positive appetite for AI
- + Existing experimentation
- + Emerging use cases

Challenges

- Governance immaturity
- Uneven capability and confidence
- Inconsistent adoption
- Lack of guidance/training

Hywel Dda demonstrates emerging AI readiness through strong awareness and growing interest in AI adoption. However, governance, capability and organisational support gaps currently limit scalable and consistent adoption.

According to the user research outcomes there is stronger evidence of interest and emerging appetite for AI adoption than there is outright resistance to AI itself. However, establishing strong AI Governance and building trust will act as core enablers to future AI adoption

Current AI Awareness & Usage

The survey findings suggest Hywel Dda has a **growing awareness** and **early adoption** of AI across the organisation, but **usage maturity** and **confidence** remain **inconsistent** across teams and roles. Survey data shows that AI is currently being used primarily for **low-risk productivity** and **administrative support** rather than embedded operational or clinical transformation.

Some teams and individuals are already **experimenting** with AI tools and **identifying opportunities** for use, while others have **limited exposure**, **low confidence**, or **uncertainty** around how AI could be applied to their role or area of work.

Current organisational understanding of AI is largely centred around **assistive**, **administrative** and **supportive** use cases rather than transformational automation.

High awareness but inconsistent adoption

- **402** respondents had heard of Microsoft Copilot and **392** had heard of ChatGPT or similar tools, indicating strong organisational awareness of AI technologies.

However, **AI usage maturity** remains inconsistent:

- **94** respondents reported using AI **daily** or **multiple times a day**
- **109** respondents reported using it **a few times a week**
- **121** respondents reported **never** using AI at work

AI use is currently focused on low-risk productivity support

The most common use cases were:

- summarising information (**230** respondents)
- drafting emails/documents (**191** respondents)
- administrative tasks (**152** respondents)

Confidence, Skills & Capability Readiness

The findings suggest that while many staff are **digitally confident** in their current systems, confidence in using AI specifically is **inconsistent** and still in **development** across the organisation.

Barriers to adoption include appear to relate less to resistance toward AI itself and more **to gaps in knowledge**, the need for **structured training**, governance **clarity**, and organisational **support**.

While there is evidence of **adoption** and **interest** in AI tools, levels of **confidence**, **understanding**, and **practical experience** **vary** significantly across teams and roles.

AI confidence is lower than general confidence

- **193** respondents reported being somewhat confident and **110** reported being very confident using **digital systems**.
- Compared to **130** respondents being somewhat confident and **47** being very confident using **AI tools** at work. .

Organisational confidence in AI adoption is mixed

- **44** respondents reported they were **not confident at all** in using AI tools
- **182** respondents were unsure what AI tools were available
- **148 respondents** identified lack of guidance/policy as a barrier

Training and enablement are major readiness gaps

The biggest barriers identified were:

- Lack of knowledge/training (**207** respondents)
- Uncertainty around available tools (**182** respondents)
- Concerns about data privacy/security (**150** respondents)
- Lack of guidance and policy (**148** respondents)

Governance, Trust & Risk Awareness

The findings suggest that **governance, trust, and risk awareness** are **significant factors** influencing AI readiness across Hywel Dda.

Data privacy and security concerns emerged as particularly significant themes throughout both the survey and stakeholder interviews. Concerns were raised around the handling of **sensitive information, reliability** of public AI tools, and the **potential risks** associated with using AI without **appropriate safeguards** or oversight.

Clearer communication around **approved tools, governance processes, security expectations**, and **safe usage guardrails** will likely be important in building organisational confidence and enabling more consistent AI adoption moving forward

Data privacy and security and major concerns

The most reported concerns relating to AI usage were:

- Data privacy/security concerns (**193** respondents)
- Lack of accuracy or reliability (**266** respondents)
- Accountability and who is responsible (**120** respondents)
- Patient safety and clinical risk (**84** respondents)

Governance awareness appears limited

Survey responses suggest many staff are unclear about approved AI usage, organisational policy, governance processes and what tools are permitted.

- **182** respondents reported they are unsure what AI tools are available to them
- **165** respondents were unsure of whether they currently have a Microsoft Copilot license
- **220** respondents reported little or no clear guidance to AI usage at Hywel Dda

Governance, Trust & Risk Awareness

Trust in AI across Hywel Dda is currently **cautious** and **conditional** rather than fully established. While many staff recognise the potential value of AI and are already using AI tools to support productivity and administrative tasks, there are **significant concerns** around **governance, reliability, accountability, and safe usage**.

Our findings suggest that **governance** and **assurance** will be **critical enablers** of future AI adoption within Hywel Dda. Building trust is likely to depend on establishing **clear organisational policies, approved toolsets, structured evaluation processes**, and **visible governance mechanisms** that allow staff to use AI safely and confidently within defined boundaries.

Trust in AI outputs is mixed

The most reported concerns relating to AI usage were:

- **181** respondents felt neutral about how much they would trust AI to support their work
- **156** respondents reported they mostly trust AI to support their work

However, building trust is still a necessity to moving towards AI adoption as:

- **109** respondents reported a lack of trust in AI outputs preventing them from using AI more in their role

Concerns were raised around:

- Accuracy and reliability
- Bias
- Overreliance on AI
- The need for human oversight/Accountability

Interview participants commonly described AI as useful for support tasks but not yet trusted for critical decision making without validation

Operational & Technical Readiness

Operational and technical readiness for AI adoption across Hywel Dda is **still developing** and remains **inconsistent** across the organisation.

While there is evidence of **growing experimentation** and practical AI usage, support structures are not yet fully established to provide staff with the knowledge of what tools they **can and can't use**, or what specific AI tooling is available to them.

Many who do have an awareness of the AI technologies in use across the organisation face **difficulty** with finding the **time to experiment** and **develop AI proficiency** and understand where it could be applied to their work and roles.

Survey responses highlighted uncertainty around available AI tooling

- **182** respondents said they were unsure what AI tools were available to them
- **165** respondents were unsure whether they had a Microsoft Copilot license
- **77** respondents reported they did not have access to Copilot

Survey results combined with interview discussions also suggested that there is **fragmented use** across departments, **uncertainty** around **approved tools**, and varying levels of **access to AI technologies**.

Time and capacity constraints contribute to adoption barriers

Operational pressures were identified as a challenge:

- **133** respondents said they do not have enough time to explore or learn AI tools

Interview participants also referenced **workload pressures** and **limited time** for experimentation **restrict AI learning and adoption**, particularly in Clinical and high-pressure service areas.

Organisational Culture & Change Readiness

Organisational sentiment towards AI is **generally positive** with staff largely viewing AI as a **supportive tool** for **improving efficiency** rather than replacing key activities or roles.

Our findings reinforced the view that the organisation's primary readiness gap is not access to technology itself as experimentation is already apparent, and **appetite for AI already exists**.

The broader capability required is to **develop mature governance, improve AI literacy and visibility**, and **support staff to build confidence** in implementing AI.

Staff interest in AI adoption is generally positive

- **76** respondents rated AI as **Extremely** valuable
- **88** respondents rated AI as **Very** valuable
- **87** of respondents rated AI as **Moderately** valuable
- **88** of respondents rated AI as **Not at all** valuable

It's important to note that some lower value perceptions may reflect users who have had limited exposure or opportunity to use AI, rather than solely negative experiences with AI itself.

Staff want Structured Support & Organisational Enablement

The most requested were:

- Formal training sessions (**305** respondents)
- Access to approved AI tools (**266** respondents)
- Written guidance and policies (**225** respondents)
- Ongoing support/Helpdesk (**208** respondents)

Opportunities & Potential Value Areas

Staff across Hywel Dda already recognise practical opportunities where AI could support **efficiency, reduce administrative burden, and improve access** to information.

Current perceptions of value are largely centred around **productivity enhancement** and **operational support** using Copilot or ChatGPT, rather than fully transformational or autonomous AI use cases.

However, certain departments and clinical services are already exploring more **advanced, workflow-specific AI capabilities**, which shows evidence of **meaningful AI opportunities** for AI-enabled improvement that targets operational and clinical workflows.

Administrative burden reduction is a large area for opportunity

There appears to be strong organisational appetite for AI solutions that reduce administrative pressures. This interest exists across both Clinical and Corporate functions with many respondents describing AI as being most useful for:

- Reducing manual administrative work
- Summarising large amounts of information
- Supporting documentation
- Improving productivity in day-to-day tasks

Clinical & Operational AI use cases are evident

Although prevalence was lower, staff show awareness and use of Clinical based tools then demonstrate opportunity for AI implementation in the future beyond tools that are solely dedicated to administrative support

- **72** respondents have heard of the AI Scribing tool
- **57** respondents have heard of the AI Diagnostic tool
- **22** respondents have heard of the AI Pathology tool
- **Totalling 15% of respondents overall**

Confidential

Acknowledgements

Stakeholder Acknowledgement

We would like to thank the following Hywel Dda colleagues for sharing their time, experience and insight during the AI Readiness Assessment.

Contributor	Role
Anthony Tracey	Digital Director
Gareth Beynon	Head of Information Services
Chris James	Head of Data Development
Gareth Jenkins	Head of Data Science
Patrycja Duszynska	Head of Information Governance
Gavin Jones	Head of Digital Operations
Claire Langdridge	Data Quality Improvement Manager
Tracy Price	Head of Performance
Leighton Phillips	Director Research, Innovation and Value
Anthony Smith	Consultant Anaesthetist / CCIO
Paula Goode	Service Director for Planned and Specialist Care
Dana Scott	Director of Midwifery and Professional Governance for Women and Children
Olwen Morgan	Assistant Director of Nursing
Ihab Abbasi	Associate Medical Director, Planned and Specialist Care
Michelle James	Head of Resourcing and Utilisation
Sarah Barnes	Workforce Manager: Systems and Workforce Intelligence
Tracy Walmsley	Assistant Director of People Planning
Nicola Henwood	People Programme Manager
Megan Morgan	PhD Student / Research-Linked Data Science
Praboda Rajapaksha	Lecturer in Data Science / Researcher
Gareth Skye	Business and Governance Manager


Appendix

Supplementary material


Questionnaire results




Which Area Do You Work In?




What Is Your Role Level?




How Confident Do You Feel Using Digital Tools And Systems Required For Your Day-To-Day Work?




Which Of These AI Tools Have You Heard Of?




How Often Do You Use AI For Work?




What AI Tool(s) Do You Use?




Do You Have A Co-Pilot Licence Provided By Hywel Dda?




What Do You Use AI For?



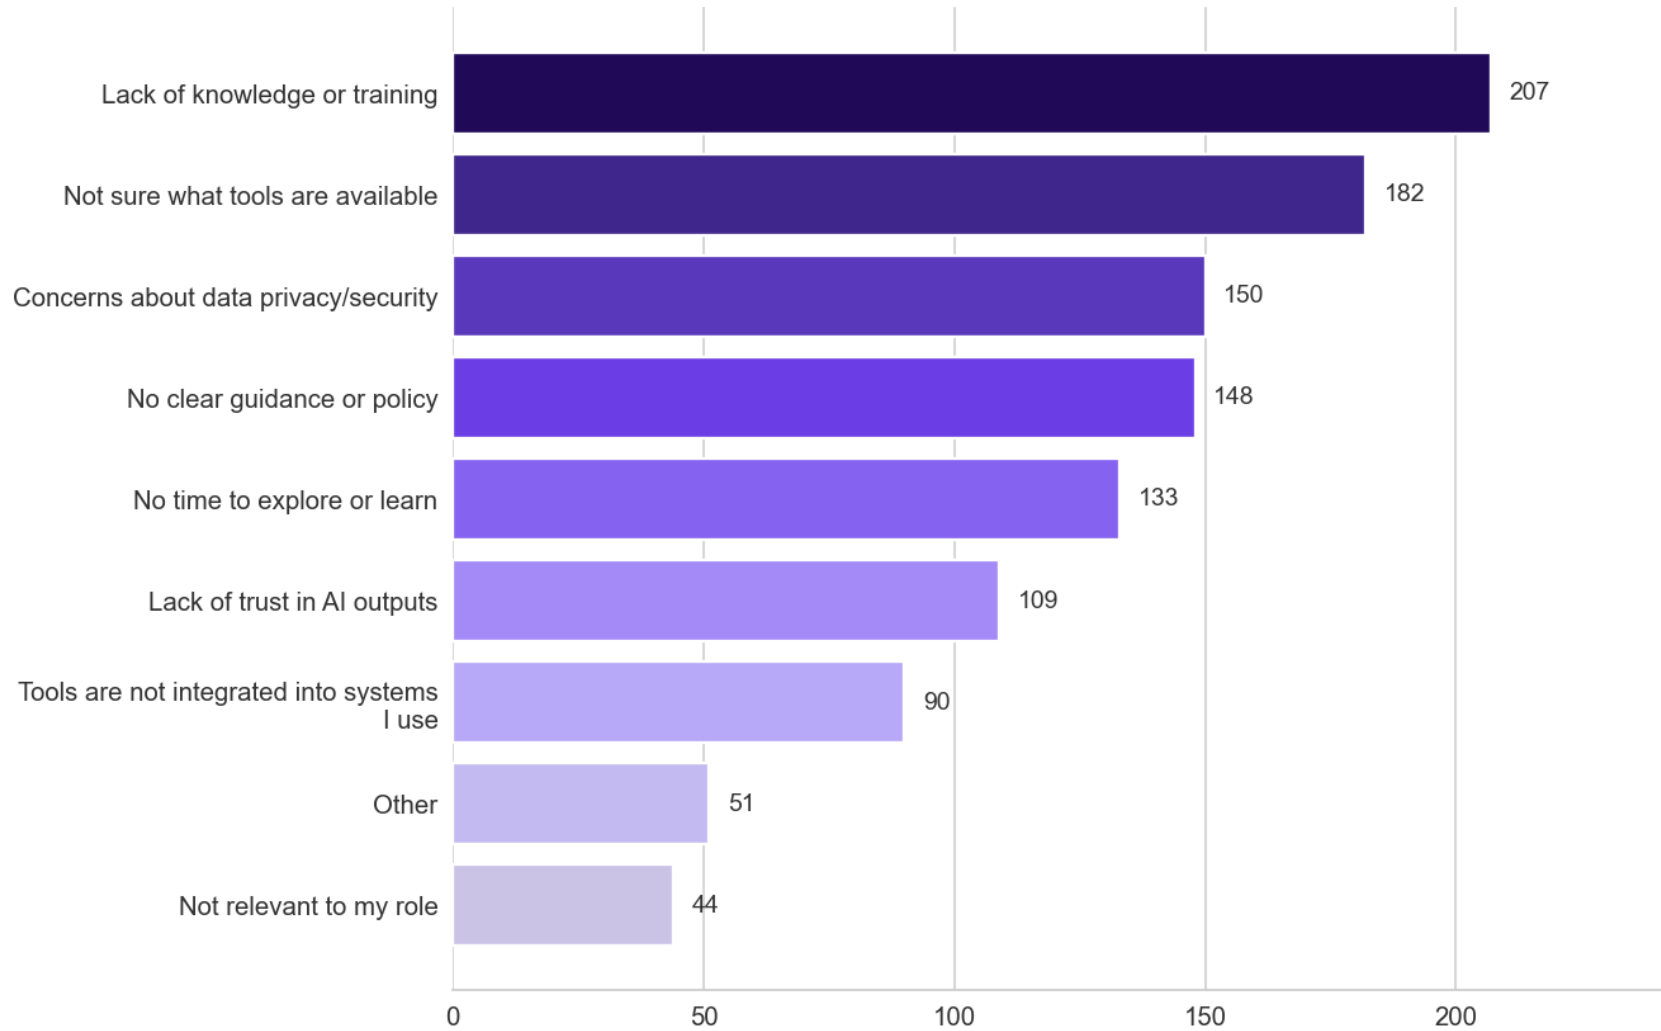
How Valuable/Useful Has AI Been In Your Work So Far?




Do You Feel Your Organisation Provides Clear Guidance On How AI Can Be Used Safely And Appropriately In Your Role?




What Prevents You From Using AI More In Your Role?




How Confident Do You Feel In Your Ability To Use AI Tools At Work?




How Much Would You Trust AI To Support Your Work?



What Is Your Main Concern About AI Being Used In Work? (Select Up To 3)



What Type Of Support Would Help You Use AI Effectively In Your Role? (Select All That Apply)



What Have Been The Main Benefits To Using AI In Your Work So Far?

