



NH TECH
ALLIANCE

NEW HAMPSHIRE AI TASK FORCE

The State of AI in New Hampshire

2026 Report

Convened under the NH Tech Alliance | October 6, 2026

Table of Contents

Foreword.....	3
About This Report.....	4
Part One: The Statewide Picture.....	7
Chapter 1. New Hampshire at the Inflection Point.....	8
Chapter 2. Where NH Organizations Stand: AI Adoption Across the State.....	11
Chapter 3. The Adoption Gap: People, Process, and the Human Side of Change.....	13
Chapter 4. The Data Readiness Problem.....	18
Chapter 5. Workforce & Talent: The Readiness Deficit.....	19
Chapter 6. AI Safety & Responsible Use Across NH.....	20
Chapter 7. Ethics & AI Governance: Organizations Can't Wait for Legislation.....	26
Chapter 8. ROI, Cost, and the Environmental Footprint of AI.....	27
Chapter 9. Access & Equity: Who Gets Left Behind.....	30
Chapter 10. Policy & Regulatory Landscape.....	31
Part Two: Sector Findings.....	32
Chapter 11. Technology & Software.....	33
Chapter 12. Financial Services & Insurance.....	36
Chapter 13. Healthcare & Life Sciences.....	39
Chapter 14. Manufacturing, Engineering & Defense.....	44
Chapter 15. Professional & Business Services.....	48
Chapter 15a. Real Estate Breakout.....	52
Chapter 16. Consumer & Small Business.....	56
Chapter 17. Public Sector.....	60
Chapter 18. Education.....	63
Chapter 19. Startups.....	66
Part Three: Recommendations.....	71
Chapter 20. A Framework for Action.....	72
Chapter 21. For State Government.....	74
Chapter 22. For NH Businesses and Institutions.....	76
Chapter 23. For Educators and Workforce Development Partners.....	77
Chapter 24. For the NH AI Task Force.....	78
Appendices.....	79
Appendix A. Task Force Members & Working Group Structure.....	80
Appendix B. AI Frameworks & Standards Reference Guide.....	82
Appendix C. Glossary.....	84
Appendix D. Chapter References.....	87

Foreword

Spend a few weeks listening to New Hampshire organizations talk about AI and you notice two conversations happening at once. In one, a hospital is documenting patient visits with ambient AI or a startup is rebuilding its product in a matter of weeks. In the other, an equally capable organization a few towns over is still asking whether it is even allowed to begin exploring AI. Both are true at the same time, and the distance between them is rarely about the technology itself.

That gap is the finding this report keeps returning to, and it is also the opportunity. The organizations out front are not working with different tools than everyone else. They have done the less glamorous work of deciding what problem they were actually trying to solve, preparing their data, setting up rules, and checking whether anything measurably improved. Most of what separates them from everyone else is learnable, and this is work New Hampshire can choose to do together rather than one organization at a time.

We tried to build this report the same way, from the ground up. Over the past year the Task Force held more than thirty roundtables, interviews, and working sessions across 8 sectors, and what you will read is grounded in what practitioners actually told us, not in predictions about where the technology is headed. It is written for two readers who rarely read the same document: the policymaker deciding what the state should do, and the operator deciding what to do on Monday. And it is written to invite curiosity instead of raising an alarm, because fear tends to produce either paralysis or bad bets, and New Hampshire cannot afford either.

I want to be clear about where the credit belongs. Every person who contributed to this report is a volunteer. They organized and ran the roundtables, pulled in colleagues who do not normally sit at the same table, and brought hard-won judgment from their own industries into a shared conversation. The candor in these pages, people willing to say plainly what is working and what is not, is the reason the findings hold up. That generosity is the real engine of this effort, and without it there is no report and no Task Force. To everyone who gave their time and their effort, thank you.

Our goal for this work is easy to say and hard to accomplish: a New Hampshire where every person, business, and organization can understand, explore, and apply these tools - not just the leading edge already racing ahead. This report is a step toward that and not the finish line. We intend to keep listening, to keep turning what we hear into something a busy leader can act on, and to keep showing up as a practical, non-partisan resource for the state. So read this report with the same candor that produced it. Tell us where we got it right, tell us where we got it wrong, and then help us close the gap it describes.

Phil Magnuszewski

Chair, New Hampshire AI Task Force

About This Report

Methodology and Data Sources

30+

roundtables, interviews, and working sessions

8

Sectors covered

12

months of evidence, Oct 2025 – Sep 2026

This report is built on what New Hampshire practitioners told the Task Force, along with key insights from the sector leaders. Its evidence base is more than thirty roundtables, interviews, and working sessions held from October 2025 through September 2026 across eight sectors: Technology and Software, Financial Services and Insurance; Healthcare and Life Sciences; Manufacturing, Engineering and Defense; Professional and Business Services (including a Real Estate breakout); Consumer and Small Business; Public Sector and Education (reported as two chapters); and Startups. Participation ranged from a handful of practitioners in a focused session to more than a dozen in the larger roundtables, and several sectors also drew on one-on-one interviews and a member survey.

Each sector chapter was drafted by its named lead or leads from those conversations, then edited into one consistent report voice without changing the underlying findings. A cross-sector synthesis reads the sectors side by side to surface what repeats, diverges, or compounds.

One ground rule shaped the work. Roundtables were candid and confidential, so no observation is attributed to a named individual without permission, and observations are reported at the sector-level or anonymized throughout.

This report also builds on the Task Force's February 2026 memo, which set out ten early signals about how New Hampshire was experiencing AI. Those signals are tested against the deeper 2026 evidence in the cross-sector synthesis and, sector by sector, in each chapter.

The Adoption Scale

To compare progress across very different sectors, the report uses a single five-stage adoption scale: Awareness (talking, not acting), Active (funded pilots and real experimentation), Operational (live in a workflow somewhere), Systemic (embedded across most workflows), and Transformational (reshaping the business model). Each sector placed both its leading edge and its typical organization on this scale, and the distance between those two points is the sector's adoption gap. Chapter 2 compares those placements across sectors.

In plain terms: each stage describes how deeply AI is woven into an organization, from first conversations to a changed business model. Most organizations move through the stages in order, and no stage represents a failure. It is simply a starting point.

The Five Stages of AI Adoption

How this report describes where an organization is on its AI journey




Figure A. The five-stage adoption scale used throughout this report. Chapter 2 shows where each New Hampshire sector places itself.

Separately, each sector rated the maturity of its AI risk governance using three plain labels: "ahead of the curve", "still figuring it out", or "not yet a priority". Those ratings appear in each chapter's AI Safety and Risk section.

Frameworks Used

The report treats the NIST AI Risk Management Framework as its common reference, because it is voluntary, widely recognized, and maps onto requirements taking shape elsewhere, which lets a small business and a state agency speak the same language at different depths. Chapter 6 layers the recognized security and management standards on top of that anchor, including the NIST Generative AI Profile, ISO/IEC 42001 and 27001, the OWASP Top 10 for LLM Applications, MITRE ATLAS, and the CISA and NCSC secure-by-design guidelines. Additional information on these frameworks is provided in Appendix B.

How to Use This Report

The report is written for two audiences, and each can read it differently.

Government & Policy

Legislators, agencies, and policy staff

START WITH

Part One, the statewide picture

THEN GO TO

Part Three, particularly Chapter 21 (For State Government)

USE IT TO

Ground legislation and agency guidance in the cross-sector findings on data readiness, workforce, safety, ethics, access, and the policy landscape

Practitioners

Business leaders, educators, and the people doing the work

START WITH

Your own sector in Part Two

THEN GO TO

Part Three, particularly Chapter 22 (For NH Businesses and Institutions)

USE IT FOR

Practical, right-sized guidance and the named tools in the appendices

Short on time? Follow this four-step path.

1

Chapter 1

For the state of things

2

Chapter 2

Skim the adoption comparison

3

Part Two

Read your own sector chapter

4

Part Three

Recommendations for your specific audience

A Note on Tone

This report is written to invite curiosity, not to overstate urgency. The goal is to help every person, business, and organization in New Hampshire understand, explore, and apply these tools thoughtfully, so the tone throughout favors honest synthesis and practical next steps over alarm. The recommendations in Part Three are a starting point for discussion, offered so the Task Force and its partners can react to them, refine them, and own them.

Acknowledgments

This report reflects the work of the sector leads and co-leads named on each chapter, the many New Hampshire practitioners who gave their time and candor in the roundtables and interviews, and the organizations that agreed to share their stories. The full Task Force roster and working-group structure appear in Appendix A.

Part One: The Statewide Picture

Chapters 1–10

Chapter 1. New Hampshire at the Inflection Point

Building on the February 2026 Memo

The February 2026 memo set out ten early signals about how New Hampshire was experiencing AI. Tested against seven more months of roundtables and the cross-sector findings, nine of the ten hold, several of them strongly, and one is under-evidenced rather than wrong. That is the headline change since February: the early read was largely correct, and the evidence base behind it is now materially deeper.

What has been confirmed. AI is already in use but adoption is uneven; data readiness is the primary bottleneck; AI is seen mostly as efficiency, not replacement, but with early replacement pressure at the entry level; workforce readiness lags tool availability; demand for practical guidance is strong; risk awareness is high while governance maturity varies; access concerns are emerging; and policy uncertainty is slowing some decisions. Each of these now has independent support from multiple sectors, not a single early reading.

What we understand more clearly. Three things are sharper than they were in February.

First. The widest gap is inside sectors, not just between them: in nearly every sector the distance between the leading-edge organization and the typical one spans two or three maturity stages.

Second. Shadow AI, meaning employees using personal and unsanctioned accounts, is moving faster than any sector's formal adoption, so the real baseline is higher than any official inventory shows.

Third. The most consistent ask across otherwise unrelated sectors is practical, plain-language guidance and shared templates: an AI acceptable-use policy, a maturity self-assessment, and checklists. A workable procurement pathway is the specific top need of the public sector, where municipal purchasing rules were not built for how AI is priced and consumed.

The one signal to revisit. Heavy reliance on vendors shows up indirectly, in complaints and small businesses' frustration with vendor noise, but it is not yet documented in depth across sectors. The report treated it as under-evidenced, and the Task Force will ask about it directly in the next roundtable cycle.

The February Signals, Checked Against Cross-Sector Findings

Nine of the ten signals from the February 2026 memo hold up against this cross-sector data; one is under-evidenced and worth a targeted follow-up question in the next roundtable cycle.

#	February Signal	Verdict	Evidence
1	AI is already in use, but adoption is uneven	Holds, strongly	The spectrum runs from Systemic/Transformational (startups) to Aware/Experimenting (government, real estate) inside the same state.

#	February Signal	Verdict	Evidence
2	Data readiness is the primary bottleneck	Holds	Named independently in Manufacturing (legacy PLM/ERP silos), Healthcare (PHI / staff / financial data location unknown), and Finance (24-hour stale-data gaps).
3	AI is seen as efficiency, not replacement	Holds, with exception	True for most sectors, but junior-role automation, specifically in Legal and Technology sectors show early replacement pressure at the entry level.
4	Workforce readiness lags tool availability	Holds	Every sector in the findings names a skills or training gap; no one reports training pipelines are ahead of tool adoption.
5	Demand for practical guidance is strong	Holds	Policy templates, playbooks, and readiness tools are requested by Government, Small Business, Manufacturing, and Education alike.
6	Risk awareness is high, but maturity varies	Holds	Regulated sectors (Finance, Healthcare, Government) are furthest along; most others remain informal or ad hoc.
7	Heavy reliance on vendors is the norm	Under-evidenced	Implied by procurement and vendor-pitch-overload complaints from Small Business, but not documented in depth across all 8 sectors; worth a targeted follow-up.
8	Access concerns are emerging	Holds	Enterprise pricing does not scale to towns under 150 employees; Education flags a widening digital divide; Small Business cites resource constraints.
9	Policy uncertainty is slowing some decisions	Holds, strongly	No comprehensive state AI statute exists; most 2026 private-sector bills failed; federal preemption is an open threat.
10	New Hampshire is at an inflection point	Holds	The scale of the maturity gap, combined with mounting regulatory and procurement pressure, describes a genuine decision point rather than steady drift.

The Case for Acting Deliberately in This Window

The combination of a large maturity gap, procurement mechanics that were not built for how AI is priced and bought, and unresolved regulatory uncertainty describes a genuine decision point, not steady drift. The finding across sectors is not that New Hampshire must rush. It is that the conditions for deliberate, coordinated action are unusually favorable right now: the needs are consistent, the highest-value levers are ones the state can realistically influence, and the organizations furthest ahead are already demonstrating what works. Acting deliberately in this window means moving on the shared enablers, practical guidance, shared policy templates, a

common maturity baseline, and a workable procurement pathway for the public sector, before the gap between leaders and laggards hardens.

Chapter 2. Where NH Organizations Stand: AI Adoption Across the State

Adoption by Sector

Read across sectors, New Hampshire's adoption runs the full length of the five-stage maturity scale at the same moment. Startups sit at the far end, at Systemic to Transformational, building AI products, software-development systems, and internal operations. In Technology, Financial Services, Healthcare, Manufacturing, Professional Services, and Education, the leading edge has reached Systemic while the typical organization is Active. Consumer and Small business and Real Estate lead at Operational with the typical organization at Active. The Public Sector shows the widest spread from Systemic at the leading edge to Awareness for the typical municipality.

Adoption by Organization Size

The most consistent finding across the sector chapters is that organizational resourcing, not sub-industry, drives the gap. Financial Services attributes the distance between leading and typical institutions directly to size, because the largest institutions have the resources to invest in AI maturity. Manufacturing shows the same divide through data infrastructure, where firms already on integrated PLM platforms move faster than those on scattered spreadsheets. The Technology sector debated which end of the size spectrum leads but converged on resourcing as the cause. In the Public Sector the divide is capacity, with smaller municipalities often lacking any dedicated data, cybersecurity, legal, or IT staff.

The Adoption Gap, Visualized

The pattern is easiest to see side by side. Where sectors marked both their leading edge and their typical organization on the five-stage adoption scale, the distance is consistently about two stages, and it is widest in the Public Sector.



Figure 2.1. The adoption gap by sector. For each sector, the typical organization (magenta) and the leading edge (blue) are placed on the five-stage adoption scale (1 Awareness, 2 Active, 3 Operational, 4 Systemic, 5 Transformational); the connecting line is the gap. Source: Part Two sector chapters.

Adoption by Geography

The findings show geography mainly as a capacity and access divide, not a mapped regional pattern. The clearest geographic signals are that smaller and more rural municipalities carry the thinnest staffing and therefore the slowest path to governed adoption, and that access concerns are emerging where enterprise pricing does not scale to small towns and small businesses. The Task Force does not yet have a finer regional breakdown of adoption across the state, a gap for the next report cycle to address.

Chapter 3. The Adoption Gap: People, Process, and the Human Side of Change

Cross-sector Analysis: Alexis Brochu

Current State

The most accurate thing to say about New Hampshire's AI adoption is that there is no single state of it. The pattern is best described as Swiss cheese: wide knowledge gaps, siloed pilots, and no coordinated statewide baseline. Startups and sophisticated law firms are moving fast, in one Task Force member's words building agentic infrastructure and rebuilding tens of thousands of lines of code in a month. The public sector and traditional small businesses sit in a wait-and-see posture, stalled less by lack of interest than by policy uncertainty and procurement mechanics that were not built for how AI tools are priced and bought.

As Chapter 1 notes, the widest gaps sit inside sectors. The examples are concrete: the typical health system documents visits with ambient listening while the typical small practice does not, and the typical commercial bank is piloting fraud detection while the typical community bank may wait on its board. Shadow AI widens the distance between claim and practice, because employees adopt tools faster than any formal program can govern them.

The table below places each sector's typical organization on the five-stage adoption scale, alongside its primary focus. Placements follow the sector chapters where one exists; Chapter 2 shows the leading edge for comparison. Legal and Retail and Hospitality come from the cross-sector change management findings.

Sector	Maturity Signal (Typical Stage)	Primary Focus
Technology & Software	Active / Systemic	Deployment and process automation, agentic development, support chatbots
Financial Services & Insurance	Active / Systemic	Fraud detection, commercial underwriting, compliance
Manufacturing, Engineering & Defense	Active / Systemic	Technical documentation, predictive inventory, PLM/ERP integration
Healthcare & Life Sciences	Active / Systemic	Ambient documentation, diagnostic assistance
Professional & Business Services	Active / Systemic (Real Estate breakout Active / Operational)	Legal research and drafting, audit and financial analysis, client communications; in Real Estate, listing, valuation, and property visualization
Government (Public Sector)	Awareness / Systemic	Policy writing, legal review, early resident-facing chatbots

Sector	Maturity Signal (Typical Stage)	Primary Focus
Education (K-12 & Higher Ed)	Active / Systemic	Curriculum development, administrative efficiency
Consumer & Small Business	Active / Operational	Marketing, admin productivity
Startups	Systemic / Transformational	Agentic infrastructure, software factories, automated QA

Key Use Cases

Six workflows show up across sectors as concrete applications running today, not roadmap items:

- **Ambient clinical documentation.** Some healthcare systems report a 50% reduction in charting time; this is in production, not pilot, in multiple NH systems.
- **Resident-facing municipal chatbots.** Portsmouth has deployed a chatbot for construction project data; state agencies use similar tools for unemployment claims processing and wildfire tracking.
- **Contract review and legal research.** Half of surveyed attorneys were using generative AI by 2025 for drafting, discovery, and research, with human-in-the-loop review as the operating norm.
- **Fraud detection and commercial underwriting.** Community banks report 98 to 99% first-contact resolution on customer service and automated audit workflows, alongside faster commercial loan underwriting.
- **Technical documentation with predictive inventory.** Manufacturers report 50 to 80% cost reduction in technical documentation and translation, and estimate that AI-driven forecasting would cut required safety stock 25 to 50%.
- **Agentic software development.** NH startups describe AI rebuilding roughly 30,000 lines of code in a month and running 40-plus scheduled workflows through custom agents connected to internal databases, on timelines that previously took a year.

Primary Barriers

Six barriers surface repeatedly across sectors, and they compound on each other:

1. **Data fragmentation and legacy systems.** Named unprompted in Manufacturing (legacy PLM/ERP silos), Healthcare (uncertainty over where PHI actually lives once vendors absorb it), and Finance (stale-data gaps of up to 24 hours).
2. **Shadow AI and the absence of formal policy.** Employees across sectors, most visibly in government, use personal accounts without IT oversight, creating data-leakage risk that no inventory captures.
3. **Procurement and budget mismatch.** Enterprise pricing does not scale to municipalities with under 150 employees or to small businesses; this is the barrier named most consistently as the one that would do the most for adoption if it went away.
4. **Regulatory and policy uncertainty.** No comprehensive state AI statute exists. Most private-sector regulation bills died in the 2026 session, and federal preemption efforts

(Executive Order 14365, the DOJ AI Litigation Task Force) add uncertainty about which rules will hold.

5. **Junior talent pipeline erosion.** Legal and Technology both name automation of entry-level tasks as a direct threat to how new professionals are trained. This barrier is specific to knowledge-work sectors, not universal.
6. **The governance-agility gap.** Heavily regulated sectors want to move faster than their regulators and boards allow, the inverse of sectors held back by under-governance.

Workforce Implications

New senior roles are emerging at the top of the org chart. Organizations are creating Chief AI Officer positions while roles in the middle are being redefined rather than eliminated. Legal associates are shifting from drafting to reviewing and prompt-engineering; teachers are shifting from content providers to facilitators and evaluators of AI output; realtors are becoming legally responsible for the accuracy of AI-generated property descriptions; finance staff are moving from manual processing toward advisory, client-facing work.

The sharpest workforce risk sits at the entry level, covered in depth in Chapter 5. Healthcare's pressing issue is building the expertise to train, oversee, administer, and troubleshoot AI as adoption increases. Across every sector represented in the findings, training pipelines lag tool adoption.

Data Readiness

Data readiness is the fundamental blocker in most sectors that discuss it, and Chapter 4 covers it in depth. The change-management point is ownership: data governance is clearer in regulated sectors with existing compliance functions and largely undefined elsewhere.

AI Safety and Risk

The NIST AI Risk Management Framework (Govern, Map, Measure, Manage) is the recommended statewide anchor, covered in Chapter 6, and it fills a real gap given New Hampshire's narrow statutory patchwork. Governance maturity varies sharply by sector. Regulated industries (Finance, Healthcare, Government) are building written policy and assigning ownership; most other sectors operate on informal norms or nothing at all. The risks that surface most differ by sector: privacy and PHI exposure in healthcare, algorithmic bias in criminal-justice and predictive-policing tools, hallucinated citations and malpractice exposure in legal, and algorithmic discrimination across all sectors in employment tools. Regulation is not yet the primary driver of governance anywhere in the state; voluntary frameworks are doing that work for now.

Ethics

The ethical tensions repeat across sectors in recognizable shapes: who could be harmed, where automation meets professional judgment, and when people should be told AI is involved. Chapter 7 covers each. For change management, the consistent thread is human review, as legal, healthcare, and law enforcement all describe professionals shifting toward reviewing and attesting to AI-assisted work.

What Is Needed Most

The one thing that would move the needle furthest across sectors is not sector-specific: a statewide, standardized procurement pathway paired with a short, usable policy template. It is the single need that recurs independently across Government, Consumer and Small Business, and Education, sectors that otherwise have little in common. It is also realistic for New Hampshire to influence the efficient procurement of AI tools for towns through the Municipal Association, and a short acceptable-use, high-risk-use, and prohibited-use template modeled on the EU AI Act approach.

Close behind it: a self-scoring maturity assessment tool (Strategy, Governance, Data/Tech, Talent, Execution) that lets an organization place itself on the adoption spectrum without a consultant, and closing the Task Force representation gap for organizations that don't fit into the current sector breakout.

Additional Observations

- **Signal 7 needs more evidence.** Vendor dependency is implied (procurement complaints, sales-pitch overload in small business) but not documented with the same depth as the other nine signals across all 8 sectors. Recommend adding a direct vendor-reliance question to the next roundtable round.
- **This chapter is a rollup, not a replacement.** Each sector's own chapter remains the authoritative, first-person source on that sector's specifics; this chapter exists to surface what repeats, diverges, or compounds when the sectors are read side by side.

Full Cross-Sector Change Management Matrix

Condensed from the Cross-Sector AI Change Management Findings (8 sectors). Full source detail, including opportunity and benefit language and citation indices, is retained in the underlying findings workbook.

Sector	Change Management Focus	Sharpest Risk	Greatest Need
Government (State & Local)	Fragmented experimentation to intentional, accountable, enterprise-wide governance	Deepfakes, algorithmic bias in justice tools, Right-to-Know compliance	Statewide procurement, model policies, task-force coordination
Professional & Business Services	Legal: ethical use and quality control of generative AI, and the shift from billable-hour to deliverable-based work. Real Estate: ethical use in property marketing	Hallucinated citations, malpractice exposure, junior pipeline erosion; in Real Estate, deceptive-marketing liability and the environmental cost of data centers	Ethics-opinion-aligned CLEs, standardized playbooks; in Real Estate, guardrail guidance, liability clarity
Education (K-12 & Higher Ed)	Theoretical literacy to operational fluency and authentic assessment	FERPA/COPPA exposure, academic integrity, digital divide	Dedicated PD funding, statewide AI mission framework

Sector	Change Management Focus	Sharpest Risk	Greatest Need
Healthcare & Life Sciences	Governing a sprawling, disparate application ecosystem with clinical transparency	PHI exposure, black-box denials, algorithmic bias in treatment	Clinical-specific governance, BAA-compliant tools, disclosure systems
Finance & Insurance	Automating high-volume work while closing the governance-agility gap	UDAP non-compliance, missing audit trails, stale-data lag	Regulatory sandboxes, sector working groups, audit-trail tooling
Manufacturing / Eng. / Defense	Legacy silos to integrated PLM/ERP systems for AI readiness	IP leakage, supply-chain attacks, loss of institutional knowledge	Practical roadmaps, communities of practice, readiness assessments
Technology & Software Industry	Billable-hour services to outcome-based software factories	Threat to consulting model, unreviewed AI slop, junior pipeline risk	Higher-ed partnerships, affordable compute, lightweight governance
Consumer & Small Business	Closing the awareness gap and analysis paralysis with bite-sized guidance	Resource constraints, vendor noise without functional knowledge	SBA-style consulting networks, NH policy templates, peer exchanges

Chapter 4. The Data Readiness Problem

Data Readiness as the Foundational Bottleneck

Across the sector findings, data readiness is the constraint named most often as the one standing between AI and real impact, and it is a foundational blocker rather than a secondary one. It was raised unprompted in Manufacturing, where legacy PLM and ERP silos keep production, maintenance, and quality data from talking to each other; in Healthcare, where organizations frequently cannot say with confidence where protected health information lives once AI vendors have processed it; and in Finance, where stale-data gaps of up to 24 hours undermine real-time decisioning. The Manufacturing chapter is the clearest case, naming data readiness as the central bottleneck in every session, more than any other topic.

Fragmentation Across Legacy Systems

The dominant pattern is fragmentation. Organizations generally do not lack data; they lack unified, accessible, well-structured data. Manufacturers already on integrated PLM platforms such as Windchill and ARAS move meaningfully faster, while those working from scattered spreadsheets are advised to clean and organize that foundation first, because the underlying risk is garbage in, garbage out. Financial institutions describe the same problem as rigid siloing across legacy systems, addressed slowly through migrations to cloud-based data lakes. Education reports substantial data assets fragmented across multiple systems and departments, while in contrast, startups treat clean, well-structured, permissioned proprietary data as a central competitive asset that requires consistent identifiers, source provenance, and defined access boundaries.

Regulatory Constraints on Data

Regulatory constraints compound fragmentation and are not uniformly well understood. Education operates under FERPA and COPPA obligations, with K-12 vendor vetting and data-protection agreements capable of delaying adoption for months. Healthcare must manage BAA compliance and PHI handling. Finance and the Public Sector both weigh privacy and security of sensitive information before data can be used at all. In the Public Sector, that legal question, meaning whether an organization is even permitted to use AI on given information, frequently surfaces before data readiness in the list of stated barriers.

Governance Ownership

Ownership of data governance is inconsistent. It is clearer in regulated sectors that already have compliance functions but largely undefined elsewhere. Larger, more regulated organizations are building procedural controls, described in Manufacturing as data-flow tunnels comparable to Sarbanes-Oxley-style measures, while smaller manufacturers often have no named data owner at all. Portsmouth's chatbot illustrates the work involved: the city spent significant time normalizing its data before the tool could move into public-facing production, underscoring that data preparation, not the model, is usually the larger task. A self-scoring maturity assessment that includes a data and technology dimension would help organizations place themselves and see the readiness work ahead.

Chapter 5. Workforce & Talent: The Readiness Deficit

The Upskilling Gap Across Sectors

Every sector represented in the findings names a skills or training gap. Roles are being redefined more than eliminated: the creation of Chief AI Officer positions, legal associates are shifting from drafting toward reviewing and prompt-engineering, teachers are moving from content providers to facilitators and evaluators of AI output, realtors are becoming responsible for the accuracy of AI-generated descriptions, and finance staff are moving from manual processing toward advisory work. Financial Services is standing up new roles such as AI program managers and specialized AI compliance leads.

The Junior Talent Pipeline Concern

The sharpest workforce risk sits at the entry level, and it was raised independently in Legal, Technology, and the startup cohort. The concern is consistent: AI is absorbing the routine work through which junior professionals traditionally build expertise, so automating that work risks a generation of workers who never get the chance to become senior. Startups frame the answer as new apprenticeship models that teach judgment, validation, security, and systems design rather than relying on repetitive production tasks. Early replacement pressure, as opposed to augmentation, also appears at the entry level and in startups deferring hires such as a dedicated QA role.

What NH's Education System Is and Is Not Providing

The Education sector shows substantial activity but highly uneven delivery, and there is a stalemate between higher education and employers: schools are unclear on the curriculum AI-era employers need and employers are unable to give useful feedback while their own businesses are changing. The Education chapter's strongest request is a coordinated statewide framework for AI literacy, guidance, and professional development that can be adapted locally.

What Employers Actually Need

What many employers struggle to find is job-function-specific AI fluency. Meaning finance-, marketing-, and underwriting- specific AI capability rather than generic prompt engineering. While, in the startup cohort, skills like systems thinking and the ability to supervise AI-supported work is considered incredibly valuable.

Where training is working, it's often informal and peer-led: train-the-trainer and internal AI-champion models, hackathons and internal user groups are where many programs start and end.

Chapter 6. AI Safety & Responsible Use Across NH

Cross-sector Analysis: David Gerry

Two Disciplines, One Objective

AI safety and AI security are often used interchangeably. They should not be. AI safety asks whether a system behaves as intended and avoids unintended harm to users and the public. AI security asks whether that same system, and the data and infrastructure behind it, can withstand a deliberate adversary. One guards against failure while the other guards against attack. Organizations that focus on one over the other are exposed on the side they neglect.

AI safety and AI security are interdependent. A model that is unsafe is easier to exploit, and a model that is insecure cannot be trusted to stay safe. ISO/IEC 42001, the first international AI management system standard, treats data governance, adversarial testing, runtime monitoring, and oversight as one system rather than separate workstreams [4]. The controls overlap, and the audience is the same, from a single-location small business to a multinational organization.

The regulatory backdrop is unsettled, which makes a stable internal posture more valuable, not less. A federal executive order signed in December 2025 created an AI Litigation Task Force at the Department of Justice to challenge state AI laws in court [9], and a National Policy Framework released March 20, 2026 asked Congress to broadly preempt state AI laws, but it is not binding and creates no obligations on its own [9]. Earlier attempts to impose a moratorium on state regulation failed in Congress, and as of this writing no federal preemption bill has passed and no comprehensive federal AI statute exists [9]. Meanwhile, state laws are in force and shifting: California's Transparency in Frontier Artificial Intelligence Act took effect January 1, 2026, while Colorado repealed and replaced its 2024 AI Act in May 2026 with a narrower disclosure law that starts January 1, 2027 [10]. The lesson is not to wait for the law to settle. It is to anchor on a recognized voluntary framework that maps onto whatever emerges at the federal level.

Today, an anchor exists in the NIST AI Risk Management Framework, published January 2023, which is voluntary and built on four functions: Govern, Map, Measure, and Manage, plus seven characteristics of trustworthy AI [1]. Its principles line up with requirements taking shape in the EU and elsewhere, which lowers the cost of meeting several regimes at once [1]. Everything that follows treats NIST as the common reference and layers the security taxonomies on top.

One priority cuts across both safety and security and deserves to lead: protecting minors as they learn and use AI. It is also the area where the law gives states the clearest room to act.

Minors Are Where New Hampshire Should Lead First

Child safety is the one area of AI policy where state authority is least in doubt. The December 2025 federal order aimed at challenging state AI laws declined to touch laws protecting children, listing child safety among the categories not proposed for preemption [9]. That carve-out is the opening on which to build a defensible state policy. While most AI rules face years of preemption fights, a state that builds clear protections for minors can move now and set a standard others copy. Every

major jurisdiction is converging on the same point: the EU's July 2026 amendments to its AI Act added an outright ban on AI tools that generate nonconsensual intimate imagery or child sexual abuse material [8].

Minors meet AI in two ways and both need to be considered. First, they use it: as a study aid, a tutor, a search tool, and increasingly as a companion that simulates friendship. Second, they are targeted by it, most painfully through AI-generated sexual imagery. New Hampshire should address both.

On the targeting side, the legal floor is now federal. The TAKE IT DOWN Act, signed May 19, 2025, criminalizes the nonconsensual publication of intimate images, including AI deepfakes, with penalties of up to three years when the depicted person is a minor, and since May 19, 2026 covered platforms must remove reported images and known copies within 48 hours of a valid request, with the FTC enforcing [11]. The FTC runs a complaint portal for platforms that fail to act, and parents may file on behalf of minors [11]. The first convictions are already in, including an Ohio case involving AI-generated abuse material [11]. New Hampshire is not starting from zero. State deepfake statutes took effect January 1, 2025, and two more laws took effect January 1, 2026: HB 143 makes it child endangerment for a chatbot operator to use AI to encourage sexual conduct, substance use, self-harm, or violence by a child, with a private right of action, and SB 300 makes AI-generated intimate images of children a felony [14].

On the use side, other states have built out disclosure and crisis-referral rules that New Hampshire can improve on. California's companion chatbot law, effective January 1, 2026, requires operators to disclose that a user is interacting with AI, to remind known minors of that at least every three hours, to warn that companion bots may not be suitable for minors, and to maintain a crisis protocol that refers at-risk users to help [12]. New York's companion-model statute took effect November 5, 2025, and Utah and Maine have added their own disclosure rules [13]. Nearly every state now has a deepfake law on the books [11]. The pattern is clear and so are the gaps: the rules are fragmented, mostly reactive, and rarely connect classroom use to image-abuse response.

That fragmentation is New Hampshire's opening. The state can offer what the patchwork does not: one coherent, practical standard for minors and AI that a small district or a large platform can follow. There are four building blocks that provide this:

- 1. Safe classroom adoption.** Guidance that lets teachers use AI under supervision while setting plain rules on student data, disclosure, and acceptable use. Students who learn both how to use AI and where its risks lie gain an edge as the technology shifts.
- 2. Age-appropriate design for youth-facing tools.** Clear AI disclosure, no sexual content to minors, and crisis referral, building on HB 143 and drawn from the strongest parts of other states' laws rather than reinvented.
- 3. Fast response to image abuse.** Align school and law-enforcement reporting with the federal takedown process so a targeted student gets content removed in days, not months, and ideally within the 48-hour statutory clock [11].
- 4. Literacy for parents, educators, and students.** Practical material on spotting unsafe AI use, understanding data collection, and reporting harm.

None of this requires waiting on Congress or the courts. It uses what New Hampshire already calls its strengths: a tight-knit business community and direct access to policymakers, to build a standard from the ground up [15]. Do it well and New Hampshire stops adopting other states' AI rules and starts writing the ones they borrow.

AI Safety: Foundational Requirements

Start with the definition of what good AI safety looks like. NIST names seven characteristics of trustworthy AI: valid and reliable; safe; secure and resilient; accountable and transparent; explainable and interpretable; privacy-enhanced; and fair, with harmful bias managed [1].

The operating model is four functions [1]:

1. **Govern** sets accountability: who owns AI risk, what the policies are, and where the approval gates sit.
2. **Map** establishes context: where AI is used, who it affects, and what could go wrong.
3. **Measure** assesses and monitors those risks.
4. **Manage** acts on them in priority order. The functions repeat as conditions change rather than running once, and they apply whether the system is a traditional classifier, a third-party model, or a generative tool.

Generative AI adds risks the original framework did not fully cover. NIST's Generative AI Profile, published July 2024, defines 12 risk categories specific to or amplified by generative systems and lists more than 200 suggested actions, spanning fabricated output, harmful content, data leakage, and supply chain exposure [2]. Agentic systems that take actions on their own, call tools, and chain steps raise the stakes again, because an error can spread before a person can step in. NIST has follow-on work underway for agentic and critical-infrastructure uses [3].

Two foundational controls deserve emphasis for this audience. Human oversight means a person can review, override, or stop the system on consequential decisions. Transparency means people are told when they are dealing with AI and when content is AI-generated, which is becoming law in the EU and several states [8][12]. For a small business this can be a one-page policy and a labeling habit. For an enterprise it is a governed program. The principle is the same at both ends.

AI Safety: Technical Requirements

Policy fails without technical practice behind it. Four practices carry most of the weight.

1. **Data governance and lineage:** Know where training and fine-tuning data came from, whether you had the right to use it, and how it was transformed. Bad data is the root of most safety failures, from biased output to leaked personal information [4][5].
2. **Pre-deployment evaluation:** Test the system against the harms you care about before it ships, including red-teaming that actively tries to make it misbehave [5][7].
3. **Monitoring and drift detection:** Models degrade as the world changes. Log inputs and outputs, watch for performance drift, and set thresholds that trigger review [7].
4. **Documentation:** Keep model or system cards that record what a model does, its limits, and its tested behavior, plus impact assessments for higher-risk uses. The EU AI Act requires technical documentation, record-keeping, and human oversight for high-risk systems. The EU's Digital Omnibus, now in force, moved the main high-risk (Annex III) deadline from August 2, 2026 to December 2, 2027, but the work still carries a compliance payoff for anyone serving that market [8].

AI Security: Foundational Requirements

AI security starts with a wider view of the attack surface. A traditional application exposes its code and infrastructure, whereas an AI system also exposes its training data, its model weights, the

pipeline that builds it, the prompts that drive it, and any tools or plugins it can call. Each of these becomes a target whether the attacker is a person or another AI system.

The reference standard is secure by design. In November 2023, CISA and the UK's NCSC, joined by 21 other agencies including all G7 members, published Guidelines for Secure AI System Development, the first international agreement of its kind [7]. It organizes security across four stages: secure design, secure development, secure deployment, and secure operation and maintenance [7]. The idea is to own security outcomes for the customer and build them in from the start rather than bolt them on later [7].

Federal policy now treats AI security as a national security matter. Executive Order 14409, signed June 2, 2026, directs agencies to harden federal systems against AI-enabled attacks, to stand up an AI cybersecurity clearinghouse in voluntary coordination with industry, and to build a benchmarking framework for the most capable models, while expressly ruling out mandatory licensing or pre-clearance for AI development [9]. It creates no obligations for private organizations, but it signals where federal expectations are heading.

Three governance hooks to implement in your organization today:

1. **Keep an inventory** of every AI system in use, including the ones employees adopted without asking.
2. **Assign an owner** to each.
3. **Extend your incident response plan** to cover AI-specific failures, so a prompt-injection incident or a poisoned model has a defined path the same way a data breach does [7].

AI Security: Technical Requirements

The field now has shared threat references, and using them beats inventing non-standard models. There are three of these threat references that matter the most:

The OWASP Top 10 for Large Language Model Applications, updated for 2025, ranks the most common risks in deployed LLM systems. The leading entries are prompt injection, sensitive information disclosure, supply chain weaknesses, and data and model poisoning [5]. It is written for the people building and shipping these applications.

MITRE ATLAS is a knowledge base of real adversary tactics and techniques against AI systems, modeled on the widely used MITRE ATT&CK framework [6]. It is the reference for red teams and threat modeling. NIST's adversarial machine learning taxonomy supplies the formal vocabulary underneath both [5]. The controls follow from the threats.

- Validate and constrain inputs and outputs so untrusted text cannot hijack the model.
- Separate untrusted external content from instructions.
- Apply least privilege to the data and tools a model can reach.

Treat third-party models like any other supply-chain component: verify provenance, scan them, and track them with a machine-learning bill of materials [5]. Defend training data with validation, anomaly detection, and version control [5]. Log activity, and require human approval for high-impact actions taken by autonomous agents.

A practical note for smaller organizations: Most of these controls are not exotic. Many ATLAS mitigations map to security controls teams already run, so much of AI security is extending existing practice rather than starting over [6].

AI also sits on both sides of the fight. Attackers use it to scale phishing, write malware, and find weaknesses faster while defenders use it to triage alerts and spot anomalies. The defining problem is speed. Whoever moves faster, attacker or defender, holds the edge.

Where Safety and Security Converge

The two disciplines share more than they differ and one AI inventory can serve both purposes. Red-teaming surfaces safety failures and security holes, and monitoring catches drift and intrusion in the same logs.

The frameworks are built to stack. Use the NIST AI RMF as the risk operating model, ISO/IEC 27001 for information security, and ISO/IEC 42001 for AI management, which is certifiable and signals diligence to enterprise buyers and regulators [4]. NIST has a draft profile connecting its widely used Cybersecurity Framework to the AI RMF, with companion SP 800-53 control overlays for securing AI systems still in draft [3]. An organization already running a cybersecurity program does not need a parallel one for AI.

Practical Path for New Hampshire Organizations

If every New Hampshire organization speaks the language of the four NIST functions, a ten-person firm and a hospital network can hold the same conversation at different depths by using the same underlying baseline. Having a shared vocabulary is worth more than any single control will provide alone.

Right-size the effort:

- **Small business:** an AI inventory, an acceptable-use policy, a short list of vendor questions, AI labeling, and basic logging.
- **Mid-size:** add red-teaming, model or system cards, and an incident playbook that names AI failures.
- **Enterprise:** a full AI management system, ISO/IEC 42001 certification where buyers expect it, and dedicated AI security testing.

Watch three regulatory tracks:

- **State:** California is live. Colorado's replacement law starts in 2027 and is far narrower than the original, which is itself a lesson in how quickly state frameworks can be unwound. New Hampshire's own 2026 session produced no new framework: a comprehensive bill (HB 1725) was defeated in February and a narrower oversight bill (SB 657) died after the House and Senate could not agree on a version [16].
- **Federal:** the December 2025 preemption order, the March 2026 policy framework, and the June 2026 order on AI security set the direction, but no federal AI statute exists and no preemption bill has passed [9].
- **International:** the EU AI Act's high-risk obligations now apply from December 2, 2027 under the Digital Omnibus, which took effect in July 2026, and they still matter for anyone serving that market [8].

Treat AI Like a New Hire

Most organizations are asking the same three questions: where to start, how to protect the business and the people in it, and how to get the value out of AI without creating new exposure. Three rules provide a good starting point for both professional and personal use cases:

- **First**, treat any new tool like a new employee. Trust must be earned, even for software.
- **Second**, require human approval for business-critical decisions. Before a bill is paid, a person reviews the amount, who is being paid, why they are being paid, and any other information available.
- **Third**, in the same way that you would not give a new employee unfettered access until trust is built, minimize what AI tooling can reach. As an individual, this means not providing access to banking, financial, username, or password information unless it is required to perform a task. As a business, it means limiting access to company financial data, banking information, and client information. And, always, review access regularly to determine whether what was granted is still required.

The immediate opportunity for New Hampshire is to lead on protecting minors. It is the area of clearest state authority, the state already has two laws on the books to build from, and it is the place where a practical New Hampshire standard would travel furthest. The state's advantages, a tight community and short lines to policymakers, are built for this kind of work [15]. New Hampshire does not have to be the biggest player in AI to be the one others copy, but it does have to be the clearest.

Chapter 7. Ethics & AI Governance: Organizations Can't Wait for Legislation

Why Organizations Need Their Own Frameworks Now

The regulatory picture makes the case on its own. No comprehensive New Hampshire AI statute exists, most private-sector regulation bills died in the 2026 session, and federal preemption efforts and constant change add uncertainty about which rules will hold. In that vacuum, voluntary frameworks are already the de facto driver of governance across the state. Which means organizations that wait for legislation are not pausing so much as operating ungoverned. Developing an internal ethical framework now is the realistic path, not a premature one.

Who Could Be Harmed

The people at risk of harm differ by sector, but the shape of the tension repeats: patients who have not consented to, or are unaware of AI's role in their care; defendants exposed to biased predictive-policing and facial recognition tools; job applicants screened by employment algorithms that can proxy for age or other protected characteristics; students navigating privacy, integrity, and equity questions; and artists whose voice and likeness can be cloned without consent. In the Finance sector, the risk of bias in underwriting and fraud detection that could disadvantage underserved populations, and in the Manufacturing sector, we observe the risk of leaving quality dispositions or engineering-change approvals to AI agents without subject-matter-expert oversight.

Automation Versus Professional Judgment

The tension between automation and professional judgment is sharpest in Legal and Healthcare, where unreviewed output carries malpractice, clinical, or civil-rights consequences. The same holds true in the Finance sector where there must be a human-in-the-loop requirement for any decision that directly affects a customer's financial standing, a pattern other high-stakes sectors are converging toward.

Disclosure as the Live, Unresolved Question

Disclosure is the ethical question sectors are wrestling with today. There's a growing consensus that people should be told when they are interacting with AI, and when AI contributes to a significant deviation, and that they should have a clear path to a human. Healthcare has a consent gap, Real Estate faces open liability questions over AI-enhanced marketing photos.

Governance for Small and Medium-Sized Organizations

Smaller organizations that lack the sophistication of large firms need practical, lightweight governance rather than enterprise-scale bureaucracy. The startup cohort's minimum controls and the small-business request for ready-made New Hampshire policy templates point the same way, as does the proposed short template distinguishing acceptable, high-risk, and prohibited uses, modeled on the EU AI Act approach.

Chapter 8. ROI, Cost, and the Environmental Footprint of AI

Make Measurable Business Impact the Standard

The sharpest realization in the findings is that enthusiasm for AI is not the same as proof that it pays. Organizations across the state describe pilots that save time and win user approval, yet the financial result often does not move. A pilot can save ten hours a week, users can like it, but finance can still see nothing, because what is missing is the conversion of operational improvement into economic value. The recommendation that follows is direct: treat every AI initiative as a business investment, and define the business outcome and its measurable success criteria before funding it, not after.

A Defensible ROI Discipline

A simple, sequential discipline makes AI initiative returns measurable and defensible. Each step produces the evidence the next one needs, and the order matters, because an organization that waits until after implementation to define ROI has already lost the baseline.

1. **Define the business objective.** Name the process being improved, the metric that should change, who owns the outcome, and what happens if nothing is done. If the objective is simply to implement AI, the initiative does not yet have a business case, and the first question is whether AI is even the right tool or whether workflow, rules, or existing systems would solve the problem better.
2. **Establish a current-state baseline.** Improvement cannot be measured without the full cost of today's process, including direct labor, errors and rework, delays and cycle time, risk and compliance exposure, and the opportunity cost of capacity that is tied up.
3. **Describe the desired future state.** State the observable operational change that would show the initiative worked, such as a cycle time that falls from five days to two, before choosing any technology.
4. **Translate improvements into financial value.** Connect the workflow change to where value actually appears in the business: cost reduction, new or protected revenue, reduced risk, better service, or higher throughput.
5. **Measure and validate over time.** Track realized results against both the baseline and the full investment, and revisit the result on a defined schedule rather than declaring victory at launch.

Time Saved Is Not ROI

The most important distinction on this topic is that time savings alone are not a return. Saving hours creates capacity; it does not by itself create value. Capacity becomes ROI only when it is deliberately converted, either by removing cost through workforce optimization or by redeploying the freed capacity into measurable revenue, service, compliance, quality, or throughput. The Task

Force's own findings make the point concrete: healthcare systems report a reduction in charting time, manufacturers report large reductions in translation cost, and public-sector staff report meaningful time saved, but each of those is capacity created, not yet value realized. The test is what the organization does with the hours, and whether the resulting gain can be tracked against the total cost of the tool.

The Full Cost of AI

ROI falls apart when the denominator is incomplete, and the license fee or subscription is rarely the total cost of an initiative. A complete accounting includes software and licensing; model, token, and consumption costs; implementation and integration; data preparation; internal business and technical staff time; training and adoption; change management; security and governance; and ongoing support and improvement. The affordability gap for the smallest organizations is treated in Chapter 9. The point here is that the ongoing cost of the compute and all associated variables behind AI is especially easy to underestimate, which is why returns should be stated plainly: economic benefit minus total investment, divided by total investment, alongside a payback period, and counting only benefits that can be quantified, tracked, and validated.

Measure Adoption and Impact Separately

Adoption and impact are different measurements, and confusing them is how initiatives stall in permanent pilot status. Adoption indicators, such as active users, frequency of use, and task completion, show that a tool fits a workflow. Operational indicators, such as cycle time, throughput, error rate, and backlog, show that the work changed. Financial indicators, such as cost reduction, cost avoidance, margin improvement, and incremental revenue, show where the value landed. High adoption is a strong leading signal, but it is not a return, and a report of hours saved is the beginning of the ROI question, not the answer.

Turn Measurement into a Funding Decision

Measurement should change what leadership funds next. Every AI initiative deserves a go / no-go decision instead of indefinite pilot status: fund it when there is clear value, an owner, a baseline, and a measurement path; optimize it when the value exists but adoption, cost, or workflow needs work; delay it when readiness or controls are insufficient; or stop it when there is no measurable outcome or credible path to one. A short executive scorecard makes this routine, asking of every investment what business outcome is being funded, who owns it, what the baseline is, which operational metric should change, how that change creates value, what the total investment is, when the result will be reviewed, and what would trigger scaling, redesign, or stopping. Where a question cannot be answered, the right move is to mark it unknown and assign an owner, not smooth it over.

Environmental Footprint

The environmental dimension appears as an ethical tension rather than a measured analysis today. The Public Sector chapter raises the energy demands of AI as a real tension for organizations that also carry climate, sustainability, and stewardship responsibilities, the cross-sector matrix records the environmental cost of data centers as a sharp risk in Real Estate, and one Technology respondent named the energy cost of AI as a source of internal pushback.

The report does not yet contain a quantified environmental analysis, a gap noted here for the next report cycle.

Chapter 9. Access & Equity: Who Gets Left Behind

The Small Business and Rural Divide

Access concerns are emerging most clearly where price and capacity meet. Enterprise pricing does not scale to small municipalities or small businesses. Small businesses also report resource constraints and analysis paralysis, wanting bite-sized guidance and SBA-style consulting networks. Smaller and more rural municipalities carry the added burden of little or no dedicated data, cybersecurity, legal, or IT staff, so the same organizations with the least capacity are the ones asked to evaluate and govern AI.

Resource Disparities in Education

Education names equity as a central concern in its own right. Students at institutions with strong AI programs may graduate with substantially different skills and opportunities than students whose schools restrict access, and the gap between paid and free tools widens those disparities further. K-12 privacy requirements and funding constraints compound the divide, since the institutions with the least capacity face the longest vendor-vetting delays.

Geographic and Demographic Dimensions

The uneven adoption documented throughout the report has a distributional edge: employees who learn to work with AI gain productivity advantages over colleagues who lack access, training, or confidence, a dynamic named explicitly in the Public Sector findings. Representation is itself an equity issue. Retail and Hospitality and Real Estate are active adopters with no guaranteed seat in the policy conversation shaping the rules they will operate under, and the cross-sector matrix flags Community Advocacy's need for stipends and inclusion so that vulnerable populations are represented where decisions are made.

The Lever Most Likely to Close the Gap

The access findings point to a practical lever the state can influence: cooperative purchasing and government pricing, so that cost does not exclude smaller communities and organizations from tools and training. For the public sector specifically, a standardized procurement pathway would remove the mechanics that municipal purchasing rules impose. The equity value of both is that they reach the organizations least able to solve the problem alone.

Chapter 10. Policy & Regulatory Landscape

What NH Organizations Are Asking For

The asks are broadly consistent across sectors, though not identical. The one that recurs most widely is practical, plain-language guidance and shared templates: a self-scoring maturity assessment and a short template distinguishing acceptable, high-risk, and prohibited uses, modeled on the EU AI Act and proposed as a three-page foundational bill, requested in some form by Government, Small Business, Manufacturing, and Education. The Public Sector's own top ask is a standardized AI procurement pathway, since municipal purchasing rules were not built for how AI is bought, along with clear guidance on how public-records and Right-to-Know requirements apply. Finance asks specifically for regulatory sandboxes or safe-harbor frameworks so it can move beyond low-risk pilots, and Manufacturing asks for a practical roadmap and a white paper for legislators.

Legislative Developments in New Hampshire

New Hampshire currently regulates AI through a narrow patchwork rather than a comprehensive statute: HB 1688 restricting state-agency AI use, HB 1432 criminalizing fraudulent deepfakes, and SB 255's data-privacy provisions. Most private-sector regulation bills died in the 2026 session. Outside those measures, state-level AI regulation is limited to narrow issues such as deepfakes and political-advertising disclosure, there is no comprehensive statewide framework, and there is no state-wide law on data centers, leaving it up to individual towns to address.

Federal Policy Context and Its NH Implications

Federal action is adding uncertainty rather than resolving it. Federal preemption efforts, including Executive Order 14365 and the DOJ AI Litigation Task Force, create open questions about which rules will ultimately hold, which complicates any state-level planning and is one reason organizations are advised not to wait for legislation before governing themselves.

Where Policy Uncertainty Is Slowing Adoption

Policy uncertainty slowing decisions is among the strongest-holding February signals. It is most visible in Finance, where a wait-and-see culture limits AI to low-risk, non-customer-facing use, and in the Public Sector, where organizations pause on the basic question of whether they are permitted to use AI on given information before they ever reach data or tooling questions. Reducing that uncertainty, through clear guidance, a shared policy baseline, records guidance, and, for the public sector, a workable procurement pathway, is where policy action would most directly unlock adoption.

Part Two: Sector Findings

Chapters 11–19

Chapter 11. Technology & Software

Sector Lead: Katie Destramp

Current State

In the Technology sector, adoption is uneven, with a wide gap between the leading edge and the “typical” organization. Participants debated where the gap falls. One view holds that smaller, nimble organizations move faster than larger, more traditional ones constrained by culture and process. A later discussion saw it the other way: adoption is highest at very large organizations and at startups, with small-to-midsize businesses lagging both because they lack dedicated people or budget to focus on AI while still running the business. Both accounts agree that organizational resourcing, not sub-industry, drives the gap; they disagree only on which end of the size spectrum leads.

The sector's general posture sees AI as more opportunity than threat. Tech leaders across conversations were more animated about what AI enables than fearful of it, even as individual participants flagged real concerns about cybersecurity and about output quality and trust. A 20-response survey of NH Tech Alliance members (mostly small, AI-native or AI-adjacent firms, gathered August through December 2025) illustrates this range: answers ran from a single company just starting to seven describing themselves as fully integrated, with the largest cluster, 8 of 20, actively implementing AI in select areas. That sample skews toward companies already engaged with AI, so it likely reads more mature than the sector as a whole, but it confirms the core point. Taken together, the roundtables and survey place the leading edge at roughly Systemic and the typical organization near Active, a two-stage gap held with moderate confidence.

Key Use Cases

- Applications named in the discussions include deployment automation; internal process automation for customer migration and onboarding; vibe coding and AI agents building internal CRM tooling and automating marketing outreach in-house; AWS-based automation of media subtitle changes; and automation of manual retail processes.
- Others include automated customer-support tooling built to reduce future headcount; AI to improve network performance and proactively flag impairments; consulting productivity gains that let a few senior engineers cover output that used to take more engineering resources.
- Retrieval-augmented-generation (RAG) chatbots for sales and support; custom computer-vision models for sports analytics; a custom-trained cybersecurity threat-detection model; and agentic development stacks (Cursor, GitHub Copilot, workflow platforms such as n8n).
- Individual participants also described an AI-driven IoT platform deployed for a Japanese client, an audit-trail product for financial-services clients, and an in-house tool that keeps employees from typing sensitive customer configuration data into external models.

Across the discussions, AI is touching nearly every desk-level job, and unreviewed AI content reaching customers or the public was repeatedly named as a real, recognized risk.

Primary Barriers

The barrier repeated most consistently, across three separate discussions, is business-model disruption to consulting and services work. Participants described direct conflict between AI efficiency and the billable-hour model and an uncomfortable shift toward outcome-based, deliverable pricing. One firm is considering productizing its own knowledge to guard against clients using AI instead of a consultant. One roundtable sharpened this into a threat: large firms cutting workforce on AI efficiency put rate pressure on smaller NH firms, and clients building their own AI solutions instead of hiring were compared to the shadow-IT phase of cloud adoption.

A second theme, across two discussions, is that AI governance and agent-management practice is still immature. Even sophisticated engineering organizations lack a standardized model for what AI agents should be allowed to access. One participant described the regulatory backdrop as the wild west, and separately raised the risk of tools or agents that only one person fully understands, a liability once that person leaves.

A third theme is budget and capital access. In the member survey, budget constraints were the single most commonly cited barrier, ahead of ROI uncertainty and not knowing where to start, with several respondents asking directly for funding or low-cost tooling.

Workforce Implications

The workforce concern with the broadest support is risk to junior and entry-level talent, raised independently in two discussions. Historically, developers have to move through a junior path to become senior, and AI hype suggesting less need for developers threatens that pipeline. A related point: there is a real difference between an experienced developer using AI as an efficiency tool and a novice using it to build an application they cannot troubleshoot, with the advice being to pair any AI-generated work with an experienced mentor.

A second theme, from two discussions, is that curriculum and training have not kept pace. Participants described a stalemate between higher education and employers: schools are unclear on what curriculum AI-era employers need, employers are not able to provide useful feedback, and new graduates struggle as a result. A broad call for structured workforce training ran across three discussions and is covered below under what the sector needs most.

Individual participants added three further observations - work is shifting toward directing and reviewing AI output instead of doing it directly; a distinct group of employees already skilled in their jobs struggle to incorporate AI into work they do well; and AI lets senior staff hand routine work to AI to focus on complex work.

Data Readiness

Data readiness was not a major theme in this sector's discussions. When it came up, it leaned toward data security and privacy, not data quality or access: organizations restricting external AI tools over PII concerns and building internal alternatives, with the member survey separately naming data security and privacy as live concerns. For technology and software firms, data is a smaller constraint than it appears to be in other sectors.

AI Safety and Risk

The one governance finding supported across several discussions is that formal AI risk management is immature, inconsistent, and not standardized. This topic showed the widest range of any in the chapter: some smaller AI-focused firms have written policies, and one NH company's legal team built a formal policy for AI-assisted coding with mandatory peer review, while at the other end a training session told attendees to just start using ChatGPT with no mention of policy.

One company sells an audit-trail product for AI-assisted decisions; a professional body (PMI) was cited as building AI governance into project-management education; one participant suggested governance could become mandatory through insurance requirements before any law requires it, similar to cyber insurance; and AI slop, meaning low-quality AI output that goes out unreviewed, was raised as a specific quality risk, echoed in a second discussion. The sector places the typical NH technology and software organization at “still figuring it out”.

Ethics

The person most at risk of harm is the entry-level worker whose skills could be replaced before they have a chance to build them, connecting to the junior-developer pipeline concern raised independently in two discussions. A related thread, from two discussions is accountability: AI output quality depends heavily on whether the user already has the expertise to judge it. One discussion made the point about unreviewed AI content; another made the same point about code. Together they raise a sector-specific question: when a client receives AI-assisted work, do they know how much was AI-generated, and who is accountable if it is wrong.

Two individual participants added related points: one connected AI ethics work to workforce impact specifically, and one raised the environmental and energy cost of AI as a source of internal pushback.

What This Sector Needs Most

Structured workforce training is the single best-supported need in the chapter, raised independently across all three discussions, which described people blindly using AI without understanding it and widespread confusion between terms such as AI, generative AI, and machine learning. Two secondary needs each had support from two sources: peer visibility, meaning seeing how other NH companies actually use AI, and funding and capital access, which was the dominant ask in the member survey. One participant suggested routing state AI support through Chambers of Commerce and the Business and Industry Association instead of building new programs, an option worth considering.

Chapter 12. Financial Services & Insurance

Sector Leads: Jim Kisch & Scott Weller

Current State

Financial institutions fall across the full adoption scale, but the leading edge tends to lean “Systemic” and the typical institution leans Active. The gap between leading and typical institutions is explained largely by size as the largest institutions have more resources to invest in AI tools, adoption and governance.

Key Use Cases

Institutions are in production or piloting a common set of use cases.

- In internal communications and policy development (Human Resources, Marketing, Administration), AI drafts routine correspondence, internal policies, and job descriptions to ensure consistency and eliminate repetitive writing.
- In regulatory compliance and audit summarization and ad and policy review (Compliance, Internal Audit, Legal), AI accelerates synthesis of complex regulatory documents and internal audit findings for faster review cycles.
- In fraud and security pattern analysis (Fraud and Risk Management, IT Security), AI replaces manual review with real-time detection of suspicious login patterns and transaction behaviors. In commercial loan underwriting and data extraction (Commercial Lending, Credit Administration), AI reduces the labor of manual financial spreading and policy verification, speeding underwriting while maintaining compliance.

Primary Barriers

The most significant barrier is a governance-agility gap rather than simply security or resource constraints. While leaders often cite regulatory uncertainty as the primary blocker, the underlying friction is the attempt to govern highly dynamic AI models using legacy, deterministic model-risk-management frameworks. That mismatch creates a wait-and-see culture, with institutions limiting AI to low-risk, non-customer-facing use cases because current compliance infrastructure cannot provide the algorithmic reliability required for high-stakes decisions. If sector-specific regulatory sandboxes or clear safe-harbor frameworks existed, this paralysis would ease, letting institutions move from internal efficiency pilots to true strategic integration. The sector's high-stakes regulatory landscape means that, absent a standardized way to prove algorithmic safety, adoption remains tactical and cautious rather than transformative.

Workforce Implications

The sector is creating new career paths such as AI program managers and specialized AI compliance leads, and is actively seeking people who combine technical AI proficiency with a deep understanding of banking regulation, viewing the talent gap as an invitation to upskill the existing workforce. Rather than widespread job loss, leadership sees AI as a collaborative partner

that automates repetitive, high-volume tasks and transforms entry-level administrative roles toward higher-value work such as strategic customer advisory and complex problem solving.

Institutions leverage internal expertise through train-the-trainer approaches in which internal AI champions mentor colleagues, and implement two-tier strategies offering foundational AI awareness for all staff alongside targeted instruction for operational roles.

Data Readiness

AI is technically viable in the sector today, but its efficacy is constrained less by the models than by the maturity of data architecture. Institutions possess rigorous data governance for security compliance, but that strength often manifests as rigid data siloing across legacy systems. The blocker is therefore not a lack of data but a lack of unified, accessible readiness. Without cleaned, normalized, and integrated foundations, which many institutions are still building through migrations to cloud-based data lakes, AI models struggle to reach the predictability required for complex banking operations. Data-governance ownership remains fragmented rather than centralized, especially at smaller institutions, and the transition from disconnected repositories to unified foundations is slow.

AI Safety and Risk

On risk, financial institutions are generally ahead of the curve, with the caveat that some are holding off on expanding AI capabilities until they receive more robust guidance from regulators. The sector integrates formal governance directly into existing information-security and compliance frameworks, establishing written policies that address data privacy, security vulnerabilities, algorithmic bias, and model hallucinations. Oversight is provided by a mix of risk subcommittees, internal workgroups, and IT departments, with priority on vendor risk management and employee-conduct training. This rigorous oversight is largely voluntary, reflecting both high internal standards and anticipation of evolving regulatory guidance.

Ethics

A prominent tension lies in balancing algorithmic efficiency with human-led professional judgment, especially in credit underwriting and fraud detection. Automated systems offer speed but risk embedding historical biases that could unfairly deny services to underserved populations. To mitigate this, institutions maintain a human-in-the-loop requirement for any decision that directly affects a customer's financial standing. There is also growing consensus that customers should be explicitly informed when AI contributes to significant decisions affecting their financial lives, honoring the customer's right to understand the factors behind a loan denial or account flag and protecting both the customer and the institution.

Impact Story

At a regional New Hampshire financial institution, AI anomaly-detection systems flagged an uncharacteristic, large wire-transfer attempt from an elderly account holder that deviated sharply from the customer's established patterns. The system triggered an immediate manual review, allowing staff to intercept the transaction and discover the client was being targeted by a sophisticated imposter scam that threatened their life savings. The bank prevented a life-altering financial crisis, illustrating that the value of AI in the sector is not merely efficiency but the preservation of dignity and security for vulnerable community members.

What This Sector Needs Most

While New Hampshire banks are moving at different speeds, the roundtables surfaced several practical priorities. Connect AI to the broader strategic plan, identifying where it can advance existing priorities rather than treating it as a standalone initiative. Establish clear governance covering acceptable use, data protection, security, oversight, and employee responsibilities, supported by acknowledgment and training. Invest in the workforce with both foundational awareness and practical training on approved tools. And learn from one another by continuing to share case studies and lessons across the industry. On the February signals, the sector reports that all still hold, with progress being made on nearly all of them.

Chapter 13. Healthcare & Life Sciences

Sector Lead: Maggie Allard

Current State

AI adoption across New Hampshire's healthcare ecosystem varies considerably, from organizations still considering implementation within the next one to two years to larger health systems running a broad array of models across multiple departments. Adoption is more advanced within larger health systems and less mature among organizations that primarily serve behavioral health and Medicaid populations. Rural health systems may offer particularly promising environments for piloting and incubating AI solutions, because they are smaller and more nimble with fewer staff and fewer operational hurdles. For most organizations, ambient AI for clinical documentation is the common entry point, a relatively low-risk, high-value starting place for broader adoption. The sector spans Awareness to Systemic with the leading edge at Systemic and the typical organization at Active.

Key Use Cases

Healthcare organizations are increasingly leveraging AI to transform clinical workflows, ranging from ambient listening tools to predictive algorithms that capture data invisible to standard human observation.

- Ambient documentation is the most widely deployed clinical use case, with one health system reporting a reduction in charting time; in behavioral health a general large language model did not work well and a more conversational model was used instead.
- AI is also advancing documentation automation by summarizing lengthy patient histories for clinicians and enhancing diagnostic accuracy through specialized applications, such as computer-aided models for mammography, real-time polyp detection during colonoscopies, and cancer-scanning tools on CT images.
- Health systems are piloting large language model (LLM) APIs integrated into electronic medical records to run patient surveys such as the PHQ-9 and improve documentation intake while carefully navigating privacy concerns, including sensitive 42 CFR Part 2 consents in behavioral health.
- Within healthcare operations, AI optimizes efficiency and patient safety by using mobile apps to match patients with medications at the point of care, tracking IV-line metrics to prevent hospital-acquired infections, and automating medical coding. It streamlines administrative and clinical workflows by triaging portal messages, detecting cybersecurity threats, providing emergency-department order prompts, and expediting radiological documentation.
- In financial operations, AI streamlines revenue-cycle management and prevents claim denials by scouring clinical documentation for missing or inappropriate codes, and it improves claims processing by prioritizing high-revenue claims and ensuring accuracy.
- Beyond clinical and financial settings, AI handles first-line threat detection, augments cyber-defense surveillance, and performs environmental scans of third-party partners, recalls, and legislation to support high reliability organization (HRO) goals.

- In value-based care, AI automates the tracking and matching of CMS and LeapFrog reporting requirements under human oversight. Within research, AI algorithms interpret complex data and assist with advanced genomics.

Primary Barriers

Adopting AI in healthcare presents a complex set of challenges spanning security, governance, finance, and workforce dynamics. Data security and compliance remain primary concerns: depending on exposure to protected health information (PHI), organizations require stringent vendor certifications such as HITRUST, SOC 2 Type II, or ISO 27001, alongside Business Associate Agreements and technical evaluations. It is exceptionally difficult to oversee data absorbed by dozens to hundreds of discrete AI applications, with third-party business partners representing significant cyber risk. Robust technical guardrails are essential to protect sensitive data while enabling medical advancement, and organizations must battle cyber threats amid cost and resource constraints.

Workforce education, training, and trust also present significant hurdles. Staff require dedicated training on how to interact with AI tools, since poor prompting and a lack of AI etiquette create noise and unnecessary administrative work. Employees must apply critical thinking to verify AI output, and organizations need specialized in-house expertise depending on whether they are building bespoke models, deploying vendor solutions, or both. Employee distrust, driven by ethical concerns or fears of job displacement, requires targeted work to frame AI as a tool that enhances human capability rather than replaces it.

Purchasing and financial sustainability add another layer. Deploying AI tools involves substantial expense, which can reach millions of dollars in larger systems, requiring rigorous financial analysis to confirm a sustainable return on investment. Organizations struggle with selecting the right agents, understanding the true value of an AI platform, formalizing vendor relationships, and knowing when to move from pilot to partnership. Specialized barriers complicate adoption: rural hospitals face disproportionate infrastructure challenges, while behavioral health integration may be delayed by the deployment of a standardized electronic medical record.

Workforce Implications

Across New Hampshire's healthcare organizations, AI appears to be filling workforce gaps and enabling employees to do more with existing resources rather than replacing staff, with no layoffs reported. A prior mindset that AI would allow doing more with fewer people did not translate into workforce reduction in healthcare; it translated into doing more with the same resources. Older employees show notable enthusiasm, because they have a baseline for how time-consuming these tasks were without AI, and organizations report opportunities to improve operational efficiency, reduce administrative burden, and decrease burnout.

Realizing continued benefits will require upskilling and realignment so that frontline staff, clinicians, operational leaders, and administrative teams understand how to use, implement, and support AI effectively. Organizations will also need to evaluate the financial impact of these investments, determining whether upskilled employees generate sufficient value and whether additional expertise or staffing is needed to oversee increasingly complex agentic AI platforms, including maintenance, oversight, optimization, and administration.

Data Readiness

Data readiness remains a significant consideration. Legacy systems can hamper readiness, particularly when organizations must connect newer AI applications with longstanding infrastructure and disparate data sources. Even as adoption surges, organizations must know exactly where PHI is located, a regulatory requirement that has become harder as EHRs, vendors, and third parties integrate AI across products with data ending up in unexpected places.

To secure healthcare data and understand where it resides, how it moves, and who can access it, organizations are rethinking the roles of IT, quality, data analytics, and decision support, and some are connecting data across APIs, discrete applications, and EHR enhancements to create an intelligence layer. Historically these capabilities lived within IT alongside legacy systems, but there appears to be a shift toward business intelligence and decision support aligning more closely with quality and data analytics, producing new organizational and staffing structures. Ultimately, data governance builds on the governance established with EHR implementation, often decades prior, and must keep maturing as AI becomes more deeply embedded, so that organizations maintain visibility into where data resides, keep it protected, generate actionable intelligence, and align expertise and staffing to keep pace.

AI Safety and Risk

Healthcare organizations are building governance instead of operating ad hoc, though maturity is still emerging, and the sector places itself at still figuring it out. Leading practices are creating robust governance frameworks and establishing oversight for the tens to hundreds of discrete AI applications currently in use. AI tools are increasingly covered by both a Business Associate Agreement and carefully crafted vendor agreements that address data and PHI along with deployment, support, and use. Health systems may require transparency into LLMs and vendor algorithms to reduce bias. They also seek disclosure on when AI is in use and may require that all data resides in the United States. Depending on the level of PHI exposure, organizations require vendors to be HITRUST, SOC 2 Type II, or ISO 27001 certified.

Health systems may operate a responsible adoption of AI policy where a multidisciplinary team convenes as needed to evaluate a vendor, using a graded framework and defined procedures. For healthcare staff, acceptable-use documents may be created and regularly updated, especially for generative AI. Additionally, organizations have established AI governance committees which meet on set cadence and are supported by executive meetings and wider town halls. These convenings allow a process by which issues are surfaced, use cases identified, value understood, and AI operations continually reviewed. Greater risks include over-reliance on AI output, which end users must validate, and insufficient insight into how vendors build and test their models, a concern sharpened by vendors activating AI inside their products before notifying healthcare customers.

Ethics

The sector's ethical priorities are patient privacy under HIPAA, data security against cyber threats, and patient safety. The consistent principle is a human, and specifically a physician, in the loop for any AI use that requires judgment about a patient, so that AI supports rather than replaces healthcare staff. Because there are no specific federal regulations governing the ethical use of AI in hospitals, organizations are setting their own guardrails: patients must consent to AI scribe use, and vendors are expected to be transparent and to permit review of how they use LLMs, AI, and machine learning in their products. The emerging model is trust-but-verify governance, in which

every generative or diagnostic model undergoes careful review and rigorous local clinical validation, with humans and physicians in the loop, to ensure patient safety, data privacy, and algorithmic equity before it scales across a system.

Impact Story

Practical AI in action: How a Rural Health System is Leading the Way. North Country Healthcare (NCH) is a strong example of a rural New Hampshire health system approaching AI adoption practically to improve clinical workflows, address workforce challenges, strengthen operations, and support patient care. Across three critical-access hospitals serving approximately 32,000 patients over a large geographic area, NCH has implemented or piloted a range of AI applications, beginning with ambient voice technology for clinical departments that includes a specialized model for behavioral health. A pilot program is underway in behavioral health which utilizes AI for form intake, documentation prompting, and survey administration initially with the PHQ-9. In the rehab department, AI is being piloted to summarize lengthy medical records so clinicians can quickly understand a patient's history. In radiology, computer-aided detection analyzes mammography images to identify potential irregularities early. In the hospital, barcode medication administration has reduced patient risk, reaching a 99.5% compliance rate over the past two years. The organization is now exploring LLM integration into its EHR, with plans to expand into patient engagement and tailored discharge instructions.

Beyond direct patient care, NCH applies AI to operations, cybersecurity, and regulatory monitoring. AI-powered prior-authorization workflows will be able to automatically identify requirements based on CPT codes, reducing administrative burden, and the organization is exploring automation of both prior authorizations and claims processing. NCH has also built a custom AI monitor that produces daily reports and monthly summaries scanning for cybersecurity risks and other threats tied to vendors, third parties, and recalls. It even monitors legislation that may affect food, pharmaceutical, and biomedical product regulation.

At the same time, NCH is building the data and quality infrastructure to support these initiatives. All 1,200 NCH employees are required to be Lean Six Sigma trained, and data analytics and decision-support functions have moved from IT to the quality department, allowing previously disparate EMR, lab, and staffing data to be bundled into a business-intelligence layer that surfaces staffing, operations, and patient-population insights. Underlying these efforts is a clear focus on responsible AI adoption: NCH has established a responsible AI adoption policy with a multidisciplinary review process and requirements for transparency into large language models, particularly to address potential bias, and it is seeking greater transparency from vendors about where and how AI is used and how data are stored. Taken together, NCH's journey illustrates a practical approach to AI adoption in a critical-access, rural environment: identify clinical and operational needs, pilot partner technologies within existing systems, build supportive data and quality infrastructure, and strengthen governance as adoption expands.

What This Sector Needs Most

To advance safe and sustainable AI, organizations require robust governance frameworks, help with vendor selection, expert staffing, and transparency legislation. They insist that humans remain in the loop for oversight and accountability on both patient-facing and administrative decisions, including diagnosis, treatment, prior authorization, utilization review, and cyber risk, and they request guidance on governance and risk management. As vendor solicitations increase, organizations insist on HITRUST, SOC 2 Type II, and ISO 27001 certification.

Practical hurdles are driven primarily by capital constraints, staff training needs, and a lack of specialized expertise. Healthcare organizations, especially those serving rural and Medicaid populations, seek support to fund technology acquisition, train their workforce, establish governance, and select vendors, and in the coming years will need specialized expertise to navigate integration, including maintenance, monitoring, optimization, and administration. Beyond stringent contracts, many also request transparency legislation requiring vendors to disclose AI use, data-collection practices, and storage locations to safeguard patient, workforce, and financial interests.

Checked against the February signals, the sector affirms most: AI is in use but adoption is uneven; AI is efficiency rather than replacement, with a human making the final judgment where patients are involved; workforce readiness lags tool availability; demand for practical guidance is strong; risk awareness is high while governance maturity varies; access concerns are emerging for rural and lower-income populations; and New Hampshire is at an inflection point, engaging AI at different levels without clear guidance or strong governance. It qualifies two signals, holding that data readiness is one bottleneck rather than the only one and that heavy vendor reliance is only somewhat true given some distrust of vendors, and it dissents on one, reporting that policy uncertainty slowing decisions was not heard in this sector.

Chapter 14. Manufacturing, Engineering & Defense

Sector Lead: Brian Stumm

Current State

Adoption is concentrated in pockets rather than spread evenly. At the leading edge, larger, better-resourced manufacturers have moved AI into real production workflows: structured authoring and machine translation are live and have cut translation costs 50 to 80%, and AI agents run around the clock to scan code repositories for security vulnerabilities ahead of EU Cyber Resilience Act deadlines. At the other end, smaller manufacturers and service providers are still in early rollout, using off-the-shelf tools such as Copilot and consumer chatbots for individual tasks like email, Excel formulas, or basic documentation, without a formal strategy or platform.

The gap between leaders and typical organizations shows up most in data infrastructure. Companies already running integrated PLM systems (Windchill, ARAS) can move faster; companies still working from scattered spreadsheets and siloed departmental data cannot get meaningful value from AI regardless of the tool. Leadership alignment follows the same divide, from encouraged experimentation through engineering-week hackathons at Hypertherm to bottom-up, individual adoption at smaller firms. Across all three sessions, participants repeatedly said the real constraint was not the AI tools but everything underneath them: data structure, training, and clear objectives. The leading edge sits at Systemic and the typical organization at Active, a placement held with moderate confidence. National context: one 2026 survey found AI adoption had reached 72% of manufacturers but only about 10% had scaled it [3], while a Census-based analysis found roughly 87% had not yet integrated AI into operations as of early 2026 [5]; weak IT/OT collaboration and unresolved cybersecurity gaps are what stall further progress [1].

Key Use Cases

- Structured technical authoring and machine translation (Hypertherm): live and in production; moving documentation to XML/DITA structured authoring made content reusable and AI-ready, cutting translation costs 50 to 80%.
- Quality documentation automation (Air Technology Corp QA team): live; AI helps generate Production Part Approval Process documentation and PowerBI queries, cutting tasks that took days down to minutes.
- AI-assisted parts and product identification in sales (SMC Corporation): live; AI helps customers narrow to the correct parts series for their environment or media.
- Predictive, forward-looking inventory management: discussed as a near-term opportunity; shifting from reactive planning to AI-enhanced forecasting was estimated to potentially cut safety-stock levels 25 to 50%.
- Regulatory and compliance monitoring (Hypertherm): live; AI agents scan software bills of materials and code repositories continuously to track EU Cyber Resilience Act compliance and flag hard-coded credentials.

- Institutional knowledge capture from a retiring workforce (Opsmate, via EAC): being piloted; aimed at capturing tacit knowledge before experienced employees leave.
- CAD and engineering-change automation (BOM manipulation, ECO generation, Creo drawing updates): actively wanted, not yet solved, with technical dependencies on parametric CAD models an open problem.

National context: a 2026 manufacturing and supply-chain survey found inventory, quality, and supply-chain management are the AI use cases with heaviest adoption nationally [2], the same cluster raised in the NH roundtables.

Primary Barriers

The most consistent barrier named across all three sessions, unprompted, was the state of the underlying data: legacy systems, department-level silos, and scattered spreadsheets that keep production, maintenance, and quality information from talking to each other. Participants were explicit that this is more of a blocker than the AI models or platforms themselves.

Second was a lack of formal training and unclear objectives; companies often do not know what they do not know when starting, and keyword-matching applicant-tracking systems get mistaken for real AI, which muddies expectations. Integration complexity is a specific technical barrier for manufacturers running multiple disconnected systems (ERP, PLM, EPDM, MES). Security and intellectual-property concerns are real but partly a misconception, with some firms over-restricting AI out of fear that standard business accounts make their data public. Cultural resistance shows up around updating legacy engineering specifications and drawings, and regulatory complexity (EU Cyber Resilience Act, life-sciences validation, state-by-state rules) adds friction for smaller firms without dedicated compliance staff. National context: skills shortages, financial constraints, and fragmented data are the leading barriers for small and midsize manufacturers internationally [9], and one 2026 survey found cybersecurity the top-cited barrier (40%) with weak IT/OT collaboration at 43% of organizations [1].

Workforce Implications

Workforce effects showed up mostly around recruiting and technical-documentation roles rather than wholesale displacement. On hiring, AI-driven applicant-tracking systems were described as a frustration: resumes not matching exact keywords get rejected automatically, and participants distinguished this keyword matching from genuine AI screening that could understand context; manual networking and direct outreach were still more effective, especially at smaller companies. On the technical side, technical writers and quality engineers are picking up structured authoring, prompt engineering, and AI-generated PowerBI queries, and one company is launching an internal AI user group focused on prompt engineering.

Initial resistance among technical staff, including software developers, was overcome once people saw tangible results, such as long-standing bugs finally getting fixed or a hackathon project replacing a graphics engine in three days with AI assistance versus an estimated six to nine months without it. Training pipelines are informal and on the job (hackathons, internal user groups) rather than formalized. A broader open question was how senior management expects KPIs to be developed to determine the ROI of AI spend.

Data Readiness

Data readiness was raised as the central bottleneck in every session, more than any other topic. The recurring pattern is legacy systems and department-specific silos: maintenance records, Statistical Process Control data, and production data exist but are trapped in disconnected systems and rarely cross-referenced, which prevents manufacturers from correlating machine health with process capability or feeding AI anything structured enough to be useful. Companies already on integrated PLM platforms (Windchill, ARIS) are meaningfully ahead; companies without one were advised, bluntly, to clean up spreadsheets and organize scattered work instructions first, since the underlying risk is garbage in, garbage out. Governance ownership of data is inconsistent: larger, more regulated organizations are building procedural controls and data-flow tunnels to validate AI behavior, comparable to Sarbanes-Oxley-style compliance, while smaller manufacturers generally have no named data owner at all. National research describes a similar pattern, with data flows and exception handling between automated systems remaining manual [1].

AI Safety and Risk

Governance is uneven and mostly reactive, not proactively designed. At the more mature end, some organizations require mandatory AI training and a defined usage policy before employees may use AI tools, including verifying AI output before use and defining specific data-handling swim lanes, and a few larger or parent-company-governed organizations run secure internal AI platforms with IT-enforced policies. At the typical end, governance is closer to informal norms or none at all, with IP and data leakage the risks named most often, including outdated fears that standard AI accounts make company data public and a well-known incident of proprietary code exposed through an early public chatbot used as a cautionary tale. Regulatory pressure, more than internal initiative, is driving formal governance in some cases: the EU Cyber Resilience Act's software-bill-of-materials requirements and life-sciences validation rules are pushing structured, auditable practices ahead of where voluntary governance has otherwise reached. The sector places itself at "still figuring it out". National context: state initiatives follow similar patterns, establishing task forces and principles while leaving technical controls to employers [10][6].

Ethics

The clearest ethical tension was in hiring and customer-facing AI. Candidates described AI-driven interview processes, such as chatbots recording unretractable responses or self-facing video interviews, as impersonal and difficult, and job seekers can be filtered out by keyword-matching systems not sophisticated enough to evaluate them fairly. Participants generally agreed that when AI is used in customer service, setting clear expectations that a person is talking to AI, with an escalation path to a human, made the experience feel less deceptive. Within manufacturing, a risk arises if decision-making is left to AI agents without oversight by subject-matter experts, such as AI-assisted quality dispositions or engineering-change approvals, where poor decisions could affect manufacturability, quality, product robustness, and ultimately customer satisfaction and brand.

Impact Story

A Hypertherm developer used Claude during an internal engineering-week hackathon to replace a graphics engine in three days, a task that would normally have taken six to nine months.

What This Sector Needs Most

The most useful intervention would be a New Hampshire Manufacturing AI Readiness and Skills Accelerator: a small, cohort-based program that provides a practical roadmap instead of another general awareness guide. It would help manufacturers move from an initial question to a governed, measurable pilot, whether they are new to AI or already running disconnected experiments without a shared operating model.

The roundtables revealed a central challenge: companies “do not know what they do not know.” A checklist can frame the right questions, but a cohort and community of practice can help answer them. The program would combine a standard readiness assessment, role-based workforce planning, foundational governance templates, manufacturing-focused technical assistance, and peer learning. Each employer would choose one focused workflow, set a baseline, identify data and skill gaps, establish human-review and security controls, run a time-limited pilot, and decide whether to scale, revise, or stop.

New Hampshire can connect employers, workforce boards, education partners, apprenticeship programs, and incumbent-worker training around real manufacturing workflows. NH MEP can serve as the manufacturing-focused entry point; the Community College System and WorkReadyNH can translate employer needs into accessible training; the Job Training Fund can support eligible employer training; and the NH Tech Alliance can convene the peer community. Rather than choose a single AI vendor or create a large bureaucracy, the state can reduce adoption barriers by aligning existing institutions around a shared readiness model and evidence from controlled pilots.

National reference models include Louisiana's LA.IO readiness assessments for manufacturing and logistics SMEs [8] and Mississippi's AI Workforce Readiness Council, which pairs a statewide skills framework with targeted grants [7].

Chapter 15. Professional & Business Services

Sector Leads: Brian Bouchard, Cameron Shilling

Current State

Across New Hampshire's professional services sector, AI adoption remains a study in sharp contrasts. Forward-thinking organizations, particularly within the state's legal sector, established governance frameworks years ago, actively experimented with enterprise-grade platforms, and fundamentally modernized their workflows. Leading New Hampshire law firms are using AI to accelerate legal research, automate routine drafting, and redesign associate training so junior attorneys build analytical depth rather than grinding through repetitive tasks or becoming overly reliant on the AI technology.

Much of the broader professional services community, notably regional CPA and wealth management firms and fractional service companies, remains tentative. Many practices are still formulating baseline acceptable-use policies and have restricted experimentation to ad-hoc, web-based large language model queries. One driver of this divide is regulatory framing: New Hampshire attorneys operate under formal professional-conduct rules that explicitly mandate technological competence, creating an institutional imperative to understand emerging legal technology. The leading edge of the sector sits at Systemic, while the typical organization sits at Active.

Key Use Cases

- **Billable research and discovery efficiency (attorneys):** AI synthesizes cross-jurisdictional case law, tax rulings, and state-specific code revisions into first-draft analysis memos. Tasks that required substantial attorney time, mostly through associates, in the pre-AI era now require a fraction of the time. Although expected, this is a very meaningful shift for an industry that has largely relied on the “billable hour” as its primary measure of revenue.
- **Improved client communications (industry wide):** Nearly everyone in the professional services industry is using AI to ghostwrite client correspondence. The obvious effect is that communication is becoming crisper and more polished, but many also reported that AI removes the barrier to important communications being sent, so clients are receiving better and more timely correspondence from their service provider.
- **Financial analysis and audit efficiency (accountants):** AI automates routine bookkeeping tasks, reconciles accounts, categorizes transactions, and generates first-draft financial statements. During audit season, AI tools accelerate workpaper preparation, extract key data from client records, and flag anomalies for review—reducing the volume of manual data entry while allowing CPAs to focus on judgment-intensive advisory work.
- **Cash flow forecasting and portfolio analysis (financial management):** Wealth managers and financial planners use AI to model retirement scenarios, stress-test investment portfolios, and generate personalized client reports. This is an impactful example of AI synthesizing significant volumes of data, while still producing something personable for the client.

- Recruiting, onboarding, and policy administration (human resources): HR professionals use AI to draft job descriptions, screen resumes for qualified candidates, and generate first drafts of employee handbooks and workplace policies. Many also use AI for compliance questions, which creates confidentiality concerns.
- Deliverable acceleration and knowledge synthesis (consulting): AI compresses research timelines by synthesizing information across multiple sources into structured frameworks, enabling consultants to spend more time on strategic recommendations and client engagement and less on initial data gathering and formatting. In many ways, AI is creating a dynamic in which experienced human judgment, as opposed to cognition, is becoming the scarce resource in the labor market.

Primary Barriers

The barriers slowing adoption are fundamentally cultural, anchored by institutional inertia and defensive habits more than a lack of technology. Five patterns emerged across the breakout sessions:

- The old-guard mentality. Entrenched leadership often treats legal and accounting work as an artisanal craft defined by the grind, so delegating drafting, analysis, or reconciliation to software feels like diluting professional rigor and commodifying services, and AI gets dismissed as an unwelcome shortcut.
- Replacement anxiety. Many employees fear that demonstrating radical efficiency will render their roles obsolete, so without explicit assurances, rank-and-file staff quietly resist adoption.
- Data security as a shield. Legitimate concerns about client confidentiality and privilege are frequently weaponized into excuses for total paralysis. Rather than evaluate closed-loop, enterprise environments that safeguard proprietary data, conservative committees use security as a conversation-ender, and some organizations simply lack the sophistication and resources to evaluate the concern and so remain stuck.
- The cart-before-the-horse fallacy. Many organizations believe they must define an airtight, proven business case before enacting governance policies or permitting sandbox experimentation. The circular logic freezes progress: staff cannot safely uncover valuable applications without ground rules, while leadership refuses to set rules without demonstrated use cases.
- Cost. Many bespoke or industry specific AI products are expensive, bordering on cost-prohibitive, and not all organizations can justify the price tag. Additionally, some organizations that have integrated AI have not been able to measure the ROI of “what has become an industry staple”.

Workforce Implications

AI is reshaping workforce dynamics across the sector, though the impact varies significantly by role and organizational culture. At the entry level, paralegals, junior associates, and staff accountants are experiencing the most pronounced transformation: tasks that once consumed the bulk of their time, including document review, initial contract drafting, tax-return preparation, and basic research, are increasingly augmented or automated. Rather than displacement, leading firms report these roles evolving toward higher-value work such as quality control of AI outputs, client communication, and earlier exposure to substantive analysis.

Training pipelines tell a divided story. Some organizations have proactively built AI literacy into junior-development programs, treating prompt engineering, output verification, and ethical AI use as core competencies alongside traditional skills. Others are encountering resistance at both entry and managerial levels: junior employees sometimes fear that embracing AI undermines their path to partnership, while mid-level managers worry about justifying their oversight roles when AI performs tasks they once delegated downward, which shows up as passive non-adoption or outright skepticism.

The skills gap most frequently cited is not technical proficiency with the tools but critical evaluation, meaning the ability to assess AI-generated work product for accuracy, completeness, and professional judgment, and employers consistently report difficulty finding candidates who can blend domain expertise with AI-augmented workflows. Fear of displacement surfaced most acutely among administrative support and bookkeeping roles, where automation potential is highest and alternative career pathways within the firm are least clear.

Data Readiness

Data quality varies widely depending on firm size and practice area. Larger law firms and accounting practices have invested in document management systems and practice management software over the past two decades, resulting in reasonably structured repositories of contracts, correspondence, and financial records. These organizations can deploy AI tools for document review, contract analysis, and research with relative confidence in their underlying data. Smaller firms, however, often operate with fragmented data across disparate systems—or even rely on paper files and informal digital filing, making AI deployment impractical without substantial data remediation.

Data access presents both technical and regulatory challenges. Most professional services firms handle extraordinarily sensitive information: privileged attorney-client communications, confidential tax returns, proprietary business strategies, and protected health information in healthcare consulting. Regulatory constraints including attorney-client privilege, IRS Circular 230, and HIPAA create legitimate guardrails that limit how data can be used to train or fine-tune AI models. Many professionals remain uncertain whether feeding client data into AI systems, even those with strong confidentiality protections, complies with their professional obligations. This uncertainty is well-founded; guidance from state bars and professional associations is still evolving and inconsistent.

Despite these challenges, data is not an absolute blocker for the sector. Many AI applications, particularly those involving general legal research, drafting assistance, or knowledge synthesis can function effectively using public data and proprietary AI training without requiring access to firm-specific client data. The critical path forward involves distinguishing between AI use cases that require firm data (and therefore robust governance) and those that can deliver value immediately using external resources and carefully managed inputs.

AI Safety and Risk

Governance maturity is a mixed bag, and the sector marked itself across all three levels at once, from “ahead of the curve” to “still figuring it out” to “not yet a priority”. Some organizations, particularly larger law firms and accounting practices with national affiliations, have implemented sophisticated protocols: formal acceptable-use policies, designated AI oversight committees,

mandatory training, and clear escalation pathways, treating AI governance as an extension of existing data-security and professional-responsibility frameworks. Others, particularly smaller regional practices, have almost nothing formalized, with AI use proceeding informally under individual discretion and unwritten norms, a significant vulnerability.

Everyone shares in AI risk, from the immediate user to supervisors to legal, compliance, and IT leadership, but the nature of the concern differs by role. Immediate users focus on data quality, ensuring inputs are accurate and outputs verified before they are relied on; supervisors worry about oversight gaps and accountability; legal and compliance teams concentrate on confidentiality, privilege, and regulatory exposure; and IT focuses on security, system integrity, and vendor management.

Ethics

For attorneys, the New Hampshire Rules of Professional Conduct intersect directly with AI use. Rule 1.1 (Competence) requires competent representation, which includes understanding the benefits and risks of the technology used in a client matter. Rule 1.3 (Diligence) demands reasonable promptness, and while AI can accelerate work, relying on flawed output without verification breaches that duty. Rule 1.6 (Confidentiality) prohibits revealing client information without consent, raising serious questions about which platforms are appropriate for sensitive matters and whether client data may be used to train third-party models.

For accountants, financial advisors, and other practitioners, parallel concerns apply. Fiduciary obligations require acting in the client's best interest, so delegating analysis to AI without appropriate oversight may compromise that duty. Client transparency about AI use is an emerging expectation, especially in high-stakes contexts, and protecting confidential information remains paramount, since inadvertent disclosure through AI tools could expose firms to liability and erode trust. A deeper tension exists between automation and professional judgment: clients engage professionals precisely for experienced human analysis, and over-reliance on AI risks commoditizing services and undermining the value of professional expertise. Finally, fee arrangements may shift, because if AI dramatically reduces time spent on certain tasks, traditional hourly billing faces pressure, yet consensus on alternatives such as value-based pricing or efficiency-sharing has not yet emerged.

What This Sector Needs Most

The single most useful step named is for the Task Force to publish a basic AI governance template or blueprint the sector can adopt.

Checked against the February signals, the leads flagged two takeaways in particular. First, AI maturity and risk awareness remain remarkably uneven across organizations. Second, and contrary to the statewide trend, AI is not seen merely as a tool for efficiency: employees at every level are increasingly concerned about AI replacement, to the point that the concern itself is making them reluctant to adopt.

Chapter 15a. Real Estate Breakout

Sector Lead: Chris Norwood

Current State

Real estate has high individual adoption and almost no organizational governance, and the gap between those two is the story. Agents are adopting AI faster than brokerages can govern it, and because the workforce is overwhelmingly independent contractors, adoption is thousands of separate experiments rather than firm strategy. New Hampshire has roughly 7,000 individual real estate businesses, each with its own approach to tools and data security.

Adoption is concentrated in pockets and very uneven. At one end are two-year power users building their own apps, community maps, and deed-research tools and working daily to keep pace with the national platforms. At the other are practitioners who were, in their own words, dragged in kicking and screaming and now use AI as a better search engine. Firm-wide use is the exception; several people are further along than their brokerages know, and brokerages often do not know which tools their agents use.

The defining feature of this sector is that AI has entered the client relationship directly. Buyers and sellers now arrive with their own AI-generated valuations and market predictions, frequently wrong for New Hampshire, and correcting that has become part of the job. The leading edge sits at Operational, with a few individuals looking Systemic; the typical practitioner is at Active, with a real tail still at Awareness.

Key Use Cases

Concrete workflows heard, mostly live in individual practice, include:

- Listing copy: drafting and, more often, proofreading MLS remarks. Several practitioners found fully AI-written descriptions too canned and shifted to using AI to polish their own writing.
- Client-ready valuation prompting: firm-level prompt templates that force local comparables from Redfin, Zillow, and Realtor.com plus local variables such as school districts, days on market, and recent price reductions, used both to price listings and to counter clients' bad AI valuations.
- Land and title research: surfacing rights of way, easements, and historical deeds, including reading 1700s cursive across counties, and building community maps and apps.
- Document summarization: extracting key data points from inspection reports, HOA documents, and multiple-offer comparisons.
- Communications triage: scanning the day's inbox for unanswered client questions and summarizing long correspondence threads.
- Field due diligence: verifying obscure facts fast, from organic-orchard certification dates to a town's sewer-line-layer license.
- Commercial estimating and business development: quantity takeoffs from drawings, and cross-referencing event attendee lists against LinkedIn and public records, tracking lease expirations and job postings as expansion signals.

- Marketing and search: blog posts written specifically to be surfaced by AI search so the agent becomes the cited source.

Primary Barriers

The barrier unique to this sector is AI misinformation reaching clients directly. Buyers and sellers bring valuations and market calls generated from national data, often pulled from California, Texas, or Florida, that are simply wrong for New Hampshire, and then argue with their own agent about them. In one case a seller fired the agent within an hour over an inflated AI valuation. Correcting a client's AI takes as much work as doing the analysis from scratch, the same pattern seen when clients bring pre-written AI contracts to a lawyer. Southern New Hampshire has ranked among the top four national markets for years, so the AI narrative that the market is cooling is actively false here and pushes clients toward lowball offers and waiting.

Behind that sits a governance vacuum made worse by the independent-contractor structure. Most brokerages have no AI policy, and enforcing one over independent contractors is hard. Unreviewed CRM and third-party integrations quietly harvest client data and can breach fiduciary duty and confidentiality without the agent realizing it. Third is the prompt-skill gap, garbage in and garbage out, and the verification tax that now falls on the end user to fact-check every output. The single unlock would be clear, enforceable guidance on what data may go into AI, paired with brokerage policy that comes with training, not just a document.

Workforce Implications

AI is transforming real estate work rather than cutting roles. Agents use it for research, marketing, and communications; managing brokers and owner-brokers carry the accountability for policy and enforcement; commercial estimators and business-development staff use it for takeoffs and lead research. The human agent as local expert and fiduciary is the differentiator, and practitioners are actively reasserting that role against client AI. Displacement fear surfaced less as job loss and more as being outrun: small brokerages worry about competing with Zillow, Redfin, and Compass, which field large AI-development teams.

The missing skills are prompt engineering, basic AI literacy, and security awareness, plus a broker-level ability to orient, train, and enforce. Training is uneven and informal. One practitioner found a beginner course far over her head and gave up for a year before trying again; another firm trained an employee who then left to start her own AI business, the same retention and return-on-training concern that came up in the small business roundtable.

Data Readiness

Data problems came up repeatedly, in two distinct forms. The first is source data: public records fragmented across county registries, historical cursive deeds, and county-by-county access differences, for example Carroll County restricting deed viewing. AI surfaces this material fast but demands heavy prompting and verification. The second, and sharper, is data input and governance: unreviewed CRM and third-party integrations automatically collect and share client data, and there are no clear rules on what non-public client information may be entered into a model. Practitioners specifically asked for state or association guidance drawing that boundary.

Regulatory constraints are real and state-specific, including fair housing, license law, and errors-and-omissions exposure, and national AI data is frequently wrong for New Hampshire.

Ownership of data governance mostly rests with nobody at the brokerage level: individuals self-manage, some very strictly, and the association has materials, but there is no clear owner across the sector.

AI Safety and Risk

Risk management is mostly ad hoc, with rising awareness and low maturity, and the sector places itself at not yet a priority. Written brokerage AI policies are rare. The NH Association of Realtors offers a sample policy manual members can adapt, but brokerage-level adoption is thin and many practitioners do not yet realize they need a policy; commercial firms are more likely to have paid-access controls in place.

The risks names include confidentiality and fiduciary breach through silent integrations, data leakage to cloud tools, hallucinated valuations, fair-housing violations in AI-written remarks, misrepresentation through virtual staging and AI-enhanced photos, loan fraud where AI-enhanced photos meet desktop appraisals, and deed fraud enabled by AI-generated identification. Ownership of AI risk lands on managing brokers and owner-brokers, but enforcement is the gap. As one veteran put it, a policy is not worth the paper it is on without orientation, training, follow-up, and enforcement. Some practitioners self-impose strict controls, storing client data locally and putting nothing in the cloud. Regulation is not driving any of this.

Ethics

The sharpest sector-specific tension is disclosure to consumers. AI-generated and AI-enhanced listing imagery, including virtual staging that shows a fixer-upper as fully rehabbed, creates misleading expectations, an amplified version of the old wide-angle-lens problem. Whether buyers should be told when imagery or remarks were AI-produced is unresolved, and misrepresentation lawsuits around descriptions and photography are already emerging.

Fiduciary duty and confidentiality are a second tension: integrations can harvest and share client data without the agent's knowledge, and one agent's tool can expose the other side's client data in a co-brokered deal. Fair housing is a third, since AI-written remarks can carry noncompliant language under the agent's name. The fraud dimension is also real: AI-generated IDs, including fake New Hampshire driver's licenses, worsen deed fraud, especially with absentee owners and vacant land, and AI-enhanced photos combined with desktop appraisals raise loan-fraud and class-action exposure. Underneath all of it is automation against professional judgment: clients trusting a national model over their licensed local agent, whose duty of accuracy cannot be delegated to a confident but frequently wrong tool.

What This Sector Needs Most

A New Hampshire brokerage AI policy starter kit that comes with a training and enforcement model, not just a document. A veteran practitioner asked directly for a starting-point policy, another stressed that a policy without orientation, training, and enforcement is worthless, and the Association already has a sample manual to build from. The clear priority is turning that raw material into something brokerages can adopt and actually enforce over independent contractors. Close behind, in rough order:

- Clear state or association guidance on data boundaries, specifically what non-public client data may and may not be entered into a model.

- Broker education on platform integrations, their data effects, and their impact on errors-and-omissions insurance, delivered through the state and local boards.
- Client-facing education tools, such as the comparative-prompt templates several agents built, to counter AI misinformation at the listing appointment.
- Objective, comparative data on the infrastructure demands of AI data centers, water and energy against golf courses and industrial users, for the public-policy conversation.

If only one thing were possible, it is the enforceable policy-plus-training starter kit.

Added Observations

The most inventive responses in the room were about client-education tools, not productivity hacks. A dissenting current worth capturing is a veteran practitioner's blunt view that much of what clients get from AI is, in his words, a lot of bunk, a warning about over-trusting outputs more than a rejection of the tool. And a structural note: with roughly 7,000 independent businesses and an independent-contractor workforce, any guidance must be adoptable at the individual level and enforceable at the brokerage level, or it will not land.

Chapter 16. Consumer & Small Business

Sector Lead: Liz Gray

Current State

Consumer and small businesses in New Hampshire are experimenting, not operating. AI is real in this sector, but adoption is individual, uneven, and thin below the surface. Most use is a single motivated person inside an organization rather than a team or company rollout, and the distance between what people say they are doing and what is actually in production is wide.

Adoption is concentrated in pockets. The furthest along is the small business support side itself: the SBDC has a named AI program manager, an internal use policy it reviews quarterly, and advisors using AI in nearly every client session, roughly 300 unique clients with only one opting out of transcription. A large independent real estate firm has rolled Microsoft Copilot to its executive team while many of its agents self-fund Claude on their own. A 130-person business software firm has enthusiastic individual power users but no company policy. A four-person Concord nonprofit is essentially one person experimenting on the side.

The leader-to-laggard gap shows up inside single organizations as much as between them: leadership dabbles while frontline staff and independent contractors improvise, sometimes around the rules. The real baseline is a sector that has crossed into real use in isolated spots but has very little that is governed, measured, or repeatable. The leading edge sits at Operational, the typical organization at Active.

Key Use Cases

Concrete workflows heard in the room, live unless noted, include:

- Marketing collateral, listing videos, and social media mashups, the most visible production use because the output is easy to see.
- Meeting transcription, notes, and agenda drafting across every organization present, mostly through Copilot; one nonprofit uses it to turn meeting notes into strategic-plan agendas.
- Client cash-flow projections in SBDC advising, using Claude with anonymized client data and personally identifiable information held back.
- Search replacement, with advisors using Perplexity in place of Google specifically because it cites sources they can pass to clients.
- Business plan creation for clients, mostly through ChatGPT.
- RFP dissection and prospect research at the software firm, via a personal Claude agent loaded with company context to break down proposals and qualify prospects.
- NotebookLM turning hour-plus conference transcripts into short internal podcasts and summaries, and assistants generating action lists from meetings.

Still experimental or failing: custom agents built to search SharePoint or read New Hampshire law, which largely did not work because the source data is buried in PDFs and unstructured stores. Direct email integration is wanted everywhere but blocked or unapproved in every organization present.

Primary Barriers

The real blocker is trust and security, not capability. Every organization in the room wants to connect AI to email and internal systems, and none has been allowed to, so people build workarounds. In one firm at least one agent forwards company email to a personal Gmail account so a model can read it, pushing company data outside the walls to get the workflow they want. There is real fear that even anonymized data, a scrubbed profit-and-loss statement for example, can be re-identified by a model that already knows the firm.

Right behind that is a governance vacuum. Policies are absent, thin, or half-written, sometimes drafted by AI because nobody knew what to put in them. The independent-contractor structure of real estate makes any policy nearly impossible to enforce, and IT support is slow: one formal request to connect Outlook has sat in research for four months. Third is accuracy and the human-verification tax: outputs are confidently wrong, described as an overconfident middle schooler, and in one case an AI-built market analysis invented data before a client caught it. Fourth is cost: budgets do not exist, spending is petty cash, token and credit pricing is volatile, and increases are expected within six months. The single unlock would be trusted, plain-language security and governance guidance, so organizations stop improvising and stop routing data around their own controls.

Workforce Implications

AI is reshaping work quietly rather than eliminating roles. Business advisors use it to augment client-facing work while insisting the human cannot be replaced in advising; independent real estate agents use it for marketing and drafting; administrative and executive-assistant roles use it to turn long meetings into action lists; and marketing is the function most visibly changed. No strong fear of displacement surfaced, and the framing throughout was augmentation, not replacement.

The sharper workforce signal is training and retention. One firm sent an employee to a six-week AI course and she left to start her own AI business, which captures both the value of the skill and the risk of building it and losing it. The missing skills are foundational: prompt engineering, basic AI literacy, and security awareness. Training so far has been fits and starts, with people hitting one bad output and giving up entirely, and what organizations want is practical, job-specific training that cuts through generic and expensive consultant classes.

Data Readiness

Data problems came up repeatedly and unprompted, but showed up as the reason custom builds fail rather than as a headline complaint. Records are fragmented and often unstructured: agents built to search SharePoint returned poor results, and an agent meant to read New Hampshire law could not, because the law is buried in PDFs. Legacy and unstructured data is the ceiling on anything beyond consumer-grade tools.

Regulatory constraints are understood as real but not well supported. A housing nonprofit operates under HIPAA because it serves people with severe mental illness; real estate agents operate under fair housing and license law; advisors work under client confidentiality. People know these lines exist; what they lack is clear guidance on how to use AI safely inside them. Governance ownership mostly rests with nobody, or defaults to whichever single person is championing AI, with the clear exception of the SBDC, which has a named owner and a quarterly review cycle.

AI Safety and Risk

Risk management is mostly ad hoc, with high awareness and low maturity, and the sector places itself at “still figuring it out”. The room readily named data leakage, re-identification of anonymized data, hallucinations, and shadow-IT workarounds, but few organizations have real controls behind any of it. Written policy is the exception: the SBDC has one it reviews quarterly, and a real estate firm is on its second or third iteration, now building from the National Association of Realtors baseline, while a 130-person software firm has none beyond a guide for one recording tool and a nonprofit has none yet.

Ownership of AI risk is unclear and usually lands on the AI champion by default. Managed service providers handle security but are frequently not equipped to advise on AI, and the caution raised was that today's IT partner may be the wrong AI partner. Regulation is not driving any of this; governance here is voluntary, reactive, and early.

Ethics

The sharpest ethical tension is confidentiality against convenience. Advisors and firms hold sensitive client data, including financials and housing and health information, and the easy path routes it through consumer tools. Anonymization is not a clean answer, because a model that has worked with a firm for months can plausibly re-identify a scrubbed profit-and-loss statement. Accountability is a second tension: when independent contractors push company data outside the firewall to get AI to work, the exposure and reputational damage roll up to firm leadership even though leadership did not authorize it.

Disclosure is a third. In advising, transcription consent is asked and honored, and one client in about 300 declined. Where AI-generated market analyses are handed to clients, it raises the question of whether the client should know, and an AI-generated analysis that invented its numbers shows the stakes when the answer is no. Underneath all of it is automation against professional judgment: the advisor's and the fiduciary's duty cannot be delegated to a confident tool that is sometimes simply wrong.

What This Sector Needs Most

Practical, trusted, plain-language guidance delivered through hands-on workshops, not sold as an expensive class. The room was skeptical of the AI consultant at every networking event whose answer is a 700-dollar course, and hungry for help that cuts through that noise. In priority order, the concrete asks are:

- A collaborative write-your-AI-policy-together working session that sends each organization home with an actual policy, since everyone is too busy to write one alone.
- A catalog of role- and function-based use cases, organized by job (finance, marketing, client relationships), so a person can see where to start rather than testing tools at random.
- Clear security do-and-do-not guidance, including what data should never be uploaded and real examples of where it has gone wrong.
- Budget and pricing guidance, and nonprofit or scholarship access to paid tiers, because small teams run on petty cash and cannot forecast token-based costs.

If only one thing were possible, it is security-grounded practical guidance paired with collaborative policy workshops mapped to job function.

Additional Observations

Three things worth putting on the record. First, a pricing warning: the sector believes it is in a subsidized honeymoon and should expect per-user or per-token increases within roughly six months, and token and credit models make budgeting nearly impossible for organizations that run on petty cash, though there is hope that increasingly capable free models will pressure paid tiers.

Second, the partner gap: existing IT and security providers are often not equipped for holistic AI advising, and starting with the wrong partner is expensive to unwind, so organizations should vet AI advisors specifically rather than assuming the current IT vendor can fill the role.

Third, a dissenting current worth capturing: strong fatigue with AI hype and the sense that most training ends in a sales pitch without a tangible product. This is not a claim that AI is overhyped, but a demand to cut the noise. It pairs with a governance caution that doubles as a warning: policy without technical controls is words on paper, because these tools will actively help an employee bypass weak rules, for example by suggesting the browser client when the desktop app is blocked. Mapping controls to policy is the maturity step this sector has not yet taken.

Chapter 17. Public Sector

Sector Leads: Melanie McDonough, Shaun Mulholland

Current State

AI adoption in New Hampshire local government remains early, uneven, and largely decentralized. Municipalities are not uniformly rejecting AI, but many are proceeding cautiously. Individual employees and departments are experimenting with generative AI for everyday productivity, while relatively few municipalities have moved to mature, organization-wide programs with governance, approved platforms, training, and formal evaluation. The recurring posture is best described as interested, but cautious. Leaders see potential value, particularly given staffing constraints, but operate in a substantially different risk environment than most private organizations, weighing public accountability, cybersecurity, procurement requirements, records retention and Right-to-Know obligations, sensitive resident information, accessibility, accuracy, and the consequences of providing incorrect government information, all under growing budget constraints.

The picture is therefore not inactivity but fragmented experimentation occurring faster than institutional governance and organizational readiness. Smaller municipalities face an added capacity problem, often lacking dedicated data, cybersecurity, innovation, legal, or IT staff, so the same organization asked to evaluate AI may already have limited technology resources.

Key Use Cases

The most practical near-term applications are augmentation tools rather than autonomous decision-making systems: drafting and editing routine communications, reports, policies, summaries, and internal documentation; summarizing meetings, lengthy documents, regulations, and public comments; research and information retrieval across policies, ordinances, procedures, and records; data analysis and administrative support; knowledge assistants and resident-information tools that make complex government information easier to search; legal and contract analysis; and a small number of vite-coded, custom internal applications built rapidly to test what AI-enabled software development can do for government. Most use cases center on productivity gains where employees perform multiple roles and budgets are constrained.

Primary Barriers

Procurement and vendor review: government processes were not designed for technology that changes as fast as modern AI platforms, and evaluating pricing, licensing, data protections, model behavior, and integrations demands staff time and expertise many do not have.

- Data security and privacy: municipalities hold personnel, financial, public-safety, resident, and infrastructure information, and staff need clear guidance about what can be entered into which systems and under what protections.
- Right-to-Know and records management: prompts, outputs, generated documents, transcripts, and logs raise unresolved questions about how existing New Hampshire public-records requirements apply.

- Accuracy and hallucinations: plausible but incorrect answers are especially problematic for authoritative government information on ordinances, taxation, permitting, elections, public safety, and benefits, and liability is a concern.
- Limited capacity: many municipalities lack AI specialists, data scientists, innovation staff, or large cybersecurity teams to evaluate vendors, develop governance, train staff, and monitor systems.
- Organizational uncertainty: uneven AI literacy means some employees experiment extensively while others are hesitant, creating both unmanaged risk and uneven access to productivity gains.

Workforce Implications

The discussion centered on augmentation more than replacement, with AI viewed as a way to improve efficiency and extend staff capacity rather than reduce positions. The more immediate problem is capability: many organizations lack staff with the technical and strategic expertise to organize adoption across the enterprise, where the harder work is deciding where AI belongs, connecting use cases to organizational goals, setting governance, evaluating vendors, managing risk, and scaling pilots. There is no evidence of roles being eliminated or new AI job titles being created; instead, existing technology, management, legal, policy, and operational roles are being asked to absorb new AI responsibilities faster than training pipelines are developing.

Municipalities continue to face workforce constraints, difficulty recruiting specialized talent, retirement-driven knowledge loss, and rising expectations for digital service delivery. A significant risk is uneven capability, where employees who learn to work with AI gain productivity advantages over colleagues who lack access, training, or confidence. Needed skills include AI literacy and responsible use, evaluating AI-generated information, understanding data sensitivity, recognizing hallucinations and bias, designing human-AI workflows, knowing when AI should not be used, and documenting and validating AI-assisted work.

Data Readiness

Data readiness is a major constraint, but not always the first barrier named, because it typically surfaces after AI is adopted; organizations still in the awareness stage often have not reached it. The gap is most visible where in-house IT capacity is limited, leaving organizations unsure what data they have, how it is structured, who owns it, and what can be shared with an AI system. Portsmouth's chatbot illustrates the work required: before launching, the city spent significant time normalizing its data so the system had a usable base. For the public sector, policy and guidance uncertainty often appears before data readiness in stated barriers, as organizations first ask whether they are even allowed to use AI on given information; those who answer that question typically find data governance is the next barrier.

AI Safety and Risk

Governance maturity is emerging and highly uneven. Some organizations have written policies, acceptable-use rules, technology-review processes, and clear oversight, while others are drafting a first policy or relying on informal norms, and in the least mature environments leaders may believe AI is not being used while having little ability to verify. The risks raised most consistently are privacy and security of public information, inappropriate use of data, bias, reliability, and the possibility of an AI-supported process producing an unacceptable result without sufficient human oversight.

New Hampshire's Right-to-Know law adds uncertainty about how prompts, outputs, and records fit existing obligations. Overall risk awareness is high, but the structures for managing risk vary considerably, and the sector places itself at “still figuring it out”.

Ethics

Bias is a central concern, because an AI-supported process could disadvantage members of the public or reproduce inequities in access, service delivery, or decision-making, and participants questioned whose voices and experiences are missing from training data. A second concern is the environmental cost of AI: for public organizations responsible for climate goals and sustainability planning, AI's energy demands create a real tension between adopting useful technology and its environmental footprint. These concerns reinforce the need for human judgment and oversight, so the question is not simply whether a tool can perform a task but whether its use preserves fairness, accountability, transparency, and public trust.

Impact Story

The City of Portsmouth offers an early example of a New Hampshire municipality moving from AI experimentation into a live public-facing use case. Portsmouth's first AI pilot was a custom chatbot trained specifically on city projects. Before it went live, the city spent significant time normalizing and preparing the underlying data. The tool now answers questions about city projects in real time, drawing on data entered by staff onto the website and pulling from the city's GIS data. It shows that a visible AI application is only part of the work: data preparation and organizational readiness came first and took extensive effort and change management before launch. It also gives other municipalities a concrete example of AI moving beyond individual productivity tools into services built around trusted local information.

What This Sector Needs Most

New Hampshire's local governments need practical support to adopt AI responsibly, and the strongest opportunity is shared resources and expertise that municipalities cannot efficiently build alone. Priorities include practical guidance and shared resources (model policies, acceptable-use rules, purchasing questions, security guidance, and implementation examples); clear public-records guidance developed with legal and records-management experts; affordable tools and training, potentially through cooperative purchasing so cost does not exclude smaller communities; reusable training and implementation assistance, with higher education supporting training and applied research; a continuing peer network or community of practice free of vendor pressure; and a clear path for responsible experimentation that distinguishes routine productivity uses from higher-risk applications. These supports would be especially valuable to smaller municipalities without dedicated technology, cybersecurity, legal, or innovation staff. On the February signals, all hold, with data readiness qualified (policy and guidance uncertainty is often named first) and the inflection-point signal holding but at risk if practical adoption capacity does not accelerate.

Chapter 18. Education

Sector Lead: Shawn Powers

Current State

The education sector in New Hampshire shows substantial AI activity, but adoption remains highly uneven. The most advanced institutions have moved beyond experimentation and integrated AI into professional development, curriculum planning, administrative workflows, student support, and governance. Pinkerton Academy, for example, described a three-year journey that includes AI policy development, faculty training, approved tools, and a shift from AI literacy toward AI fluency. UNH, Rivier, community colleges, and SNHU have also launched structured efforts around AI literacy, governance, and instructional integration. The leading edge sits at Systemic while the typical organization is at Active.

Key Use Cases

Current use cases focus more on productivity, instructional support, and student services than full automation.

- In curriculum and lesson development, faculty use AI to create lesson plans, rubrics, learning activities, assessments, and instructional materials.
- For administrative efficiency, educators draft communications, create reports, summarize documents, organize information, generate formulas, and analyze institutional data. In student learning support, institutions are experimenting with AI tutoring systems, writing assistance, study guides, flashcards, and learning companions.
- For career and workforce preparation, AI is integrated into career-readiness programs, resume optimization, internship coordination, and skill development, as employers increasingly expect graduates to have AI literacy.
- And in governance and institutional planning, universities and colleges use AI to support policy development, governance evaluations, risk management, and operational decision-making.

Primary Barriers

The strongest barrier is institutional readiness, a combination of limited training, inconsistent guidance, funding constraints, and uncertainty about appropriate use. Data-privacy requirements create a particularly acute challenge for K-12 education, where New Hampshire's privacy requirements and vendor-vetting processes can delay adoption for months and restrict access to otherwise useful tools. A second barrier is workforce readiness: faculty and staff often receive little formal preparation and must learn independently while carrying existing responsibilities, and many institutions lack dedicated personnel for AI implementation. A third is inconsistency, with policies varying across institutions, departments, and classrooms, so students frequently encounter conflicting expectations. If one barrier disappeared tomorrow, participants suggested dedicated funding for training, staffing, and implementation support would accelerate adoption most.

Workforce Implications

AI is transforming nearly every educational role rather than eliminating positions. Faculty are increasingly expected to use AI for course design, feedback generation, planning, and content development; administrative staff use it to analyze data, draft communications, and streamline reporting; and career-services professionals incorporate it into workforce-readiness activities. The most significant skills gaps involve AI literacy, critical evaluation, prompt development, data literacy, and understanding AI limitations, with a consistent emphasis that students and employees need to learn to collaborate with AI rather than simply consume its outputs. Job displacement did not factor into participants' perspectives; most described AI as a workforce amplifier that increases productivity while preserving the need for human expertise, judgment, mentorship, and interpersonal skills.

Data Readiness

Data readiness is mixed. Institutions generally possess substantial data assets but many struggle with fragmentation across multiple systems and departments, with challenges around access, integration, governance ownership, and interoperability. Data privacy remains a dominant concern, particularly in K-12 environments where student information is highly protected, and vendor vetting, compliance reviews, and data-protection agreements frequently determine whether a tool can be adopted. Most organizations view AI as workable today, but success depends on disciplined governance and responsible handling of sensitive information.

AI Safety and Risk

Risk awareness is high throughout the sector, but governance maturity varies significantly. A small number of institutions have formal policies, governance structures, approved platforms, and training programs, while many others remain in early stages of policy development or rely on informal guidance. The most commonly cited risks include privacy, security, hallucinations, inaccurate information, bias, misuse by students, and overreliance on AI during learning. Participants generally favored governance approaches emphasizing transparency, responsible use, and human oversight over outright restriction, and the sector at “still figuring it out”.

Ethics

The sector's primary ethical concern is preserving meaningful learning while integrating AI, as participants repeatedly questioned how educators can maintain curiosity, critical thinking, creativity, and productive struggle when answers are increasingly available through AI. Equity is another significant concern: students at institutions with strong AI programs may graduate with substantially different skills and opportunities than students whose schools restrict access, and differences between paid and free tools further increase disparities. Participants also raised bias, safety, privacy, developmental appropriateness, and the need for transparency when AI influences educational decisions or outputs.

Impact Story

One of the most compelling examples comes from Pinkerton Academy's four-year AI journey. School leadership began exploring AI almost immediately after generative AI tools emerged. They established a task force which later morphed into a Faculty AI Think Tank, created an AI policy,

expanded faculty training, piloted Khan Academy's Khanmigo before broader state adoption, and later revised their policy as institutional understanding matured.

Rather than focusing on AI detection, Pinkerton shifted attention to the learning process. The school created an "AI continuum" that allows educators to define acceptable levels of AI use (if any) for assignments and encourages conversations about how to think with, think without, or think about AI based on the learning objectives of the lesson. Faculty describe moving from AI literacy toward AI fluency, while students are being introduced to AI literacy through approved tools and structured learning experiences across the curriculum.

What This Sector Needs Most

The single most important need is a coordinated statewide framework for AI literacy, guidance, and professional development that can be adapted locally. Participants consistently requested practical guidance, shared best practices, training opportunities, examples of successful implementations, and clearer expectations about responsible AI use. Funding remains important, but many emphasized that alignment, coordination, and shared resources could deliver substantial value even in constrained budgets. On the February signals, most strongly hold, with data readiness only partially supported because privacy, training, and funding compete as equally important barriers. A recurring closing theme was that AI should not be framed primarily as a replacement technology but as an amplifier that raises the importance of judgment, creativity, mentorship, communication, and real-world experience; the strongest concern was not that AI will replace education but that institutions may fail to prepare learners to use it responsibly and effectively.

Chapter 19. Startups

Sector Lead: Jack Potvin

Current State

New Hampshire's high-growth startup sector is already moving well beyond occasional chatbot use. The leading companies are building AI into products, software-development systems, internal operations, data access, marketing, and decision-making, and increasingly describe AI as a coworker or operating layer rather than a standalone application. Adoption among the participating startups is advanced, broad, and embedded in normal operations, spanning product development, engineering, quality assurance, research, finance, accounting, operations, sales, and marketing.

The conversation frequently centered on two concepts: the AI harness (the surrounding infrastructure used to operate and control a model or agent, managing prompts, tools, context, permissions, retries, tests, evaluations, and operating rules) and the software factory (the broader end-to-end system for producing software repeatedly and systematically). Leading startups are working toward systems in which people define objectives and guardrails while multiple agents generate code, review changes, run tests, reproduce bugs, and verify results, a significant shift from using AI merely to autocomplete code. Another emerging pattern is the rise of AI coworkers: persistent agents connected to Slack, company knowledge, databases, code, and operational tools that can investigate requests and carry out work. This cohort appears to be among the most advanced AI adopters in New Hampshire, and its primary challenge is no longer access to AI but building the organizational systems needed to use AI reliably and safely at high speed.

Key Use Cases

- AI-enabled software development and quality assurance: agents draft product requirements and technical specifications, generate and refactor code, review pull requests, reproduce reported bugs, run unit/smoke/UI tests, and verify fixes. One company built a closed-loop QA process around Cursor cloud agents and runs roughly 50 historical regression scenarios before releases and model changes; it postponed hiring a dedicated QA employee, although a person still tests every release.
- Action-oriented AI coworkers: one startup deployed an internal OpenClaw agent named Odin through Slack. Connected to Notion, the codebase, databases, marketing materials, and external APIs, Odin investigates employee requests, submits pull requests, assembles evidence for high-confidence Stripe disputes (addressing roughly \$1,000 per month in disputed charges), drafts press releases, contacts journalists, indexes company systems, and operates roughly 40 recurring automated jobs. Its operating principle is effectively ask Odin first, automating work first and formalizing it into permanent code later only if it proves valuable.
- Proprietary AI products and natural-language data access: a New Hampshire sports-technology company uses proprietary computer-vision models to turn game video into statistics, analytics, and shareable highlights, and has built an MCP connection giving customers read-only access to their data through an AI assistant. A digital-identity

intelligence company uses AI to translate and analyze regulatory materials across countries, and an investment-data company applies AI to summarize lengthy SEC filings and surface red flags.

- Marketing and content operations: startups produce and repurpose written, visual, and video content, conduct competitive analysis, and support executive communications. One company reported producing more marketing content in roughly 20 minutes per day, at a cost of roughly \$4, than its previous full-time manual process. Results were not universally positive: one company found AI-generated outbound sales emails did not convert, because prospects were inundated with similar automated messages.

Primary Barriers

The primary barriers are not a lack of tools. They are workforce capability, validation, governance, security, context, and unrealistic expectations. AI proficiency is becoming a learned professional skill: merely providing access does not guarantee gains, since effective users know how to scope tasks, supply context, inspect outputs, create tests, manage agents, and recognize when conventional methods are more appropriate. Testing and guardrails often lag, because AI can generate output faster than a team can safely review it, and code, analysis, and agent actions still require evaluation suites, permissions, human review, monitoring, and rollback, especially when an agent can access production data, contact customers, modify code, or initiate financial workflows.

Unrealistic expectations and weak scoping are a recurring hazard: because AI can complete some tasks in minutes, leaders may assume every feature should be finished immediately, and measuring progress by lines generated rather than customer value can make productivity illusory. Security and governance gaps grow as nontechnical employees build dashboards and automations and interact directly with code, APIs, and sensitive data; one company reported an employee accidentally committing an API key to a shared repository. Finally, reliability and context limitations persist: general-purpose models can produce plausible but incorrect answers when they lack proprietary data or domain context, and one company found general models achieved only roughly 50 to 60% accuracy on sports footage, acceptable for a casual summary but not where precise statistics matter.

Workforce Implications

AI is changing job design faster than it is eliminating entire functions, shifting engineering work from manual production toward supervision, system design, validation, and exception handling: defining requirements and architecture, clarifying work for agents, creating harnesses and tool permissions, writing tests and evaluation suites, reviewing AI-generated changes, monitoring production behavior, and managing technical debt. The roundtable included examples that affected or deferred hiring, such as a postponed QA hire, automated marketing production, and reduced need for multi-geography specialists. Human roles remain necessary where judgment, accountability, and domain expertise matter, but fewer people may be required for repetitive production work, and remaining roles carry higher expectations for AI fluency.

One technical leader said AI was reducing the importance of hiring for a specific language or framework, letting the company weigh systems thinking, problem solving, architecture, and the ability to supervise AI-supported development. That could broaden the talent pool but raises questions about the entry-level pipeline, since junior professionals traditionally gain expertise by performing routine work; if AI absorbs that work, startups will need new apprenticeship models that teach judgment, validation, security, and systems design. Adoption is ultimately a

management and training challenge, and outside experts should build internal capability rather than create long-term dependence on consultants.

Data Readiness

Data readiness varies widely, but the strongest startups treat proprietary data and context as central competitive assets. Usable AI depends on more than collecting data; companies need clean and well-structured records, consistent identifiers, reliable permissions, source attribution and provenance, current documentation, defined access boundaries, feedback loops and evaluation datasets, and human review for high-consequence information. Technical debt remains a significant obstacle: one participant described a case in which an AI coding tool recommended matching records by stable identifiers rather than text, where the underlying issue was not model intelligence but an inherited data and software-design compromise. The sector's leading companies are increasingly creating knowledge layers over fragmented systems, but governance maturity is uneven and many startups remain dependent on informal documentation, founder knowledge, and rapidly changing codebases.

AI Safety and Risk

AI-safety maturity is mixed: technical awareness is relatively high, but formal governance often lags the pace of experimentation. The most mature practices discussed include human review of AI-maintained data, read-only database connections for customer-facing assistants, smoke/unit/UI/regression tests, evaluation suites built from historical examples, model-vendor portability testing, agent instruction files and behavioral rules, high-confidence thresholds before autonomous action, incremental permissioning, and continued human testing of releases. Material gaps remain, including credential exposure by inexperienced builders, AI-generated code expanding faster than teams can audit it, agents gaining broad access to internal systems, model upgrades silently changing behavior, premature deployment, fragile AI-generated changes reaching existing customers, and reputational risk from automated external communications. Governance should focus on what an AI system can access and do, not merely which model it uses, and the sector would benefit from lightweight, startup-appropriate controls: an inventory of AI systems, named owners, access reviews, secrets management, logging, evaluation requirements, human-approval thresholds, incident procedures, and the ability to disable or roll back an agent quickly.

Ethics

The roundtable did not include an extended standalone ethics discussion, but several sector-specific tensions emerged. On productivity versus worker displacement, AI has let startups defer hiring and automate work, creating real efficiency but also raising questions about worker transition and whether gains are shared through better jobs, training, or opportunity; startups should distinguish eliminating low-value manual work from eliminating the pathways through which people develop expertise. On transparency and accountability, when an agent submits code, disputes a charge, summarizes a legal filing, or contacts a journalist, responsibility still rests with the company, so startups need clarity about who reviews the action, who can override it, and how affected people can contest errors. On access to sensitive and proprietary data, agents become more capable as they gain access to customer records, databases, and code, which increases privacy and security risk, so organizations should apply data-minimization and least-privilege principles. And on accuracy in high-consequence contexts, a 50 to 60% accurate

sports-video summary may be entertaining, but inaccurate statistics or an incomplete SEC-filing summary could influence real decisions, so companies must define acceptable accuracy by the consequence of error, not by whether the output appears impressive.

Impact Story

New Hampshire Product Rebuild (speed, scope, and fragility): a New Hampshire startup undertook a major AI-assisted rebuild of its online product after learning of another application that generated roughly 30,000 lines of code in two months, comparable work estimated at about 12 months before AI. Although the founder was not a software engineer, they used Claude to draft a nine-page product-requirements document, which the engineering team expanded into a roughly 40-page technical specification and provided to an AI coding workflow; by the next morning the system had generated an initial 6,000 lines of code. The initial result was promising but unstable and required substantial human refinement, and the scope expanded toward roughly 150,000 lines. The case captures both sides of AI-assisted development: dramatic acceleration that let a nontechnical founder participate directly in product creation, alongside scope growth and pressure on validation, architecture, and expectations, with real operational consequences because the company already had customers.

SportsTech startup (specialized AI outperforms general models): a New Hampshire sports-technology company uses proprietary computer vision to turn ordinary sports video into statistics, analytics, and highlights, and lets customers ask natural-language questions about their data through a read-only AI connection. It reported substantial internal marketing-efficiency gains and uses AI across software, finance, operations, accounting, sales, and research. Its experiments also revealed a limit: general-purpose models produced only roughly 50 to 60% accuracy on raw sports video, insufficient for reliable statistics, so the company continues to rely on proprietary models, specialized training data, testing, and human oversight.

Autonomous QA (deferring a hire without eliminating oversight): another startup uses a cloud coding agent (Cursor) to reproduce bugs logged in Slack, create tickets, generate patch fixes, test the application, and verify results, which let the company postpone hiring a dedicated QA employee at its current stage. The company did not eliminate human review; a person still tests every release, and the team maintains a regression suite of roughly 50 historical scenarios, illustrating a likely near-term pattern in which AI reduces routine QA workload while humans remain responsible for release confidence and edge cases.

Odin (automating work too small for the engineering roadmap): a startup's OpenClaw agent, Odin, handles roughly 40 recurring jobs and can investigate requests directly from Slack; one workflow reviews customer data and prepares high-confidence Stripe disputes, addressing roughly \$500 per month in losses. That amount may be too small to justify a traditional engineering project but is large enough to matter to a resource-constrained startup, and Odin changes the economics by letting the company automate first and formalize the workflow later only if necessary.

Bad and ugly outcomes: participants also reported failures and warning signs, including AI-generated outbound sales messages that failed because recipients were overwhelmed by automated outreach, an inexperienced employee exposing an API key in a code repository, AI-generated product code that became fragile as project scope expanded, leaders developing unrealistic expectations that every feature should be completed immediately, and small software and messaging edge cases that created outsized customer impact in operations coordinating thousands of people across locations and time zones.

What This Sector Needs Most

- More affordable compute and model access: participants identified cost and usage capacity as immediate constraints and pointed to shared compute resources, startup credits, university partnerships, and pooled purchasing or infrastructure programs.
- Lightweight governance templates: practical controls that improve safety without the full compliance burden of a large enterprise, including AI-system inventory templates, data-classification and acceptable-use policies, agent-permission and access checklists, model and vendor risk assessments, human-review thresholds, evaluation standards, incident-response playbooks, secure-coding and credential-management guidance, and customer disclosure and transparency templates.
- Shared evaluation and benchmarking support: many startups know they should evaluate AI systems but lack the time or expertise to build rigorous test sets; shared technical assistance could help create domain-specific evaluations, measure hallucination and failure rates, test model upgrades, and establish production-readiness thresholds.
- Peer learning and expert intervention: the roundtable itself demonstrated the value of founder-to-founder learning, and New Hampshire could support recurring technical roundtables, implementation workshops, shared case studies, and short expert engagements.
- Better access to proprietary and public data: clearer pathways to appropriate public datasets, shared data standards, secure research environments, and guidance on data provenance and responsible reuse.

New Hampshire's leading startups are not waiting for AI adoption to become conventional; the next phase will be defined by connected agents, proprietary data, AI harnesses, software factories, and systems that take action under defined constraints. For the startup sector, the central policy and ecosystem challenge is not persuading companies to experiment but helping them build the skills, infrastructure, safeguards, and talent pipelines needed to convert rapid experimentation into durable and responsible economic value.

Part Three: Recommendations

Chapters 20–24

Chapter 20. A Framework for Action

This part turns the report's findings into action. It is organized by audience so each group can find its asks quickly, and every recommendation traces back to something the report actually heard, not to outside opinion. Like the chapters that follow, it is a starting point for the Task Force to react to, revise, and own, not a settled position.

How to Read the Recommendations

Part Three is four chapters for four audiences: State Government (Chapter 21), New Hampshire Businesses and Institutions (Chapter 22), Educators and Workforce Development Partners (Chapter 23), and the Task Force itself (Chapter 24). Many recommendations touch more than one audience; each is placed with the group best positioned to act on it, with cross-references where a second group has a role.

One theme runs through all four. The state's highest-leverage role is to remove the friction that keeps responsible adoption from spreading, through shared guidance and templates, a common framework, and, for the public sector, a workable procurement pathway, rather than through broad new regulation. The single clear exception, where New Hampshire can and should lead now, is protecting minors.

Tiered by Audience

- **State Government:** the enablers only the state can provide, including a standardized procurement pathway, a short foundational statute, agency coordination, and leadership on child safety.
- **Businesses and Institutions:** practical tools and right-sized governance, including the AI Implementation Checklist, the AI Self-Assessment Tool, and treat-it-like-a-new-hire controls scaled to organization size.
- **Educators and Workforce Development Partners:** a coordinated statewide literacy and fluency framework, and employer-education partnerships that protect the talent pipeline.
- **The NH AI Task Force:** turning the recurring cross-sector asks into shared assets, and standing up as a durable, non-partisan resource for the state.

Prioritization Criteria

Recommendations are meant to be sorted, not adopted all at once. Four criteria help the Task Force and its audiences decide what to do first:

- **Time horizon.** Some moves are realistic within twelve months: a standardized procurement pathway and short policy template, model policies and a maturity self-assessment, child-safety guidance building on HB 143 and SB 300, data-boundary guidance, and hands-on policy workshops. Others are longer-horizon, such as a foundational acceptable-use statute moving through the legislative process, new apprenticeship models, and maturing the annual survey-and-roundtable cycle.
- **Influence.** Prioritize what New Hampshire can actually control, its own procurement, guidance, and convening power, over what depends on federal action or the unsettled question of preemption.

- **Reach.** Favor shared enablers that help the least-resourced organizations, small towns, small businesses, and rural providers, because that is where the adoption gap is widest.
- **Agreement and representation.** Weight actions that reflect cross-sector consensus and that close the Task Force's own representation gap, so the recommendations carry the authority of the whole group.

A note on sequencing: with federal preemption unsettled and the 2026 legislative session having produced no framework, the report favors leading with defensible, high-agreement moves, child safety, shared guidance and templates, and public-sector procurement, while staging the broader statutory work for when the legal landscape is clearer.

Chapter 21. For State Government

The findings point to a state role that is less about regulating AI and more about removing the friction that keeps responsible adoption from spreading. The highest-leverage moves are ones New Hampshire can realistically make in the next twelve months, and they recur, independently, across sectors that otherwise share little.

For the Governor's Office

- **Endorse a common statewide framework.** Adopt and promote the NIST AI Risk Management Framework as New Hampshire's shared reference, so a ten-person firm and a state agency can hold the same conversation at different depths.
- **Direct a standardized procurement pathway.** Task the relevant offices, working with the NH Municipal Association, to build a standardized AI procurement pathway and a cooperative-purchasing option, so enterprise pricing and purchasing rules stop excluding municipalities under 150 employees and small businesses. Procurement is the public sector's most consistently named barrier, and the lever the state is best positioned to move.
- **Make protecting minors the flagship.** Child safety is the area of clearest state authority. Direct a coordinated minors-and-AI standard that builds on HB 143 and SB 300, covering safe classroom use, age-appropriate design for youth-facing tools, and fast response to image abuse.

Legislative Priorities

- **A short, foundational acceptable-use statute.** Advance a concise bill, on the order of three pages, modeled on the EU AI Act's tiered approach and distinguishing acceptable, high-risk, and prohibited uses, rather than a comprehensive statute.
- **Vendor transparency requirements.** Require vendors to disclose when AI is used, what data is collected, and where it is stored, a request that recurred independently in Healthcare and Real Estate.
- **Right-to-Know clarity for AI.** Clarify how existing public-records and retention law applies to AI prompts, outputs, transcripts, and logs, the barrier the public sector named first.
- **Child-safety guardrails.** Build on the existing NH deepfake statutes, HB 143, and SB 300 with age-appropriate-design and image-abuse-response provisions aligned to the federal TAKE IT DOWN Act.

Context for sequencing: New Hampshire's 2026 session produced no framework (HB 1725 and SB 657 did not become law), and federal preemption is unsettled (Executive Orders 14365 and 14409 and the DOJ AI Litigation Task Force). That argues for narrow, defensible action now, child safety above all, over a broad statute that may be preempted or contested.

State-Agency Guidance and Coordination

- **Publish model policies and a maturity self-assessment.** Provide model acceptable-use policies and a short self-scoring maturity assessment (Strategy, Governance, Data and Technology, Talent, Execution) that any agency or organization can use without a consultant.

- **Establish a coordination point.** The government is already creating senior AI roles; a state-level coordinating function would help agencies share policies, vendor evaluations, and lessons rather than each solving the same problems alone.
- **Offer a regulatory sandbox or safe harbor.** Give heavily regulated sectors such as finance a supervised way to move beyond low-risk pilots.
- **Issue plain data-boundary guidance.** State clearly what non-public data may and may not be entered into a model, the gap named in Real Estate and Small Business.

Most of these steps can be taken through guidance and coordination, without waiting on new statute, and most are realistic within twelve months. Together they would give agencies and municipalities a shared starting point, spare each office from writing its own policy and vetting the same vendors alone, and give the state a clearer picture of what is working. That groundwork also puts the legislative priorities above on firmer footing once the federal picture settles.

Chapter 22. For NH Businesses and Institutions

The report's clearest business-facing message is that most of what organizations need is practical and shared, not exotic. The two named tools below, paired with right-sized governance, let an organization start without a large build-out.

Two Practitioner Tools

1. **The AI Implementation Checklist.** A short plain-language checklist of what to address when implementing AI, plus categories of service providers to seek out.
2. **The AI Self-Assessment Tool.** A four-phase, roughly twelve-question diagnostic an organization can use to gauge where it is and what to focus on next.

Right-Size the Effort

- **Small business:** an AI inventory, an acceptable-use policy, a short list of vendor questions, AI labeling, and basic logging.
- **Mid-size:** add red-teaming, model or system cards, and an incident playbook that names AI failures.
- **Enterprise:** a full AI management system, ISO/IEC 42001 certification where buyers expect it, and dedicated AI security testing.

Governance and Security Essentials

- **Treat any new tool like a new hire.** Trust is earned; require human approval for business-critical decisions, and minimize what AI tooling can reach, reviewing access regularly.
- **Inventory everything, including shadow AI.** Keep an inventory of every AI system in use, assign each an owner, and extend the incident-response plan to cover AI-specific failures such as prompt injection or a poisoned model.
- **Do not route data around your own controls.** Get plain security do-and-do-not guidance on what should never be uploaded, so staff stop building risky workarounds.

Where to Start

- **Use a role- and function-based catalog of use cases** (finance, marketing, client relationships) rather than testing tools at random.
- **Invest in job-function-specific fluency,** not generic prompt engineering, and pair AI-generated work with an experienced reviewer.
- **Vet AI advisors specifically.** Today's IT or security partner may not be the right AI partner, and starting with the wrong one is expensive to unwind.

Chapter 23. For Educators and Workforce Development Partners

Education is both a sector adopting AI and the pipeline every other sector draws from, so its recommendations carry double weight. The through-line is coordination: the raw activity exists, but it is uneven and uncoordinated.

K-12 and Higher Education

- **Build a coordinated statewide AI-literacy and professional-development framework** that can be adapted locally. This was the education sector's single strongest ask.
- **Move from AI literacy to AI fluency.** Use approaches such as assignment-level AI continuums and authentic assessment, so students learn to collaborate with AI rather than either avoid it or over-rely on it.
- **Streamline privacy-compliant vetting.** Provide FERPA- and COPPA-aligned vendor-vetting support so K-12 adoption is not delayed for months.
- **Close the access and equity gap.** Address the divide between institutions with strong AI programs and those that restrict access, and between paid and free tools, so students do not graduate with sharply different readiness.
- **Adopt safe classroom-use guidance.** Let teachers use AI under supervision with clear rules on student data, disclosure, and acceptable use, consistent with the minors-first agenda.

Workforce Development and Employer-Education Partnerships

- **Fix the curriculum-employer stalemate.** Create structured feedback loops so schools teach what AI-era employers actually need and employers give usable input.
- **Build new apprenticeship models.** Protect the junior pipeline in knowledge-work sectors such as legal, technology, and startups, where AI is absorbing the entry-level tasks through which people used to build expertise.
- **Deliver hands-on, job-specific training,** not generic or expensive consultant classes, using train-the-trainer and two-tier models like those working in finance and the SBDC.
- **Leverage higher-education partnerships and communities of practice** to share applied research, training, and real implementation examples across manufacturing, technology, and beyond.

Chapter 24. For the NH AI Task Force

The report is a milestone, not the destination. The Task Force's own next steps are what turn a snapshot into a standing capability for the state.

What the Task Force Commits To

- **Turn the recurring asks into shared assets.** Produce the assets the report calls for most: shared policy templates, a maturity self-assessment, the AI Implementation Checklist, starting points for an organizational AI ethics statement, and a standardized procurement pathway for the public sector.
- **Champion the minors-first agenda** with the Governor's Office and the Legislature, as the area where a practical New Hampshire standard would travel furthest.
- **Close the representation gap** by adding Retail and Hospitality, Real Estate, and community-advocacy voices to the Task Force itself.

Ongoing Impact-Story Tracking

- **Maintain a living record of real New Hampshire AI stories,** the good, the bad, and the ugly, as a recurring feature that outlives this report and keeps the findings concrete and current.

Next Report Cycle and Survey Cadence

- **Run the listen, synthesize, translate, and track cycle on an annual cadence.** Strengthen the evidence base with a second Consumer and Small Business roundtable, a direct question on vendor reliance (the one February signal still under-evidenced), and broader sector coverage.

Positioning as a Standing Resource

- **Be the trusted, non-partisan bridge** between policymakers who often do not use the technology they govern and technologists who often do not see the communities affected, and serve as a standing sounding board for state government and the New Hampshire business community.

Appendices

Tools, references, and acknowledgments

Appendix A. Task Force Members & Working Group Structure

This report reflects the work of volunteers from across New Hampshire's business, government, academic, and nonprofit communities. Members led sector working groups, convened roundtables, and drafted chapters and practitioner tools. The tables below show who contributed to each part of the report.

Leadership

Name	Organization	Role
Phil Magnuszewski	Infused Innovations	Chair; Executive Committee
Julie Demers	NH Tech Alliance	Executive Committee
Scott Merrick	Darby Field Advisors	Executive Committee

Sector Working Groups

Sector	Lead(s)	Organization
Technology & Software	Katie Destramp	Fastly
Financial Services & Insurance	Scott Weller / Jim Kisch	EnFi / Passumpsic Bank
Healthcare & Life Sciences	Maggie Allard	Community Health Design Corps
Manufacturing, Engineering & Defense	Brian Stumm	Hypertherm
Professional & Business Services	Brian Bouchard / Cameron Shilling / Chris Norwood	Sheehan Phinney / McLane Middleton / NAI Norwood Group
Consumer & Small Business	Liz Gray	NH Small Business Development Center
Public Sector & Education	Shawn Powers	Southern New Hampshire University
	Melanie McDonough	City of Lebanon
	Shaun Mulholland	Town of Londonderry
Startups	Jack Potvin	Business Automation Systems

Cross-Sector Chapters and Practitioner Tools

Area	Contributor(s)	Organization
Adoption: People, Process, and Change	Alexis Brochu	Upstart-Labs
AI Safety & Responsible Use	Dave Gerry	Bugcrowd
Legislative Monitoring	DJ Burke	BIA

Program Support

Name	Organization	Role
Cam Chicko	NH Tech Alliance / UNH Student	Operations, scheduling, and outreach tracking

Past Contributors

The Task Force thanks members who contributed earlier in the 2026 cycle.

Name	Organization	Contribution
Jason Alexander	ChiefAI / Benson Consulting	Co-Chair through mid-2026
David Smith	New Balance	Consumer & Small Business sector lead
Shawn Stapp	Elliot Health System	Healthcare

Appendix B. AI Frameworks & Standards Reference Guide

This appendix is a short, plain-language reference for the frameworks and standards the report relies on, drawn from Chapter 6 and the About This Report section. Each entry says what the framework is, what it covers, and how this report uses it. None requires adoption in full; they are reference points an organization can draw on at the depth of its size and risk warrant.

NIST AI Risk Management Framework (AI RMF 1.0)

Published by the U.S. National Institute of Standards and Technology in January 2023. A voluntary framework for identifying and managing AI risk, built around four functions, Govern, Map, Measure, and Manage, and seven characteristics of trustworthy AI: valid and reliable; safe; secure and resilient; accountable and transparent; explainable and interpretable; privacy-enhanced; and fair, with harmful bias managed. The functions are designed to repeat as conditions change rather than run once, and they apply whether the system is a traditional classifier, a third-party model, or a generative tool.

This report treats the AI RMF as its common reference because it is voluntary, widely recognized, scales from a small business to a state agency, and maps onto requirements taking shape in the EU and elsewhere, which lowers the cost of meeting several regimes at once.

NIST Generative AI Profile (NIST AI 600-1)

Published by NIST in July 2024. A companion to the AI RMF that addresses risks specific to or amplified by generative AI. It defines twelve risk categories, spanning fabricated output, harmful content, data leakage, and supply-chain exposure, and lists more than two hundred suggested actions mapped to the RMF's four functions. An organization already using the AI RMF can adopt the profile without starting over.

ISO/IEC 42001:2023, AI Management System

Published by ISO and IEC in 2023. The first international management-system standard for AI. It treats data governance, adversarial testing, runtime monitoring, and human oversight as one system rather than separate workstreams, and it is certifiable, which lets an organization signal diligence to enterprise buyers and regulators. It complements the AI RMF: the RMF is the risk operating model, and 42001 is the management system that runs it.

ISO/IEC 27001, Information Security Management

Published by ISO and IEC. The widely adopted international standard for information-security management systems. It is not AI-specific, but because most AI risk rides on top of data and systems security, an organization that already runs a 27001-aligned security program has much of the foundation AI governance needs and does not have to build a parallel one.

OWASP Top 10 for LLM Applications (2025)

Maintained by the Open Worldwide Application Security Project and updated for 2025. A ranking of the most common security risks in deployed large-language-model applications, with prompt injection, sensitive-information disclosure, supply-chain weaknesses, and data and model poisoning among the leading entries. It is written for the people building and shipping AI applications and serves as the practical checklist for securing them.

MITRE ATLAS

Maintained by the MITRE Corporation. A public knowledge base of real adversary tactics and techniques against AI systems, modeled on the widely used MITRE ATT&CK framework. It is the reference for red-teaming and threat-modeling AI, and many of its mitigations map to security controls teams already run, so much of AI security is extending existing practice rather than starting over. NIST's adversarial-machine-learning taxonomy supplies the formal vocabulary underneath it.

Guidelines for Secure AI System Development

Published by U.S. CISA and the UK NCSC with 21 other agencies in November 2023, the first international agreement of its kind on building AI securely. It organizes security across four stages, secure design, secure development, secure deployment, and secure operation and maintenance, and advances a secure-by-design principle: own security outcomes for the customer and build them in from the start rather than bolt them on later.

How These Fit Together

These standards are complementary, not competing. The NIST AI RMF is the risk operating model and the common language. ISO/IEC 42001 is the certifiable management system that runs it, with ISO/IEC 27001 covering the information-security foundation beneath it. The OWASP Top 10 for LLM Applications and MITRE ATLAS supply the threat references for securing AI, and the CISA and NCSC guidelines set the secure-by-design expectations for how AI is built. An organization does not need all of them at once; the report's recommendation is to anchor on the NIST AI RMF and add the others as scale, risk, and buyer or regulatory expectations require.

Appendix C. Glossary

Definitions of the technical, framework, and regulatory terms used in this report. Where the report uses a term in a specific sense, that sense is noted.

42 CFR Part 2: A federal regulation giving heightened confidentiality protection to substance use disorder treatment records, stricter than HIPAA in several respects. Referenced in behavioral-health AI pilots.

Adoption scale (Awareness, Active, Operational, Systemic, Transformational): The five-stage scale this report uses to place a sector, running from talking but not acting (Awareness), to funded pilots (Active), to live in a workflow (Operational), to embedded across most workflows (Systemic), to reshaping the business model (Transformational).

Agentic AI (AI agent): An AI system that can take actions on its own, call tools, and chain multiple steps toward a goal rather than only answering a single prompt.

AI coworker: A persistent AI agent connected to an organization's systems, such as chat, documents, databases, and code, that can investigate requests and carry out work rather than only respond to questions.

AI harness: The surrounding infrastructure used to operate and control a model or agent, managing prompts, tools, context, permissions, retries, tests, evaluations, and operating rules.

AI safety: Whether an AI system behaves as intended and avoids unintended harm to users and the public. Distinct from AI security.

AI security: Whether an AI system, and the data and infrastructure behind it, can withstand a deliberate adversary.

Ambient clinical documentation (ambient listening): AI that listens to a clinical encounter and drafts the visit note, reducing the time clinicians spend charting.

Anonymization and re-identification: Removing identifying details from data (anonymization), and the risk that a model can reconstruct who or what the data refers to (re-identification), even from a scrubbed record.

Business Associate Agreement (BAA): A HIPAA-required contract that binds a vendor handling protected health information to safeguard it.

CAD: In healthcare, computer-aided detection, meaning AI that flags possible abnormalities in medical images such as mammograms. In manufacturing and engineering, computer-aided design, the software used to model parts and products.

Comparative market analysis (CMA): A real estate valuation of a property based on comparable recent sales and local conditions.

COPPA (Children's Online Privacy Protection Act): A federal law restricting the online collection of personal information from children under 13, relevant to K-12 education tools.

Data governance: The policies, ownership, and controls that determine how an organization's data is collected, stored, accessed, and used.

Data lineage (provenance): A record of where data came from, whether the organization had the right to use it, and how it was transformed.

Deepfake: Synthetic image, audio, or video generated or altered by AI to depict a real person doing or saying something they did not.

Drift (model drift): The gradual decline in a model's accuracy as real-world conditions change from those it was built on.

Electronic health record and electronic medical record (EHR / EMR): The digital system of record for patient information in a healthcare organization.

Enterprise resource planning (ERP): Core business software that integrates functions such as finance, inventory, and operations.

Errors and omissions (E&O) insurance: Professional liability coverage for mistakes or omissions in services; the report notes AI use is beginning to affect E&O exposure in real estate.

EU AI Act: The European Union's comprehensive AI regulation, which classifies systems by risk and sets obligations for high-risk uses. Its high-risk obligations apply from December 2, 2027 under the Digital Omnibus.

Fair housing law: Federal and state law prohibiting discrimination in housing; the report flags that AI-written listing remarks can carry noncompliant language under an agent's name.

FERPA (Family Educational Rights and Privacy Act): A federal law protecting the privacy of student education records, relevant to education AI tools.

Generative AI: AI that produces new content such as text, images, audio, or code, typically built on large language or diffusion models.

Hallucination: A confident but false or fabricated output from an AI system.

HIPAA (Health Insurance Portability and Accountability Act): The federal law governing the privacy and security of protected health information.

HITRUST: A widely used security-and-privacy certification framework that healthcare organizations often require of vendors handling sensitive data.

High reliability organization (HRO): An organization that sustains near-error-free operations in high-risk conditions through disciplined process and oversight.

Human-in-the-loop: A control requiring a person to review, approve, override, or stop an AI system on consequential decisions.

ISO/IEC 27001: The international standard for information security management systems.

ISO/IEC 42001: The first international standard for AI management systems, treating governance, data, testing, monitoring, and oversight as one system. It is certifiable.

Large language model (LLM): An AI model trained on large text datasets that generates and interprets natural language.

Machine-learning bill of materials (ML-BOM): An inventory of the components behind an AI system, such as data, models, and dependencies, used to track and secure the supply chain. See also SBOM.

Manufacturing execution system (MES): Software that tracks and controls production on the factory floor.

MITRE ATLAS: A public knowledge base of real adversary tactics and techniques against AI systems, modeled on MITRE ATT&CK, used for red-teaming and threat modeling.

Model Context Protocol (MCP): An open standard that lets an AI assistant connect to external data and tools through defined, permissioned interfaces.

Model card / system card: Documentation that records what a model does, its limits, and its tested behavior.

Model risk management (MRM): The banking discipline for validating and governing models used in decisions. The report notes it was built for static, deterministic models and strains against dynamic AI.

NIST AI Risk Management Framework (AI RMF): A voluntary US framework for managing AI risk, organized around four functions, Govern, Map, Measure, and Manage, plus seven characteristics of trustworthy AI. The report's recommended common anchor.

NIST Generative AI Profile: A 2024 NIST companion to the AI RMF defining risk categories and suggested actions specific to generative AI.

OWASP Top 10 for LLM Applications: A community-maintained ranking of the most common security risks in deployed LLM applications, such as prompt injection and data poisoning.

Product lifecycle management (PLM): Software that manages a product's data and process from design through manufacturing and service. Examples include Windchill and ARAS.

Production Part Approval Process (PPAP): A standardized manufacturing quality process for approving parts before production.

Prompt engineering: The skill of writing effective instructions and context to get reliable output from an AI system.

Prompt injection: An attack that hides malicious instructions in text an AI reads, causing it to ignore its rules or act against the user's intent.

Protected health information (PHI): Individually identifiable health information protected under HIPAA.

Red-teaming: Deliberately testing a system by trying to make it fail or misbehave, to find weaknesses before deployment.

Retrieval-augmented generation (RAG): A technique that grounds an AI's answers in an organization's own documents or data retrieved at the time of the query.

Safety stock: Extra inventory held to buffer against demand or supply variability. The report notes AI-driven forecasting can reduce it.

Secure by design: Building security into a system from the start rather than adding it later; the basis of the CISA and NCSC guidelines.

Shadow AI: Employees using personal or unsanctioned AI accounts for work, outside IT oversight and any inventory.

SOC 2 Type II: An independent audit report on how well an organization's data-security controls operate over a period of time; often required of vendors.

Software bill of materials (SBOM): An inventory of the components in a piece of software, used to track and secure the supply chain.

Software factory: An end-to-end system for producing software repeatedly and systematically, potentially combining coding agents, testing, CI/CD, deployment, monitoring, and governance.

Structured authoring: Writing content in a structured format such as XML/DITA so it can be reused and made machine- and AI-ready. The report cites large translation-cost reductions from it.

TAKE IT DOWN Act: A 2025 federal law criminalizing nonconsensual intimate images, including AI deepfakes, and requiring covered platforms to remove reported images within 48 hours.

UDAP (unfair or deceptive acts or practices): The legal standard prohibiting unfair or deceptive conduct toward consumers; a compliance concern for AI in finance.

Vibe coding: Building software largely by describing what you want to an AI and iterating on what it generates, rather than writing most of the code by hand.

Virtual staging: Digitally furnishing or enhancing property photos. The report flags AI-enhanced staging as a disclosure and misrepresentation risk in real estate.

Appendix D. Chapter References

Any sources that have been referenced in specific chapters are noted below.

Sources (Chapter 6)

1. NIST, AI Risk Management Framework (AI RMF 1.0), January 2023.
<https://www.nist.gov/itl/ai-risk-management-framework>
2. NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1), July 2024. <https://doi.org/10.6028/NIST.AI.600-1>
3. NIST, Cybersecurity Framework Profile for AI (NIST IR 8596, preliminary draft, December 2025) and SP 800-53 Control Overlays for Securing AI Systems (COSaIS, in development). <https://csrc.nist.gov/pubs/ir/8596/iprd> and <https://csrc.nist.gov/projects/cosais>
4. ISO/IEC 42001:2023, Artificial Intelligence Management System (AIMS). <https://www.iso.org/standard/42001>
5. OWASP, Top 10 for LLM Applications 2025.
<https://owasp.org/www-project-top-10-for-large-language-model-applications/>
6. MITRE ATLAS (Adversarial Threat Landscape for Artificial-Intelligence Systems). <https://atlas.mitre.org>
7. CISA and UK NCSC, Guidelines for Secure AI System Development, November 2023.
<https://www.cisa.gov/news-events/news/dhs-cisa-and-uk-ncsc-release-joint-guidelines-secure-ai-system-development>
8. EU AI Act, Regulation (EU) 2024/1689, as amended by the Digital Omnibus on AI (published in the Official Journal July 24, 2026, in force July 27, 2026); high-risk (Annex III) obligations postponed from August 2, 2026 to December 2, 2027; new Article 5 prohibition on AI-generated nonconsensual intimate imagery and child sexual abuse material. <https://artificialintelligenceact.eu>
9. Executive Order 14365 (December 11, 2025), Ensuring a National Policy Framework for Artificial Intelligence, establishing the DOJ AI Litigation Task Force; White House National Policy Framework for Artificial Intelligence, March 20, 2026 (nonbinding legislative recommendations); Executive Order 14409 (June 2, 2026), Promoting Advanced Artificial Intelligence Innovation and Security.
<https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>
10. Colorado SB 26-189 (signed May 14, 2026), repealing and replacing the 2024 Colorado AI Act with a narrower automated decision-making disclosure law effective January 1, 2027; California Transparency in Frontier Artificial Intelligence Act (effective January 1, 2026).
11. TAKE IT DOWN Act, Pub. L. No. 119-12, signed May 19, 2025; covered-platform notice-and-removal obligations enforceable from May 19, 2026, with FTC enforcement and complaint portal at TakelDown.ftc.gov. Congressional Research Service summary: <https://www.congress.gov/crs-product/LSB11314> and FTC press release, May 19, 2026:
<https://www.ftc.gov/news-events/news/press-releases/2026/05/ftc-begins-enforcing-take-it-down-act>
12. California SB 243, Companion Chatbots (effective January 1, 2026).
13. New York AI Companion Models statute, N.Y. Gen. Bus. Law sec. 1700 et seq. (effective November 5, 2025); Utah and Maine AI chatbot disclosure laws (2025).
14. New Hampshire RSA 507:8-j and RSA 638:26-a, fraudulent use of deepfakes (effective January 1, 2025); HB 143 (Chapter 270, Laws of 2025), AI chatbots and child endangerment (effective January 1, 2026); SB 300 (2025), AI-generated intimate visual representations of children (effective January 1, 2026); NH Department of Information Technology, Use of Artificial Intelligence Technologies Policy.
15. New Hampshire Tech Alliance, New Hampshire AI Task Force.
<https://nhtechalliance.org/new-hampshire-ai-task-force/>
16. New Hampshire General Court, 2026 session: HB 1725 (regulation of AI technologies, Inexpedient to Legislate, February 19, 2026) and SB 657 (as amended by the House, use of information technology and AI

systems by state agencies; did not become law). <https://legiscan.com/NH/text/HB1725/id/3292332> and https://gc.nh.gov/bill_status/billinfo.aspx?id=2483

Sources (Chapter 13)

AI in Healthcare at Southern New Hampshire Health. Patch, 8 Apr. 2026.
<https://patch.com/new-hampshire/nashua/ai-healthcare-southern-new-hampshire-health-advancing-safe-patient-centered>

Bruce, Giles. How Dartmouth Health Built AI to Respond to Patients. Becker's Hospital Review, 26 June 2025.
<https://www.beckershospitalreview.com/healthcare-information-technology/ai/how-dartmouth-health-built-ai-to-respond-to-patients/>

Embracing Innovation: Elliot Hospital Leverages AI to Enhance Patient Care. Elliot Hospital, 2 Oct. 2024.
<https://www.elliethospital.org/about-us/newsroom/news/elliott-hospital-leverages-ai-enhance-patient-care>

Leading the Way: Cottage Hospital's Vision for Rural Healthcare. Cottage Hospital, 3 Jan. 2025.
<https://www.cottagehospital.org/blog/leading-the-way-cottage-hospital-s-vision-for-rural-healthcare>

Sources (Chapter 14)

- [1] Cisco, 2026 State of Industrial AI Report (manufacturing sector) and related coverage of cybersecurity and IT-OT barriers.
- [2] Fictiv, 2026 State of Manufacturing Report.
- [3] Parsec, The 2026 State of the Manufacturing Industry.
- [4] Deloitte, The State of AI in the Enterprise (physical AI in manufacturing, logistics, and defense).
- [5] Digit Software, analysis of U.S. Census Bureau data on AI adoption in U.S. manufacturing.
- [6] Wisconsin Department of Workforce Development, Governor's Task Force on Workforce and Artificial Intelligence.
- [7] Mississippi AI Network, AI Workforce Readiness Council.
- [8] Boston Consulting Group, Are US States Ready for the AI Economy? The AI Maturity Matrix.
- [9] OECD, AI adoption by small and medium-sized enterprises.
- [10] National Governors Association, AI and the Future of Work (state AI task force patterns).
- [11] ERP Today, Manufacturers Struggle With AI Readiness Despite Widespread Exploration.

Sources (Chapter 15)

ABA Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512: Generative Artificial Intelligence Tools (July 29, 2024).
https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/aba-formal-opinion-512.pdf

New Hampshire Rules of Professional Conduct,
<https://www.courts.nh.gov/sites/g/files/ehbemt471/files/inline-documents/sonh/nhrpc.pdf>

Thomson Reuters, 2025 Generative AI in Professional Services Report.
<https://legal.thomsonreuters.com/en/insights/reports/2025-generative-ai-in-professional-services>

SHRM, Navigating AI in the Workplace 2026,
<https://www.shrm.org/topics-tools/topics/artificial-intelligence-in-the-workplace>