WHITE PAPER

Risk and Governance Strategies for AI in Healthcare

Navigating regulatory volatility

Caleb Pfanstiel, GRC Healthcare Consultant,
MBA | SECURITY+

COALFIRE®

Table of Contents

AI Governance is a Competitive Advantage		**Page 02**
Case Study: HHS AI Strategic Plan Overview	Organization of the Plan	**Page 03**
	Goals of the Plan	
	Promoting Ethical and Responsible Use of AI	
	What's Next?	
AI Governance in Healthcare	Defining AI Governance in Healthcare	**Page 07**
	AI Risk Management Frameworks	
	Algorithmic Impact Assessments	
	Ensuring Transparency, Accountability, and Security	
Risk Mitigation Strategies for AI in Healthcare	Identifying AI Risks	**Page 10**
	Proactive Risk Management Approaches	
	AI Testing and Validation	
	Incident Response Preparedness	
	Embedding Risk Mitigation into AI Governance	
Regulatory Compliance and Ethical Considerations	Existing Regulatory Landscape	**Page 13**
	Aligning AI Governance with Federal Guidelines	
	Ensuring Transparency, Accountability, and Ethical AI Use	
	Future Regulatory Trends	
Conclusion		**Page 15**

AI Governance is a Competitive Advantage

Artificial intelligence (AI) is transforming the healthcare industry, offering unprecedented opportunities to enhance patient outcomes, streamline administrative processes, and drive medical innovation. For instance, a recent AI model developed by Harvard Medical School demonstrated up to 94% accuracy in detecting various cancer types, surpassing existing methods by up to 36%[1]. Additionally, AI-enabled conversational agents have been associated with a threefold increase in engagement with therapeutic activities compared to control groups[2]. Moreover, a meta-analysis published found that patients undergoing robotic-assisted surgery experienced a 10% reduction in 30-day postoperative complications compared to those undergoing laparoscopic procedures. These advancements highlight AI's early-stage yet transformative potential in healthcare.

Alongside the challenges that come with developing, maintaining, and implementing these advancements come the complex governance challenges to meet fundamental requirements for patient security and privacy within AI development and use. Effective governance therefore goes beyond compliance and focuses on the stakeholders impacted by a given model. This shift from viewing AI governance as just a regulatory requirement to recognizing it as a competitive advantage enables organizations to see patient security and privacy as a market-driven strategy to gain and retain consumer trust, mitigate risks, and sustain future-oriented compliance.

Studies indicate that organizations leveraging structured AI governance frameworks see significant improvements in operational efficiency and financial performance. According to the World Economic Forum[3], companies that proactively manage AI risks gain a competitive advantage over other organizations, fostering long-term trust and success with their consumers. Beyond enhancing consumer trust, the Economist Intelligence Unit studied the benefits and costs of adopting responsible AI practices and found that these practices can build competitive advantage through factors such as product quality, ability to attract and retain talent, and expanding revenue streams. These studies highlight the fact that healthcare organizations integrating governance and risk mitigation strategies into their AI initiatives are better positioned to maximize the benefits of AI while maintaining compliance and trust.

A robust AI governance program also empowers companies to navigate and adapt to shifting regulatory landscapes with confidence and agility. By implementing strong governance frameworks, organizations can proactively address compliance challenges, reducing the risk of disruptions caused by evolving laws and policies. Rather than reacting to new regulations as they arise, businesses with well-structured AI governance programs can anticipate changes, implement adaptable policies, and integrate compliance measures seamlessly into their AI initiatives. This proactive approach not only minimizes regulatory risks but also ensures operational continuity, fostering trust among stakeholders and positioning the organization as a leader in responsible AI adoption.

A recent example of the evolving regulatory landscape is the now-rescinded U.S. Department of Health and Human Services' (HHS) AI Strategic Plan, which initially sought to establish a structured approach for AI oversight in healthcare. While the plan's withdrawal reflects concerns that regulatory frameworks might stifle innovation, it also underscores the volatility of AI policy. This regulatory uncertainty highlights the importance of proactive AI governance and how organizations that establish robust internal risk management and compliance frameworks will be better positioned to adapt to future regulatory shifts without disrupting operations or innovation.

The debate over AI regulation often centers on a perceived trade-off between fostering innovation and enforcing oversight. However, effective AI governance does not have to be an obstacle to progress; rather, it serves as a tool for responsible and effective AI adoption. By embedding governance, risk mitigation, and compliance measures into AI strategies from the outset, healthcare entities can safeguard patient data, reduce liability, and enhance stakeholder confidence while driving sustainable AI-driven advancements.

As regulatory priorities continue to evolve, the organizations that take a forward-thinking approach to AI governance will not only meet compliance expectations but also strengthen their competitive edge. Investing in governance continuity and risk mitigation today ensures resilience in the face of policy shifts tomorrow. The following section outlines the structure of the annulled plan and what it aimed to achieve. While rescinded, this example of regulation provides insight into potential future regulation and best practices.

[1] https://www.nature.com/articles/s41586-024-07894-z
[2] https://www.theverge.com/news/622528/microsoft-dragon-copilot-ai-healthcare-assistant
[3] https://www.weforum.org/stories/2024/09/ai-governance-trends-to-watch/

Case Study: HHS AI Strategic Plan Overview

The U.S. Department of Health and Human Services (HHS) developed its AI Strategic Plan in response to Executive Order 14110 (which has since been rescinded). However, despite this rescission, the plan can still offer valuable insights that healthcare organizations can leverage to improve their AI strategies to enhance public health outcomes, streamline human services, and promote health equity and accessibility.

The HHS AI Strategic Plan, as previously published, placed emphasis on risk management, transparency, and the ethical use of AI, signaling the importance of a commitment to harnessing AI's transformative potential while ensuring its ethical, equitable, and responsible integration. By embracing AI innovations like predictive disease models, drug discovery tools, and personalized medicine, HHS envisioned a future where AI-driven advancements improved the quality of life for individuals nationwide while increasing access to underserved populations.

The plan also recognized that harnessing these opportunities comes with significant challenges. Risks such as algorithmic bias, misleading outputs, and unequal access to AI technologies could worsen disparities in healthcare outcomes. Addressing these concerns requires a human-centered design approach that prioritizes fairness, accountability, and transparency. HHS emphasized building AI systems that adhere to principles of fairness, appropriateness, validity, effectiveness, and safety (FAVES), as well as ensuring that systems are tested in real-world settings and are trusted by communities.

The HHS AI Strategic Plan outlined a proactive framework for AI governance, advocating for cross-agency collaboration, updated regulations, and strategic investments to adapt to emerging technologies. By balancing innovation with safeguards, HHS aimed to lead in advancing AI solutions that serve public health while mitigating potential harm. This plan offered a roadmap to responsibly leverage AI's potential, prioritizing inclusivity and equity to ensure its benefits are accessible to all. This approach established a benchmark for ethical AI use in health and human services that healthcare organizations can learn from today, fostering trust and innovation for societal benefit.

Organization of the Plan

The HHS AI Strategic Plan outlined the agency's vision and goals for using AI across health, human services, and public health. It established a unifying framework consisting of seven domains to align HHS policies, programs, and activities involving AI:

Primary Domains (core parts of the HHS value chain):

1. **Medical Research and Discovery:** Includes fundamental research, clinical applications, and healthcare delivery.
2. **Medical Product Development, Safety, and Effectiveness:** Includes drug and device development, clinical trials, regulatory approval, and safety monitoring.
3. **Healthcare Delivery:** Includes diagnosing, treating, and preventing diseases while promoting health and financing care.
4. **Human Services Delivery:** Addresses providing social services to address basic needs and improve well-being.
5. **Public Health:** Addresses enhancing population health through prevention, life prolongation, and promotion of healthy choices.

Cross-Cutting Domains (supporting functions across primary domains):

6. **Cybersecurity and Critical Infrastructure Protection:** Addresses securing systems critical to health and human services.
7. **Internal Operations:** Addresses adapting HHS policies, infrastructure, and programs to enable AI implementation and rapid innovation.

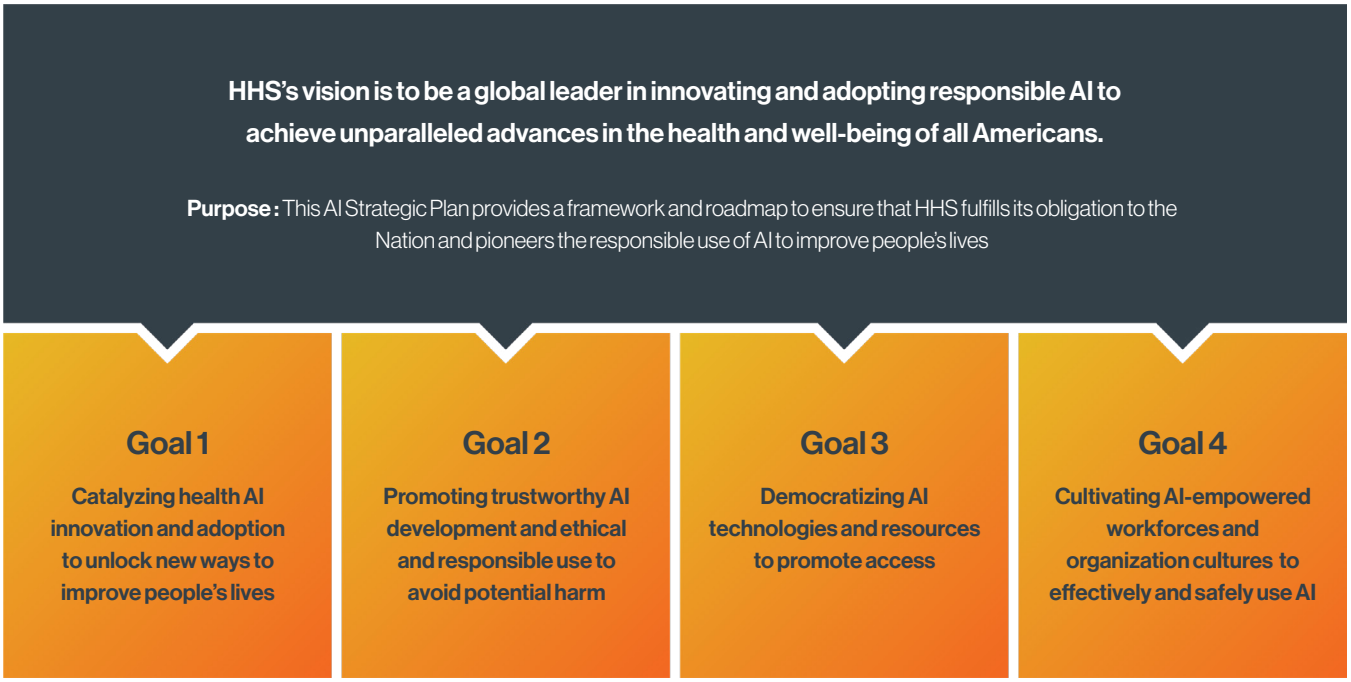
Each domain was structured with chapters that covered AI-related context, stakeholders, opportunities, trends, use cases, risks, and an action plan. This framework promoted comprehensive and coordinated AI adoption across HHS's diverse functions based on the nuance of each sector.

Goals of the Plan

HHS aimed to foster a coordinated public-private approach to enhance the quality, safety, efficiency, accessibility, equity, and outcomes of health and human services through the responsible use of AI. This objective was guided by four key goals:

1. **Catalyzing AI Innovation:** Drive health AI adoption to discover new ways to improve lives.
2. **Promoting Trustworthy AI:** Ensure ethical, responsible, and safe AI development to prevent harm.
3. **Democratizing AI Access:** Expand access to AI technologies and resources.
4. **Empowering Workforces:** Build AI-empowered cultures and equip workforces for effective and safe AI use.

HHS planned to implement its AI-focused goals within each Primary Domain through a variety of actions and tools, including regulations, policies, grants, funding, public education, and internal infrastructure improvements. The plan emphasized updating existing approaches to address emerging AI concerns while maintaining compliance with current regulations, such as patient privacy.



(With detailed action plans in each chapter)

Primary Domains:

- **Medical Research and Discovery**
- **Medical Product Development, Safety, and Effectiveness**
- **Healthcare Delivery**
- **Human Services Delivery**
- **Public Health**

Additional Domains:

- **Cybersecurity and Critical Infrastructure Protection**
- **Internal Operations**

Promoting Ethical and Responsible Use of AI

Promoting ethical and responsible use of AI in healthcare requires balancing innovation with risk management. HHS recognized the significant opportunities that AI presents for improving health and wellness outcomes while acknowledging the challenges associated with its adoption. The AI Strategic Plan emphasized the need to foster trustworthy AI, promoting that its deployment must align with ethical and responsible principles.

One of the primary concerns outlined in the plan was addressing the risks associated with AI-driven decision-making. Improper implementation could lead to misdiagnosed patient conditions, breaches of patient confidentiality, misallocation of resources during emergencies, the development of harmful medical products, and the exacerbation of social or environmental inequalities. To mitigate these risks, HHS prioritized governance strategies that enhance oversight, transparency, and risk management in AI applications. These are concepts that, separate from the rescinded plan, should be applied in healthcare today.

Promoting equity was another key focus of the plan, with proposed efforts aimed at reducing biases in AI models and increasing healthcare accessibility, particularly for underserved populations. This includes ensuring that AI benefits reach geographically remote communities and individuals with disabilities. To guide responsible and equitable AI development and deployment, HHS incorporated well-established frameworks such as the National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF)[4] and the Office of the National Coordinator for Health Information Technology (ONC) HTI-1 Final Rule[5]. These frameworks emphasize transparency, risk mapping, measurement, and governance to ensure that AI technologies align with ethical and regulatory standards. The agency also adheres to the FAVES principles, which define key attributes of responsible AI:

- **Fair (F):** Avoiding prejudice or favoritism
- **Appropriate (A):** Ensuring suitability for specific contexts and populations
- **Valid (V):** Producing accurate and expected results
- **Effective (E):** Demonstrating real-world benefits
- **Safe (S):** Minimizing risks so that benefits outweigh potential harm

HHS committed to ongoing adaptation in AI governance, recognizing that responsible AI implementation requires continuous evaluation. By engaging with stakeholders and updating strategies in response to technological advancements, the agency sought to ensure that AI in healthcare remained safe, fair, and effective. Despite the withdrawal of the HHS plan, these values remain a business imperative today as the responsibility for ethical AI use extends beyond compliance and across all stakeholders, reinforcing the collective duty to uphold high standards in AI-driven healthcare solutions.

[4] https://www.nist.gov/itl/ai-risk-management-framework

[5] https://www.healthit.gov/topic/laws-regulation-and-policy/health-data-technology-and-interoperability-certification-program

What's Next?

While the HHS AI Strategic Plan has been rescinded, its framework and goals provide valuable insights for healthcare organizations seeking to navigate the uncertain regulatory landscape. The rapid shifts in AI policy highlight the need for healthcare entities to take a proactive approach to governance, risk mitigation, and compliance, rather than relying solely on government directives. Organizations can learn from the plan's emphasis on structured oversight, ethical AI development, and cross-sector collaboration to establish their own resilient AI governance strategies.

By adopting key elements of the rescinded plan, such as risk management frameworks, transparency measures, and workforce empowerment, healthcare entities can future-orient their AI initiatives. This approach not only helps streamline compliance with any forthcoming regulations but also positions an adopting organization as a leader in responsible AI adoption. The following sections explore additional guidance on governance, risk management, and compliance strategies that healthcare entities can implement to maintain continuity, adapt to evolving regulations, and mitigate the risks associated with AI integration.

AI Governance in Healthcare

Defining AI Governance in Healthcare

Achieving the goals outlined in the HHS AI Strategic Plan requires more than policy direction; it necessitates active engagement from healthcare organizations to translate strategic priorities into operational safeguards. While HHS provided a framework for what the outcome should look like, the onus for responsibly implementing AI falls on healthcare entities and technology developers. Organizations must establish AI governance structures that integrate risk management, compliance strategies, and ethical oversight to ensure that AI aligns with regulatory requirements and patient safety standards.

AI governance in healthcare encompasses the policies, structures, and processes that guide the responsible deployment of AI, fostering trust, compliance, and risk mitigation. A well-defined governance framework is essential for upholding principles such as transparency, accountability, fairness, and security. These principles integrate seamlessly with broader risk management frameworks, including the NIST AI RMF, HITRUST AI[6], International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 38507 Information technology — Governance of IT — Governance implications of the use of artificial intelligence by organizations[7], and Organization for Economic Cooperation and Development (OECD) AI Principles for AI governance[8], providing structured approaches to managing AI-related risks while aligning AI technologies with healthcare's ethical and regulatory imperatives.

[6] https://info.hitrustalliance.net/hitrust-ai

[7] https://www.iso.org/standard/56641.html

[8] https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

AI Risk Management Frameworks

The NIST AI RMF offers a structured methodology for identifying, assessing, and mitigating AI risks throughout the system lifecycle. It is built on four key functions:

1. **Govern:** Establishes policies, roles, and accountability structures for AI oversight.
2. **Map:** Identifies potential AI risks and contextual factors influencing system performance.
3. **Measure:** Evaluates AI models for accuracy, robustness, and fairness, ensuring they operate as intended.
4. **Manage:** Implements continuous improvements and mitigation strategies to address evolving risks.

In healthcare, the NIST AI RMF is instrumental in evaluating model reliability, detecting biases in AI-driven diagnostics, and ensuring transparency in clinical decision-making tools.

Building on NIST's foundation, HITRUST AI extends AI governance by integrating risk-based, certifiable security and privacy controls tailored to healthcare. HITRUST's AI framework provides a standardized approach to assessing AI trustworthiness, aligning with NIST AI RMF while incorporating controls for data protection, algorithmic transparency, and bias mitigation. Given HITRUST's strong adoption across the healthcare industry, its AI framework serves as a bridge between technical AI risk management and broader compliance efforts, supporting organizations in meeting HIPAA, the EU General Data Protection Regulation (GDPR), and other regulatory requirements.

Similarly, ISO/IEC 38507 provides governance principles that help organizations align AI strategies with business objectives while maintaining compliance with healthcare regulations such as HIPAA and the 21st Century Cures Act. It emphasizes risk-informed decision-making to ensure that AI applications align with ethical, regulatory, and operational requirements. In healthcare, this means establishing clear accountability mechanisms for AI-based clinical support tools and ensuring that AI-driven automation does not compromise patient safety or regulatory compliance.

The OECD AI Principles, adopted globally by governments and regulatory bodies, offer additional guidelines for responsible AI development. These guidelines emphasize the following principles:

- **Human-centered values and fairness to prevent discrimination and bias**
- **Transparency and explainability to ensure stakeholders understand AI-driven decisions**
- **Robustness and safety to minimize unintended consequences**
- **Accountability to establish clear responsibilities and governance structures**
- **Inclusivity to promote equitable AI access and application**

By embedding these principles into governance structures, healthcare organizations can foster trust in AI applications and ensure that AI-driven decisions are interpretable, unbiased, and aligned with the best interests of patients. The integration of HITRUST AI, ISO/IEC 38507, and NIST AI RMF provides a comprehensive, multi-layered approach to AI risk management, balancing innovation with compliance and ethical considerations.

Algorithmic Impact Assessments

A critical component of AI governance is the use of Algorithmic Impact Assessments (AIAs) to evaluate AI systems before deployment. AIAs serve as a structured mechanism to identify ethical, legal, and technical risks early in the AI lifecycle, allowing organizations to implement proactive risk mitigation strategies.

AIAs are modeled after established risk assessment frameworks, such as Data Protection Impact Assessments (DPIAs) in the GDPR and cybersecurity risk assessments, both of which have demonstrated success in improving compliance and security across industries. Given the complexity of AI in healthcare, these assessments offer a systematic approach to ensuring responsible AI adoption. Specifically, AIAs provide a framework to:

- **Document AI decision-making processes to enhance transparency and accountability.**
- **Ensure compliance with evolving regulatory and ethical standards.**
- **Identify potential biases in AI models that may lead to unfair or inaccurate patient outcomes.**
- **Assess security vulnerabilities to protect sensitive health data and prevent misuse.**

In healthcare, the stakes of AI failures are extraordinarily high — misdiagnoses, biased algorithms, and opaque decision making can directly impact patient safety and public trust. AIAs can help mitigate these risks as they are particularly valuable in evaluating:

- **AI-driven diagnostics:** Ensuring accuracy, minimizing false positives/negatives, and preventing biases
- **Clinical decision-support tools:** Validating AI recommendations to align with medical best practices
- **Automated resource allocation systems:** Preventing disparities in care distribution and patient prioritization

Real-world examples have shown the risks of inadequate AI governance. Studies have shown instances where AI models, trained on incomplete or biased datasets, have led to discriminatory outcomes in patient care. For example, a widely used healthcare algorithm to identify patients who would require additional care was found to disproportionately allocate fewer resources to Black patients compared to white patients with similar health conditions[9]. Implementing AIAs before AI deployment can help detect and correct such biases before they cause harm.

As healthcare organizations increasingly rely on AI, regulatory uncertainty further underscores the need for AIAs. The rescission of the HHS AI Strategic Plan does not eliminate the need for oversight — it highlights the importance of voluntary governance frameworks to anticipate potential future regulations. Organizations that integrate AIAs into their AI strategy today will be better positioned to adapt to future regulatory shifts, in addition to strengthening patient trust and ensuring ethical AI deployment.

[9] https://pubmed.ncbi.nlm.nih.gov/31649194/

Ensuring Transparency, Accountability, and Security

To maintain public trust, AI governance must embed transparency and accountability into its framework. Transparency ensures that AI-driven decisions are interpretable and auditable, while accountability establishes clear responsibilities for AI oversight. Implementing explainability mechanisms, such as Shapley Additive Explanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME),[10] can help organizations to enhance interpretability and regulatory compliance. Regular audits, validation checkpoints, and continuous monitoring should be integrated into governance workflows to identify and rectify potential issues in real time.

Additionally, security and privacy safeguards must be central to AI governance. AI systems should be designed to comply with cybersecurity best practices to help ensure resilience against data breaches and adversarial attacks. By establishing a well-defined governance framework, healthcare organizations can align with current and future HHS strategic objectives while responsibly innovating and safeguarding public trust.

By integrating NIST AI RMF, OECD AI Principles, and AIAs into AI governance strategies, healthcare organizations can create a comprehensive framework for responsible AI use. This approach enhances regulatory compliance, fosters trust, and ensures that AI innovations contribute meaningfully to improved patient outcomes while minimizing risks to safety, privacy, and equity, aligning with the original goals of the rescinded HHS AI Strategic Plan.

[10] https://arxiv.org/html/2305.02012v3

Risk Mitigation Strategies for AI in Healthcare

Identifying AI Risks

The integration of AI in healthcare presents a range of risks that organizations must proactively manage to ensure safe and effective deployment. A governance, risk, and compliance (GRC)-driven approach is essential for identifying, assessing, and mitigating these risks within a structured framework, particularly as AI technologies continue to evolve.

One of the most significant concerns for AI integration in healthcare is data privacy and security. AI systems in healthcare rely on vast amounts of patient data, particularly in the context of electronic health records (EHRs) and diagnostic tools. These systems are prime targets for cyber threats given their volume of enriched datasets on patients and ensuring that AI models are secure is critical to compliance with HIPAA regulations. Failures in securing these models could result in data breaches, regulatory penalties, and reputational damage. Further, third party AI vendors and external datasets introduce additional vulnerabilities, requiring healthcare organizations to evaluate supply chain risks as part of their AI governance policies.

Another major challenge is the explainability of AI models and the ability for clinicians and patients to trust them. Many AI systems, particularly those based on deep learning, operate as "black boxes," making it difficult to interpret their decision making processes.

This lack of transparency can lead to clinical hesitancy, increased liability, and regulatory scrutiny from agencies like the Food and Drug Administration (FDA) and ONC. For healthcare organizations, developing clear model governance policies and ensuring transparency are essential for fostering trust among providers and patients.

Algorithmic bias also remains a pressing concern. AI models that are trained on imbalanced datasets may inadvertently reinforce health disparities, leading to misdiagnoses, incorrect treatment recommendations, and inequitable healthcare access for marginalized populations. The rescinded HHS AI Strategic Plan highlighted making fairness and equity as central tenets of AI, making bias mitigation not only a regulatory requirement but also an ethical imperative. As AI regulations evolve — mandating transparency, risk management, and bias mitigation — healthcare organizations must ensure their AI governance frameworks remain agile and compliant with these shifting standards.

Proactive Risk Management Approaches

To mitigate AI-related risks, organizations should implement structured risk management frameworks that integrate AI governance into their enterprise risk management and compliance programs. A critical component of this approach is bias mitigation, which requires organizations to prioritize the use of diverse and representative training datasets, conduct ongoing fairness audits, and implement algorithmic transparency measures. These initiatives align with the previous HHS AI strategic goals to foster trustworthy AI models and democratize AI access, ensuring equitable benefits for all users.

Data security and compliance measures must also be reinforced to protect patient information and AI infrastructure. Organizations should implement Zero Trust security models, enforce multi-factor authentication, and apply role-based access controls to restrict unauthorized access. End-to-end encryption should be a standard practice to safeguard sensitive data, while regular penetration testing and third-party risk assessments can help identify and address vulnerabilities in AI-driven healthcare applications. Additionally, organizations should establish clear policies for the use of third party AI vendors, ensuring that external providers adhere to the same security and compliance standards as the organization.

Explainability and clinician adoption are essential for building trust in AI systems. To facilitate transparency, organizations should implement model governance policies that require explainability techniques such as SHAP, LIME, and counterfactual explanations. These techniques help clinicians interpret AI-generated insights and enhance their confidence in AI-driven recommendations. Human-in-the-loop (HITL) systems should also be adopted to ensure that AI outputs are reviewed by medical professionals before final decisions are made, reducing the risk of errors and reinforcing accountability in patient care.

AI Testing and Validation

Before deployment, AI systems should undergo rigorous testing and validation to ensure they meet safety, efficacy, and regulatory standards. Preclinical validation should include simulated environment testing, where AI models are assessed under controlled conditions to evaluate their accuracy and potential clinical impact. Sandbox testing should also be conducted to detect and resolve any issues before AI systems are introduced into real-world patient care environments. Adhering to established regulatory frameworks, such as the FDA premarket review for software as a medical device, can help ensure that AI technologies meet necessary compliance and performance requirements.

Once deployed, continuous monitoring is essential to maintain AI system integrity and effectiveness. Organizations should implement automated performance tracking tools to detect "concept drift," where AI models lose accuracy over time due to shifts in clinical data patterns. Establishing governance protocols for ongoing model auditing can help identify performance degradation and ensure that AI systems remain aligned with clinical best practices. Real-world performance audits should be conducted at regular intervals to assess AI effectiveness, ensure compliance with regulatory standards, and incorporate feedback from healthcare practitioners. This continuous evaluation process is vital for maintaining AI reliability, positive patient outcomes, and consistent consumer trust.

Trust in AI-driven healthcare solutions is not just a regulatory necessity but also a critical competitive advantage. As AI adoption expands, organizations that prioritize transparency, reliability, and patient safety will differentiate themselves in the market. Maintaining consumer confidence through robust validation, continuous oversight, and ethical AI practices fosters stronger adoption rates and long-term brand loyalty. By proactively managing AI risks and ensuring compliance, organizations can secure their position as trusted leaders in the evolving healthcare AI landscape.

Incident Response Preparedness

Healthcare organizations must develop AI-specific incident response plans to effectively address failures, biases, and security threats. AI performance degradation protocols should be in place to rapidly detect and correct system malfunctions, preventing inaccurate or unsafe AI-driven actions from affecting patient care. Establishing regulatory reporting mechanisms aligned with guidelines from HHS, FDA, and ONC will help ensure that AI-related incidents are managed transparently and in compliance with federal requirements.

Beyond technical failures, organizations must also prepare for ethical concerns related to AI decision-making. Ethical escalation frameworks should be integrated into AI governance structures, allowing organizations to respond proactively to concerns about fairness, safety, and unintended harm. Designating AI ethics committees to oversee these issues can help ensure that AI systems are deployed responsibly and in alignment with both legal and ethical obligations. By fostering a culture of accountability and continuous improvement, healthcare organizations can enhance trust in AI technologies while minimizing risks to patients and healthcare providers.

Embedding Risk Mitigation into AI Governance

By integrating these risk mitigation strategies into their AI governance frameworks, healthcare organizations can align AI deployments with the best practices suggested in the rescinded HHS AI Strategic Plan, helping ensure that AI-driven innovations remain safe, equitable, and effective. A structured GRC approach can strengthen regulatory compliance, foster stakeholder trust, and ensure AI systems contribute to responsible and sustainable healthcare transformation.

Regulatory Compliance and Ethical Considerations

Existing Regulatory Landscape

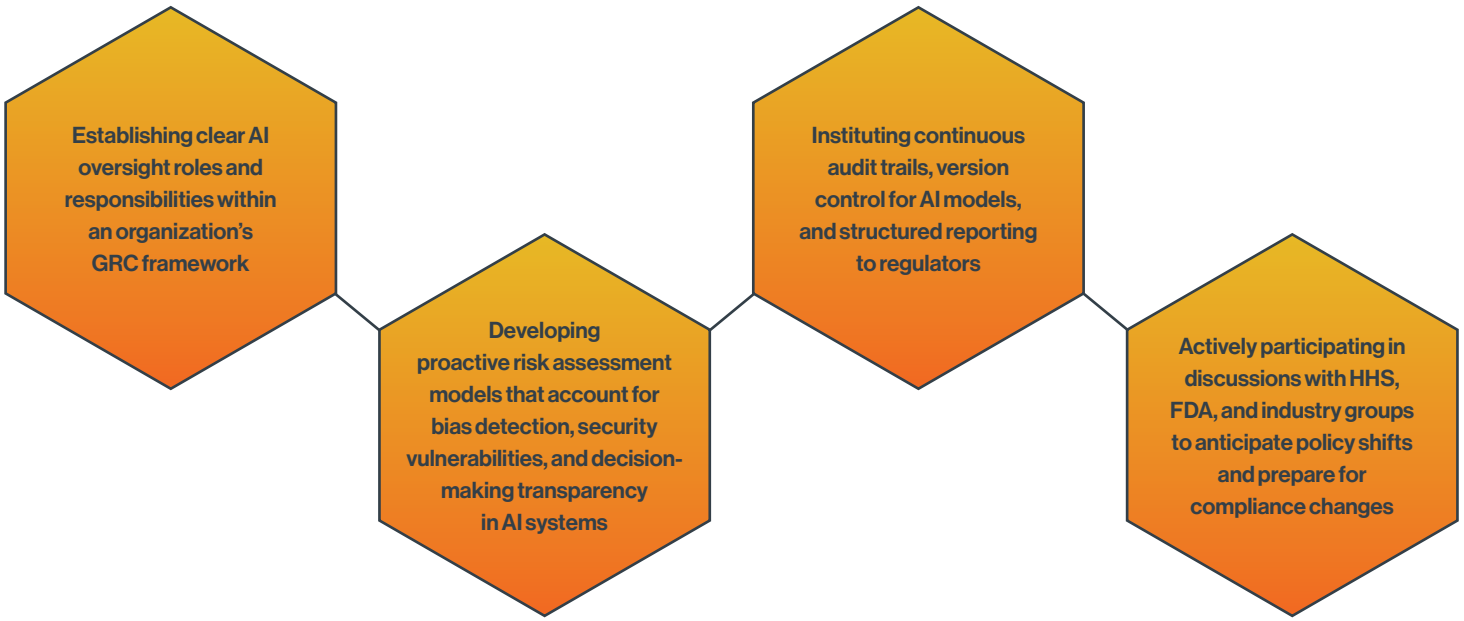
AI adoption in healthcare is subject to a complex regulatory environment designed to protect patient safety, ensure data privacy, and provide for ethical AI use. HIPAA mandates stringent security and privacy measures for patient data, ensuring that AI-driven healthcare applications uphold confidentiality and integrity. The FDA oversees AI and machine learning (ML)-based Software as a Medical Device (SaMD), requiring compliance with safety, efficacy, and validation standards before deployment. Additionally, the Health Information Technology for Economic and Clinical Health (HITECH) Act expands security and privacy protections, promoting responsible electronic health information sharing.

Beyond these existing regulations, frameworks such as 42 CFR Part 2[11] Confidentiality of Substance Use Disorder Patient Records and state-specific AI laws increase regulatory complexity, necessitating a structured GRC approach to ensure organizations remain compliant across multiple regulatory domains.

[11] https://www.ecfr.gov/current/title-42/chapter-I/subchapter-A/part-2

Aligning AI Governance with Federal Guidelines

Implementing AI regulatory compliance programs that align with the objectives of the rescinded HHS AI Strategic Plan can position organizations to more ably respond to evolving regulations. As such, these programs should include:



Ensuring Transparency, Accountability, and Ethical AI Use

Regulatory compliance is only one aspect of responsible AI adoption — ethical considerations must also be a core component of AI governance. To meet current and future regulatory expectations and objectives, organizations should:

- **Prioritize explainability through adopting interpretable AI techniques, such as SHAP or LIME, ensuring clinicians and regulators understand AI-driven decisions and models.**
- **Implement AI audit trails and logging mechanisms that track AI decision-making processes, enabling accountability and compliance validation.**
- **Monitor bias and fairness by leveraging AI governance frameworks such as AI RMF, HITRUST AI, and ISO/IEC 38507 to ensure fairness, minimize bias, and uphold ethical standards.**

These measures support an ethical AI governance structure that not only meets regulatory requirements but also fosters patient trust, clinician adoption, and responsible AI deployment.

Future Regulatory Trends

AI regulations in healthcare are rapidly evolving, with bias mitigation mandates, risk management requirements, and transparency obligations becoming key areas of focus. The White House's Blueprint for an AI Bill of Rights[12] and HHS's AI Strategy highlight the growing emphasis on equity, privacy safeguards, and algorithmic accountability.

Additionally, global AI governance frameworks, such as the EU AI Act[13], are likely to influence U.S. policies, necessitating a broader GRC approach that accounts for cross-border regulatory harmonization. Forward-thinking organizations should prepare by:

- **Establishing AI governance structures that incorporate global best practices**
- **Enhancing AI risk management to meet future compliance requirements**
- **Engaging with policymakers to shape responsible AI regulations that balance innovation with patient safety**

By proactively addressing regulatory compliance and ethical considerations through a strong GRC framework, healthcare organizations can align with HHS strategic objectives, ensure legal compliance, and drive responsible AI adoption in a way that is both trustworthy and future oriented. As AI regulations evolve, organizations must remain proactive in adapting their governance frameworks to align with emerging legal and ethical standards. Looking ahead, continuous engagement with regulatory bodies and adherence to best practices will be key to ensuring AI remains a force for good in healthcare.

[12] https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/

[13] https://artificialintelligenceact.eu/

Conclusion

The integration of AI in healthcare offers transformative potential to enhance patient outcomes, streamline operations, and drive medical innovation. However, these advancements come with significant risks, including algorithmic bias, data privacy concerns, and regulatory challenges. The HHS AI Strategic Plan's emphasis on the importance of ethical principles, transparency, and accountability provides a foundational framework for responsible AI adoption that can benefit organizations despite its recission.

To effectively address these challenges and achieve the goals outlined in the HHS AI Strategic Plan, healthcare organizations must implement practical risk management and governance strategies that go beyond the strategic plan. Establishing robust regulatory frameworks is essential to ensure compliance with ethical standards and legal requirements. Promoting transparency and accountability through regular audits, clear documentation, and stakeholder involvement can build trust in AI systems. Ensuring data privacy and security is paramount, requiring robust protection measures and policies that promote data minimization and anonymization.

Interdisciplinary collaboration is crucial for addressing the complex challenges of AI implementation. By fostering collaboration between healthcare professionals, AI developers, ethicists, and policymakers, organizations can develop more effective and ethical AI solutions. Continuous monitoring and adaptive governance are necessary to keep pace with evolving AI technologies, ensuring they remain aligned with the goal of enhancing patient care.

The successful integration of AI in healthcare requires a balanced approach that prioritizes innovation while safeguarding ethical considerations and patient safety. By aligning AI strategies with regulatory mandates and industry best practices, healthcare organizations can harness AI's full potential, ensuring that AI-driven healthcare remains safe, equitable, and effective. Future research and continuous refinement of governance models will be essential to navigate the rapidly changing landscape of AI in healthcare, ultimately fostering trust and innovation for societal benefit.

Additional Coalfire Thought Leadership

Navigating the AI Security Landscape: The Federal Push for Responsible AI Adoption

Navigating the AI Security Landscape: From Executive Orders to Cyber Resilience

AI Risk and Governance Obligations: Informational Guide

Mastering AI Risks: Navigating the NIST AI RMF Core with Coalfire

Impending Threats From AI: The Problem with Trust

Legal Disclaimer

This white paper is provided by Coalfire Systems, Inc. or its subsidiaries ("Coalfire") for informational purposes only. This white paper is the property of Coalfire and is protected by U.S. and international copyright laws. Unauthorized use, reproduction, or distribution of this white paper, in whole or in part, is strictly prohibited. Factual information included in this white paper has been taken from sources that Coalfire believes to be reliable, but its accuracy, completeness, or interpretation cannot be guaranteed. Information is current as of the date of this white paper only and is subject to change without notice. This white paper is provided "as-is" with no warranties, including any warranty of merchantability, fitness for a particular purpose, and non-infringement. Coalfire expressly disclaims all liability arising from or relating to the use of any information or material included in this white paper for any purpose, including any actions taken or not taken based on the contents of this white paper. You are solely responsible for making your own independent assessment of the information in this white paper and for the development, implementation, and execution of your information security program. For questions regarding any legal or compliance matters referenced in this white paper, you should consult your legal counsel, security advisor, or the relevant standard authority.

About the Author

Caleb Pfanstiel, GRC Healthcare Consultant

Caleb Pfanstiel has extensive experience in risk management, compliance, and securing complex healthcare operations. He is passionate about bridging the gap between technical challenges and business needs, helping organizations navigate the ever-evolving cybersecurity landscape. With a focus on healthcare, Caleb combines a strong understanding of regulatory frameworks with practical strategies to protect sensitive data and critical systems. His work ensures organizations can innovate confidently while maintaining robust security postures.

About Coalfire

The world's leading technology infrastructure providers, SaaS companies, and enterprises – including the top 5 cloud service providers and 8 of the top 10 SaaS businesses – rely on Coalfire to strengthen their security posture and secure their digital transformations. As the largest firm dedicated to cybersecurity, Coalfire delivers a comprehensive suite of advisory and managed services, spanning cyber strategy and risk, cloud security, threat and vulnerability management, application security, privacy, and compliance management. A proven leader in cybersecurity for the past 20 years, Coalfire combines extensive cloud expertise, advanced technology, and innovative approaches that fuel success. For more information, visit **Coalfire.com.**

Copyright © 2025 Coalfire. All rights reserved. The information in this document is subject to change at any time based on revisions to applicable regulations and standards. Any forward looking statements are not predictions and are subject to change without notice. Coalfire is not responsible for any errors or omissions.

WP_HealthAIGov_031025