Governing the Unpredictable: Ensuring Quality in the Age of Large Language Models

Rahul Ravel

Principal Program Manager, Nike, Inc. rravelpm@gmail.com

Abstract

AI Platforms are based on Large Language Models ("LLMs"). These LLMs represent a transformative departure from traditional software systems. Unlike conventional programs defined by deterministic logic, LLMs are trained on massive datasets and operate on probabilistic principles. This shift from deterministic to generative AI introduces profound implications for enterprise IT governance. Traditional governance frameworks rooted in Command-and-Control ("C2") mechanisms are ill-equipped to address the dynamic, emergent, and sometimes unpredictable nature of LLM outputs.

This whitepaper explores how enterprises can adapt IT governance to the unique demands of LLMs. We examine the limitations of current governance approaches, propose an observe-and-respond strategy, and analyze how organizations can maintain quality, accountability, and compliance in the context of probabilistic systems. We also outline key challenges and offer a roadmap for navigating the evolving landscape of AI-integrated enterprise environments.

Biography

Rahul is a Principal Governance, Risk, and Compliance Program Manager with Nike and is Past President of ISSA Portland. He has used his 20+ years of technical experience in helping drive enterprise compliance programs focused on areas including SOX, PCI, ISO, and Privacy (GDP, California Privacy). His current focus is development of a cybersecurity program management framework which addresses impact to scope if there is a dynamic state of program risk. In his free time, you can find him managing personal risk while rock climbing at Smith Rock, running trails in Forest Park, or climbing up Mt Hood.

1. A Paradigm Shift in IT Systems

Traditional software systems are designed with deterministic logic, i.e., given the same input, they reliably produce the same output. But LLMs defy this norm. They generate responses based on statistical patterns in data, meaning outputs may vary with each interaction. This probabilistic behavior introduces variability, complexity, and uncertainty. These conditions fundamentally alter how IT systems behave and ultimately present implications in how they should be governed.

So how do we deal with this? How do we tame this dragon? First, these risks should not be taken as adoption blockers nor be perceived as friction points preventing an organization's usage of LLMs. There are benefits. For example, a large corporate enterprise may benefit from cost and time efficiency with LLMs providing value from areas such as customer support (LLMs serving as a first line of query), supply chain (cost optimization modeling), and human resources (job applicant filtering).

2. Considerations for LLM Adoption

2.1. Probabilistic Behavior and the Nature of LLMs

LLMs operate through a process of probabilistic token generation. Their outputs are not fixed; but influenced by a range of factors, including model weights, input context, and randomness. This variability enables creativity and adaptability, but also poses challenges for predictability, repeatability, and control.

Key Characteristics of LLMs:

- **Stochastic Outputs**: The same input can produce different responses.
- **Emergent Behavior**: Capabilities may arise unpredictably as model scale increases.
- **Opaque Reasoning**: Decision-making processes are not transparent.

These characteristics make it difficult to apply traditional software quality and governance models that assume deterministic logic and traceable decision paths.

Compare this with the classic Command and Control ("C2") model. C2 assumes that outputs are relatively known and standards and compliance controls can be easily established. But the C2 model becomes challenged when it comes to the probabilistic nature of LLMs. The C2 model then becomes inflexible, costly, and difficult to manage. In the realm of compliance, this can become an issue. What if we do nothing? These are not perceived risks. They are a reality. For example, the international organization, The Open Worldwide Application Security Project ("OWASP"), has already identified that LLM security risks as part of their OWASP Top Ten Risks list. [1]

2.2. The Implications of "Vibe" Coding

If your organization does software development, there is a high probability that developers are trying a "vibe" approach. This utilizes an LLM to generate software based on a request ("prompts") to the LLM. But this also means that the resulting LLM-generated software may have cybersecurity risks:

- **Insecure code patterns**: AI models are trained on vast datasets of code from public repositories, which can include insecure or outdated patterns. These patterns can introduce vulnerabilities like SQL injection and cross-site scripting (XSS) into new applications.
- **Hardcoded secrets**: In the pursuit of rapid development, developers may include API keys or other credentials in their prompts, and the AI may embed these "secrets" directly into the code. This is particularly dangerous if the code is pushed to a public repository, leaving sensitive information exposed to attackers.

- **Insufficient access control**: A core security principle is verifying that a user is authorized to perform a specific action. Vibe coding often generates quick-and-dirty code that may lack or misconfigure authorization checks, allowing authenticated users to access or modify data they should not have rights to change.

2.3. Enterprise Governance

What about compliance to ensure uniform governance across a large enterprise? Enterprise IT governance has traditionally relied on common prescriptive controls, predefined rules, and compliance checks to safeguard at scale. From a governance perspective, it made it easier to write controls and capture audit logs. However, with LLMs, challenges now include:

- **Difficulty in Predefining Rules**: LLMs can encounter unforeseen scenarios which may not be explainable.
- **Accountability Ambiguities**: It is unclear who is responsible when a model generates harmful or biased output.
- **Post-Hoc Oversight Limitations**: Reviewing outputs after deployment is insufficient given real-time interaction requirements.
- **Unintentional Lack of Compliance**: Interaction with an LLM may cause non-compliance and exposure of sensitive data.

This necessitates a fundamental rethinking of governance strategies. There must be a strategy that enables governance to operate frictionless at the scale and speed of the enterprise.

2.4. Industry Momentum: A Software Test Case for LLM Adoption

A recent article in Shift Asia [2] calls it an “AI revolution.” The article nicely summarizes several areas where software quality is positively impacted by and can enhance LLM adoption. Software development is becoming increasingly complex and business continues to ask for faster release cycles. Use of LLMs in software testing allows developers to develop and automate frameworks which can optimize test execution ensuring that software meets quality standards before it is released. Another opportunity is where LLM platforms can learn from previous testing cycles, improving over time and making the testing process more efficient. This can lead to better software performance and reliability.

3. Observe-and-Respond Strategy

Instead of static controls, enterprises must adopt dynamic governance models. The Observe-and-Respond strategy continuous monitoring, feedback loops, and adaptive controls that evolve an enterprise’s capacity to adopt an LLM platform.

The Four Core Principles of Observe-and-Respond:

- **Real-Time Monitoring**: Track system outputs for anomalies, bias, or risk.
- **Auditability**: Implement mechanisms for recording and reviewing interactions.
- **Iterative Adjustment**: Fine-tune models and governance policies based on observed behavior.
- **Matrix of Stakeholders**: Identify a vertical and horizontal group of LLM-knowledgeable stakeholders who will set direction, be accountable, and resolve issues.

This model aligns with modern DevOps and AIOps practices, integrating quality assurance into continuous deployment pipelines.

3.1. Cross-Functional Stakeholder Management

Before strategies are established, tools are procured, and tactical plans are developed, there should be a cross-functional team of vested stakeholders. These stakeholders must be in alignment with each other, especially if they share corporate strategic objectives. These stakeholders should meet on a recurring basis with specific objectives on LLM adoption and ownership. Figure 1 outlines an example Stakeholder Group and the responsibilities of each group.

Figure 1: Example Stakeholder Group

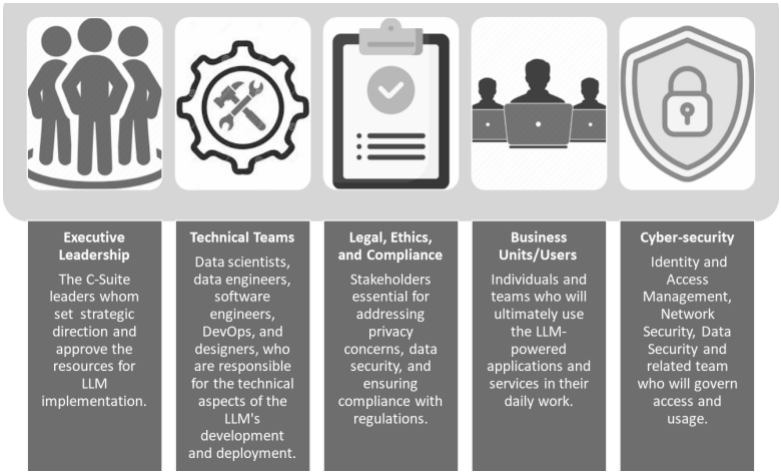

4. Challenges in LLM Governance Adoption

If only adoption were that easy. Compliance at scale, especially a large (i.e., number of people, geographically dispersed) enterprise. This is evidenced by the number of case studies for large Change Management or Enterprise Transformation projects. While there are a number of areas to discuss, this whitepaper will focus on Software Assurance.

4.1. Implications for Software Assurance

IT quality management traditionally emphasizes stability, reliability, compliance, and metrics. These principles must be reinterpreted for LLMs:

- **Stability** becomes about bounding variance rather than eliminating it.
- **Reliability** focuses on maintaining acceptable performance across variable outputs.
- **Compliance** requires tracking not just inputs and code, but also model behavior.
- **Metrics** should evolve to include:
    - Output consistency rates
    - Responsiveness to edge cases
    - Interpretability of results
    - Frequency and severity of harmful or biased outputs
    - Training to ensure that enterprise roles understand LLM flow and impact

4.2. Integration Risks

Integrating LLMs into enterprise systems introduces several risks:

- **Accountability Gaps**: When systems generate content independently, who bears responsibility?
- **Ethical Risks**: Bias, discrimination, and misinformation can emerge unexpectedly.
- **Regulatory Uncertainty**: Existing laws may not clearly apply to generative AI.
- **Security Concerns**: LLMs may leak sensitive data or be manipulated through prompt injection.

Mitigating these risks requires proactive strategies that blend technical, legal, and ethical expertise.

4.3. Awareness of Regulations

The regulations are growing for AI adoption. The European Union ("EU") is one of the first to establish compliance areas impacted by AI. In addition to the EU, the state of California is also leading AI-compliance. The following is a partial list. It is important to understand what the regulations are and how they will impact your adoption.

- **GDPR** ("General Data Protection Regulation") [3]: This EU regulation governs the collection, processing, and protection of personal data and is highly relevant to AI systems that handle sensitive information.
- **European Union AI Act** [4]: This EU regulation imposes strict compliance requirements on AI systems based on their risk level, with a focus on transparency, accountability, and human oversight for high-risk applications.
- **United States / California**: California has historically been a national leader in technical regulation and recently they have passed a series of AI regulations:
    - AB 2013 requires AI developers to disclose the data used to train LLMs.
    - AB 2355 mandates disclosures for political ads generated or significantly altered by AI, aiming to prevent voters being misled.

5. A Roadmap for Adaptive Governance

As mentioned earlier, a Observe-and-Respond strategy was advocated. One such example is an Adaptive Governance framework. Adaptive Governance supports the 'Observe' by constantly monitoring the enterprise for LLM impacts. Areas including risk, innovation, or compliance. 'Respond' takes the 'Observe' and allows a structured approach which enables a safe adoption of LLM usage at scale and speed.

5.1. Strategies for Effective Enterprise Integration

To manage LLMs effectively, enterprises should consider the following strategies:

- **Hold Yourself Accountable**: Ask questions on LLM usage, owners, outcomes, and regulation. Everyone is trying to understand their role when it comes to LLM adoption and should ask fundamental questions which may impact their responsibilities. Figure 2 has an example set of questions which can be asked at any phase of the software lifecycle.

Figure 2: LLM Implications Per Software Development Lifecycle

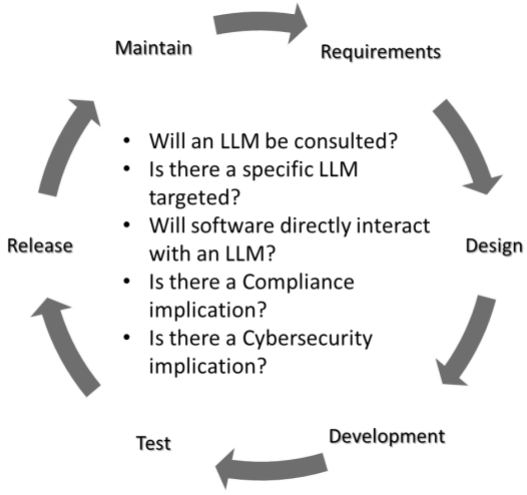

- **Understand Existing Frameworks**: There are a growing number of frameworks which can help provide guidance on LLM adoption and regulation:
    - **NIST AI Risk Management Framework** [5]: This US framework provides guidance for managing risks associated with the design, development, deployment, and use of AI systems, promoting principles such as trustworthiness, safety, security, privacy, and fairness.
    - **ISO/IEC 42001**[6]: This international standard provides a framework for AI governance, streamlining the process of documenting AI systems, assessing risks, and tracking compliance.
    - **EU AI ACT Compliance Checker** [7]**:** In the case of AI compliance for the EU, it is worth noting that it may not be clear where there are compliance implications. To address this, the EU has created an AI Compliance Checker tool and is recommended to understand its capabilities:

- **Build a Controlled Model:** A recommended reading is a recent whitepaper, 'Vulnerabilities in Deep Learning Language Models: Security Risks and Mitigation in Non-Federated, Federated and Decentralized Training.' by PNSQC members John Svetko and Bhushan Gupta. This whitepaper outlines several strategies for building an LLM model based on controlled data, essentially leading to a 'sanitized' LLM. The cost may be having a data set which will not have the global exposure of an LLM but a controlled model ensures that resulting responses are accurate and provide sound basis for key business decisions.

- **A Federated Landscape on LLM Model**s: There could be an opportunity for industry collaboration, where multiple companies want to share their data to create a more robust LLM model (strength thru numbers). The governance model would be different depending on level of participation or regulatory restrictions. The local governance would feed into the consortium governance.

- **Cross-Functional Governance Teams**: Include legal, compliance, data science, and IT stakeholders. Establishing cross-functional governance teams is essential for the responsible and efficient deployment of large language models within an enterprise. There are frameworks defined in this whitepaper which identify specific areas. These areas (e.g., technical, ethical, and regulatory considerations) can be used to help identify a stakeholder team or identify a gap if there is no specific stakeholder team.

- **Implement Tiered Quality Controls**: Apply different governance rigor based on application criticality. Implementing tiered quality controls involves applying varying levels of governance and oversight based on the criticality of the application. For example:
    - Executive Committee (Top Tier): Responsible for high-level decision-making, approving policies, allocating resources, and ensuring alignment with overall business goals.
    - Data Governance Council/Office (Middle Tier): Implements and oversees data governance policies, standards, and procedures. They also monitor data quality and compliance.
    - Data Stewards/Working Groups (Bottom Tier): Consist of representatives from different business units who are responsible for the day-to-day management and quality of specific data domains.

- **Adopt Transparent Model Management**: Document training data sources, fine-tuning methods, and version histories. Adopting transparent model management practices is vital for fostering trust and accountability in enterprise AI initiatives. This involves meticulously documenting aspects such as training data sources, data curation processes, fine-tuning techniques, and model version histories. Clear documentation facilitates reproducibility, enables effective troubleshooting, and supports audits or reviews.

  Transparency also helps stakeholders understand how decisions are made, which is crucial for addressing ethical concerns, ensuring regulatory compliance, and improving model performance over time. Maintaining comprehensive records ensures that model development and deployment processes are clear, auditable, and aligned with organizational standards.

- **Engage in External Audits**: Use third-party assessments to validate compliance and performance. Engaging in external audits provides an additional layer of validation for enterprise AI systems, ensuring they meet external regulatory and industry standards. These assessments, conducted by third-party experts, evaluate model compliance, robustness, and performance under real-world conditions.

5.2. Adaptive Governance

Adaptive Governance falls into the realm of enterprise transformation. These are multi-quarter, multi-year initiatives with specific outcomes. Thereby, an Adaptive Governance approach is recommended. An example is listed in Figure 3. A key benefit of this approach is that it addresses time constraints if an organization is pressured to adapt to a technical change. For example, these time constraints could be due to a new regulatory compliance or a Board of Directors directing a company to investigate LLM adoption by a certain date.

Figure 3: LLM-Focused Adaptive Governance Approach

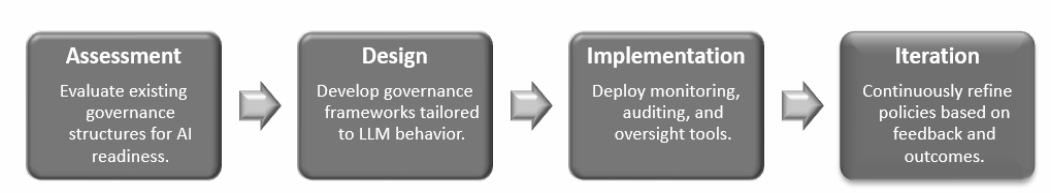

During the Assessment phase, policymakers gather comprehensive unexpected challenges or opportunities that emerge. The Assessment phase identifies such factors as current conditions, stakeholder needs, and system performance, enabling evidence-based decision-making rather than relying on assumptions or outdated information. The Design phase leverages these insights to create flexible policies that incorporate multiple scenarios and built-in adjustment mechanisms. This ensures solutions can evolve with changing circumstances.

Implementation becomes more effective because policies are grounded in real-world understanding and designed with adaptability in mind, allowing for responsive modifications. If we apply a compliance-focused Adaptive Governance approach, it should be able to ramp at a velocity where the organization can adapt to changes.

The Iteration phase completes the cycle by systematically evaluating outcomes, capturing lessons learned, and feeding insights back into the next assessment, creating a continuous learning loop that improves governance quality over time.

5.3. A Compliance-Focused Adaptive Governance Approach

The Information Systems Audit and Control Association ("ISACA") is one of the leading global organizations focused on IT governance through a framework comprising security, risk management, and assurance.

In January 2025, an ISACA leader, Goh Ser Yoong outlined four key recommendations when implementing AI [8]:

- **Involve stakeholders early, with continuous integration:** Stakeholders should be involved in the initial stages, providing input on security requirements, risk assessments and mitigation strategies. This early involvement will ensure that their considerations are embedded into the design, integration and development of an enterprise's LLM framework.

- **Upskilling and Training**: LLM stakeholders should be provided with the necessary training and upskilling opportunities to stay abreast of the latest AI developments and security challenges. This will enable them to contribute effectively to LLM usage.

- **A cross-functional collaborative culture**: Organizations should foster a culture of collaboration and communication. This can be achieved through regular meetings, joint workshops and shared training programs. Cross-functional collaboration will ensure that all teams understand each other's perspectives, share knowledge and work together to achieve common security goals.

- **Proper selection and usage of training data**: Organizations should prioritize the proper selection and usage of training data for AI models, including both properly sourced real-world data and the generation of synthetic data. This will ensure the development of robust and effective LLM solutions while addressing potential biases and privacy concerns.

5.4. ISACA and AI Compliance

ISACA has identified the need for governance when it comes to LLM adoption. They provide resources on understanding LLM platforms from the focus from an audit perspective (https://www.isaca.org/resources/artificial-intelligence).

ISACA also has a new certification, Advanced in AI Security Management™ [9]. It is the first and only AI-centric security management certification designed to help experienced IT professionals reinforce the enterprise's security posture and protect against AI-specific threats. This knowledge is seen as a value-add to the skills of those in a software-related role.

6. Conclusion

When ChatGPT was launched in November 2022, it quickly became one of the fastest-growing consumer applications in history. Within a few months, it reached millions of users, and by early 2025, it had close to 800 million weekly active users!

But this accelerated ramp left many enterprises scrambling to govern the adoption. Industry and national regulatory bodies are trying to implement governance but are trying to meet the velocity of adoption. This means there are other stakeholders on your journey. This will be a journey but there are many other travelers whom you can learn from so you're not fighting dragons but taming them. Here are components of a plan to achieve this.

6.1. Learn the Technology

It is hard not to find training resources so here is a group of focus areas which should provide a fundamental understanding to begin engagement in your enterprise:

- Fundamentals of how an LLM works
- Understanding Argentic AI and how vendors are incorporating it in their products
- Compliance. Several sources have been cited in this paper.
- Cybersecurity vulnerability introduced by LLM use.
- Development frameworks which address governance.

6.2. Learn From the Pathfinders

There are industry resources which support the case for an Adaptive Governance approach and a few are listed below. This approach enables adoption to co-exist with innovation both at the speed the business can successfully execute:

- "Generative AI Needs Adaptive Governance" [10] - This whitepaper outlines Traditional vs Adaptive Governance and make an objective Adaptive Governance model works well and where it doesn't with respect to AI.

- AI Governance Framework. Personal Data Protection Commission [11] - This is an extensive framework but has a vector for the level of human interaction required. This is a great starting point to help build a mode where governance can be scaled commensurate with human interaction.

- "Adaptive Governance with AI: A New Paradigm for Corporate Sustainability in High-Uncertainty Scenarios" [12] - This whitepaper outlines a phased approach for establishing governance. This is helpful if an enterprise needs to lay out a program timeline to implement the governance model.

6.3. Develop a Roadmap

Utilize the Adaptive Governance approach and begin to engage with your group of stakeholders. AI courses should provide the base knowledge to begin discussions. And as discussion continue, let your roadmap continue to develop with more details and longer-term milestones.

And finally, ensure you have checkpoints to ensure you understand if there is a positive return on your LLM investment. You should not have a Fear of Missing Out but you should have a Fear of Over-Investment if you quantitively and qualitatively identify where the business benefits are.

7. References

[1] Open Worldwide Application Security Project. “Top 10 Web Application Security Risks.” https://owasp.org/www-project-top-ten/ (accessed July 10, 2025)

[2] The AI Revolution in Software Testing and Quality Assurance,” Shift Asia. https://shiftasia.com/column/the-ai-revolution-in-software-testing-and-quality-assurance/ (accessed Jul 5, 2025)

[3] “General Data Protection Regulation,” Intersoft Consulting. https://gdpr-info.eu/ (accessed July 24, 2025)

[4] “The EU Artificial Intelligence Act,” Future of Life Institute. https://artificialintelligenceact.eu/ (accessed July 24, 2025)
- Updates can be subscribed thru: https://artificialintelligenceact.substack.com/

[5] “AI Risk Management Framework,” National Institute of Standards and Technology (NIST). https://www.nist.gov/itl/ai-risk-management-framework (accessed June 19, 2025)

[6] “Information technology — Artificial intelligence — Management system,” International Standards Organization. https://www.iso.org/standard/42001 (accessed June 21, 2025)

[7] “EU AI Act Compliance Checker,” Future of Life Institute. https://artificialintelligenceact.eu/assessment/eu-ai-act-compliance-checker/ (accessed June 25, 2025)

[8] Yoong, Goh Ser. “Securing Artificial Intelligence: Opportunities and Challenges,” ISACA. https://www.isaca.org/resources/news-and-trends/newsletters/atisaca/2025/volume-1/securing-artificial-intelligence-opportunities-and-challenges (accessed June 14, 2025)

[9] Advanced in AI Security Management,” ISACA. https://www.isaca.org/credentialing/aaism (accessed Jul 23, 2025)

[10] Reuel, Anka, Arne Undheim, Trond. “Generative AI Needs Adaptive Governance,” Cornell University. https://arxiv.org/abs/2406.04554 (accessed August 20, 2025)

[11] “Singapore’s Approach to AI Governance,” Personal Data Protection Commission - Singapore. https://www.pdpc.gov.sg/help-and-resources/2020/01/model-ai-governance-framework (accessed August 20, 2025)

[12] Porto Alegre De Almeida, Alex. “Adaptive Governance with AI: A New Paradigm for Corporate Sustainability in High-Uncertainty Scenarios,” ResearchGate. https://www.researchgate.net/publication/392693783_Adaptive_Governance_with_AI_A_New_Paradigm_for_Corporate_Sustainability_in_High-Uncertainty_Scenarios (accessed August 12, 2025)

Google AI, 2025. Gemini (Version 2.5), URL: https://search.google/