

From Risk to Opportunity: A Governance Blueprint for Agentic AI

Oct 2025



About

The Global Finance & Technology Network (GFTN) is a Singapore-headquartered organisation that leverages technology and innovation to create more efficient, resilient, and inclusive financial systems through global collaboration. GFTN hosts a worldwide network of forums (including its flagship event, the Singapore FinTech Festival); advises governments and companies on policies and the development of digital ecosystems and innovation within the financial sector; offers digital infrastructure solutions; and plans to invest in financial technology startups through its upcoming venture fund, with a focus on inclusion and sustainability.



For more information, visit www.gftn.co

Terminal 3 is a Hong Kong-based Web3 startup building data infrastructure for a decentralised future. The company's solutions are an alternative to centralised data systems that deprive users of privacy and saddles enterprises with compliance and security concerns. Terminal 3 leverages decentralised and privacy-enhancing technologies to empower an equitable Web3 where fully private user data is freely composable. The company's founders are successful corporate executives and entrepreneurs, who have built, scaled, and transformed some of the world's most important companies. Terminal 3 is also backed by world-class investors including Illuminate Financial, CMCC Titan Fund, Animoca Brands, IDG Blockchain, Progression Fund, Cherubic Ventures, 500 Global, Consensys Mesh, Bixin Ventures, BlackPine, Hard Yaka, and Bored Room Ventures.



For more information, visit our: [Official Website](#) | [X](#) | [LinkedIn](#)

Contents


About	02
Executive Summary	04
1. Introduction: The Imperative for Privacy in an Agentic AI World	05
Four Pillars for Accelerating the Adoption of Privacy-Preserving AI	
A. Modernising Regulatory Frameworks	
B. Building Industry Standards and Interoperability	
C. Enabling an Ecosystem of Investment and Innovation	
D. Education and Workforce Development	
2. Privacy-Preserving AI in Practice: Cross-Industry Applications	07
1. Healthcare: Collaborative Research Without Compromising Patient Privacy	
2. Financial Services: Secure Fraud Detection Across Institutions	
3. Consumer AI: Protecting User Prompts in Personal Assistants	
3. Implementation Roadmap: A Phased Approach	08
4. Conclusion: Privacy as a Competitive Advantage	09
References	10
Contributors	11

Executive Summary

Driven by billions of daily user prompts, the agentic AI market is projected to reach \$93.2 billion by 2032.¹ The explosive growth of agentic AI has created an unprecedented privacy crisis. These systems require access to our most intimate personal information to function effectively, resulting in risks such as unauthorised data access and unregulated agent-to-agent data trading. Yet, 97% of organisations lack adequate AI access controls², creating an urgent need for robust governance frameworks.

This report presents actionable solutions across four domains: modernising regulatory frameworks to address AI-specific privacy challenges, establishing industry standards for interoperability, enabling an ecosystem that fosters innovation and investment in privacy-enhancing technologies, and developing a specialised workforce needed for implementation.

It seeks to equip policymakers, industry leaders, and technologists with a pragmatic roadmap to create a future where privacy-preserving AI is not merely a compliance obligation but a distinct competitive advantage and a cornerstone of trust in the digital ecosystem.



1 Introduction: The Imperative for Privacy in an Agentic AI World

The proliferation of agentic AI systems marks a pivotal technological inflection point. It has introduced unprecedented capabilities for automation, personalisation, and efficiency. However, this transformative potential is shadowed by privacy risks that legacy frameworks are ill-equipped to manage.

The daily torrent of sensitive user data into these systems, combined with the capacity of AI agents to operate autonomously and interact with each other, creates complex privacy challenges. ChatGPT users now send 2.5 billion prompts every single day, representing a 150% increase in just eight months.³ When combined with hundreds of other consumer AI products and enterprise AI systems, trillions of new data tokens are generated annually.

Unlike traditional training data, prompt data contains the most intimate details of human thought, intention, and behaviour. Users share credit card details, passport numbers, health data, and personal relationships, a level of access that exceeds what most humans would share with their closest friends. Yet the infrastructure protecting this data remains dangerously inadequate.

The current landscape is characterised by eroding user trust, outdated regulations, and a growing number of security incidents. IBM's 2025 Cost of a Data Breach Report, based on data from 600 organisations globally that experienced that experienced a data breach, found that 60% of AI-related security incidents led to compromised data, while 31% resulted in operational disruption. AI application breaches have occurred in at least 13% of the surveyed organisations.⁴

This necessitates a fundamental shift in our approach to AI governance, moving from a reactive, compliance-focused posture to one that is proactive, collaborative, and centered on privacy-by-design. With privacy-forward technologies becoming ready for market implementation, like zero-knowledge proofs and multi-party computation, we can build the technical and ethical foundations for a sustainable agentic AI ecosystem—one that unlocks the immense value of this technology while safeguarding the fundamental right to privacy.

Four Pillars for Accelerating the Adoption of Privacy-Preserving AI

To accelerate the adoption of privacy-preserving agentic AI, a multi-stakeholder effort is required across four key domains.

A. Modernising Regulatory Frameworks

Current privacy regulations were designed for a pre-AI world and must be updated to address the novel challenges posed by agentic AI systems. To stay relevant, policymakers can prioritise the development of AI-specific privacy regulations that explicitly govern prompt data, AI-to-AI communications, and autonomous decision-making.

The European Union's AI Act provides a useful starting point, establishing risk-based classifications and transparency requirements. Nonetheless, further international efforts can help create harmonised standards. Meanwhile, the United States can explore having a comprehensive federal AI privacy legislation that provides clear guidance, to harmonise the fragmented landscape of different policies like Health Insurance Portability and Accountability Act (HIPAA), Gramm-Leach-Bliley Act (GLBA), and state laws like California Consumer Privacy Act (CCPA).

In modernising regulatory frameworks, regulatory bodies worldwide are increasingly recognising the importance of privacy enhancing technologies (PETs). The UK Information Commissioner's Office (ICO) has published guidance explicitly endorsing PETs. The Monetary Authority of Singapore (MAS) emphasises explainability, transparency, and customer data protection through its Veritas initiative. The European Data Protection Supervisor (EDPS) has called PETs "an essential tool for privacy-by-design," while US Federal Trade Commission (FTC) commissioners increasingly reference the need for technical safeguards in AI systems.

Regulatory frameworks can actively encourage PET adoption through the creation of safe harbour provisions, regulatory sandboxes, and other incentives that reward organisations for implementing privacy-by-design principles. Organisations that can cryptographically prove their adherence to privacy standards can benefit from reduced compliance burdens, lower penalties for breaches, and expedited approval processes, thereby creating a clear business case for investment in PETs. Additional incentives can include government procurement preferences that prioritise PET-enabled vendors and public certification programmes creating "Privacy-Preserving AI" trust marks that provide marketing advantages and meet enterprise procurement requirements.

However, one commonly proposed incentive—cross-border data transfer exemptions—should be approached with caution. While international data transfers under GDPR, CCPA, and China's Personal Information Protection Law (PIPL) create compliance challenges, data should remain sovereign and be transferred through as few hands as possible. In place of data transfer exemptions, a combination of decentralised storage, Oracle systems, and PETs can better enable the cross-border **verification** of data without actual transfer, safeguarding national identity while enabling global AI collaboration.

B. Building Industry Standards and Interoperability

The AI industry must collaboratively develop comprehensive standards for privacy-preserving AI systems, addressing technical implementation, security requirements, interoperability, and governance frameworks. Professional bodies like the Institute of Electrical and Electronics Engineers (IEEE) and the International Organisation for Standardisation (ISO) can establish dedicated working groups to formalise these standards, bringing together technical experts, policymakers, privacy advocates, and industry practitioners.

Industry leaders can champion transparency by publicly reporting on their privacy protection measures. This will not only help to build user trust but also foster a competitive environment where robust privacy protections become a market differentiator.

Industry consortiums for best practices can facilitate knowledge sharing and collaborative problem-solving. Drawing inspiration from successful interoperability initiatives in other sectors, such as the Fast Healthcare

Interoperability Resources (FHIR) standard in healthcare, the AI industry can create a common language for privacy that accelerates innovation and adoption. A prime example is in financial services where data sharing frameworks have enabled institutions to collaborate on fraud detection while maintaining competitive privacy.

Existing industry collaboration efforts provide models for standardisation. The Confidential Computing Consortium advances trusted execution environment technologies. The OpenMined community has pioneered open-source tools for federated learning. The Partnership on AI develops best practices for responsible AI development. These initiatives demonstrate that collaborative standard-setting creates shared infrastructure benefiting all participants.

C. Enabling an Ecosystem of Investment and Innovation

Accelerating PET deployment requires increased public and private investment. While the PETs market shows substantial growth, with a projected compound annual growth rate of 25.4%,⁵ targeted funding is needed to address challenges in performance, cost, and scalability.

Public funding agencies are recommended to establish dedicated research programmes and public-private partnerships focused on privacy-preserving AI. Likewise, the National Science Foundation, the European Research Council, and similar organisations can establish funding programmes supporting both fundamental research into novel cryptographic techniques and applied research for real-world AI applications.

Concurrently, private sector investment, from venture capital to corporate R&D, can be channeled towards startups and open-source projects that are pushing the boundaries of what is possible with PETs. Fostering a vibrant ecosystem of innovation is critical to ensure that these powerful technologies become accessible and practical for a wide range of AI applications.

Academic-industry partnerships are essential for translating research breakthroughs into deployable technologies. Universities possess deep expertise in cryptography and privacy-preserving computation, while industry partners understand practical constraints of production AI systems. Collaborative research initiatives, joint faculty appointments, and industry-sponsored PhD programmes can unlock this opportunity.

D. Education and Workforce Development

Successful PET implementation hinges on the availability of a workforce skilled in cryptography, privacy-preserving computation, and secure AI systems design. A significant talent gap currently exists in these fields, creating an adoption bottleneck. To close this gap, educational institutions, professional organisations, and employers can collaborate to develop comprehensive training programmes.

Universities can integrate these topics into their computer science, engineering, and business curricula. For instance, undergraduate courses can introduce fundamental cryptographic privacy concepts, while graduate programmes can offer specialised tracks in privacy-preserving machine learning, secure multi-party

computation, and applied cryptography. These curricula can also address ethical, legal, and policy dimensions.

Beyond academia, professional certification programmes and continuing education initiatives are needed to upskill the existing AI workforce. The National Initiative for Cybersecurity Education (NICE) Framework has successfully standardised role definitions and competencies. Similar frameworks for privacy-preserving AI can accelerate training programme development and workforce planning.

Public awareness campaigns are essential to educate consumers about AI privacy risks and benefits of PETs. An informed public that demands privacy protection will create powerful market incentives for organisations to prioritise the development and deployment of privacy-preserving AI technologies.

2 Privacy-Preserving AI in Practice: Cross-Industry Applications

PETs are already enabling innovative, privacy-preserving solutions across multiple sectors, including healthcare, financial services, and consumer applications.

1. Healthcare: Collaborative Research Without Compromising Patient Privacy

Federated learning, combined with secure multi-party computation, allows multiple hospitals to collaboratively train AI models for disease detection or treatment efficacy analysis without ever exposing sensitive patient data. By keeping the data decentralised and only sharing model updates, this approach enables groundbreaking research while upholding the strictest patient confidentiality standards.

A consortium of research hospitals can develop an AI system for early and accurate detection of rare diseases by training it on a diverse, multi-institutional dataset—a task that would be impossible under traditional data-sharing agreements. For instance, zero-knowledge proofs enable patients to prove clinical trial eligibility without revealing specific medical conditions, and insurance verification without exposing detailed medical histories.

2. Financial Services: Secure Fraud Detection Across Institutions

Each financial institution has valuable transaction data and fraud intelligence, but competitive pressures and regulatory requirements often prevent data sharing. With PETs, banks can contribute transaction data to collaborative analysis that helps identify fraudulent activities across the industry without revealing actual transactions.

By performing computations on encrypted data, banks can detect sophisticated, cross-institutional fraud patterns that would be invisible to any single organisation. Privacy-preserving machine learning ensures the AI learns to detect fraud patterns without memorising specific transactions or customer information. This enhances security for the entire financial system while protecting both customer privacy and proprietary business information.

3. Consumer AI: Protecting User Prompts in Personal Assistants

As AI agents become more integrated into our daily lives, protecting the privacy of our interactions becomes paramount. Technologies like Apple's Private Cloud

Compute and Google's Confidential AI ensure confidentiality by processing user requests in secure, isolated environments.

Another approach is exemplified by platforms like the Terminal 3 Network, which utilise a decentralised architecture combining blockchain for verification, decentralised storage, and a PET-enabled processing layer.

This tri-network architecture allows AI systems to perform tasks without ever having direct access to the underlying sensitive data, ensuring that personal conversations, financial queries, and confidential information remain private. The network's design enables AI agents to interact and coordinate while maintaining cryptographic guarantees about data handling, addressing both prompt privacy and emerging AI-to-AI communication risks.

3 Implementation Roadmap: A Phased Approach

The journey towards a privacy-preserving AI ecosystem is a marathon, not a sprint. It requires a phased, collaborative approach with clear milestones and metrics for success. The table below shows how this phased implementation approach can be planned and how its success can be measured.

Timeline	Focus	Milestones	Key Success Metrics
Near-Term (6-12 Months)	Foundation Building	• Convene government-led working groups to educate stakeholders and create industry standards • Launch pilot projects in regulated sectors (healthcare, financial services) • Map existing AI and data regulations across jurisdictions	• Number of pilot projects initiated across sectors • Participation rates in working groups • Documented cost-benefit analyses
Medium-Term (1-2 Years)	Ecosystem Development	• Establish industry consortium for PET interoperability to publish standards and reference architectures • Work with 3-5 key jurisdictions to pass safe harbour provisions • Develop PET curricula and certification programmes with higher education institutions	• At least three major industry consortiums established • PET-friendly regulatory provisions in major jurisdictions • Measurable increases in PET adoption rates
Long-Term (3-5 Years)	Mainstream Adoption	• Achieve regulatory harmonisation establishing PETs as privacy standard compatible worldwide • Facilitate widespread adoption making PETs the default for private data and agentic AI	• PETs prevalent in commercial AI applications • Established career pathways and professional certifications • Emergence of new privacy-preserving business models



4 Conclusion: Privacy as a Competitive Advantage

Privacy-preserving AI represents a distinct competitive advantage in an increasingly privacy-conscious market. Organisations that move first will build deeper trust with users increasingly aware of privacy risks. They will navigate regulatory compliance more efficiently, avoiding costly retrofitting when regulations inevitably tighten. They will unlock new collaborative opportunities impossible under traditional data-sharing constraints, accessing larger and more diverse datasets that improve AI performance. They will attract and retain top talent who want to work on

responsible AI that serves both technological advancement and human values.

The future of AI privacy requires coordinated action from different stakeholders, as shown in the table below.

By working together to make PETs more powerful, accessible, and efficient, we can choose to build a future where innovation and privacy are not competing values but mutually reinforcing pillars of a trustworthy and equitable digital world.

Stakeholder	Possible Actions Towards the Future of AI Privacy
Polymakers	• Establish regulatory sandboxes for privacy-preserving AI • Convene multi-stakeholder working groups on a shared goal towards AI governance • Commit to education, industry workshops, and funding around PETs
Industry Leaders	• Champion privacy as a core business value • Integrate privacy-by-design into AI development lifecycle • Conduct privacy risk assessments of current AI systems • Launch pilot projects in non-critical applications • Join industry consortiums and standards bodies • Establish cross-industry partnerships
Technologists	• Educate themselves on privacy-enhancing technologies • Evaluate PET libraries and frameworks • Advocate for privacy-preserving architecture in their organisations • Contribute to open-source PET projects



References

- ¹ **Agentic AI Market – Global Forecast by 2032, June 2025**
[MarketsandMarkets](#)
- ^{2,4} **IBM Cost of a Data Breach Report 2025**
[IBM Reports](#)
- ³ **OpenAI Usage Statistics, July 2025**
[TechCrunch](#)
- ⁵ **Grand View Research, Privacy Enhancing Technologies Market Size Report 2030**

Authors/Contributors

Gary Liu

Co-Founder & CEO
Terminal 3

Lauren Ho

Chief of Staff
Terminal 3

Krithi Sundar

Contributor
Terminal 3

Kaitlyn Thinn

Head of Strategy & Research
Global Finance & Technology Network (GFTN)

Bernice Neo

Customer Success Manager
Global Finance & Technology Network (GFTN)

Production

Sachin Kharchane

Graphic Designer

Global Finance & Technology Network (GFTN)

6 Battery Road, #28-01, Singapore 049909
gftn.co | hello@gftn.com

This document is published by Global Finance & Technology Network Limited (GFTN) as part of its FutureMatters insights platform. The findings, interpretations, and conclusions presented in GFTN Reports reflect the views of the author(s) and do not necessarily represent those of GFTN, its Board, management, stakeholders, or any individual participant and their respective organisations.

© 2025 Global Finance & Technology Network Limited, All Rights Reserved.
Reproduction Prohibited.