

Multi-Cloud Cybersecurity for LLMs in Banking: Governance and Threat Surfaces Across Financial AI Systems

Mubashir Ali Ahmed

Department of Computer Science, University of the People, CA, USA

Abstract: Large-scale neural networks referred to as LLMs are quickly transforming the world of banks with functionalities such as intelligent fraud detection, customer services, credit scoring, compliance management, AML, and financial forecasting among others. But then, given the increasing deployment of these artificial intelligence technologies in cloud computing ecosystems including Amazon Web Services, Microsoft Azure and Google Cloud Platform the cybersecurity risks are becoming very complex and difficult to manage. Multi-cloud computing offers many advantages regarding scalability, resilience, partial regulatory flexibility and business continuity among others; but at the same time increases the attack vectors and surfaces for cyber-attacks spread across distributed architectures, APIs, identity systems, vector databases, prompt pipelines and model registry infrastructure. Financial data is the type of data that any form of attack, whether it be a prompt injection, model poisoning and several others will be after illegally or improperly. This paper focuses on discussing some of the cybersecurity threats associated with multi-cloud deployments of LLMs within banking organizations together with the challenges surrounding their governance, followed by proposing a security framework that is governance-based.

Keywords: Multi-cloud security, large language models, banking cybersecurity, AI in finance systems, governance of AI, prompt injection, securing model regulators.

I. INTRODUCTION

The application of LLMs in banking institutions is rapidly increasing to enhance operations' efficiency, customer engagements, fraud detection, and monitoring & compliance with AML and credit risk assessment [1]. AI models allow banks to perform extensive processing of structured as well as unstructured data, thus accelerating decision-making for the organization. Various financial institutions are deploying LLM applications within multi-clouds such as AWS, Microsoft Azure, and Google Cloud Platform in order to achieve scalability, resiliency, and flexibility [2]. The multi-cloud environment facilitates availability and business continuity; however, this brings about a significant challenge with regard to the security and governance of LLMs since sensitive information passes through numerous platforms, leaving itself exposed to issues such as timely injection, data leakage, model poisoning, access control, and identity assertion through the process. The necessity to have audits and comply with various regulations further complicates the issue of security. For this reason, there is an urgent need to put appropriate governance measures in place to safeguard the LLMs used by banks against possible vulnerabilities and dangers.

II. LITERATURE REVIEW

A. Multi-Cloud Security in Banking

Among the main changes that banks implemented recently is moving towards multi-cloud computing, as that enables them to improve the following areas: availability of the service, disaster recovery, vendor neutrality and workload flexibility [3]. Banks are not splitting their critical services between external parties such as Amazon Web Services, Microsoft Azure and Google Cloud Platform in order to reduce reliance on one single provider for running operations. At the same time, there is a concern that the use of multi-cloud infrastructure poses cybersecurity challenges to the banks. Due to the fragmentation of identity, access control, monitoring tools and compliance policies across multiple platforms used, the risks of data breaches and any unauthorized access increase. This happens because of misconfigured APIs, improper Identity and Access Management implementation across devices, low visibility and complex environment and, finally, because of the absence of standard end-to-end encryption policies. The banking industry processes sensitive information about customers, and any vulnerability can lead to major consequences such as reputational loss or significant financial damage. Multi-cloud cybersecurity research points to centralized governance, unified monitoring and Zero Trust Architecture to minimize attack vectors within cross-platforms environment and make banks more resistant to non-cyber threats attacks [4]

B. LLM Security Challenges in Financial Systems

LLMs bring about a new kind of cyber-attack that differs from the typical software vulnerability. For example, LLMs are able to detect fraud, provide assistance through chatbots, assist in onboarding, evaluate loans, and conduct inspections for money laundering within the banking sector [5]. These systems require reliance on prompt engineering, RAG, vector databases, and API access to live data. Threats identified by researchers are prompt injection, prompt leakage, adversarial manipulation, hallucinations, and model poisoning. In this regard, prompt injection allows a model to behave abnormally and act in an unsafe manner, whereas prompt leakage may cause leaking of customer or institutional information. Hallucinations, on the other hand, will cause misinformation regarding financial advice or wrong compliance advice. These issues will directly affect customer trust and compliance needs. Example: Model poisoning during training or fine tuning will affect either fraud detection or predictions of lending [6]. The security of prompts, training pipelines, and outputs should be prioritized since financial institutions function in high-risk environments.

C. Compliance with AI Governance and Regulation

Protection of the use of LLMs in banking Yet another crucial function is AI governance. Banks are obliged to abide by strict laws regarding data protection, audit, and compliance with fairness and explainability standards. Control measures for increased governance over AI systems under frameworks like the National Institute of Standards and Technology AI Risk Management Framework, ISO 42001, PCI-DSS, GDPR, and central bank regulations [7]. It becomes challenging for governance to operate in a multi-cloud environment where the burden of security is shared between several parties, with varying degrees of compliance among them. It was reported that the governance models developed for single-cloud environments in the past tend to be less effective for the decentralized distribution of LLMs. Effective governance must be characterized by model transparency, access logging, quick validation, explainable decision-making systems, and consistent monitoring of compliance. Effective governance frameworks help banks gain credibility, reduce legal liabilities, and ensure that AI systems are compliant with internal security policies and external regulations.

Table 1: Literature Review Summary

Study	Focus Area	Key Contribution	Limitation
Reece et al. (2023)	Multi-cloud Risk Management	Applied STRIDE and DREAD models to rank cloud threats and identify IAM as the highest-risk area in distributed systems.	Limited discussion on LLM-specific risks.
George et al. (2024)	Banking + LLM + Multi-cloud	Examined the strategic use of LLMs across banking cloud platforms and highlighted governance concerns.	Focused more on adoption than technical security.
Yang et al. (2024)	Threat Modeling for LLMs	Proposed automated threat modeling for LLM applications used in sensitive sectors like finance.	Assumed mostly single-cloud deployment.
Rashid et al. (2026)	AI Lifecycle Security	Introduced unified security controls for training, deployment, and monitoring of AI systems.	Industrial focus with fewer banking examples.
ISACA (2026)	AI Governance in Multi-cloud	Emphasized governance-first security and policy consistency across cloud providers.	Limited practical BFSI case studies.

III. PROBLEM AND MOTIVATION

A. Problem Statement

Multi-cloud LLMs used by banks for a wide variety of applications like fraud detection/prevention, anti-money laundering (AML) compliance monitoring, customer support automation, and loan approval processes. However, these features come with serious cybersecurity concerns. The nature of the information stored on the cloud (such as security policies, identities management, encryption protocols etc.) varies depending on the cloud service provider. It's because each provider offers its own approach when it comes to access control mechanisms. The fragmentation of governance leads to poor governance, low-level transparency, and incoherence when managing financial AI-based systems [8]. Sensitive customer data, prompts, model predictions, and communication within the APIs is going to flow between multiple clouds [9]. There is a huge risk platform from which someone might violate the privacy rights of the individuals or break a certain law. That's why contemporary cybersecurity practices pay more attention to infrastructure-related matters while neglecting prompt-layer, model-based attacks, and other unique issues related to AI. If the banks do not ensure that they harmonize any governance measures across both public and private cloud networks, they will face a range of new threats, including regulatory infractions that lead to financial loss.

B. Motivation

The adoption of LLMs within banks should also come with better governance and security protocols compared to regular cloud computing. The adoption of AI in banks entails that the banks should abide by different regulatory processes like use of transparent and explainable models (PCI-DSS, GDPR, RBI guidelines and Basel standards). This means that there will be direct impacts of adopting AI in decision making when dealing with fraud detection, lending and compliance functions. In multi-cloud scenarios, it is harder to obtain these risks since the logging system and policies may come from several providers. On the contrary, decentralized governance poses a challenge when ensuring accountability and auditing of the security measures. The motive behind this research project is to narrow the widening gap between these two scenarios through understanding threat surfaces and creating governance-based recommendations to promote trust in the adoption of AI and ensure better cybersecurity in LLM-based banking applications operating in multi-clouds [10].

C. Main problems

1. Identity Inconsistencies Between Cloud Providers

The result is that banks do not have a robust authentication, authorization and privilege control as the identity and access management framework are specific to every cloud vendor [11]. This increases the chances of unauthorized accesses and insider threats.

2. Prompt injections and prompt leaks

Furthermore, bad prompting may corrupt model outputs, causing them to leak sensitive information related to the bank [12]. This directly impacts fraud detection systems and other sensitive customer information.

3. Model Drift & Model Registry

Improperly secured model registry may be exploited by adversaries for injecting polluted versions, leading to continuous model drift and inaccurate predictions [13].

4. Regulation Differences

There is a difference in regulation standards in terms of different cloud vendors and thus regions/jurisdictions [14]. This complicates matters for banks to establish uniform governance and audit mechanisms.

5. Third Party API Security

The majority of LLM solutions rely heavily on APIs and plugins which are susceptible to security vulnerabilities.

6. No Coordinated Response Mechanism

Security breaches in AWS, Azure and GCP clouds are often managed independently resulting in delays and inefficiencies in responding to cyberattacks [15].

7. Weak Audit Trails

With no consolidated log mechanism in place, banks cannot perform proper forensic analyses and reporting to regulators.

Table 2: Problem Analysis

Challenge	Impact	Banking Risk
IAM inconsistency	Weak access control	Unauthorized account access
Prompt injection	Manipulated model output	Fraudulent financial decisions
Data leakage	Exposure of confidential data	Compliance violations
Model poisoning	Incorrect fraud detection	Lending and risk assessment errors
API misconfiguration	Service compromise	Payment fraud and service disruption
Compliance fragmentation	Inconsistent reporting	Regulatory penalties
Weak monitoring	Delayed threat detection	Increased operational loss


Figure 1: Major Threat Severity in Multi-Cloud Banking LLM Systems

IV. RELATED WORK

Although there is extensive literature available on topics like cloud security, and AI governance of LLM protection in finance, very little research has been conducted regarding multi-cloud cybersecurity for banking-driven (or supply intelligent) LLM implementations. In the context of multi-cloud threat modelling using STRIDE and DREAD, Reece et al. highlighted Identity & Access Management (IAM) as the most vulnerable Section within distributed cloud infrastructure environments [16]. While providing comprehensive infrastructure-based security analyses, the authors fail to consider prompt-layer security vulnerabilities and models. Yang et al. proposed Threat Modelling-LLM, which is a novel automatic threat modelling approach for automatically recognizing potential risks posed to LLM applications, including the detection of spoofed input injection and vector-space attacks resulting in data leakage within critical industries, such as the banking sector. Nevertheless, the framework can primarily be used in single-cloud environments. Rashid et al. presented AI security control mechanisms with a lifecycle-based approach, addressing all stages of implementation, including training, deployment, and monitoring. The researchers focus their attention on practical applications in the industrial sector and overlook significant governance aspects relevant to the BFSI domain.




Figure 2: Related work

V. METHODOLOGY / PROPOSED FRAMEWORK

A. Governance-Centric Secure Multi-Cloud Framework

For our purposes, this paper describes the design and implementation of a cybersecurity architecture for deploying LLMs within multi-cloud banking settings that are focused on governance [17]. The purpose of such an architecture is to ensure that security, compliance, transparency, and operational trust are maintained within distributed infrastructures comprising AWS, Azure, and GCP. Such governance layers are built directly into prompts, models, vector DBs, APIs, and inference pipelines rather than only being added as protective layers to the underlying infrastructure. The governance engine will run centrally, continuously grooming identity access control, prompt behavior management, model management, and compliance reporting across all cloud providers.

B. Elements of the framework

1. Identity Governance Layer

This layer offers you authentication, authorization and privilege management across all your cloud platforms with Zero Trust architecture and centralized Identity & Access Management (IAM). RBAC - Adopt RBAC, MFA and privileged account management to avoid any unauthorized use and possible insider threats to the network [18].

2. Security Layer Prompt

Prompt Safety The purpose of prompt safety is to protect against prompt injections, where an attacker uses a prompt injection technique instead of the desired prompt to manipulate the output. In addition to this, it would also involve leaking any statistics or information from the grave yard prompts [19]. Therefore, the constraints for such interactions include firewall, input validation, prompt filtering and behavior monitoring to ensure that such interaction does not lead to any leakage of banking information by the LLMs

3. Governance Model Layer

This layer ensures security for model registries, training pipeline and inference engine. Here you can implement model integrity verification, preventing changes to the deployed model, change detection, model versioning, model validation and detection of poisoned model (mitigation of hallucinations) [20].

4. Compliance & Regulatory Layer

This layer represents all the activities related to the AI that need to be done as part of the PCI-DSS compliance, GDPR guidelines, RBI guidelines, and Basel compliances and standards. Proof of Compliance, Continuous Auditing, Explain ability Reporting, and Regulatory Traceability enhance the governance function to take updates on any compliance breaches [21].

5. Unified Observability Layer

It will help in creating links and monitor the security incidents in terms of file activity, access logging, and output generated by models in multiple clouds interconnected through cross-cloud SIEM system [22]. It helps in enhancing the incident detection, forensics, and real-time response mechanism.

C. Proposed Framework Workflow

Via banking applications like fraud detection, AML systems, or even chatbots performing customer support, the prompts will be generated and pushed to the LLM engine. Prompt Security Layer – This layer validates prompts as well as identifies any malicious prompts. Verified prompts are then sent to the Governance Engine, which performs identity verification, policy validation, and compliance matching. Afterward, prompt request is subjected to the Model Governance Layer, where model integrity and inference safety is ensured. Lastly, any anomalous response is generated in isolation and logged in the Unified Observability Layer [23].

Table 3: Proposed Framework Components and Functions

Framework Layer	Primary Function	Security Benefit
Identity Governance Layer	IAM, RBAC, MFA, Zero Trust	Prevents unauthorized access
Prompt Security Layer	Prompt filtering and validation	Stops prompt injection attacks
Model Governance Layer	Model registry validation	Prevents poisoning and drift
Compliance Layer	Regulatory mapping and auditing	Ensures legal compliance
Unified Observability Layer	SIEM + centralized monitoring	Improves detection and response




Figure 3: Proposed Multi-Cloud Governance Framework

VI. EXPERIMENT / CASE STUDY / SIMULATION

A. Case Study Scenario

One of the big banks has developed an AML and fraud detection system using a large language model deployed in a multi-cloud architecture (AWS + Azure). The customer's transaction logs, compliance data, suspicious activities reports and risk assessment data are all processed through the same retrieval-augmented generation (RAG) pipelines. The language model was used to generate risk alerts, suspicious transaction reports and compliance suggestions. Since the telecom case study involves running of the system in several cloud infrastructures, identity management security, prompt security and regulatory compliance become important concerns. In practice, many financial firms face a similar challenge when trying to ensure both security and compliance of critical business information and customer data [24].

B. Threat Simulation

Assessing the effectiveness of the proposed governance-based architecture is done by launching a prompt injection attack against the AML model using the information input from the perpetrator customers. This attempt tries to confuse the language model in order to overlook the rules for detecting any fraudulent behavior by introducing false prompts that manipulate the scoring of any suspicious transaction [25]. The other types were unauthorized API access, model poisoning during re-training, and privilege escalation attacks using clouds. Prompt Security Layer for Adversarial Instructions Identity Governance Layer ensures the legitimacy of access. The Model Governance Layer ensures the authenticity of the model before making inference, while the Compliance Layer keeps track of all the actions for investigation purposes.

C. Metrics of evaluation

The experiments evaluate the system performance from different aspects of cybersecurity and governance metrics [26]. Prompt attack detection precision means if the current prompt to send to the LLM engine is really malicious. Litmus Test – Unauthorized Access Alert represents how effective is the IAM and Zero Trust framework. Audit Coverage reflects compliance with regulatory logging requirements. Mean Incident Response Time or mean time to detect (MTTD) represents system capability in detecting the attacks and taking actions on time. Poisoning and Model Drift Prevention: Metrics to validate the model integrity. These metrics can give us an assessment based on both governance and technical evaluation of multi-cloud banking LLM systems protected by the framework.

D. Results & Discussion

Findings reveal that the suggested model performs well when improving performance. Prompt detection of injection was achieved between 61% to 95%, and unauthorized access discovery was recorded at 58% to 93%. In line with some improvements regarding regulatory requirements, compliance audit reached 65% to 97%. Response time for security incidents was significantly reduced from 14 minutes to 4 minutes due to centralization and cross-cloud SIEM integrations. Moreover, there was a considerable increase in validation accuracy models, reducing chances of making wrong judgments when lending and detecting errors. Overall, these results indicate that the security architecture-based design offers higher security than the traditional approach based on the infrastructure perspective [27].

Table 4: Simulation Results Comparison

Metric	Without Framework	With Proposed Framework
Prompt Attack Detection	61%	95%
Unauthorized Access Alerts	58%	93%
Compliance Audit Coverage	52%	97%
Mean Response Time	14 min	4 min
Model Integrity Validation	64%	96%
Fraud Decision Accuracy	69%	94%



Figure 4: Simulation Workflow for Secure Multi-Cloud Banking LLM System

VII. CONCLUSION

LLM applications within multi-cloud platforms have revolutionized banking operations through enabling the detection of fraudulent transactions, anti-money laundering and compliance checking, automation of customer services and financial decision making. Nevertheless, there is an array of cybersecurity and governance issues associated with the digital transformation. Identity management and access control limitations, prompt injection threats, model poisoning attacks and API insecurity are some of the challenges that arise from multi-cloud deployments. Traditional security mechanisms that protect infrastructure from malicious attacks cannot provide adequate protection for AI-based banking systems. In this study, we define the core threat domains of multi-cloud LLMs in banking applications and introduce a governance-based cybersecurity architecture that includes identity governance, prompt security, model validation, compliance audits and comprehensive observability. Our empirical evaluation reveals satisfactory improvement in terms of attack detection, compliance coverage and response management.

REFERENCES

- [1] M. F. Ansari, A. Mohammed, K. R. Janamolla, S. W. Khadri, S. A. Khader and M. A. Raheem, "The Double-Edged Sword: Navigating AI's Role in Banking Cybersecurity," 2025 International Conference on Transformative Computing Technologies (ICTCT), Bangalore, India, 2025, pp. 294-300, doi: 10.1109/ICTCT69201.2025.00062
- [2] KASHIF, MOHAMMED, ABDUL RAHMAN JIBRAN SYED, and MUBASHIR ALI AHMED. "ADVANCING ANOMALY AND FRAUD DETECTION IN BIG DATA WITH ARTIFICIAL INTELLIGENCE."
- [3] Khadri, Waheeduddin, Janamolla Kavitha Reddy, Abubakar Mohammed, and T. Kiruthiga. "The Smart Banking Automation for High Rated Financial Transactions using Deep Learning." In 2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC), pp. 686-692. IEEE, 2024.
- [4] Ansari, Meraj Farheen, and Syed Sharik Ali. "AI-driven zero-trust architecture for enhanced cybersecurity in dynamic network environments." International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering 13 (2025): 12.
- [5] RAHEEM, MOHD ABDUL, and MOHAMMED AZMATH ANSARI. "INTELLIGENT AND TRUSTWORTHY 6G: AI-DRIVEN ARCHITECTURES, APPLICATIONS, AND SECURITY FRAMEWORKS."
- [6] Reddy, Balavardhan, and Amir Ahmed Ansari. "AI-ENHANCED NETWORK TRAFFIC ANALYSIS FOR PREVENTING FRAUD PAYMENT IN BANKS."
- [7] AI, NIST. "Artificial intelligence risk management framework (AI RMF 1.0)." URL: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai> (2023): 100-1.
- [8] Lourenço, Rui Pedro. "Government transparency: Monitoring public policy accumulation and administrative overload." *Government Information Quarterly* 40, no. 1 (2023): 101762.

- [9] Verma, Manish. "The Future of Data Science in the Internet of Everything (IoE) by Prompt Science Analysis." *Future* 8, no. 5 (2024).
- [10] Ansari, Mohammed Azmath, Shaik Aqheel Pasha, and Narendar Kandula. "Reinforcement Learning for Adaptive Network Defense in Financial Institutions."
- [11] Chittoju, S. R., and Siraj Farheen Ansari. "Blockchain's Evolution in Financial Services: Enhancing Security, Transparency, and Operational Efficiency." *International Journal of Advanced Research in Computer and Communication Engineering* 13, no. 12 (2024): 1-5.
- [12] Liu, Chris Y., Yaxuan Wang, Jeffrey Flanigan, and Yang Liu. "Large language model unlearning via embedding-corrupted prompts." *Advances in Neural Information Processing Systems* 37 (2024): 118198-118266.
- [13] Dong, Tim, Shubhra Sinha, Ben Zhai, Daniel Fudulu, Jeremy Chan, Pradeep Narayan, Andy Judge et al. "Performance drift in machine learning models for cardiac surgery risk prediction: retrospective analysis." *JMIRx Med* 5, no. 1 (2024): e45973.
- [14] Tyra, Alexandra T., Siobhan M. Griffin, Thomas A. Fergus, and Annie T. Ginty. "Individual differences in emotion regulation prospectively predict early COVID-19 related acute stress." *Journal of Anxiety Disorders* 81 (2021): 102411.
- [15] Khaja, Moin Uddin, and Balavardhan Reddy. "Securing Agentic AI in Software-Defined Networks: A Policy-Driven Framework for Governance, Monitoring, and Incident."
- [16] Kashif, M., & Ansari, A. A. (2026). Building a unified AI-driven analytics pipeline for real-time anomaly detection in high-velocity data streams. *IJREEICE*, 14(1), 66–75. <https://doi.org/10.17148/ijreeice.2026.14111>
- [17] Janamolla, Kavitha, Ghousia Sultana Sultana, Fnu Mohammed Aasimuddin, Abdul Faisal Mohammed, and Fnu Shaik Aqheel Pasha Pasha. "Integrating Blockchain and AI for Efficient Trade Exception Handling: A Case Study in Cross-Border Settlements." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 24-30.
- [18] Ghadge, Nikhil. "Optimizing Identity Management: Key Strategies for Effective Governance and Administration." *International Journal of Security, Privacy and Trust Management* 13, no. 3 (2024): 01-11.
- [19] Khader, Shuaib Abdul, Amir Ahmed Ansari, and Syed Sharik Ali. "Zero-Day Exploit Prediction Using Graph-Based Deep Learning on Vulnerability and Threat Intelligence Data."
- [20] Bruno, Alessandro, Pier Luigi Mazzeo, Aladine Chetouani, Marouane Tliba, and Mohamed Amine Kerkouri. "Insights into Classifying and Mitigating LLMs' Hallucinations." *arXiv preprint arXiv:2311.08117* (2023).
- [21] Mohammed, Abdul Faisal, and Mohammed Akifuddin Ghori. "AI-Enhanced Safety for Heavy Load Construction Vehicles: An Integrated Embedded C++ Software Approach."
- [22] Ryu, Jung-Hwa, Seo-Yi Kim, Ri-Yeong Kim, Yeeun Kim, Seongmin Kim, and Il-Gu Lee. "Interoperable Security Information and Event Management Framework for Multi-cloud Environment." In *International Conference on Security and Privacy in Communication Systems*, pp. 61-83. Cham: Springer Nature Switzerland, 2024.
- [23] Kashif, Mohammed, Mohammed Aasimuddin, Mubashir Ali Ahmed, Laxmi Bhavani Cheekatimalla, Eraj Farheen Ansari, and Ahwan Mishra. "AI-DRIVEN CTI FOR BUSINESS: EMERGING THREATS, ATTACK STRATEGIES, AND DEFENSIVE MEASURES."
- [24] Sultana, Ghousia, Siraj Farheen Ansari, Mohammed Imran Ahmed, Abdul Faiyaz Shaik, Moin Uddin Khaja, and Bibhu Dash. "RESPONSIBLE AI ANALYTICS FOR REAL-WORLD IMPACT: NAVIGATING ETHICS, PRIVACY AND TRUST."
- [25] Clasen, Mathias, Jens Kjeldgaard-Christiansen, and John A. Johnson. "Horror, personality, and threat simulation: A survey on the psychology of scary media." *Evolutionary Behavioral Sciences* 14, no. 3 (2020): 213.
- [26] Melaku, Henock Mulugeta. "A dynamic and adaptive cybersecurity governance framework." *Journal of Cybersecurity and Privacy* 3, no. 3 (2023): 327-350.
- [27] Mohammed, Akheel, Zubair Ahmed Mohammed, Naveed Uddin Mohammed, Shravan Kumar Gunda, Mohammed Azmath Ansari, and Mohd Abdul Raheem. "AI-NATIVE WIRELESS NETWORKS: TRANSFORMING CONNECTIVITY, EFFICIENCY, AND AUTONOMY FOR 5G/6G AND BEYOND