

Speaking notes of the keynote address of Sh. Ananth Narayan, WTM, SEBI, on ***How Technology is Reshaping the Role of Regulatory Authorities***, at ET Legal World dialogue, 2025 – 12th June 2025

Good evening, to the many distinguished luminaries and participants in the ET Legal World conference. I thank the organizers of this timely and topical conference for the opportunity to be here amongst you.

Bill Gates once observed, “The first rule of any technology used in business is that automation applied to an efficient operation will magnify the efficiency. The second is that automation applied to an inefficient operation will magnify the inefficiency”.

Those words resonate deeply with every regulator. Our mandate is clear: as artificial intelligence sweeps across finance and law, we must collectively ensure that the magnification effect strengthens the delivery of our mandate of sustained capital formation, and does not detract from it.

In general, at the Securities and Exchange Board of India, we speak of two kinds of errors. A *Type I* error is the failure to detect or prevent wrongdoing: an insider-trading nexus we miss, a mis-selling scandal or cyber breach that hollows out investor faith. A *Type II* error is the opposite: in our zeal to prevent mischief we erect guard-rails so cumbersome, that legitimate innovation never lifts off. One error destroys confidence; the other smothers progress.

Striking the balance between them is too important to be left to regulators alone. It demands co-creation of regulations and the active participation of the lawyers, technologists, entrepreneurs and other stakeholders in this hall, as trusted advisors. Today I will outline how technology—particularly Artificial Intelligence and Large Language Models—is reshaping the way regulatory authorities perform their role, the compliance expectations we set, and the guard-rails we must collectively build. I will draw on SEBI’s own journey, discuss the opportunities, lay out some of the risks, and describe the principles guiding our next steps.

SEBI's Pre-AI Tech Journey

Let me start with SEBI's Pre-AI Tech Journey, where our initiatives have laid the foundation for AI adoption in our Reg-Tech and Sup-Tech.

On Market surveillance: SEBI has invested in integrated market surveillance systems, data warehouses, business intelligence tools, data lakes and high-speed analytics to detect anomalous trading patterns. Algorithms—continuously refined and evolved by SEBI and exchange officers—ingest billions of order and trade messages daily, analyse them and throw up aberration alerts. With terabytes of data now at our disposal, the next stop is to harness AI/ML to detect patterns and alerts that help maintain trust in the ecosystem.

On Compliance and reporting: The SEBI Intermediary Portal has moved registration and correspondences online. By replacing mountains of paper with structured data, we slashed processing times and created a searchable record—an essential precursor for effective AI.

On investor grievance redressal: Timely analysis of complaints can serve as an early warning to systemic issues in the ecosystem. Our revamped SCORES platform converts complaints into data workflows. SCORES has prepared us for the next wave of predictive analytics—identifying systemic issues before they erupt.

Automated inspections: Working with exchanges, clearing corporations, and depositories, we have automated large portions of routine inspections with significant investments into Sup-tech. This Sup-tech database also lends itself to analytics by AI/ML.

Investigations: The journey of any case from the start of an investigation until the issuance of an order is digitally captured in SEBI's Case Management System, creating a repository of violators, violations, and regulatory actions taken. This repository is now ready for deeper analysis.

Taken together, all these foundational initiatives have established the digital plumbing for AI to work through.

Riding the AI Wave: Opportunities and Challenges

Algorithmic trading, robo-advice, and smart KYC are now mainstream realities. AI promises lower costs, hyper-personalised products, and faster inclusion. Yet the same algorithms can accelerate mis-selling, data breaches, and flash-crash-like dislocations. SEBI's approach has therefore been calibrated and incremental, weighing both Type I and Type II errors.

SEBI's 2019 AI/ML Reporting Framework: It requires that every intermediary deploying AI in trading, advisory, or customer interfaces must report the model's purpose, data sources, and risk controls. Sunlight is indeed the best disinfectant.

With the February 2025 Accountability Amendments: We have now hard-wired accountability into multiple regulations: if you deploy an AI tool, *you remain wholly responsible* for its data privacy, outputs, and compliance with all applicable laws, regardless of vendors or open-source code.

Within SEBI, internal experimentation is underway: We are leveraging large-language models deployed on-premises to process bulky documents such as DRHPs and PPMs. Early results suggest a significant reduction in manual effort, yet every output still passes through human review, reinforcing the principle that machines assist, not replace, judgment.

Our own SEBI investor-education app: SAARTHI 2.0. now carries an AI-enabled chatbot that answers queries and explains investment concepts to retail investors.

Regulatory and Compliance Implications of AI Adoption

First, we aim for principle-based rules, anchored by clear accountability. Technology evolves too quickly for line-by-line prescriptions. Instead, we will set outcomes - fairness, transparency, data-protection and pin responsibility on the entity that deploys

and profits from the tool. This will hopefully keep the rulebook future-proof and yet enforceable.

Second, we need robust data governance. India's Digital Personal Data Protection Act is the umbrella under which all financial-sector AI must operate. Consent, purpose limitation, and the right to be forgotten are *not* optional extras. They are design requirements that need to be baked in from day zero.

Third, we must have auditability and explainability. Black-box models cannot be a black hole for accountability. Intermediaries using AI systems must maintain “model logs” that capture inputs, parameters, and decisions. If an investor alleges bias or a flash crash triggers an investigation, we must be able to reconstruct the algorithmic chain of events—step by step. We should strive to move towards explainable AI systems⁵⁷.

Finally, we need cyber-resilience and systemic stability. AI systems feed on data; they are also vulnerable to adversarial data poisoning and supply-chain attacks⁵⁹. Our cyber-security framework now compels real-time monitoring, segregation of training and production environments, and immutable audit trails⁶⁰.

The Way Forward: Innovation Sandboxes and Collaborative Oversight calibrated experimentation

SEBI's Innovation Sandbox: It allows start-ups, brokers, and academic labs to trial AI-enabled products on anonymised or synthetic market data. These products can then be tested in a live but contained environment under the Regulatory sandbox. By observing these trials, regulators learn alongside innovators, reducing the friction between rules and progress.

Industry collaboration: Given the rapid pace of AI development, there is perhaps a case for a AI-Risk Consortium of exchanges, clearing corporations, fintech associations, and academic experts could deal with regulatory aspects like shared taxonomies for incidents, stress scenarios, and open-source testing libraries.

Capacity building: Rules without expertise are like radars without screens. SEBI has therefore been recruiting more staff with IT skills while also upskilling existing staff through online courses and expert sessions. Regulators must speak the same mathematical language as the models they supervise.

International alignment: AI risk travels at fibre-optic speed across borders. We actively engage with IOSCO, the FSB, and the Global Financial Innovation Network to harmonise disclosure templates and incident-reporting protocols. Divergent rules only create regulatory arbitrage and weaken everyone. We are currently working on a consultation paper on AI guardrails for the Indian securities market based on IOSCO recommendations for AI use by intermediaries.

What We Expect from Market Participants

Know Your Algorithm. Perform rigorous pre-deployment testing—bias audits, adversarial pen-tests, scenario analysis. Document it, sign it, own it.

Disclose it to investors. If using AI in any investor-facing application, inform the investors.

Explain It to Investors. If a robo-advisor rebalances a portfolio, disclose the inputs in plain language, such as: “We use your age, income, and market volatility”. Clarity is not merely ethical; it is now regulatory.

Respect user autonomy. AI systems should be designed so users can override their decisions.

Monitor in Real Time. Models drift; markets mutate. Continuous validation is not a “best practice”—it is a survival necessity.

Report Material Incidents Promptly. If an AI-driven trading strategy malfunctions or a data leak occurs, the clock for regulatory notification starts immediately. Early transparency often turns a crisis into a solvable problem.

And finally, engage with regulators. We acknowledge the risk that an excessive focus on *Type I* errors can open us up to *Type II* errors. Please engage with us, as trusted advisors, to get the balance right. Even as you point out the *Type II* errors and seek relaxations, you must explain and think through how *Type I* errors can be avoided or kept to an acceptable level. If we lose credibility and mutual trust, the capital markets ecosystem will suffer immensely—we must not kill the goose that lays the golden eggs.

Reflections Beyond Finance: The Legal Sector's Parallel Journey

While my focus has been capital markets, the same forces are reshaping legal practice. Contract-review bots, predictive sentencing tools, and generative drafting assistants all raise similar questions: Who audits the model? Who owns the hallucination? Which confidentiality norms apply when your brief is processed by an external interface?

Jurisdictions worldwide are gravitating toward three universal principles:

First – human-in-the-loop accountability: No final legal opinion or regulatory order may be issued without a qualified professional's sign-off. Some prefer “Machine in the loop” to emphasize that humans are and will always be in charge and accountable, while machines assist or refine human decisions.

Second – Data provenance: Training a model on client-privileged material without explicit consent is a breach, period.

Third – Audit trails: Every prompt, output, and subsequent edit must be logged for potential discovery.

SEBI's own legal affairs department is experimenting under these foundational principles.

Conclusion: A Call to Collective Stewardship

Technological change is indifferent to the ambitions of regulators or the anxieties of incumbents; it advances because someone, or now something, somewhere, writes the next line of code. Our responsibility is to ensure that the line of code contributes to inclusive growth, efficient markets, and protected investors.

Let me return to Bill Gates: automation magnifies the underlying state of a system—good or bad. SEBI's mission, and that of every forward-looking regulatory authority, is to ensure the underlying state is *good*—anchored in robust governance, transparent markets, ethical data practices, and a relentless focus on investor welfare.

For that, I emphasise, we need partnership. Traditionally, regulators provide the guard-rails, industry provides the engine, academia provides the map, and investors provide the destination. However, if we move together and co-create—openly, responsibly, and innovatively—AI will not be a net disruptor of trust but a multiplier of it.
