


Hari Kishan

Director of Cloud Engineering · Enterprise AI Architect · Author · Researcher

- Author, The Thinking Voice
- 5 published papers across GenAI, voice agents and planning robustness
- Speaker across AI4, ClueCon, RE•WORK, Corinium and executive forums
- Focused on agentic AI, Voice AI, cloud architecture and governance in regulated enterprises

“Technology becomes transformational only when intelligence meets responsibility.”



Scan for my book, research &
speaking

thethinkingvoice.com


About me

Governing AI Agents in Finance


Tactics, strategies and tools for a new era

**THE FASTEST SYSTEMS IN THE WORLD
ARE FAST BECAUSE THEY ARE CONTROLLABLE.**

F1 drivers do not win by avoiding the brakes.
They win by knowing exactly how late they can use them.



The moment AI can act, governance changes



A wrong answer is a model problem. An unauthorized action is a governance problem.

2:17 A.M. The agent is still working.

- 01

————

Detects

\$250M in idle cash
- 02

————

Reasons

43 bps better yield elsewhere
- 03

————

Acts

Prepares transfers across 17 accounts
- 04

————

Explains

“I optimized liquidity.”

THE QUESTION IS NOT

“Was the AI right?”

THE QUESTION IS

“Who gave it permission?”

Intelligence does not equal authority

A teenager may understand exactly how a Ferrari works.

That does not mean you hand over the keys and say,
“See you Sunday.”

**Authority should be based on identity, context,
consequence, reversibility and limits.**

MODEL CAPABILITY

≠

AGENT AUTHORITY

The world has noticed

FSB: Build the guardrails

12 sound practices finalized Oct. 15. Board and senior-management accountability, lifecycle risk management, human oversight, cyber risk, third-party dependencies.

G7: Crash-test the guardrails

Finance leaders agreed to tabletop exercises simulating AI-enabled cyberattacks against financial systems, testing resilience and response protocols.

Regulators: Show the controls

Recent scrutiny is moving toward agent limits, monitoring, approvals, auditability and shutdown mechanisms, not just model documentation.

AI governance is moving from policy documents into operational resilience.

Fast cars need more control, not less

Formula 1 has:

Track limits

Runoff areas

Pit-lane speed limits

Yellow / red flags


Race control

Telemetry


Does F1 have these controls
because it hates speed?

No.

It has them because the
cars are fast.




Assume something will eventually fail



BLAST RADIUS

If this agent fails right now, what is the maximum damage possible before intervention?



A \$10 hallucination and a \$10M action are not the same risk

\$10

Low-value credit
reversible
small blast radius

\$10,000,000

Treasury movement
potentially irreversible
high blast radius

Same model. Different authority. Completely different governance.


$RISK \approx \text{authority} \times \text{consequence} \times \text{reversibility}$

Governance has to become executable

HUMAN-SPEED GOVERNANCE



MACHINE-SPEED GOVERNANCE



Policy-as-code for delegated authority

Put governance in the path of action



Airport analogy: legitimate traveler, legitimate destination, still passes through checkpoints before boarding.

Do not give agents keys. Give them visas.

AGENT AUTHORIZATION PROFILE

Treasury-Rebalancer-07

OWNER	Global Treasury
CAN	Draft transfers ≤ \$500k
CANNOT	Add beneficiaries / cross-border
VALID	Weekdays 06:00–18:00 CT
ESCALATE	> \$250k or low confidence
KILL	Immediate credential revocation

A visa defines:

- Where you can go
- Why you can go
- For how long
- Under what conditions
- How authority is revoked

A key opens the door.
A visa defines the boundaries.


“Human in the loop” is not a strategy

**If a human approves everything, you have automation theatre.
If a human approves nothing, you have unmanaged autonomy.**

AUTO-EXECUTE	Low consequence, reversible, inside tight limits
AUTO + POST-REVIEW	Moderate consequence, strong evidence
PRE-APPROVAL	High value, customer impact, uncertainty
PROHIBIT	Irreversible or outside delegated authority

The goal is minimum necessary human friction.

Do not make incident response an archaeology project



Could you reconstruct a consequential agent decision six months from now?

If the answer requires 3 engineering teams, 4 vendors and 6 weeks, you do not have auditability. You have archaeology.

Boards need telemetry, not vanity metrics

AUTONOMY RATE

How much consequential work happens without pre-approval?

POLICY BLOCK RATE

How often do runtime controls stop or challenge actions?

PRIVILEGE DRIFT

Are effective permissions exceeding approved authority?

ORPHAN AGENTS

Which production agents have no accountable owner?

BLAST RADIUS

Maximum \$ / customers / systems one agent can affect

TIME TO KILL

Seconds to disable an agent and revoke credentials everywhere

Winning is not “how many agents did we deploy?”

Your turn: would you let the agent do this?

TREASURY AGENT REQUEST

Move \$4.8M

to a newly suggested money-
market fund

+37 bps yield
94% confidence

A EXECUTE

B ASK A HUMAN

C BLOCK

Then reveal:

- New beneficiary
- Fund not approved
- Cross-entity transfer
- Agent limit = \$500k
- Recommendation came from a webpage

BLOCK: not because the model is “bad,” but because the action exceeds delegated authority.

Demo Live

Let's run the exact same \$4.8M request through governance.

WATCH FOR THREE THINGS

1 What the agent wants to do

2 Which policy checks fire

3 What evidence is left behind

DEMO FLOW

1. Submit request
2. Evaluate authority
3. Score / classify risk
4. Block or escalate
5. Inspect audit trail

5–6 min demo

Do not collect frameworks. Assemble a control system.

BOARD / ENTERPRISE	FSB sound practices · ISO/IEC 42001	Who owns AI, what is the risk appetite, how is it governed?
AI RISK	NIST AI RMF · internal taxonomy	How do we map, measure and manage AI risk?
AGENT AUTHORITY	MAS SAFR · agent authorization patterns	What can this agent do right now, under what conditions?
REGULATION	EU AI Act · DORA · privacy · conduct · AML	What legal and sector obligations apply?

Steering wheel, brakes, tyres and telemetry solve different problems. So do these frameworks.

Sources: FSB 2026 · NIST AI RMF · ISO/IEC 42001 · MAS SAFR · EU AI Act and sector rules

A 3 phase path to governed autonomy

FIND THE CARS

Inventory agents, owners, tools,
credentials and consequential actions.
Remove orphan access.

DEFINE THE TRACK

Set authority profiles, risk tiers, transaction
limits, escalation triggers and prohibited
actions.

INSTALL RACE CONTROL

Runtime policy checks, evidence,
emergency shutdown, red-team tests and
board telemetry.

Do not start with autonomous everything. Start with safe autonomy inside known boundaries.

NO BRAKES. NO SPEED.

NO OVERSIGHT. NO SCALE.

Poor governance says “NO.”

Immature governance says “ASK US FIRST.”

Mature governance says “YES, WITHIN THESE
BOUNDARIES.”



Intelligence creates opportunity. Authority creates risk.
Governance creates scale.

THE QUESTION TO TAKE HOME

How much authority
have we given our AI
agents...

...and how quickly
could we take it
back?

thethinkingvoice.com

Selected references

Financial Stability Board	Sound Practices for Responsible Adoption of Artificial Intelligence, final report, 15 Oct 2026
Reuters	G7 to launch exercises to guard against AI cyber attacks, 16 Oct 2026
Reuters	Banking watchdogs ramp up AI scrutiny as stability fears mount, 13 Oct 2026
Monetary Authority of Singapore	Safeguards for Agentic Finance at Runtime (SAFR), 2026
NIST	AI Risk Management Framework and AI Resource Center
ISO	ISO/IEC 42001 Artificial Intelligence Management System
European Commission	AI Act Service Desk and enforcement timeline

Presentation claims should be interpreted in the context of the applicable jurisdiction, use case and institutional risk framework.