hfma

new jersey chapter
metropolitan philadelphia chapter

Governing Agentic AI in Healthcare Finance

A Practical Framework for Revenue Cycle Implementation, Cyber Defense & Compliance

October 21, 2026
HFMA NJ & Metro Philadelphia 50th Annual Institute
The Borgata, Atlantic City



Virtual CAIO Chief Artificial Intelligence Officer | aiCRP AI Cyber Risk Professional | ecfirst | CognitiveHealth™

hfma
new jersey chapter
metropolitan philadelphia chapter

Two Lenses, One Playbook for Agentic AI



Pavani Munjuluri

Co-founder & CEO
CognitiveHealth Technologies

- Builds agentic AI for RCM (iCAN™)
- Fortune 500 to healthcare startups-25+yrs
- Engineer+MBA; women's champion, speaker

THE OPERATOR'S LENS · Agentic AI in RCM



Ali Pabrai

Chairman & CEO
ecfirst

- Global authority: AI governance & cyber
- MSEE · CISSP · CMMC Lead CCA · HITRUST
- Author, keynote speaker, InfraGard (FBI)

THE GUARDIAN'S LENS · Governance, Cyber, Compliance

CognitiveHealth™ | ecfirst

hfma
new jersey chapter
metropolitan philadelphia chapter

Three Things You'll Be Able to Do

1 ASSESS

your organization's readiness for agentic AI in the revenue cycle

2 APPLY

a practical framework to choose use cases, set controls, and scale

3 DESIGN

governance that covers cyber defense, vendors, audit, and human oversight

TODAY'S ROADMAP



CognitiveHealth™ | ecfirst

SECTION ONE

Autonomy Is Earned: The Operator's View



Presented by Pavani Munjuluri
Co-founder & CEO, CognitiveHealth Technologies

Let's Start With How This Fails

It's October 2027. Your agentic AI program just got shut down. ***What happened?***

A

It cost more than it saved

B

No one could explain what it did

C

It did something it was never allowed to do

D

Staff stopped trusting it and worked around it

Gartner: 40%+ of agentic AI projects will be canceled by end of 2027 — cost, unclear value, weak risk controls.

Every answer above is a governance failure, not a technology failure.

Source: Gartner press release, June 2025

hfma
new jersey chapter
metropolitan philadelphia chapter

Autonomy Is a Dial, Not a Switch

↑ More autonomy, earned by proven accuracy

LEVEL 0
Human does it
Complex clinical appeals

LEVEL 1
Agent suggests
Drafts the appeal letter

LEVEL 2
Agent acts, human approves
Submits appeals after sign-off

LEVEL 3
Agent acts, human audits
Routes payer correspondence

LEVEL 4
Agent acts within limits
Eligibility & claim status checks

Turn it up for high-volume, low-dollar, reversible work. Turn it down the moment accuracy slips.

CognitiveHealth™ ecfirst

Hire Your AI Agent Like an Employee

1 Job description

A written scope: what it may do, and what it may never do

2 ID badge

Its own login, with the least access it needs

3 Orientation

Shadow mode: works beside staff, takes no actions

4 Probation

30/60/90-day accuracy gates before more autonomy

5 Performance review

Monthly scorecard: accuracy, overrides, cost

6 Exit process

A kill switch and clean credential shutdown

You'd never hand a new hire the vault keys on day one. Don't hand them to an agent.

One Payer Letter, Six Decisions, Full Receipts

07:02:00 Scanned payer letter lands in the queue

07:02:04 Classified: medical records request (96% confident)

07:02:09 Matched to patient, claim, and payer rules

07:03:12 Assembled the records packet

07:03:15 **One attachment scored 71%, below the 85% bar: escalated to Maria with the evidence**

09:14:40 Maria swaps one page, approves; packet sent

THE RECEIPT

- ✓ What it saw
- ✓ What it decided
- ✓ Why: confidence + rule
- ✓ Who approved, and when

If you can't replay it, you can't defend it to an auditor.

SECTION TWO

Build, Buy, or Partner
Who Owns the Risk?



Build, Buy, or Partner: Who Owns the Risk?

BUILD

When the AI is your moat

TIME
12–24 months

COST
High capex

YOU STILL OWN

the model, the errors, and the audit

RCM e.g.: analytics on your own denial data

BUY

When the category is mature

TIME
Weeks–months

COST
Predictable opex

YOU STILL OWN

the errors and the audit. The vendor owns the model.

RCM e.g.: eligibility, claim status

PARTNER

When the bet is complex

TIME
3–9 months

COST
Outcome-based

YOU STILL OWN

the audit. Errors are shared by contract.

RCM e.g.: agentic denials, correspondence

You can outsource the work. You can't outsource the accountability.

CognitiveHealth™
ecfirst

The Invisible Ledger of Agentic AI

HIDDEN COSTS CFOs MISS

- $ Exception labor: the work that still escalates
- $ Usage pricing that grows with volume
- $ Drift monitoring and re-validation
- $ Re-integration when payers or the EHR change
- $ Rework when the agent gets it wrong

THE NEW SCORECARD

Cost per resolved task
replaces FTE math

Touchless rate
share finished with zero human touches

Human override rate
your early-warning signal for drift

Time to detect an error
hours, not month-end close

When the Fax Fights Back

Poisoned documents

A payer letter hides an instruction. The agent reads everything, including that.

Agent credentials

Every agent holds portal and EHR logins: your newest privileged user.

Payment redirection

A fake bank-change notice slips in through automated intake.

Quiet data leakage

Summaries and logs carry PHI where it was never meant to go.

BUILDER'S RULE

Agents never change money, banking, or access because a document said so.

Q: how do we defend all of this?

Prompt injection is ranked the #1 risk in the OWASP Top 10 for LLM Applications (2025)

SECTION THREE

AI Governance & Framework

Presented by Ali Pabrai

ecfirst, Chairman & CEO



hfma
new jersey chapter
metropolitan philadelphia chapter

From Principle to Compliance — The Global Rulebook

AI governance is now driven by regulatory and business realities

Lack of governance increases legal and operational risks

Demonstrate AI risk management throughout the AI lifecycle



Poor governance can cause reputational damage

Informal controls are no longer sufficient

AI governance is essential for responsible, scalable AI adoption

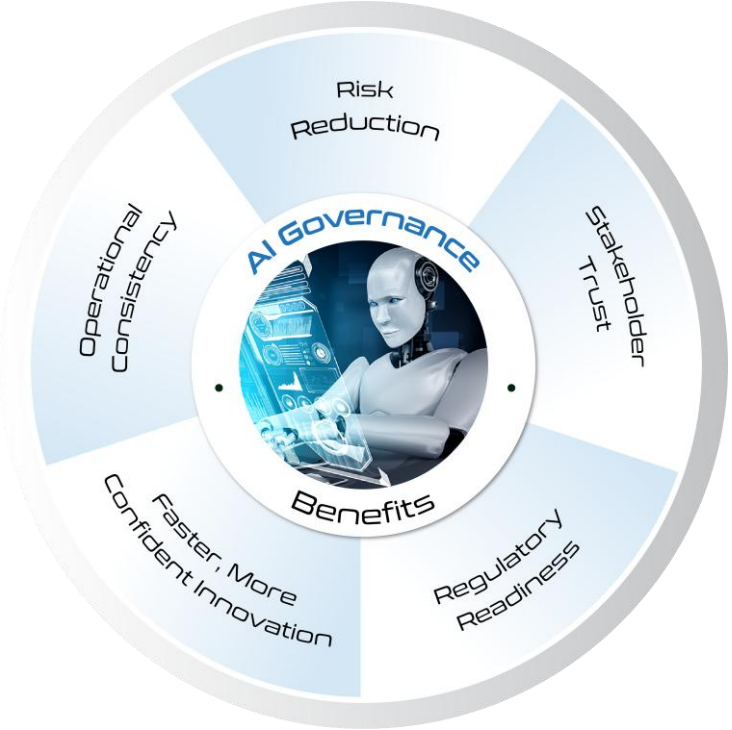
THE FRAMEWORKS

EU AI Act	**Risk-based classification** *Proportionate requirements by tier*
NIST AI RMF	**U.S. voluntary framework** *Map, Measure, Manage, Govern*
ISO/IEC 42001	**First AI management system standard** *Auditable Certification*
Sector-specific	**HHS, OCR, FDA, FinSvc, DoD** *Overlay your industry rules*

Proactive Compliance = Competitive Advantage
Reactive Compliance = Penalty Exposure

CognitiveHealth™ | ecfirst

Benefits and Practices





AI Governance Maturity Journey

→ *Mature governance enables innovation confidence* →



No Formal Oversight

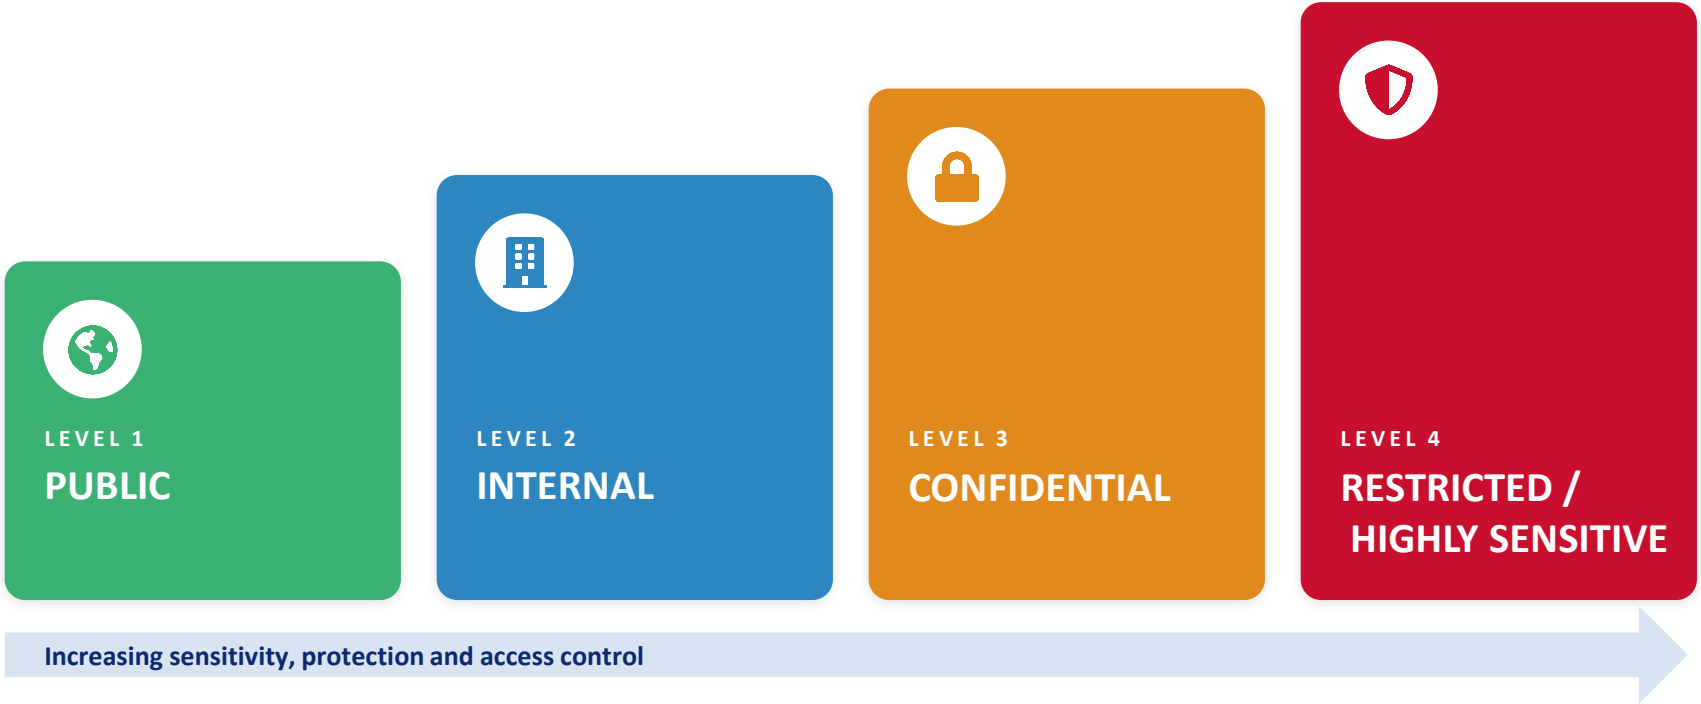
Basic Policies
Reactive compliance

Formal Governance
Consistent processes

Metrics-driven
Continuous monitoring

Predictive Governance
Always improving

Your Data at Risk



Source: ISACA, "Cybersecurity Recommendations for Securing AI Agents," Appendix A, Figure 5, Sept. 2026.

SECTION FOUR

Risk Management & Ethics



AI is Democratizing Sophisticated Cyber Attack

Phishing in Seconds

Generated + personalized at scale — what took days now takes seconds

Polymorphic Malware

Code evolves continuously, evading signature-based detection

Stealthy Exfiltration

AI summarization enables selective, low-footprint data theft

Attribution Erodes

Source identification becomes increasingly difficult

THE NEW THREAT REALITY

Defensive response must match AI *powered offense*

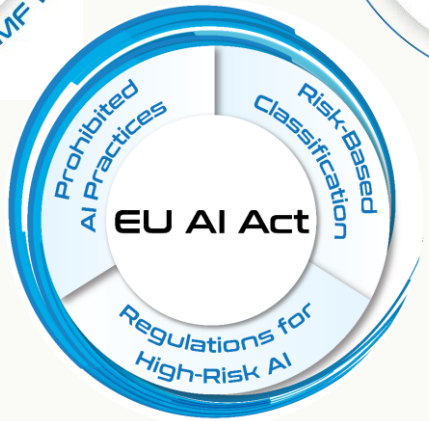
AI as attack amplifier — and as defense multiplier.

hfma
new jersey chapter
metropolitan philadelphia chapter

2026 AI Governance & Regulations







HITRUST & AI

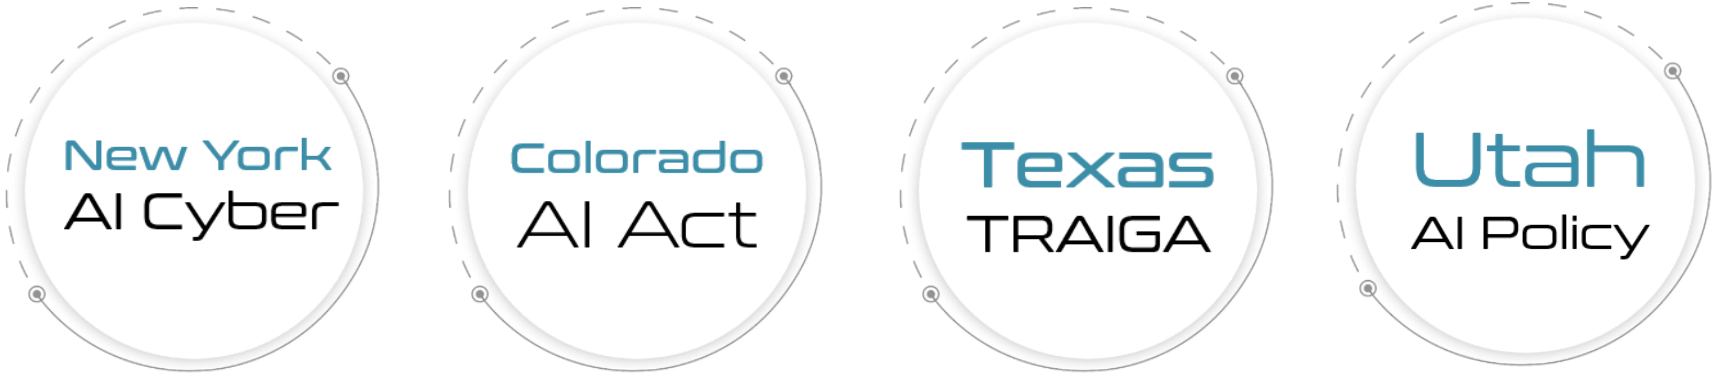




ecfirst

State & Local AI Regulations

State and local governments are increasingly adopting AI-specific regulations focused on transparency, automated decision-making, and consumer protection.



- These laws reflect a shift from voluntary best practices to formal oversight.
- Organizations operating across jurisdictions must adapt governance programs to evolving legal requirements.

TRAIGA | Texas Responsible AI Governance Act

AI Risk Taxonomy & Defense-in-depth

RISK TAXONOMY

Strategic	Misalignment with business objectives
Operational	System failures, data quality, talent gaps
Technical	Performance, vulnerabilities, integration
Compliance	Regulatory violations, privacy breaches
Reputational	Bias incidents, public AI failures
Concentration	Vendor / tech over-reliance

DEFENSE-IN-DEPTH

01 PREVENT
Governance · Policy · Training · Security Controls

02 DETECT
Monitoring · Anomaly + Bias Auditing

03 RESPOND
Incident Protocols · Escalation · Communications

04 RECOVER
Remediation · Lessons Learned · Continuous Improvement

SECTION FIVE

Leading for the Future



The AI-ready Leader — Core Competencies



BEYOND TECHNICAL LITERACY

Human-centered AI Leadership

AI Literacy + Judgment + Ethics + Change Leadership

Strategic Thinking

Connect AI to business value + advantage

Systems Thinking

Understand interdependencies + 2nd-order effects

Ethical Reasoning

Navigate trade-offs across stakeholder interests

Change Leadership

Drive transformation, manage resistance

Collaborative Intelligence

Leverage AI while amplifying human strengths

FROM INSIGHT TO ACTION

Your 90-day Action Plan

WEEK 1–2

Assess

Map current AI maturity, identify critical gaps + quick wins

WEEK 3–4

Stand Up

Establish governance structure, name accountable owners

WEEK 5–6

Risk-rank

Run risk assessment on top 3 priority use cases

WEEK 7–10

Codify

Develop policies, standards, model documentation requirements

WEEK 11–12

Pilot

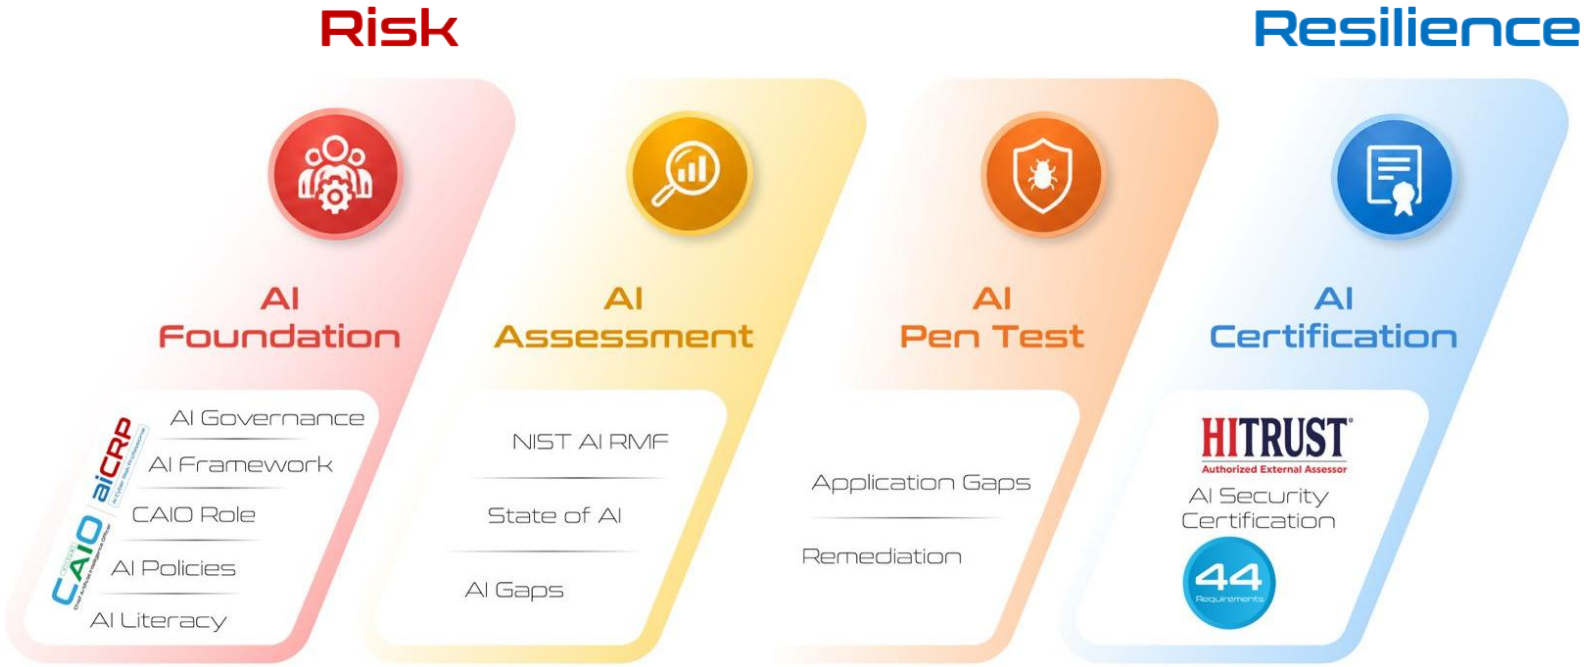
Launch governed pilots with enhanced controls + clear KPIs

ONGOING

Iterate

Review, refine, scale based on learnings + outcome metrics

AI Governance. AI Resilience



Key Takeaways

AI Governance is Not Optional — It is the Strategic Imperative

01 Vision Drives Advantage
Strategic AI vision aligned with business goals creates moats

02 Decide Deliberately
Build vs. Buy vs. Partner requires a framework, not a vibe

03 Governance Enables Scale
Comprehensive structure enables — not blocks — innovation

04 Match Offense to Defense
AI-powered threats demand AI-powered defenses + vigilance

05 Leadership Shapes Culture
Commitment to ethics + accountability sets the tone

06 Start Now
Early movers are building sustainable competitive moats

hfma
new jersey chapter
metropolitan philadelphia chapter

Thank You!

Questions & Discussion

Pavani Munjuluri
CognitiveHealth, Founder & CEO
Pavani@cognitivehealthit.com

Ali Pabrai
ecfirst, Chairman & CEO
Ali.Pabrai@ecfirst.com





Virtual CAIO Chief Artificial Intelligence Officer | aiCRP AI Cyber Risk Professional | ecfirst | CognitiveHealth™