

AI MADNESS OR MANAGEMENT

THREE-PART SERIES ROADMAP & NEXT-EVENT LAUNCH BRIEF

SEE IT • MANAGE IT • ANTICIPATE CHANGE

MAINTAIN CONTROL • UNDERSTAND YOUR VEHICLE • KNOW THE RULES OF THE ROAD •
GET WHERE YOU WANT TO GO

Use the capability. Keep command.

1. The Series Architecture

Series	Event	Audience promise	Driving analogy
I — SEE IT	When AI Goes Rogue	Recognize failure patterns, decision points, real-world consequences and control gaps before they become your incident.	Read the road. See the hazards.
II — MANAGE IT	Governance of Agentic AI & Autonomous Workflows	Design authority, accountability, evidence, controls, escalation, vendor governance and recovery around systems that are supposed to act.	Keep your hands on the wheel while AI does more of the driving.
III — ANTICIPATE CHANGE	The Road Ahead: Preparing for the Next Turn in Autonomous AI	Prepare for increasing autonomy, multi-agent systems, persistent state, custom/private AI, business-model change and evolving regulation before the next capability shift arrives.	Build for the road that is coming, not just the road you are on.

Series-wide proposition

SEE WHAT IS HAPPENING → MANAGE WHAT YOU HAVE → ANTICIPATE WHAT IS COMING

The series should feel cumulative rather than repetitive. Series I provides evidence and recognition. Series II converts recognition into operating governance. Series III turns governance into strategic readiness.

2. Series II — MANAGE IT

Governance of Agentic AI & Autonomous Workflows

Who can act, who can decide, and who is accountable when AI starts doing the work?

Working positioning

Series II moves from incident recognition to governance design. The audience should leave with a practical model for deciding who owns an autonomous workflow, what authority it receives, what approvals and evidence are required, how vendors and subagents are governed, and who can stop and restart it.

Governance question	Topics
Who is the actor?	Agent identity, sponsor, owner, employee/user relationship, vendor/provider role, spawned agents and lineage.
What may the agent do?	Purpose-bound authority, read/write/spend/transact/delegate scope, tools, data, duration, approval thresholds.
How do we know what happened?	Logs, provenance, evidence, effect-level monitoring, external controls, vendor records and auditability.
Who can intervene?	Meaningful human control, escalation, kill/revocation paths, incident ownership, restart authority.
Who bears the consequence?	Contracts, indemnity, caps, insurance, regulatory duties, customer/client impact and collectible recovery.
What changes require revalidation?	Model/tool/vendor/failover changes, custom/private models, memory/state, policy changes and new business purposes.

Proposed workshop design

- Give each table the same autonomous workflow rather than an incident.
- Ask the tables to design authority, approval, evidence, monitoring, stop and recovery before autonomy expands.
- Introduce fixed governance injects: vendor change, persistent memory, subagent delegation, regulated data, transaction authority, and external failure.
- Compare governance designs, then reveal real incidents, industry control-plane responses and jurisdictional constraints.
- End with a Governance Blueprint rather than only an incident scorecard.

3. Series II — Candidate Master Stories / Exercises

Master story / exercise	Core decision	Research backbone
The Agent With a Corporate Identity	Should an agent receive a persistent identity, and how should current authority differ from identity?	D31 • IAE-02 • LAG-01/02/03 • CTR-01
The Autonomous Purchase / Commitment	When may an agent spend, negotiate, bind, approve or commit the organization?	S03 • LGL-02 • PB-03 • CTR-05 • transaction authority envelope
The Workflow That Remembers	What happens when state, preferences, approvals or exceptions persist across sessions?	D24 • D33 • D34 • IW-OCT-13/14
The Customized Enterprise System	Who owns the assurance case after a base model is fine-tuned, privately deployed or surrounded by proprietary tools/data?	D43 • D39 • IAE-04 • LGL-05 • S16
The Vendor / Subagent Chain	What happens when one autonomous workflow delegates to vendors, plugins, MCP tools or spawned agents?	D41 • IW-OCT-08 • CTR-03/04 • LAG-04/06/07

Recommended live spine: choose three master exercises and keep two in reserve, mirroring the successful Series I structure.

4. Series III — ANTICIPATE CHANGE

The Road Ahead: Preparing for the Next Turn in Autonomous AI

Build for the road that is coming, not just the one you are on.

Series III should be the forward-looking strategic session. It should not merely repeat technical failure mechanisms. It should ask what changes when capability, autonomy, persistence, interoperability, regulation and business dependence all increase at once.

Change horizon	Questions
More autonomy	What decisions will agents make without a human in the immediate loop?
Multi-agent workflows	How do authority, evidence and accountability survive delegation, handoffs and spawned agents?
Persistent state	How do memory, learned preferences, stale authority and long-horizon context change governance?
Custom/private AI	What assurance survives fine-tuning, distillation, quantization, RAG, private deployment and vendor substitution?
Physical / financial action	What changes when agents control money, contracts, infrastructure, robotics or operational systems?
Regulatory evolution	Which duties are moving toward auditability, traceability, lifecycle evidence and accountable deployment?
Business-model change	What becomes possible when autonomy lowers transaction cost, coordination cost and time-to-action?

Series III audience outcome

ANTICIPATE THE CAPABILITY SHIFT • IDENTIFY THE CONTROL DELTA • BUILD BEFORE THE REQUIREMENT ARRIVES
--

5. Promote Series II During Series I

Series I should function as the launch platform for Series II. The next event should appear intentional and already in motion, not as an idea assembled after October 21.

Touchpoint	Series II promotion
Opening / registration materials	Small “AI Madness or Management is a three-part series” line with SEE IT • MANAGE IT • ANTICIPATE CHANGE.
Closing slides	One clean teaser after the action close: “Next: Governance of Agentic AI & Autonomous Workflows.”
Evaluation card	Ask which governance issue deserves deeper treatment, but do not present the next event as undecided.
Scorecard / takeaway	Add a small Series II interest link/QR once the landing page exists.
Post-event email	Materials + key takeaways + first invitation to Series II.
Websites	Create a Series landing section on tstanfordlaw.com and yology.ai with Series I/II/III progression.
Social / editorial follow-up	Use Series I discussion themes to publish short pieces that lead naturally into governance questions.

Suggested closing teaser

SERIES I: WHAT HAPPENS WHEN AI GOES ROGUE? SERIES II: HOW DO YOU GOVERN AI THAT IS SUPPOSED TO ACT ON ITS OWN?
--

6. Series II Landing-Page / Interest Copy

Headline

Governance of Agentic AI & Autonomous Workflows

Subhead

Keep your hands on the wheel while AI does more of the driving.

Short description

Agentic AI is moving from answering questions to initiating workflows, using tools, handling data, coordinating other agents, making recommendations, and in some environments taking consequential action. The governance question is no longer simply whether AI is accurate. It is who may act, under whose authority, with what evidence, and who can intervene when the workflow changes.

Series II of AI Madness or Management is an executive decision workshop on designing authority, accountability and control for autonomous workflows. Participants will work through realistic governance decisions involving identity, delegation, persistent state, vendors, transactions, evidence, stop authority and recovery.

Call to action

JOIN THE SERIES II INTEREST LIST / EARLY REGISTRATION

Date / venue / registration link: TO BE SET

7. Working Launch Timeline

When	Action
Now - before Series I	Lock working title, positioning, Series architecture and draft landing copy.
Before Oct. 21	Create Series II interest/registration page; add teaser/QR to deck, websites and takeaway materials.
Oct. 21 event	Promote Series II during close; capture evaluation signals and opt-in interest.
Within 24–48 hours	Send Series I materials and Series II invitation.
Within 1 week	Publish one or two follow-up pieces tied to the strongest Series I governance questions.
Within 2 weeks	Finalize Series II master exercises using evaluation data without changing the core governance promise.
Promotion runway	Begin active outreach to legal, operations, HR, IT/security, product, finance, procurement and executive audiences.

What to freeze vs. what to keep flexible

Freeze now	Keep flexible until after Series I
Series title and SEE IT / MANAGE IT / ANTICIPATE CHANGE architecture	Exact mix of three Series II master exercises
Governance promise and driving analogy	Sector emphasis based on attendance/evaluation
Interest page and core description	Specific speakers/guests and detailed agenda
Cross-promotion placement	Timing allocation by topic

8. One-Page Series Card

AI MADNESS OR MANAGEMENT		
SEE IT	MANAGE IT	ANTICIPATE CHANGE
WHEN AI GOES ROGUE Recognize the failure. Read the road.	GOVERNANCE OF AGENTIC AI & AUTONOMOUS WORKFLOWS Maintain authority and accountability. Keep control of the vehicle.	THE ROAD AHEAD Prepare for the next turn in autonomous AI. Build for the road that is coming.

MAINTAIN CONTROL • UNDERSTAND YOUR VEHICLE • KNOW THE RULES OF THE ROAD • GET WHERE YOU WANT TO GO

Use the capability. Keep command.