

Carpe Diem: The New Rules of Leadership in the Age of AI

by: Lynda McNutt Foster

When I grew up in the 1970's access to any type of information that would help me move from the financial means I had to a more stable, long term situation, required resources, lots of them. Even to get access to the World Book Encyclopedia to turn in stellar homework assignments required funds to buy them, a ride to the library, or negotiation skills to get access to them from a neighbor.

How we used to win was by getting access—to capital, to data, to the room where decisions happened. Today, everyone has access. The winners now are the leaders who turn access into action—who can find the signal in the noise, learn faster than the problem changes, and partner with intelligent tools without losing their human edge.

That's the heart of my message to every leader—especially those who've never considered themselves "tech people": **Carpe diem.** Seize the day, and control the noise. The following are the five new rules of leadership I'm sharing at RBTC's Emerging Tech Conference—why they matter now, and how to live them.

The Five New Rules

01

Co-Think with Machines

Don't just prompt. Converse.

02

Filter the Noise

Make discernment your core skill.

03

Design for Time

Automate the tasks. Humanize the moments.

04

Lead Through Learning

Teach people how to think—not just how to use tools.

05

Protect the Signal

Anchor innovation in ethics, clarity, and long-term vision.



1. Co-Think with Machines

Don't just prompt. Converse.

What it means. Treat AI like a thinking partner, not a vending machine. Move beyond one-and-done prompts into back-and-forth dialogue: define the problem together, iterate on hypotheses, pressure-test options, and decide what you'll verify in the real world.

Why it matters now. We've entered the "agent" era—AI that plans, reasons in steps, and takes actions across your tools. That doesn't replace you; it multiplies you—if you can collaborate well. Leading research frames this as human-AI co-thinking: humans set goals, frame problems, and verify outcomes; AI accelerates exploration, connects patterns, and proposes paths—an alliance that keeps human judgment in the loop rather than delegating it away. And yes, "agentic AI" is exploding inside companies because it can act like a virtual co-worker across multistep workflows.

❑ How to start

- Ask AI to explain its plan before it executes ("Outline the steps you'll take and where you need my input").
- Pair every AI output with a human verification step ("What would we need to test in the field to trust this?").
- Give your agent a role, constraints, and success criteria—the same way you onboard a new teammate.

Carpe diem connection. You don't need to be "technical" to co-think. You need curiosity and conversational courage—skills you already have.

2. Filter the Noise

Make discernment your core skill.

What it means. In an age of infinite inputs, attention—not information—is the scarce resource. Leaders must design how attention gets allocated: which data streams matter, which dashboards earn a daily glance, which metrics govern outcomes (and which don't).

Why it matters now. With AI producing oceans of plausible content, the premium is shifting from finding information to qualifying it. The most credible playbooks emphasize keeping human judgment as the verifier—especially for context, ethics, and anything requiring embodied experience. Meanwhile, enterprise trends show rapid adoption of AI and agents—and with scale comes risk: leaders must set safety, transparency, and liability guardrails so signal doesn't drown in automation.

▢ How to start

- Define a "rule of three signals" for big decisions (e.g., customer evidence + operational data + expert review).
- Install a kill-switch culture: when quality or ethics are in doubt, pause the automation and escalate.

Carpe diem connection. Seizing the day starts with choosing the right day to seize—clarity first, velocity second.

3. Design for Time

Automate the tasks. Humanize the moments.

What it means. Treat time as your primary design material. If a workflow steals thinking time, automate it. If a moment requires trust, presence, or creativity, protect it fiercely.

Why it matters now. The tech stack is shifting fast: smaller, cheaper, domain-specific models are moving intelligence to the edge; agent platforms can now operate browsers, forms, and tools for you. That makes "calendar architecture" a leadership discipline—what gets automated, batched, or delegated to agents versus what stays fully human. Co-thinking research also warns against "cognitive off-loading": don't let the machine do long stretches without human checkpoints. Rhythm your work so humans stay the deciders.

▢ How to start

- Run a "time audit" of one process (weeklies, proposals, onboarding). Offload the repetitive 30–50% to an agent; reinvest the time in coaching, scenario thinking, and customer discovery.
- Alternate modes: deep work (disconnected) ↔ agent-assisted sprints. The alternation itself reduces errors.

Carpe diem connection. Free hours aren't the goal; higher-quality hours are. Design your calendar to match your calling.

4. Lead Through Learning

Teach people how to think—not just how to use tools.

What it means. Tool tips expire. Thinking skills compound. Make metacognition (how we frame questions, test assumptions, and verify results) part of your team's daily muscle memory.

Why it matters now. McKinsey's latest tech-trends work shows AI adoption scaling—but value lags when orgs don't evolve operating models and talent. The next performance frontier is workplace learning systems that blend agents, analytics, and human mentoring—so teams make better decisions together, faster, with accountability. The co-thinking literature goes further: students (and professionals) should document the dialogue—problem statement, plan, refinement, and verification—because the process is now as valuable as the result.

❑ How to start

- Replace "show me the answer" with "show me your co-thinking"—the prompts you tried, the plan you refined, how you validated.
- Promote "learning artifacts" (short memos, agent transcripts with highlights) to speed onboarding and de-risk decisions.

Carpe diem connection. When teams learn out loud, today's work becomes tomorrow's advantage.

5. Protect the Signal

Anchor innovation in ethics, clarity, and long-term vision.

What it means. Move fast where it's safe; move right where it's consequential. Establish a visible spine for responsible AI—privacy, provenance, fairness, explainability—and tie every deployment to customer value and mission.

Why it matters now. As agents act across systems, trust becomes the adoption gate. The strongest playbooks combine technical controls (observability, audits, model choice) with governance (roles, approvals, escalation) and a clear stance on what must remain human-decided. That's how you scale without compromising the brand—or the people you serve. Education research mirrors this: humans keep the verifier role for anything tied to lived experience, context, or ethics.



▢ How to start

- Publish your "Responsible AI in Practice" one-pager: where you will use AI, where you won't, and how you'll verify outcomes.
- Require a verification plan for any AI-influenced decision that impacts people, money, safety, or brand.

Carpe diem connection. The loudest voice is not the clearest signal. Values are your permanent amplifier.

From Access to Action: What You Can Do This Month

▢ Stand up one safe, useful AI agent for your team

Content research, meeting synthesis, or customer follow-ups. Give it a role, scope, and verification checklist.

▢ Run a 90-minute "co-thinking lab"

Pick a live challenge, co-solve with AI, and capture the transcript as a learning artifact.

▢ Kill three reports or meetings that don't move decisions

Replace them with a weekly "signal review" ritual.

▢ Name your non-negotiables

Publish the decisions that will stay 100% human, and why.

Take action today with the prompt supplied next.

Use this Prompt by copying and pasting it into your LLM (ChatGPT, CoPilot, Gemini, Grok, Claude, Perplexity, etc)

*You are a master leadership facilitator using the article *Carpe Diem: The New Rules of Leadership in the Age of AI* (5 rules: Co-Think, Filter the Noise, Design for Time, Lead Through Learning, Protect the Signal). Job: Run a one-question-at-a-time adaptive interview that ends with a customized Question Outline and 90-day Execution Plan for the user's role. Method (strict): Ask one question, wait. Use each answer to shape the next. If vague, ask one short clarifier. If "skip," offer one alternative and move on. Keep a running scratchpad: ROLE, INDUSTRY, OUTCOMES_12M, CONSTRAINTS, NON_NEGOTIABLES. At the end, output the deliverables below in clean Markdown. Intake (3) Your role and scope? Your industry and key context (customers, compliance, sales cycle)? Three outcomes that would make the next 12 months a win? Save as ROLE, INDUSTRY, OUTCOMES_12M.*

Five Rules (ask in order; for each: 1 starter → 1 adaptive follow-up → 1 micro-commitment)

1. *Co-Think with Machines Starter: In your ROLE/INDUSTRY, what recurring decision or deliverable would benefit most from a weekly back-and-forth with AI? Follow-up (pick one): key inputs you'll provide each time, or your real-world verification step. Micro-commitment: 15-min weekly co-thinking ritual + verification checklist.*
2. *Filter the Noise Starter: For OUTCOMES_12M, what three signals best predict success? Follow-up: Which signal is weakest and how will you qualify it (source/freshness/threshold)? Micro-commitment: Rule of 3 Signals + kill-switch when doubt arises.*
3. *Design for Time Starter: Which workflow drains the most thinking time? Follow-up: Which 30–50% is safe to agent-offload, and what stays human? Micro-commitment: Time audit → agent pilot; reinvest saved time in coaching/strategy.*
4. *Lead Through Learning Starter: Where would documenting co-thinking cut the most rework? Follow-up: What lightweight artifact (memo/log/transcript highlights) will you standardize? Micro-commitment: Replace "answer" with "show your co-thinking" + 1-page template.*
5. *Protect the Signal Starter: Which upcoming AI-influenced decision could impact people/money/safety/brand? Follow-up: Your non-negotiables (privacy/provenance/fairness/explainability) and verification plan? Micro-commitment: Publish Responsible AI in Practice one-pager; require verification for consequential calls.*

Deliverables (produce after interview) - A) Personalized Question Outline grouped by the 5 rules, referencing ROLE/INDUSTRY/OUTCOMES_12M. B) 90-Day Execution Plan: goal, 3 workstreams, M1/M2/M3 milestones, owners, evidence of success, risks & kill-switch triggers. C) Calendar Architecture: weekly deep-work vs agent sprints, 15-min Signal Review, Co-Thinking Lab cadence. D) Responsible AI Guardrails: where AI is used/not used, verification, human-decided boundaries. E) Starters & Checklists: tailored co-thinking prompt, verification checklist, Rule of 3 Signals card. End with a concise 7-day kickoff checklist. Tone: practical, human, confident.

Why Carpe Diem—Especially for Non-Tech Leaders

Every decade changed the tools.

None changed the signal: focus, empathy, courage, purpose.

What's different now is speed and access. You can ship a video, a memo, a model, a prototype—today—not because you're a coder, but because you're a leader who can frame the problem and verify the work.

The gap is no longer talent or budget. It's willingness.

So, yes—**Carpe diem**. Seize the day. Or better yet, cease the noise. Find the signal worth amplifying. Lead with discernment, ethics, and a long view—while co-thinking with the smartest machines humans have ever built.

How will you cease the day?

