EF INTERNATIONAL ADVISORS

DISCUSSION BRIEF

AI and the Future of K–12 Education: From Control to Capacity in a Rapidly Changing World

For K–12 Educators, Leaders, Boards, Policymakers, and Cross-Sector Partners

As artificial intelligence becomes embedded in daily life and work, education systems face new expectations and new risks. This brief surfaces the big questions leaders must grapple with now—how capacity is built, how innovation is stewarded, and how learners are prepared for realities that are already emerging.

Published: December 13, 2025
Prepared by: E. Ted Fujimoto - Managing Director, EF International Advisors, LLC and Dr. Antonio Corrales - Director, Doctoral Program in Educational Leadership, University of Houston–Clear Lake

Introduction

We originally intended to release this discussion brief three months ago. Each iteration, however, felt less like a revision and more like a restart. That shift was driven by three realities: the speed at which AI capabilities have evolved; insights from early, real-world experiments integrating and mandating AI use—among the earliest such efforts in the country for educational leaders—within a strategic planning course across three cohorts in the Educational Leadership doctoral program at the University of Houston–Clear Lake; and lessons learned from our own internal use of AI at EF International Advisors.

Our intention is to help educators—teachers, leaders, board members, and policymakers—navigate an unprecedented moment in history. AI is now at the forefront of nearly every sector's consciousness, with real consequences for how work gets done, which skills become essential, and which skills rapidly lose value. This applies both to how education organizations are managed and to how learning happens in classrooms.

As we developed this brief, the pace of change raised more questions than answers. We realized that trying to resolve the question so many others are asking—how or whether to use AI in education, what the best practices are, or how to regulate classroom use—may actually be the wrong place to start. The landscape is shifting too quickly, and the impacts are too profound. What may be more helpful is a different framing altogether—one that guides decision-making and policy-making without pretending we can fully define or control the future.

That realization pushed us to look across sectors and into the messy, real world of how AI is actually being adopted.

Over the years, we have refined a process called Design Studio—a systems-level evolution of design thinking that unlocks team capacity by intentionally incorporating neuroscience. It has been transformative for hundreds of teams, more than half of them in education. A central focus of this work is understanding the brain's biases and barriers to deep understanding. When planning, people naturally gravitate toward what they understand and can control, while underestimating—or ignoring—what they do not.

AI exposes this bias clearly. None of us fully understand it—no one does, including many of those building it—and none of us truly control it. Even its creators openly debate what safeguards are needed. Faced with that uncertainty, the instinct is predictable: ban it, restrict it, or attempt to tightly control it.

While K–12 education appears especially challenged by this moment, virtually every sector is grappling with the same core tension at the board level. On one hand, leaders must figure out how to leverage AI to remain competitive and relevant—a time-sensitive issue requiring retooling, new understanding, and real capital investment. On the other hand, they must determine how to safeguard its use. Safeguarding is not the same as control or restriction. Every powerful technology requires safeguards—placed in the right hands, with the right capabilities—to generate benefit rather than harm.

We have airplane technology, but putting an aircraft in the hands of an unlicensed pilot would be irresponsible and dangerous. The answer is not banning planes or pilots. We have automobiles, along

with licensing, traffic laws, and safety standards. These are safeguards—not blanket restrictions on who can drive or where society is allowed to travel.

We invite you to consider this framing: AI is an automobile that is not going away. Our responsibility is to build a system of roads and highways that enables its use, with safeguards at critical intersections. We must equip educators, leaders, and students with the capabilities to use AI powerfully and responsibly, rather than leaving outcomes to chance. Some individuals—those with advanced skill and judgment—will take AI off the main roads and onto racetracks or off-road terrain to push boundaries and experiment. Even there, safeguards and specialized licenses apply, just in different forms.

Importantly, being licensed to drive does not mean understanding how to engineer a car or program its software. Most people do not know the math behind transmission design or vehicle control systems, yet that does not prevent them from being excellent drivers. Ironically, many educators and business leaders claim they are “not good at technology,” while seamlessly using smartphones, navigation apps, social media, ride-sharing, online banking, and voice assistants every day. I am old enough to remember when word processors required specialized training because of obtuse command structures. Today, no one thinks twice. Increasingly, I find myself speaking to my phone to do things that once required typing on a laptop. It raises a legitimate question: are we over-investing in certain legacy skills while under-investing in more critical ones?

In the 1990s, I was featured in an Atlantic magazine article titled “Computer Delusion,” which questioned the value—and even the potential harm—of putting personal computers in classrooms. At the time, I was helping design one of the first public high schools with a one-to-one computing model: Napa New Technology High School, which later became the New Tech Network and expanded to more than 350 public schools. The absurdity of the concern is obvious in hindsight. Personal computers did not go away. The internet did not go away. Mobile devices did not go away. AI will not go away.

Every transformative technology can be used for progress or misused for harm. History reminds us that fear-driven control often misses the mark. During the early days of the printing press, people were imprisoned or killed for using it. When it could not be stopped, books were banned instead—an irony that still echoes in modern education debates. Safeguarding and control are fundamentally different motivations. One exists to enable safe, positive use while preventing harm. The other exists to suppress change out of fear.

Our hope is that this brief helps you think differently—so you can lead your organizations to create conditions that promote progress, safeguard innovation, and build the lifelong capabilities students and educators need not for some distant future, but for the present moment.

E. Ted Fujimoto, Managing Director

Executive Summary

Artificial intelligence is no longer a future consideration for education. It is already reshaping how work is done across nearly every sector, altering which skills are valued, how decisions are made, and where human judgment remains essential. For K–12 education, this moment presents both risk and opportunity—not because of the technology itself, but because of what it reveals about how learning systems are currently designed.

This discussion brief does not attempt to answer the increasingly narrow question of *how* AI should be used in classrooms or organizations. The pace of change makes prescriptive guidance quickly obsolete. Instead, **the brief offers a reframing: AI should be understood as a stress test for education systems, exposing long-standing misalignments between what schools reward and what the world now requires.**

Across sectors, AI is accelerating the value and high priority need of skills such as inquiry, sensemaking, ethical reasoning, collaboration, and adaptability—while compressing the value of routine execution and task-based proficiency. These shifts challenge traditional approaches to instruction, assessment, accountability, and workforce preparation. Attempts to respond primarily through restriction, surveillance, or compliance mechanisms may reduce short-term discomfort, but they often undermine long-term readiness and equity.

The brief argues for a distinction between safeguarding and stifling. Safeguarding acknowledges that powerful tools will be used and focuses on building judgment, capability, and responsibility over time. Stifling, by contrast, relies on control and fear, often pushing use underground and widening gaps between those with access and those without.

A central theme is that education cannot navigate this transition alone. Much of the most relevant knowledge about AI now resides inside organizations, embedded in proprietary workflows and real-world decision contexts. Meaningful business–education collaboration—focused on translating thinking rather than transferring tools—is essential if schools are to prepare students for current realities, not outdated assumptions.

The brief also addresses the distinct roles of educators, school and district leaders, policymakers, boards, and business leaders. Each has a different form of leverage, and misalignment among these groups creates fragmentation rather than progress. Effective leadership in this moment requires creating conditions for learning, not chasing certainty or final answers.

Ultimately, this is not a technology brief. It is a leadership brief. AI accelerates a question education has long deferred: are we primarily designing systems for compliance and coverage, or for thinking and judgment? The answer to that question—and the actions that follow—will shape whether education evolves with the world or increasingly finds itself out of step with it.

Move Past the Hype to Clarity and Action

There is hype around AI—just as there was around personal computers, the internet, and smartphones. Creators and investors amplify potential because it serves their interests. In each case, the technology ultimately became embedded in how organizations operate and how society functions. At the same time, much of the early hype overstated near-term impact or underestimated how long meaningful adoption would take. We talked about self-driving cars for decades before they became usable on real roads.

To move forward productively, we need to separate hype from what is already happening.

AI has been used for years in call center operations to detect emotional tone in customer interactions, flagging conversations that require intervention or improvement. For more than a decade, AI has supported physicians in reviewing CT scans for certain cancers, significantly improving detection rates when used alongside top clinicians—often outperforming the average physician alone.

In software development, recent advances have radically compressed timelines. Tasks such as code review and refactoring that once took weeks now take days, and work that previously required months can often be completed in weeks. With appropriate safeguards—such as limiting data sources to specific jurisdictions—legal firms are using AI to draft contracts that are then reviewed by senior counsel, reducing reliance on large teams of junior associates.

Consulting firms have relied on large numbers of entry- and mid-level analysts whose primary role was research and synthesis. Today, with careful direction and citation requirements, AI can conduct deep research, evaluate source quality, and synthesize findings in minutes and hours rather than weeks.

Our internal experience mirrors this shift. A software platform project originally scoped for five to nine months was reduced to three to four months last year using early AI tools. With more recent releases, we scrapped three months of work, started over, replicated our progress in under a week, and moved to a delivery timeline of only a few months. We have saved hundreds of thousands of dollars in legal costs over the past two years by using AI to draft agreements aligned with our intent, followed by targeted legal review that required only minor revisions. We also use AI as a learning partner—not to give answers, but to help us develop understanding. By asking it to challenge our thinking and surface blind spots, we've uncovered insights that would have required far more iterations on our own.

In the University of Houston-Clear Lake educational leadership doctoral program, students in a strategic planning course for the first time were required to use AI to develop their plans. Their final reflections revealed a key insight: **AI is not a shortcut. It deepened the work, but only when students knew how to ask strong questions. Critical thinking, inquiry, and communication were essential for AI to be useful.**

How you "talk" to AI matters. It is like addressing a crowd of a million people. If you are vague, the response will be generic. Asking the same question of a stand-up comic, a leadership coach, or a specialized attorney produces entirely different outcomes. Precision, context, and intent shape results.

The bottom line is this: AI is already reshaping work in ways that may not be visible from the outside. This is not a future-state issue. It is happening now. Leaders and boards increasingly see a workforce divided into three broad groups: those with deep institutional knowledge and relationships; those with strong critical thinking and adaptability who can assemble resources and navigate constant change; and those whose roles consist primarily of routine tasks that AI can already perform as well or better.

What makes this moment different from past technology cycles is that AI does not merely automate tasks. It analyzes, synthesizes, and generates insight. The value it produces depends less on the tool itself and more on how it is engaged. A prompt that asks AI to think like a CFO balancing efficiency with customer experience yields a fundamentally different result than a CFO whose mindset is to just cut costs. In an AI-enabled environment, asking the right questions of the right persona matters as much as the answers themselves.

This raises a critical question for education: what skillsets and thinking capacities do we need to build in our teams and students, where are they today, and how quickly can we get them there? Time is not on our side.

From Tools to Thinking: What AI Exposes About Our Education System

AI is not just another instructional tool or productivity platform. It is a stress test for how our education systems think about learning, work, and value.

What AI exposes—often uncomfortably—is not a technology gap, but a thinking gap. Many of the structures and practices that dominate K–12 education were designed for a world where knowledge was scarce, information moved slowly, and proficiency meant mastering procedures. In that world, education systems optimized for compliance, task completion, and predictability. Those assumptions are now colliding with a reality where information is abundant, synthesis matters more than recall, and the ability to frame good questions often outweighs the ability to produce quick answers.

This is why AI feels so destabilizing. It performs exceptionally well at many of the things schools have traditionally rewarded: summarizing information, generating written responses, following rules, completing structured tasks, and executing well-defined procedures. When a tool can do these things faster—and often better—than students or staff, it forces an uncomfortable question: were we measuring what actually matters?

AI makes visible a long-standing imbalance in education between *knowing* and *thinking*. Knowing is necessary, but it is no longer sufficient. Thinking—especially critical thinking, sensemaking, ethical judgment, and the ability to operate in ambiguity—is far harder to automate. Yet these capacities have often been treated as secondary outcomes rather than core design goals.

This exposure is not an indictment of educators. It is the predictable result of systems that were built to scale efficiency, not adaptability. Teachers are asked to cover extensive standards, prepare students for standardized assessments, and operate within tightly constrained accountability frameworks. Under those conditions, it is rational to emphasize task completion and right answers. AI simply reveals the limits of that model.

The disruption is also uneven. Students who already possess strong inquiry skills, curiosity, and confidence in navigating uncertainty tend to use AI as an amplifier of their thinking. Students who have been trained primarily to comply, follow instructions, and produce acceptable responses are more likely to use AI as a shortcut—or be displaced by it entirely. The same pattern holds for adults in organizations.

In this sense, **AI is less a revolution than a mirror. It reflects the degree to which an education system has invested in developing thinkers rather than task completers.** It exposes where learning has been reduced to outputs instead of understanding, and where success has been defined by meeting requirements rather than making meaning.

This reframing matters. If leaders approach AI primarily as a tool to manage—something to allow, ban, or monitor—they will miss the deeper issue. The more consequential question is what AI reveals about the capacities our systems are actually building, and whether those capacities align with the world students are entering now, not the one education was designed for decades ago.

Before asking how AI should be used in classrooms, districts, or schools, it may be more productive to ask a prior question: *What kind of thinking does our system currently reward—and what kind of thinking does the world now demand?4*

Safeguarding Without Stifling: A Leadership Responsibility

Once AI is understood as a mirror rather than merely a tool, the leadership challenge becomes clearer—and more difficult. The issue is not whether AI should exist in schools. It already does. The real question is how leaders safeguard its use without suppressing learning, innovation, or professional judgment.

In moments of rapid change, organizations often default to control. Restrictions feel decisive. Bans feel safe. But history shows that control-based responses rarely produce the outcomes leaders intend—especially when the underlying forces are external, powerful, and accelerating. They tend to drive use underground, widen inequities, and signal mistrust rather than responsibility.

Safeguarding is different. Safeguarding assumes the tool will be used and focuses instead on *who*, *how*, and *under what conditions*. It is not about eliminating risk; it is about managing it intelligently while building capacity over time.

In K–12 systems, this distinction is especially important because leaders are responsible not only for protecting students, but also for preparing them. **Shielding students from powerful tools may reduce short-term discomfort, but it often increases long-term vulnerability.** Students who never learn how to use AI responsibly will still encounter it—just without guidance, context, or ethical grounding.

Effective safeguarding starts with role clarity. Not everyone needs the same level of access, responsibility, or autonomy. Students, teachers, administrators, and external partners each engage AI differently, and policies should reflect those differences. Expecting a single, universal rule to govern all use cases oversimplifies the problem and invites unintended consequences.

Safeguarding also requires transparency. Hidden rules, vague prohibitions, and inconsistent enforcement erode trust. Clear expectations—what is permitted, what is not, and why—allow educators and students to make informed decisions. Transparency shifts the conversation from “What will I get in trouble for?” to “What is responsible use in this context?”

Another core element is capability-building. Safeguards are ineffective if the people expected to follow them lack the skills to do so. This applies as much to adults as to students. Many districts invest heavily in student-facing policies while underinvesting in educator learning. When teachers are unsure how AI works or how it can support deeper learning, safeguards become blunt instruments rather than enabling structures.

Importantly, safeguarding should be dynamic. AI capabilities are evolving too quickly for static policies to remain relevant. Leaders should expect to revisit assumptions, update guidelines, and learn from real use rather than attempting to finalize rules upfront. Treating safeguards as living frameworks—subject to review and refinement—models the very adaptability schools are trying to build.

Finally, safeguarding must be anchored in values, not fear. Decisions driven primarily by liability avoidance or public pressure tend to overcorrect. Decisions grounded in a clear vision for learning and long-term readiness are more likely to strike the right balance. This does not mean ignoring risks such as data privacy, bias, or misuse. It means addressing them directly while keeping the larger purpose in view.

The leadership task, then, is not to answer every question about AI today. It is to create conditions where powerful tools are used with judgment, where mistakes become learning opportunities rather than crises, and where responsibility grows alongside capability. That is far harder than issuing a ban—but it is also far more aligned with the mission of education.

What Skills Actually Matter Now

If AI is already embedded in how work gets done across nearly every sector, then the most urgent question for education is no longer whether students should encounter it, but whether schools are developing the capacities that make people effective in an AI-enabled world.

This requires separating enduring skills from legacy ones.

Many of the skills traditionally emphasized in K–12 education were rational responses to past constraints. Accuracy, speed, memorization, procedural compliance, and individual task completion mattered when information was scarce and work was narrowly defined. AI now performs many of these functions effortlessly. Continuing to center them as primary indicators of readiness creates a growing mismatch between what schools reward and what the world demands.

The skills gaining value are different—and more demanding.

First is the ability to frame good questions. AI is only as effective as the inquiry that guides it. Students and educators who can articulate purpose, context, constraints, and desired outcomes consistently outperform those who cannot, regardless of technical sophistication. Question quality has become a multiplier of thinking, not a soft skill.

Second is sensemaking—the ability to synthesize information, recognize patterns, evaluate credibility, and draw meaningful conclusions. AI can surface options and insights, but it cannot determine relevance or consequence on its own. Human judgment remains central, especially in complex, values-laden contexts such as education.

Third is ethical reasoning. AI introduces real questions about authorship, bias, equity, privacy, and responsibility. Navigating these issues requires more than rule-following. It requires principled judgment, perspective-taking, and an understanding of tradeoffs. These capacities are developed through discussion, reflection, and exposure to ambiguity—not through compliance-based instruction.

Fourth is communication and collaboration. AI increasingly acts as a partner within teams, not a standalone tool. The ability to articulate thinking, test ideas with others, and integrate diverse perspectives becomes more important as technical barriers fall. Clear communication is no longer just about expression; it is about directing and refining intelligence—both human and artificial.

Fifth is adaptability. AI tools evolve on a monthly, sometimes weekly, basis. The most valuable individuals are not those who master a specific platform, but those who can learn, unlearn, and relearn continuously. This includes comfort with experimentation, iteration, and revision—skills that are often unintentionally discouraged in high-stakes, perfection-oriented environments.

None of these skills are new. What is new is their relative importance. AI compresses the value of routine execution and expands the value of judgment, curiosity, and integration. This shift challenges long-standing assumptions about rigor, assessment, and readiness.

It also challenges how time is used in schools. If students spend the majority of their learning time practicing tasks that AI can already perform, they are being prepared for a world that no longer exists. This does not mean foundational knowledge is irrelevant. It means that knowledge must increasingly serve thinking, not replace it.

The uncomfortable implication is that developing these capacities takes intentional design. They do not emerge automatically from exposure to technology. They require learning environments that reward inquiry over compliance, depth over coverage, and growth over performance. Without those conditions, AI will widen gaps rather than close them.

The question for K–12 leaders is not whether these skills are important—most would agree they are. The harder question is whether current structures, incentives, and accountability systems are actually aligned to develop them at scale.

Implications for Teaching, Assessment, and Accountability

Once we accept that AI fundamentally shifts which skills matter most, the implications for teaching, assessment, and accountability become unavoidable. This is where the conversation often becomes uncomfortable—not because answers are unclear, but because long-standing structures begin to feel misaligned.

Many instructional practices in K–12 education are built around the assumption that individual work products reliably reflect individual understanding. Homework, essays, problem sets, and projects have traditionally served as proxies for learning. AI complicates this assumption. When a student can generate a well-written essay or a complete solution with minimal effort, the artifact itself becomes a weaker signal of thinking.

This does not mean learning has disappeared. It means our measurement systems are no longer aligned with what we claim to value.

Teaching in an AI-enabled environment shifts from delivering content to designing experiences that make thinking visible. **Educators increasingly need to focus on how students frame problems, justify decisions, revise ideas, and reflect on their reasoning.** These processes are harder to standardize and more time-intensive to assess, but they are also far more indicative of real understanding.

Assessment, therefore, becomes less about detecting misuse and more about redefining evidence. If AI is available outside of school—and it is—then the question is not how to eliminate its influence, but how to design tasks where its use either adds value or becomes irrelevant. Oral defenses, iterative drafts, real-world application, collaborative problem-solving, and reflective explanations provide richer insight into learning than static submissions alone.

Accountability systems add another layer of complexity. Many educators operate under policies that reward coverage, pacing, and test performance, often at the expense of deeper learning. When these systems remain unchanged, teachers face an impossible bind: develop higher-order thinking that may not be immediately measurable, or prioritize compliance with existing metrics. AI intensifies this tension by exposing how easily surface-level performance can be produced without corresponding understanding.

None of this suggests abandoning standards, rigor, or accountability. It suggests reexamining what rigor actually means. In a world where AI can generate answers, rigor increasingly lies in reasoning, judgment, and the ability to apply knowledge in unfamiliar contexts. Accountability, in turn, must evolve to recognize growth in these areas rather than relying exclusively on easily automated outputs.

For leaders, the key challenge is resisting the urge for quick fixes. Plagiarism detectors, surveillance tools, and blanket rules may provide short-term reassurance, but they do little to address the underlying misalignment. More importantly, they risk signaling to students and educators that compliance matters more than learning.

The deeper question is not how to protect existing systems from AI, but whether those systems still serve their intended purpose. AI forces education to confront a long-deferred conversation: are we primarily assessing what students know, or how they think? The answer to that question has implications far beyond technology—it goes to the heart of what schools are designed to do.

Preparing Educators, Not Just Students

Much of the AI conversation in K–12 education focuses on students: what they should be allowed to use, how misuse should be addressed, and how learning should be protected. Far less attention is given to the readiness of the adults responsible for designing, guiding, and safeguarding learning. This imbalance is one of the greatest risks in the current moment.

Education systems do not change through policy alone. They change through the daily decisions of teachers, principals, and district leaders. If those adults are uncertain, underprepared, or unsupported in navigating AI, no amount of student-facing guidance will produce coherent outcomes.

Many educators are being asked to respond to AI while simultaneously managing existing pressures: standards coverage, assessment demands, staffing shortages, and accountability constraints. In that environment, AI can easily feel like an additional burden rather than an enabling tool. When professional learning focuses primarily on rules, warnings, or technical demonstrations, it reinforces anxiety instead of building confidence.

Preparing educators begins with reframing their role. Teachers are not expected to become AI experts, just as they are not expected to be software engineers or data scientists. Their value lies in judgment, pedagogy, and relational expertise. AI should be positioned as a thinking partner that supports planning, differentiation, feedback, and reflection—not as a replacement for professional expertise.

Adult learning around AI must therefore prioritize understanding over mastery. Educators need space to explore how AI works at a conceptual level, what it does well, where it fails, and how bias and limitations show up in practice. Without this foundation, safeguards become abstract and inconsistent, and use becomes either overly cautious or overly dependent.

Equally important is psychological safety. Educators need permission to experiment, make mistakes, and learn in public without fear of punishment or reputational damage. In systems where missteps are penalized, innovation goes underground or disappears altogether. Leaders who want responsible AI use must model learning themselves—asking questions, acknowledging uncertainty, and revising assumptions as understanding evolves.

This preparation also extends to leadership teams and boards. **Strategic decisions about AI cannot be delegated entirely to IT departments or compliance offices.** Leaders need enough fluency to ask the right questions: How does this align with our learning priorities? What risks are we managing, and which are we accepting? Where are we building capability versus enforcing compliance?

Finally, **preparing educators requires time.** AI adoption is often discussed as if it should be immediate, but meaningful integration—especially in education—requires reflection, dialogue, and iteration. Systems that rush implementation without investing in adult learning tend to experience fragmentation, resistance, or superficial use.

If students are expected to develop higher-order thinking, adaptability, and ethical judgment, educators must be supported in developing those same capacities in relation to AI. The credibility of any student-facing expectation depends on adult readiness behind it.

Preparing educators is not a secondary consideration. It is the enabling condition for everything else this brief has discussed.

The Critical Role of Business–Education Collaboration

One of the least discussed—but most consequential—realities of the AI era is where knowledge actually lives. Increasingly, the most advanced understanding of AI systems, workflows, and applications is not in the public domain. It is proprietary, embedded inside companies, protected by intellectual property, and evolving through real-world use rather than academic publication.

Education institutions, particularly in K–12, are structurally disadvantaged in this environment. They are not designed to operate at the leading edge of rapidly commercialized technology. They lack the capital to continuously invest in cutting-edge tools, the incentives to iterate at market speed, and the access to real-time operational learning that occurs inside firms deploying AI at scale.

This is not a failure of schools. It is a mismatch of purpose.

Expecting education systems to independently define "future-ready" skills without meaningful connection to the organizations shaping that future is unrealistic. The pace of change alone makes this impossible. By the time a capability is formalized into standards, curriculum, or assessments, industry practice has often moved on.

This creates a widening gap. Students are asked to prepare for a world that schools cannot fully see, while companies develop capabilities behind closed doors that rarely translate cleanly into educational settings. Without intentional collaboration, education risks preparing students for yesterday's version of work—while believing it is being forward-looking.

Business–education collaboration is not about outsourcing education's mission or allowing corporate interests to dictate curriculum. It is about shared visibility and mutual learning. Companies understand emerging tools, workflows, and expectations. Educators understand learning, development, ethics, and long-term human capacity. Neither can replace the other, but each is incomplete without the other.

In an AI-driven economy, this collaboration must go deeper than guest speakers, advisory boards, or episodic internships. It requires structured mechanisms for translating proprietary knowledge into transferable capabilities—without violating intellectual property or turning schools into training centers for specific firms.

Examples might include co-developed problem spaces rather than prescribed solutions; exposure to real decision-making contexts without revealing sensitive data; and joint exploration of how judgment, inquiry, and ethical reasoning operate inside complex organizations. The goal is not to teach specific tools, but to surface how thinking actually happens when stakes are real and ambiguity is unavoidable.

This collaboration also addresses a hard truth: education systems cannot, and should not, attempt to compete with industry on tooling or capital investment. Schools will always lag in access to the newest platforms and infrastructure. Their advantage lies elsewhere—in cultivating adaptable thinkers who can enter organizations and learn quickly. But doing so requires clarity about what adaptability actually looks like in practice, and that clarity increasingly sits inside companies.

For policymakers and system leaders, this raises important questions. How do we create partnership structures that are reciprocal rather than extractive? How do we ensure access to these learning opportunities does not reinforce existing inequities? How do we protect educational independence while acknowledging economic reality?

Ignoring these questions does not preserve education's integrity. It isolates it.

In a world where critical knowledge is increasingly proprietary and experiential, the future of education depends less on owning information and more on building bridges to where learning is actively occurring. AI accelerates this need. Without intentional collaboration between education and business, the gap between schooling and real-world readiness will continue to widen—regardless of how much technology is added to classrooms.

What Helps—and What Hurts: Guidance for Policymakers and Boards

Policymakers and boards play a decisive role in how education systems respond to AI—not through day-to-day implementation, but through the signals they send, the constraints they impose, and the conditions they create. In periods of rapid change, governance decisions can either enable thoughtful adaptation or unintentionally freeze systems in place.

What helps begins with clarity of purpose. Boards and policymakers are most effective when they articulate *why* AI matters in relation to learning, workforce readiness and long-term public value. When the conversation is framed narrowly around risk avoidance or compliance, systems optimize for defensibility rather than growth. When it is framed around capability-building and readiness, leaders have permission to learn, iterate, and improve.

Helpful governance also distinguishes between principles and prescriptions. Clear principles—such as transparency, age-appropriate use, data protection, and ethical responsibility—provide direction without overconstraining practice. Highly prescriptive rules, especially those tied to specific tools or use cases, tend to become obsolete quickly and force educators into workarounds rather than thoughtful implementation.

Time and flexibility matter. Policymakers often underestimate how long it takes organizations to build real capacity. Short timelines, pilot fatigue, or mandates without corresponding support create compliance behaviors rather than learning behaviors. Boards that explicitly allow for phased experimentation, reflection, and recalibration create space for responsible adoption rather than rushed deployment.

Investment signals are another powerful lever. Funding that supports professional learning, cross-sector collaboration, and redesign of assessment and instructional models is far more impactful than funding tied exclusively to purchasing tools or software licenses. Technology without accompanying capacity-building almost always underdelivers.

Equally important is what hurts.

Blanket bans on AI use—especially those enacted quickly in response to public concern—often produce the opposite of their intended effect. They do not eliminate use; they push it out of sight, reduce transparency, and widen inequities between students who have informal access and those who do not. Bans also signal mistrust of educators' professional judgment, which undermines morale and innovation.

Surveillance-heavy approaches are similarly problematic. Tools designed to detect misuse may provide short-term reassurance, but they reinforce a culture of compliance and suspicion. Over time, they discourage experimentation and honest conversation, both of which are essential for learning how to use powerful tools responsibly.

Another common pitfall is treating AI as a technical issue rather than a strategic one. Delegating decisions entirely to IT, legal, or compliance teams without broader educational leadership involvement narrows the lens. AI affects pedagogy, assessment, workforce preparation, and organizational design—not just infrastructure and risk.

Finally, **inconsistency hurts.** Conflicting guidance across agencies, frequent policy reversals, or unclear enforcement standards create confusion and erode trust**.** Educators are left guessing which expectations matter and which will change next year. Stability in direction, even amid uncertainty, is more valuable than premature certainty.

For boards and policymakers, the goal is not to "get AI right" immediately. It is to govern in a way that allows systems to learn faster than the technology changes. That requires humility, patience, and a willingness to resist pressure for simple answers to complex problems.

The most helpful posture is not control, but stewardship—setting direction, protecting core values, and enabling educators to build the capacities that students and society now require.

Practical Next Steps by Role: Where to Focus Now

AI feels overwhelming precisely because it cuts across roles, responsibilities, and long-held assumptions. The most productive next steps are not about doing everything at once, but about concentrating attention where it has the highest leverage. What follows is not a checklist. It is a set of role-specific priorities that clarify where focus matters most right now.

District Leadership: Start with Conditions, Not Tools

The core responsibility at the district level is coherence—clarity of purpose, shared expectations, and system-level conditions that enable learning and adaptation. AI should be framed as a strategic issue tied to learning, workforce readiness, and equity, not as a technology rollout or a compliance exercise.

Disproportionate investment should go into adult learning. Not one-off trainings, but sustained, collective learning for leaders and educators about how AI works conceptually, where it adds value, where it introduces risk, and how it intersects with pedagogy, assessment, and professional judgment. Without this shared foundation, policies fragment and guidance dissolves at the school level.

District leaders should also examine where existing accountability structures quietly undermine stated goals. If inquiry, judgment, and adaptability matter, do pacing guides, assessments, and evaluation systems actually allow for them? AI does not create these misalignments—it exposes them.

Finally, **districts should pursue meaningful cross-sector partnerships**—not for products or sponsorships, but for insight into how thinking, learning, and decision-making operate in real organizations using AI. The goal is not job training, but the translation of real-world cognitive demands into transferable capabilities for students.

School Leadership: Build Trust and Instructional Coherence

At the building level, the priority is culture. Fear must be reduced and trust deliberately built, especially among teachers. Leaders should clearly signal that AI is not a “gotcha” issue and that responsible experimentation is expected, supported, and learned from—not punished.

Instruction and assessment deserve more attention than enforcement. The most important questions are instructional: How is student thinking made visible? Where does AI deepen understanding, and where does it bypass it? These questions are best explored collaboratively through real student work and classroom scenarios, not abstract rules.

Time is a non-negotiable condition. Teachers need protected time to plan differently, redesign tasks, talk with one another, and reflect. Without it, even well-designed initiatives collapse under daily operational pressure.

School leaders also set the tone by modeling learning themselves—using AI transparently in leadership work, asking questions publicly, and sharing what they are learning along the way. In periods of uncertainty, how leaders learn matters as much as what they decide.

Classroom Practice: Design for Thinking, Not Policing

The most productive shift for teachers is away from policing AI use and toward designing for thinking. The starting point is clarity: What should students understand, reason through, or be able to explain? Learning experiences should be designed so AI either supports that goal or becomes largely irrelevant to it.

Greater attention should be paid to the quality of student questions, reasoning, and explanations—not just final products. How students frame problems, revise ideas, and justify conclusions reveals far more about learning than polished outputs. AI can assist with drafts, examples, and feedback, but it cannot replace a student’s explanation of why something makes sense to them.

Transparency with students matters. Open conversations about when and why AI is appropriate builds responsibility and trust. Students are already navigating these tools outside of school. Ignoring that reality does not protect learning—it leaves students without guidance.

Teachers should also allow space for not having everything figured out. Teaching in an AI-enabled world is genuinely new work. Curiosity, reflection, and professional judgment matter more than technical mastery.

A Shared Thread

Across all roles, the most important shift is the same: move from control to capacity. AI rewards systems that learn faster than they regulate, and it exposes systems that confuse restriction with responsibility.

The question is no longer whether AI belongs in education. The real question is whether education systems are willing to evolve how learning, leadership, and accountability work—so that powerful tools are met with equally strong human judgment.

The Role of Business Leaders in Supporting Education

Business leaders who genuinely want to help education prepare students for an AI-enabled world must begin with realism about where knowledge now lives. Much of the most important learning about AI—how it is actually used, where it fails, and how judgment is applied—exists inside companies. It is proprietary, experiential, and shaped by real stakes. Education systems cannot independently access or replicate this environment, nor should they be expected to.

This reality creates a responsibility, not an opportunity for control.

The role of business is not to push tools into classrooms or to dictate what schools should teach. It is to help translate real-world thinking into learning contexts that educators can adapt responsibly. The most valuable contribution business can make is not technology or funding in isolation, but visibility into how complex decisions are made when AI is part of the process.

Effective support focuses on exposing how work actually happens. How do teams frame problems? How do they test assumptions? How do they weigh speed, risk, ethics, and customer impact? Where does AI accelerate insight, and where does human judgment remain essential? These moments sit at the core of readiness, yet they are rarely visible outside organizational walls.

Time matters as much as capital. One-off donations, pilots, or guest lectures may generate goodwill, but they rarely build capacity. Sustained relationships—with districts, schools, and educator teams—allow trust to develop and learning to compound on both sides. Education gains insight into evolving realities; business gains a deeper understanding of how people learn, adapt, and grow.

Respecting education's independence is essential. Supporting education does not mean turning schools into workforce training pipelines or aligning curriculum to the immediate needs of individual companies. It means contributing to the development of transferable capabilities—judgment, inquiry, collaboration, and ethical reasoning—that allow people to succeed across contexts, not just within a single organization.

Fair access is not optional. If meaningful business–education collaboration is confined to well-resourced systems, AI will deepen existing disparities. Leaders who care about long-term talent pipelines and social stability must take responsibility for who has access—and who does not.

Why This Matters to Business

Engaging with education in this way is not altruism alone; it is enlightened self-interest.

The gap between how schools prepare people and how work is actually done is already costly. Organizations invest heavily in onboarding, retraining, and remediation—not because people lack intelligence, but because they have not been asked to develop judgment, inquiry, and adaptability at scale. AI accelerates this cost. Companies that depend on task execution alone will struggle. Those that require thinkers, integrators, and ethical decision-makers cannot afford to wait for education systems to adjust on their own.

There is also a growing risk dimension. AI amplifies both capability and error. Poorly prepared talent increases operational, reputational, and ethical risk. Supporting education in building stronger thinking capacity is therefore a form of long-term risk management for the economy as a whole.

Many business leaders benefit from education systems they did not design but now depend on. The choice is whether to remain passive consumers of talent or become active stewards of the conditions that produce it. For organizations intended to endure, this is not abstract. The quality of thinking in the next generation will directly shape customers, workforces, communities, and institutional legitimacy.

Finally, **there is a generational responsibility. AI is not a passing cycle; it is a structural shift.** Business leaders are among the few actors with direct visibility into what is changing fastest. With that visibility comes an obligation to share insight responsibly—not to control education, but to help it evolve with clarity and foresight.

In an AI-driven future, education does not need business leaders to provide answers. It needs them to share context, surface reality, and engage as long-term partners. That is in education's interest—and it is equally in the interest of business.

A Different Starting Point

This discussion brief is not an argument for embracing AI uncritically, nor is it a warning to resist it. It is an invitation to start in a different place.

Too many conversations about AI in education begin with the question, *How do we control it?* That question is understandable—but it is also limiting. It narrows attention to tools rather than capacities, compliance rather than learning, and fear rather than responsibility. It assumes that stability is still the goal, when adaptability has already become the requirement.

A more productive starting point is this: *What kind of thinking does our system need to cultivate now—and what conditions must exist for that thinking to develop?*

Answering that question shifts the work. It moves leaders away from chasing policies that will quickly become outdated and toward designing environments where judgment, inquiry, and ethical reasoning can grow alongside powerful tools. It reframes safeguarding as stewardship, collaboration as necessity, and uncertainty as something to be navigated rather than eliminated.

AI does not diminish the role of educators, leaders, or institutions. It raises the bar. It demands clearer purpose, stronger coherence, and deeper investment in human capacity. It exposes misalignments that have existed for years and accelerates the consequences of ignoring them.

No single group can respond to this moment alone. Educators need support and trust. District and school leaders need flexibility and time. Policymakers and boards must govern for learning rather than certainty. Business leaders must share insight responsibly rather than operate behind closed doors. Each plays a distinct role, but the work only holds together if it is shared.

The choices made now—often incrementally and under pressure—will shape whether AI widens gaps or expands opportunity, whether schools remain relevant or drift further from the realities students will face, and whether innovation is guided by values or driven by fear.

This moment does not call for perfect answers. It calls for better questions, clearer intent, and the courage to design systems that can learn as fast as the world is changing.

That is not a future challenge. It is the work in front of us now.

> ***On Crafting This Brief*** *- AI was used in the development of this discussion brief to support copyediting, improve flow, and iterate on structure and ideas. It did not replace authorship or judgment, but functioned as a thinking partner—helping test clarity, surface alternatives, and refine language. Final decisions about framing, emphasis, and conclusions remained human and intentional.*

At EF International Advisors, LLC, we translate strategy into measurable results by closing the gap between vision and execution.

EF INTERNATIONAL ADVISORS

For more than 30 years, EFIA has partnered with executive teams and boards to navigate pivotal moments, manage risk, and convert uncertainty into opportunity. Operating with the agility of a boutique firm and the depth of seasoned expertise, we work across finance, real estate, education, technology, media, and insurance to help organizations move from planning to sustained impact.

In the education sector, the execution gap is especially pronounced. Ambitious goals, policy shifts, and reform initiatives frequently outpace an organization's capacity to implement them coherently and consistently. In response, EFIA formally launched Built to Deliver over the past year as both a published framework and a series of Design Studio–based engagements for education leadership teams. The work is designed to help districts and schools align strategy, culture, systems, and daily practice—ensuring that adopted plans translate into operational reality rather than remaining aspirational.

Built to Deliver is grounded in EFIA's broader body of work, including the Agile Action Strategy Process (AASP), Design Studio methodology, Strategy Dojo, and mission-aligned business intelligence tools. Together, these elements support systems-level decision-making, disciplined execution, and adaptive leadership—equipping K–12 organizations with the structure and resilience required to perform in complex, rapidly changing environments.

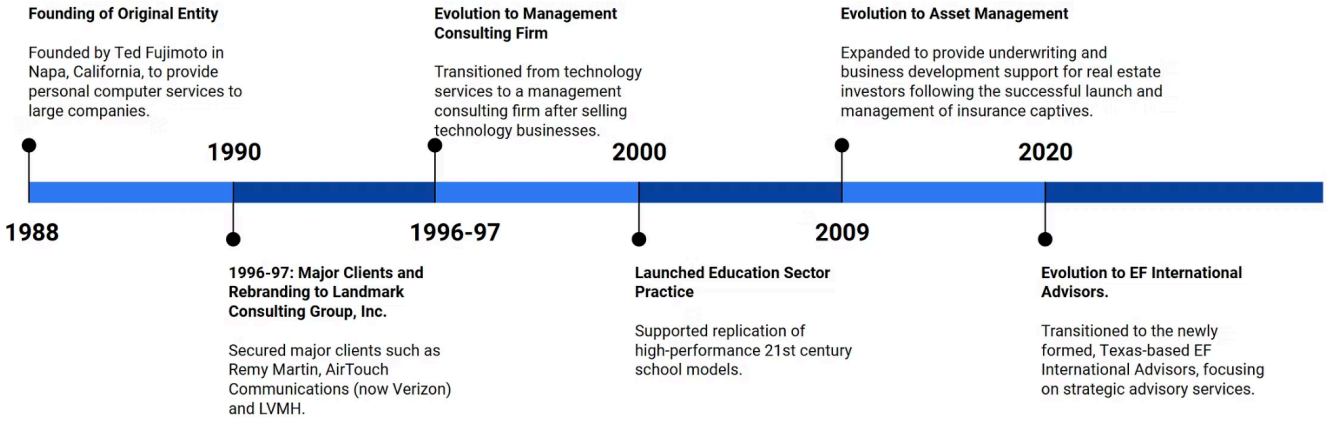

Contact: E. Ted Fujimoto, Managing Director
Email: tedf@efinternationaladvisors.com
Website: www.efinternationaladvisors.com

Ted Fujimoto - Managing Director, EF International Advisors, LLC

E. Ted Fujimoto is a nationally recognized strategist in education system design, innovation scaling, and execution discipline. For over three decades, he has supported leadership teams across public education, finance, real estate, media, technology, and insurance—advising organizations such as LVMH, IBM, Verizon, Sony, the Bill & Melinda Gates Foundation, and the Walton Family Foundation.

He co-founded the original organization and flagship school that became the New Tech Network, helped shape Big Picture Learning's strategy, and has advised school, district, and state partners through his Design Studio. His support has helped turn bold ideas into sustainable delivery—shaping student-centered experiences and leadership capacity. His broader work has supported the deployment of over $1.6 billion in investment capital and over $500 million in education-focused philanthropy.

At 18, Ted founded his first company and later co-founded Napa New Technology High School (1996), a national model for project-based learning. His design approach is grounded in his late-1990s experience when a global consulting firm—initially a client—invited him to join as an equity partner. The firm was renowned for leading large-scale system redesigns that shifted organizational culture from compliance to purpose-driven performance.

He and his work have been featured by the World Bank's IFC, WISE, Education Dive, Black Enterprise, the Partnership for 21st Century Learning, and a range of other national and international education and innovation platforms. He is the co-author of the book Built to Deliver and the originator of the Agile Action Strategy Process (AASP) framework and system on which it is based.

Dr. Antonio Corrales - Director, Doctoral Program in Educational Leadership at the University of Houston–Clear Lake

Dr. Antonio Corrales is a nationally respected leader in educational transformation, systems leadership, and executive coaching. With over 25 years of experience, he has served as a teacher, campus administrator, central office leader, and professor—supporting schools and organizations nationwide to strengthen leadership, instruction, and culture. For the past decade, he has served as Director of the Doctoral Program in Educational Leadership at the University of Houston–Clear Lake, helping education leaders navigate complex challenges through strategic planning and evidence-based practices. (***Affiliation listed for identification only. Views expressed are the authors' and do not represent the University.)***

Before arriving in the U.S., he ran for governor in Venezuela at age 27, which led him to leave his country for political reasons. Today, his leadership focuses on school redesign, strategic innovation, and executive coaching for superintendents and central office leaders.

Dr. Corrales has authored several books on school turnaround, leadership, and personal growth. His recent work focuses on executive function development in PK–12 students, with an emphasis on character, resilience, and readiness for college, careers, and life. He is also the co-author, with Ted Fujimoto, of the book Built to Deliver. He actively connects with educators worldwide through multiple influential outlets, books, and research articles, sharing insights on critical topics for parents, students, and education professionals.