

# The Private Capital Opportunity in AI-Enabled Climate and Sustainability Sectors

May 2026

By Daniel Oehling, Greg Fischer, Andrey Berdichevskiy, Tariq Nanji, Vinay Shandal, Benjamin Sheridan, Reshma Jayaprabha, Franziska Zimmermann, Alvin Low, Jasmine Teo, Kelly Chen, and Gladys Tan

# Contents

- 03** The Mutually Reinforcing Benefits of AI
- 07** AI Is Creating Value Across a Broader Landscape Than Most Climate Investors Recognize
  - Deep Dive 1: Climate Risk Modeling (including Insurance)
  - Deep Dive 2: Industrial Equipment and Systems Efficiency
  - Deep Dive 3: Grid, Storage, and System Flexibility Management
  - Deep Dive 4: Inclusive Education
  - Deep Dive 5: Materials Discovery
- 23** What Could Slow Adoption
- 24** Looking Ahead
- 25** About the Authors



# The Mutually Reinforcing Benefits of AI

Sustainability is, at its core, about resource efficiency: achieving the same or better outcomes with less energy, fewer materials, less waste, and—where natural systems have been degraded—restoring them to productive capacity. AI’s fundamental capability is optimizing how resources are used. Where these two domains meet, financial returns and sustainability outcomes are positively correlated. Not by coincidence. By design.

We estimate that deploying current AI capabilities across climate and sustainability sectors could generate approximately \$600 billion globally in annual value. That figure is conservative; the broader AI opportunity in each of these sectors is substantially larger. Industrial systems, the knowledge economy, and health care are all undergoing AI-driven transformation that extends well beyond climate and sustainability applications.

This report focuses specifically on the intersection where AI-driven efficiency gains produce both financial returns and measurable environmental or social outcomes. That is why we discuss industrial resource efficiency but not broader manufacturing automation; insurance risk pricing but not general fintech; and climate-linked health applications but not the wider health care AI opportunity.<sup>1</sup>

To define this landscape, we developed an AI-enabled Climate and Sustainability Investment Opportunity Map covering more than 40 subsectors across three domains: climate and energy transition, natural capital and resource management, and social systems and livelihoods. Each was assessed for subsector attractiveness—reflecting current market size and expected growth—and positive AI impact to identify the most promising near-term opportunities. This report examines five priority subsectors in depth:

- climate risk modeling (including insurance)
- industrial equipment and systems efficiency
- grid, storage, and system flexibility management
- inclusive education
- materials discovery

While these are not the only areas likely to attract investors, they offer the most immediately evident combined financial and sustainability impact from AI.

1. For readers interested in AI’s broader impact on health care, financial services, and industrial systems, the following BCG reports are an excellent starting point: “[GenAI Is a Productivity Game-Changer in Health Care Operations](#)”; “[Convergence: Human + AI for the Next Era of Finance](#)”; and “[How Physical AI Is Reshaping Robotics Today—and What Comes Next](#).”

The opportunities span the full capital spectrum. Venture capital suits AI-native companies. Growth equity fits scaling platforms. Buyout strategies apply where AI integration can improve margins in established businesses. And infrastructure capital enters where AI is embedded in physical assets such as grids, storage systems, and data centers. The case studies in this report are not based on optimistic projections. They rely on measured outcomes from the deployment of AI across a wide variety of applications, from cement kilns to insurance underwriting to disaster early warning technology. Each case demonstrates that commercial performance and sustainability gains can be mutually reinforcing.

The complexity of embedding AI in these systems is considerable. It requires domain-specific data, deep integration with legacy infrastructure, operational expertise, and significant change management. For investors, that complexity creates durable competitive advantage. For businesses, it means early deployment is compounding, because each cycle generates data and integration depth that later entrants cannot shortcut.

## Sustainability and Financial Returns Are Linked by Design

Why does this report treat seemingly disparate subsectors like climate hazard intelligence, industrial process optimization, grid orchestration, and inclusive education as parts of the same investment landscape? Because AI connects them through a common mechanism.

Climate and sustainability challenges, whether in energy systems, manufacturing, agriculture, or education, share a structural feature: scarce resources being used inefficiently. Energy is wasted in industrial thermal processes. Renewable generation is lost to grid congestion and curtailment. Climate-driven insurance losses arise from risks that existing models cannot price accurately. Students are falling behind because instruction cannot adapt to their pace. Natural capital—forests, wetlands, soils—is degraded to the point where it provides a fraction of the ecosystem services it could. In each case, the gap between current performance and what is achievable is large, persistent, and costly, in both financial and environmental or social terms.

AI narrows these gaps. And because the resource being optimized is the same whether the outcome is measured in dollars or emissions, the financial and sustainability results move together. When a cement plant's thermal energy consumption falls 3% through AI-driven process control,

the manufacturer's fuel costs drop and its Scope 1 emissions decline. When an AI platform orchestrates a fleet of residential batteries, the asset owner earns more from energy markets while the grid displaces fossil peaker plants. When an adaptive learning platform delivers 23 months of learning progress in 17 months at \$20 to \$25 per student per year, student retention improves and the education system reaches students it was previously failing. In each case, one action produces two positive outcomes.

We do not claim that every AI application delivers sustainability benefits. Many do not. We are making a specific observation about a specific class of applications—those that improve how energy, materials, capital, and human capacity are used in systems where inefficiency carries both a financial cost and an environmental or social cost. In these typically large systems, the link between financial performance and sustainability outcomes is structural. This relationship is the organizing principle of our report. Each subsector warrants inclusion by demonstrating a clear chain from AI capability to commercial value to sustainability outcome. Subsectors in which that chain cannot be demonstrated are not included.


We do need to take a clear-eyed view of AI's own resource demands. Compute-intensive models require significant electricity, and data centers place demands on land, water, and local energy infrastructure. Communities in several regions have raised legitimate concerns about rising utilities bills and emissions. The relevant question, however, is whether the systems that could be optimized by AI would consume fewer resources in aggregate than they currently do. Across the sectors in this report, the evidence suggests this is the case. But this is an empirical question that should be evaluated sector by sector and revisited over time.

AI is expanding what counts as climate and sustainability investing. Sectors not historically viewed through a climate or sustainability lens—such as industrial process control, catastrophe risk pricing, and grid congestion management—are becoming central to the landscape as AI makes their environmental and social performance both measurable and improvable. The \$600 billion in annual global value represents an estimate of the efficiency gains, cost reductions, and new revenue enabled if current AI capabilities were deployed at scale across 40 subsectors analyzed in this report. **(See Exhibit 1.)** It is a directional indicator of the opportunity's magnitude, not a forecast of realized returns. Moreover, the pace of innovation is accelerating and can unlock previously unanticipated opportunities.

## EXHIBIT 1

# AI Can Unlock \$600 Billion in Value Annually Across the Climate and Sustainability Ecosystem by 2028

Annual value from applying AI to climate and sustainability (\$US B, 2028)<sup>1</sup>



Source: BCG analysis.

<sup>1</sup>Calculated based on 2028 global market size. Value uplift potential reflects total commercial value that could be unlocked by users, based on applicable AI use cases.

The estimate of annual global value aggregates three types of economic benefit: cost reductions (such as lower energy spend or reduced maintenance costs); revenue enablement (such as new insurance products made possible by AI-driven risk scoring); and improved asset utilization (such as deferred capital expenditure from extended equipment life). (See “Methodology.”)

Several forces are converging to make these opportunities investable now. Modeling capabilities have advanced rapidly. The cost of compute and sensors continues to fall. The data infrastructure needed to train and deploy sustainability-specific AI models is maturing. In selected markets, disclosure requirements are reinforcing demand for greater levels of sustainability, though the underlying drivers are economic, not compliance driven. Rising energy costs, tightening resource constraints, and the operational reality that climate hazards are already reshaping business performance across sectors are creating non-discretionary demand for the solutions outlined in this report.

# Methodology



The \$600 billion bottom-up estimate of annual global value is expressed in US dollars and was derived from an assessment of more than 40 climate and sustainability subsectors. For each subsector, our team:

- identified the AI applications most relevant to its core challenges,
- estimated the size of the addressable market using 2028 projected global market sizes based on industry sources and capital market data (including S&P Capital IQ), and
- assessed the potential efficiency gain or revenue improvement based on published case studies, expert interviews, and AI solution provider-reported outcomes.

Uplift ranges were expressed as potential improvements in EBITDA margin, operating cost reduction, or revenue enablement, depending on which mechanism dominates in each subsector.

The value figure for each subsector comprises three components: cost reductions accruing to businesses deploying AI, new revenue enabled by AI-created products or services, and improved asset utilization including deferred capital expenditure. The 2028 horizon reflects the period over which current-generation AI capabilities could plausibly reach meaningful, though not universal, adoption across these sectors. The figures do not assume full deployment or full success; they represent the value that would be created if commercially proven AI applications were adopted at rates consistent with current trajectory and observed scale-up patterns.

The estimate assumes broad deployment of current AI capabilities. It is not a forecast of adoption rates or realized returns. It is best interpreted as a directional indicator of the opportunity's scale, useful for comparing relative attractiveness across sectors, not for projecting specific financial outcomes. The figure represents the full economic benefit from AI deployment, the majority of which accrues to businesses deploying AI rather than to AI solution providers, who typically capture an estimated 15% to 25% of the value created.

The \$600 billion figure is conservative by construction. It reflects only the climate-and-sustainability-specific slice of each sector's AI opportunity—the value that arises where AI-driven improvements in resource efficiency, in the productive capacity of degraded natural systems, or in the unit economics of delivering essential services produce both financial returns and measurable environmental or social outcomes.

The five priority subsectors examined in this report account for approximately \$423 billion of the total:

- climate risk modeling (\$75 billion)
- industrial equipment and systems efficiency (\$300 billion)
- grid, storage, and system flexibility management (\$32 billion)
- inclusive education (\$13 billion)
- materials discovery (\$3 billion)

The remaining value is distributed across additional subsectors including renewable power generation, water management, precision agriculture, and carbon management.

Each subsector was assessed on two dimensions. Sector attractiveness reflects the size of the current market and its expected growth trajectory over three to five years. AI impact is a composite of the current investment activity (the volume and stage of private capital already deployed into AI solutions in the subsector) and the potential AI uplift, reflecting how well-suited current AI capabilities are to the subsector's core problems and the magnitude of value they could create. Subsectors scoring high on both dimensions were classified as priority areas. Those scoring high on one were classified as attractive. Those scoring low on both were classified as opportunistic.



# AI Is Creating Value Across a Broader Landscape Than Most Climate Investors Recognize

To identify where AI creates the most investable climate and sustainability value, we need a structured view of the full landscape. Existing classification frameworks, developed primarily for policymakers or technologists, do not map cleanly to the questions investors and businesses need answered: *Where is the value, how large is it, which subsectors are ready for deployment, and what investment strategies fit which opportunities?*

The subsectors are organized into three domains: climate and energy transition, natural capital and resource management, and social systems and livelihoods. (See **Exhibit 2.**) Each subsector represents a distinct market with identifiable companies, competitive dynamics, and investment activity—from grid flexibility management to precision agriculture to climate risk modeling.

## Prioritizing Subsector Readiness for AI

Not all climate and sustainability subsectors are equally suited for AI deployment or investment. Some combine large markets, strong structural growth, and high AI impact. Others are nascent, policy-dependent, or lack commercial traction. Investors and businesses need a way to distinguish between them.

We assess each subsector on two dimensions. Sector attractiveness reflects the size of the current market and its expected growth over three to five years. AI impact is a composite of current investment activity and the potential magnitude of AI-driven value creation: how well-suited current AI capabilities are to the subsector's core problems and how large the resulting improvement could be.



## EXHIBIT 3

# Each Subsector Can Be Sorted into Three Types of Opportunities

	Opportunistic	Attractive	Priority	
Sector attractiveness	High	<ul style="list-style-type: none"> <li>1.1 Renewable power generation</li> <li>7.4 Water and wastewater treatment and recycling</li> <li>3.5 Low-carbon fuels</li> <li>2.2 Building energy management systems</li> <li>10.4 Microfinancing<sup>1</sup></li> <li>11.2 Green financing</li> </ul>	<ul style="list-style-type: none"> <li>1.3 Grid, storage, and system flexibility management</li> <li>10.2 Inclusive education</li> <li>2.4 Industrial equipment and systems efficiency</li> <li>11.3 Climate risk modeling (including insurance)</li> <li>11.4 Materials discovery</li> <li>5.1 Waste collection, sorting, and logistics</li> </ul>	
	Medium	<ul style="list-style-type: none"> <li>3.3 Electric rail and infrastructure</li> <li>7.3 Water quality monitoring</li> <li>4.1 Carbon capture, utilization, and storage</li> </ul>	<ul style="list-style-type: none"> <li>1.4 District heating and cooling</li> <li>7.1 Water sourcing and harvesting</li> <li>6.3 Alternative and sustainable proteins (including dairy)</li> <li>6.1 Low-carbon fertilizers and soil input</li> <li>3.2 Road EV charging infrastructure</li> </ul>	<ul style="list-style-type: none"> <li>10.1 Disaster preparedness, response, and recovery</li> <li>3.6 Sustainable logistics</li> <li>9.3 Climate-responsive treatment and care</li> <li>6.2 Precision agriculture and digital tools</li> <li>11.1 Sustainability data, reporting, and compliance</li> <li>4.3 Nature-based carbon removal</li> </ul>
	Low	<ul style="list-style-type: none"> <li>3.4 Marine fuel switch and systems</li> <li>4.2 Methane abatement and non-CO<sub>2</sub> gas mitigation</li> </ul>	<ul style="list-style-type: none"> <li>8.2 Land and forestry conservation</li> <li>2.3 Industrial heat and thermal efficiency</li> <li>8.1 Marine and wetland conservation</li> <li>9.1 Climate-driven disease surveillance and detection</li> <li>10.3 Inclusive workforce participation and economic mobility<sup>1</sup></li> </ul>	<ul style="list-style-type: none"> <li>7.2 Water efficiency and management</li> <li>9.2 Climate-linked disease prevention</li> </ul>
	Low	Medium	High	

AI impact

✓ Selected deep-dive

Source: BCG analysis.

<sup>1</sup>Subsectors focused on low-income (LIC) and low- and middle-income countries (LMIC), where unmet demand remains large and populations are more vulnerable to just transition and climate-related risks. In these markets, AI-enabled models can be transformative in improving access and unit economics, whereas in high-income markets the impact is more incremental.

## Who Captures the Value

The value figures in this report represent the full economic benefit from deploying AI across climate and sustainability sectors, not the share captured by any single participant. In practice, the companies using these solutions—whether utilities, insurers, or industrial manufacturers—capture the majority of value through lower costs, improved asset performance, and new revenue. AI solution providers typically capture a smaller share through subscription fees,

licensing, and outcome-based contracts, estimated at roughly 15% to 25% based on observed pricing models across the sectors in this report. Both cost reduction and revenue growth contribute to the financial case. This dynamic is also giving rise to a growing ecosystem of specialized AI solution providers—companies purpose-built to serve climate and sustainability sectors—representing an investable category in its own right.

## Matching Opportunities with Investment Strategies

AI-driven climate and sustainability opportunities span the full capital spectrum, but different strategies target different parts of the landscape.

*Venture capital* typically funds AI-native companies, purpose-built around specific climate and sustainability challenges with strong technology but unproven unit economics at scale. The thesis is classic: to back domain-specific teams with proprietary data and deep technical capabilities before incumbents replicate the offering. Climate risk modeling startups, adaptive learning platforms, and AI-driven grid planning tools all fit this profile.

*Growth equity* targets companies that have demonstrated commercial traction and repeatable deployments across multiple customers or geographies. These are businesses with proven returns on investment, expanding customer bases, and clear pathways to profitability that need capital to accelerate market expansion. Growth investors may also find opportunities in established businesses that are repositioning their products around climate and sustainability as a growth vector.

*Buyout strategies* fit established companies in climate-relevant industries where AI-driven transformation can improve margins or reposition product portfolios. In hazard intelligence, for example, several platforms have reached commercial scale with proven cash flows, but their AI capabilities can be deepened to strengthen pricing power and customer retention. In insurance, scaled platforms offer profiles where AI-enabled underwriting can materially improve combined ratios. In industrial systems, automation incumbents with OEM distribution represent opportunities where AI integration differentiates against commodity competitors.

*Infrastructure capital* enters where AI is integral to how a physical asset performs. Infrastructure investors do not need to buy an AI company. They need to own assets such as battery storage fleets, grid infrastructure, water networks, and waste concessions where AI improves the economics. In these cases, the commercial benefit from AI accrues predominantly to the asset owner, not the software provider.

The deep dives that follow apply this logic to several priority subsectors, examining what is driving demand, where AI creates value, and what deployed proof points look like today. **(See Exhibit 4.)**

### EXHIBIT 4

## Each of the Five Investable Subsectors Offers Varying Degrees of Potential Value, AI Applicability, and Impact on Sustainability

Subsector	Estimated annual value	AI readiness	Policy dependence	Sustainability impact
11.3 Climate risk modeling (including insurance)	~\$75B	High	Moderate	Expanded insurance coverage; improved early warning
2.4 Industrial systems and equipment efficiency	~\$300B	Medium	Low	Reduced emissions (by ~0.6 Gt); lower waste; safer workplaces
1.3 Grid, storage, and system flexibility management	~\$32B	High	High	Higher renewable integration; reduced curtailment
10.2 Inclusive education	~\$13B	Medium	Low	Reduced learning gaps; greater inclusion of underserved students
11.4 Materials discovery	~\$3B	Medium	Low	Accelerated clean energy transition; lower R&D waste

Source: BCG analysis.

# Deep Dive 1: Climate Risk Modeling (including Insurance)

Climate hazards are intensifying, and the economic costs are no longer abstract. Between 2000 and 2025, the frequency of billion-dollar natural disasters more than doubled. Insured losses from natural hazards exceeded \$100 billion in both 2023 and 2024. Climate-related disruptions cost airlines up to 3% of revenues, utilities face escalating outage restoration costs, and logistics companies absorb billions in delay-related losses annually.

The demand for accurate climate risk modeling is no longer driven by voluntary sustainability commitments. The driving factors today are the operational need to reduce losses, accurately price risk, and comply with tightening disclosure requirements. Traditional hazard models, built on historical data and coarse geographic assumptions, are increasingly unable to capture evolving climate dynamics at the resolution that businesses and insurers require. AI addresses this by integrating diverse data sources—including satellite imagery, sensor networks, and atmospheric models—to generate asset-level, forward-looking risk assessments that historical approaches cannot produce.

## AI Application Areas

AI-enabled climate risk modeling falls into three application archetypes, each serving different end users and offering distinct investment characteristics.

### Operational Hazard Intelligence

This technology translates weather and hazard data into site-specific, time-sensitive triggers, enabling logistics companies to reroute shipments, utilities to pre-position repair crews, and public agencies to issue targeted warnings. Public weather forecasts typically operate at resolutions of 10 to 50 kilometers and cannot be translated into the hyperlocal, asset-level intelligence that businesses need to act. AI-enabled platforms close this gap.

The financial and sustainability benefits of these tools are direct: reduced weather-related losses for businesses and extended early warning protection for communities. DTN's Storm Risk Analytics, for example, predicted the location and number of utility customers affected by Hurricane Ian within 3% accuracy a day before landfall, enabling utilities to pre-stage crews and materially reduce outage duration. Companies in this space, including Tomorrow.io, Spire, Climavision, and DTN, have achieved commercial scale. Proprietary data assets such as satellite constellations and deep integration into end-user operating systems have created significant barriers to replication.

The potential value from AI-enabled hazard intelligence totals approximately \$30 billion, distributed across logistics (\$11 billion), utilities (\$8 billion), the public sector (\$7 billion), and aviation (\$4 billion). The logic is consistent across each industry: a large revenue or spending base, a meaningful share exposed to climate-related disruption, and an AI-addressable portion where better forecasting can reduce those losses. In logistics, for example, climate-related delays affect a measurable share of the \$3.5 trillion global road and rail freight market, but AI-enabled route-level predictions can reduce excess fuel costs and losses due to rerouting. In the public sector, the more than \$100 billion in annual disaster response spending could be reduced by up to 30% through early warning systems that most countries still lack. Utilities and aviation follow similar patterns: large weather-exposed revenue bases where AI forecasting can reduce disruption costs.

### Asset-Level Risk Analytics

This AI application combines computer vision on aerial imagery with climatological data to generate property-level risk scores, transforming how insurers underwrite climate-exposed assets. Traditional catastrophe models rely on zone-level assumptions and historical loss data that increasingly understate emerging climate shifts. AI-enabled platforms analyze the structural characteristics of individual properties, simulate damage under varying hazard intensities, and produce hazard-specific scores at a resolution legacy models cannot match.

The impact on underwriting is substantial. One US insurer achieved a 4.4% reduction in its combined ratio—the core measure of insurance profitability, comparing claims and expenses paid out against premiums collected—in the first year of deploying ZestyAI's platform, a result that reflects both better risk selection and fewer surprise losses. Beyond pricing accuracy, these capabilities can expand insurance availability to regions and hazard types previously considered uninsurable, supporting the underwriting of an estimated 15 to 20 million additional policies while maintaining pricing that reflects asset-specific risk rather than blanket geographic exclusions.

The potential value from AI-enabled asset-level risk analytics totals approximately \$30 billion. This value is concentrated in property and casualty insurance, some \$800 billion of whose premium base is exposed to climate risk.

Providers include AI-native platforms such as ZestyAI and large-scale property intelligence platforms like Cotality and Nearmap that serve the broader property ecosystem. Neara, which currently markets its asset-level analytics to non-insurance sectors, has capabilities readily transferable to insurance underwriting.

## Portfolio Climate Stress Testing

The third AI archetype enables financial institutions and real estate investors to project chronic climate shifts across portfolios over decades, stress-test for concentration risk, and automate reporting aligned with disclosure frameworks. This segment is earlier-stage, with an estimated \$15 billion in potential value, and is served predominantly by pure-play providers including AlphaGeo, Jupiter Intelligence, Mitiga Solutions, and Climate X.

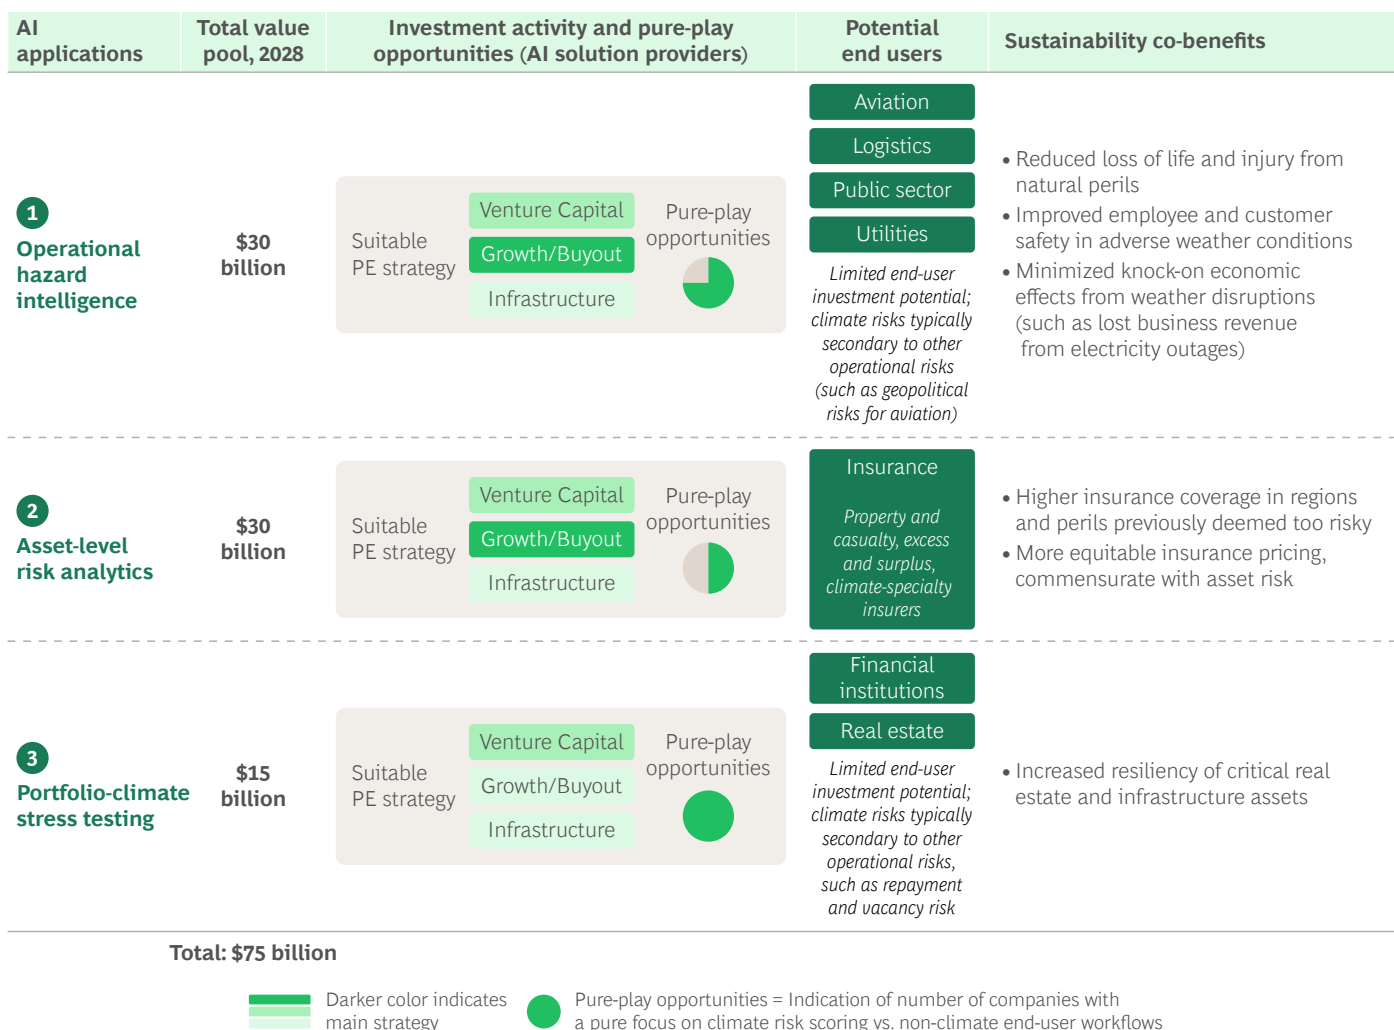
For example, Stoneweg, a real asset investment group with €11 billion in assets under management, partnered with AlphaGeo to embed physical climate resilience into its investment operations. AlphaGeo's AI-driven inferential downscaling translates global climate models into 100-meter resolution hazard projections, asset-specific risk scores, and cashflow-level metrics, enabling more targeted resilience spending and fewer valuation surprises at exit.

## Investable Opportunities

Taken together, AI-enabled climate risk models have the potential to generate approximately \$75 billion in annual value by 2028. (See Exhibit 5.) The commercial benefits, including reduced disruption costs for businesses, improved pricing accuracy for insurers, and better-informed capital allocation for financial institutions, will increase as adoption expands. The sustainability benefits are equally significant. AI-based early-warning systems could extend protection from climate disasters to millions of people globally, enable precise risk assessments for regions previously deemed uninsurable, and inform climate-adaptive interventions across critical real estate and infrastructure.

### EXHIBIT 5

## Operational Hazard Intelligence and Asset Risk Analytics Offer the Largest Value Pools and the Most Growth or Buyout Opportunities



Source: BCG analysis.

AI providers offer the broadest range of investment entry points. Operational hazard intelligence platforms are largely at commercial scale, making growth equity the natural vehicle for geographic expansion and sensing capability development. Stonecourt Capital and HarbourVest Partners, for example, co-led a \$175 million equity round in Tomorrow.io in 2026 to fund satellite network expansion into data-sparse regions.

Asset-level risk analytics features a mix of AI-native platforms suited to venture capital and broader property intelligence platforms suited to buyouts. Thoma Bravo's AUD \$1.0 billion take-private of Nearmap in 2022, followed by Nearmap's acquisition of Betterview's AI risk-scoring platform, illustrates the potential for buyouts. Portfolio stress testing remains earliest-stage, favoring venture capital.

Insurance offers the most compelling entry point among end users. Two segments stand out. Excess and surplus insurers and climate-specialty carriers are established businesses well-suited to buyouts; Ryan Specialty Holdings' \$525 million acquisition of Velocity in 2025 is a notable example. These companies have limited in-house AI capabilities, and the deployment of AI-driven risk models represents a tangible post-acquisition value creation opportunity. Parametric insurers, by contrast, develop AI in-house and represent a faster-growing field. Descartes Underwriting is the category leader, with Battery Ventures investing at a premium to its \$120 million Series B in 2022 to fund global expansion. Earlier-stage players such as Arbol, Kettle, and Understory are suited to venture capital.

The broader ecosystem offers selective opportunities tied to AI-identified risk priorities. In insurance, AI-driven analytics can expand the addressable market for brokers connecting high-risk clients to carriers. In infrastructure and real estate, AI stress testing can identify specific resilience gaps, creating demand for resilience manufacturers and installers that directly benefit from the resulting intervention spend. In supply chain and logistics, AI can unlock additional opportunities to minimize losses and optimize costs and insurance underwriting.

# Deep Dive 2: Industrial Equipment and Systems Efficiency

Industrial activities involving the extraction, processing, and manufacture of physical goods account for roughly a quarter of global carbon emissions. The production of primary materials such as iron and steel, cement, and industrial chemicals, along with process manufacturing in plastics and food and beverage, contributes approximately 90% of those industrial emissions. These are continuous, energy-intensive processes requiring heavy machinery and complex thermochemical operations. They consume immense amounts of energy, generate significant materials waste, and expose workers to dangerous conditions.

These characteristics make industrial systems especially well suited to AI-driven optimization. They also make the dual-value link particularly direct. Every unit of energy waste eliminated is simultaneously a cost saved and an emission reduced. Every equipment failure predicted and prevented is both downtime avoided and material not scrapped. The financial and sustainability cases are not parallel arguments. They are the same argument.

The scale of the climate and sustainability case across industrial systems matches the scale of the financial case. Deploying current AI capabilities across these industries could reduce Scope 1 and 2 emissions by an estimated 0.6 gigatons annually, an amount roughly equivalent to the total emissions of Germany. In steel alone, AI-driven safety systems could prevent approximately 20,000 workplace injuries per year. These are not incidental co-benefits. They are the direct, measurable consequences of the same interventions that generate the \$300 billion in commercial value.

To be clear, not all industrial AI applications belong in a climate and sustainability report. Scheduling optimization that improves throughput without reducing resource intensity is an operational improvement, not a sustainability intervention. The applications in this section are included because they demonstrably reduce energy consumption, emissions, material waste, or workforce harm—environmental and social outcomes that are measurable and that scale with deployment.

Four pressures are converging to accelerate demand. Energy, labor, and commodity cost volatility is driving businesses to seek solutions that strengthen cost resilience. Expanding carbon pricing mechanisms and disclosure mandates are pushing industrial companies toward lower-carbon production. Customers are demanding materials that combine higher performance with lower carbon intensity. And tightening labor and safety regulations are intensifying scrutiny of workforce hazard exposure.

## AI Application Areas

AI-enabled applications in industrial systems cluster into six archetypes spanning equipment, production processes, and workforce safety. Three account for the majority of value and illustrate the range of investment opportunities.

### Predictive Maintenance

This archetype offers the largest potential value, at \$79 billion. Industrial equipment is typically serviced reactively or on fixed schedules, leading to unexpected failures, production disruptions, and high emergency repair costs. Machine learning models trained on high-frequency sensor data (vibration, temperature, and acoustics) can detect anomalies before failures occur and predict time-to-failure with increasing precision. The sustainability co-benefit is direct: Extended equipment life means fewer replacement parts manufactured, less material waste, and lower resource consumption.

One of the world's leading cement manufacturers, for example, deployed Nanoprecise's MachineDoctor sensors on roller press bearing points to continuously monitor vibration, temperature, and overall condition. The system's AI analytics detected developing gear faults and lubrication contamination before the machine failed—findings subsequently confirmed through laboratory oil analysis—thus avoiding a catastrophic gearbox failure, extending the machine's service life, and preventing over \$500,000 in costs. Other providers include Tractian, Augury, AssetWatch, and Infinite Uptime, a mix of venture-backed pure plays and OEM-backed platforms.

### Adaptive Process Optimization

This technology offers \$59 billion in potential value by continuously adjusting operating conditions (temperature, pressure, oxygen flow, and fuel rate) using real-time optimization models and reinforcement learning. The solution addresses the limited ability of conventional process control to manage multiple parameters simultaneously, which leads to unstable production, excess fuel consumption, and avoidable yield losses.

Japan's Tokuyama Cement, for example, deployed ABB's Ability Expert Optimizer to optimize its kiln and plant operations in real time. The system uses multivariable control algorithms to continuously adjust temperature, airflow, and fuel rate, reducing process fluctuations and improving combustion efficiency. The result has been a 3% reduction in thermal energy consumption and a 70% reduction in manual operator interventions. Other providers include Imubit, Aleph, and Conundrum, all deep-tech companies with highly domain-specific models and proprietary training data.

## Energy Management Systems

These systems offer \$56 billion in potential value by continuously monitoring and optimizing plant-wide energy flows with equipment-level visibility. They detect and eliminate energy waste from inefficient subsystems, compressed-air leaks, missed waste-heat recovery, and other sources that fragmented, static monitoring cannot see. In one example, a leading steel manufacturer deploying C3.ai's energy forecasting and optimization platform achieved \$14 million in annual energy cost savings, a 1.8% increase in onsite power generation use, and a 40MW reduction in monthly utility charges.

Guidewheel, backed by Decarbonization Partners and Breakthrough Energy Ventures, illustrates a different deployment model. Its clip-on power sensors attach to any machine in minutes, feeding a cloud-based AI platform that monitors uptime, throughput, and energy consumption across entire fleets. Using the platform, a large US packaging manufacturer reduced the gap between machine failure and maintenance response from two hours to 15 minutes. Its cross-fleet learning capabilities also create a data advantage that compounds with scale by calibrating each new machine against tens of thousands already in the system.

The three remaining archetypes—quality control and assurance (valued at \$47 billion), advanced planning and scheduling (\$47 billion), and workplace safety (\$12 billion)—complete the picture.

## Adding Up the Impact

Altogether, AI applications across industrial equipment and systems could generate approximately \$300 billion in annual value, with food and beverage manufacturing accounting for roughly two thirds of the total. (See Exhibit 6.) Primary materials industries—iron and steel (\$33 billion) and cement (\$16 billion)—capture value primarily through predictive maintenance and process optimization.

The sustainability impact is equally concrete: an estimated 0.6 gigaton reduction in industrial Scope 1 and 2 emissions through energy optimization, lower material waste through real-time quality monitoring—and a potential 15% reduction in recordable workplace injuries in steel alone, or approximately 20,000 fewer injuries per year.

Adoption patterns for these AI-powered industrial systems vary by region. In the US, reshoring pressures and an aging manufacturing workforce are accelerating demand for AI-driven productivity solutions. In parts of China, factory automation is more advanced, but AI adoption by external solution providers remains constrained by a preference for in-house development. In parts of India and Southeast Asia, low labor costs reduce the urgency for automation but create opportunities for AI platforms that can demonstrate rapid payback on energy and waste reduction, benefits that are less sensitive to wage levels.

## Investable Opportunities

AI providers span the full investment spectrum, with entry points shaped by application maturity. Predictive maintenance is the most commercially proven archetype. Augury closed a \$75 million round in 2025 backed by Lightrock, Insight Partners, and others, maintaining its unicorn valuation, and subsequently acquired Seebo to move beyond machine health into end-to-end production optimization. Growth equity and buyout strategies are well suited here. Adaptive process optimization and energy management are served by a mix of scaled horizontal platforms such as C3.ai and ABB Ability as well as by earlier-stage pure plays better suited to venture capital and growth equity. A range of emerging and established players across stages offer quality control, scheduling, and safety solutions.

Industrial end users represent compelling opportunities in their own right. Food and beverage manufacturers offer a particularly strong entry point given the sector's large global market (more than \$8 trillion), high AI uplift potential from complex production processes, diverse product mixes, and relatively attractive margins. Within food and beverage, protein and dairy processing and processed food manufacturing are the highest-impact segments.

The broader ecosystem presents additional opportunities. Industrial sensor and edge hardware vendors benefit directly as AI adoption drives the attachment of hardware to existing equipment fleets. System integrators that configure and deploy AI across complex plant environments are capturing deployment revenue as demand for integration increases. Sixth Street Growth's backing of Seeq, an advanced analytics and AI platform for industrial process data, in a \$50 million Series D investment illustrates the range of capital being deployed across the ecosystem.

**EXHIBIT 6**

# AI-Enabled Predictive Maintenance for Machinery Has the Largest Potential Value Uplift for Industrial Manufacturers

AI applications	Total value pool, 2028	Investment activity and pure-play opportunities (AI solution providers)	Potential end users	Sustainability co-benefits
<b>1</b> <b>Asset performance and predictive maintenance</b>	\$79 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Iron and steel Cement Industrial chemicals F&B manufacturing Plastics processing	<ul style="list-style-type: none"> <li>Improved efficiency of raw materials and input use; lower environmental footprint with lower resource use</li> </ul>
<b>2</b> <b>Adaptive process optimization</b>	\$59 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Iron and steel Cement Industrial chemicals F&B manufacturing Plastics processing	<ul style="list-style-type: none"> <li>Lower direct and indirect emissions</li> <li>Reduced material waste and production losses; improved efficiency of input use</li> </ul>
<b>3</b> <b>Quality control and assurance</b>	\$47 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Iron and steel Cement Industrial chemicals F&B manufacturing Plastics processing	<ul style="list-style-type: none"> <li>Avoided emissions from material waste and rework</li> <li>Reduced material waste and production losses; improved efficiency of input use</li> </ul>
<b>4</b> <b>Energy management solution</b>	\$56 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Iron and steel Cement Industrial chemicals F&B manufacturing Plastics processing	<ul style="list-style-type: none"> <li>Lower direct and indirect emissions</li> <li>Improved efficiency of input use</li> </ul>
<b>5</b> <b>Advanced planning and scheduling</b>	\$47 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Iron and steel Industrial chemicals F&B manufacturing Plastics processing	<ul style="list-style-type: none"> <li>Lower environmental footprint with reduced input and other industrial resources</li> </ul>
<b>6</b> <b>Workplace safety</b>	\$12 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Iron and steel Cement Industrial chemicals F&B manufacturing Plastics processing	<ul style="list-style-type: none"> <li>Reduced workplace injuries and fatalities; improved safety in high-risk industrial environments</li> </ul>

**Total: \$300 billion**

*Infrastructure plays span all use cases through AI-embedded smart industrial factories*

Darker color indicates main strategy
  Pure-play opportunities = Indication of number of companies with a pure focus on industrial equipment and systems efficiency

Source: BCG analysis.  
 Note: F&B = food and beverage.

# Deep Dive 3: Grid, Storage, and System Flexibility Management

The global electricity system is under more strain than at any point in its history. Demand is rising from electrification, urbanization, and data center expansion. Renewable penetration is accelerating, introducing greater volatility. Distributed energy resources, including rooftop solar, behind-the-meter batteries, and electric vehicles, are turning grids into bidirectional systems that must coordinate millions of small, variable assets. An estimated 25% to 30% of US transmission and distribution assets are approaching the end of useful life. And extreme weather, geopolitical disruption, and cyber threats are raising the bar for reliability.

The traditional grid investment thesis was physical: Build more transmission, deploy more storage, connect more renewables. AI shifts the value equation toward orchestration: how existing assets are monitored, predicted, dispatched, and coordinated. In deployed cases, the results are already substantial. A battery storage fleet with AI-driven dispatch optimization has earned 25% to 30% more revenue from the same installed hardware. A transmission network with AI-enabled dynamic line ratings can unlock 10% to 30% additional capacity without new construction. These are not efficiency gains at the margin. They change the economics of the underlying assets.

The need is global, but uneven. In developed markets, aging infrastructure and rising renewables penetration drive demand for AI-enabled grid orchestration. In emerging markets, where grids are expanding rapidly but often lack the monitoring infrastructure of mature systems, AI can leapfrog legacy approaches, deploying predictive analytics on newer assets from the start rather than retrofitting decades of installed equipment.

The dual-value link is direct. AI-optimized dispatch displaces fossil-fueled peaker plants while increasing revenue for storage operators. Predictive monitoring reduces wildfire ignition risk from power lines while extending asset life and deferring capital expenditures. Better demand forecasting reduces curtailment of renewable generation while lowering balancing costs. In each case, the grid becomes low-carbon and more reliable through the same intervention that makes it more profitable.

## AI Application Areas

AI applications span three interconnected parts of the electricity system: transmission and distribution networks, utility-scale battery storage, and virtual power plants. Five application areas emerge, of which two account for the largest near-term value.

### Predictive Maintenance and Asset Monitoring

This represents the largest opportunity, at \$22 billion. Today's transmission and distribution operators rely on reactive maintenance and static line ratings (fixed thermal capacity limits based on conservative assumptions). AI-enabled applications replace these with real-time condition monitoring, failure prediction, and dynamic line ratings that unlock latent capacity based on actual operating conditions.

Power lines may ignite fires when interacting with surrounding vegetation. AI-driven monitoring helps detect and prevent these ignition risks by identifying where line conditions and vegetation proximity create the greatest danger.

Neara's physics-based digital twin platform illustrates the impact. Deployed by a large US utility facing escalating wildfire exposure and equipment-driven outages, the platform consolidated GIS data, LiDAR profiles, field inspection logs, and operational datasets into a unified model of the network. Engineering simulations assessed actual operating stresses for every pole and span. The result: 80% of transmission poles reclassified based on actual condition rather than incomplete records, a projected 8% reduction in average outage duration, and more than \$5 million in recoverable annual revenue from unauthorized attachments and clearance violations. Other providers of this technology include AiDash, Overstory, Ampacimon, and LineVision.

### AI-Enabled Virtual Power Plant Orchestration

This area represents \$2 billion in near-term value, the smallest application by dollar amount today, but structurally important for the long-term energy transition. Virtual power plants (VPPs) aggregate and coordinate thousands of distributed energy resources—including residential batteries, rooftop solar, EV chargers, and smart thermostats—to function as a single flexible power asset. Traditional rule-based dispatch systems cannot manage this coordination at scale, but AI-driven platforms can. An emerging frontier extends this logic to data centers themselves: NVIDIA and Emerald AI have partnered with leading energy companies to pioneer “flexible AI factories” that use AI to shift computational workloads in response to grid conditions. This approach is effectively turning data centers from rigid power consumers into dispatchable grid assets that can absorb excess renewable generation and reduce load during peak demand.

Sunrun, which operates one of the largest fleets of residential solar-plus-battery systems in the US, chose Lunar Energy's Gridshare platform to orchestrate its distributed batteries. Gridshare predicts household load, solar production, and battery availability at the device level, aggregates thousands of units, and optimizes dispatch across energy, capacity, and grid services while balancing customer bill savings against backup reserve needs. The deployment increased energy delivered from the same fleet by approximately 60% compared to rule-based systems, providing peak demand reduction and fast response services across multiple markets. Other providers in this space include Octopus Energy, Uplight, and GridBeyond.

This subsector includes three additional application areas: grid infrastructure planning and design (valued at \$3 billion annually), dynamic grid stability and congestion management (\$3 billion), and market bidding and real-time dispatch for utility-scale battery storage (\$2 billion). Chile's national grid operator partnered with Tapestry to deploy an AI-enabled simulation platform that runs 30 times more scenarios and completes simulations 86% faster than legacy tools, an example of how AI-driven planning is accelerating grid modernization beyond North America.

## Investable Opportunities

AI applications in grid, storage, and system flexibility management could generate a total of approximately \$32 billion in annual value by 2028. (See Exhibit 7.) Vertically integrated utilities are expected to capture the largest share, at roughly \$20 billion, followed by distribution operators (\$6 billion) and transmission operators (\$4 billion). Battery energy storage system (BESS) operators and VPP platforms will capture smaller absolute amounts, but their margin improvement—up to 5 or 6 percentage points—can be proportionally larger.

AI providers are concentrated at venture and growth stages. VPP orchestration and predictive maintenance represent the more mature end. Octopus Energy Group, for example, announced in 2026 that it will acquire a majority stake in Uplight, in partnership with Schneider Electric, creating a vertically integrated demand-side platform spanning customer engagement, DER orchestration, and grid operations. Grid planning and congestion management remain predominantly venture-stage, although EQT-backed IFS's acquisition of Copperleaf Technologies for approximately \$740 million in 2024 shows that buyout-scale opportunities exist.

Among end users, BESS operators and independent power producers have been typical entry points for private capital. The interest is evident: KKR's \$1 billion acquisition of Zenobē and Brookfield's €6.1 billion acquisition of Neoen both reflect the thesis that AI-enabled asset optimization makes the underlying infrastructure economics more attractive.

The broader ecosystem offers additional opportunities in hardware vendors, where AI-driven improvements in storage economics can accelerate battery deployment, and in smart meter and IoT connectivity providers that form the data layer on which grid AI depends.

**EXHIBIT 7**

# AI-Enabled Grid Predictive Maintenance Is Expected to Generate the Largest Value Uplift for Market Players by 2028

AI applications	Total value pool, 2028	Investment activity and pure-play opportunities (AI-embedded players or AI solution providers)	Potential end users	Sustainability co-benefits
<b>1</b> <b>Grid infrastructure planning and design</b>	\$3 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Transmission system operators Distribution system operators Vertically integrated utilities	<ul style="list-style-type: none"> <li>Reduced outage risks, enhanced grid resilience and flexibility</li> <li>Higher renewable penetration</li> <li>Reduced environmental disruption through smart siting and planning</li> </ul>
<b>2</b> <b>Predictive maintenance and asset monitoring</b>	\$22 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Transmission system operators Distribution system operators Vertically integrated utilities	<ul style="list-style-type: none"> <li>Reduced outage risks, enhanced grid resilience and flexibility</li> <li>Improved safety with lower wildfire and infrastructure risk</li> <li>Improved protection of critical public services</li> </ul>
<b>3</b> <b>Dynamic grid stability and congestion management</b>	\$3 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Transmission system operators Distribution system operators Vertically integrated utilities	<ul style="list-style-type: none"> <li>Improved system balancing and enhanced grid flexibility</li> <li>Reduced renewable curtailment with integration of distributed energy resources</li> <li>Improved access to clean energy technologies and broader participation in flexibility programs</li> </ul>
<b>4</b> <b>Market bidding and real-time dispatch</b>	\$2 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	BESS operators Independent power producers	<ul style="list-style-type: none"> <li>Improved system balancing and enhanced grid flexibility</li> </ul>
<b>5</b> <b>AI-enabled VPP orchestration</b>	\$2 billion	Suitable PE strategy Venture Capital Growth/Buyout Infrastructure Pure-play opportunities	Transmission system operators Distribution system operators Vertically integrated utilities	<ul style="list-style-type: none"> <li>Improved system balancing and enhanced grid flexibility</li> <li>Integration of distributed energy resources</li> <li>Improved access to clean energy technologies and broader participation in flexibility programs</li> </ul>

**Total: \$32 billion**

Darker color indicates main strategy    
 Pure-play opportunities = Indication of number of companies with a pure focus on grid, storage, and system flexibility

Source: BCG analysis.

Note: VPP = virtual power plant; BESS = battery energy storage system.

# Deep Dive 4: Inclusive Education

Education gaps are widening globally, and the consequences extend well beyond the classroom. In low-income countries, primary school completion rates stand at just 43%, compared with 99% in high-income economies. Pupil-to-teacher ratios average 71:1 in low-income countries versus 17:1 in high-income countries. These disparities are not only social challenges; they are economic ones. Workforce readiness and long-term social mobility depend on education systems that can deliver quality learning at scale. Yet the structural constraints facing those systems—including underfunding, teacher shortages, and rigid curricula—are intensifying, not abating.

Why does education belong in a report on AI-enabled climate and sustainability? Because the dual-value test applies. AI-enabled education technology improves access and outcomes for underserved communities, a measurable sustainability gain, while generating commercial value for education providers through reduced costs, improved retention, and expanded reach. The mechanism is the same as in industrial systems or grid management: AI optimizes a scarce resource (in this case, teacher capacity and instructional time), and the financial and social returns flow from the same intervention.

## AI Application Areas

AI applications in education span four areas across the learning value chain: content generation, adaptive learning, personalized assessment and feedback, and teacher and staff copilots. Two applications carry the bulk of the near-term impact.

### Adaptive Learning Platforms

These assess each student's level and dynamically adjust content, pacing, and difficulty, reaching students whom uniform instruction fails. ConveGenius deployed its adaptive learning platform across government schools in India as part of the Personalized Adaptive Learning program. The platform's adaptive engine assesses each student's baseline and continuously adjusts content difficulty, supplemented by gamification mechanics to sustain motivation. Students achieved between 22 and 23 months of learning progress in just 17 months, more than five months of acceleration compared with peers outside the program. The initiative has scaled to more than 15,000 schools and 2 million students, at a cost of \$20 to \$25 per child per year, demonstrating that AI-enabled personalization can be delivered affordably even in resource-constrained settings. Other platform providers include Squirrel AI and QANDA.

### Teacher and Staff Copilots

This technology delivers the largest directly measurable commercial impact, at \$7.5 billion. These applications automate administrative tasks such as reporting, scheduling, communication, and first-line student queries that consume a disproportionate share of educator time and are a leading driver of teacher attrition. Texas State Technical College, experiencing rapid growth across its 11 campuses, deployed Element451's AI-enabled student engagement platform. Multi-modal AI agents handle student questions around the clock, with nearly one-third of all conversations occurring after hours, while behavior-triggered personalization automates outreach tailored to each student's program. The result: 33% enrollment growth over five years sustained without proportional increases in admissions staff, 29 weeks of advising capacity saved, and a 92% application satisfaction rate. Other providers in this space include Ellucian, GoGuardian, and Teachmint.

## Adding Up the Impact

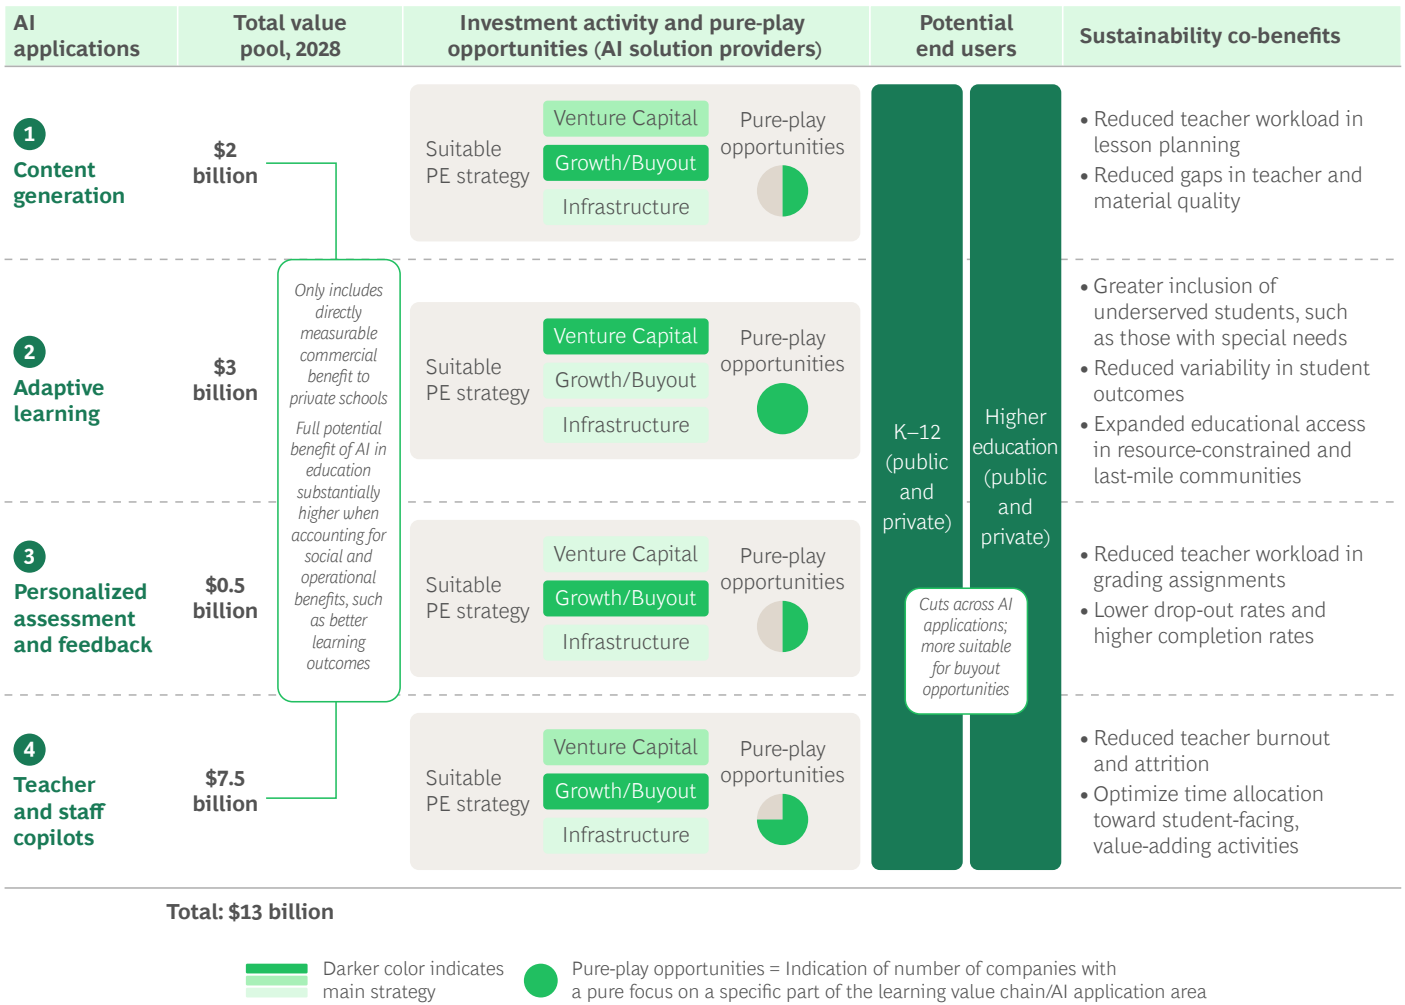
In total, AI applications in inclusive education could generate approximately \$13 billion in directly measurable commercial value for private schools by 2028. **(See Exhibit 8.)** Private higher education institutions will capture the larger share at roughly \$8 billion, reflecting higher per-student spending and greater operational complexity, with \$5 billion in value for private K–12. These figures are conservative as they capture only the directly commercial component. The full value is substantially larger when accounting for redeployable teacher capacity, improved learning outcomes, and expanded access, benefits that can be reinvested into school operations or accrue to society over the long term.

The sustainability case is equally impressive. AI-enabled adaptive learning can reduce variability in outcomes across students by tailoring content to individual baselines. It enables greater inclusion of underserved learners, including low-income, special needs, and rural students, through low-cost delivery models. And by reducing administrative burdens, AI directly combats teacher burnout, a key driver of attrition globally.

Risks remain real, however. AI capability in education is still maturing; models trained primarily on English-language data underperform in other linguistic and cultural contexts, and the evidence base for long-term learning outcomes is still developing. Access is also uneven. The populations who would benefit most are often those with the least connectivity and device infrastructure. That said, AI-based education is structurally cheaper per learner at scale than traditional approaches, and as costs continue to fall, the technology will become increasingly affordable in settings where it is needed most.

**EXHIBIT 8**

# Teacher and Staff Copilots Offer the Largest Value Pool, with Opportunities in Growth and Buyout Stages



Source: BCG analysis.

## Investable Opportunities

Private equity has long been active in education through school operator acquisitions and vendor platforms. AI is reinvigorating deal activity across three plays.

Established vendors and edtech platforms are compelling buyout targets. In 2024, Bain Capital paid \$5.6 billion for PowerSchool, a cloud-based K-12 platform with AI capabilities spanning personalized tutoring, adaptive assessments, and content generation across more than 90 countries. McGraw-Hill, acquired by Platinum Equity for \$4.5 billion in 2021, has developed internal AI content tools that recouped their initial investment in the first year. For these companies, AI enables both premium pricing and accelerated product development.

AI-native providers represent a distinct venture and growth equity opportunity. KKR acquired a majority stake in Education Perfect for \$345 million in 2021 through its Global Impact Fund, and PSG Equity invested \$175 million in Element451 in 2024. These companies' architectural flexibility—where AI is built in from the start rather than retrofitted onto legacy systems—gives them a structural advantage in product development, although they require capital to build distribution.

School operators remain the most established buyout and roll-up play. AI reinforces the thesis by improving learning outcomes, student retention, and enrollment growth. Brookfield's \$825 million investment in Spring Education Group in 2025 illustrates the model. The group has launched an online K-2 school targeting mobile student populations such as military families.

# Deep Dive 5: Materials Discovery

The slow pace and high cost of materials R&D has historically constrained the development of viable low-carbon alternatives across energy, carbon management, agriculture, and life sciences. Discovering a new battery chemistry or carbon capture material typically requires years of laboratory experimentation, most of which fails. AI is changing the economics of discovery by narrowing the search space, predicting material properties before synthesis, and reducing the number of failed experiments.

The near-term quantified value is modest: approximately \$3 billion by 2028, distributed across energy materials such as battery chemistry (\$1.1 billion), agricultural and biological materials such as microbial strain discovery (\$0.9 billion), carbon management materials such as sorbent and solvent screening (\$0.6 billion), and pharmaceutical applications targeting climate-linked diseases (\$0.4 billion). But the strategic importance is larger than the near-term figure suggests. Breakthroughs in battery chemistry, carbon capture materials, or nitrogen-fixing microbes could reshape entire industries and accelerate the climate transition in ways that current cost curves do not capture.

## AI Application Areas

Mitra Chem illustrates the potential. Its AI-enhanced acceleration platform cuts lab-to-production timelines for next-generation iron-based cathode materials for batteries by over 90%, testing tens of thousands of variations annually through machine-learning-guided synthesis and automated characterization. Citrine Informatics takes a cross-sector approach: Its AI-driven materials informatics platform enables 5 to 10 times faster product development across industries including batteries and specialty chemicals. Other companies include CuspAI in carbon capture materials, Ginkgo Bioworks in microbial strain discovery, and Kebotix and Schrödinger as cross-sector platforms.

## Investable Opportunities

The space is predominantly venture and early growth stage. For investors, materials discovery represents a longer-duration position adjacent to industrial decarbonization and renewable energy, a bet on the enabling science rather than the deployment layer. The returns, if they come, will be nonlinear.



# What Could Slow Adoption

The value described in this report is not guaranteed. Deploying AI in physical-world climate and sustainability systems is harder than deploying it in digital-only workflows, and several constraints could slow the adoption of AI or limit its impact. Countering these forces will be essential in realizing the full benefits of AI.

## Explainability and trust

When AI outputs affect underwriting decisions, grid operations, or industrial process control, users need to understand why a model reached its conclusion. Where AI cannot be validated against trusted baselines or expert judgment, adoption slows, particularly in regulated industries. The strongest implementations maintain human decision rights over high-stakes actions and document model limitations clearly.

## Data access and quality

AI performance is dependent on operational data—which can often be fragmented, siloed, or incomplete, particularly in emerging markets and legacy industrial environments. Companies that invest in building a focused, high-quality data foundation before attempting broad deployment consistently outperform those that do not.

## Model drift

Weather patterns shift. Equipment degrades. Market rules change. Industrial processes evolve. AI models trained on historical conditions will lose accuracy over time unless they are continuously monitored and retrained. Post-deployment quality assurance is not optional; it is a core operational requirement.

## Liability and accountability

When an AI recommendation contributes to a pricing error, an outage, or a compliance failure, responsibility for the incident must be clear. The most resilient deployments maintain explicit decision ownership, audit trails, and contractual risk allocation, and avoid full autonomy in high-stakes contexts.

## Policy and regulatory uncertainty

Disclosure rules, grid market design, and AI governance frameworks are evolving at different speeds across jurisdictions. Companies and investors operating across markets face an uneven regulatory landscape that can accelerate adoption in some regions and delay it in others. Building products that can adapt to multiple regulatory environments is a practical necessity, not a competitive luxury.

These constraints are real, but they are not unique to using AI in climate and sustainability; they apply to AI deployment broadly. What distinguishes this space is the tailwind of structural demand: The problems AI addresses in these sectors are growing more urgent, not less. The question is not whether adoption will happen, but how quickly the constraints can be managed.



# Looking Ahead

AI will not solve every climate and sustainability challenge. And it will not replace the physical infrastructure, policy frameworks, and capital commitments that these challenges require. But across the sectors examined in this report, AI is proving to be a powerful force multiplier, one that improves how scarce resources are used. And in doing so, it delivers both financial value and sustainability outcomes from the same interventions.

The opportunity is large and growing. Our estimate of \$600 billion in global value reflects today's AI capabilities applied to today's sustainability challenges. As models improve, data infrastructure matures, and the physical impacts of climate change intensify demand for solutions, that figure will likely grow. So will the number of sectors where AI can expand what counts as climate and sustainability investing.

But scale is not inevitable. The constraints described in this report—the challenge of embedding AI in physical systems, data fragmentation, model drift, regulatory uncertainty, change management, and adoption resistance—are real and will determine the pace of adoption. Early deployment is not a guarantee of advantage; it is an opportunity to build the proprietary data, integration depth, and operational expertise that compound over time. Companies and investors that approach this space with both urgency and discipline will be best positioned to capture its potential.

Done well, AI enables companies to cut emissions, reduce waste, price risk more accurately, and reach people they could not previously serve. Done at scale, it reshapes how capital flows to climate and sustainability challenges, toward solutions that work, in sectors where financial performance and sustainability outcomes move in the same direction by design.

# About the Authors

**Daniel Oehling** is a managing director and partner at BCG's Singapore office. You may contact him by email at [oehling.daniel@bcg.com](mailto:oehling.daniel@bcg.com).

**Andrey Berdichevskiy** is a partner and associate director at BCG's Singapore office. You may contact him by email at [berdichevskiy.andrey@bcg.com](mailto:berdichevskiy.andrey@bcg.com).

**Vinay Shandal** is a managing director and senior partner at BCG's Toronto office. You may contact him by email at [shandal.vinay@bcg.com](mailto:shandal.vinay@bcg.com).

**Greg Fischer** is a partner and director at the London office of BCG. You may contact him by email at [fischer.greg@bcg.com](mailto:fischer.greg@bcg.com).

**Tariq Nanji** is a partner and associate director at BCG's Toronto office. You may contact him by email at [nanji.tariq@bcg.com](mailto:nanji.tariq@bcg.com).

**Benjamin Sheridan** is a managing director and senior partner at the Singapore office of BCG. You may contact him by email at [sheridan.benjamin@bcg.com](mailto:sheridan.benjamin@bcg.com).

**Reshma Jayaprabha** is a project leader in BCG's Singapore office. You may contact her by email at [jayaprabha.reshma@bcg.com](mailto:jayaprabha.reshma@bcg.com).

**Franziska Zimmermann** is Managing Director of Sustainability at Temasek.

**Alvin Low** is Director of Ecosperity and Director of Macro Strategy at Temasek.

**Jasmine Teo** is Director of Sustainability Strategy at Temasek.

**Kelly Chen** is Assistant Vice President of Sustainability Strategy at Temasek.

**Gladys Tan** is Assistant Vice President of Ecosperity at Temasek.

## For Further Contact

If you would like to discuss this report, please contact the authors.



### **Boston Consulting Group**

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

*This document stems from an engagement to write a commissioned report whereby BCG was compensated by Temasek.*

*This document is based on primary qualitative and quantitative research executed by BCG and Temasek. BCG and Temasek do not provide legal, accounting, or tax advice. Readers are responsible for obtaining independent advice concerning these matters. This advice may affect the guidance in the document. Further, BCG and Temasek have made no undertaking to update the document after the date hereof, notwithstanding that such information may become outdated or inaccurate. BCG and Temasek do not provide fairness opinions or valuations of market transactions, and this document should not be relied on or construed as such. Further, any evaluations, projected market information, and conclusions contained in this document are based upon standard valuation methodologies, are not definitive forecasts, and are not*

### **Temasek and Ecosperity**

Temasek is a global investment company, with a net portfolio valued at S\$434 billion (US\$324 billion) as at 31 March 2025. Our Purpose “So Every Generation Prospers” guides us to make a difference for today’s and future generations. We seek to build a resilient and forward-looking portfolio that will deliver sustainable returns over the long term.

Ecosperity is our key platform for sustainability engagement and advocacy. The word “Ecosperity” twins ecology with prosperity, reflecting our belief that doing good and doing well can—and must—go together. We work with global leaders from the private and public sectors, academia, and civil society to exchange views, share best practices and push the agenda on sustainable development.

*guaranteed by BCG and Temasek. BCG and Temasek have used data from various sources and assumptions provided to BCG and Temasek from other sources. BCG and Temasek have not independently verified the data and assumptions from these sources used in these analyses. Changes in the underlying data or operating assumptions will clearly impact the analyses and conclusions.*

*This document is not intended to make or influence any recommendation and should not be construed as such by the reader or any other entity. This document does not purport to represent the views of the companies mentioned in the document. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by BCG and Temasek.*

For information or permission to reprint, please contact BCG at [permissions@bcg.com](mailto:permissions@bcg.com). To find the latest BCG content and register to receive e-alerts on this topic or others, please visit [bcg.com](https://www.bcg.com). Follow Boston Consulting Group on [LinkedIn](#), [Facebook](#), and [X \(formerly Twitter\)](#).

