



ESG VIRTUAL SYMPOSIUM

AI, ESG & Responsible Innovation

Understanding the Environmental and Governance Implications of AI

For business leaders, sustainability professionals, security practitioners, and technology teams | July 2026

Panelists

Moderated by Angela Casesa
Wavelynx Technologies



Dr. Alka Khungar
Johnston Controls International



Audrey Singer
Brivo



What we'll cover

01

The environmental footprint of AI

Energy, emissions, and water behind training and inference

02

Why measurement is so hard

Data gaps, disclosure gaps, and a real-world case study

03

Responsible AI governance

Transparency, accountability, and the regulatory landscape

04

Balancing innovation and sustainability

A practical framework for ESG-aligned AI adoption

From environmental footprint to a full ESG lens

Responsible AI adoption touches every pillar of ESG — not just the environmental one.

E

Environmental

Energy use, GHG emissions, water consumption, e-waste, and resource extraction across AI's lifecycle

S

Social

Algorithmic bias, workforce impact, data privacy, misinformation risk, and equitable access to AI's benefits

G

Governance

Transparency, accountability, ethical use policies, oversight structures, and regulatory compliance

Today's focus has been the “E” — the remainder of this session turns to the “G” that helps manage all three.

Panel Questions

- E: Where in your organization is AI usage growing fastest — and are there gaps to allow you to track footprint?
- S: How can the security industry address bias and ensure fairness in AI-driven surveillance, access control, and threat detection systems?
- G: How do we balance the benefits of AI-powered security with growing expectations around privacy, consent, and responsible data stewardship?
- BONUS: How are companies at large considering the broader impacts across the value chain related to AI?
- BONUS: What are best practices that you can share around AI usage?

Executive Discussion Questions

1. Human-Centric Security (Executive Level)

"As AI becomes embedded in security operations, how do we ensure it enhances human safety, privacy, and trust rather than creating new social risks?"

2. Workforce Impact

"What responsibility do security organizations have to reskill and empower employees as AI automates more monitoring, analysis, and response functions?"

3. Bias and Ethical Use

"How can the security industry address bias and ensure fairness in AI-driven surveillance, access control, and threat detection systems?"

4. Trust and Transparency

"What level of transparency should organizations provide regarding how AI is used to make security decisions affecting employees, visitors, and customers?"

5. Data Privacy and Protection

"How do we balance the benefits of AI-powered security with growing expectations around privacy, consent, and responsible data stewardship?"

AI is no longer a side project — it's infrastructure

Which means its risks and impacts now belong inside the ESG conversation, not outside it.

1

Operations

Embedded in core workflows, forecasting, and logistics

2

Decision-making

Informing strategic and financial choices at scale

3

Security

Powering threat detection and access technologies

4

Productivity

Everyday copilots, assistants, and content tools

As adoption accelerates across all four areas — e.g., Brivo's Eeva and Genius Mobile Agent (Security) or Johnson Controls' OpenBlue building AI (Operations) — organizations must expand ESG strategies to address both the opportunities and risks.

The AI energy equation: training vs. inference

Every model carries two distinct energy costs — and inference is quickly becoming the larger one.

TRAINING

A one-time, massive computational burst — weeks or months of server farms running continuously.

1,287

MWh to train GPT-3 — powering ~120 U.S. homes for a year

~552

tons of CO₂e generated training that single model

INFERENCE

Ongoing energy use every time a model responds — multiplied across millions of daily queries.

~5x


more electricity per ChatGPT query than a standard web search

Weeks

to months before usage emissions exceed training emissions for popular models

Data center electricity demand is accelerating

Generative AI is a major driver behind the fastest period of data center growth on record.



4%

of global electricity use could come from data centers, crypto, and AI by 2026 — roughly Japan's total consumption (IEA)

7–8x

more energy an AI training cluster can consume vs. a typical computing workload (MIT / MCSC)

3x

the CO₂ data centers are projected to emit by 2030 vs. a no-AI-boom baseline (Morgan Stanley)

Beyond electricity: the hidden footprint

Energy gets the headlines, but water use, hardware waste, and material extraction all carry AI's costs too.



H₂O

Water

Training GPT-3 in Microsoft's U.S. data centers is estimated to have consumed ~700,000 liters of freshwater. A mid-sized data center can use ~300,000 gallons a day for cooling — comparable to 1,000 households.



E-W

E-waste

Short hardware lifespans and rapid GPU refresh cycles are projected to generate 16 million cumulative tons of AI-related e-waste by 2030, adding to an already fast-growing global waste stream.




Ore

Raw materials

Manufacturing GPUs and other AI hardware depends on mining lithium, cobalt, and rare earth elements — processes linked to habitat disruption, water pollution, and toxic-chemical use.

Not all AI use is equal

Task type and model efficiency drive dramatic swings in energy intensity.



Text is the least energy-intensive output; a single AI-generated image can use as much energy as half a smartphone charge, and generating video is far more intensive still.

Efficiency is improving fast

33x

reduction in median energy per Gemini Apps text prompt over 12 months

44x

reduction in carbon footprint per prompt over the same period

~9 sec

of TV viewing — the energy of one median Gemini text prompt

Case study: measuring your own AI footprint

CNaught, a carbon management company, tried to measure its own team's AI usage emissions in 2026 — and hit real obstacles.

2.5%

of CNaught's total organizational emissions came from AI usage in 2025 — not yet material, but rising

~4x

potential growth in that share if every employee used AI as heavily as the team's top user does today

Why it was so hard

- No hourly or daily usage granularity — only monthly summaries, and only a 90-day export window
- No location or grid-routing data, so emissions factors must rely on averages, not actuals
- Model-specific energy use per query is undisclosed — estimates lean on third-party research, not provider data
- AI embedded inside other software (copilots, summarizers) isn't captured at all

Bottom line: even a motivated, sophisticated team could only produce an estimate — not a verified figure.

The disclosure gap is the real bottleneck

Most organizations know AI has an environmental cost. Very few can quantify it.

24%

of businesses had integrated generative AI into their processes as of 2024

12%

of executives using generative AI are measuring its environmental impact

42%

of executives say AI use is forcing them to re-examine previously set climate goals

What would close the gap?

- ✓ Time-based usage data (hourly/daily) to match compute against grid cleanliness
- ✓ Grid routing and location data for accurate, non-averaged emissions factors
- ✓ Model-specific energy-per-query disclosure, alongside existing performance model cards

Emerging good practice in measurement

A handful of organizations are showing what better transparency can look like.

Salesforce

Publishing environmental metrics directly on AI model cards, alongside performance benchmarks.

Hugging Face

AI Energy Score Leaderboard rates models on energy efficiency — letting teams compare before they deploy.

Google

Published a full methodology for measuring energy, water, and carbon per AI inference, with year-over-year efficiency data.

CNaught

Released an open-source tool (Carbonlog) to help organizations track emissions from their own AI coding tool usage.

Responsible AI governance: the core pillars

Across every major framework — NIST, the EU AI Act, ISO 42001 — the same principles recur.

1 Transparency

Disclosing how AI systems are used, what data trains them, and their known limitations

2 Accountability

Clear ownership for AI outcomes, with escalation paths when systems cause harm

3 Ethical use

Guardrails against bias, misuse, and applications that conflict with organizational values

4 Human oversight

Meaningful human review built into high-stakes or high-risk AI decisions

5 Regulatory alignment

Active tracking of evolving rules across the jurisdictions where AI is deployed

6 Data protection

Safeguarding the privacy and security of data used to train and run AI systems

Privacy by design: five questions before you build

Johnson Controls builds these five questions into its AI Ethics Principles for every AI-enabled security deployment.

- 1 What specific security outcome are we trying to achieve?
- 2 What data is truly required to achieve it?
- 3 Who can access the data and the resulting insights?
- 4 How long will the information be retained?
- 5 How will individuals be informed, and where appropriate, given meaningful choices?

Organizations should also distinguish detection from identification — anonymized, de-identified, aggregated, or event-based data can often meet the security goal with less privacy risk.

The regulatory and framework landscape

Three frameworks anchor most organizations' AI governance programs today.

NIST AI RMF

United States — voluntary

Risk-management framework built around four functions: Govern, Map, Measure, Manage. Flexible enough to fit any industry or size — a practical starting point for internal programs.

EU AI Act

European Union — binding law

Risk-based classification with strict obligations for “high-risk” systems and human oversight requirements. Obligations are phasing in through 2026–2027.

ISO/IEC 42001

International — certifiable

A certifiable AI management system standard, similar in structure to ISO 27001. Increasingly treated as a credible signal of AI governance maturity to customers and investors.

Stakeholder expectations are rising fast

AI governance is moving from a technical nice-to-have to a board-level expectation.

Investors & ESG raters

Increasingly asking whether AI risk and environmental impact are addressed in ESG disclosures and ratings

Regulators

Expanding disclosure and human-oversight requirements as the EU AI Act and similar rules phase in globally


Customers & partners

Requesting evidence of responsible AI practices in vendor due diligence and procurement

Employees

Expecting clear, ethical policies on how AI tools are used, monitored, and governed internally

The AI value chain: upstream, enterprise, downstream



Responsible deployment reaches beyond IT — procurement, sustainability, legal, privacy, security, HR, operations, and product all have a stake.

Upstream

Before the tool is built

Energy, water, infrastructure, hardware, minerals, manufacturing, labor, and supplier practices required to develop and operate AI systems.

Within the Enterprise

While the tool is in use

Workforce impacts, data governance, cybersecurity, process redesign, training, and who benefits from or may be adversely affected by automated decisions.

Downstream

After the answer is given

How customers deploy the technology, possible misuse, maintenance and support, data retention, equipment replacement, e-waste, and what happens when a model or product is retired.

Balancing innovation and sustainability: a practical path

Six steps to bring AI inside your existing ESG strategy — without slowing innovation.

- 1 Inventory current AI usage across teams, tools, and vendors — you can't manage what you haven't mapped
- 2 Start tracking usage data monthly, even with imperfect granularity — waiting for perfect data means losing history
- 3 Favor providers who disclose energy, water, and carbon data, and match model size to the task at hand
- 4 Stand up an AI governance function or council with clear ownership for transparency and accountability
- 5 Fold AI-related environmental and social metrics into existing ESG reporting, rather than a separate track
- 6 Monitor the regulatory landscape (EU AI Act, NIST AI RMF, ISO 42001) and adjust policy as rules phase in

Practical habits from the panel

Five habits panelists say apply across organizations of any size.

- 1 **Know it:** maintain an inventory of AI tools and uses
- 2 **Classify it:** apply controls based on potential impact
- 3 **Protect it:** minimize data and use approved tools
- 4 **Verify it:** test outputs, performance, and bias
- 5 **Own it:** assign human accountability and monitor continuously

Quick win: turn off unused AI agents like devices left plugged in; and try downgrading the model or shrinking context once a workflow works — smaller footprint, smaller privacy risk.

Key takeaways

- ✓ AI's environmental footprint is real and growing — driven by energy-hungry data centers, water-intensive cooling, and short hardware lifecycles
- ✓ Inference, not just training, is becoming the dominant and ongoing source of AI's environmental impact as adoption scales
- ✓ Measuring that footprint accurately is currently very difficult — usage data, location data, and per-model energy data are all limited
- ✓ Responsible AI governance — transparency, accountability, oversight — is the throughline connecting environmental, social, and governance risk
- ✓ Stakeholders across investors, regulators, customers, and employees now expect visible AI governance, not just good intentions
- ✓ Organizations that start measuring and governing AI now — even imperfectly — will be far better positioned as disclosure expectations tighten

Thank you

Want to keep exploring AI adoption in security?

Register for "Driving Trustworthy AI Adoption in Security,"

Date: Friday, September 25th

Time: 11:00 a.m. to 12:00 p.m. EDT

It'll also be virtual via Zoom — you can find details and register on the SIA website by clicking below.

[Register Now](#)