AI Governance, from Innovation to Accountability

Artificial intelligence (AI) is consistently becoming a smarter investment for businesses across every vertical, from healthcare to the military, finance, and the arts. But even as AI becomes a must-have tool, it is vital to remember that this emerging technology presents both new frontiers and new opportunities for exploitation by bad actors, with regulatory and even consequences for companies that don't provide proper guardrails. These risks are why thoughtful and deliberate governance is so vital for any business leveraging AI in today's regulatory landscape.

As a worldwide hub for AI development, the Bay Area hosts multiple thought leaders ready to help businesses create winning governance strategies. The San Jose Business Journal recently hosted several leading voices in AI governance at a vibrant roundtable, including Chief Information Security Officers (CISOs) and Chief Legal Executives for trailblazing Silicon Valley companies. Topics included legal accountability, cybersecurity, and a roadmap for ensuring thoughtful governance for all AI solutions.

Ray Cheung opened the discussion by framing a tension many organizations are facing today: AI is moving faster than any governance model we've ever built, yet the cost of getting it wrong has never been higher. He then posed a central question to the group:

ALL PHOTOS BY NOEL RABINOWITZ

From left to right: David Tugwell, Agilent; Leslie Stevens, Agilent; Wei Li, Intel; Rick Orloff, Pure Storage; Lindsay White, Applied Materials; Tim Howard, Freshfields; Jannie Affeld, Google; Ray Cheung, Crowe; Serge Jorgensen, Crowe; Aman Sirohi, Cyberhaven

> "As AI scaling risks outpace governance, success belongs to the company with the most trusted AI, not the company with the most AI. At Crowe, we've established four pillars for governance. Our first pillar is determining if AI enhances innovation. The second pillar is the ability to assess AI readiness. The third pillar is organizational value and how AI policy aligns with governance principles. Security and ethics have to inform your AI strategy. The fourth pillar is continuity and being able to maintain governance over a long period of time."
>
> **RAY CHEUNG**
> Crowe LLP

RAY CHEUNG, TECHNOLOGY, MEDIA & TELECOMMUNICATIONS INNOVATION LEADER CROWE LLP: What are some key concerns regarding legal accountability and regulatory risk when it comes to AI?

TIM HOWARD, PARTNER, FRESHFIELDS: As innovation moves rapidly, it is important to be aware that it is easy to get caught up in a DOJ investigation. I can say this from my 12 years as a federal prosecutor at SDNY where I was the Chief of the Complex Frauds and Cybercrime Unit and led countless investigations into misuse and harms arising from technology. AI is newsworthy, scares the public, and its sexy for a DOJ office to open a case about the misuse of AI. A mainstream news article describing something wrong with AI controls signals a race between DOJ components to a subpoena. So proper governance is important to reduce the risk and impact from such an event that can carry negative publicity, bring concerns from auditors, and lead to substantial costs from even responding to an investigation, even if you prevail the government to not bring an enforcement action.

When it comes to securities fraud, especially as concerns AI, prosecutors are only getting started. Any statement your CEO or CISO makes to the public, or anything you include on your website, can be evidence of misrepresentations. Be careful of statements that may exaggerate or over inflate your use and power of your AI, qualify statements and avoid absolutes, so as to avoid ending up on the wrong end of a lawsuit or a criminal securities fraud investigation.

AMAN SIROHI, SVP, CHIEF SECURITY OFFICER & IT, CYBERHAVEN: AI governance is no longer optional. CISOs need to partner closely with legal, compliance, and the business to ensure innovation doesn't outpace accountability. This means embedding governance into the lifecycle from model selection and data sourcing to deployment, monitoring, and incident response not treating it as an afterthought. The organizations that win won't be the ones that move fastest but the ones that move responsibly.

TIM HOWARD: Exactly. The best way to think about reducing costs from potential investigations is thinking backwards from a potential presentation to regulators or enforcement authorities – how can we develop a record today to demonstrate that we are responsible. It's about explainability, explaining reasonableness. Bad things may happen with dynamic technologies. You need someone whose job is to think about these eventualities and how to build a defense against accusations. If considered risk-based decisions are wrong, that may be a bad judgment call, but not the basis for criminal liability.

Companies leaning heavily into developing AI platforms should consider their relationship and reporting protocol to law enforcement on the platform because that can be powerful evidence of good faith. When something goes wrong, when a problem slips past your radar, regulators can see the 50 other times you reported to the proper agencies when you've caught misuse, or potential harm to others, to demonstrate good faith. Evidence of cooperation negates criminal intent.

Note that a corporation is liable under

US law criminally for the acts of any of its agents, and prosecutors and the Department of Justice can argue that the actions of an AI agent count. So even if no individual goes to jail, an AI agent that ends up causing criminal harm could lead to that corporation facing criminal exposure.

WEI LI, VP/GM, AI & ANALYTICS (FMR), Intel: Drawing on experience leading large-scale AI and analytics platforms at Intel, collaborating with clinicians at Stanford University School of Medicine, and advising executives and boards on responsible AI adoption, I often come back to the idea of Precision AI Governance. It's a framework that aligns validation, oversight, and accountability with the impact and autonomy of each AI use case, enabling enterprises to scale AI safely, reliably, and responsibly.

A good example is Air Canada, where a chatbot gave a customer information that was not true. The chatbot made it up and kept the customer from getting a refund. Air Canada couldn't say "not my problem." They went to court, and the customer proved that the chatbot misunderstood the airline's bereavement policy, and Air Canada was liable. Companies are held responsible for the actions of their agents, including AI agents, no matter who produced the AI tools.

RAY CHEUNG: As AI scaling risks outpace governance, success belongs to the company with the most trusted AI, not the company with the most AI. At Crowe, we've established four pillars for governance. Our first pillar is determining if AI enhances innovation. The second pillar is the ability to assess AI readiness. The third pillar is organizational value and how AI policy aligns with governance principles. Security and ethics have to inform your AI strategy. The fourth pillar is continuity and being able to maintain governance over a long period of time.

How do those pillars resonate, and how are governance frameworks a guardrail rather than a roadblock within your organizations?

JANNIE AFFELD, VICE PRESIDENT, FINANCIAL ENGINEERING, GOOGLE: Control and governance are always top-of-mind for AI. We're leveraging whichever application access control and system access control frameworks we can to contain questionable activity, trying to not to let anything go.

RICK ORLOFF, VP, CISO, EVERPURE: On the governance side, AI is in its infancy. It's fantastic, but looking back five years from today, we're going to say this was just the beginning.

I look at AI as a swim lane with two guardrails. The guardrail on the left represents governance, risk, and compliance, plus policy, procedure, standards, and guidelines. The guardrail on your right becomes technical controls and monitoring capabilities. Businesses can swim as fast as they want between those guardrails, but security owns the rails. If those two guardrails are in place, they may have been able to mitigate the risks that have manifested in the AI space.

AMAN SIROHI: These pillars resonate to me and my organization. Innovation and readiness are foundational organizations need to ask if AI truly adds value and whether they're prepared to support it. Alignment ensures AI policies reflect risk tolerance, regulatory needs, and business priorities. And adoption is key for success. Done right, governance frameworks act as guardrails, not roadblocks. They create clarity, speed up decision-making, and give teams confidence to innovate safely.

SERGE JORGENSEN, PRINCIPAL, CYBER CONSULTING, CROWE LLP: Something we're just starting to see is companies getting frustrated by vendors telling them two months later, six months later, "Hey, we had this incident, and we're working to figure out what your data is." Companies are pushing back as part of their vendor agreements to insist that vendors inform them and keep them informed of all processes. One question coming up from a legal perspective is, "Would you be willing to sign a joint defense agreement?" But then vendors are pushing back, saying, they'll happily share data if there's no insistence on a joint agreement. All that has to be negotiated as part of the risk assessment process and the contractual assessment.

As the conversation continued, the group discussed how many organizations are approaching AI governance by adapting existing technology review frameworks rather than building entirely new systems from scratch. Several leaders noted that reviewing AI through the same lens as other enterprise technologies allows companies to stay agile while still accounting for unique AI-related risks.

Participants emphasized the importance of incorporating AI-specific considerations into established cybersecurity, legal, and compliance procedures, particularly around vendor assessments, deployment safeguards, and governance reviews. Rather than creating separate governance structures, many organizations are refining proven internal processes to evaluate AI responsibly while maintaining speed and flexibility.

"Control and governance are always top-of-mind for AI. We're leveraging whichever application access control and system access control frameworks we can to contain questionable activity, trying to not to let anything go."

JANNIE AFFELD
Google

DAVID TUGWELL, CISO, AGILENT: We need to build and apply new safeguards into our development cycle, so we can have space to experiment before production.

RAY CHEUNG: Cybersecurity concerns are difficult to navigate because businesses can pay the price for an external bad actor's behavior. What are some ways governance can help mitigate cybersecurity risks?

RICK ORLOFF: In some areas we are being prescriptive on the documentation and requirements, and in others we are intentionally not being prescriptive because we don't want to create unnecessary liability. There are two types of CISOs: risk-averse, and those who understand how to accept calculated risk. A risk-averse CISO generates a lot of policies and procedures the company's not going to follow. Doing so may increase the liability for the company. It's best to maintain a realistic policy portfolio and not create policy just for optics.

DAVID TUGWELL: A company needs one risk owner. We might be in conversation about some sort of risk, and if there are multiple owners, we're making decisions without understanding upstream or downstream impacts. You're not going to get that big picture view unless somebody understands the entire horizon, across the entire enterprise. It's not just cyber risk or enterprise risk. If there's a market risk for not releasing the product, you have to weigh that with the other risks.

LINDSAY WHITE, MANAGING DIRECTOR, CHIEF OF STAFF, APPLIED MATERIALS: We're constantly reassessing what data types we're pulling in, so that

we can control the data more securely and block unwanted information before it gets pulled into an application we don't want it in. This is especially important now that people can bring in data from any number of sources.

AMAN SIROHI: Cyber risk often comes from outside the organization, but governance helps you stay in control. It creates clear ownership, enforces consistent security practices, and builds readiness to respond quickly. Governance isn't a blocker it's a guardrail that helps organizations reduce risk and operate with confidence.

LINDSAY WHITE: How much power the vendor has in various situations will factor into how everyone can play in the same sandbox together. As a tech company looking at all these different AI vendors, many are very young and don't have robust security infrastructure. In a more conservative, risk-averse environment, we're going to look first at the players who have been around for a long time and who may be developing something similar. Bigger vendors tend to have the security protocols we require.

RAY CHEUNG: According to Crowe TMT clients' interviews, alongside Gartner's survey, 70% of companies are scaling AI, but fewer than 30% feel they have adequate governance in place. That disconnect is going to exacerbate risk in the next few years.

Can you describe some ways you've shaped company culture to help team members assist with governance and minimize risks?

TIM HOWARD: One point about culture can be analogized to what happens with cybersecurity breaches - there's no such thing as perfect cybersecurity. If you're at the forefront of AI, things are going to go wrong. You need a culture that accepts that, and that everyone is in it together. Everyone will work together to mitigate issues. Maybe someone madea the wrong risk-based decisions in retrospect, playing Monday morning quarterback. But the whole team is going to do their best together.

You are in a bad place facing a government investigation if you have a culture where someone's head is going to roll because something bad happened in a race, because then it becomes an internal race to protect yourself. You need a culture that stays reasonable when bad things happen and works together to fix the problem as management figures out what went wrong — but in a mature way, without pointing fingers. People freaking out will only create more risk.

LINDSAY WHITE: How do you inform your global employee population of a problem, the need to educate versus causing panic? Oftentimes, employees don't realize these issues can happen at your company. It's not just something made up in the news. There's a balance between calling something a threat versus a risk while educating. The goal is that when they are targeted, they don't click the link or respond to the phone call, or otherwise risk security.

DAVID TUGWELL: Being able to tell people what they can and can't safely do is so important. We have a new CFO who recently made it very clear that we will never ask our employees for money. It's simple stuff, but it helps team members understand that if they get a call like that, no matter who it is from, they can just hang up.

RICK ORLOFF: Many years ago, I saw people panic because they received a text and thought the CEO was asking them to buy gift cards. Literally gift cards, for our CEO. I finally sent an email to the entire company that said, "Our CEO is absolutely never ever going to ask you for money, never going to ask you for your private information, and never going to ask for gift cards. If they ever do, don't buy them, call me personally."

I coordinated with the CEO, who then responded back to the entire company and said, "That is absolutely right."

RAY CHEUNG: Share any closing thoughts on how to install guardrails to ensure proper AI governance.

SERGE JORGENSEN: In 2017, Russian president Vladimir Putin said that whoever figures out AI is going to rule the world. So we have to put in guardrails to make our systems safe and defensible while also encouraging people to move faster. Healthcare is the perfect example of how we keep innovation going without exposing ourselves to undue risk. CISOs and cybersecurity and legal teams are always trying to strike that balance.

"I look at AI as a swim lane with two guardrails. The guardrails on the left are governance, risk, and compliance, plus policy, procedure, standards, and guidelines. The guardrail on your right becomes technical controls and monitoring capabilities. Businesses can swim as fast as they want between those guardrails. But we own the rails. Look at some of the risks that have manifested in the AI space. If those two guardrails had been in place, they could have mitigated a lot of those risks."

RICK ORLOFF
Everpure

But I would encourage everyone to err on the side of accepting risk and qualifying risk in order to have a good story to tell about it and plan for failure instead of just building the walls higher.

WEI LI: My current interest is AI in the healthcare field, and that balance is very important. I've been working with the school of medicine at Stanford doing heart surgery. Providers have to consider: What is the governance model for AI assisting with major surgery? What is the infrastructure software for that? How do you apply AI to real applications and how much is the surgeon's responsibility, and how much is the AI's?

RICK ORLOFF: I'm after technical controls to enforce, and I'm continuing to expand visibility and observability because we can only control what we know. I want the best customer experience. I don't want so many false positives that I can't spot the real threats to defend my customers. AI and governance are going to mature. Visibility will get better. A lot of startup companies are focusing on exactly that problem.

JANNIE AFFELD: It's vital to share best practices on why governance is foundational and central for AI enterprise. Ultimately, governance should always add to a business's structure, not create more tension within the structure.

RAY CHEUNG: Every company today can build AI fast. Very few can govern it at scale. That's where the real competitive divide is emerging. The companies that win the AI race won't be the fastest – they'll be the most trusted.

Crowe