

# BLOOM ETHICAL AI CONSULTING

ARCHITECTING ETHICAL AI

© 2026 Bloom Ethical AI Consulting, LLC. Proprietary and confidential. No part of this publication may be reproduced without the express written consent of the author.

## MORAL ARCHITECTURE NEWSLETTER

Week Ending May 30, 2026 | Issue No. 9

Translating AI Governance Into Everyday Insight | For professionals, practitioners, and the curious

### FIVE QUESTIONS TO CONSIDER BEFORE YOU READ THIS WEEK

- 1 When the CEO of Google DeepMind says AGI could arrive as early as 2029, does your organisation have a plan for what that means for your workforce, your decisions, and your accountability structures?
- 2 If the FBI's 2025 Internet Crime Report documented 22,364 AI-related fraud complaints totaling \$893 million in losses, has your organisation briefed employees on how to verify voice communications before acting on urgent requests?
- 3 Can you name every AI agent your organisation has deployed, every account it operates through, and every action it is authorised to take without a human in the loop?
- 4 When Google announces that the shift from chatbots to autonomous agents is no longer a future roadmap but a present deployment, has your governance framework been updated to match that reality?
- 5 If a voice clone of your CEO called your finance team today with an urgent wire transfer request, what verification procedure would stop it from succeeding?

Three developments defined this week in AI governance. Google DeepMind CEO Demis Hassabis, speaking at Google I/O 2026, narrowed his AGI prediction to 2029 or 2030 and described today's AI agents as a societal practice run for far more powerful systems. The FBI's 2025 Internet Crime Report documented AI-related fraud losses of \$893 million in a single year, with voice cloning now driving scams that impersonate family members, executives, and government officials. And Google's I/O keynote declared the shift from AI chatbots to autonomous agents complete, announcing Gemini 3.5 and a full agentic enterprise platform. The convergence is not coincidence. The governance gap between what AI can do and what organisations are prepared to manage is widening in real time.

### FEATURE

#### ARTIFICIAL GENERAL INTELLIGENCE

## Google DeepMind CEO Narrows AGI Timeline to 2029, Calls Agents a Practice Run

At the sidelines of Google I/O 2026 in mid-May 2026, Google DeepMind CEO Demis Hassabis told Axios co-founder Mike Allen that humanity is standing in the "foothills of the singularity," drawing a sharp reaction from the audience. Hassabis said he now expects AGI, the point at which AI systems can match human-level intelligence across a broad range of tasks, to arrive around 2030, with 2029 emerging as a genuine possibility. "I've been saying, recently, around 2030, plus or minus a year, I think is a reasonable estimate, from what I'm seeing now," he said. This represents a significant narrowing of his prior prediction of 2030 to 2035.

Hassabis described the current wave of AI agents as a societal stress test. "You can imagine the agentic era in this next year is a little bit like a practice run," he said. He added that one looming milestone is recursive self-improvement, AI systems capable of materially accelerating their own development. "All the leading labs are quite focused on that," he said. Hassabis also warned that the conversation about AI's society-reshaping impact remains largely confined to technology circles. "My economist friends, I feel, are still not taking this seriously enough," he said.

### WHY THIS MATTERS TO YOU

When a Nobel laureate and the head of the world's most decorated AI lab says AGI could arrive in three years, the governance question is not whether to believe him. It is whether your organisation has even begun to think through what that means for human oversight, accountability, and decision authority. Moral Architecture holds that governance frameworks must be built before the systems they are meant to govern arrive. Waiting for AGI to prompt preparation is not caution. It is abdication.

Sources: Axios, May 26, 2026 (axios.com); Computerworld, May 28, 2026 (computerworld.com); Sherwood News, May 28, 2026 (sherwood.news); Digit.in, May 28, 2026 (digit.in).

### TECHNOLOGY AND RISK

#### FRAUD AND IDENTITY

## FBI Documents \$893 Million in AI Fraud Losses, Voice Cloning Now a Primary Vector

The FBI's 2025 Internet Crime Report, released in April 2026, includes a dedicated AI section for the first time in its nearly 25-year history. The IC3 documented 22,364 complaints with losses of \$893 million, attributed substantially to voice cloning attacks. A convincing clone requires as little as three seconds of audio. CNN reported on May 29, 2026 that a California mother lost thousands of dollars to a scam using a cloned version of her daughter's voice. Congressional researchers estimate fewer than 5 percent of victims report, meaning actual losses are vastly higher. Deloitte projects total US AI fraud losses could reach \$40 billion annually by 2027.

### WHY THIS MATTERS TO YOU

The FBI documented \$893 million in losses and estimates only 5 percent of victims report, meaning the real damage is orders of magnitude larger. Voice cloning is a current, documented attack surface. Does your organisation have an out-of-band verification protocol, or are urgent requests still acted on because a voice sounds familiar?

Sources: FBI Internet Crime Complaint Center, 2025 Internet Crime Report (fbi.gov, April 2026); CNN, May 29, 2026 (cnn.com); GovTech, April 9, 2026 (govtech.com); CybelAngel, April 27, 2026 (cybelangel.com); TechTimes, May 29, 2026 (techtimes.com).



AI AGENTS AND GOVERNANCE

## Google Declares the Agentic Era Has Arrived, Deploying Autonomous AI Across Enterprise

At Google I/O 2026 in mid-May, CEO Sundar Pichai declared the company is now in the "agentic Gemini era," signaling a fundamental shift from AI systems that respond to questions to AI systems that act autonomously on behalf of users and organisations. Google announced Gemini 3.5 Flash, described as running four times faster than prior models and optimised for low-latency agentic workflows, as well as Gemini Omni, a multimodal model designed to generate any output from any input. Google Cloud simultaneously launched the Gemini Enterprise Agent Platform, positioning Google Cloud as, in the words of CEO Thomas Kurian, "the new front door for AI in the workspace."

Google also introduced Gemini Spark, a 24-hour AI agent that runs continuously in the background across devices and apps, completing long-running tasks without active user direction. AI Mode in Google Search has crossed 1 billion monthly users, while AI Overviews reach 2.5 billion users per month. Pichai told the audience that Google has been AI-first for a decade, but that the current moment represents the shift from AI as a feature to AI as the operating model. The developer keynote confirmed that Google's Antigravity agent platform now enables full agentic orchestration at the enterprise tier.

### WHY THIS MATTERS TO YOU

When the world's most widely used technology platform declares that autonomous agents have replaced chatbots as the primary AI deployment model, the governance question is urgent and concrete: who authorised each agent your organisation is now running or about to run, what can it access, and who is accountable when it acts unexpectedly? Google's shift is not a product announcement. It is a governance event. Most organisations have not yet built the oversight infrastructure to match the deployment pace.

*Sources: Google Blog, I/O 2026 keynote, May 2026 (blog.google); Google Cloud Blog (cloud.google.com); Google Developers Blog (developers.googleblog.com); Medianama, May 2026 (medianama.com); Medium / Krishna, May 2026.*

### MORAL ARCHITECTURE INSIGHT

#### THIS WEEK'S PRINCIPLE

*Three stories, one truth: the pace of AI capability is now outrunning the pace of governance preparation. The head of the world's leading AI lab sets a three-year AGI horizon. The FBI documents nearly a billion dollars in AI fraud losses in a single year. The world's largest technology company declares autonomous agents the new default. Governance built for yesterday's AI will not hold tomorrow's AI accountable. Build before you are required to respond.*

#### ONE QUESTION TO ASK IN YOUR ORGANISATION THIS WEEK

**If an AI agent in your organisation acted outside its authorised scope today, would you know within the hour, and could you terminate it, trace its actions, and explain what happened to a regulator?**

### ALSO WORTH KNOWING

- **EU Article 50 Comment Deadline: June 3, 2026.** The stakeholder comment window on the European Commission's draft Article 50 transparency guidelines closes June 3. August 2 enforcement is weeks away. Chatbot disclosures, deepfake labelling, and emotion recognition obligations are live on that date.
- **Connecticut SB 5 Signed May 29, 2026.** Governor Lamont signed one of the most comprehensive US state AI laws. AEDT obligations take effect October 1, 2026. Companion chatbot safety and frontier model whistleblower protections are now law. Inventory your AI tools now.
- **Colorado AI Act: Enforcement Stayed, Obligations Remain.** A federal court has stayed enforcement in *x.AI LLC v. Weiser*. The stay may lift. Continue ADMT reclassification work and do not treat a litigation stay as a governance holiday.

### CLOSING QUESTION FOR THOUGHT

*As AGI timelines shorten, AI fraud losses reach nearly a billion dollars in a single year, and autonomous agents replace chatbots as the default enterprise AI model, ask yourself honestly: is your organisation building governance infrastructure at the speed AI is being deployed, or are you still writing policies for systems that were already obsolete before the ink dried?*

### ABOUT THE AUTHOR

**Andrew Bloom, Founder and Principal, Bloom Ethical AI Consulting.** Andrew Bloom has built original AI governance frameworks from ethical first principles. As an ordained Rabbi and founder of Moral Architecture, his work is grounded in formal ethical and theological training. Author of *The Ten Commandments of Ethical AI* and *Technology and Theology*. Founder of Bloom Ethical AI Consulting and Rabbi of Congregation Ahavath Sholom in Fort Worth. This newsletter is separate from the board-level Executive Brief. Advisory enquiries: 817-616-1313.

© 2026 Bloom Ethical AI Consulting, LLC. All rights reserved. No part of this newsletter may be reproduced without the express written consent of the author. Moral Architecture™ is a framework developed by Andrew Bloom, Bloom Ethical AI Consulting, LLC.