The RISE of the CHIEF AI OFFICER

TYPE IN A BASIC SEARCH FOR "CHIEF AI Officer" into LinkedIn and you get 1,800 results, a number that rises by the day. From NASA to L'Oréal, from national oil companies to nation states, the new executive position to have is "Chief AI Officer." Companies like Accenture and Microsoft have multiple CAIOs. Last June, Dubai announced plans for 22 Chief AI Officers across its government, and the US previously mandated that every federal agency and military service had to hire chief AI officers.

This explosive growth underscores the rising demand for AI leadership across industries. Notably, 2024 marked a key milestone, with the number

Brunswick's CAROLINE DANIEL talks to leading CAIOs about growing demand and expectations for the new senior executive role.

of CAIOs surpassing 1,000 (up from 250 in 2022)— a tipping point that mirrors the trajectory of Chief Digital Officers (CDOs) and Chief Data Officers (CDOs) before them and signaling the mainstream adoption of AI leadership as a critical priority for organizations worldwide.

The rise of the role has been swift. The first media reference was in 2016, but interest pretty much flatlined until the launch of ChatGPT in early 2023.

By the end of that year, Forbes was setting out "The Case for the Chief AI Officer—A Role whose Time has Come." A few months later, the CDO Club held the first CAIO Summit at Northeastern University in Boston hosted by the Institute for Experiential AI and the D'Amore-McKim School of Business;

500 CAIOs attended, most clustered around technology companies.

So, what's driving this shift? I asked David Mathison, Chairman, CEO and Co-founder of the CDO Club and CAIO Summit, and a leading authority on Chief AI, Analytics, Data, and Digital Officers. Some of this, he says, is title inflation: "Fully half of the 1,800 CAIO profiles on LinkedIn are unqualified people at mom-and-pop shops, startups, at AI companies, or people who are no longer in the role using the title as clickbait to attract investors, headhunters, analysts or the media."

A second, and in some ways equally misleading, driver is an enthusiasm for Generative AI. At his recent conference, Mathison showed a job description to his audience of chief data and AI officers. "It said they wanted a 'chief AI officer.' But what they really wanted was a GenAI person. People are confused, even at the highest levels of companies."

The skills needed in the role are challenging and varied. What companies should be asking, Mathison says, is: "Where have they deployed AI? What have they learned from their mistakes. What teams can you bring to the table, because being able to attract talent to the organization is critical. You want to be able to parachute someone into a company and bring in a dozen of your top AI leaders."

On the technical side, many of the best CAIOs have a Ph.D. or a master's degree. "Like Chief Data Officers," Mathison says, "CAIOs have the most Ph.D.s of any other job title on Earth." Absent those degrees, the best have strong mathematics, statistics, AI, a good grounding in data science and a minimum of 10 to 15 years of experience in deploying machine learning models, leading data science and AI teams and using cloud technologies.

Also marking exceptional CAIOs are their softer skills like culture change management and the ability to deploy AI responsibly, he says.

"You also need to be able to talk to the C-suite, to business managers, and find out where this drives business value, and have the soft skills of delivering responsible, ethical and trustworthy AI. Some people might be really good at GenAI, but terrible at delivering it across the enterprise. That's why it's important to get people who understand both the business implications of AI and then understand responsible, trustworthy, ethical AI.

"AI is an exponential technology, evolving at a pace that human skills simply can't keep up with. The toothpaste is out of the tube, and this could lead to serious reputational harm for companies. We're in uncharted territory."

DAVID MATHISON
Chairman, CEO
and Co-founder, the CDO
Club & CAIO Summit

> "AI IS AN EXPONENTIAL TECHNOLOGY, EVOLVING AT A PACE THAT HUMAN SKILLS SIMPLY CAN'T KEEP UP WITH. … THIS COULD LEAD TO SERIOUS REPUTATIONAL HARM FOR COMPANIES. WE'RE IN UNCHARTED TERRITORY."

These "softer skills" are no longer that soft. Getting real value out of AI is increasingly dependent on the ability to create the right AI culture, along with the right understanding of the real impact of AI on their organization.

There is a significant trust gap between CEOs and their workers. CEOs are excited. They trust AI's potential: around half consider it a top priority, in typical surveys. Employees, however, trust AI far less. The World Economic Forum reported last year that only 55% of employees are confident their organization would ensure AI is implemented in a responsible and trustworthy way.

How CAIOs help to build that trust is essential. Part of it is addressing the jobs concerns head on and ensuring employees have access to the best tools, technology and training, allowing them to start to experiment and understand how it can be useful in their work.

Daniel Hulme is one of the early Chief AI Officers. He joined WPP in 2021 when it bought his AI company, Satalia. He was recognized by *AI Magazine* as one of the Top 10 CAIOs globally in 2023. "Part of my job is to give people an understanding about what these technologies are, what they can and can't do, and what they might be able to do," he says. "Once empowered with that understanding, people feel better about how they can direct their own destiny."

Beyond his own technological credentials—with a Master's and Doctorate in AI at London-based international university UCL—Hulme recognized that responsible use of AI is fundamental to any successful deployment. "My job is to figure out what's our AI strategy over the next three years, make sure that we're tracking the trajectory of the technology and placing the right bets in terms of governance. I'm very interested in how we deploy these technologies safely and responsibly within our organizations."

There are three questions he advises organizations to ask when implementing AI. "First, is the intent appropriate? Many people have rebranded themselves as AI ethicists. I would controversially argue there's no such thing as AI ethics. The difference between AI and humans is that humans have intent. AIs don't. There are well-established frameworks set up to scrutinize intent—you don't need a new AI ethics committee."

His second question is: "Are my algorithms explainable? The difference between software and AI is that AIs tend to be opaque in terms of how they make their decisions. So when we build systems, particularly those that have a material impact on

people's lives, we try to make sure they are explainable, to mitigate any risks."

Even so, there are real challenges in ensuring that the outputs of the current large language models (LLMs) can ever be truly explainable, given they are trained on billions of data sources.

One solution is to use AI itself to help. WPP doesn't want to remove humans from the loop altogether, Hulme stresses. But he offers an example that could help address the challenge of a world in which people can create ads in seconds, together with a mechanism to test whether those ads are safe and responsible.

"We can actually build LLMs that represent different corners of society. I can build an agent that represents a political party or a newspaper, or a culture or minority group, or even a food compliance framework, or ad compliance framework, or sustainability framework. We can show these rapidly created ads to thousands of 'experts', and see if we're going to trigger any communities, break any laws or cause any harm. We're trialing issues of greenwashing to identify if ads might be greenwashing, for example."

The third question Hulme suggests organizations ask is a strange one, he admits. "What happens if my AI goes very right? As engineers, when we build systems, we try to identify and mitigate failure points, but now what we have to ask ourselves is, what happens if we overachieve our goal and it starts to cause harm elsewhere? There are lots of examples where an AI has massively achieved the KPI that we've given it, but caused harm in other KPIs."

One example Hulme provides is personalized marketing and the risk of it reinforcing the human bias that people have to engage with those who look and sound like us. "If we let AI loose to optimize marketing content, you might end up with a world of ads selling just to you. That might enforce bias, bigotry and social bubbles. What happens if we create a post-truth world? What happens in terms of the impact on jobs and how to retrain?"

Hulme's thinking on these subjects recently led him to found Conscium, an AI research lab, separate from WPP, that brings together leaders in neuroscience, evolutionary computation and deep learning, to explore how to build safe AI that benefits humanity. At the end of 2024, Conscium launched a new app, Moral Me, to learn more about human morality and how people feel about having AI even more integrated into their lives and the ethical questions that arise as it starts to take on more human-like roles—a growing topic as we enter an era of AI-agents—or
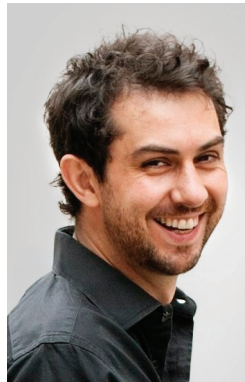
start to create "digital twins of ourselves" acting on our behalf.

A more immediate concern is getting access to the right CAIO talent right now. "There are very few true chief AI officers out there, like Daniel," says Mathison. "Folks that have 10 or 15 years of experience, seasoned veterans, with successes and failures, in delivering enterprise-wide AI. I call them unicorns. Salaries for them are through the roof. I've been tracking them since 2020 when there were 250. Now there are 1,000."

Demand for CAIO's is also shifting geographically fast. "There were no chief AI officers in the Middle East. There's none in South America, there's practically none in Canada. Across Europe, there's just a few. Then suddenly the Middle East has moved—Dubai is now second or third as a region, given the 2024 mandate to hire CAIOs across Dubai's government agencies," he says.

One consequence of that: There is likely to be even more of a coming disparity emerging between the companies with the best AI talent and CAIOs and those who are catching up. Small- to medium-sized businesses, nonprofits, regulators and government agencies all risk falling further behind as they struggle to attract the best talent.

Or, as Hulme warns, they end up hiring the wrong people. "A lot of people get excited about emerging technologies—currently, that's Generative AI. What is happening is that people are rebranding themselves and saying, 'I'm going to now build a team and a career around this.' But would you hire somebody that's just passed the Bar and had a few years of experience to be your general counsel and dictate business strategy? People end up focusing on one sort of exciting set of technologies. They then apply those technologies to solving the wrong problems and blame the technologies rather than themselves."

As the race for talent tightens, unconventional approaches are on the rise, says Mathison. "Traditional executive search for this title is completely broken, which is why the CDO Club has now launched a new service on fractional chief AI officers. Companies need to get somebody in … today."

Hulme agrees that the fractional approach may be the best route for many companies in the short term. "It all depends on how disruptive you think AI is going to be to your organization. For the most part, the CIO, the CTO, can get knowledgeable enough to make sure that they're placing the right technology bets. But if your industry is reliant heavily on AI, you're going to need to have a spokesperson." ◆

> ## "I WOULD CONTROVERSIALLY ARGUE THERE'S NO SUCH THING AS AI ETHICS … YOU DON'T NEED A NEW AI ETHICS COMMITTEE."

DANIEL HULME
CAIO, WPP

CAROLINE DANIEL, a Partner based in London, is part of Brunswick's AI Client Impact Hub. Caroline is a former Editor of *FTWeekend*, senior editor and correspondent with the *Financial Times*.