

Establishing a Learning Officer Aligned with the CTO in a Midsized Business and Augmenting Your Workforce with AI Avatars

Dr. Anastassia Lauterbach

ORCID ID 0009-0006-6988-7406

October 2026

Executive Summary

Every incumbent company has a handful of people who carry the business in their heads: the service engineer every customer asks for, the underwriter who never misprices a risk, the teacher whose students always pass. For decades, these people could only be in one place at a time. AI now offers a way to build *digital humans* from them: conversational replicas that speak with their knowledge, and sometimes with their voice and face.

In a conversation on AI Snacks podcast, Sam Huber, who leads Napster's enterprise business from Dubai, described the goal as the *elastic organization*: a company that can scale its workforce up and down the way cloud computing scales servers.¹ Digital humans built from the best people are one of the most powerful, and most misunderstood, instruments of that elasticity.

This white paper is not about AI agents. Generic agents, built to book meetings, triage tickets, or write code, are a different discussion. Our focus here is narrower and more personal: digital humans modeled on *specific, identifiable people*, built with their consent, from their knowledge, their stories, and often their likeness.

edutAlnment proposes five positions, ordered by urgency:

1. **No AI is a reflection of the full person. It is a reflection of a person on an average day.** A digital human reproduces how someone usually explains, advises, and decides. It does not reproduce the rare, high-stakes moments that made that person exceptional.
2. **Digital Warren Buffetts do not belong in the boardroom.** A replica of a famous investor or a retired chairman offers the average of their public record, frozen in time, without accountability, and with a built-in invitation to fraud.
3. **Training, onboarding, and knowledge capture are where digital humans already create value.** Companies from Network Rail to GEMS Education show what works: narrow scope, clear consent, and the original expert in the loop.
4. **A digital human is an IT asset of the highest protection class.** It combines biometric data, personal knowledge, and company secrets in one system. Storage, backup, and cybersecurity must be designed for that before the first recording is made.
5. **Someone must own the capability.** We recommend a Chief Learning Officer reporting to both the CHRO and the CTO, responsible for AI literacy, AI fluency, and the knowledge inside every digital human. Midsized companies

that are too small for this role should share it through their ecosystem: private equity owners, industrial associations, and chambers of commerce.

From Agents to Digital Humans: What We Are Talking About

The vocabulary in this market is loose, so precision matters. An *AI agent* is software that performs tasks. A *digital human*, in the sense of this paper, is an AI system built to represent a particular person: their knowledge, their way of reasoning, and optionally their voice and appearance. It answers as *them*, not as a generic assistant.

Rob LoCascio, who founded LivePerson and pioneered web chat for customer service, has spent the past two years building exactly this. His company, first called Eternos and now Uare.ai, raised \$10.3 million in seed funding in November 2025. Its “Human Life Model” is built only from an individual’s own data, combining their professional knowledge with their life story, values, and decision-making style. LoCascio’s design principle is ownership: “You own the model, and you can share it and monetize it.” The company also says its AI will not fill gaps with general large language model knowledge and will admit when it does not know an answer.² When Rob joined me on our podcast, we talked about cloning top performers, digital employees, and why data ownership must stay with the individual.³

Sam Huber approaches the same market from the enterprise side. Napster, which he represents in the Middle East and Africa, offers teams of avatar-fronted AI agents that companies onboard with their own data. Huber’s point is that roughly 90% of the value sits in a company’s first-party data, not in the underlying model, and that growth should come from doing more with more people and AI, not from cutting headcount.¹ In Dubai, Napster is now building digital twins of teachers for GEMS Education, which brings his work squarely into the territory of this paper.⁴

They are not alone. The market for digital humans of specific people has become crowded within two years:

Company	What it builds	Built from	Typical use
Uare.ai (formerly Eternos)	Personal AI of an individual, owned by that individual ²	Life story, values, professional knowledge, by text, voice, or video	Professionals, creators, top employees
Napster	Avatar-fronted AI agents and digital twins of teachers ⁴	Company data; teachers’ own course materials	Education, sales, support
Viven	A digital twin of every employee ⁵	Email, Slack, and documents, with pairwise privacy controls	Querying colleagues who are unavailable
Delphi	“Digital minds” of experts ⁶	Books, articles, podcasts, and posts by the expert	Coaching, advice, audience engagement
Twin1 AI	A digital twin for each professional ⁷	Work communication, with permission controls	Sharing expertise across an organization

Synthesia	Video avatars, stock or personal ⁸	1–5 minutes of footage plus live consent	Training and communication videos
-----------	---	--	-----------------------------------

Why the Elastic Organization Needs Its Best People Twice

In the white paper on Implementing the Future of Work in a Midsized OEM, we showed how demography is emptying the factory floors and engineering offices of Europe and North America: in Germany alone, 12.9 million economically active people reach retirement age by 2036, and Switzerland expects a shortfall of more than 460,000 workers within ten years.⁹ Hiring alone will not close that gap.

The elastic organization answers with two levers. The first is automation of tasks. The second, the subject of this paper, is the replication of *expertise*: making the knowledge of the best people available at the moment and place it is needed, even when the person is asleep, on another continent, or retired.

There is solid evidence that this works when the knowledge is captured well. Erik Brynjolfsson and colleagues found that an AI assistant raised the productivity of novice customer-support agents by 34%, precisely because it disseminated the tacit best practices of the most able workers.¹⁰ A digital human of a top performer is a more personal version of the same mechanism. A large life insurance company tested a replica of one of its top employees built on LoCascio's platform, and the pilot led to a larger deal to expand it across 20,000 life insurance agents in the field.¹¹

César Hidalgo reminds us why this matters: know-how is embodied in people, and a single person can hold only a limited amount of it, a “personbyte.”¹² An organization that loses its best people loses capabilities it cannot buy back on the market. A well-built digital human does not replace those people. It keeps part of their personbyte inside the company.

No AI Is the Full Person. It Is the Person on an Average Day.

This is the central claim of this paper, and every decision about digital humans should start from it.

What the Machine Actually Learns

A digital human learns from what a person leaves behind: documents, emails, transcripts, interviews, videos. That record overwhelmingly describes the person on ordinary days, explaining things calmly, after the fact, to an audience. The decisions that made the person exceptional are rare by definition. They happen under time pressure, in a crisis, in a negotiation, or in a moment of intuition the person could not fully explain afterwards. Michael Polanyi's observation that “we can know more than we can tell” applies to the training data of every digital human.¹³

The mathematics reinforces this. Language models generate the most plausible continuation of a pattern. Plausible means typical. When political scientists

compared synthetic survey answers from ChatGPT with real human responses, the synthetic answers were far less varied: in one example, a standard deviation of 16.1 against 31.4 among real respondents.¹⁴ A model built from a person inherits the same pull toward the middle of that person's distribution.

What the Best Research Shows

The most rigorous study so far confirms both the promise and the limit. Researchers at Stanford and Google DeepMind interviewed 1,052 people for about two hours each and built a generative agent of every participant. The agents matched participants' answers in the General Social Survey about 85% as well as the participants matched their own answers two weeks later, and reached about 80% correlation on Big Five personality traits. On behavioral economic games, however, where real choices with stakes are made, the correlation fell to about 66%.¹⁵

In other words: the closer we come to what people say about themselves, the better the replica. The closer we come to what people *do* when something is at stake, the weaker it gets.

The Replica Is Frozen in Time

A person keeps learning. A digital human, once trained, does not, unless it is deliberately updated. As we argued in our Responsible AI white paper, current AI systems cannot learn continually after training without the risk of forgetting what they already knew.¹⁶ A digital human of a retired engineer will still give the advice of 2026 in 2031, with the same confident voice.

Design Principles That Follow

- **Use digital humans for the repeatable, not the exceptional.** Onboarding, explanations, standard procedures, and frequently asked questions are average-day work. That is where replicas shine.
- **Teach the replica to say “I don't know.”** Uare.ai's refusal to fill gaps with general model knowledge is the right default for every enterprise deployment.²
- **Keep the original in the loop.** As long as the expert is available, they review what their replica says, especially where it is used for decisions.
- **Date-stamp the knowledge.** Every digital human should tell users how recent its knowledge is.

Why a Digital Warren Buffett Does Not Belong in the Boardroom

The most seductive idea in this market is also the most dangerous: put a digital replica of a legendary investor, a founder, or a retired chairman into the boardroom and let it advise on strategy. If a top salesperson can be cloned, why not Warren Buffett?

A Short History of Machines at the Board Table

The idea is not new. In 2014, the Hong Kong venture firm Deep Knowledge Ventures announced that it had appointed an algorithm called VITAL to its board.

Commentators widely dismissed it as a publicity stunt, the company later described its status as an observer without voting rights, and by 2019 it was reported to be no longer in use.¹⁷ In 2024, Abu Dhabi's International Holding Company gave an AI system called Aiden Insight a non-voting observer seat on its board, built by G42 with Microsoft to analyze business records, financial data, and market trends.¹⁸

The difference between a tool and an avatar matters. Aiden Insight is an analytical tool with a board seat. A digital Buffett would be something else: a replica of a person, presented as that person's judgment. That is the step we advise against.

Seven Reasons Against Digital Buffetts

1. **You get Buffett on an average day.** His public corpus consists largely of shareholder letters and interviews: carefully worded explanations written after the fact. The value of Buffett lies in the rare moments, such as the deals he struck at the height of the 2008 financial crisis, when he acted against the crowd. A model of the average of his words cannot reproduce the exception in his behavior.
2. **The public record is not the decision process.** What a person publishes about their reasoning is a curated story. A replica trained on it learns the story, not the reasoning.
3. **The replica is frozen.** Markets change, and a digital Buffett trained on past letters advises on today's questions with yesterday's knowledge.
4. **Nobody is accountable.** Corporate law requires directors to be real people. In Delaware, "each director shall be a natural person"; in Germany, only a natural person with full legal capacity can sit on a supervisory board.^{19,20} A replica cannot carry fiduciary duties, be liable, or be questioned by a court. If the board follows its advice, the responsibility is diffused into nobody.
5. **Consent is not optional, and the real person may object.** Berkshire Hathaway has publicly warned that AI-generated videos impersonating Buffett are spreading, calling the problem a "spreading virus" and stating that Buffett is "deeply concerned."²¹ Buffett himself called AI-enabled scamming potentially "the growth industry of all time."²² A digital Buffett built without his consent is an impersonation; built with it, it is still an average-day Buffett.
6. **Boards need dissent, not an echo.** Because models compress variance, a persona consulted by many directors pushes the board toward the plausible middle. The value of an independent director is the uncomfortable question no model would rate as the most probable one.
7. **A board-level avatar is a template for fraud.** In 2024, an Arup employee in Hong Kong transferred HK\$200 million (US\$25.6 million) after a video conference with people posing as senior officers of the company.²³ The same year, criminals used a voice clone and YouTube footage of WPP's CEO Mark Read in a Microsoft Teams meeting to try to extract money and information; the attempt failed thanks to vigilant staff.²⁴ Normalizing executive avatars in board processes makes such attacks more plausible, not less.

What to do instead: use AI in the boardroom as an analyst, not an oracle. Let it prepare scenarios, surface data, and check assumptions, with a named human accountable for every recommendation. And invest in the AI literacy of directors, so they understand what the tools in front of them can and cannot do.

Where Digital Humans Work: Positive Cases

The boardroom is the wrong place for replicas of people. Training rooms, help desks, and onboarding programs are the right ones. The following cases show what success looks like.

Organization	What was built	What it achieved or aims to achieve
Network Rail	AI avatars and voices for training, briefings, and safety videos (Synthesia)	Production time cut by about 95%; 500+ training videos in about six months across three departments ²⁵
Large U.S. life insurer	Replica of one top employee (Eternos, now Uare.ai)	Pilot expanded toward 20,000 field agents ¹¹
GEMS Education (Dubai)	Digital twins of teachers, trained on their own course materials (Napster)	Round-the-clock answers for students; feedback to teachers on where students struggle; student data kept in-country ⁴
Genpact, Eightfold	Digital twins of employees (Viven)	Colleagues can query an unavailable person's twin about shared projects ⁵
A 45-person HR research firm	A twin for every employee, tested for six months (Viven)	Faster access to expertise; knowledge retained when people leave ²⁶
Kaltura	Conversational digital twin of CEO Ron Yekutieli	Personalized updates to employees and customers in 30+ languages ²⁷
Zoom	AI avatar of CEO Eric Yuan in an earnings video	Openly disclosed as an avatar by Yuan himself ²⁸
Michael Bommer (Berlin)	Personal legacy AI (Eternos), built from 300 sentences of voice and 150 life stories	Conversations with his family in German and English, designed to continue after his death ²⁹

Three lessons connect these cases.

- **Narrow jobs win.** Network Rail did not build a digital managing director. It built training videos. GEMS did not replace teachers. It extended their reach between lessons, and Huber himself acknowledges that current models cannot replace the primary teacher-student relationship.⁴
- **Consent and control are built in, not added later.** Viven lets employees see every query to their twin, which discourages inappropriate questions.⁵ Delphi requires photo ID and does not allow users to create digital versions of other people.⁶ Synthesia requires a live consent recording from the person shown in the footage.⁸
- **Disclosure builds trust.** Eric Yuan introduced his own digital avatar to the Zoom platform. Klarna's CEO also presented results with an AI avatar, at a company whose CEO had earlier admitted that AI customer service delivered lower quality than humans.²⁸ The lesson: an avatar used transparently is a communication tool; an avatar used to suggest presence is a liability.

Who Owns It: A Chief Learning Officer Between the CHRO and the CTO

Digital humans fail for the same reason most AI projects fail: nobody owns the human side of the technology. MIT's Project NANDA found that 80-95% of enterprise generative AI pilots deliver no measurable return, and the main cause is organizational learning gaps, not model quality.³⁰ In our previous white paper, we described what happens when AI literacy is left to HR alone: purchasing decisions are made by people who have never studied AI or run an AI project, and training turns into a compliance circus.⁹ When it is left to IT alone, companies get platforms full of technology and empty of human knowledge.

The EU AI Act has made the question urgent. Since 2 February 2025, providers and deployers of AI systems must ensure a sufficient level of AI literacy among the staff who operate and use them, and supervision begins on 2 August 2026.³¹ An obligation without an owner becomes a certificate in a drawer.

The Role and Its Reporting Lines

We recommend a **Chief Learning Officer (CLO)** with two reporting lines. The concept was introduced for the first time in 2025 in a US-based OEM. The reporting lines were as follows:

- **To the CHRO**, because the CLO's work is about people: consent, dignity, career paths, incentives, works council agreements, and the retirement schedule that tells the company whose knowledge is at risk first.
- **To the CTO**, because knowledge only survives if it lives in systems: retrieval databases, knowledge graphs, avatar platforms, and their security. Without the CTO, learning programs float free of the company's systems. Without the CLO, the systems contain no human knowledge.

The CLO is responsible for two different things that are often confused. *AI literacy* is understanding what AI is, what it can and cannot do, and when not to use it. *AI fluency* is the practical ability to work with AI productively, including building the harness around a language model. Literacy comes first. Fluency without literacy produces confident users who cannot detect failures.¹⁶

The profile is demanding: technological fluency without being a technologist, a deep grasp of learning science, and a strategic understanding of the business.

Four Mandates

1. **Map and prioritize strategic knowledge.** A *Knowledge Risk Matrix* scores each expert on uniqueness, concentration, succession depth, and strategic leverage. An expert who scores high on all four is the first candidate for a digital human.
2. **Build the digital human program with the CTO.** The CLO owns knowledge elicitation, consent and governance, validation, and integration into training. The CTO owns platform selection, data pipelines, infrastructure, security, and integration with existing systems.

3. **Align strategy and training.** A quarterly capability roadmap translates strategy into capability needs, quantifies skill gaps, decides whether to build, buy, or borrow skills, and embeds learning into daily work.
4. **Protect knowledge sovereignty.** Capture knowledge before people leave, structure it for AI retrieval, govern shadow AI use, and keep it current. In our survey of SMEs, 67% reported shadow AI use, that is, employees using AI tools the company has not approved.

Guardrails and Rhythm

Three governance decisions belong to the CLO and cannot be delegated to the CTO: which decisions always require a human; how experts' consent and dignity are protected, including their right to review their replica and a sunset protocol when they leave; and the framing of the program as capability amplification, not workforce reduction. The operating rhythm is simple: a 30-minute weekly sync between CLO and CTO, a 90-minute monthly strategy alignment, and a half-day quarterly knowledge audit.

Months	Milestone
1–2	Knowledge Risk Matrix; platform evaluation with the CTO
3–4	First knowledge elicitation sessions; first internal prototype
5–6	Validation by the experts; deployment in internal training
7–8	Capability roadmap version 1; onboarding program
9–10	Governance framework; security audit
11–12	Strategic review; knowledge graph version 1; annual report to the board

Has Anyone Done This?

Chief Learning Officers are established in large companies, but they almost always report to HR alone. When Cognizant appointed a new Chief Learning Officer in September 2025, he reported to the Chief People Officer.³² The structural innovation we propose is the second line to technology. The nearest precedent is Moderna, which in 2025 merged its HR and technology functions under a single Chief People and Digital Technology Officer, Tracey Franklin, because decisions about which work is done by people and which by AI should not be split across the organization.³³ We have not found a documented case of a CLO formally reporting to both a CHRO and a CTO. Moderna shows that the logic behind it is already reaching the executive suite. We suggest a double reporting line in every project we are working on.

The Midsized Business: Too Small for a Chief Learning Officer?

Most midsized companies have no learning department. HR consists of one or two generalists who handle payroll, recruiting, and contracts. A full-time CLO with a technology background is out of reach, both in cost and in employer attractiveness. Yet these are exactly the companies that suffer most when an expert retires, because there is rarely a second person who knows the machine.

The Part-Time CLO

The answer is not to skip the role, but to share it. The model already exists for technology leadership. Researchers at TU Dresden and OTH Regensburg studied 62 *fractional CIOs* in ten countries: experienced IT leaders who serve several SMEs part-time, giving them access to senior expertise at lower cost and without long-term commitments.³⁴ Consultancies now offer *fractional Chief Learning Officers* on the same basis.³⁵ A part-time CLO working two days a month for a 200-person company can run the Knowledge Risk Matrix, the quarterly roadmap, and the governance framework, while the company's own people carry out the elicitation sessions.

Ecosystems: Sharing a CLO

The greatest leverage comes when several companies share one CLO, and one curriculum, through the ecosystem they already belong to:

- **Private equity firms** already share operating expertise across their portfolios. Vista Equity Partners' agentic AI hackathons in 2025 drew more than 300 participants from nearly half of its portfolio companies.³⁶ One private equity firm trained more than 2,000 employees in AI and extended reusable materials to its portfolio companies.³⁷ A CLO at the fund level, serving every portfolio company, is the logical next step: one Knowledge Risk Matrix method, one literacy curriculum, and one governance framework, applied to each company's own experts.
- **Industrial associations** already run academies. Swissmem Academy, for example, trains more than 1,200 industry specialists a year in around 70 courses and seminars.³⁸ An association-level CLO could add AI literacy, AI fluency, and knowledge capture to that offer for members who could never hire the role alone.
- **Chambers of commerce** are moving in the same direction. German chambers (IHKs) inform their members about the AI Act's literacy obligation, and offer certificate courses such as the "KI-Scout (IHK),: which trains apprentices to bring AI ideas into their companies.^{39,40}
- **Public networks** provide the infrastructure. Germany's Mittelstand-Digital network runs 29 centres with about 100 AI trainers who support SMEs free of charge.⁴¹ The federal labour ministry has funded 53 continuing-education networks (*Weiterbildungsverbünde*) in which SMEs pool resources for training on digitalization and skilled-worker shortages.⁴²
- **Employer groups** offer a legal template. In France, *groupements d'employeurs* allow several companies to employ people jointly and share them on a time-split basis, including executives.⁴³

Ecosystem	What exists today	What a shared CLO would add
Private equity firm	Operating partners, portfolio hackathons, firm-wide AI training ^{36, 37}	One CLO for all portfolio companies; knowledge capture as part of value creation
Industrial association	Member academies with technical courses ³⁸	AI literacy and fluency curricula; shared Knowledge Risk Matrix method

Chamber of commerce	AI Act information; certificate courses such as KI-Scout ^{39,40}	Part-time CLO pool for members; real literacy instead of compliance certificates
Public SME network	Free AI trainers; continuing-education networks ^{41,42}	Continuity: one accountable person per group of firms
Employer group	Legally shared, time-split employees ⁴³	A contractual vehicle for employing a shared CLO

Has a shared CLO ever been implemented? There is some partial implementation of building blocks: shared executives through employer groups and fractional leaders, shared training through public networks and associations, and portfolio-wide AI programs in private equity. We have not found a documented case of a Chief Learning Officer shared across several companies with a mandate for AI literacy, AI fluency, and knowledge capture. That is the gap this paper proposes to close.

A shared CLO needs a few rules to work: a steering committee of the participating CEOs; a transparent cost split; strict separation of each company's data and digital humans; a common curriculum with company-specific knowledge capture; and a common AI literacy assessment, so each company can compare where its people stand against the wider database before training begins.⁹

How to Build a Digital Human from Your Best People

The method builds on the capture approach we described for video-to-workflow in our previous paper.⁹ What changes is that the output is not an instruction, but a conversational representation of a person.

1. **Select for knowledge criticality, not charisma.** Start with people whose knowledge is scarce, frequently requested, and at risk of retirement, not with the most visible executives.
2. **Treat consent as a contract.** Define scope, purpose, compensation, the right to revoke, and what happens when the person leaves the company or dies. Decide upfront who owns the model. LoCascio's answer, that the individual does, deserves serious consideration.²
3. **Capture stories, not only facts.** Structured interviews about decisions, mistakes, and difficult customers teach the replica more than any manual. Bommer's legacy AI was built from 150 stories, not from a document archive.²⁹
4. **Write a job description for the replica.** List what it may answer, what it must refuse, and when it must hand over to a human. Never let a digital human authorize payments, sign documents, or approve exceptions.
5. **Let the original validate, and let skeptics test.** The expert reviews answers before go-live. Pair the expert with colleagues who distrust AI to probe its limits, as in our skeptic-tinkerer model.⁹
6. **Label it.** From 2 August 2026, the EU AI Act requires that people be informed when they interact with an AI system, and that deployers disclose deepfakes of real people.⁴⁴ Good practice goes further: every answer should identify itself as coming from the replica, with its knowledge date.

7. **Maintain and retire.** Schedule updates, version every release, and define when a replica is retired, for example when the knowledge is outdated or the person withdraws consent.
8. **Involve the works council early.** Digital humans of employees combine knowledge capture with personal data; in co-determination countries, purpose limitation must be agreed in writing before recording starts.⁹

What It Takes from IT: Networks, Storage, Backup, and Cybersecurity

A digital human looks like a chatbot with a face. Underneath, it is one of the most sensitive data assets a company can create: the voice, face, knowledge, and judgment of an identifiable person, combined with confidential company information. IT must design for that from the first day.

What Has to Be Stored

- **Raw capture:** video and audio of interviews, consent recordings, documents. Even a basic personal video avatar requires 1-5 minutes of continuous footage and a live consent video.⁸ A full knowledge-capture program for one expert can easily run to dozens of hours of video. As an illustrative estimate, at a typical HD bitrate of about 5 Mbit/s, one hour of video takes roughly 2 GB, so 20 hours of interviews mean about 40 GB per expert before copies and backups.
- **Derived data:** transcripts, curated knowledge bases, and vector embeddings used for retrieval.
- **Model artifacts:** voice clones, avatar models, fine-tuned weights, system prompts, and configuration.
- **Logs:** every question asked and every answer given, which are needed for quality control and for the person's own oversight, as Viven's query history shows.⁵

All four layers must be classified at the highest protection level. Voice and facial data can constitute biometric data under the GDPR (in case you are a European business, or serve customers in Europe), and processing biometric data to uniquely identify a person falls under the special categories of Article 9.⁴⁵ Embeddings are not anonymous either: researchers reconstructed 92% of texts exactly from their embeddings, and recovered 89% of full names from a clinical-notes dataset.⁴⁶

Networks and Residency

Real-time video avatars require low latency and stable bandwidth at the point of use: a training room, a classroom, or a shop floor tablet. Where avatars are delivered over video chat, as Napster's platform does, the network becomes part of the user experience.⁴⁷ Data residency matters just as much. In the GEMS project, student data stays in the country, hosted through G42 in the UAE and through Lenovo and Detasad in Saudi Arabia.⁴ European companies should expect works councils and data protection officers to ask the same question: where does my colleague's digital self live?

Backup and Recovery

A digital human that took months of an expert's time to build is not easily rebuilt, especially if the expert has retired. CISA's ransomware guidance applies directly: maintain offline, encrypted backups of critical data, maintain golden images of critical systems, and regularly test whether backups can actually be restored.⁴⁸

Backups create a paradox. If the person withdraws consent, their data must be deleted, including from backups. A practical answer is to encrypt each person's data with its own key, so that destroying the key makes every copy unreadable. Design this before the first backup is made, not after the first deletion request.

Cybersecurity

Digital humans inherit every known weakness of large language model applications. The OWASP Top 10 for LLM Applications 2025 lists, among others, prompt injection, sensitive information disclosure, data and model poisoning, system prompt leakage, vector and embedding weaknesses, and excessive agency.⁴⁹ For a digital human, each has a specific face:

Risk	What it looks like for a digital human	Control
Prompt injection	A user tricks the replica into revealing what the expert told it in confidence	Strict separation of public and private knowledge; output filtering; red-teaming before launch
Sensitive information disclosure	The replica repeats salaries, health details, or customer secrets it saw in emails	Curate training data; pairwise access rights as in Viven ⁵ ; regular audits
Embedding inversion	Stolen vector databases are turned back into readable text ⁴⁶	Encrypt embeddings at rest; restrict access like source data
Data and model poisoning	Manipulated documents change what the replica “believes”	Version and sign training data; approve every update
Impersonation and deepfakes	Criminals use the voice and face to stage fake meetings, as at Arup and WPP ^{23,24}	Never let avatars authorize actions; call-back verification for payments; watermark and label all output
Excessive agency	The replica sends emails or makes commitments on the person’s behalf	Read-only by default; human approval for every external action

Cybersecurity of digital humans is not only an IT question. It is a governance question: who may build a replica of whom, who may query it, and who is alerted when it is misused.

Legal and Ethical Guardrails

Law is catching up quickly. The EU AI Act's transparency obligations for chatbots and deepfakes apply from 2 August 2026.⁴⁴ The GDPR treats biometric identification data as a special category.⁴⁵ Denmark has proposed amending its copyright law to give citizens rights over their own face and voice and to ban sharing deepfakes of them without consent; the bill was expected to pass in 2026.⁵⁰ Companies should not wait for each jurisdiction to legislate. The principle is already clear: a person's likeness and knowledge are theirs, and a company may use them only with informed, revocable consent and for a defined purpose.

There is also an ethical line that no contract fully resolves. Employees may feel unable to refuse when their employer asks to clone them. Consent given under such pressure is weak. We recommend that participation be voluntary, compensated, and reversible, and that no one's career depends on agreeing to it.

Where This Approach Sits Relative to the Hype

Dimension	Hype	This approach
Scope	Clone anyone, do anything	Specific people, narrow jobs, clear refusals
Fidelity claim	"It is me"	It is me on an average day, as of a stated date
Boardroom	Digital Buffett as adviser	AI as analyst; humans accountable for every decision
Consent	Terms of service	Contract with scope, compensation, and revocation
Disclosure	Avatar suggests presence	Every interaction labeled as AI
IT	A chatbot project	Highest protection class: biometric, personal, and confidential data
Ownership	An IT project or an HR training budget	A Chief Learning Officer reporting to the CHRO and the CTO, shared where necessary
Success metric	Number of avatars	Knowledge reused, time to competence, incidents avoided

Conclusion

The elastic organization will not be built on generic agents alone. Its most valuable resource is the knowledge of its best people, and digital humans (or avatars) are a new way to multiply it. Used well, they let a retiring engineer keep teaching, a teacher answer questions at midnight, and a top salesperson coach twenty thousand colleagues.

Used carelessly, they create a dangerous illusion: that a person's judgment can be bottled and placed at the board table. It cannot. No AI is a reflection of the full person. It is a reflection of a person on an average day, and the moments that matter most in a company's life are rarely average.

Build digital humans from your best people for training, onboarding, and knowledge capture. Keep the originals in the loop as long as they are with you. Treat their data as the most sensitive asset you own. Give the capability an owner, a Chief Learning Officer between people and technology, and if you are too small to hire one, share one with your peers. And keep the boardroom human.

Sources

1. Lauterbach A., 88: *AI Workforce*, with Sam Huber, *AI Snacks with Romy & Roby* (podcast), August 18, 2026, <https://ai-snacks-with-romy-roby.cohostpodcasting.com/episodes/88-ai-workforce-with-sam-huber>.
2. Temkin M., *Immortality Startup Eternos Pivots to a Personal AI That Sounds Like You*, *TechCrunch*, November 11, 2025, <https://techcrunch.com/2025/11/11/immortality-startup-eternos-pivots-to-a-personal-ai-that-sounds-like-you/>.
3. Lauterbach A., 59: *Kiddie Chatbots, Clone Employees, Digital You: Language AI with Rob LoCascio*, *AI Snacks with Romy & Roby* (podcast), January 27, 2026, <https://ai-snacks-with-romy-roby.cohostpodcasting.com/episodes/59-kiddie-chatbots-clone-employees-digital-you-language-ai-with-rob-locascio>.
4. Schafer J., *Napster Returns as an AI Startup Cloning Teachers for Dubai Classrooms*, *AI Chat Daily*, September 18, 2026, <https://www.aichatdaily.com/ai-business/napster-returns-ai-startup-cloning-teachers-dubai-classrooms>.
5. Temkin M., *Eightfold Co-Founders Raise \$35M for Viven, an AI Digital Twin Startup for Querying Unavailable Coworkers*, *TechCrunch*, October 15, 2025, <https://techcrunch.com/2025/10/15/eightfold-co-founders-raise-35m-for-viven-an-ai-digital-twin-startup-for-querying-unavailable-coworkers>.
6. McCracken H., *Meet Delphi, the AI Startup That Lets Experts Turn Themselves into Chatbots*, *Fast Company*, June 24, 2025, <https://www.fastcompany.com/91356476/delphi-ai-digital-mind>.
7. *Twin1 AI*, news item, *Daily Research News Online (DRNO)*, August 24, 2026, <https://mrweb.com/drno/news40234.htm>.
8. Synthesia, *Create an Avatar*, *Synthesia Help Center*, accessed October 9, 2026, <https://docs.synthesia.io/docs/personal-avatars>.
9. Lauterbach A., *Pairing People and Machines: Implementing the Future of Work at an Incumbent Manufacturing OEM*, white paper, October 2026, <https://www.aiedutainment.ai>
10. Brynjolfsson E. et al., *Generative AI at Work*, *The Quarterly Journal of Economics* 140, no. 2 (2025): 889–942, <https://doi.org/10.1093/qje/qjae044>.
11. Smith A., *The Godfather of the AI Chatbot's Big Bet on Personal AI*, *AI Odyssey*, August 28, 2025, <https://aiodyssey.io/p/the-godfather-of-the-ai-chatbot-big-bet-on-personal-ai-9ccf64eb9802cb9c>.
12. Hidalgo C.A., *Why Information Grows: The Evolution of Order, from Atoms to Economies* (Basic Books, 2015); reviewed by Mark Buchanan, *Why Information Grows*, *Physics World*, September 22, 2015, <https://physicsworld.com/a/why-information-grows/>.
13. Polanyi M., *The Tacit Dimension* (University of Chicago Press, 1966; reissued 2009).

14. Bisbee J. et al., *Synthetic Replacements for Human Survey Data? The Perils of Large Language Models*, *Political Analysis* (2024), <https://doi.org/10.1017/pan.2024.5>.
15. Joon Sung Park et al., *Generative Agent Simulations of 1,000 People*, arXiv (submitted November 15, 2024), <https://doi.org/10.48550/arXiv.2411.10109>; Miller K., *AI Agents Simulate 1,052 Individuals' Personalities with Impressive Accuracy*, Stanford HAI, January 21, 2025, <https://hai.stanford.edu/news/ai-agents-simulate-1052-individuals-personalities-with-impressive-accuracy>.
16. Lauterbach A., *What Should Responsible AI Be About - If We Put the Noise Aside?*, white paper, May 2026, <https://www.aiedutainment.ai>
17. VITAL (Machine Learning Software), Wikipedia, accessed October 9, 2026, [https://en.wikipedia.org/wiki/VITAL_\(machine_learning_software\)](https://en.wikipedia.org/wiki/VITAL_(machine_learning_software)).
18. International Holding Company, *AI Advisor Takes Observer Seat on IHC's Board, the UAE's Largest Publicly Traded Entity*, press release, February 26, 2024, <https://ihcuae.com/media-centre/press-releases/ai-advisor-takes-observer-seat-on-ihcs-board-the-uaes-largest-publicly-traded-entity>.
19. Delaware General Corporation Law, Title 8, § 141(b), <https://delcode.delaware.gov/title8/c001/sc04/>.
20. Aktiengesetz (AktG) § 100(1), https://www.gesetze-im-internet.de/aktg/_100.html.
21. Ticong L., *'It's Not Me': Buffett Denounces Rise of AI Impersonation Scams*, eWeek, November 10, 2025, <https://www.eweek.com/news/ai-deepfakes-warren-buffett/>.
22. Pia Singh, *Warren Buffett Says AI Scamming Will Be the Next Big Growth Industry*, CNBC, May 4, 2024, republished by NBC10 Philadelphia, <https://www.nbcphiladelphia.com/news/business/money-report/warren-buffett-says-ai-scamming-will-be-the-next-big-growth-industry/3850065/>.
23. Gerrard N., *Arup Employee Falls Victim to US\$25 Million Deepfake Video Call Scam*, *Construction Briefing*, May 21, 2024, <https://www.constructionbriefing.com/news/arup-employee-falls-victim-to-us-25-million-deepfake-video-call-scam/8037338.article>.
24. West T., *WPP CEO's Identity Usurped by Deepfake Scammers*, *Marketing Beat*, May 13, 2024, <https://www.marketing-beat.co.uk/2024/05/13/mark-read-deepfake-wpp-ai/>.
25. Synthesia, *Network Rail*, customer case study, accessed October 9, 2026, <https://www.synthesia.io/case-studies/network-rail>.
26. Bersin J., *Arriving Now... The Digital Twin*, Josh Bersin (blog), 2025, <https://joshbersin.com/2025/10/arriving-now-the-digital-twin/>.
27. *Kaltura Puts the AI CEO into Production with Ron Yekutieli's Digital Twin, Using the Same Agentic Avatar Tech It Offers Enterprises*, Yahoo Finance, April 22, 2026, <https://finance.yahoo.com/sectors/technology/articles/kaltura-puts-ai-ceo-production-120000122.html>.
28. *The CEOs of Zoom and Klarna Have Presented Earnings Calls Using AI Avatars*, TechSpot, May 27, 2025, <https://www.techspot.com/community/topics/the-ceos-of-zoom-and-klarna-have-presented-earnings-calls-using-ai-avatars.292482/>.
29. Mocerino M., *Man Creates Digital Identity to Engage with Family After His Death*, *Interesting Engineering*, July 7, 2024, <https://interestingengineering.com/innovation/man-digital-identity-engage-family>.

30. Challapally A. et al., *The GenAI Divide: State of AI in Business 2025* (MIT Project NANDA, 2025),
https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
31. Smith T., *The EU AI Act's AI Literacy Requirement: Key Considerations*, July 29, 2025, <https://www.traverssmith.com/knowledge/knowledge-container/the-eu-ai-acts-ai-literacy-requirement-key-considerations/>.
32. Cognizant Appoints Thirumala Arohi Mamunooru as Chief Learning Officer, *The People's Board*, September 23, 2025,
<https://www.thepeoplesboard.com/news/cognizant-appoints-thirumala-arohi-mamunooru-as-chief-learning-officer/>.
33. Kabra N., *Moderna Merging Tech and HR Because AI and People Are the Same Thing? AI IQ* (Substack), May 24, 2025, discussing reporting by *The Wall Street Journal*, <https://aiiq.substack.com/p/moderna-merging-tech-and-hr-because>.
34. Kratzer S. et al., *Traction with Fraction: Strategic IS Management in SMEs through Fractional CIOs*, *International Journal of Information Systems and Project Management* 12, no. 1 (2024): 5–16,
<https://doi.org/10.12821/ijispm120101>.
35. *Fractional Chief Learning Officer: Consulting Services*, *Futurist Speakers*, accessed October 9, 2026, <https://www.futuristsspeakers.com/fractional-chief-learning-officer/>.
36. Vista Equity Partners, *Vista Equity Partners Names Winners of Agentic AI Hackathons*, press release, September 15, 2025,
<https://www.vistaequitypartners.com/news/vista-equity-partners-names-winners-of-agentic-ai-hackathons/>.
37. Ajaia, *Private Equity Workforce AI Upskilling Initiative*, case study, accessed October 9, 2026, <https://ajaia.ai/ai-insights/private-equity-workforce-ai-upskilling-initiative>.
38. Swissmem, *Training Programmes*, accessed October 9, 2026,
<https://www.swissmem.ch/en/added-value-for-your-company/services/training-programmes.html>.
39. IHK Hannover, *Verpflichtungen im AI-Act: KI-Kompetenzen im Unternehmen sichern*, accessed October 9, 2026,
<https://www.ihk.de/hannover/hauptnavigation/digitalisierung/aktuell-digitalisierung/ai-act-unternehmen-muessen-mitarbeitende-in-ki-schulen-6470624>.
40. IHK Köln, *KI-Scout (IHK)*, certificate course, accessed October 9, 2026,
<https://www.ihk.de/koeln/hauptnavigation/weiterbildung/zertifikatslehrgaenge/ki-scout-ihk--6173432>.
41. Federal Ministry for Economic Affairs and Climate Action (BMWK), *Die KI-Tour von Mittelstand-Digital*, press release, September 10, 2024,
<https://www.bundeswirtschaftsministerium.de/Redaktion/DE/Pressemitteilung/en/2024/09/20240910-die-ki-tour-von-mittelstand-digital.html>.
42. forum wbv / f-bb, press release on the 53 continuing-education networks funded by the Federal Ministry of Labour and Social Affairs (BMAS), January 10, 2023,
https://www.f-bb.de/fileadmin/user_upload/PM_Webseite_forum_wbv.pdf.
43. French Labour Code (Code du travail), Articles L1253-1 ff. on employer groups (*groupements d'employeurs*),

- https://www.legifrance.gouv.fr/codes/section_lc/LEGITEXT000006072050/LEGISCTA000006195698/.
44. Secure Privacy, *EU AI Act Article 50: Transparency Obligations for Chatbots and Deepfakes* (2026), July 9, 2026, <https://secureprivacy.ai/blog/eu-ai-act-article-50-transparency-obligations-for-chatbots-and-deepfakes-2026>.
 45. Regulation (EU) 2016/679 (General Data Protection Regulation), Articles 4(14) and 9, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
 46. Morris J.X. et al., *Text Embeddings Reveal (Almost) As Much As Text*, *Proceedings of EMNLP 2023*, <https://doi.org/10.48550/arXiv.2310.06816>; Marquart S., *Researchers Win Award for Study on Text Embedding Privacy Risks*, Cornell Tech, January 30, 2024, <https://tech.cornell.edu/news/researchers-win-award-for-study-on-text-embedding-privacy-risks>.
 47. *Napster Relaunches as Generative AI Company with Agentic Video Platform*, *Digital Commerce 360*, May 16, 2025, <https://www.digitalcommerce360.com/2025/05/16/napster-relaunches-generative-ai-company-agentic/>.
 48. Cybersecurity and Infrastructure Security Agency (CISA) and MS-ISAC, *#StopRansomware Guide* (September 2023), <https://www.cisa.gov/stopransomware/ransomware-guide>.
 49. Abbasi S., *AI Under the Microscope: What's Changed in the OWASP Top 10 for LLMs 2025*, *Qualys Blog*, November 25, 2024, <https://blog.qualys.com/vulnerabilities-threat-research/2024/11/25/ai-under-the-microscope-whats-changed-in-the-owasp-top-10-for-llms-2025>.
 50. Brooks J., *Denmark Eyes New Law to Protect Citizens from AI Deepfakes*, *Associated Press*, November 6, 2025, <https://www.news4jax.com/tech/2025/11/06/denmark-eyes-new-law-to-protect-citizens-from-ai-deepfakes/>.