

# Global Innovation

July 2025



## Innovation Insights from Japan & the US – Q2 2025



Storm Uru

Co-Head of the  
Liontrust Global Innovation Team



Clare Pleydell-Bouverie

Co-Head of the  
Liontrust Global Innovation Team

It has been an eventful second quarter for the Liontrust Global Innovation team. Over the past two months, we journeyed through Japan, Texas, and Silicon Valley, engaging with approximately 40 management teams across diverse industries. Despite traversing vast distances – both geographically and conceptually – every interaction we had distilled into three pivotal structural themes that we believe will dominate market dynamics over the next three years:

1. **Physical AI is in its ascendancy** – this encompasses robotics, which explains our visit to Japan, a nation that has almost single-handedly propelled robotics adoption over the past half-century. The central question we are addressing is whether Japanese incumbents can sustain their first-mover advantage amid rising competition.
2. **The AI infrastructure buildout** – which we regard as the fourth industrial revolution – is emerging as the most expansive infrastructure project in history. Sovereign AI represents an underappreciated demand driver: being designed-in is crucial, and in our view, traditional hyperscalers face inherent disadvantages.
3. **Value is moving up the tech stack** – this shift is now underway, accelerated by AI models' capacity for reasoning. Revenue opportunities abound for first movers, yet these largely involve a fresh cadre of companies, distinct from the software victors of the cloud SaaS era.

### Physical AI is in its ascendance – the opportunity set

If 2024 was the year of the agent, 2025 is poised to be the year of the robot. Three technological breakthroughs have converged to make this possible: Firstly, in 2024 we saw a major breakthrough with vision transformers: robots can now finally see. Just as large-language models predict next word in a sentence, these predict the next pixel, enabling full-fidelity simulation – from factory layouts for new cars to software-defined robotics.

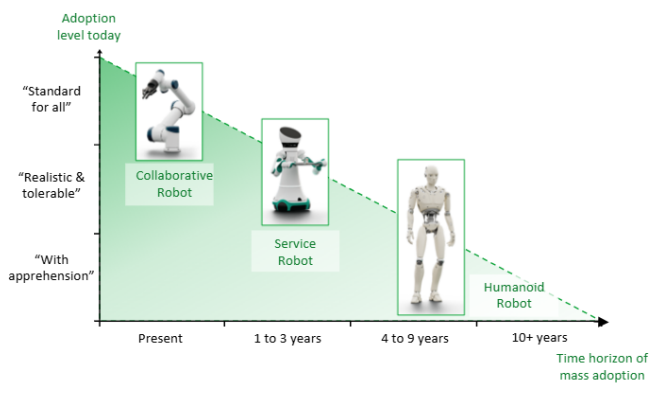
Synthetic data generation has been the next unlock: robots have to be trained on an inordinate amount of both real and synthetic data to perform a task at human-level quality. A surprise finding for us in Q4 2024 was that simulations trained solely on synthetic data (now incorporating structured and unstructured inputs) already match – and often exceed – the performance of models trained on real-world data.

Finally, is reinforcement learning – these models require extensive post-training. Whereas reinforcement learning once demanded armies of human data label-workers, it is now performed by a combination of AI and human supervisors – collapsing the iteration loop.

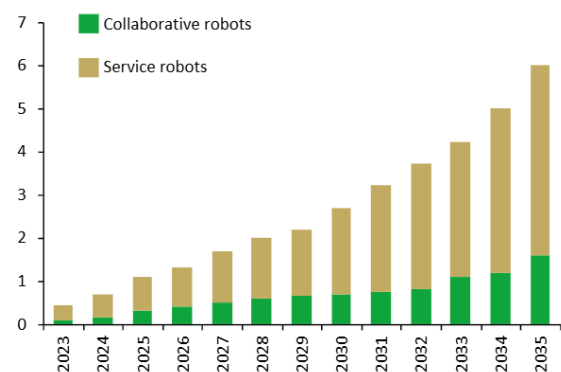
We believe that identifying the winners of the autonomy ecosystem today presents an opportunity analogous to spotting identifying the winners in the AI infrastructure ecosystem two to three years ago. Between now and 2035, more than six million robots are expected to be deployed globally, representing a 25% compound annual growth rate (CAGR) in the world's installed base of industrial robots.

Collaborative robots (cobots) are the immediate opportunity – robots that work with people. They are safe to stand up, easy to redeploy, and pay-back periods are now falling below 18 months. Over a one-to-three year horizon, service robots take centre stage. These robots actually perform a task for you, often in spaces that were never designed for automation – think warehouse tugs, hospital runners, retail shelf-scanners. Beyond that come humanoid robots. The technology is real, legged human scale machines are capable of dancing and even driving a forklift today, but the economics need another design cycle (likely two) before mass adoption.

Pace of robotic mass adoption



Addressable market - Over 6 million robots by 2035



Source: Melexis 2024. All use of company logos, images or trademarks in this presentation are for reference purposes only.

## Physical AI – who has the right to win?

Japanese incumbents enjoy a powerful first-mover advantage: robotic powerhouses Fanuc, Yaskawa and Kuka control roughly 36% of the world's industrial robot units, with Fanuc alone commanding 16%. However, low-cost Chinese competitors investing aggressively in innovation, and becoming increasingly visible. Incumbent management teams face the innovator's dilemma; those who acknowledge the threat are pivoting, those who do not will be disrupted.

What does pivoting the business look like? Firstly, ramping R&D spend to match Chinese counterparts. Fanuc, for example, has already accelerated its product-launch cadence 1.5x in the past 18 months alone. Management is under no illusion of the threat China poses to the dominant maker of yellow robots found on almost every factory floor.

Nidec, the world's number one motor manufacturer, is going one step further and re-shaping its portfolio to align with new vectors of demand. While its heritage lies in the auto sector, under a new CEO, Nidec is pivoting the business towards high-growth industrial drives, generators and battery-storage systems – critical back-up power for hyperscale data centres. The strategy is bearing fruit: order backlogs for these businesses now stretch beyond 2030, ROIC has climbed from single digits to 12%, and Nidec is also applying its motor expertise to supplying reduction gears and motors for humanoid robot prototypes.

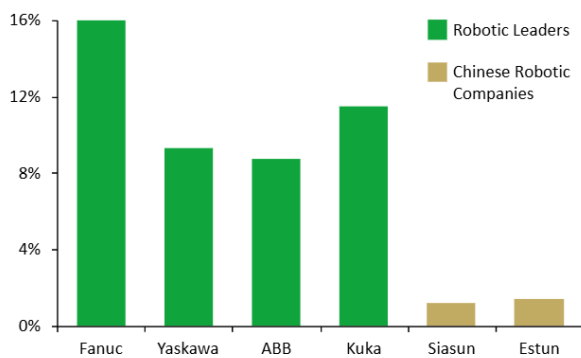
## Physical AI – unique insight – value accrues to different layers of the robotics stack

The resounding takeaway from our meetings across the entire ecosystem was that low end robotics is commoditising at an eye-watering pace. Value does remain in high-end, high-touch approaches (companies like Fanuc and Keyence). The middle will be squeezed.

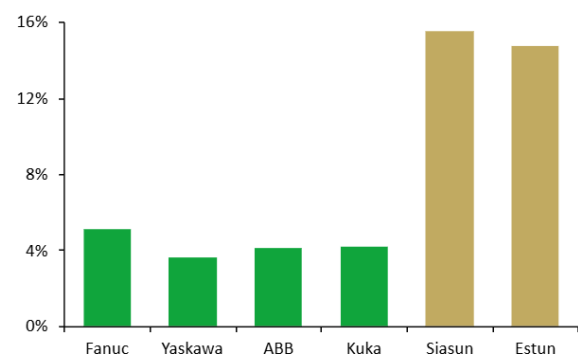
Ultimately though, value will accrue to different layers of the robotics stack in this next cycle – this is where Nvidia and Tesla, two companies outside the traditional robot league, come in with their full-stack approaches. Nvidia’s recently launched Project GROOT, a foundation model for humanoid robots baked into the Isaac platform, is essentially a ‘GPT-for-motion’. Whoever ships the best silicon and the best pretrained policy model will capture the software royalty layer across every original equipment manufacturer (OEM).

With its vision-only autonomy stack and neural networks, Tesla commands a formidable data advantage when it comes to autonomous driving – and an equally significant compute advantage when it comes to humanoid robots, given that the same stack underpins both domains.

Robotic industry market share



Investing in innovation (R&D/Sales)



Source: Company reports 2024. Liontrust 2025. All use of company logos, images or trademarks in this presentation are for reference purposes only

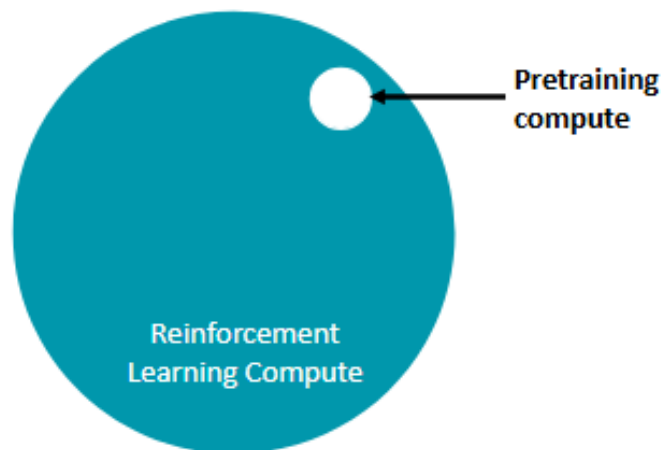
## The AI infrastructure buildout – the opportunity set

We have lost count of the number of times market participants have predicted ‘peak Nvidia earnings’. From our vantage point, this is only the beginning as AI rolls out across the economy. 2025 marks the transition of inference demand from experimentation to deployment; compute requirements are now accelerating exponentially.

Up until Q4 of 2024, there was one scaling law driving demand for AI compute – pre-training of AI models. Now, we have three: Pre-training, Post-training (reinforcement learning), and Inference-time reasoning.

Post-training and inference time reasoning demand multiples of the compute used in pre-training; the latter is 100-1000x more compute intensive than “one shot answer” AI chatbots. We face a supply-constrained demand environment for compute and have only just started moving into production – supply is unlikely to catch up to demand for at least three years, and potentially much longer if sovereign customers continue to sign projects that absorb 20-30 % of Nvidia’s annual graphics processing unit (GPU) capacity (as with Stargate).

Reinforcement learning is very compute intensive



Source: OpenAI, 2025

### The AI infrastructure buildout – who has the right to win?

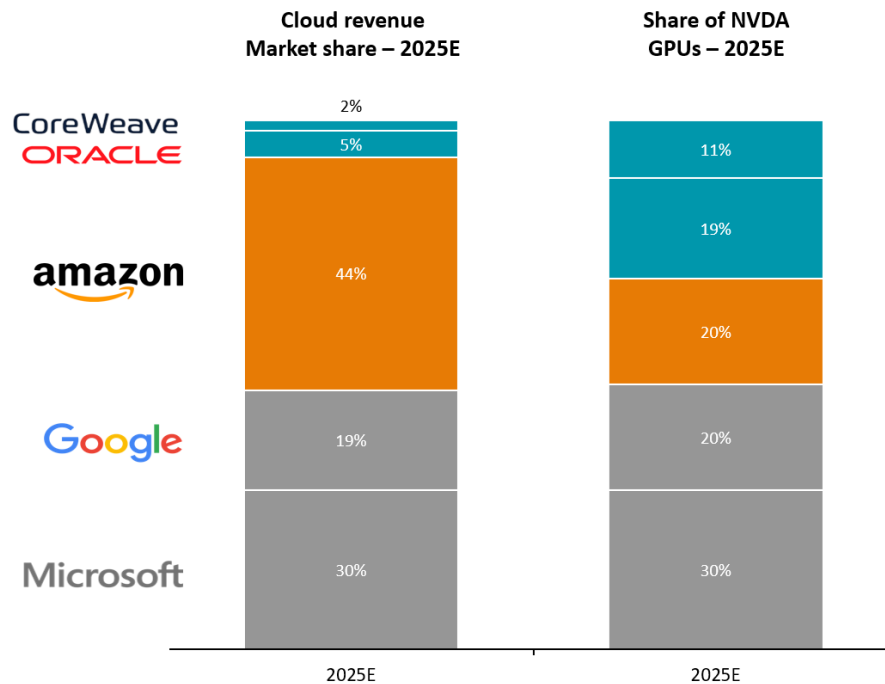
As frequent readers of our literature will be aware, we are of a firm belief that the winners of the next decade will be markedly different to the winners of the last decade. We do not hold Microsoft, Alphabet, Amazon or Apple across any of our three funds.

As with every technology platform shift, technology infrastructure demands are changing and a different cohort of cloud providers are taking share of AI workloads. Notably, Oracle, who holds just c.5% cloud revenue market share, commands c.19% of Nvidia GPU allocations, was selected to provide the accelerated-compute backbone for Stargate over Azure and AWS. CoreWeave is winning a similarly outsized share. Why? Because both firms designed their clouds specifically for accelerated compute: they are cheaper, faster and more reliable.

Oracle has engineered its Gen-2 data centres around massive, remote direct memory access (RDMA)-linked GPU clusters. Clustering breaks the old rule that a data-centre's performance is capped by the size of the individual server, and as such, delivers AI training that is 50-60 % cheaper, 30-40 % faster and materially more reliable than the incumbents. No other hyperscaler has managed to master this to-date.

CoreWeave took this concept even further. Every single rack in CoreWeave's data centres is optimised for GPUs – dense power, direct-liquid cooling, low-latency fabrics. The result? Up to 40x better performance and 40-65% cheaper than Azure and AWS.

## Estimated share of NVDA GPUs and share of Cloud Rev



Source: Company reports & Liontrust 2025.

What does this decoupling mean going forwards? Sovereign AI represents a \$1.5 trillion opportunity – the massive projects announced in the Middle East and Stargate are only the beginning of this build out. Being designed in matters and is in no one's numbers yet.

These sovereign programmes are a blank canvas: contracts are awarded purely on technology and price/performance, unhindered by legacy architecture. We therefore expect the legacy hyperscalers to cede further share to best-of-breed challengers.

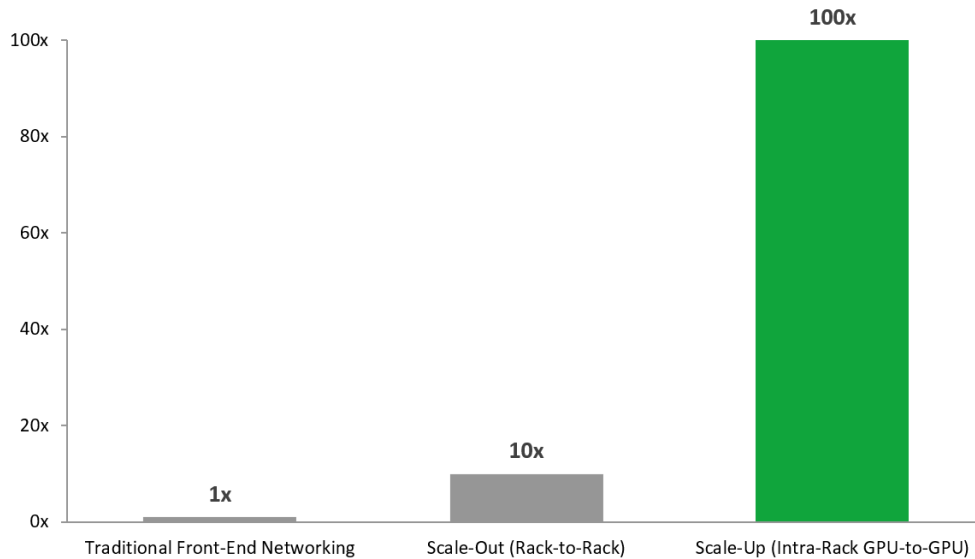
## The AI infrastructure buildout – unique insight – networking is the new battlefield

As compute demand scales and AI clusters move to 1 million GPUs, networking becomes the key battleground – the Nvidia ecosystem versus traditional standards. The opportunity for 'scale out' (rack-to-rack connectivity) is 10x that of traditional front-end networking in the data centre; the opportunity for 'scale up' (intra-rack connecting GPUs with GPUs) is 10x that of scale out. No wonder competition is heating up when we think about the size of the prize – new and growing moat – an underappreciated growth lever in our opinion which completes the company's full-stack compute offering. Already the de facto standard for AI compute and software, Nvidia is doubling down on becoming the de facto networking fabric for both scale up (with its NVLink solution) and scale out networking (with its Spectrum X solution).

Connectivity players are broadly divided into those that support Nvidia's closed ecosystem (although the company is beginning to open up its fabric architecture), and those who support traditional open architectures (traditional ethernet, for example). The likes of Amphenol, who provides the spine for NVLink in Blackwell compute architectures, are directionally linked to the success of the former; the likes of Credo, Astera Labs, Broadcom and Arista largely fall in the latter category.

Who will win? We invest in all of the companies hitherto mentioned. Ultimately, the addressable market for AI networking needs to be supported by multiple players given the scale of demand and the rate at which this demand is accelerating. Back-end networking for AI clusters is a completely fresh market set to compound over 30% through the end of the decade. As the number of GPUs in training clusters hit 1 million – likely before the street expects – we think these estimates could prove conservative.

## Scale-up versus traditional networking demand



Source: Broadcom 2024. All use of company logos, images or trademarks in this presentation are for reference purposes only

## Value moving up the stack – the opportunity set

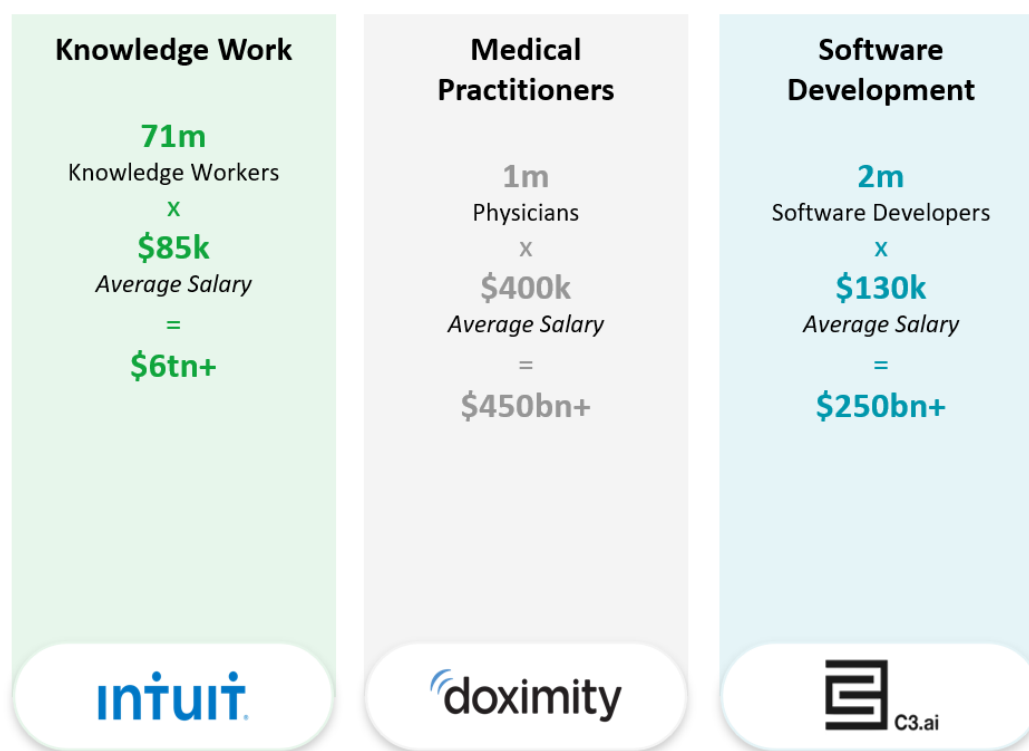
If there was one single most important observation from our trips, it was that inference demand has decisively shifted from experimentation to deployment. After months of debate over where the return on investment (ROI) on AI spend would appear, companies are now pushing their pilot schemes into production. The productivity improvements will not be monumental on day one, but – as with any compounding effect – a faint whisper can rise to a deafening roar.

Revenue-generating opportunities are materialising for first movers who are seeing an ROI on AI spend come through in a meaningful way. The mood music we encountered in Silicon Valley was akin to the exuberance at the start of the AI-infrastructure build-out two and a half years ago; the difference is the inclusion of the whole stack this time, not just the infrastructure layer.

Although not all AI deployments are created equal, the speed of diffusion across early use cases is astounding – even for the most optimistic supporters of the technology. The significant breakthroughs achieved last year in Q3 in inference-time reasoning and reinforcement learning are enabling use cases to broaden out and encourage significant adoption and investment across every industry. Just like electricity, AI will soon be table stakes for every enterprise.

The largest near-term opportunity for agentic AI is in software and services – an addressable market measured in the trillions, not the billions. We see the greatest potential in those companies targeting the services profit pool and, in doing so, creating entirely new markets.

## U.S. labour spend across key roles



Source: Company Reports 2025. Liontrust 2025. Sequoia 2025. All use of company logos, images or trademarks in this presentation are for reference purposes only

Take Dexterity, the 'LinkedIn for doctors' transforming into the 'digital office for doctors'. The company operates a double-sided platform serving doctors on one side (over 80% of US physicians are members), monetised through pharma marketing on the other. Dexterity GPT (its AI assistant) is proving capable of cutting administrative workflows for doctors in half, and as a result growing at a 5x clip year-over-year. This is a new market that did not exist before; Dexterity is targeting the \$450 billion labour pool of medical practitioners in the US, and it is doing so by driving down costs for its customers. If doctors can save hours a day, at, say, \$100 an hour, the ROI is significant and enables more patient visits.

## Value moving up the stack – who has the right to win?

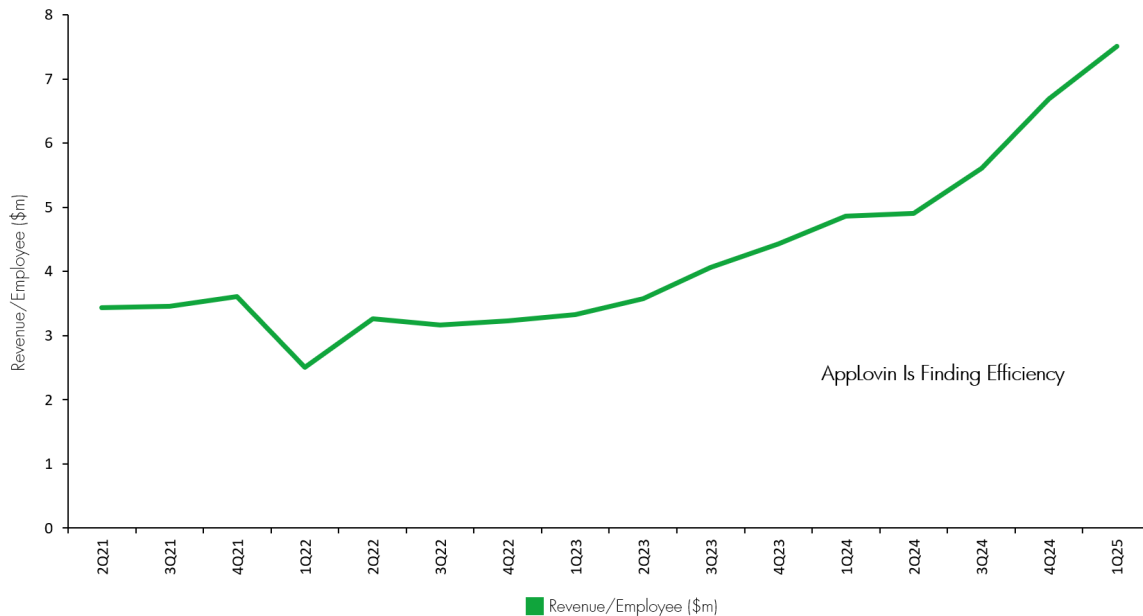
The generation of SaaS companies born in the era of the cloud are giving way to new growth companies pioneering the AI transition. These companies are built on accelerated computing, not traditional computing, and their business models look very different from those of the past – just as the early emerging leaders of the internet age, such as Google, Facebook, and Amazon, looked completely different from those of the personal computing age.

Agents themselves will be a commodity. Enterprise software companies sitting on a fraction of an organisation's data and unable to reach the rest are disadvantaged, as are those that operate as 'systems of record' where their customer base is in fact the owner of the data residing in their platforms. Value capture, we believe, accrues to two types of companies: those with a unique, proprietary, real-time data differentiation, or those serving as the backbone platforms for the orchestration of agents.

For the former, think of Intuit's QuickBooks accounting software platform, which incorporates real-time payments and book-keeping data for seven million online businesses, or Planet Labs' satellite imagery, which captures the entire world's surface every single day. Data is AI's fuel, but not all data assets are created equal. As for the latter, very few companies have the correct data architecture to serve as the orchestration platform for agents right across an organisation – C3.ai and Palantir are in a category of two.

The companies that can unlock value for themselves and their customers at scale will have two distinct characteristics very different from those of the past – one, higher relative revenue per employee compared to competitors; and two, a low marginal cost of growth. These traits are starting to reveal themselves in companies that have moved quickly to become first adopters of AI across different industries – these are often nimble enterprises with founder-operator management teams, not marketing heavyweights. Take Applovin: not only is its revenue-per-employee metric extraordinary (c.\$8 million), but the acceleration in these metrics over the past two years (from c.\$3.5 million in 2021) is breathtaking and creates significant shareholder value for investors.

### Applovin revenue per employee



Source: Company Reports & Liontrust 2025.

### Value moving up the stack – unique insight – agential cyber

As in the late 1990s hand-off from one tech cycle to the next, the winners of the previous cycle are becoming vulnerable to new competition. The once-robust barriers to entry incumbents have enjoyed across key verticals such as search, digital advertising, mission-critical software, and cloud computing, are coming to an end.

New winners are starting to reveal themselves, and they comprise a mixture of companies that missed out on the structural growth of the last cycle, such as Oracle; new companies formed over the past 5-10 years, such as OpenAI; or some companies like Shopify that succeeded in the last cycle but also have management teams willing to pivot the business quickly to adopt new technology.

The winners of the next decade are emerging



Source: Liontrust 2025. All use of company logos, images or trademarks in this presentation are for reference purposes only

A structural standout is cybersecurity, where demands tied to agential AI open up significant opportunities – every agent is a new identity, a new threat. However, not every cyber company, we believe, has the right technology, architecture, or end market to capitalise equally.

Traditional network- and firewall-centric models – where everything inside the network is implicitly trusted – collapse in an agent-driven world. Instead, zero-trust, cloud-based architectures connect users, devices, apps, or agents on a strict one-to-one policy basis.

This shift displaces the winners of the last cycle in the cyber market, such as Palo Alto Networks and Fortinet, and elevates three essential cyber pillars in this new era: cloud, identity, and endpoint. These pillars will define the cyber leaders of the next decade, in our view. First-mover advantage matters, as does having a data advantage at the platform-level to train security agents to hunt threats autonomously.

For cloud, that's Zscaler, which has the majority of the world's cloud security data and generates an astonishing 80% of the world's total security logs. For endpoint, that's CrowdStrike, which has the majority of the world's endpoint threat detection data. The company's agents are installed on millions of devices from which they crowdsource threat information to constantly have the latest threat protection – the result is approximately 700% faster threat detection versus peers. And for identity, that's SailPoint – the company manages identity entitlements across all applications in a company, which can reach up to 27,000, so they know exactly who or which agent should touch what.

---

## Key Risks

Past performance does not predict future returns. You may get back less than you originally invested. We recommend any fund is held long term (minimum period of 5 years). We recommend that you hold funds as part of a diversified portfolio of investments.

The Funds managed by the Global Innovation Team:

- May consider environmental, social and governance ("ESG") characteristics of issuers when selecting investments for the Funds.
- May hold overseas investments that may carry a higher currency risk. They are valued by reference to their local currency which may move up or down when compared to the currency of a Fund.
- May have a concentrated portfolio, i.e. hold a limited number of investments, or have significant sector or factor exposures. If one of these investments or sectors / factors fall in value this can have a greater impact on the Fund's value than if it held a larger number of investments across a more diversified portfolio.
- May encounter liquidity constraints from time to time. The spread between the price you buy and sell shares will reflect the less liquid nature of the underlying holdings.
- Do not guarantee a level of income.

The risks detailed above are reflective of the full range of Funds managed by the Global Innovation Team and not all of the risks listed are applicable to each individual Fund. For the risks associated with an individual Fund, please refer to its Key Investor Information Document (KIID)/PRIIP KID.

## Disclaimer

Non-UK individuals: This document is issued by Liontrust Europe S.A., a Luxembourg public limited company (société anonyme) incorporated on 14 October 2019 and authorised by and regulated as an investment firm in Luxembourg by the Commission de Surveillance du Secteur Financier ("CSSF") having its registered office at 18, Val Sainte Croix, L-1370 Luxembourg, Grand Duchy of Luxembourg and registered with the Luxembourg trade and companies register under number B.238295.

UK individuals: This document is issued by Liontrust Investment Partners LLP (2 Savoy Court, London WC2R 0EZ), authorised and regulated in the UK by the Financial Conduct Authority (FRN 518552) to undertake regulated investment business. It should not be construed as advice for investment in any product or security mentioned, an offer to buy or sell units/shares of Funds mentioned, or a solicitation to purchase securities in any company or investment product. Examples of stocks are provided for general information only to demonstrate our investment philosophy. The investment being promoted is for units in a fund, not directly in the underlying assets.

The document contains information and analysis that is believed to be accurate at the time of publication, but is subject to change without notice. Whilst care has been taken in compiling the content of this document, no representation or warranty is given, whether express or implied, by Liontrust as to its accuracy or completeness, including for external sources (which may have been used) which have not been verified.

**This is a marketing communication.** Before making an investment, you should read the relevant Prospectus and the Key Investor Information Document (KIID) and/or PRIIP/KID, which provide full product details including investment charges and risks. These documents can be obtained, free of charge, from [www.liontrust.co.uk](http://www.liontrust.co.uk) or direct from Liontrust. If you are not a professional investor please consult a regulated financial adviser regarding the suitability of such an investment for you and your personal circumstances. All use of company logos, images or trademarks in this document are for reference purposes only.