

## BrainChip – Annual General Meeting CEO and Chairman’s Address

---

**Sydney, 21 May 2024:** BrainChip Holdings Ltd (ASX:BRN), appends the address and presentation of the Chief Executive Officer and Chief Technology Officer to the 2024 Annual General Meeting, in accordance with ASX Listing Rule 3.13.3.

**This announcement is authorised for release by the BRN Board of Directors.**

---

### About BrainChip Holdings Ltd (ASX:BRN)

BrainChip is the worldwide leader in edge AI on-chip processing and learning. The Company's first-to-market neuromorphic processor, Akida™, mimics the human brain to analyze only essential sensor inputs at the point of acquisition, processing data with unparalleled efficiency, precision, and economy of energy. Keeping machine learning local to the chip, independent of the cloud, also dramatically reduces latency while improving privacy and data security. In enabling effective edge compute to be universally deployable across real world applications such as connected cars, consumer electronics, and industrial IoT, BrainChip is proving that on-chip AI, close to the sensor, is the future for its customers' products as well as the planet. Explore the benefits of Essential AI at [www.brainchip.com](http://www.brainchip.com).

Follow BrainChip on Twitter: [https://www.twitter.com/BrainChip\\_inc](https://www.twitter.com/BrainChip_inc)

Follow BrainChip on LinkedIn: <https://www.linkedin.com/company/7792006>

### Company Contact:

Kim Larkin

Company Secretary

Ph: +61 7 3010 9393

Email: [kim.larkin@boardroomlimited.com.au](mailto:kim.larkin@boardroomlimited.com.au)

## **Chief Executive Officer and Chief Technology Officer's Address**

### **Chief Executive Officer**

Hello and thank you for attending BrainChip's 2024 AGM. Joining me for this year's address is our Chief Technology Officer, Dr Tony Lewis. He and I are looking forward to this session to share an update on our company, the progress we have made since our last AGM, our plans and most importantly why I am very confident in BrainChip.

Let me acknowledge right up front that our revenue numbers are not there yet. The evaluation and design cycle are longer, deeper, and more complex than we anticipated as customers plan their strategic roadmaps. However, after hearing my prepared remarks, hopefully you too will share my conviction that we are a much stronger and a better positioned company and share my optimism in our near-term and long-term success.

But before I dive into what we are doing at BrainChip let me make some very brief comments about the market which I believe is important context.

#### **The Edge AI Market:**

Three key trends are having significant impact, accelerating demand and growth in the edge AI space.

The meteoric rise of generative AI has moved to the edge, as the cost of cloud computing for large language model (LLM) processing in data centres is inefficient and cost prohibitive. Industry leading chip manufacturers have silicon plans that enable GenAI and LLMs on devices without connection to the cloud. The push to the edge for LLM's is being fuelled by the need for privacy, reduced latency and to lessen the burden on data centres. All of these areas are where I contend BrainChip has market leading capabilities that we will not only maintain but extend.

Secondly applications and use cases for edge AI continue to grow. Natural language interaction with computers, once science fiction, is becoming commonplace. Denoising and automatic speech recognition at the endpoint or device are becoming increasingly sophisticated, requiring performant and efficient AI processing. Also, advances in medical devices enable remarkably sophisticated detection and analysis requiring distributed computing and new in-device semiconductor solutions. You will hear in a moment about our new neural network architecture we call Temporal Event based Neural Networks or TENNS and its quantifiable impact in areas such as audio denoising and vital signs analysis providing BrainChip a leadership position in the edge market. I firmly believe TENNS is an immensely important innovation and a critical differentiator for BrainChip. I look forward to the impact this technology will have on the market.

Finally, key vertical markets, including industrial – for surveillance, vision, GenAI predictive maintenance, automotive – for in cabin, ADAS, connected sensors, and general medical – for predictive diagnostics, continue to push the limits of edge-first computing. In these use cases the ability to scale up and down to exceptionally small configurations is a key requirement. Akida has the unique capability to scale down to less than single node configurations making it the ideal choice for this market.

I know BrainChip is positioned to be the market leader to support all these trends. You will hear from Tony in a moment that we are a market driven organization and are refining our offerings to capitalize on these trends.

**A year in review:**

Now let's look at the past year.

I'm sure you have seen our many announcements around new engagements, ecosystem expansions, and publications, we have announced a lot lately because a lot is happening. But let me share some of the critical activities we have undertaken that are perhaps less visible and that greatly increase the commercial muscle of BrainChip.

Last year, just prior to our AGM we announced our Generation 2 Akida IP platform. In my comments last year, I spoke of the importance of the Generation 2 platform as it allowed us to address many more use cases, attracting more customers, and improvements that make it much easier to adopt and implement. I must stress this was not an optional upgrade, we had to have it to win in this market. As I review our current pipeline, it is predominantly anchored on the Generation 2 offering. Over this past year, we spent thousands of engineering hours completing that platform and extending the original specification even further to make it even more robust. Improvements such as including a programmable activation function to future proof our platform for models in the coming years and robust support for PyTorch and ONNX runtime.

I want to express my gratitude for all the hard work put in by the employees of BrainChip to bring this critical offer forward.

Important and substantial advancements were made in evolving our patented TENNS algorithm. This is a huge differentiator for BrainChip. You can expect me to direct more investment into growing its commercial interest. The impact of TENNS on many use cases such as audio denoising, automated speech recognition, and LLM's is not to be understated. TENNS provides our customers with equal or better results with a fraction of the operations, parameters, and power on the inference side of AI and the ability to train on much smaller data sets on the training side. All these benefits equal tremendous cost savings for our prospects and customers. TENNS is a potential

For personal use only

replacement for the transformer, the foundation used by nearly every LLM today including OpenAI's GPT2-GPT4, Anthropic's Claude, Meta's LLAMA, and all Google LLMs. In our internal testing compared to GPT-2 (which has the same basic transformer architecture as GPT4) we're seeing 5,000x less computation required, and training is at least 10-60x faster for a comparable model. In my opinion this is the definition of a game changer and can be a catalyst for the adoption of BrainChip technology. We are currently engaged in critical discussions with several potential key licensees. based on the differentiated combination of TENNS plus our Generation 2 platform. Tony will elaborate more on TENNS during his address in a few minutes.

On the sales readiness front, last year I spoke of the geographical sales footprint we had developed. This year was all about providing our team with the capabilities needed to engage confidently and respond quickly.

- We have automated the ability to simulate a customer model's performance on Akida. What was once a manual process for a sales engagement which took weeks to complete, now takes less than a minute. This is something that sounds simple but required vast amounts of work and is critical to our success. The ability for customers to quickly evaluate model performance on Akida is foundational to adoption. This is something BrainChip will never be satisfied with and will undergo continuous improvement. It requires ongoing investment and improvements, something I will ensure happens.
- Investments were made in silicon design tools, like Synopsys that allow our engineers to provide precise specifications for area and speed for any process nodes, a foundational requirement for any engagement.
- We have developed a formal competitor benchmarking program, where we acquire competitor products, run models and measure results against Akida, a critical step in all sales engagements. This benchmarking program continues to give us confidence that Akida has a material competitive advantage in critical areas such as power used per inference and compute density.
- We have integrated Akida runtime into several microcontroller environments, demonstrating the ease of adoption and the immediate impact to the edge AI market. There are literally billions of microcontrollers in the world, the market potential is huge. Proactively integrating into their leading environments is an absolute necessity for BrainChip to bring its technology into this space. I am very encouraged by our progress here.

You have heard me say in the past, in an industry where smart people matter a lot, BrainChip continues to attract the best talent. This is an area where I spend many important cycles of my time as we continue to make your company stronger with every hire. Our engineering, CTO and Sales Engineering departments all made critical hires that have made them substantially deeper. BrainChip priorities technical talent, with our team comprised of over 80% engineers and over 15% PhD graduates from leading research institutions.

At the executive level, while we did recently lose a few executives which is to be expected in this highly competitive labour market we participate in, my recruiting efforts are focused on always attracting the absolute best talent and make each functional area even stronger. This year we saw world class executive talent join BrainChip with Dr Tony Lewis coming on as our CTO and Steve Thorne taking on the role as Vice President of Sales. Seeing the quality of candidates coming across my desk for the recently opened Chief Marketing Officer (**CMO**) role, I have no doubt that we will hire a world-class CMO for BrainChip very soon.

These are just a sample of the initiatives we have undertaken to make us a stronger organization anchored on an industry-leading product portfolio. I will come back to the podium in a moment to share our plans for 2024, but I would now like to ask Tony to come up and go deeper into our product plans for this and next year as well as his long-term technology vision for BrainChip.

Welcome Tony.

### **Chief Technology Officer**

Thank you, Sean. Good morning, valued shareholders. When Sean invited me to speak at this AGM, he specifically asked me to discuss three key topics:

- The uniqueness of our Generation 2 Akida IP;
- Our groundbreaking technology called TENNs; and
- Our future technology direction both near and far for BrainChip.

Before I get into these exciting topics let me share a little bit about myself. I joined BrainChip with extensive experience in neuromorphic engineering, artificial intelligence, and robotics. When BrainChip approached me, I immediately recognized the chance to be right in the middle of the next revolution in AI. That is the principal reason I joined. Over the last 6 months of my tenure, my belief in the future of BrainChip's technology has profoundly deepened. In the rest of my presentation, I will explain why.

### **Akida Architecture:**

After examining our Akida 2.0 IP, I was impressed by its efficient digital architecture. Akida 2.0 is designed with five key attributes:

1. Event-Based Computation, which conserves energy by computing only when necessary;
2. Akida uses near-memory compute, designed to minimize unnecessary data movement;
3. Akida 2.0 performs inference independently of any application processor. These three architectural features are key to low power performance;

4. All of this runs in the smallest area, reducing manufacturing costs for our customers; and finally
5. Ease of model simulation, development, and deployment.

### TENNs exponentially reduces energy use

TENNs is a unique BrainChip innovation that replaces the wildly popular transformer architecture used in LLMs like ChatGPT. Practically, this means a TENNs network can be built with fewer parameters and significantly reduced energy consumption, a fundamental benefit of the Akida architecture. In some cases, this reduction of energy is truly phenomenal as we will see in a moment.

TENNs is the Swiss Army Knife of neural networks. This versatile structure not only excels in denoising but can seamlessly transition to tasks like automatic speech recognition. This technology could transform your car into a smarter companion, enhance your home security system, or even revolutionise baby monitors. The possibilities for integrating such systems into the fabric of our daily lives are nearly boundless and very exciting.

Today, I'll also discuss its role in text generation, which underpins large language models. Imagine a compact device, similar in size to a hockey puck, that combines speech recognition, LLMs, and an intelligent agent capable of controlling your home's lighting, assisting with household repairs and much more—all without needing constant cloud connectivity or having to worry about privacy and security concerns, a major barrier to adoption particularly in an industrial setting.

Let's look at some practical applications of TENNs and where we are today:

- **Language Modeling:** Our TENNs architecture drastically outperforms traditional models like GPT 2.0 in terms of efficiency. Versus GPT-2, a predecessor to GPT-4, we see an astonishing 5000 times reductions in power. As you might know, power being used for AI systems is becoming an economic and ecological problem. TENNs may unlock the possibility of moving high performing LLMs to the edge. Let me show you a demo. On the left, is a TENNs network completing a sentence generation task. On the right, is a GPT-2 network. We can generate words at 1575 per minute. While GPT2- limps along at a mere 8 words per minute. Both networks had identical compute available to them.
- **Audio Denoising:** TENNs achieves near state-of-the-art results in audio denoising with dramatically lower energy consumption. Remarkably, TENNs can process raw audio signals directly, eliminating the need for the significant data transformations required by conventional neural networks. We can achieve excellent denoising on a standard benchmark data set but using 1/10 of the power. Let me show you a demo. Here we see an example of an audio clip with a noisy background. And here we will hear this same clip with noise

removed by our TENNs network. This kind of denoising is particularly important for use cases like low power hearing aids.

### **Future Technology Direction:**

We will Lower Energy of computation, by leveraging more aggressively ideas from the brain like sparse computation to achieve efficiencies that seem to be beyond the limits of physics. Neuromorphic is one of the only pathways to automatically, on the fly compute only what is needed. We will also leverage advanced mathematics to help us create compute engines with unprecedented energy per calculation ratios. What that means is that we will see powerful AI models running on battery power for weeks or months at a time. Some say that's not practical today.

We will achieve scale with multi-chip systems, to allow our customers to address the rapidly increasing needs of edge computing. So, we can eventually reach new markets like autonomous driving, smart robots, and if TENNs proves out, datacenters will be in sight.

We will broaden AI model support. New network architectures are introduced literally every day. We are committed to supporting advanced architecture of high value to our customers and knock down any remaining barrier to Akida adoption.

We will offer a Zero-Code programming option. We envision a web-programmable system that will remove the need for highly skilled engineers to use our hardware. Anyone from a clever high schooler to a medical doctor wanting to invent, say, a pocket device to detect skin cancer.

We will double down on TENNs because the combination of TENNs and Akida 2.0 will be a combination that is flat out unbeatable. The combination gives us differentiation that simply no one else can match today. And we must flesh out the use cases to allow our customers to see the future with us.

I have been putting intelligence into chips for over a decade. Even when I worked for huge cell phone chip companies, nothing has excited me more than the new possibilities of TENNs + Akida 2.0. Nothing makes me more excited to come to work and see the components of technology coming together so quickly. BrainChip is positioned to be at the centre of the next revolution in AI.

### **Chief Executive Officer**

Thank you, Tony. At this point you have heard a brief overview of the state of the market, a recap of some the important initiatives of this past year, and Tony just shared our technology plans.

So let me bring my remarks to a close in the last few minutes by sharing our core areas of focus for the remainder of this year and into next.

For personal use only

- First and foremost, I know we need to close license sales. We have several engagements that have been in evaluation for over 1 year that are near closing in on a decision. These engagements are with leading companies in audio, IOT, and microcontroller segments. We have put a tremendous amount of energy, handholding and resources into these relationships and evaluations as we lead the change from traditional architectures with big brand backing, to the future, a more performant offering. It is imperative we win more than our share. Based on direct prospect feedback I believe we are well positioned in some of the more critical engagements. In a number of cases, we have made it through the down selection process that eliminates many competitors and focuses on only 2 or 3 vendors for final selection. I have personally spent extensive time with C level executives in these accounts, an investment in time and resources that would not have occurred if they did not consider BrainChip a potential strategic partner. We now have the product, the ability to excel in detailed evaluations, the team and deal structure to win strategic and lucrative deals.
- BrainChip will invest in and extend the lead of our TENNS offering. You heard earlier from Tony about the promise TENNS represents. BrainChip will spend the remainder of this year refining the technical and commercial market push of TENNS.
- I will focus extensively on increasing our marketing efforts. This past year we built out a marketing engine. 2024 is a year of marketing acceleration as we focus on world-class product marketing, demand and lead generation activities to drive our sales funnel further.
- I will aggressively drive our team to complete our Generation 3 Product requirements for the next generation Akida which will focus on adaptability, programmability, usability and continue to extend our leadership in areas such as power, area, and performance. We must continue to invest in ourselves and innovate, this is how BrainChip remains at the forefront of the industry.
- Along with VVDN, our partner for the Edge Box, we will co-sell aggressively to increase the workloads on Akida. We will also judiciously explore other hardware vehicles to augment our long-term business model of IP licensing. These hardware offerings serve two primary purposes; they accelerate the adoption of Akida use cases which creates more demand, and also provide additional revenue streams as we build the IP business.

Let me bring my remarks to a close so we can take some questions. You can see the many advancements we have made not only in our product offering, but also in our go-to-market capabilities over this past year. You have heard of the substantial progress we have made on many key engagements and our unrelenting focus on doing everything possible to win. Tony shared his views on the unique advantages of our Akida architecture and how we will continue to enhance to meet future AI market demands.



Finally, I laid out a set of initiatives we are focused on in 2024 into 2025. All this work is supporting one simple goal; BrainChip winning business. I can assure you that I remain committed, and the entire company is dedicated, aligned, and working to make this happen, no matter how challenging the environment.

The board brought me on with a critical mission: to transform a promising technology into a successful product for a challenging market, while establishing the sales and marketing capabilities to drive its adoption.

We are doing the right things. I know we are on the right trajectory. When I pair all these actions with the emergence of a true edge AI market, I am more confident than ever that we are on the cusp of generating sustainable revenue streams. BrainChip has the momentum in place and potential to emerge as the undisputed leader in the exciting and continuously growing Edge AI market.

That is why you are invested in BrainChip. I and the entire board thank you.

- **END** -

For personal use only

# Annual General Meeting

OTCQX: BCHPY | ASX:BRN

Sean Hehir  
CEO

May 2024

# brainchip

Essential AI



# Forward Looking Statements

Certain views expressed here contain information derived from third parties or publicly available sources that have not been independently verified. This presentation includes certain statements, projections and estimates of the anticipated future financial performance of **BrainChip Holdings Ltd.** and the size, growth and nature of future markets for the company's products.

Such statements, projections and estimates reflect various assumptions made by the directors concerning anticipated results, which assumptions may or may not prove to be correct. **BrainChip Holdings Ltd.** and its subsidiaries have not sought independent verification of information in this presentation.

While the directors believe that they have reasonable grounds for each of the assumptions, statements, projections and estimates and all care has been taken in the preparation of this presentation, no warranty of representation, express or implied is given as to the accuracy, correctness, likelihood of achievement, or reasonableness of assumptions, estimates, statements and projections that are contained in this presentation. Such assumptions, estimates, statements and projections are intrinsically subject to significant uncertainties.

To the maximum extent allowed by law, none of **BrainChip Holdings Ltd.**, its directors, employees nor any other person accepts any liability arising out of any error, negligence or fault for any loss, without limitation, arising from the use of information contained in this presentation.

# Opening Comments...

ersonal use only



# The Edge AI Market

Generative AI at the Edge. New requirements and opportunity.

ersonal use only

# The Year in Review

Products enhanced. Customer focused.

ersonal use only

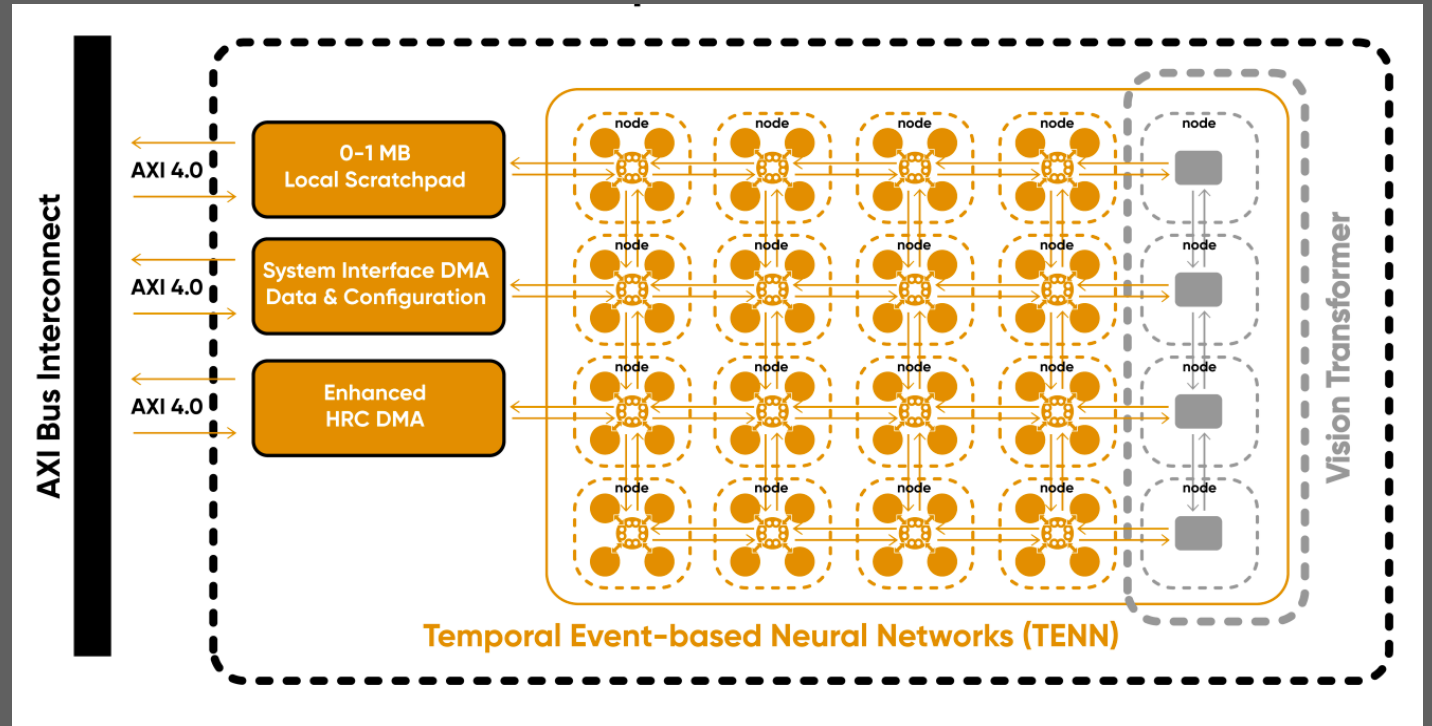
# CTOs Perspective

Dr. Tony Lewis

# Akida Architecture

## Key hardware differentiators


1. Event-Based Computation
2. Near-memory compute
3. Akida 2.x "independent execution"
4. Smallest footprint
5. Ease of model of deployment





# TENNs exponentially reduces energy use

- Replacement for many Transformer tasks
  - Language models
  - Time-series data
  - Spatiotemporal data
- Dramatically lowers energy requirements across all compute platforms



# TENNs vs GPT2

## Sentence Generation Task



TENNS  
> 1575 words/minute  
10.3 Performance Score



GPT2  
< 8 words/minute  
5700x MAC Energy  
10.2 Performance Score

Single thread CPU performance | 11th Gen Intel i7 - 3.00 GHz

# Audio Denoising

Personal use only



# Future Technology and beyond

- Lower energy of computation
- Scale up with chained chips to handle demanding LLM tasks
- Broaden Model support
- Zero-Code Environment

# The Next 12 Months

Executing, scaling...

Thank you

Q&A

ersonal use only

