



Welcome to the revolution

BrainChip Holdings (ASX: BRN) is an Australian, ASX-listed semiconductor company that is producing a revolutionary neuromorphic processor, called Akida. This chip is the first-ever digital computer chip that functions the same way the human brain does, i.e. it processes so-called spikes instead of conventional computer data.

Three solutions in one

The Akida processor can learn autonomously, even when the chip is already deployed in the field. It also improves the overall performance of artificial intelligence (AI)-dominated systems by substantially reducing latency and power consumption. Additionally, given that Akida can do most of its processing on-chip rather than rely on data centers, there's no need for continuous internet connections.

In other words, Akida helps solve three big problems, that today's AI solutions struggle with, all at once. Specifically, Akida needs no, or much lower, internet bandwidth, it uses substantially less power, and it takes Edge Computing to the next level through its much lower latency.

Wide range of application areas

With the advent of AI-driven autonomous devices, Akida has significant utility across a range of applications, including autonomous vehicles, internet-of-things (IoT) devices, drones, robotics, medical diagnostics, etc., making BRN a highly valuable proposition, in our view.

First Akida chips were delivered in August 2021

BRN, in collaboration with SocioNext, a leader in ASIC solutions, has secured commercial production for its Akida chip (AKD1000). This first production chip has been manufactured by TSMC in Taiwan and was delivered in August 2021. Following the commercial launch, we expect BRN, through its Early Access Program, to attract some of the biggest strategic names in the semiconductor industry as well as product end customers, like electronics companies. We expect this will likely be a major catalyst for BRN's share price.

Valuation of A\$ 1.50 per share

Similar to the way we valued BRN at A\$0.43 per share in our previous report from May 2019, we are using recent M&A valuations that have been paid for industry peers to derive a value for BRN shares. Based on the US\$2BN that Intel paid for Habana Labs in 2019, we value BRN at A\$1.50 per share fully diluted.

Please see page 27 for risks associated with BrainChip.

Share Price: A\$0.495

ASX: BRN


Sector: Software & Services

20 August 2021

Market Cap. (A\$ m)	821.6
# shares outstanding (m)	1,659.7
# shares fully diluted (m)	1,808.6
Market Cap Ful. Dil. (A\$ m)	895.3
Free Float	83.6%
52-week high/low (A\$)	0.97-0.10
Avg. 12M daily volume ('1000)	17,941
Website	https://brainchipinc.com/

Source: Company, Pitt Street Research

Share price (A\$) and avg. daily volume (k, r.h.s.)



Source: Refinitiv Eikon, Pitt Street Research

Valuation using M&A transactions for peers	
Value per share (A\$)	1.64
Value per share fully diluted (A\$)	1.50

Source: Pitt Street Research

Analyst: Marc Kennis

Tel: +61 (0)4 3483 8134

marc.kennis@pittstreetresearch.com



Table of Contents

Introducing BrainChip Holdings (ASX: BRN)	3
<i>Enhancing Edge Computing through neuromorphic processors.....</i>	<i>3</i>
<i>First production chips have just been received.....</i>	<i>3</i>
<i>Early wins.....</i>	<i>4</i>
<i>Near term share price catalysts</i>	<i>4</i>
Akida – The Artificial Intelligence Revolution	5
<i>(Far) Edge computing is Akida's sweet spot</i>	<i>6</i>
<i>Three key issues with Edge Computing that Akida solves.....</i>	<i>7</i>
Akida – Biologically Inspired, Digitally Engineered	10
<i>On-chip learning with Akida.....</i>	<i>10</i>
<i>Akida way ahead of competitors' solutions</i>	<i>12</i>
Paths to revenue clearly visible	13
<i>A very wide range of application areas for Akida.....</i>	<i>14</i>
<i>Akida to generate four separate revenue streams.....</i>	<i>17</i>
<i>Volume production of Akida commences</i>	<i>19</i>
<i>Secured production slots with TSMC.....</i>	<i>19</i>
End markets for Neuromorphic computing are vast and diverse	20
Valuation of A\$1.50 per share	25
<i>The proof of the pudding....</i>	<i>25</i>
<i>What could a post-revenue BRN be worth?.....</i>	<i>25</i>
<i>Value of A\$1.50 per share fully diluted.....</i>	<i>26</i>
<i>Funding and capital structure.....</i>	<i>26</i>
<i>Risks.....</i>	<i>27</i>
Summary and Conclusion	28
Appendix I – Introduction to AI, machine learning and deep learning	29
Appendix II – Neuromorphic computing 101	34
Appendix III – BrainChip's Solution for Neuromorphic Computing	37
Appendix IV – Analyst Qualifications	44
General advice warning, Disclaimer & Disclosures	45



Akida's neuromorphic processors power AI-architecture with improved latency, adaptive processing and rapid learning

BrainChip is the only company in the world to have combined neuromorphic processing with event-based convolution

As connected devices are expected to accelerate at a CAGR of 13.8% over 2020–2025 to reach c.41bn, the usability of neuromorphic processors expands accordingly

Introducing BrainChip Holdings (ASX: BRN)

BrainChip Holdings (ASX: BRN), an ASX-listed semiconductor company, is in the early stage of commercialisation of its Neuromorphic System-on-a-Chip called Akida. Akida is a revolutionary technology for neural networking applications. The chip provides high-performance neuromorphic computing while consuming very little energy, solving one of the biggest issues with today's Artificial Intelligence (AI) systems, namely very high-power consumption. Akida can perform event-based convolution as well, enabling the chip to run both Convolutional Neural Networks (CNN, see Appendix I) and Spiking Neural Networks (SNN). BrainChip is the only company to have combined neuromorphic processing with event-based convolution.

The Akida technology is well suited as an individual accelerator or as a co-processor, for instance within advanced driver assistance system (ADAS), autonomous vehicles (AV), audio sensors, the industrial internet-of-things (IoT), drones, vision-guided robotics, surveillance, medical diagnostics, and machine vision systems.

Enhancing Edge Computing through neuromorphic processors

As opposed to most chips used in AI applications today, which are general-purpose chips, such as Central Processing Units that need regular communications with the data center, Akida is an event-based processor, similar to the biological brain (see appendices for an in-depth explanation of BRN's technology). Like the brain, it consumes very little electrical power. As such, Edge Computing is likely to be Akida's largest market opportunity, in our view, because most Edge devices rely heavily on sensors that capture, and process real-world data and power sources are limited. Most of the chips used in Edge applications today are power-hungry general-purpose chips, such as Graphics Processing Units (GPUs) that consume approximately 1,000x more power than the Akida chip. Or they are low-power Digital Signal Processors (DSP's) that use a similar amount of power to Akida, but with about 1/60th the performance of Akida. And these DSP's have very high latency, so they can only perform relatively simple tasks.

Akida's low power consumption, reduced need for internet bandwidth and improved latency provide faster response times within a system, making Akida an ideal solution for Edge devices within the Internet-of-Things (IoT), such as sensors in autonomous vehicles.

BRN plans to licence the Akida IP (intellectual property) to customers generating revenues from multiple sources: 1) non-recurring engineering (NRE) work related to specific integration needs; 2) one-off licencing fees; 3) recurring royalties from each chip sold by the company's customers; and 4) revenues from chip and module sales.

First production chips have just been received

BRN has collaborated with SocioNext, a global developer of advanced ASIC (Application Specific Integrated Circuits) solutions, for product development and manufacturing of Akida. SocioNext and BRN have secured commercial production at TSMC (Taiwan Semiconductor Manufacturing Company) in Taiwan with the first batch of production chips received in August 2021. These chips are currently being tested. Once this Quality Control process is completed, Akida should be commercially available in October/November 2021.



Early wins

Although BRN has only just received its first production chips, the company has had early wins in the run-up to the commercialisation phase that has now started:

- First license agreement with a major ASIC manufacturer announced in December 2020. This license was paid at the beginning of 2021 and royalties are expected in 2023.
- About 15 EAP customers under the 'Early Access Program' (EAP) have received and paid for Akida engineering samples in 2020. These include major car manufacturers and NASA. It is likely some of these will convert to IP licenses and others to chip sales.
- The company has also signed a proof-of-concept agreement with a supplier of electronic components to the Automotive industry. This a project to deploy the Akida chip in LiDAR 3D object recognition.
- BRN has signed a research project with Biotome for the identification of COVID-19 antibodies in blood samples using a sensor and Akida. Biotome has paid for this research project. The aim is to develop a commercial medical device for antibody classification for a range of diseases.

Near term share price catalysts

With the first Akida chips becoming available after quality control testing, projected to be in October, BRN is emerging from the development phase and has moved into its commercialisation phase. We believe that demonstration of Akida's commercial viability, i.e. when the first commercial chip sales are announced, will be the main catalyst for BRN's share price.

Investors have supported the company to the point where we believe further share price gains will require commercial agreements leading to meaningful revenues. Given that we are very close to that point, we see substantial upside potential for BRN stock.

We value BRN at A\$1.50 per share fully diluted.



Akida – The Artificial Intelligence Revolution

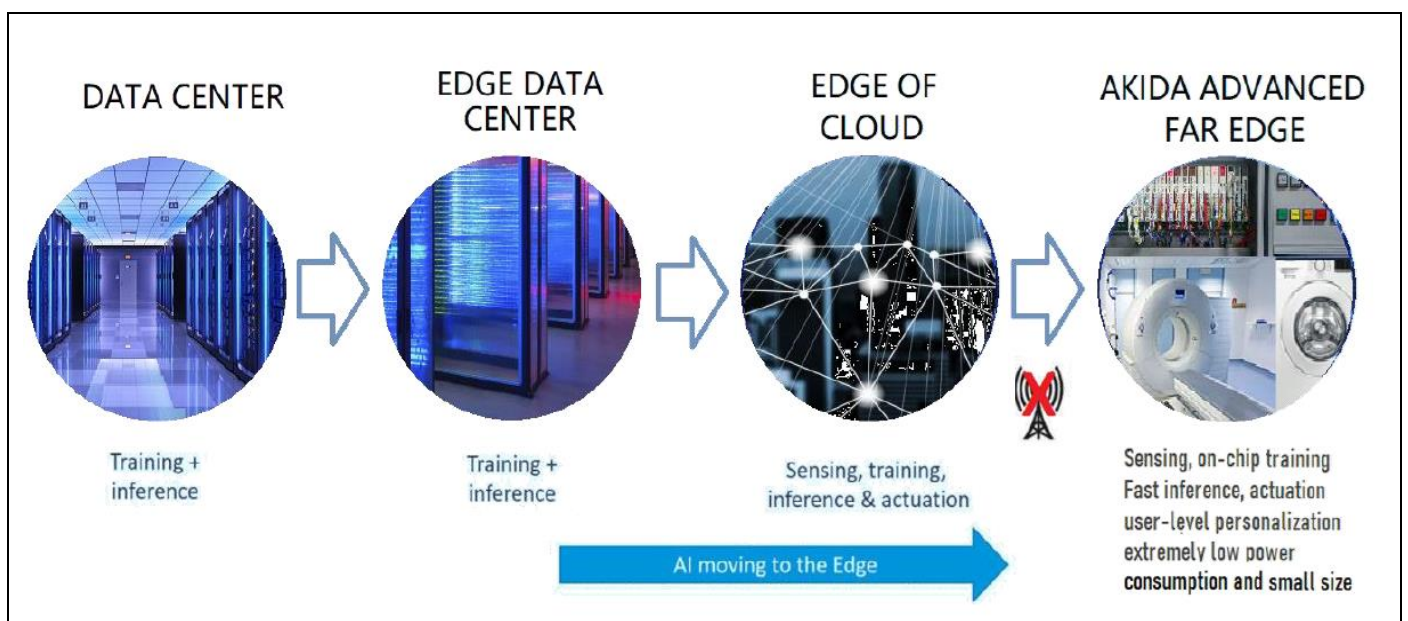
BRN is in the final stages of commercialising Akida, which is the Greek word for spike. The architecture and behaviour of this disruptive, AI-based computer chip are similar to that of a biological neuron, like the ones found in the human brain. This neuromorphic processor is expected to unlock the full potential of edge computing through Akida's 8 key differentiators that will be addressed in this report:

- On-chip learning
- On-chip event-based convolution
- Very low power consumption
- Very low latency
- Akida is stand-alone, i.e. it performs inference independent of the CPU
- Available to customers as IP (Intellectual Property) through IP licenses
- No internet connection is required
- Improved security because no images are uploaded for processing

The company has developed a neuromorphic processor, Akida, which aims to unlock the full potential of edge computing

Bringing AI to the Edge. 'Edge' is a buzzword in the technology industry, which typically means geographical distribution. Simply put, 'Edge' is where the action happens and the 'Edge devices', which include IoT controllers, smart sensors, mobile phones, laptops, medical monitors and other devices, are at the edge of the Internet or cloud where data is acquired. Likewise, 'Edge' in AI means that the AI algorithm runs on the device itself and not on a cloud-based server/data center. In the simplest terms, the sensor/IoT device is made intelligent so that it can perform data analytics locally, on the device. Akida enables data processing at the **far edge** (Figure 1) — at the periphery of the network — where sensor inputs can be analysed at the source of generation without an internet connection.

Figure 1: AI moving to the Edge



Source: Company



Notably, Akida is a high-performance, small chip, which consumes ultra-low power. Other features include on-chip training, learning and inference. Typically, inference refers to the process of using a trained machine learning algorithm to make a prediction.

Traditionally, edge devices send data to the cloud through the Internet so that it gets processed by a data center. As per the estimate of the International Data Corporation (IDC), IoT devices are expected to generate over 90 trillion gigabytes of data by 2025. Sending this huge amount of data to centralised data centers will lead to issues such as insufficient internet bandwidth and very high power consumption (416 TW in global power consumption annually in 2016). Notably, the AKD1000 chip comfortably addresses these issues, given that most data will be processed on-chip at the edge.

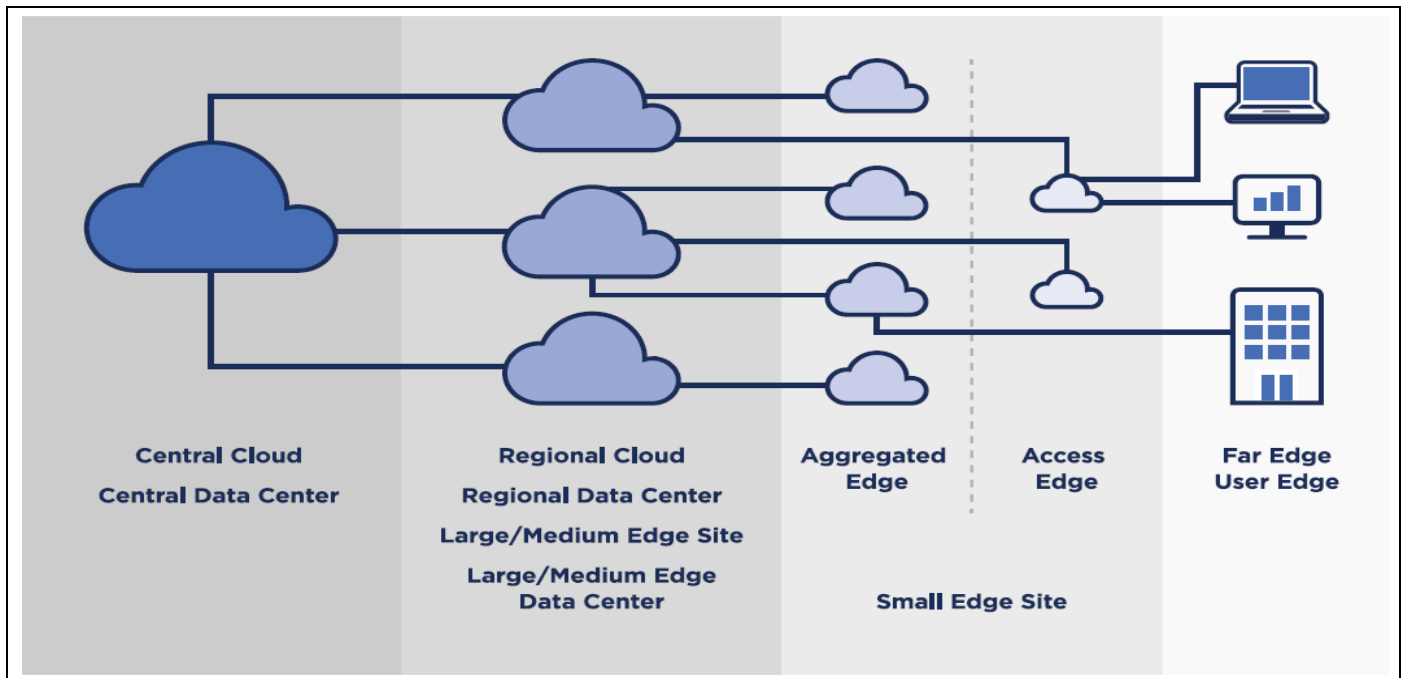
(Far) Edge computing is Akida’s sweet spot

Edge computing refers to computer processing that takes place at a location closest to the end users, and furthest from the data center. In an industrial setting, Far Edge refers to devices, such as sensors, actuators and dataloggers, which interact with the real world (Figure 2). These devices could be connected to a local/client/edge server, which is often referred to as the ‘Edge’.

Similarly, in a residential setting, Far Edge refers to devices such as smart TVs, smart thermostats, smart speakers and smart lights (Figure 2). In this case, Edge represents the wireless router, which helps connect these devices to the Internet.

Processing that takes place closest to the end user

Figure 2: Putting Edge and Far Edge in context of the infrastructure



Source: 'The Road to Open Edge Computing', Pipeline Publishing (September 2020)



The many advantages of (Far) Edge computing

Carrying out computation at the Far Edge eliminates the need for data to be transferred to, and processed by data centers. This not only helps reduce latency related to transferring of data, but also helps lower the burden on the data centers. There is a rising focus on energy consumption of data centers and their carbon footprint. In 2016, data centers globally consumed ~416 terawatts of power, ~3% of the world’s electricity. Notably, the energy consumption by data centers was 30% more than the energy consumed by the UK.

Moreover, current AI models also require vast amounts of energy. For instance, a project that involved a robotic hand trained to manipulate a Rubik’s Cube was estimated to require around 2.8 GWh of electricity — enough to power hundreds of homes for a year¹. Neuromorphic technologies can potentially provide significant power savings for these types of applications.

Far Edge computing offers several benefits

Three key issues with Edge Computing that Akida solves

Internet bandwidth

According to estimates by IoT Analytics, the number of IoT devices is expected to witness a three-fold rise during 2019–2025E. Thus, insufficient internet bandwidth is predicted to result in serious inefficiencies. With 30.9bn (Figure 3) IoT devices and 10.3bn computers connected to the Internet over the next four years, there will be a rising demand for internet bandwidth leading to delays and capacity constraints for IoT devices. However, Akida will help to reduce internet bandwidth constraints as data stays on the device. Further, certain applications that Akida facilitates, such as remote sensing, do not require an ongoing internet connection at all. They work independently of the cloud, performing the analytics on or near the device, the point of data acquisition.

The number of IoT devices is expected to witness a three-fold jump during 2019–2025E

Figure 3: Total number of connected devices (in billion) are expected to accelerate rapidly

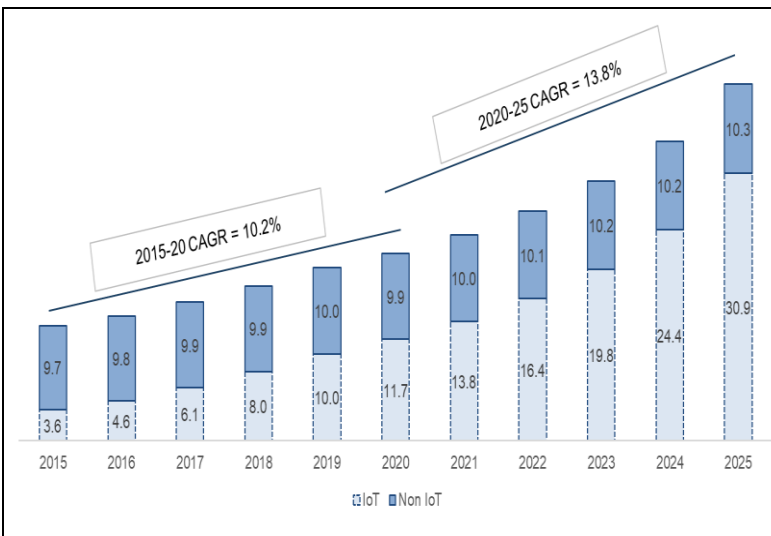
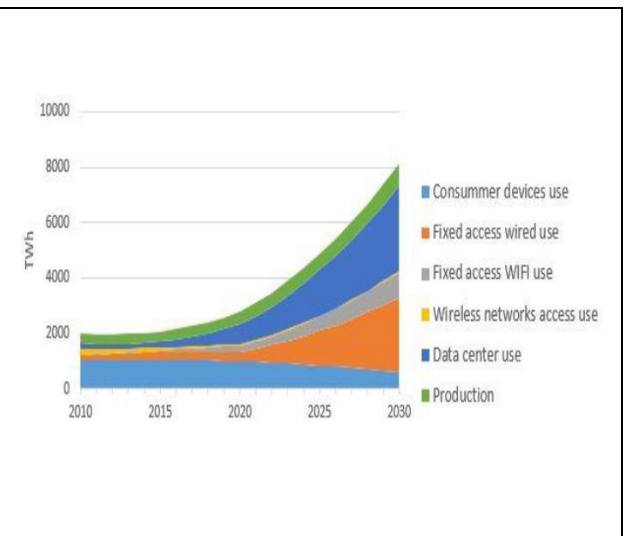


Figure 4: World electrical energy usage, projected from 2010 to 2030



Source: Source: ‘Cellular IoT connectivity & LPWA Market Tracker 2010-2025’, IoT Analytics (Nov 2020); IoT Analytics, Andrae

¹ ‘Driving intelligence at the edge with Neuromorphic Computing’, Accenture (2021)


Akida is ~97% more energy efficient compared with processing the same task on a central data center

Power consumption

The proliferation of data centers is leading to very strong growth in global energy consumption, which directly contributes to climate change. Energy consumption by data centers is on the rise (Figure 4) and by 2030, it is estimated that data centers will account for ~20% of all electricity consumption globally, vs. 3% in 2016. This substantial growth in energy consumption during the next 10 years will lead to a jump in carbon emissions. Notably, a single instance on Amazon Web Services (AWS) can consume 20 Wh. The AKD1000 chip, on the other hand, has ultra-low power consumption, making embedded AI applications at the edge possible, in the range of 100 microwatt to 300 milliwatt depending on workload (Figure 5). Thus, the AKD1000 is ~97% more energy efficient when compared to processing the same task in a data center.

We believe Akida is a green technology as it reduces a key driver of climate change, i.e., carbon emissions due to growth in energy consumption. Akida should be capable of diminishing energy demand as it enables edge computing, thereby drawing away compute requirements away from data centers.

Figure 5: Akida’s ultra-low power consumption range



Source: Company

Unlike a traditional processor that processes data, Akida processes spikes, or events

Computing power

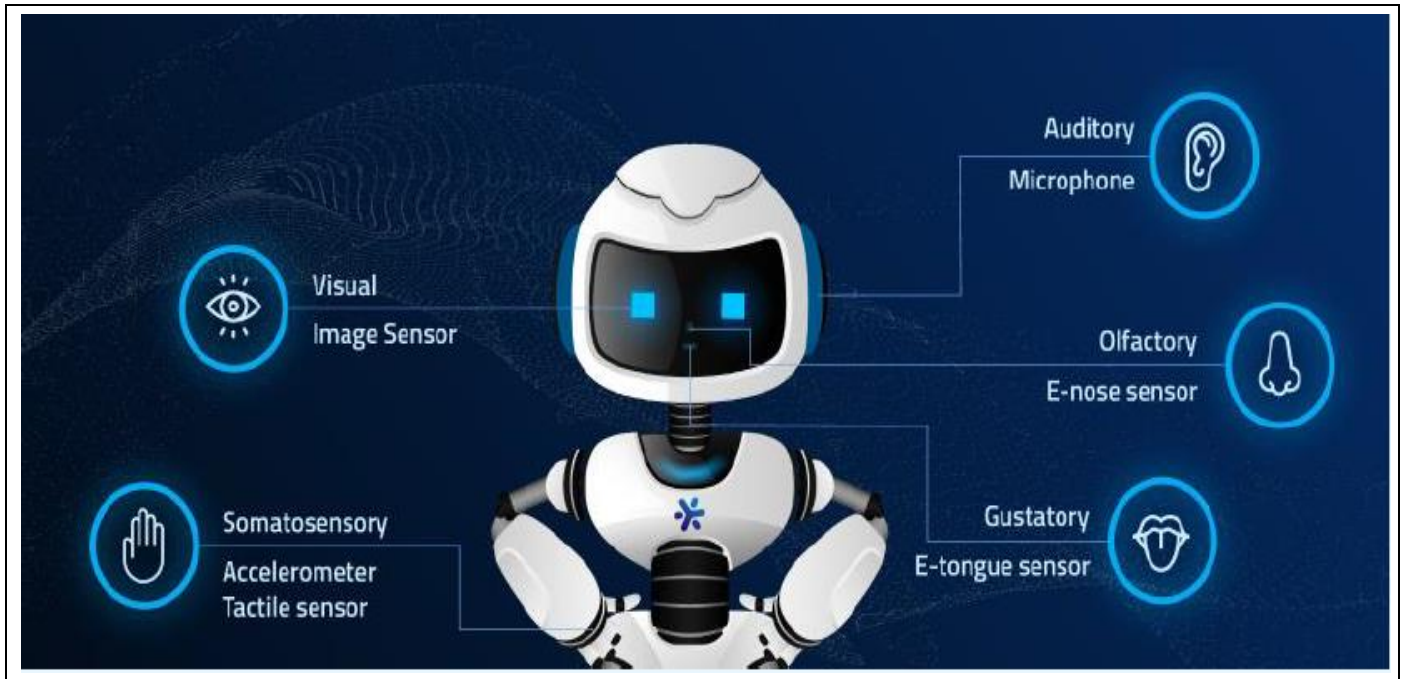
While edge devices depend on large data centers for data processing, Akida makes it possible to decentralise it. Additionally, there is lesser need for data transmissions, which improves data security.

Compared with a traditional processor, Akida processes data in a different way. While a traditional processor reads the instructions in a computer program and executes them, the AKD1000 is configured to learn a specific task. BRN’s neural processor is an event-based technology, known as Spiking Neural Network (SNN), which behaves like the cells in the human brain. Notably, Akida processes spikes, or events, instead of data and can perform 1.5 trillion operations per second while consuming very little power. We will elaborate on the exact workings of Akida later on. Suffice it to say that Akida analyses data, including images (Figure 6), sounds, data patterns and other sensor data, and extracts useful information generated from events in the data.

However, Akida can also perform on-chip convolution, which means that Akida can run deep learning CNN’s entirely on-chip as well. In other words, Akida tailors to customers’ existing CNN’s as well.



Figure 6: Akida enables efficient processing of all sensor modalities



Source: Company

As such, the Akida chip illustrates a game-changing technology for ‘Edge’ AI systems. After spending over a decade on the development of the technology, the company is now on the verge of commercialising this chip. The SNN technology can learn autonomously, evolve and relate information just like the biological brain. Given the capacity for on-chip learning, even after the chip is already operating “in the field”, Akida is expected to be well-suited for remote environments, such as space and agricultural areas, where internet access is limited or unavailable.

In December 2020, the National Aeronautics and Space Administration (NASA) placed an order under BRN’s ‘Early Access Program’ (EAP) to test Akida’s capacity for potential use in spaceflights. Notably, the EAP is only available to a select group of customers that require early access to the Akida chip, evaluation boards and dedicated support.

NASA ordered an Akida ‘Early Access Evaluation Kit’ to evaluate Akida for potential use in space

Akida has a neural fabric that comprises 1.2 million neurons and 10 billion synapses

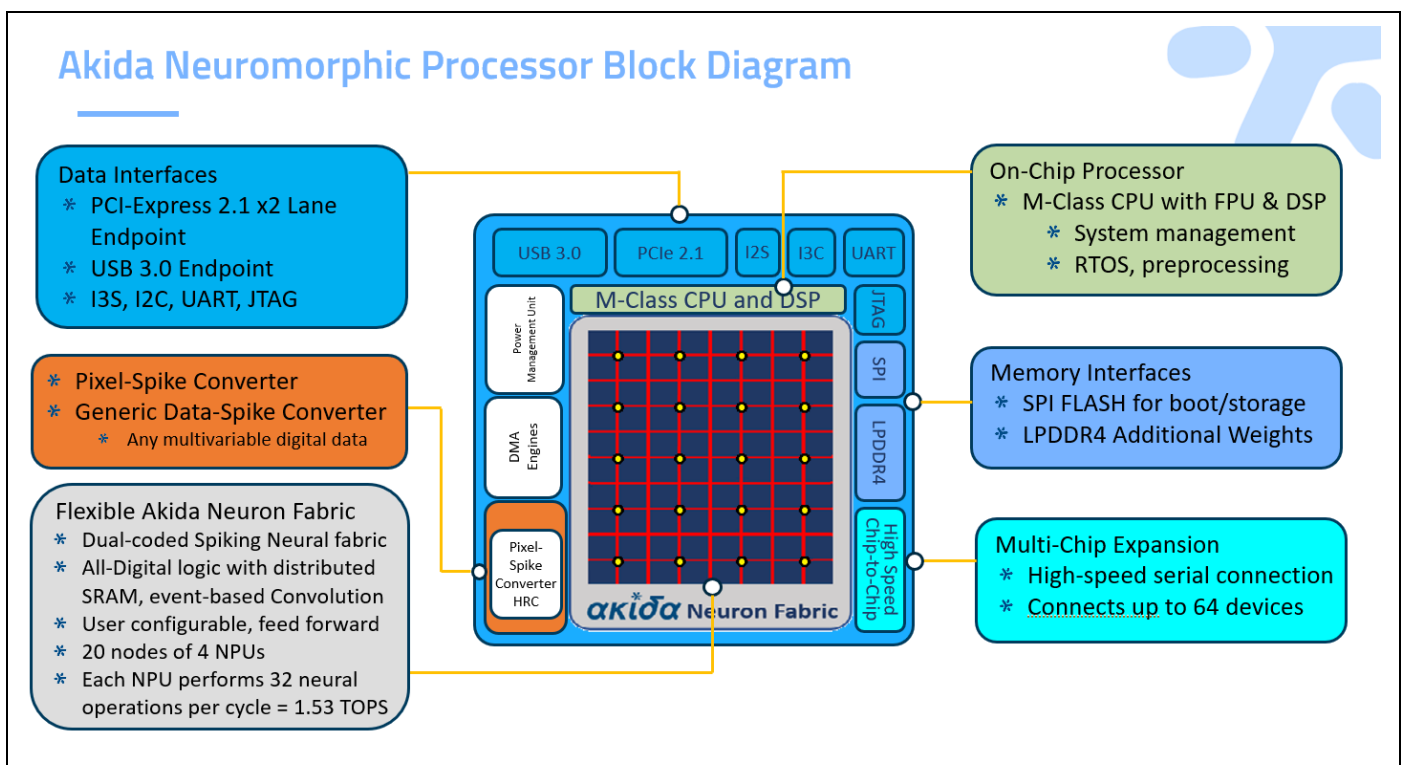
Akida – Biologically Inspired, Digitally Engineered

Akida offers a revolutionary new type of neuromorphic processing for edge devices. While the architecture of the Akida chip is almost identical to that of a biological neuron, it is implemented using a mainstream digital logic process. It has a neural fabric (Figure 7) that comprises 1.2 million neurons and 10 billion synapses.

The Akida neural fabric consists of a number of nodes and each node contains four Neural Processing Units (NPU). All of these nodes can communicate with each other simultaneously through a mesh network, which is what facilitates the parallel processing that in turn facilitates Akida’s brain-like learning capabilities.

Additionally, the chip has 8MB (Megabyte) of on-board SRAM (Static Random Access Memory), controllers and multiple interfaces to facilitate co-processing and interconnectivity with other devices. Akida also has embedded converters that convert binary data into spikes and spike trains. Using compilers and software, these converters take binary data, such as video data, and convert it into spikes that Akida can process in its neural core.

Figure 7: Akida architecture



Source: Company

On-chip learning with Akida

BRN’s neural processor allows for incremental learning, i.e., the ability to learn new tasks without forgetting old ones, and one-shot and continuous learning without the need for re-training. Although Akida learns autonomously, the technology can also be configured for a specific task.

During development the chip is trained using a known data set. This data set can be very small or very large, depending on the application and the type of network. As mentioned previously, Akida supports both CNN and SNN

Akida allows incremental, one-shot and continuous learning



networks as well as Deep Learning in addition to on-chip learning. Once trained, the application is deployed by copying what was learned to millions of chips.

A single Akida chip can run many different neural networks in the field, not just one. It is a matter of uploading the learned task, which is done in the field, not in the factory. Just like an Intel processor can run many different programs, Akida can run many different networks. Training is done in TensorFlow – a well-known development platform for neural networks. BRN’s MetaTF interfaces directly with TensorFlow.

By transferring the learnings from the trained chips to other Akida chips, these other chips will not all have to be trained individually. This transfer of learning is nothing more than copying the 10 billion synaptic weights to the Flash memory of an untrained chip.

Even though chips in inference mode are essentially in production mode, they may still fine-tune their synaptic weights “on-the-job”, meaning individual chips will keep on learning and get better at their specific tasks over time.

Figure 8: Comparing the brain, neuromorphic chip and GPU in AI inference mode

	Human brain	Neuromorphic chip	Deep learning on GPU
Power consumption	~20W	Micro to milliwatts	100s W
Processing speed	Milliseconds	Nanoseconds	Milliseconds
Efficiency (sparsity)	High	High	Variable
Learning rule	Local (we believe)	Local	Global
Event based processing	Yes	Yes	Less suitable

Source: Kisaco Research, Company

Akida has several other key advantages:

- **On-chip Convolution:** Convolution is a technique used in deep learning networks. The traditional software-based neural networks, such as Convolutional Neural Networks, can be efficiently run on BRN’s SNN. In April 2021, BRN launched MetaTF, which is a versatile Machine Learning (ML) framework that permits people working in the CNN space to easily switch to neuromorphic computing without the need to learn anything new. This development environment leverages TensorFlow for industry-standard neural network development.
- **TensorFlow Compatible:** Akida Chip is compatible with TensorFlow, which is an open-source platform for ML. As most data scientists are familiar with this environment, they can start using Akida straightway.
- **No requirement for an internet connection:** Akida doesn’t need a continuous internet connection to connect to the Cloud, as most data is processed on the chip itself. Thus, Akida has the potential to revolutionise Edge Computing — which involves managing complex AI tasks at the Edge of the network instead of sending data to the Cloud. Consequently, the reduction in system latency provides faster response times that can positively impact an application’s performance in real time. A great

CNNs can be run on Akida processor efficiently



example of this is a sensor in an ADAS implemented in an autonomous vehicle, where the sensor is scanning for objects in the car's path, e.g. pedestrians. Instead of having to send data back to the Cloud for processing, losing valuable time, Akida can make a brake or avoid decision fully autonomously in case anything is in the car's path.

Akida way ahead of competitors' solutions

BRN has a distinct competitive advantage over its peers as learning is autonomous in Akida chips. Various other neural networks developed by peers, such as Intel and IBM, still require either software code or human interaction to learn. While the most prominent processor chips in this field include IBM's TrueNorth and Intel's Loihi, both these chips are currently still lab projects.

IBM's TrueNorth: This chip is a neural network that contains one million neurons and 256 million synapses on a 28-nanometer chip. With 5.4 billion transistors, TrueNorth is IBM's largest chip to date in transistor count, and therefore very costly. The 256 synapses on each neuron need to be programmed, restricting the learning capability of the chip in real-time. TrueNorth needs to be trained off-chip. Moreover, to use this chip one needs to learn a new language, called 'Corelet'.

Intel's Loihi: This is Intel's fifth-generation self-learning neuromorphic research chip. Loihi is considered to be the closest competitor of Akida. The test chip contains 130,000 neurons. It's 128-core design is based on a specialised architecture that is optimised for SNN algorithms and fabricated on 14nm process technology. Loihi lacks convolution so running CNN's on Loihi is very inefficient. Adding to the inefficiency is Loihi's learning method, SLAYER, a slow, modified version of error backpropagation.

Both TrueNorth and Loihi are currently still lab projects and not commercially available



During 2020, the company signed multiple agreements under its EAP to gain market validation for its technology

Paths to revenue clearly visible

Akida has progressed well beyond development and is close to commercialisation. With the launch of the Early Access Program (EAP) in 2020, the company started on its path to commercialisation. The EAP focuses on specific customers in a diverse set of end markets to ensure availability of engineering samples and evaluation boards for key applications.

Since the launch, BRN signed multiple agreements under the EAP, to receive market validation for its technology. We have listed some of them below:


- In May 2020, the company inked an agreement with the Ford Motor Company to evaluate the Akida device for Advanced Driver Assistance Systems (ADAS) and Autonomous Vehicle (AV) applications.
- In June 2020, the company signed a joint development agreement with Valeo Corporation, a Tier-1 European automotive supplier of sensors and systems, for ADAS and AV applications.
- In September 2020, the company announced a collaboration with a privately held firm, VORAGO Technologies, to support a Phase I of a NASA program for a neuromorphic processor that meets spaceflight requirements.
- In December 2020, BRN signed its first IP license deal, with Renesas (see below).

BRN commenced shipments of the Akida evaluation boards to EAP partners in November 2020.

BRN signed its first IP licence agreement with Renesas Electronics America in December 2020

First IP licencing agreement with Renesas. BRN is expected to drive revenues through IP licencing agreements, royalties as well as chip and module sales, coupled with the design partner programme (Figure 9). Management signed its first IP licence agreement with global semiconductor manufacturer Renesas Electronics America in December 2020. This agreement should help Renesas gain a single-use, royalty-bearing licence to use the Akida IP. If this agreement turns out to be fruitful, we believe BRN can potentially attract dozens of IP licensing customers going forward.

Figure 9: Opportunity type — Sales process



Source: Company, Pitt Street Research



A typical semiconductor sales process incorporates seven stages

Ahead of the expected availability of commercial chips in October/November 2021, BRN has been marketing Akida IP, specifically engineering samples, through various channels.

Create awareness: In addition to potential direct sales, BRN has been highlighting the unique characteristics of Akida (ultra-low power, fast execution, on-chip learning, etc. while participating in trade shows to generate awareness. It has also been advertising online and in printed media.

Consideration: At this stage, it is determined whether the AKD1000 technology meets customers' requirements.

Evaluation: The customer engages and evaluates the Akida technology for application in its product.


Support: During the evaluation and early development stages, BRN provides support to customers for the integration of the Akida technology into their designs (how the Akida chip architecture fits into their system).

IP Licensing: Post the evaluation process, customers either purchase an IP licence (Figure 10) or place orders for physical AKD1000 devices.

Development and testing: Post product development and the integration of Akida technology, customers proceed to product testing, which can include internal product testing, market testing, target group evaluation, FCC qualification, etc.

Production and sales: BRN either sells AKD1000 chips or receives royalties from the customer.

Figure 10: Several paths to revenues



Source: Company

A very wide range of application areas for Akida

Akida is well suited for ADAS, AV, mobile phones, drones, vision-guided robotics, surveillance and machine vision systems

Akida is targeted at high-growth markets with multi-billion-dollar business opportunities, and we believe there are likely to be multiple prospective customers in each addressable industry vertical. BRN is already engaged with leading firms in major market segments, which highlights the many applications that will require AI-based microchips by 2025. Notably, we believe the Akida technology is particularly well-suited to ADAS, AV, mobile



phones, drones, vision-guided robotics, surveillance and machine vision systems.

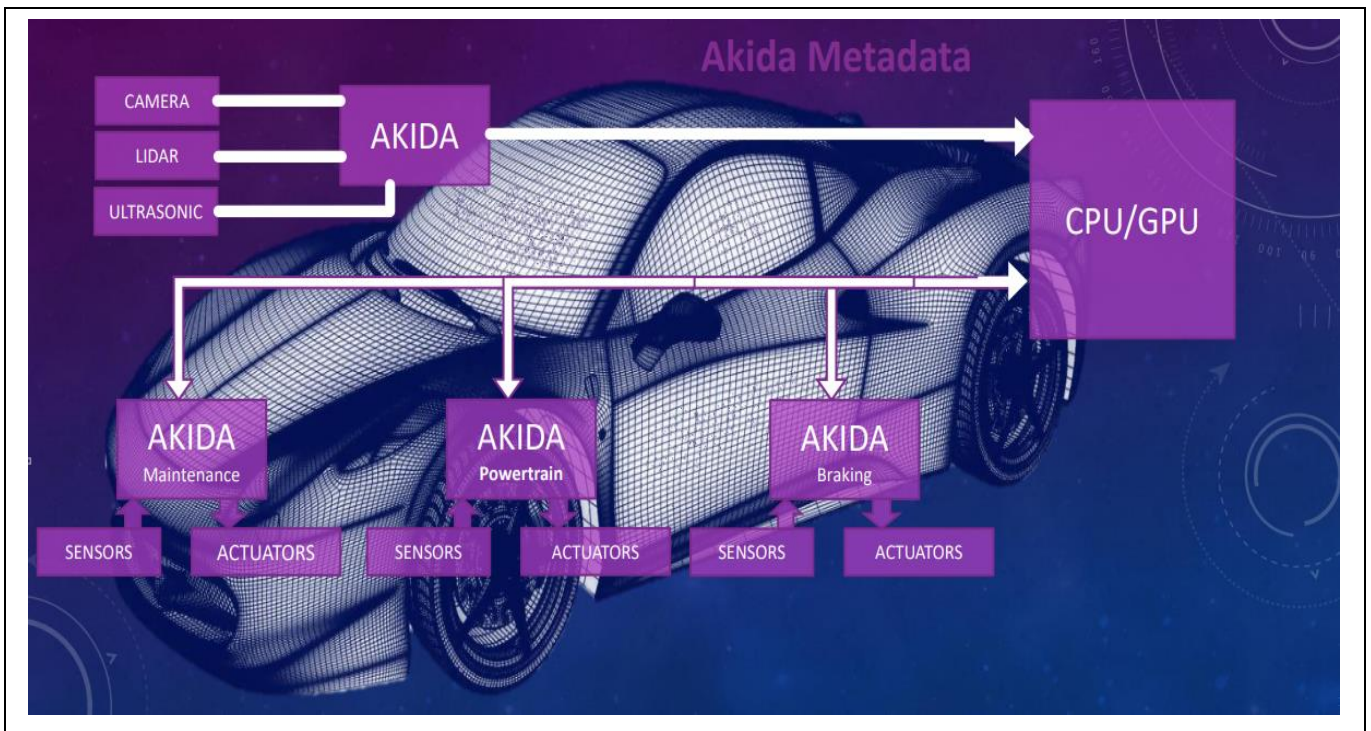
With the commercialisation of AKD1000 imminent, we believe BRN is well positioned to secure a variety of customer types for Akida, including cell phone manufacturers, semiconductor foundries, Automotive Original Equipment Manufacturers (OEMs) for ADAS and AV, third-party semiconductor IP providers, IDMs (Integrated Device Manufacturers) and firms in the Imaging space.

ADAS and AV among high-priority target markets. BRN is currently targeting customers and applications in embedded vision, cyber security and general IoT. As autonomous vehicles are required to process continuously changing data, integrating analytics in the sensor module is the Holy Grail for the automotive market (Figure 11). In our view, the embedded vision platforms for the automotive industry should be considered of high interest when it comes to commercial opportunities for Akida. Simply put, embedded vision platforms, which deploy SNNs, provide a car with a set of eyes in the form of multiple cameras and image sensors. ADAS and AV require substantial capabilities in the areas of object recognition and processing of radar and LiDAR (Light Detection and Ranging) data. Notably, real-time processing of imaging and other sensor data is crucial for safety, preventive care as well as dependability of autonomous systems.

The unique attributes of the AKD1000 processor offer a broad range of potential solutions to complex problems, including driver behaviour analyses and real-time object detection. Thus, we believe BRN’s engagements with Ford and Valeo Corporation under the EAP are very interesting and would aid the company in signing deals with additional automotive-related customers going forward.

Integrated analytics in the sensor module is the Holy Grail in Automotive

Figure 11: Akida opportunity in ADAS



Source: Company



AI is revolutionising the mobile phone industry. While embedded vision offers limitless applications, another high-growth segment for embedded vision is mobile phones. Embedded vision opportunities in cell phones entail functions such as gesture recognition, face detection and recognition. Mobile handset manufacturers are increasingly incorporating AI into mobile phone camera software. However, we believe this could be complemented (or potentially be largely displaced) by hardware-based neuromorphic computing inside the mobile phone. Additionally, we also see potential for software-based mobile applications, such as Google Assistant and machine learning applications, which currently require Cloud-access to function, to be replaced by Akida.

Smart Home. BRN is already in discussions with OEMs in the ‘Smart Home’ market. Akida technology is expected to be used in a wide range of home appliances that will improve the quality of life. For instance, a refrigerator equipped with Akida can provide an overview of the food in it – it could smell if food is about to expire and alert users. Likewise, video doorbells can recognise the person at the door, alerting the occupants of the house. Similarly, an intelligent washing machine could sense washing load, temperature, and wastewater.

Akida technology is expected to be deployed in emerging fields, such as Smart Home and Smart Healthcare

Figure 12: Expanding opportunities



Source: Company

Akida can detect COVID-19 with 93% accuracy and faster than PCR tests

Smart Healthcare: There is a plethora of opportunities in the ‘Smart Healthcare’ segment (Figure 12), which is an emerging field, and it needs AI in edge devices for fast and accurate diagnosis of several diseases. Akida can smell and classify a variety of diseases with the help of the right sensors. For instance, the company has collaborated with NaNose Medical to make a COVID-19 test kit that deploys AI to analyse exhaled breath samples. While the NaNose Medical sensor array collects data, Akida performs the diagnosis. Akida can detect COVID-19 (and cancers) with 93% accuracy, and faster than RT-PCR tests, from a breath sample or a drop of blood. Moreover, the Akida technology can be used in robotic surgery systems, X-ray and ultrasound



image interpretation and patient monitoring systems. Additionally, Akida can be deployed in the industrial segment, which includes process control, warehouse security, industrial drones, production robots and pollution monitoring.

Preventative Maintenance and structure analysis

Akida also has applications in Preventative Maintenance. It can analyse the vibrations or sounds emitted by bearings to determine if the bearing is faulty or going to fail soon. This is important in all sorts of moving stock, including trains and metros. If a bearing seizes up in a train wheel it will tear up the track causing the train to derail. Additionally, Akida can be used in structural analysis, e.g. bridges. Using a data set collected by vibration sensors on a bridge, Akida can determine the safety of the structure and its pylons by learning the resonance patterns of the bridge, predicting future failures.

Akida to generate four separate revenue streams

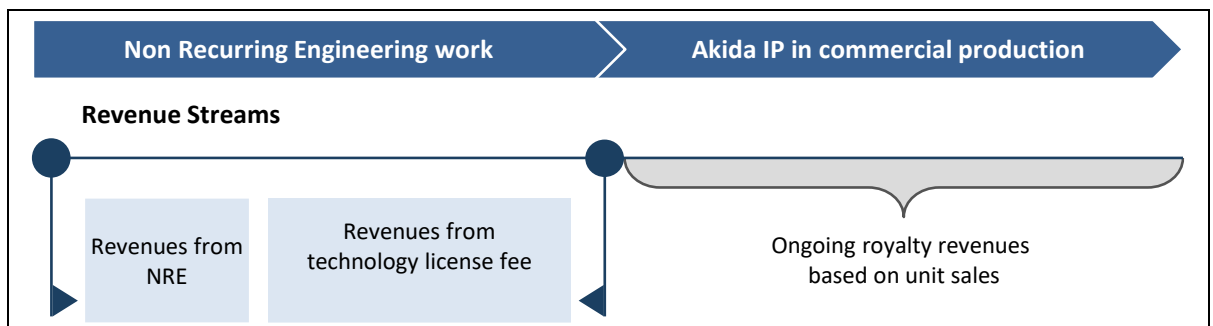
To leverage the substantial commercial potential of the technology, BRN will licence its Akida IP to other providers in the semiconductor industry. The company is expecting to generate revenues from three separate, but related streams:

- 1) Chip sales
- 2) Module sales
- 3) one-off licencing payments and
- 4) royalties from each chip sold by the company’s customers (Figure 13).

BRN also generates revenues from non-recurring engineering (NRE) work related to specific integration needs. While prospective customers have their own specific requirements for the integration of Akida chip architecture into their proprietary chipset (design-in) to create a complete System on a Chip, customisation requires intensive design work that can take 6–18 months. The one-off NRE fees cover BRN’s costs to adapt the technology to customer-specific needs. NRE fees vary widely across the industry, from several hundreds of thousands to millions of dollars. NRE fees are not scalable, however, like chip sales and IP licencing.

BRN is expected to generate revenues from four separate streams

Figure 13: Predicted revenue streams from Akida Intellectual Property



Source: Pitt Street Research

Chip and module sales: In addition to IP related revenues, we expect BRN to generate revenues from sales of Akida chips that it has manufactured by third parties, i.e. chip foundry TSMC in Taiwan. Foundries are third-party chip manufacturers that manufacture chips designed by other companies. For instance, a company like Apple could design a new mobile phone chipset using its own IP in addition to IP from design houses like ARM and Broadcom.



It can then outsource production of the actual chips to a foundry, like TSMC. By the way, we believe all of these types of players (design houses, foundries and electronics companies) are potential customers for the Akida IP. Going forward, BRN will have TSMC manufacture Akida chips, which it can then sell directly to customers.

IP licencing fees: Once the NRE work is complete and customers want to move into commercial production, they typically pay a one-off licence fee for the use of the technology. License fees can vary broadly and will be different for each customer, depending on the intended application areas, expected production volumes, the amount of IP to be used, etc. License fees typically vary from \$1m for single-product use to \$3m for multi-product use. Non-exclusive licencing will help the company sign multiple licensees to expand its revenue-generating potential.

Royalties: We believe the most lucrative future revenue stream for the company will be royalties paid by customers for each product they sell that includes Akida IP. These royalties are usually a percentage of the customer's revenue from sales and typically range from 2% to 15%, again depending on the intended application areas, the amount of IP used and expected production volumes. Notably, royalty percentages also depend on the uniqueness of the IP that is being licensed. As the specifications and features of Akida are quite unique vs. other technologies, including Intel's Loihi and IBM's TrueNorth, this may help the company charge higher-than-average royalty percentages for Akida.

Other royalty revenue models simply use a fixed dollar amount per chip sold. This is a preferred model for many high-volume production companies, including cell phone manufacturers.

The most lucrative future revenue stream for BRN will be royalties paid by customers for each chip they sell that includes Akida IP

Example: Assume BRN entered into a licencing agreement with a customer to licence its technology with specific integration requirements. Terms could include US\$500k of NRE fees, US\$1m of one-time licencing fees and 5% royalty on sales from each of its chips that incorporates Akida. Further assume that this customer will sell 1m, 5m and 10m of these chips in years 1 through 3, respectively, at US\$25 each.

In this example the company will recognise the following three revenue streams from this customer:

- 1) NRE fees of US\$0.5m.
- 2) One-time licencing fees of US\$1m before the start of production.
- 3) Royalties of US\$1.25m (1m x US\$25 x 5%), US\$6.25m and US\$12.5m, respectively, in the first three years of production.

Certified Solution Providers to ramp up sales capacity

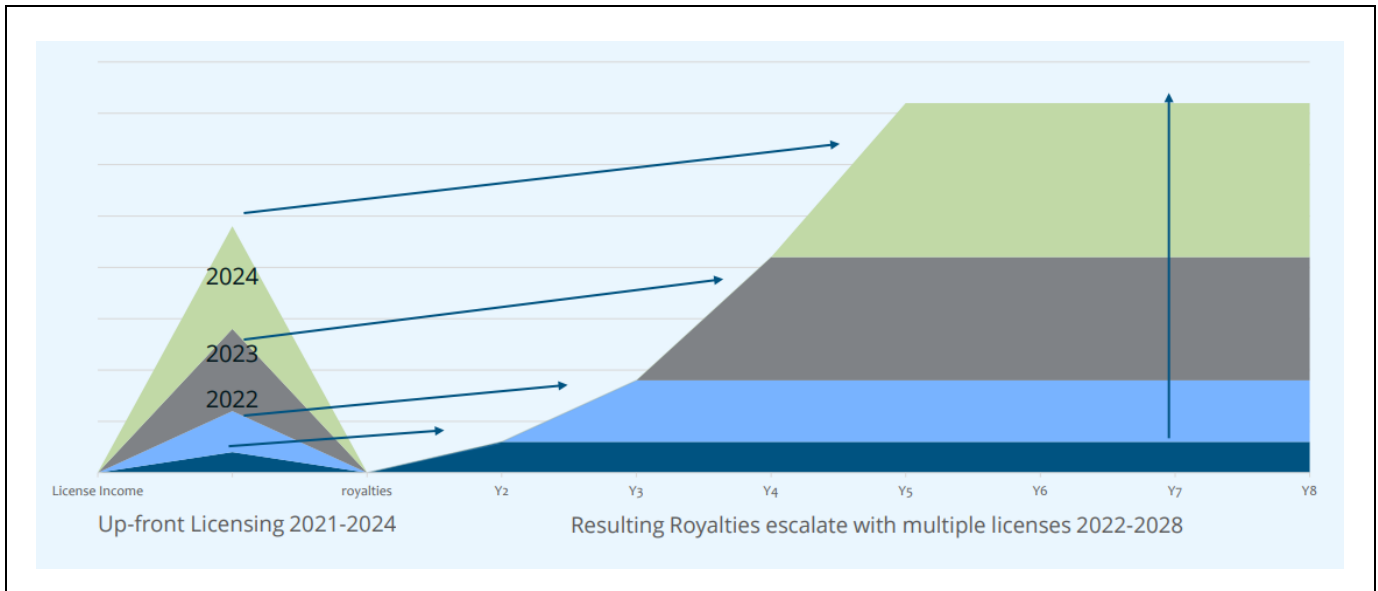
In our view, the example above clearly depicts the scalability of the company's revenue model as well as the revenue potential for Akida. Notably, the recurring royalties are expected to surge with multiple licensees (Figure 14).

Going forward, BRN may also set up a network of "Certified Solution Providers" (CSP) in a range of countries to ramp up the Akida sales capacity. These CSP's will be trained by BRN to use the Akida chip. These CSP's will already understand neural networks, deep learning and the engineering challenges to deploy such systems in the field. Working with Akida involves interfacing to sensors, data collection and scaling as well as processing of the inference results from Akida in addition to training the network with clean



data. All this could all be handled by CSP's who will be provided the tools, know-how and hardware to enable them to sell Akida. Training CSP's could become a fifth revenue stream for BRN in the future.

Figure 14: Conceptual IP licencing and royalties model



Source: Company

Volume production of Akida commences

Owing to the huge costs involved in building semiconductor manufacturing facilities, BRN is collaborating with SocioNext, a global developer of advanced ASIC solutions, for product development and manufacturing of the Akida AKD1000. The company achieved a key milestone after SocioNext released the engineering layout of the production version of the AKD1000 chip to TSMC, which also completed a multi-design wafer (MDW) previously to test the Akida design and identify any potential flaws.

While TSMC has commenced volume manufacturing, the production units are expected to be available for testing in Mid-August 2021. BRN is expected to allocate the chips to clients in strategic end markets that have bought the Akida evaluation systems for various applications.

Secured production slots with TSMC

With delivery of the first batch of Akida chips now imminent, we believe commercial agreements are likely to follow, backed by the EAP. Typically, the delivery time for large orders is usually about 3–4 months. Despite the current manufacturing capacity shortages on the part of chip manufacturers everywhere, especially the foundries, BRN's partner SocioNext has already reserved production slots with TSMC, which will certainly help BRN avoid the worst of the capacity shortages as it is working to ramp up production.

Gartner predicts that the global semiconductor shortages will likely persist until the middle of 2022.

Its wafer fabrication partner, TSMC, has recently commenced volume production



End markets for Neuromorphic computing are vast and diverse

As compared to conventional field-programmable gate arrays (FPGA), central processing units (CPUs) and GPUs, neuromorphic chips offer improved energy efficiency, reduced latency and, in the case of Akida, have on-chip learning.

Given these advantages of neuromorphic technologies, they have immense potential for application in various end markets, including industrial, aerospace & defence, healthcare, automotive, consumer goods, etc.

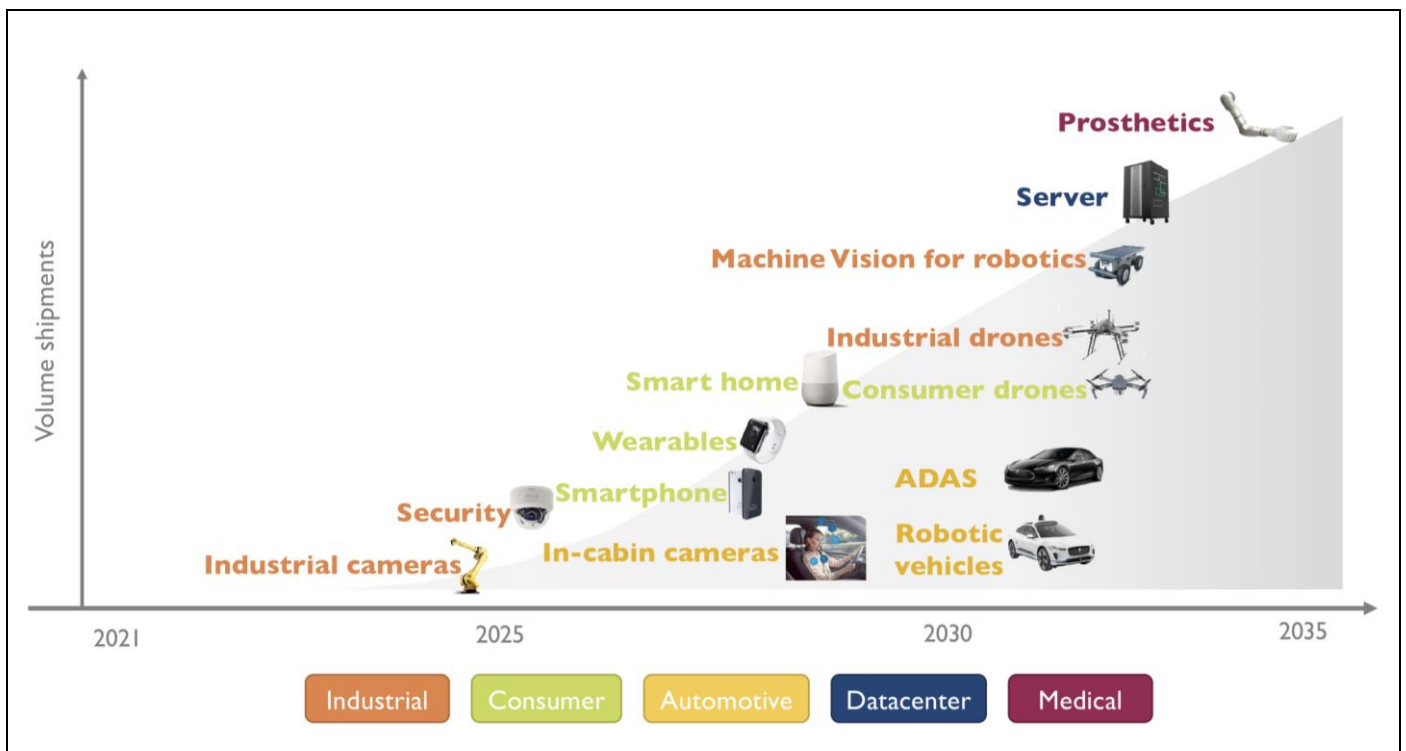
According to the 'Neuromorphic Computing and Sensing 2021' report by Yole Développement, a leading semiconductor market research firm, industrial applications are expected to be the first to see neuromorphic computing implemented (Figure 15). This will be facilitated by the low latency of these technologies, coupled with their autonomous and offline learning capabilities. Notably, BRN is one of the few players in the market, along with Prophesee and Nepes Corp (General Vision Inc.), that already have products targeting this segment, which we believe provides it an early mover advantage. However, neither Prophesee nor Nepes have on-chip convolution, while both have only very low numbers of neurons.

In terms of the speed at which neuromorphic computing will roll out, Yole has based its conclusions on SNN's without convolution. However, Akida has on-chip convolution and can therefore run applications that are in existence today, not 2025 as illustrated in Figure 15. So, we'd argue that Yole is being too conservative in its forecasts.

Neuromorphic technologies have numerous applications in industrial, consumer and automotive markets

Yole is being conservative

Figure 15: Neuromorphic computing adoption by end markets between 2021-2035



Source: Neuromorphic Computing and Sensing 2021 report, Yole Development (2021)



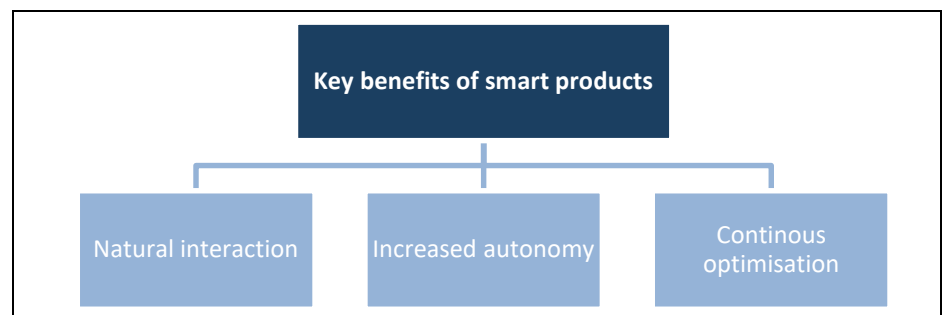
The consumer market is also expected to benefit from **neuromorphic computing**, as it brings AI capabilities to Edge. This also provides added advantages related to ensuring the privacy and safety of personal data. Furthermore, cybersecurity and fraud detection use cases will also benefit from neuromorphic computing by leveraging low latency and online learning capabilities of these technologies. Moreover, given that neuromorphic computing offers considerable power efficiency, it could also help to limit the power consumption growth in data centers.

BRN targets lucrative end markets

We believe the company has the potential to target highly lucrative end markets, such as industrial, consumer, mobile and automotive. There is a rising demand for smart products, primarily driven by the various advantages offered by these products (Figure 16):

- **Natural Interaction:** Voice and gesture-controlled products in home entertainment, automotive interiors and industrial equipment will increasingly replace products with physical buttons and control panels.
- **Increased autonomy:** Smart products, such as self-guiding vacuum cleaners, allow users to focus on other tasks by operating independently with minimal need for intervention.
- **Continuous optimisation:** Smart products are able to collect and analyse data and accordingly make adjustments in order to improve performance and better serve the needs of the business or customer. For example, a thermostat that adjusts the room temperature for different situations, or a robot that develops efficient ways to navigate a warehouse. Companies providing these products can also use this information to develop new products based on the enhanced understanding of the needs of the users.

Figure 16: Smart products provide various benefits to users



Source: 'Driving intelligence at the edge with Neuromorphic Computing', Accenture

Demand for smart products is expected to rise with growing concern for safety, security, and convenience

Ushering in the era of smart products

Smart products are expected to see an increase in demand as consumers become more willing to integrate technology in their daily activities. According to Accenture², the smart home market alone is expected to be worth US\$135bn by 2035. Using Akida technology, businesses can leverage the power of neuromorphic computing to enhance the performance of their smart products by reducing latency, improving energy efficiency and allowing for continuous, unsupervised learning.

² 'Driving intelligence at the edge with Neuromorphic Computing', Accenture (2021)



- **Aerospace and defence:** According to MarketsandMarkets³, Aerospace & Defence is expected to hold the largest market share for neuromorphic computing (around 32%) in 2021. Given the higher processing speed of these technologies, it is expected to help the military and defence industry in handling sensitive information related to battlefields, including weapons and resources management. The need for image recognition solutions in this vertical is also expected to drive growth of neuromorphic technologies. Additionally, neuromorphic computing can enhance the performance of drones, as well as help with surveillance and security, tracking and targeting.
- **Automotive:** This is another segment that offers vast growth potential for neuromorphic computing technologies. Currently, autonomous vehicles use CPUs and GPUs, which rely on external systems, such as the Cloud for data processing. However, this results in latency issues for on-board ADAS, which are required to make split-second decisions, such as avoiding objects by interpreting images.

Furthermore, facial recognition becomes important in automotive use cases, not just for the external environment, but for inside the vehicle as well, e.g., to automatically adjust a car's settings for a specific driver. Additionally, through facial recognition, Akida can, for instance, be used to see if the driver is paying attention to the road. This helps in mitigating the risk of road accidents.
- **Industrial:** As mentioned previously, industrial applications are expected to be the early adopters of neuromorphic computing technologies. There are many potential use cases in this vertical, such as warehouse security, industrial drones, production robots, process control, etc. Currently, programming robots is a slow and tedious process. However, by incorporating AI with neuromorphic computing in an industrial setting, we believe businesses will be able to leverage continuous on-board learning to improve the performance of robots in the factory.
- **Healthcare:** Neuromorphic computing also has application potential in a range of use cases in the healthcare segment. Some of these include analysis of breath compounds, blood samples and patient statistics, as well as robot-assisted surgery and detection of fraud in health insurance. BRN's management is already working with Biotome Pty Ltd. to develop a COVID-19 antibody test. The aim is to use Akida with Biotome's handheld device to interpret the sensor responses for the detection of antibodies and providing results in seconds.
- **Smart home:** Rising disposable income in developing countries and consumers becoming receptive to integrating technology in their daily lives are expected to drive demand for smart home products. Smart speakers, such as Amazon's Alexa and Google Assistant, have already witnessed tremendous uptake and the market is expected to continue to grow at a CAGR of 17.1% over 2020–2025 to reach US\$15.6bn by 2025, according to MarketsandMarkets⁴. By combining the advantages of neuromorphic computing with AI, this vertical is expected to see a host of new products, such as smart thermostats, lighting, video doorbells, refrigerators, washing machines, etc.
- **Smart city:** Driven by rising demand for sustainable infrastructure as the global population and urbanisation continue their upward trend, smart city applications are expected to grow strongly. Smart cities not only

Industrial sector is expected to be among the first to adopt neuromorphic computing technologies


³ 'Neuromorphic Computing Market', MarketsandMarkets (May 2021)

⁴ 'Smart Speaker Market', MarketsandMarkets (June 2020)



include digital interfaces in traditional infrastructure, but also leverage the insights generated from data to improve the quality of life of the residents (Figure 17). Using neuromorphic technologies in use cases, such as real-time monitoring for security, cyber-crime, environment, traffic and healthcare infrastructure, is expected to significantly improve the quality of life in these cities.

Figure 17: Smart cities infrastructure



Source: 'Smart Cities', Mckinsey Global Institute Analysis (June 2018)

Underpinned by its advantages over other processing technologies currently in use, neuromorphic computing is positioned to benefit from the boom in investment for the development of smart cities, i.e., autonomous devices. According to Forbes⁵, the market size for the global smart city industry is expected to almost double to reach US\$821bn by 2025, from US\$411bn in 2020.

Combining neuromorphic processing with AI unlocks vast potential

Currently AI is being driven by brute force, i.e. exhaustive training. This means that the system processes all the data available to it in order to give an output. Neuromorphic computing solutions such as Akida, however, only process spikes, thereby significantly reducing the computational effort required. Additionally, in CNN inference, all data is processed, even if it contains zeros. In event-based CNNs, however, only relevant data is processed, e.g. data that is not zero. When a zero is produced there is no spike, and a spike defines an event. In "traditional" computational CNNs all data is processed, even if it is zero, which means higher energy use and slower outcomes.

Moreover, in order to unlock the full potential of AI, the computing process will need to overcome the challenges posed by the dynamics of Moore's Law, as well as issues related to overheating. Moore's Law, proposed by Gordon E

Neuromorphic technology could help AI unlock its true potential

⁵ 'Investing In Innovation: The Rise Of The Smart City', Forbes (December 3, 2020)



Moore, (the co-founder of Intel), states that the number of transistors on a microchip doubles every two years. However, experts now believe that computers will reach the physical limits of this Law soon. This is evident in the slowdown in the performance improvements of CPU's. On the other hand, AI computational power requirements are increasing at a rate faster than traditional processor development, according to research conducted by Stanford University⁶.

As a result, the two major bottlenecks that hinder growth in processing power are the physical limitations of silicon and the so-called von-Neumann computer architecture, which operates under program-control and stores all results in memory, which causes a bottleneck in the processing chain.

Akida is a non-von Neumann computing device, which eliminates the memory bottleneck and does not depend on program instructions. Instead, each core is dedicated to processing sensory data in the form of spikes that indicate events. The cores can be arranged to execute any form of neural network.

In other words, the problems described above can be solved by switching to non-von Neuman computing architectures, such as Akida.

Further, in order for AI to leverage the 5G, IoT and robotics revolution, it will also need to overcome challenges related to the data wall. With the increase in the speed of the Internet and the number of connected IoT devices, the volume of data increases exponentially. Moreover, the volume of data needed to train AI models, such as those in Natural Language Processing (NLP) and Speech Learning, has also increased exponentially⁷.


This is where neuromorphic computing, and specifically Akida, comes in, with its non-von Neumann architecture and its ability to process only spikes, or events, instead of processing all available data. Underpinned by the ability to overcome these challenges, neuromorphic computing is expected to account for 18% of the total AI computing and sensing market by 2035 (Figure 18), according to Yole Développement.

Akida is a non-von Neumann computing device

Yole Développement expects the neuromorphic sensing market to reach up to US\$5bn by 2030, (CAGR of 116% over 2025–2030)

Neuromorphic computing market is expected to grow at CAGR of 88% over the same period, to reach US\$2bn in 2030

Figure 18: Revenue evolution for Neuromorphic in AI sensing and computing



Source: Neuromorphic Computing and Sensing 2021 report, Yole Development (2021)

⁶ 'AI Index 2019' Annual Report, Stanford University

⁷ 'AI and Memory Wall', Medium (2021)



Valuation of A\$1.50 per share

We previously wrote about BrainChip in a report published on 15 May 2019, available here: <https://www.pittstreetresearch.com/brainchip>. In this report we looked at several acquisitions in the neuromorphic computing space in order to properly value BrainChip. Specifically, we looked at Intel's (NASDAQ:INTC) acquisitions of Nervana and Movidius in 2016, which were both completed at valuations of approximately US\$400m, or A\$600m, at the time. Applying this valuation to BrainChip yielded a value per share of A\$0.44 per fully diluted BRN share when the stock was trading around A\$0.04. In 2021 we have seen the shares largely trade in a \$0.40-0.60 range.

The proof of the pudding....

However, the starting point for our \$0.43 valuation two years ago was very different from where BrainChip is at today. Back then, the company was still fully in development mode and still more than two years away from having a commercially viable product in hand.

Today, the company has received its first batch of production chips and has fully tested the previous batch of engineering samples that were produced in the 3rd quarter of 2020. This will kick off Akida's commercialisation and should lead to first revenues for BrainChip. So, the company's valuation premise is very different from two years ago.

But until the purchase orders and licence agreements start coming in and that revenue actually shows up in the P&L, many investors will still regard BRN shares as pre-revenue and highly speculative. Consequently, they will continue to value the company as such. And despite the company's first IP license deal with Renesas, that may not change until the day the company announces its first revenues. In other words, the company is pre-revenue until it's not. Consequently, we expect a step-change in BRN's valuation if and when that first revenue-generating announcement is made.

What could a post-revenue BRN be worth?

In anticipation of that first deal we need to see what BRN shares could be worth as a revenue-generating company. While it's easy enough to find funding rounds for AI-related companies, like Neuralink (raised US\$205m in a series C round in July 2021), Mythic (US\$70m in series C) and Untether (US\$125m in series C), the actual valuations at which these companies raise these funds is hardly ever disclosed.

So, once again, we'd like to resort to looking at Intel's acquisition targets, Intel being one of the semiconductor industry's bellwethers. In an effort to keep up with Nvidia (NASDAQ:NVDA), a leading player in GPU's (Graphic Processing Unit) for gaming, AI and SoC's for mobile and Automotive markets, Intel acquired Israel-based start-up Habana Labs for US\$2BN in 2019. Nvidia's GPU's are increasingly used in Machine Learning applications and the company has been an aggressive acquirer of semiconductor IP, including Mellanox, a high speed network / HPC company, which it acquired for US\$6.9BN in March 2019, and the intended US\$40BN acquisition of ARM. The latter deal is still awaiting approval from regulators around the world.

When Intel acquired Habana Labs, this company had already launched the Goya chip for inference processing in September 2018 and the Gaudi chip for training in June 2019, but was not generating significant revenues just yet. So, in terms of commercialisation, we believe BrainChip today is getting close to where Habana Labs was about two years ago.

BRN's valuation premise is very different from two years ago

Habana Labs acquired by Intel for US\$2BN in 2019

BrainChip today is getting close to where Habana Labs was two years ago



Value of A\$1.50 per share fully diluted

Using Habana Labs valuation of US\$2BN and an AUD/USD exchange rate of 0.735 to value BRN, we arrive at a value of A\$1.64 per outstanding BRN share and A\$1.50 fully diluted.

Funding and capital structure

As per 30 June 2021, BRN held US\$17.7m in cash while the company's cash burn averaged US\$3.4m per quarter in the first six months of FY21 (ends December).

In August 2020, BRN signed a Put Option Agreement (POA) with LDA Capital that provided BRN access to A\$29m in potential equity funding with BRN committing to drawing down a minimum of A\$10m. In October 2020, the total funding amount was increased to A\$45m with BRN's minimum commitment increasing to A\$20m.

Although the workings of the various agreement types vary, this POA is similar to Controlled Placement Agreements (CPA) and At-the-Market (ATM) agreements in that they enable companies to sell their own shares on market through a third party. They provide companies with control over the amount of shares they wish to sell and at which minimum price they wants to sell at.

On 16 August 2021, BRN announced a capital call with LDA Capital in which it indicated the sale of 8.75m BRN shares to LDA under the POA. The price BRN will receive depends on several factors, including the 10-day VWAP (volume weighted average price). But assuming the sale price will be A\$0.50, this transaction will bring in approximately A\$4.4m in funding.

BRN still has A\$34.2m of funding remaining under the agreement with LDA, with a minimum remaining commitment on the part of BRN to draw down another A\$9.2m by 22 October 2021 at the latest.

While the advantage of this POA is that BRN has a minimum guaranteed level of equity funding for the duration of the agreement, the downside of this POA specifically (as opposed to CPA's and ATM's) is that LDA can sell the shares it acquires from BRN at any given moment, without regard to market conditions. Investors need to be mindful that this can cause sudden and downward movements in BRN's share price, although LDA is bound by maximum daily volume limits.

Including the pro forma issue of shares to LDA, the fully diluted number of BRN shares (including outstanding options) amounts to 1,808.6m (Figure 19).

Figure 19: Capital structure

Class of security	In millions	% of fully diluted
Quoted Securities		
Ordinary shares (BRN)	1,659.7	91.8%
Pro forma LDA capital call 16 August 2021	8.75	0.5%
Unquoted		
BRNAD	135.7	7.5%
BRNAJ	4.5	0.2%
Fully diluted shares	1,808.61	

Source: Appendix 2A - June 2021 and announcement dated 16 August 2021



Risks

We see five main risks related to BRN's investment thesis and traded shares:

- 1) Execution risk: Any delay in Akida's commercialisation plan will negatively impact BRN's valuation.
- 2) Lower adoption rate of Akida technology: While Akida is expected to be widely used across various end markets, lower than anticipated adoption rates by customers may hamper BRN's growth.
- 3) Funding risk: BRN may need to raise additional capital to support its development activities until it reaches cash flow break even, diluting current shareholders.
- 4) Technology risk: Although we believe BRN is well-ahead of its competitors, with on-chip learning and on-chip convolution, which no other company has been able to do. The company also owns the patents on those technologies. As with all types of technology, though, we need to mention that the rapid pace of development in the Artificial Intelligence space may lead to BRN's technology becoming obsolete.
- 5) Share overhang: LDA Capital can sell the shares it owns in BRN (other than collateral shares) on market at any time, depending on certain daily volume restrictions, which may result in sudden and adverse share price movements.



Summary and Conclusion

As BRN moves closer towards the commercialisation of AKD1000, it seems well-poised to target high growth end markets, such as Industrial, Consumer, Mobile and Automotive, with multi-billion-dollar business opportunities. The first revenue opportunity has already opened up as BRN signed the IP licence agreement with Renesas Electronics America in December 2020.

If this agreement turns out to be fruitful, we believe BRN can potentially attract dozens of IP licensing customers going forward. Given an asset-light licensing model with recurring royalty revenues, we expect high gross margins once Akida sales ramp up.

Amidst chip shortages globally, the first batch of Akida chips is expected to be available in August 2021 through its wafer fabrication partner, TSMC. We expect BRN to convert current prospects through its EAP pilot projects.

Additionally, current discussion partners in each addressable industry verticals can potentially also be converted into paying customers.

Akida is geared to unlock the full potential of edge computing backed by the key differentiators, which set Akida apart from its peers:

- **On-chip learning:** Even after the chip has been deployed in the field, Akida can continue to learn fully autonomously, which is a key differentiator.
- **Ultralow power consumption:** The AKD1000 chip has ultra-low power consumption in the range of 100 microwatt to 300 milliwatt depending on workload.
- **Neuromorphic processing has an edge over traditional processing:** While a traditional processor processes data, AKD1000 processes spikes, or events. Akida can perform 1.5 trillion operations per second.
- **On-chip convolution:** The traditional software-based neural networks, such as Convolutional Neural Networks (CNN), can be run on SNN efficiently.
- **TensorFlow compatible:** Most data scientists are familiar with the TensorFlow environment and this enables them to start using Akida chip straightaway.
- **Standalone usage possible:** Akida doesn't need an ongoing internet connection to connect to the Cloud as most data is processed on the chip itself.
- **Green Technology:** Akida is ~97% more energy efficient when compared to processing the same task in a data center.
- **Increased security:** As Akida makes it possible to decentralise data processing, there is much less need for data transmissions, improving data security.
- **Minimise latency:** As most of the processing is done on the chip itself, there is a reduction in system latency, which provides faster response times.

Using fairly recent industry M&A valuations, we value BrainChip at A\$1.51 per share fully diluted.



Appendix I – Introduction to AI, machine learning and deep learning

Artificial Intelligence comes in different shapes and forms

Artificial Intelligence (AI) is the science of training systems to perform human tasks through learning and automation. AI makes it possible for machines to learn to apply logic, adjust to new inputs and reason to gain an understanding from complex data. In simple words, AI provides machines with the ability to learn from data it receives by processing and recognizing patterns in that data.

AI is an overarching term and essentially consists of foundational building blocks and key elements, namely machine learning and deep learning, computer vision, natural language processing (NLP), forecasting and optimization, and machine reasoning. These building blocks or elements can be used independently or combined to build AI capability. Several AI capabilities and their use cases in a business context are illustrated in Figure 20.

Figure 20: AI use cases

AI capability	Use case
Pattern Recognition	For fraud detection by understanding typical trends/behaviors within customer financial transactions data or spot anomalies in an account’s spending data
Prediction	For improving energy consumption forecasts by capturing short- and long-term variability in data
Classification	For supporting wildlife conservation efforts by examining animal track changes and grouping them by species type
Image recognition	For determining if nodes on a raw CT scan are malignant or benign; other applications include predictive diagnostics and biomedical imaging
Speech-to-Text	For transcribing voice messages to text for sentiment detection and further analysis in call centers
Cognitive search	For offering personalized recommendations to online shoppers by matching their interests with other customers who purchased similar items
Natural language interaction (NLI)	For generating automated financial reports on sales revenue predictions, which otherwise would be generated by the user itself
Natural language generation (NLG)	For automated generation of summaries after analyzing large sets of documents

Source: Pitt Street Research

These AI capabilities can be used either independently or combined with each other, depending on users’ objectives and underlying data. For instance, in the banking industry, combinations of these capabilities are used for credit and risk analysis and to provide market recommendations by creating automated financial advisors. In the healthcare industry, such combinations are used for processing data from past case notes, biomedical imaging, health monitoring, etc. Other industries such as manufacturing and retail are also utilizing AI capabilities to optimize supply chains or to offer personalized

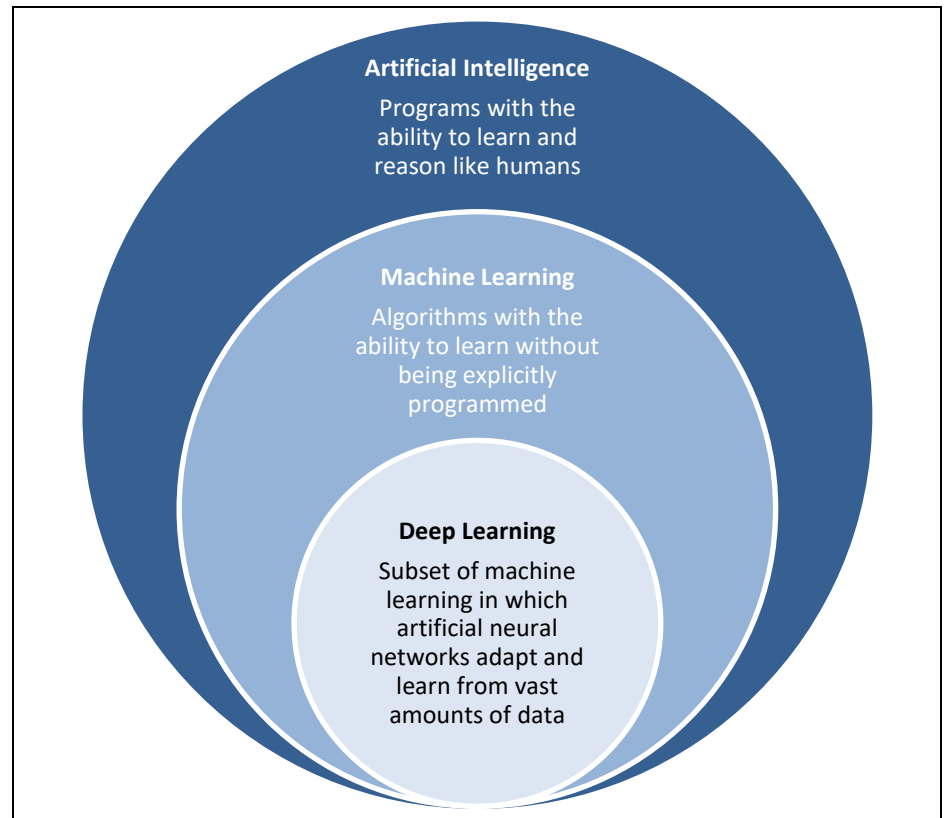


shopping experiences and customized recommendations. In addition, governments across the globe are focusing on building smart cities and utilizing capabilities such as facial recognition for use in law enforcement.

Machine Learning and Deep Learning

While AI comprises all techniques that make machines perform tasks that require intelligence, Machine Learning specifically imitates how humans learn. Basically, Machine Learning is a subset of AI (Figure 21) and consists of the techniques that enable machines to learn from the data without being explicitly programmed to do so. Conversely, other AI techniques could be classified as rules-based or expert systems, which work on a pre-defined algorithm or logic – like performing accountancy tasks, in which the system runs the information through a set of static rules.

Figure 21: Machine Learning and Deep Learning – Subsets of AI



Source: Argility, Pitt Street Research

Machine Learning is dynamic in nature and has greatly evolved over the years

One aspect that separates Machine Learning from rules-based expert systems is the ability to modify itself when exposed to more data, i.e. machine learning is dynamic and does not require human intervention to make certain changes.

Though Machine Learning has evolved a lot over the years and is used to tackle many problems, for a long time it was still difficult for machines to perform many tasks such as speech, handwriting and image recognition, and more mundane tasks such as counting the number of items in a picture. The concept of Artificial Neural Networks (ANN) kick-started the development of Deep Learning, which provides machines the capability to perform tasks such as image recognition, sound recognition and recommender systems with much greater accuracy and speed.




Deep Learning itself is essentially a subset of Machine Learning and is all about using neural networks comprising artificial neurons, neuron layers and interconnectivity. Instead of organizing data to run through predefined equations, Deep Learning sets up basic parameters around the data and trains the computer to learn on its own by recognizing patterns using many layers of computer processing.

Artificial Neural Networks learn like the human brain does

Artificial Neural Networks (ANNs) are computing systems with a large number of interconnected nodes that work almost like neurons in the human brain. They use algorithms to recognize hidden patterns and correlations in raw data and then cluster and classify that data to solve specific problems. Over time, neural networks continuously learn from new data and apply those learnings to make future decisions.

A simple neural network includes an input layer, an output (or target) layer, and a hidden layer in between. The artificial neurons (or nodes) in these layers are interconnected and form a network termed as a neural network of interconnected nodes (Figure 22). As the number of hidden layers within a neural network increases, deep neural networks are formed. A simple ANN might contain two or three hidden layers, while deep neural networks can contain as many as 100 hidden layers.

Figure 22: Neural Network basic diagram



Source: ExtremeTech

In a typical neural network, a node is patterned after a neuron in a human brain. These nodes get activated when there are sufficient stimuli or inputs (just like neurons in a human brain). This activation spreads throughout the network, creating a response to the stimuli (output). The connections between these artificial neurons act as simple synapses, enabling signals to be transmitted from one to another. Signals across layers travel from the first (input layer) to the last (output layer) and get processed along the way.

While solving a problem or addressing a request, data such as text, images, audio and video, is fed into the network via the input layer, which communicates to one or more hidden layers. Each neuron receives inputs


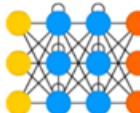

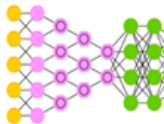


from the neuron to its left, and the inputs get multiplied by the weights of the connections they travel along. These input-weights are then summed up. If the sum is higher than a certain threshold value, the neuron fires and triggers the neurons it is connected to on the right. In this way, the sum of the input-weight product determines the extent to which a signal must progress further through the network to affect the final output.

Many types of neural networks

Over the past several years, many neural networks with different architectures and specifications have emerged. Feedforward Neural Networks (FNNs) are the simplest form of ANNs. For specific tasks, more complex ANNs have been invented, including the Convolutional Neural Networks (CNNs), which aim to mimic the human visual system, as well as the Recurrent Neural Networks (RNNs), which are used to interpret sequential data such as text and video. These major types of ANNs are described in Figure 23.

Figure 23: Types of Neural Networks Other Than Spiking Neural Networks

Neural Network	Description	Applications	Network Image
FNNs	Each perceptron (simplest and oldest form of neurons) in one layer is connected to every perceptron from the next layer. Information is fed forward from one layer to the next in the forward direction only. There are no feedback loops. Thus, the data is processed, and the results are calculated on every input sequence. This network may or may not have hidden layers.	Primarily used for animal recognition, digit recognition, cheque recognition, medical diagnosis, etc.	
RNNs	Use sequential information such as time-stamped data from a sensor device or a spoken sentence, composed of a sequence of terms. Unlike FNNs, inputs to RNNs are not independent of each other, and the output for each element depends on the computations of the preceding elements.	Primarily used in forecasting and time series applications, sentiment analysis and other text applications.	
Long Short-Term Memory (LSTM)	A type of RNN that is explicitly designed to hold information for long periods of time and process the incoming data, along with the previously calculated results. LSTMs contain their information in a memory and can read, write and delete information from its memory.	Primarily used for text classification, machine translation, dialog systems, speech recognition, translating videos and images to natural languages, etc.	
CNNs	Typically contain five types of layers: input, convolution, pooling, fully connected and output (more recent versions tend to be deep with more than seven or nine layers). Each layer has a specific purpose, like summarizing, connecting or activating.	Primarily used for image classification and object detection. Other applications include language processing, computer vision and video analytics.	

Source: Medium, SAS, Pitt Street Research

Supervised Learning versus Unsupervised Learning

Since the advent of Machine Learning, different algorithms or methods have been developed to process both structured and unstructured data. However, all Machine Learning methods can be broadly classified into either supervised learning or unsupervised learning (Figure 24), though supervised learning is the most commonly used form of Machine Learning.

With supervised learning, each input fed to the system is labelled with a desired output value. A supervised learning algorithm analyses the data and compares its actual output with desired output to find errors and modify the model accordingly. Supervised learning is commonly used in applications where future events are predicted based on historical data, e.g. determining fraudulent credit card transactions and predicting insurance customers likely to file claims.

In unsupervised learning, the training set submitted as input to the system is not labelled with the historical data or a desired outcome. In simple words, unsupervised learning is used against data that has no historical labels. Therefore, the system itself develops and structures the data, identifies common characteristics, and modifies it based on knowledge gained during the process.

This form of Machine Learning is commonly used to segment customers with similar attributes who can then be treated similarly in marketing campaigns. It can also identify the main attributes that separate customer segments from each other. Other applications include segmentation of text topics, image recognition, pattern recognition in financial markets data, identification of data outliers, sound analysis, e.g. to detect anomalies and potential problems in jet engines etc.

Input dataset is labelled in supervised learning while in unsupervised learning the system structures the data

Figure 24: Supervised Learning Vs. Unsupervised Learning

Parameter	Supervised Learning	Unsupervised Learning
Type of Input Data	Labeled	Unlabeled
Degree of Computational Complexity	High	Low
Accuracy of Results	High	Low to moderate
Timeliness of Analysis	Off-line	Real time
Commonly Used Algorithms	Random Forests, Linear Regression, Decision Trees, Naïve Bayes, Support Vector Machines, Neural Networks	Clustering (K-means, SVD, PCA, etc.), Association Analysis (Apriori, FP-Growth), Hidden Markov Model
Key Use Cases	Prediction and classification	Grouping and data interpretation

Source: Pitt Street Research

Convolutional Neural Networks are widely used today

CNNs are among the most widely used ANNs today given that they can learn unsupervised and require relatively little pre-processing. CNNs are used in a range of areas, including statistics, natural language processing as well as in signal and image processing, e.g. for medical image analysis.

However, CNNs are rather impractical for visual imagery classification given the large data sets that need to be processed, which consumes enormous



amounts of energy while CNNs are relatively slow. With the advent of autonomous vehicles and the stringent requirements on image recognition capabilities by Advanced Driver Assistance Systems (ADAS) in cars, today's CNNs may not be the best solution.

Pros and Cons of Machine Learning and Deep Learning

In summary, Machine Learning and Deep Learning have many applications, and organizations use these applications to drive automation for specific tasks and processes, e.g. to save cost, bring products to market faster, improve operational efficiencies, prevent fraud, gain new insights into data and enable new technologies to be deployed faster. Homeland Security (HLS) and law enforcement are other application areas for AI.

While Machine Learning supplements data mining, assists decision making and enables development of autonomous computers and software programs, Deep Learning, on the other hand, performs complex computations and is widely used for difficult problems that require real-time analysis, such as speech and object recognition, language translation and fraud detection.

However, these AI technologies do have their own limitations. Both Machine Learning and Deep Learning are susceptible to errors and whenever they make errors, diagnosing and correcting them can be difficult. In addition, it is impossible to make immediate accurate predictions with these technologies as they require substantial computational power and can be difficult to deploy, especially in real time.

Furthermore, the outcomes generated by these technologies are prone to hidden and unintentional biases, including racial biases, depending on the data provided to train them. Also, these technologies cannot always provide rational reasons for a prediction or decision.

Nevertheless, the utilization of Machine Learning and Deep Learning is anticipated to rise substantially as the potential of neural networks to solve problems, make predictions and improve decision-making are unparalleled.

The next iteration in neural networks is the emergence of Spiking Neural Networks that have multiple advantages over CNN's, including speed and power consumption, as we will elaborate on below. BRN is at the very forefront of Spiking Neural Network development and future commercialisation.

Appendix II – Neuromorphic computing 101

In our discussion of CNN's and ANN's so far we haven't specifically mentioned that these AI models are purely software-based. Given that these algorithms are so large and generally can't be executed locally, most queries such as Google Assistant and Siri queries on a mobile phone, need to be sent to the Cloud to be processed. The results then need to be sent back to the device. This takes time and processing such queries in the Cloud consumes tremendous amounts of energy.

However, in our view the most significant drawback of software-only neural networks is that the algorithms are designed by humans, i.e. software engineers, and hence the scope of the neural network is limited to the imagination of whoever designed the particular algorithm. We touched on this briefly when we discussed the differences between supervised and unsupervised learning, i.e. unsupervised learning allows the network to learn without any restrictions in what to look for.

Machine Learning and Deep Learning have various applications such as in defence and law enforcement

Cons of Machine Learning and Deep Learning include susceptibility to errors and unintentional biases




Therefore, we believe the logical evolution of neural networks is a hardware-only solution that allows for unsupervised learning. Enter Neuromorphic Computing.

Neuromorphic chips mimic the biological brain in a hardware implementation

In simple terms, a neuromorphic chip tries to emulate the structure and functions of neurons found in nature, for instance in the human brain. But rather than using software, the artificial neurons are hardwired in digital chips. The human brain consists of roughly 86 billion interconnected neurons (Figure 25) that send and receive electric impulses, or spikes, to and from neighbouring neurons.

Figure 25: The human brain comprises billions of neurons




Source: Penn State, Pitt Street Research

A single neuron system comprises of the actual neuron, axons, dendrites and synapses. As we alluded to earlier, electric impulses are sent from one neuron to the next through connections known as axons. Before reaching the next neuron, though, each impulse arrives at the next neuron's synapse first (Figure 26), which can be seen as a gatekeeper. The human brain has approximately 150TR synapses (trillion = 1,000 billion).




Figure 26: Synapses in between two neurons



Source: whfreeman.com, Pitt Street Research

Synapses decide whether or not to pass on a signal to the next neuron depending on the weight stored in the synapses (Figure 27). The electrical impulse must reach a certain intensity threshold in order to be considered relevant and be allowed to pass on to the next neuron. As a particular input passes through the network of synapses, this network acts as a filter, activating some neurons while others remain inactive.

Figure 27: Impulses must exceed a certain threshold to be propagated



Source: BC Campus, Pitt Street Research



The threshold function of the synapse plays a key role in neuromorphic chips

The key function in the process above is performed by the synapse that decides which impulses to propagate and which impulses to terminate. That decision is based on the synaptic weight, which attributes a value to each incoming impulse upon which the decision whether or not to propagate the impulse is based.

This threshold function is equally critical in biological neurons as it is in hardwired artificial neurons as we will elaborate on below.

Appendix III – BrainChip’s Solution for Neuromorphic Computing

BRN has developed a hardware version of the biological neuron


Company’s development work in the last ten years has focused on creating a hardware implementation of the biological neuron described above, resulting in the design of a neuromorphic chip called Akida, which is Greek for spike. The cell architecture and cell behaviour are similar to that of a biological neuron.

The architecture is such that incoming impulses, also known as spikes, are received by artificial synapses, which can autonomously decide whether or not to propagate the spike (Figure 28), based on how strong, or relevant, the spike is.

Similar to biological synapses, artificial synapses work with a threshold function, i.e. the sum of impulses received from multiple, connected, neurons must reach a certain threshold level in order to be considered relevant and to be fed forward to other neurons in the system.

Artificial synapses work like biological synapses in terms of threshold function

Figure 28: Neuron structure in a Spiking Neural Network



Source: AGH University of Science and Technology, Pitt Street Research



On the scale of the entire chip, BRN has a fabric of Neural Processing Cores (NPCs), where each NPC emulates 1,000s of the biologically-inspired neurons. These NPCs are connected in a network, so that they can emulate the multiple layers similar to the structure of the artificial neural networks we described previously (see Figure 22). This enables spikes to move through multiple layers of neurons before an output spike is generated. This type of neural network is known as a Spiking Neural Network (SNN).

Spike-Timing-Dependent-Plasticity: Learning by weights

The way in which an individual neuron in an SNN “learns” is by increasing or decreasing the weight attributed to each of its artificial synapses, depending on the relevance of the spikes the neuron received through each synapse. This relevance depends on how often and how many spikes are received. This process is known as Spike Timing Dependent Plasticity (STDP), i.e. in addition to a neuron actually firing, the frequency of spikes and the actual number of spikes fired is also relevant information. Based on this information, artificial synapses will start to favour certain connections, while they may inhibit others over time.

BRN has significantly expanded upon the base STDP learning rule for its technology. The use of STDP, the learning method of the brain, is unique in Artificial Neural Networks. Most ANNs use Deep Learning, which is a successive approximation method. STDP leads to much faster results and needs much less data than Deep Learning. The reason that CNNs use deep Learning is because they compute with numbers rather than spikes. BRN delivers a tool that can convert standard CNNs to fast-learning, low power spiking CNNs for execution on Akida, called CNN2SNN. This tool is free and delivered with MetaTF.

The memory of each neuron, i.e. what it has “learned”, is embedded in the weight of its synapses, the so-called synaptic weights. Synaptic weights are dynamic and can change over time based on the relevance of the spikes it receives from preceding neurons in the network.


Spiking Neural Networks hold advantages over CNNs

As opposed to traditional, software-based CNNs, that use integer data types (ones and zeros) as inputs, SNNs process spikes (Figure 29). Spikes are simply an encoding method of real-world sensory data points, such as sound and vision. A sequence of spikes is known as a spike train. Spike trains offer the ability to process this real-world sensory data in the same way the human brain does.

Spiking Neural Networks are much faster and use substantially less power



Figure 29: Akida's sequential processing of spikes



Source: Company, Pitt Street Research

Given that Akida's neurons receive incoming spikes from multiple other neurons at the same time, Akida can process incoming information in parallel, i.e. many spike trains are being processed by Akida simultaneously. This is a key distinction from software-based CNNs, which process data sequentially, i.e. one mathematical calculation after another.

Additionally, given that most of the memory of an SNN can be found in the synaptic weights, SNNs don't require much access to external memory, such as DRAM (Dynamic Random-Access Memory) to retrieve information on the weight of a particular synapse or to temporarily store the outcome of a calculation.

Parallel processing and "on-board" memory inside the synapses provide Akida with a tremendous speed advantage compared to traditional CNNs. For instance, training a neural network for a specific task, such as image recognition, which might take an SNN several hours, could take a CNN several days or even longer.

In addition to speed advantages, the fact that SNNs don't need to move very much data from the processor to the memory and back with each calculation also results in SNNs using substantially less power compared to CNNs. The latter require very substantial traffic of data between processing units and memory units.

Due to less and faster processing and less data traffic to and from memory units, power consumption of SNNs can be up to 95% lower compared to CNNs. This makes SNNs ideally suited for IoT edge applications that are not permanently connected to a power source, such as mobile phones, Electric Vehicles and sensors. But energy consumption is also extremely important in large scale data centers. In other words, we believe hardware-only SNN's have several key competitive advantages compared to CNNs.

Akida is expected to help take AI to the next level

Currently, AI processing is centralised, with connected devices relying on data center services, such as Amazon Web Services (AWS), to carry out

Parallel processing and "on-board" memory provide speed advantage over traditional CNNs

SNNs consume significantly less power as compared to CNNs



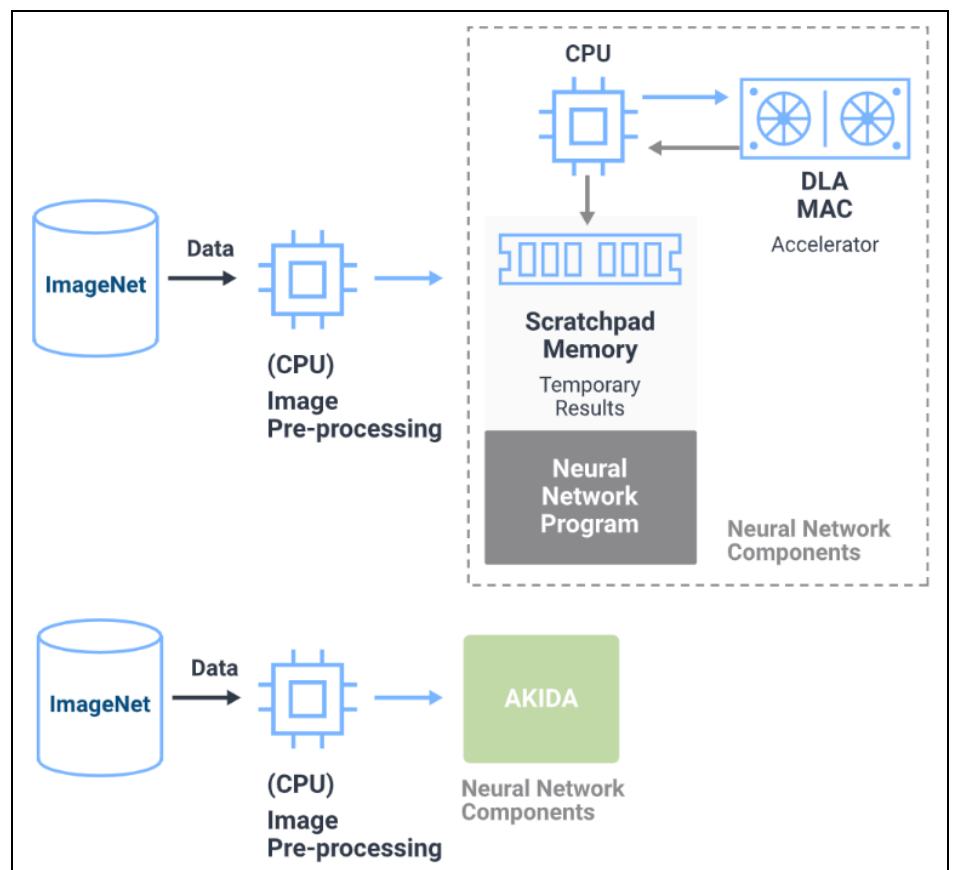
calculations. However, Akida has taken a different approach by decentralising AI processing and carrying out computations at the edge of the connected devices.

This provides various advantages such as data privacy and reducing the dependency on internet bandwidth. Given that in current AI solutions, such as Siri and the Google Assistant, the data has to be first transferred to the cloud for processing at the data center, it makes the personal information in the dataset vulnerable to cybercrimes.

Furthermore, Akida is a fully-integrated, purpose-built neural processor (Figure 30). Traditional neural processing solutions require a CPU for running the neural network algorithm, a deep learning accelerator such as a GPU for carrying out multiplication and addition calculations (known as MACs), and memory for storing network parameters.

However, by integrating these three functions into a single neural processor, can reduce energy consumption needed in traditional solutions to interact with the three components. Moreover, the consolidation also helps Akida reduce its size.

Figure 30: Traditional neural processing solution vs Akida



Source: Company



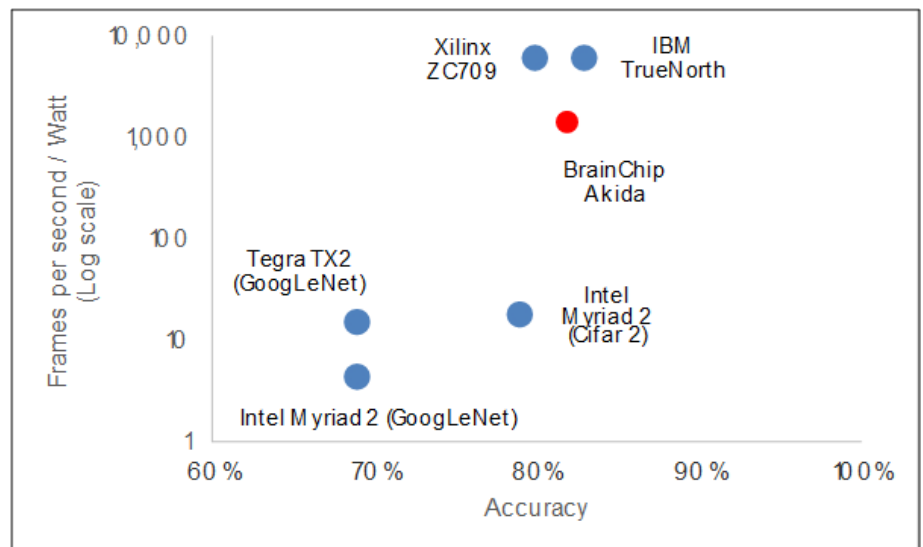
Akida stacks up very favourably compared to other technologies

In addition to Akida, there are a number of other neuromorphic technologies being developed. The most comparable in terms of performance are Intel's Loihi and IBM's TrueNorth NSoCs. Intel's research test chip Loihi has 128 neuromorphic cores and has been manufactured at a resolution of 14 nanometers (nm). It has 131,072 neurons and 130 million synapses. IBM's TrueNorth NSoC has just over 1 million neurons and 268 million synapses and has been used in configurations of 64 chips resulting in a massive 16 billion combined synapses. As mentioned earlier, Akida has 1.2 million neurons and 10 billion synapses.

In comparing the performance of neural networks, standardized datasets such as CIFAR-10 are typically used. CIFAR-10 is a collection of 60,000 images in 10 different classes, including cars, dogs, ships, birds, airplanes etc., and is specifically used to train neural networks in object recognition.

In processing CIFAR-10 datasets, TrueNorth has shown to be less accurate than Akida (89% versus 82%) in recognizing objects in pictures while being able to process 6,000 frames per second per watt. This compares to 1,000 frames per second per watt for Akida (Figure 31). No data for Intel's Loihi is available.

Figure 31: Akida compared to other technologies



Source: Company, Pitt Street Research

However, at a price of US\$8,000 per chip, TrueNorth is over 500x more expensive than the targeted commercial price point for Akida of just ~US\$15. We believe this low price point will be very beneficial to BRN as it comes to driving adoption of the Akida technology.

As shown in Figure 32, Akida holds several other advantages over other competitive offerings in the market. It is compatible with the TensorFlow development environment, which is expected to drive adoption as most data scientists are already familiar with it and hence do not need to learn a new programming language to use Akida.



Figure 32: Akida is well-positioned against competition

	Micro- to Mw Power use	Real-time on-chip learning & training	TensorFlow Compatible	Stand-alone possible (No CPU required)	On-chip Convolution	Available as IP	Green Technology
BrainChip Akida AKD1000	✓	✓	✓	✓	✓	✓	✓
IBM TrueNorth	✓	NONE	LEARN COREL	NO	NO	NO	✓
Intel Loihi	✓	PROGRAM	LEARN NEF	✓ <small>copack</small>	NO	NO	✓
Google Coral TPU	2-5W	Math chip	✓	NO	NO	NO	NO
DLAs (Nvidia, others)	>5-10W	Math chip	✓	NO	NO	NO	NO

Source: Company, Pitt Street Research

Moreover, the Akida chip has the capacity to work as a standalone embedded accelerator or as a co-processor. This is facilitated by sensor interfaces on the chip for image, audio, and analogue signals, as well as high-speed data interfaces such as PCI-Express and USB. The chip also includes data-to-spike convertors which aid in converting traditional data formats into spikes, so that it can be used to train and can be processed by the Akida Neuron Fabric.

Akida also offers on-chip convolution which helps in transferring existing Deep Learning networks to the AKD1000. Akida is also the only technology available as an IP to the customers, which we believe should help drive adoption. Furthermore, it provides the company an opportunity for early commercialisation of its technology and another revenue stream in the form of licensing fees.


What does the road ahead look like for Akida?

AKD1000 is only the beginning of what the company expects to be a robust product portfolio (Figure 33):

- **AKD500:** BRN is in the planning stages for AKD500, which is a low-cost version of the AKD1000, for the consumer products market.
- **AKD1500:** BRN is also in the process of developing and prototyping AKD1500, which will have additional features to execute Long Short-Term Memory (LSTM) and Transformer networks. LSTM is a form of recurrent neural network which can learn order dependence in sequence prediction problems. It finds applications in machine translation and speech recognition. Transformer networks, on the other hand, find their applications in natural language processing (NLP) due to their ability to resolve the vanishing gradient problem – the time for which the memory is retained by the network, which is critical in order to process lengthy texts.
- **AKD2000:** It is the optimised version of AKD1500 and the company has already started work on a prototype at its lab in Perth.



Figure 33: Product pipeline for Akida



Source: Company, Pitt Street Research

- **AKD3000:** Further down the road, management plans to integrate even more functions of the brain into Akida through the AKD3000. This version will target capsule networks, which helps solve the spatial problems inherent in current CNNs. While CNNs are only focussed on the components, for example features in a human face, the capsules are able to encode spatial information, which helps in detecting the pose of the face. These networks help in image rendering and are currently at the research stage at Google.
Additionally, AKD3000 is also expected to target recurrent cortical neural networks, which draw inspiration from cortical networks in the brain.



Appendix IV – Analyst Qualifications

Marc Kennis, lead analyst on this report, has been covering the semiconductor sector as an analyst since 1997.

- Marc obtained an MSc. in Economics from Tilburg University (the Netherlands) in 1996 and a post graduate degree in investment analysis in 2001.
- Since 1996, he has worked for a variety of brokers and banks in the Netherlands, including ING and Rabobank, where his main focus has been on the technology sector, including the semiconductor sector.
- After moving to Sydney in 2014, he worked for several Sydney-based brokers before setting up TMT Analytics Pty Ltd, an issuer-sponsored equities research firm.
- In July 2016, with Stuart Roberts, Marc co-founded Pitt Street Research Pty Ltd, which provides issuer-sponsored research on ASX-listed companies across the entire market, including technology companies.

General advice warning, Disclaimer & Disclosures

Terms & Conditions

The information contained herein ("Content") has been prepared and issued by Pitt Street Research Pty Ltd ACN 626365615 ("Pitt Street Research"), an Authorised Representative (no: 1265112) of BR Securities Australia Pty Ltd. ABN 92 168 734 530, AFSL 456663. All intellectual property relating to the Content vests with Pitt Street Research unless otherwise noted.

Disclaimer

Pitt Street Research provides this financial advice as an honest and reasonable opinion held at a point in time about an investment's risk profile and merit and the information is provided by the Pitt Street Research in good faith. The views of the adviser(s) do not necessarily reflect the views of the AFS Licensee. Pitt Street Research has no obligation to update the opinion unless Pitt Street Research is currently contracted to provide such an updated opinion. Pitt Street Research does not warrant the accuracy of any information it sources from others. All statements as to future matters are not guaranteed to be accurate and any statements as to past performance do not represent future performance.

Assessment of risk can be subjective. Portfolios of equity investments need to be well diversified and the risk appropriate for the investor. Equity investments in a listed or unlisted company yet to achieve a profit or with an equity value less than \$50 million should collectively be a small component of an individual investor's equity portfolio, with smaller individual investment sizes than otherwise. Investors are responsible for their own investment decisions, unless a contract stipulates otherwise.

Pitt Street Research does not stand behind the capital value or performance of any investment. Subject to any terms implied by law and which cannot be excluded, Pitt Street Research shall not be liable for any errors, omissions, defects or misrepresentations in the information (including by reasons of negligence, negligent misstatement or otherwise) or for any loss or damage (whether direct or indirect) suffered by persons who use or rely on the information. If any law prohibits the exclusion of such liability, Pitt Street Research limits its liability to the re-supply of the Information, provided that such limitation is permitted by law and is fair and reasonable.

General advice warning

The Content has been prepared for general information purposes only and is not (and cannot be construed or relied upon as) personal advice nor as an offer to buy/sell/subscribe to any of the financial products mentioned herein. No investment objectives, financial circumstances or needs of any individual have been taken into consideration in the preparation of the Content.

Financial products are complex, entail risk of loss, may rise and fall, and are impacted by a range of market and economic factors, and you should always obtain professional advice to ensure trading or investing in such products is suitable for your circumstances, and ensure you obtain, read and understand any applicable offer document.

Disclosures

Pitt Street Research has been commissioned to prepare the Content. From time to time, Pitt Street Research representatives or associates may hold interests, transact or hold directorships in, or perform paid services for, companies mentioned herein. Pitt Street Research and its associates, officers, directors and employees, may, from time to time hold securities in the companies referred to herein and may trade in those securities as principal, and in a manner which may be contrary to recommendations mentioned in this document.

Pitt Street Research receives fees from the company referred to in this document, for research services and other financial services or advice we may provide to that company. The analyst has received assistance from the company in preparing this document. The company has provided the analyst with communication with senior management and information on the company and industry. As part of due diligence, the analyst has independently and critically reviewed the assistance and information provided by the company to form the opinions expressed in the report. Diligent care has been taken by the analyst to maintain an honest and fair objectivity in writing this report and making the recommendation. Where Pitt Street Research has been commissioned to prepare Content and receives fees for its preparation, please note that NO part of the fee, compensation or employee remuneration paid will either directly or indirectly impact the Content provided.