



# A<sup>2</sup>I<sup>2</sup> Chooses WEKA for AI Research Data

*Deakin University Applied Artificial Intelligence Institute deploys WEKA storage solution from XENON to stay in front of increasing performance demands from AI Researchers.*

## FOCUS ON HUMAN CENTRIC AI

The Deakin Applied Artificial Intelligence Institute (A<sup>2</sup>I<sup>2</sup>) opened in 2019 merging two groups at the University – the Pattern Recognition and Data Analytics team (PRaDA) and the Deakin Software and Technology Innovation Laboratory (DSTIL). The merger combined machine learning expertise with the ability to deliver complex ideas and systems in user friendly software, web and mobile applications.

A<sup>2</sup>I<sup>2</sup> takes a multi-disciplinary approach to projects, and works on “human in the loop” problems where AI solutions enhance uniquely human capabilities. Partnering with government and industry, A<sup>2</sup>I<sup>2</sup> works

across diverse sectors including health, defence, education, finance, manufacturing and security.

To drive the machine learning processes of data analytics, training and inference, the Institute operates a state of the art high performance computing (HPC) cluster, which includes the latest NVIDIA DGX System as well as other compute nodes. There are over 80 researchers at A<sup>2</sup>I<sup>2</sup>, and at any given time over 50 of these are active users of the HPC cluster. Robert Ruge, Systems and Network Engineer, and Josh Cole, AI Systems Officer, shared their story of meeting the Institute’s increasing need for faster access to data.





## COMPUTER VISION DEMANDS HIGHER STORAGE PERFORMANCE

The HPC cluster was originally established with 1.7PB of storage deployed with a Linux based parallel file system across eight storage servers. This was set-up with volume mirroring, so effective working space was 850TB. Since that was commissioned, a number of researchers had moved into computer vision based AI, which was placing a heavy load of input/output processing (IOPS) on the storage. This was having a knock-on effect across the Institute, with the heavy IOPS load effecting all jobs which were running slower and slower.

The team started evaluating alternative storage solutions. Robert and Josh took a broad look at the current high performance storage offerings from major vendors

and new emerging vendors, looking for “highest IOPS we could get for our budget”, supportability, modern design, and minimizing complexity were key considerations. The team settled on WEKA, which was able to offer the highest IOPS for their budget, and, as Robert pointed out, “WEKA runs on commodity hardware we could buy from anyone, and we could grow it incrementally as we have funds – which is what we’ve done ... allowing us to keep one step ahead of the researchers.” WEKA also stood out for ease of operation, support, and Robert noted “WEKA was built with new technology in mind. NVME based, it doesn’t carry any baggage from the hard disk era. That was also an advantage.”

▲The Virtual Dementia Experience (VDE) is an immersive, emotive and interactive experience that aims to capture and simulate the experience of living with dementia.

## MOVE ENCOURAGES ARCHIVING

The initial roll-out was 436TB raw NVME, across ten servers. Taking a cautious approach, the original plan was to stand up the WEKA storage and use it in parallel with the existing Linux clustered storage. The team took advantage of the move to the new storage to spring clean the existing data with the researchers, who were able to archive more than half their data. “As we did our clean-up and pulled our accounts across, we soon realized we could squeeze it all in the new system. Which was good for us, we didn’t have to run two stacks,” explained Robert.

The old Linux system had also implemented mirroring and included a lot of overheads from the hard drive based architecture. As a consequence, Robert explained “monthly back-ups were taking 35 days, so we couldn’t guarantee we would have a good back-up. Since we transitioned to WEKA the back-up is now taking 6 days!” Josh added, “the team is now implementing weekly back-ups on top of the monthlies”, which weren’t possible before.

WEKA delivers

**10x**  
improvements for  
A<sup>2</sup>I<sup>2</sup> researchers  
in AI epochs.

WEKA  
demonstrates  
best performance  
per dollar, with  
**100x**  
faster IOPS.



## BENCHMARKING PERFORMANCE

With the new WEKA system in place, Josh ran a series of benchmarks to compare it to the old system. Josh found “a single host on the new system could get much higher IOPS than the entire cluster on the old system ... and that’s including using NFS, UDP” which are the slowest WEKA protocols.

Josh provided the benchmarks below on the original install with ten servers.

Commenting on the benchmarks, Robert noted that “we’ve not reached the limits of the storage system,” and Josh added, “the performance we’re seeing with the current storage is much higher than the old storage system could have handled. But we still have plenty of space to spare, and we’ve got plenty of overhead for them to play with. So, storage is no longer a bottleneck which is good.”

▲ Enroute Trauma Management Project: Helping paramedics make the right decisions to save lives.

### WEKATESTER RESULTS

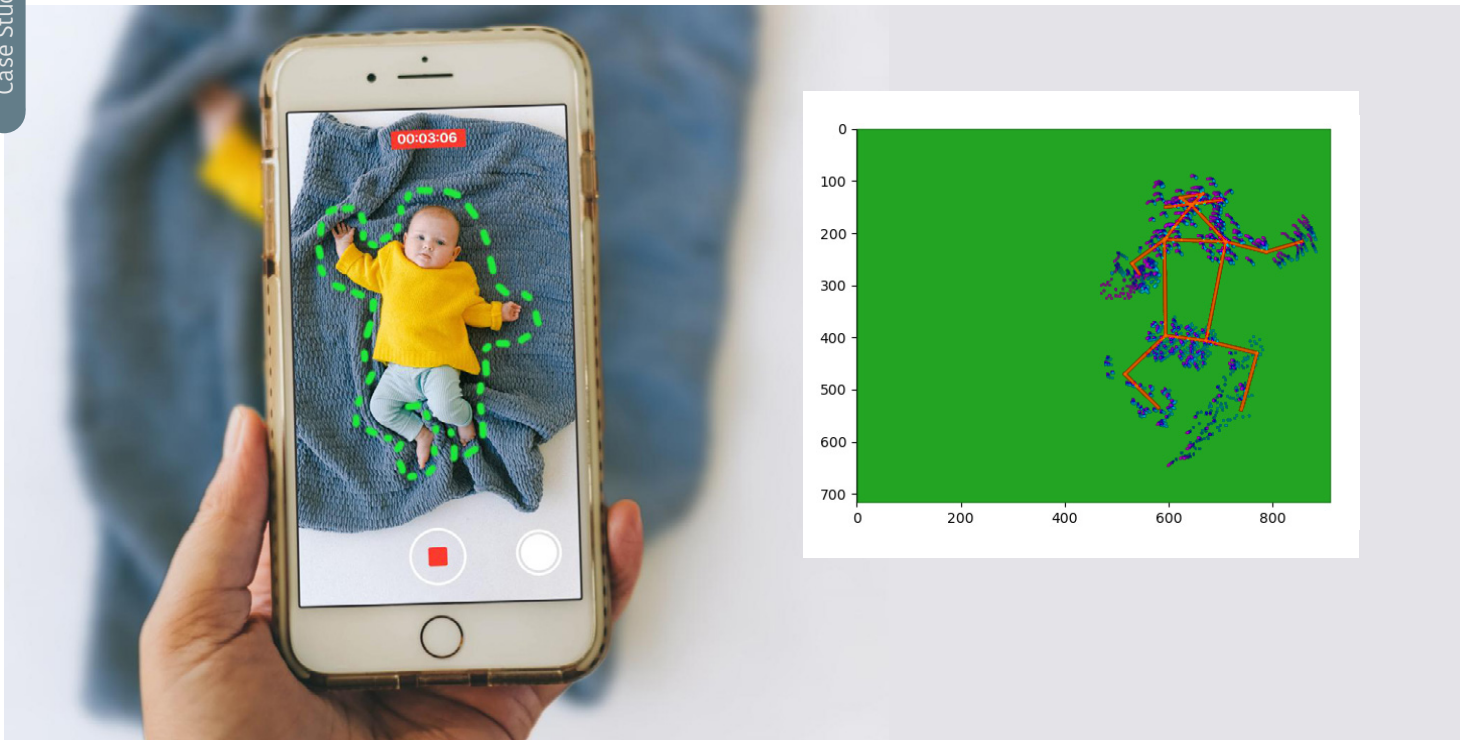
TEST	READ BANDWIDTH		WRITE BANDWIDTH		READ IOPS		WRITE IOPS	
	TOTAL	HOST	TOTAL	HOST	TOTAL	HOST	TOTAL	HOST
<b>1CORE</b>	20.40 GiB/s	1.70 GiB/s	24.67 GiB/s	2.06 GiB/s	1,418,104/s	118,175/s	807,163/s	67,263/s
<b>2CORE</b>	37.49 GiB/s	3.12 GiB/s	42.22 GiB/s	3.52 GiB/s	2,983,072/s	248,589/s	806,883/s	67,240/s
<b>NFS</b>	17.00 GiB/s	1.42 GiB/s	8.69 GiB/s	741.78 MiB/s	181,209/s	15,100/s	282,161/s	23,513/s
<b>UDP</b>	6.52 GiB/s	556.33 MiB/s	4.41 GiB/s	376.20 MiB/s	687,806/s	57,317/s	513,190/s	42,765/s
<b>LEGACY LINUX</b>	5.57 GiB/s	475.15 MiB/s	2.27 GiB/s	193.53 MiB/s	18,622/s	1,551/s	6,550/s	545/s

The experience of the researchers matches the benchmark numbers, with Josh providing the following quotes from the A<sup>2</sup>I<sup>2</sup> team:

*“I would like to inform you that my code is running very fast on the server. In particular, I usually need from 40 minutes to 1 hour to complete one epoch, but now it only takes around 6 minutes per epoch.”*

*“Thanks for implementing the new WEKA file system. Performance is very much improved.”*

*“WEKA is great! It is now pretty much instantaneous to start running jobs, whereas before I would sometimes wait 10 minutes with jobs blocked in a disk wait state.”*



## WEKA PROVES A SOLID FOUNDATION FOR A<sup>2</sup>I<sup>2</sup>

WEKA has been a game changer for A<sup>2</sup>I<sup>2</sup>. Robert noted, “price compared to performance compared to other solutions we looked at ... it definitely delivered the performance we needed and we still have headroom to push the system as researchers get more sophisticated. And in the less than 12 months that it has been in, we’ve managed to upgrade it with extra money that’s become available.” This was a key capability that stood out, allowing the team to “expand incrementally as we have funds available, and doing so increase the performance and keep ahead of the researchers.”

Robert reflected on the early part of the decision, and noted that “one concern we had, [WEKA] was a new company to us, a small company, and fairly new on the global market, and that was quite a concern for us but we decided to take a punt and so far so good and it’s been great to see the support roll out in Australia, and staff being employed”, with WEKA growing

from one staff in APJ to twenty now. “The fact that XENON was the implementation partner, and stood by WEKA meant a lot to us as well, as we’ve always had successful engagement with the XENON team.”

Josh highlighted the support, and admin interface that makes his job easier, “with WEKA the user interface has a lot more information that I can make use of and having the support is nice and just better documentation,” and when there are issues, they’ve been able to access the WEKA developers directly and resolve issues on the first call. That combined with the local support team has made for a smooth roll-out and implementation across the Institute.

A<sup>2</sup>I<sup>2</sup> researchers are constantly pushing the boundaries in AI research, data, model size, so the way of the future includes further incremental expansion, faster storage performance, and continuing to “stay one step ahead of our researchers with WEKA”.

▲ Early Cerebral Palsy Screening with Deep Learning.

Learn more about the work of the A<sup>2</sup>I<sup>2</sup> team at [www.a2i2.deakin.edu.au](http://www.a2i2.deakin.edu.au)

*"Stay one step ahead...with WEKA."*

Projects, descriptions and images of A<sup>2</sup>I<sup>2</sup> work in this case study provided courtesy of A<sup>2</sup>I<sup>2</sup>. For more information on these projects please see their website.

