



# A Customer Centric Approach to Managing Big-data and Processing in Cryogenic Cryo- electron Microscopy

Dr James Bouwer , Dr Simon Brown, Jason Andrade, Josh Silver  
17<sup>th</sup> October 2023

A collaboration of cross-disciplinary approaches to extending our  
understanding of life science

MOLECULAR  
HORIZONS



UNIVERSITY  
OF WOLLONGONG  
AUSTRALIA

# U O W

## Building the world's fastest\* CryoEM processing platform (\*under \$100K)

Jason Andrade, Dr James Bouwer , Dr Simon Brown, Josh Silver  
17<sup>th</sup> October 2023

James wants the fastest thing we can build today that we can afford and he wants it ASAP because researchers are extending our understanding of life science and they want to do it quicker.

MOLECULAR  
HORIZONS



UNIVERSITY  
OF WOLLONGONG  
AUSTRALIA

# Executive summary

Find a good partner – you'll need more than one!

- Compute
- Storage
- Networking
- Environment (hosting)



Familiarize yourself with Excel – the data analytics tool of choice, in 2023

Be realistic in your negotiations

This is going to take 3+ months

Take some risks!



"For the greatest benefit to mankind"  
*Alfred Nobel*

## 2017 NOBEL PRIZE IN CHEMISTRY

Jacques Dubochet  
Joachim Frank  
Richard Henderson



2017 Chemistry Laureates. Ill: N. Elmehed.  
© Nobel Media 2017

### 2017 Nobel Prize in Chemistry

The Nobel Prize in Chemistry 2017 was awarded to Jacques Dubochet, Joachim Frank and Richard Henderson "for developing cryo-electron microscopy for the high-resolution structure determination of biomolecules in solution".

\* More about the 2017 Chemistry Prize



### The Nobel Prize in Chemistry

Awarded to 178  
Nobel Laureates  
since 1901

*"The said interest shall be divided into five equal parts, which shall be apportioned as follows: / - - / one part to the person who shall have made the most important chemical discovery or improvement ..."*

(Excerpt from the will of Alfred Nobel)

Chemistry was the most important science for Alfred Nobel's own work. The development of his inventions as well as the industrial processes he employed were based upon chemical knowledge. Chemistry was the second prize area that Nobel mentioned in his will.

The Nobel Prize in Chemistry is awarded by the Royal Swedish Academy of Sciences, Stockholm, Sweden.

### Most Popular Chemistry Laureates

1. Joachim Frank
2. Jacques Dubochet
3. Richard Henderson
4. Emil Fischer
5. Ernest Rutherford
6. Hans Fischer
7. Marie Curie
8. Irène Joliot-Curie
9. Mario J. Molina
10. Sir J. Fraser Stoddart




# Molecular Horizons



June 2020

# Scales in Biology

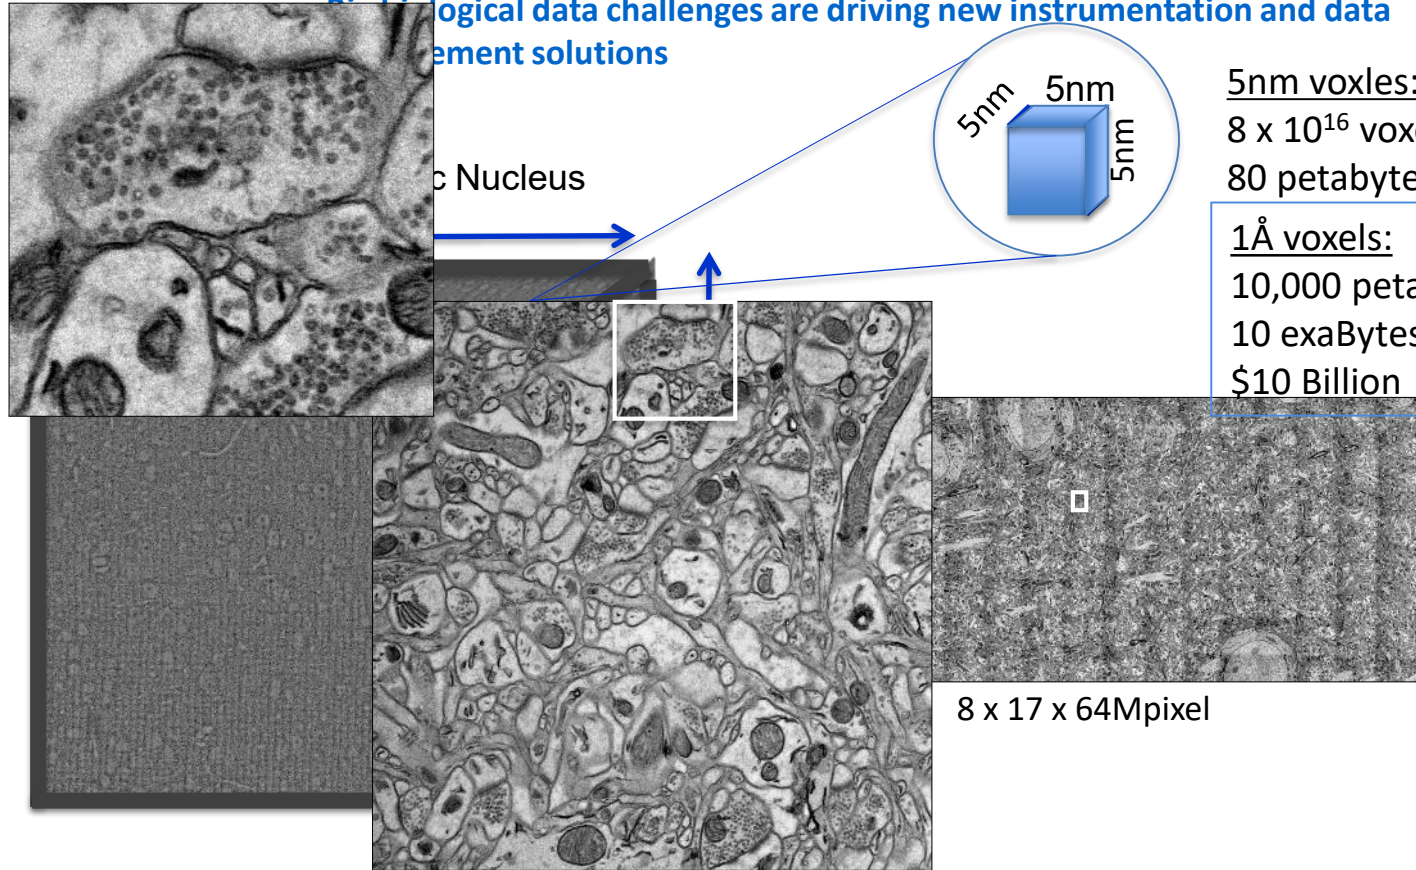
## Integrated Multi-scale Biology



Modified from slide of John A.G. Briggs of the EMBL

# Imaging Across Length Scales Spanning 6 Orders of Magnitude

Biological data challenges are driving new instrumentation and data management solutions

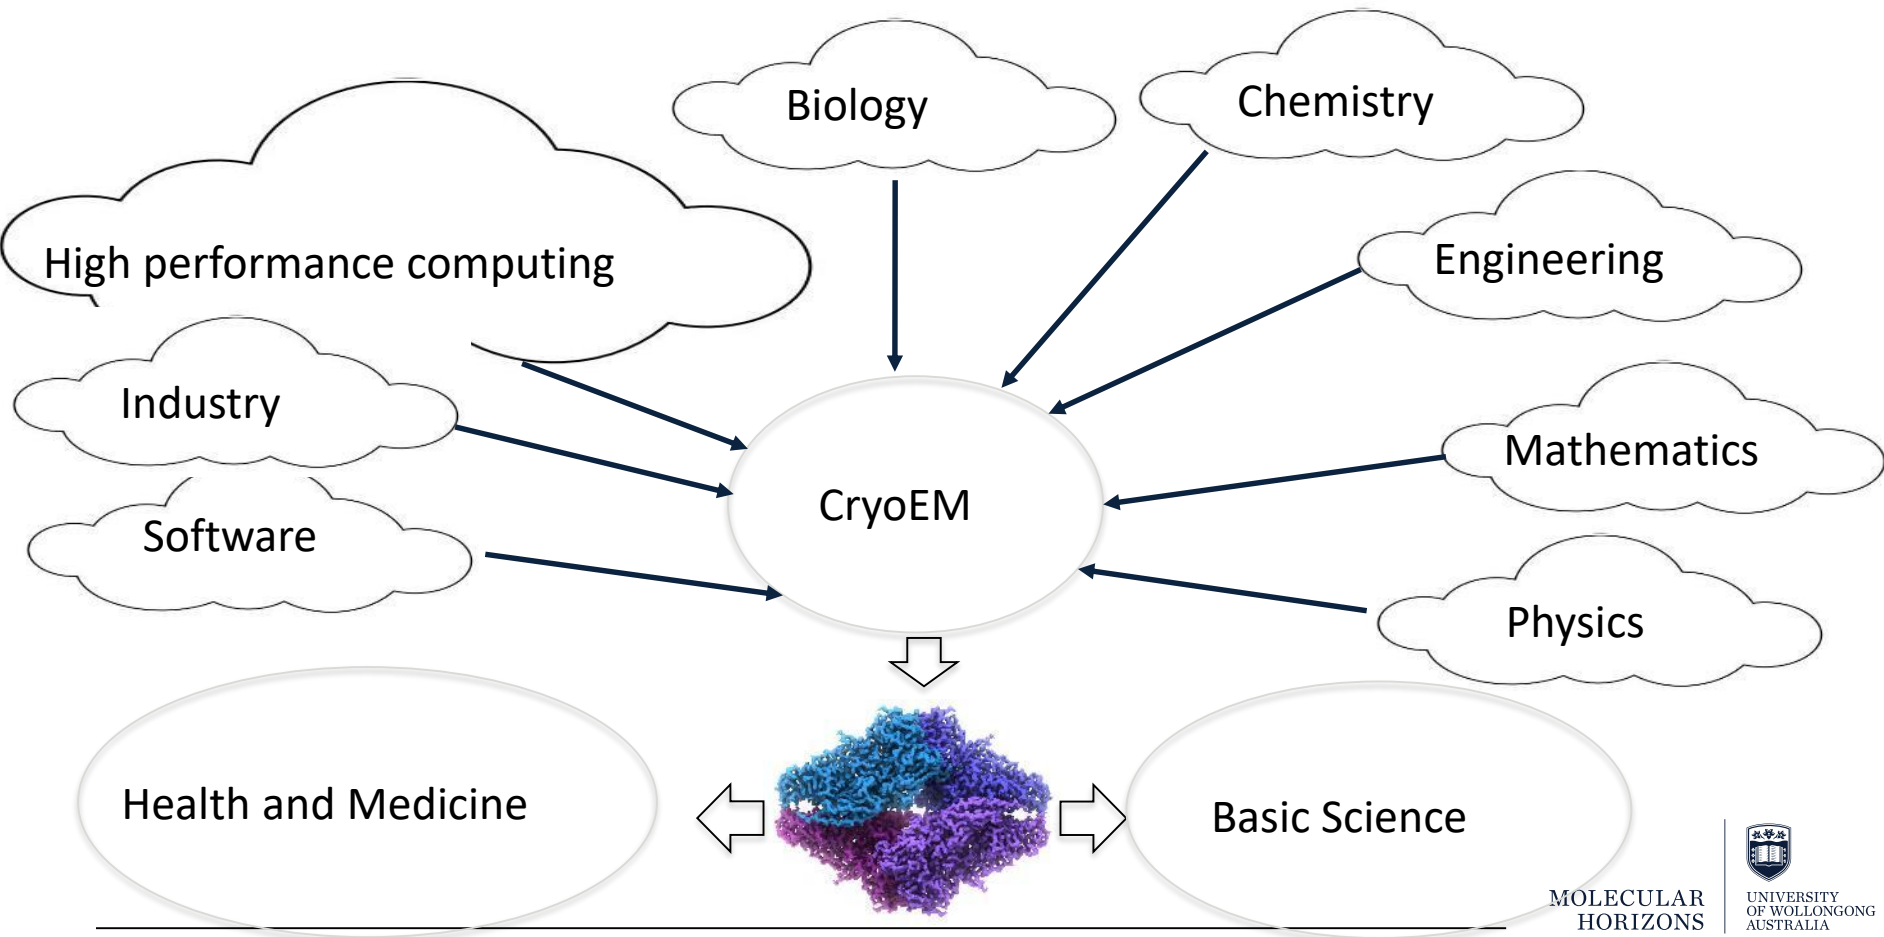


5nm voxles:  
8 x 10<sup>16</sup> voxels  
80 petabytes at 8bit

1Å voxels:  
10,000 petabytes  
10 exaBytes  
\$10 Billion


Ultra Wide-field High Resolution reconstructions require intelligent management of data and sophisticated instrumentation development

# Multi-disciplinary Approach to Collaborative Big Data Science and Workflows



# Molecular Horizons Suite of TEMs

A complete single particle workflow solution



- K2 and Falcon 3 Direct Detector
- AFIS upgrades
- Full Automation
- Structures to  $< 2.2 \text{ \AA}$
- Full Production Mode

- Falcon 3 Detector
- Energy Filter
- Gatan K3 Camera Upgrade in January 2022
- STEM equipped
- Full Automation
- Structures to  $1.2 \text{ \AA}$

- LaB6 Filament
- Training and Negative Stain
- 2k x 2k CMOS RIO4


- VitroBots
- Controlled Freezing Conditions

# Radiation Hardened CMOS Direct Detection Devices Produce Big Data

## Gatan K3 direct detection device:

- Electron counting devices
- 1500 fps to FPGA rack (16bit) - 72 GBytes/second
- 24 Mpixel at 40fps to the user 1.92 GBytes/second
- Typical Usage delivers 50 frames over a 2 second exposure at 500 images per hr = 25k frames / hr = 1.2TB per hr
- Tiff lossless compression = 120G/hr or 2.3TBytes/day


Typical Data Sets = 5-15 TBytes



# University of Wollongong CryoEM – An end-to-end workflow solution



Sample Prep



Training and  
Negative Stain




Intermediate  
Resolution




Atomic  
Resolution



Local HPC



Remote HPC +  
Local HPC



SP, Tomography, MicroED and Soft-  
materials solution



10 Gb/sec Networking

*B. subtilis* RNA polymerase elongation complex



EMD-21920, PDB:6WVJ

**For Access, google: "UOW CryoEM"**

<https://www.uow.edu.au/research-and-innovation/our-research/research-institutes-and-facilities/molecular-horizons/facilities/cryo-em/>

MOLECULAR  
HORIZONS



# Local GPU Compute Workstations

## XGPU1 (Real time processing):

6x NVidia 11GB RTX 2080 Ti

2x 28 Core Xeon's

## XGPU2 (Real time processing):

4x NVidia 11GB RTX 1080 Ti

2x 12 Core Xeon's

## V100 (Local jobs):

2x NVidia 32GB V100

2x 12 Core Xeon's

Desktop workstations



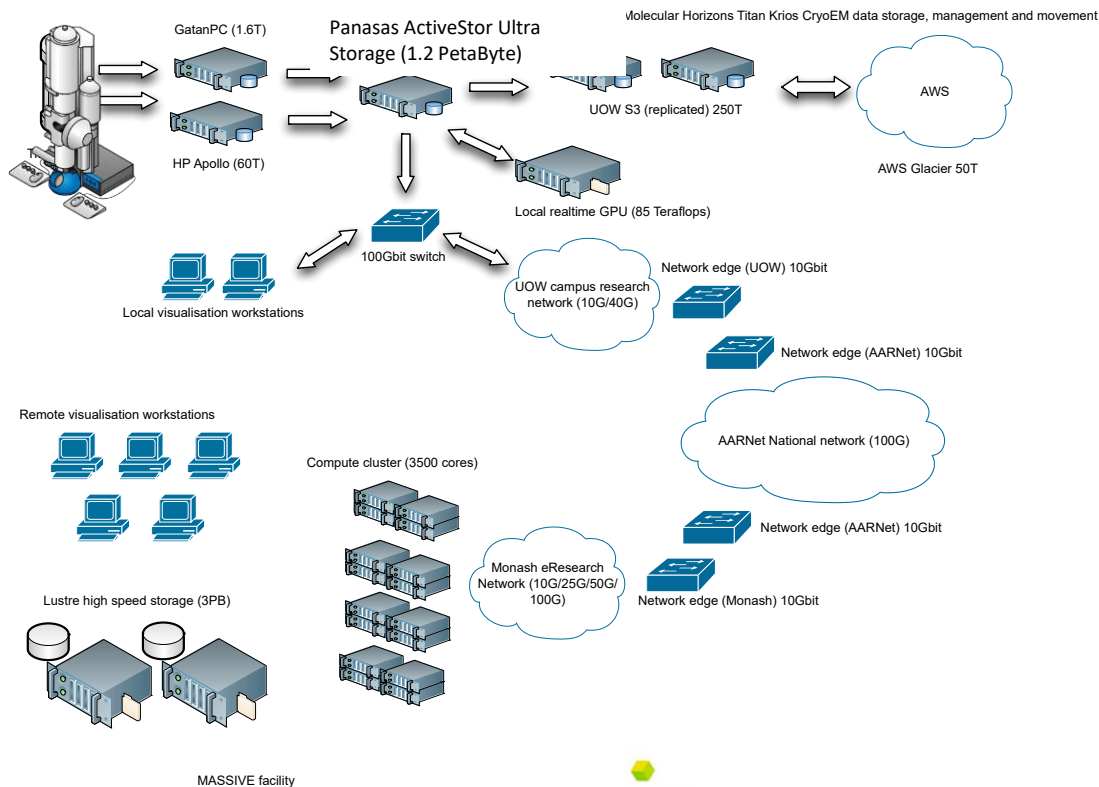
**XENON**<sup>®</sup>  
High Performance Computing

MOLECULAR  
HORIZONS



# Titan Krios and Arctica Network and Compute

## Networking Layout – UOW Research Network



**XENON**<sup>®</sup>  
High Performance Computing



**NVIDIA**



**MONASH**  
University



**MASSIVE**

**PANASAS**<sup>®</sup>

**MOLECULAR**  
**HORIZONS**



**UNIVERSITY**  
**OF WOLLONGONG**  
**AUSTRALIA**

# Panasas Partnership – Developing CryoEM data management and workflows

## Esky (Panasas ActiveStor Ultra):

3x Director Nodes

12x Storage Nodes

Capacity: 1.2PB

Data Network: 2x 25Gbit/sec per Storage Node

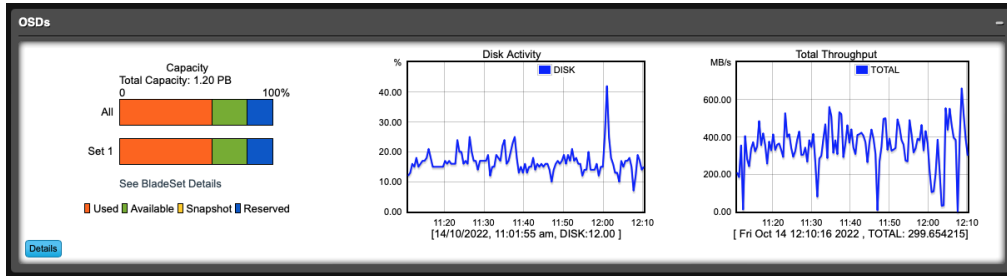
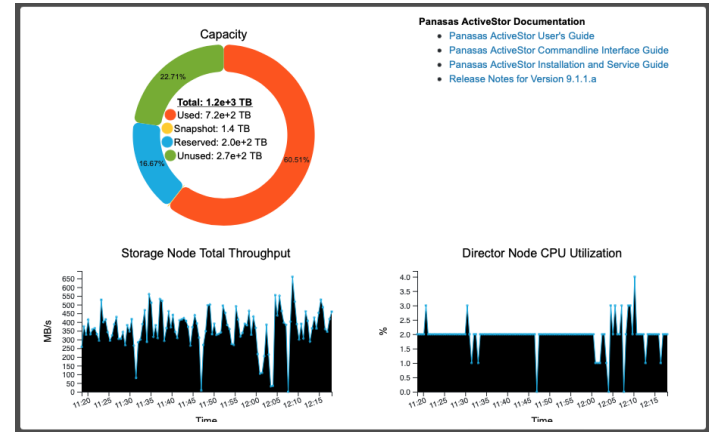
UOW Network: 2x 100Gbit/sec uplinks

- Develop CryoEM specific workflow solutions
- Optimize data transfer rates through DirectFlow
- Optimize CryoSparc and Relion cacheing
- Build tools for easy monitoring of the system
- Develop cryoEM specific customer support solutions
- Developing SSD layer for Cryo specific tasks

## Transfers to MASSIVE:

Average 700-800 MB/s

Peak at over 1GB/s using Globus.



PANASAS

MOLECULAR  
HORIZONS



# UOW pathway from 2017 -> 2023

- What did we have in 2017 ? Well.. nothing really?
- We looked around..
- Real time analytics capability from which we learned a lot
- How we specified and “designed” XGPU16
- Very detailed feedback from Simon (important operational data)
- Feedback James - future direction (strategic feedback)
- Relationship with Nvidia, Xenon for compute
- Relationship with Panasas – data storage is everything
  
- “Customer centric design, architecture and implementation”
  - This is not the same as “the customer is always right”

# Nvidia GPU comparisons (so many spreadsheets and web sites)

GPU	released	core	Ram	ram bandwidth	TFLOP (single prec)	watts	TFLOP/watt
RTX 4090	2022	AD102	24	1008	73	450	0.16
RTX 4080	2022	AD103	16	717	43	320	0.13
A5000	2021	GA102	24	768	28	230	0.12
3080 12GB	2022	GA102	12	912	22	350	0.06
A4000	2021	GA104	16	448	19	140	0.14
V100	2017	GV100	32	900	14	250	0.06
2080 Ti	2018	TU102	11	616	12	250	0.05
1080 Ti	2017	GP102	11	484	11	250	0.04
Titan X	2016	GP102	12	480	10	250	0.04
3060 12GB	2021	GA106/104	12	360	9	170	0.05
A2000	2021	GA106	12	288	8	70	0.11
A2	2021	GA107	16	200	4.5	50	0.09

# XGPU 16 specifications (after many many quotes)

- 16 x NVidia A4000 GPUs in 4 banks of 4 GPUs
- 1 Terabyte of DDR4 memory
- 2 x AMD CPUS (24 cores ea, 48 cores total, 96 w/ HT)
- 16 Terabytes of nVME SSD storage (cache)
- 100 Gigabit Mellanox ethernet + 10 Gigabit ethernet + 1 Gigabit LOM
- Dual PSU (2400W x 2) – 2RU chassis!
- Noise like you wouldn't believe



# Nvidia A4000 specifications (AMPere – now replaced with ADA)

- GPU memory 16GB GDDR6
- Memory interface 256-bit
- Memory bandwidth 448 GB/s
- NVIDIA Ampere architecture based CUDA Cores 6,144
- NVIDIA third-generation Tensor Cores 192
- NVIDIA second-generation RT Cores 48
- Single-precision performance 19.2 TFLOPS
- System interface PCIe 4.0 x16
- Power consumption Total board power: 140 W
  
- <https://www.nvidia.com/content/dam/en-zz/Solutions/gtcs21/rtx-a4000/nvidia-rtx-a4000-datasheet.pdf>

# Nvidia A4000 as installed in XGPU16 (aggregate performance)

- 256G of GDDR6 memory
- CUDA Cores 98,304
- 3<sup>rd</sup> Gen Tensor Cores 3,072
- 2<sup>nd</sup> Gen RT Cores 768
- Single-precision performance 307.2 TFLOPS
- Power consumption Total board power: 2240 W (2.2kW)
- Less than \$100,000 (< \$90K? <\$80K??)

# Nvidia A4000 as installed in XGPU16

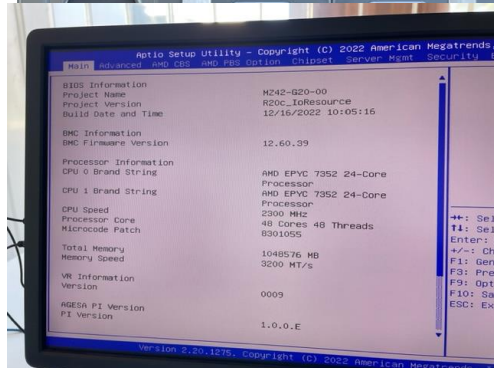

- Linux operating system (Rocky Linux)
- CryoSPARC application
- Multiuser access managed within CryoSPARC
- Everything is remote access via the web
- Users can use 2, 4, 6, 8 or possibly all 16 GPUs if required
- Expected to see 3-4 GPUs per user/job



# Terascale (memory and disk)

- A terabyte of memory
  - Full memory bandwidth with all slots filled
  - Expansion options still exist
- Optimized for multi-user (customer centric!)
- Linux is very efficient in how it uses memory
- Cheapest way to over-engineer performance
  
- Boot disks – 2 x 960G SATA SSD
- “Data” disks aka scratch – 4 x 4T nVME SSD (14T usable)
  - Configured as RAID0
  - All used as cache for CryoSPARC

# Terascale (CPU, memory, disk)




# Network and power

- 100G Mellanox is simple – copper approach for cost
- Reusing existing networking switching
- Allows us to build more data intensive future
  
- Dual power is a must
- UPS vs non UPS ?
- Hot replaceable PSU
- 10A vs 15A
- 2PSU vs 3PSU vs 4PSU
- Green power ?
- Sustainability ?



# Network and power




# UOW & NVidia

## Developing an industry partnership with Nvidia

- Understanding the roadmap for current and future GPUs
- Leveraging Nvidia's ability to talk to CryoSPARC developers
- Having hardware support for their Enterprise GPUs for 5 years
- Future work in CryoEM evaluation of new GPUs including benchmarking work
- Case study for how we are using GPUs to solve different analysis needs (including real time)
- Access to Nvidia staff, resources and expertise
- Why we needed to swap from consumer to enterprise GPUs




**NVIDIA**



# UOW & Xenon

Long term partner we had developed a relationship with

- Our first real time analysis platform vendor (XGPU1)
- In our procurement system – approval via finance system
- Worked with us to develop the specification guided by a very iterative process
- Jason Lee had to put up with a lot of emails and phone calls as we chopped and changed things
- Worked with us to find hardware that would give us a very dense platform
- On site maintenance (contract)
- Worked within a pretty tight budget!



**XENON**<sup>®</sup>  
High Performance Computing

MOLECULAR  
HORIZONS



UNIVERSITY  
OF WOLLONGONG  
AUSTRALIA

# UOW & Panasas

Storage is key to our workflow

- Long history of building high performance HPC filesystems
- Able to give us Petascale capacity with sufficient performance
- Spent a lot of time to help us optimize data throughput
- A key to allowing us to run with all flash compute
- “Standard” hardware with a simple software stack to manage
- Read performance critical to populate caches on compute nodes




# Giving your research/operations staff the right access and right tools

## Netdata:

- Provides instrumentation for what your compute is doing
- Crucial to allow operations/research staff to see what is happening AS they are developing/optimizing workflows
- Free software – minimal training/learning needed
- Helps to diagnose configuration issues
- Visualizes analytical data


Saves a lot of back and forth!

The individual metrics can be found in the [kernel documentation](#).




## load

Current system load, i.e. the number of processes using CPU or waiting for system resources (usually CPU and disk). The 3 metrics refer to 1, 5 and 15 minute averages. The system calculates this once every 5 seconds. For more information check [this wikipedia article](#).



## System Overview

Overview of the key system metrics.




### cpu

Total CPU utilization (all cores). 100% here means there is no CPU idle time at all. You can get per core usage at the [CPUs](#) section and per application usage at the [Applications Monitoring](#) section.

Keep an eye on [lowlat](#) ( 0.07%). If it is constantly high, your disks are a bottleneck and they slow your system down.

An important metric worth monitoring, is [softirq](#) ( 0.05%). A constantly high percentage of softirq may indicate network driver issues. The individual metrics can be found in the [kernel documentation](#).

Total CPU utilization (system.cpu)



### load

Current system load, i.e. the number of processes using CPU or waiting for system resources (usually CPU and disk). The 3 metrics refer to 1, 5 and 15 minute averages. The system calculates this once every 5 seconds. For more information check [this wikipedia article](#).


System Load Average (system.load)



# ram

System Random Access Memory (i.e. physical memory) usage.

System RAM (system.ram)



# System Overview

Overview of the key system metrics.




## cpu

Total CPU utilization (all cores), 100% here means there is no CPU idle time at all. You can get per core usage at the [CPUs](#) section and per application usage at the [Applications Monitoring](#) section.

Keep an eye on [lowlat](#) (0.5%). If it is constantly high, your disks are a bottleneck and they slow your system down.

An important metric worth monitoring, is [softirq](#) (1.00%). A constantly high percentage of softirq may indicate network driver issues. The individual metrics can be found in the [kernel documentation](#).

Total CPU utilization (system.cpu)




percentage Wed, 05 Jul 2023 | 16:10:10

## load

Current system load, i.e. the number of processes using CPU or waiting for system resources (usually CPU and disk). The 3 metrics refer to 1, 5 and 15 minute averages. The system calculates this once every 5 seconds. For more information check [this wikipedia article](#).

System Load Average (system.load)



load Wed, 05 Jul 2023 | 16:10:10

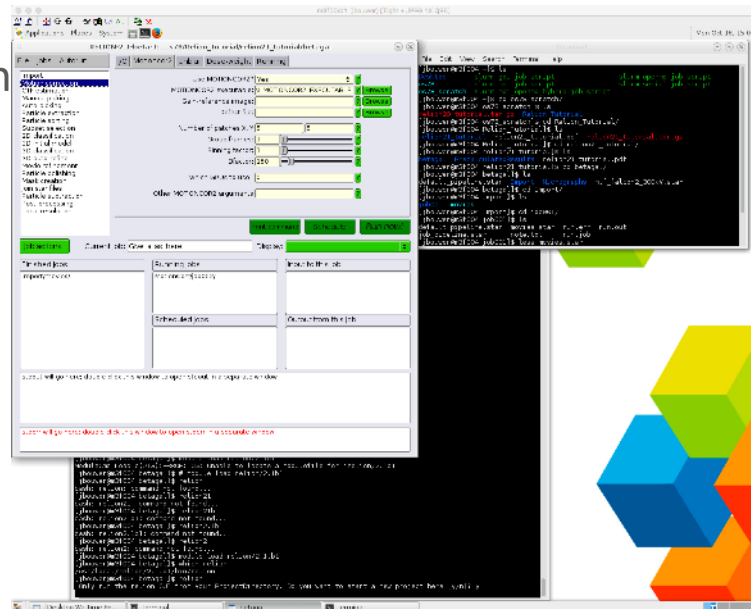


# HPC Partnerships vs Cloud vs Local HPC (aka On Prem)

- You don't have to pick one (and it might be better if you don't)
- Each has different strengths and weaknesses
- Risk and governance, budget, service level
- What is your funding ? Capex vs Opex ?
- How fast is compute moving / developing ?
- How fast is storage and networking doing the same ?
- Where do you find your staff and can you keep them ?

# Ongoing Partnership with MASSIVE at Monash University

- UOW CryoEM continues to maintain a partnership with MASSIVE
- Active customers have access to MASSIVE partnership resources
- Applications and hardware updates, new installs, and data management are abstracted from the customers.
- MASSIVE's desktop application allow UOW customers control of a supercomputer from a home or work PC.



# Compute strategy ?

- Local realtime analysis (CryoSPARC Live)
  - One node per microscope
- Remote HPC (MASSIVE partnership) continues
- Local HPC (single node, scales to 2!)
  - HPC batch queueing is a real problem for cache
- Future cloud work (AWS? GPC? Azure?)
- Future HPC (NCI ? Pawsey ? Others ?)
- Staff!
- Inflexion points – when we do spend money ? Replace equipment ?

## Where do we go from here ?

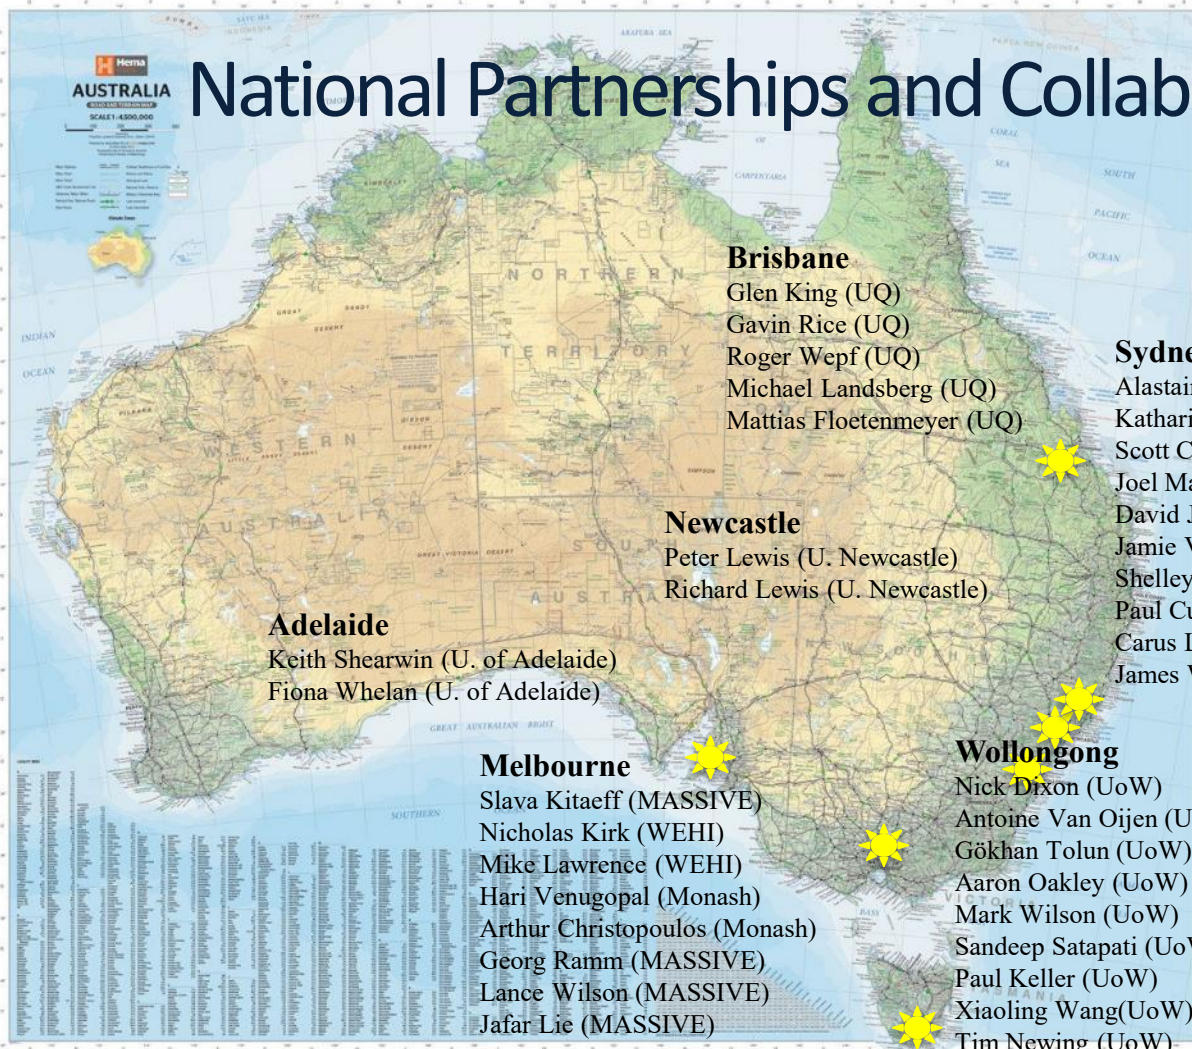
- Continuing partnership with MASSIVE
- We need to update our real time analysis nodes (Nvidia Enterprise GPUS?)
- Another compute node ?
- Will the big HPC providers ever build applications for this ?
- Cloud options – build our own in AWS/Google/Azure?
- CryoSPARC Cloud provider already exists – Simon evaluating this for our customers

# Cryo-EM projects

## Collaborations at UoW

- **Gokhan Tolun (UoW) – CryoEM Lead Academic**
  - Viral DNA Recombination
- **Alan Munn & Ishtiaq Ahmed (Griffith University, AU)**
  - Endosomal Sorting Complex Required for Transport (ESCRT) – Yeast
- **Harshad Ghodke (UOW)**
  - Transcription (RNAP complexes) – *E. coli*
- **Heath Ecroyd (UOW) & Elizabeth Vierling (University of Massachusetts Amherst, USA)**
  - Small Heat Shock Protein – Wheat
- **Jenny Beck, Julian Harrison, & Celine Kelso (UOW)**
  - Venom components – Snake
- **Jenny Martin (UOW) & Camila Cotrim (Griffith University, AU)**
  - Zn-transporters – Human
- **John Bruning (University of Adelaide, AU)**
  - Nuclear receptor (Peroxisome proliferator activated receptor  $\gamma$ ) in antidiabetics and insulin sensitivity – Human
- **Justin Yerbury & Luke McAlary (UOW)**
  - ALS (Amyotrophic lateral sclerosis or Lou Gehrig's disease) associated fibrils – Human
- **Mark Wilson & Sandeep Satapathy (UOW)**
  - Clusterin secreted chaperon – Human
- **Mikhail Kashlev & Lucyna Lubkowska (NCI/NIH, USA)**
  - Transcription elongation complexes – *E. coli*
- **Nick Dixon, Aaron Oakley, Antoine van Oijen, Zhi-Qiang Xu, Slobodan Jergic, Simon Brown, Timothy O'Shea (UOW)**
  - DNA replication complexes – *E. coli*
- **Aaron Oakley, Paul Keller,, and Jayden Gaston**
  - MicroED

# National Partnerships and Collaborations



## Brisbane

- Glen King (UQ)
- Gavin Rice (UQ)
- Roger Wepf (UQ)
- Michael Landsberg (UQ)
- Mattias Floetenmeyer (UQ)

## Sydney

- Alastair Stewart (UNSW – Victor Chang)
- Katharine Michie (UNSW)
- Scott Cohen (CMRI)
- Joel Mackay (U. Sydney)
- David Jacques (UTS)
- Jamie Vandenberg (Victor Chang)
- Shelley Wickham (U. Sydney)
- Paul Curmi (UNSW)
- Carus Lau (Victor Chang)
- James Walsh (Victor Chang)

## Newcastle

- Peter Lewis (U. Newcastle)
- Richard Lewis (U. Newcastle)

## Adelaide

- Keith Shearwin (U. of Adelaide)
- Fiona Whelan (U. of Adelaide)

## Melbourne

- Slava Kitaeff (MASSIVE)
- Nicholas Kirk (WEHI)
- Mike Lawrence (WEHI)
- Hari Venugopal (Monash)
- Arthur Christopoulos (Monash)
- Georg Ramm (MASSIVE)
- Lance Wilson (MASSIVE)
- Jafar Lie (MASSIVE)

## Wollongong

- Nick Dixon (UoW)
- Antoine Van Oijen (UoW)
- Gökhan Tolun (UoW)
- Aaron Oakley (UoW)
- Mark Wilson (UoW)
- Sandeep Satapati (UoW)
- Paul Keller (UoW)
- Xiaoling Wang (UoW)
- Tim Newing (UoW)

## New Zealand

- Peter Mace (U. of Otago)
- Kurt Krause (U. of Otago)
- Mihnea Bostina (U. of Otago)
- Alok Mitra (U. Auckland)
- Emma Andrews (U. of Waikato)

# THE WORLD



# U

## UOW:

Eileen McLaughlin  
Nick Dixon  
Antoine van Oijen  
Gökhan Tolun  
Alan Champion  
Aaron Oakley  
Simon Brown  
Claudia Kielkopf  
Jason Andrade  
Joshua Silver  
Craig Muller  
Tim Newing  
Andrew Worth  
Paul Keller  
Jayden Gaston

# O

## Monash and MASSIVE:

Slava K  
James Whisstock  
Lance Wilson  
Jafar (Dynamite) Lie  
Georg Ram  
Hari Venugopal  
Hans Elmlund

## ANSTO:

Ken Short  
Daniel Oldfield

## Victor Chang :

Alastair Stewart  
Meghna Sobti  
Jamie Vandenberg  
Carus Lau  
Mark Hunter

## Panasas:

Adam Marko  
Jeff Whitaker  
Chris Sassone  
Alex Johnson

# W

## ThermoFisher Scientific:

Marc Peeters  
Eric Chen  
Ryan Shaw  
Xi Jiang Pan  
Shiny Shi

## Nvidia:

Wei Fang  
Gary Mancuso  
Johan Barthelemy  
Suda Ramachandran  
Tomasz Bednarz

The End...

MOLECULAR  
HORIZONS



UNIVERSITY  
OF WOLLONGONG  
AUSTRALIA