MONASH University

MONASH
eRESEARCH

# Monash High Performance Computing

Gin Tan
Senior HPC Consultant

GROUP OF EIGHT AUSTRALIA

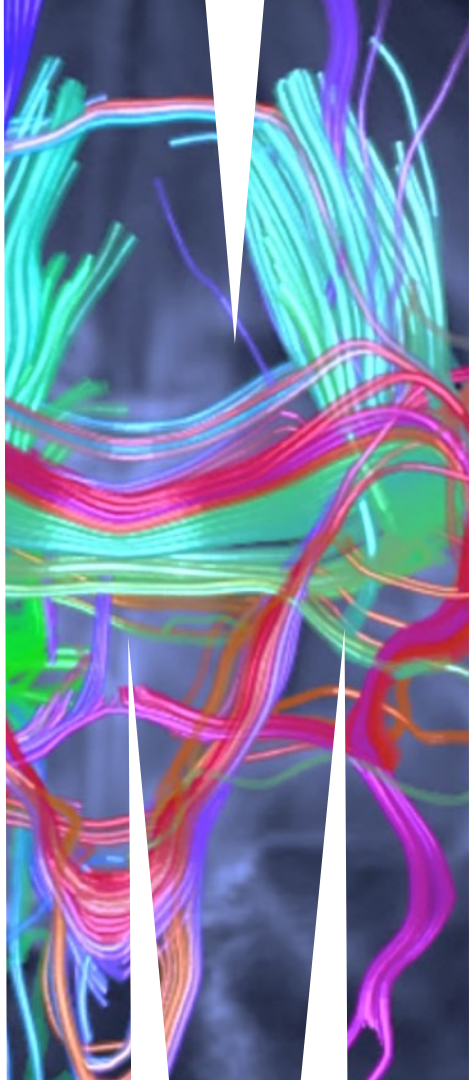

# Monash High Performance Computing

MONASH University

## MeRC (Monash eResearch)

Monash HPC Infrastructure
- MASSIVE
- MonARCH
- Characterisation VL and Instruments

MASSIVE-3

# MeRC Infrastructure

MONASH University

**63+ million**

**CPU-core hours p.a. of computing time for Monash researchers**

**Monash University is the largest user of national merit allocated supercomputing time**

**40+ peak instruments integrated at Monash**

**National Instrument integration program:**

**60+ instruments across Australia ($250M+ capital)**

**5x oversubscription**

**MASSIVE time requests through national merit is 5x what is available**

**$3.4m p.a. of research cloud access from contribution of $250k p.a.**

**10+ petabytes of research storage**

# Monash High Performance Computing

MONASH University

MeRC (Monash eResearch)

Monash HPC Infrastructure

- NCI and NCMAS
- MASSIVE
- MonARCH
- Characterisation VL and Instruments

MASSIVE-3

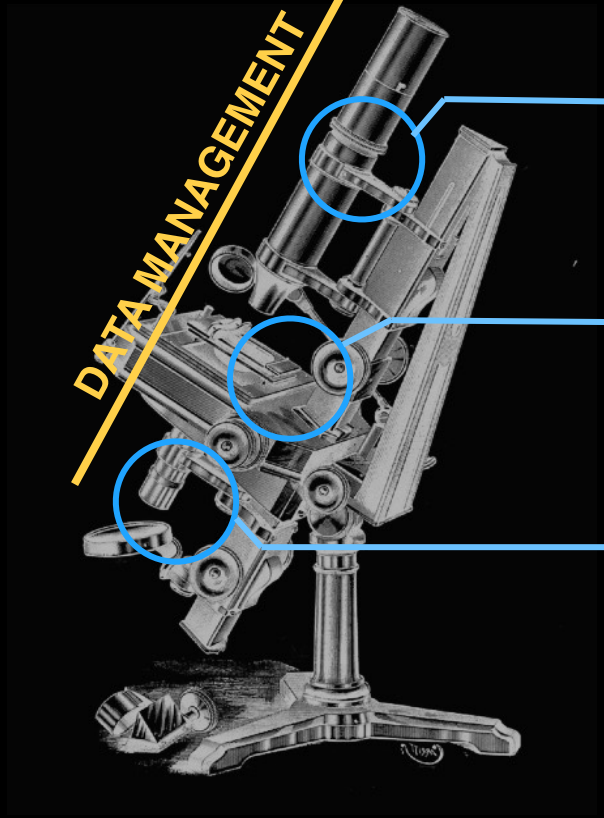


**INSIGHT**
Lens
**Visualisation techniques and infrastructure**

**ANALYSIS**
Filters
**HPC and cloud Workbenches**

**CAPTURE**
Light Source
Samples
**The instruments**

# HPC Infrastructure

## Peak, Specialised and Community

MONASH University

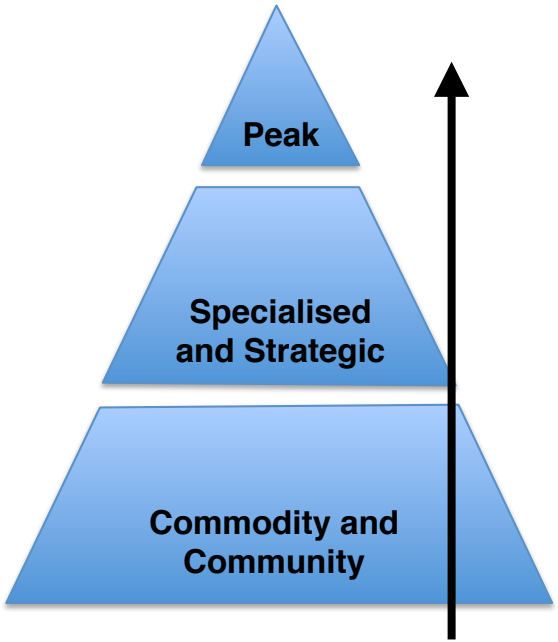


**National Computational Infrastructure**
Big engineering and science
Existing big users
Well established requirements
Next generation of big HPC users

**MASSIVE**
Leadership in Characterisation
Instrumentation and Accessible HPC
HPC for new communities

**Monash Campus Cluster**
Long tail
Undergrad and Postgrad Education
Research group solutions
Play pen

## Outcomes

1. Monash has built a respected capability in high performance computing and is very strong in peak areas of computational science;

2. Monash is a leader in NCRIS characterization informatics, which is providing significant benefit to researchers, infrastructure and future investment / leverage;

4. Monash is unique in building strong capability for the long tail of non-traditional HPC users (in particular life sciences);

5. Monash is now consistently the top merit allocation user of the NCI;

6. Monash researcher have access to dedicated local expertise and resources;

# MASSIVE

MONASH University

## HPC for Characterisation
## Specialised Facility for Imaging and Visualisation

**~$2M per year funded by partners and national project funding**

### Partners

Monash University
Australian Synchrotron
CSIRO
University of Wollongong

### Affiliate Partners

ARC Centre of Excellence in Integrative Brain Function
ARC Centre of Excellence in Advanced Molecular Imaging

### HPC

150+ active projects
2,000+ user accounts
100+ institutions across Australia

### Interactive Vis

600+ users

### Instrument Integration

Integrating with key Australian Instrument Facilities.
– IMBL, XFM, MX2
– CryoEM
– MBI
– NCRIS: NIF, AMMRF

### Large cohort of researchers new to HPC

## 10+ Big Data and Big Collection Generating Instruments at Monash University

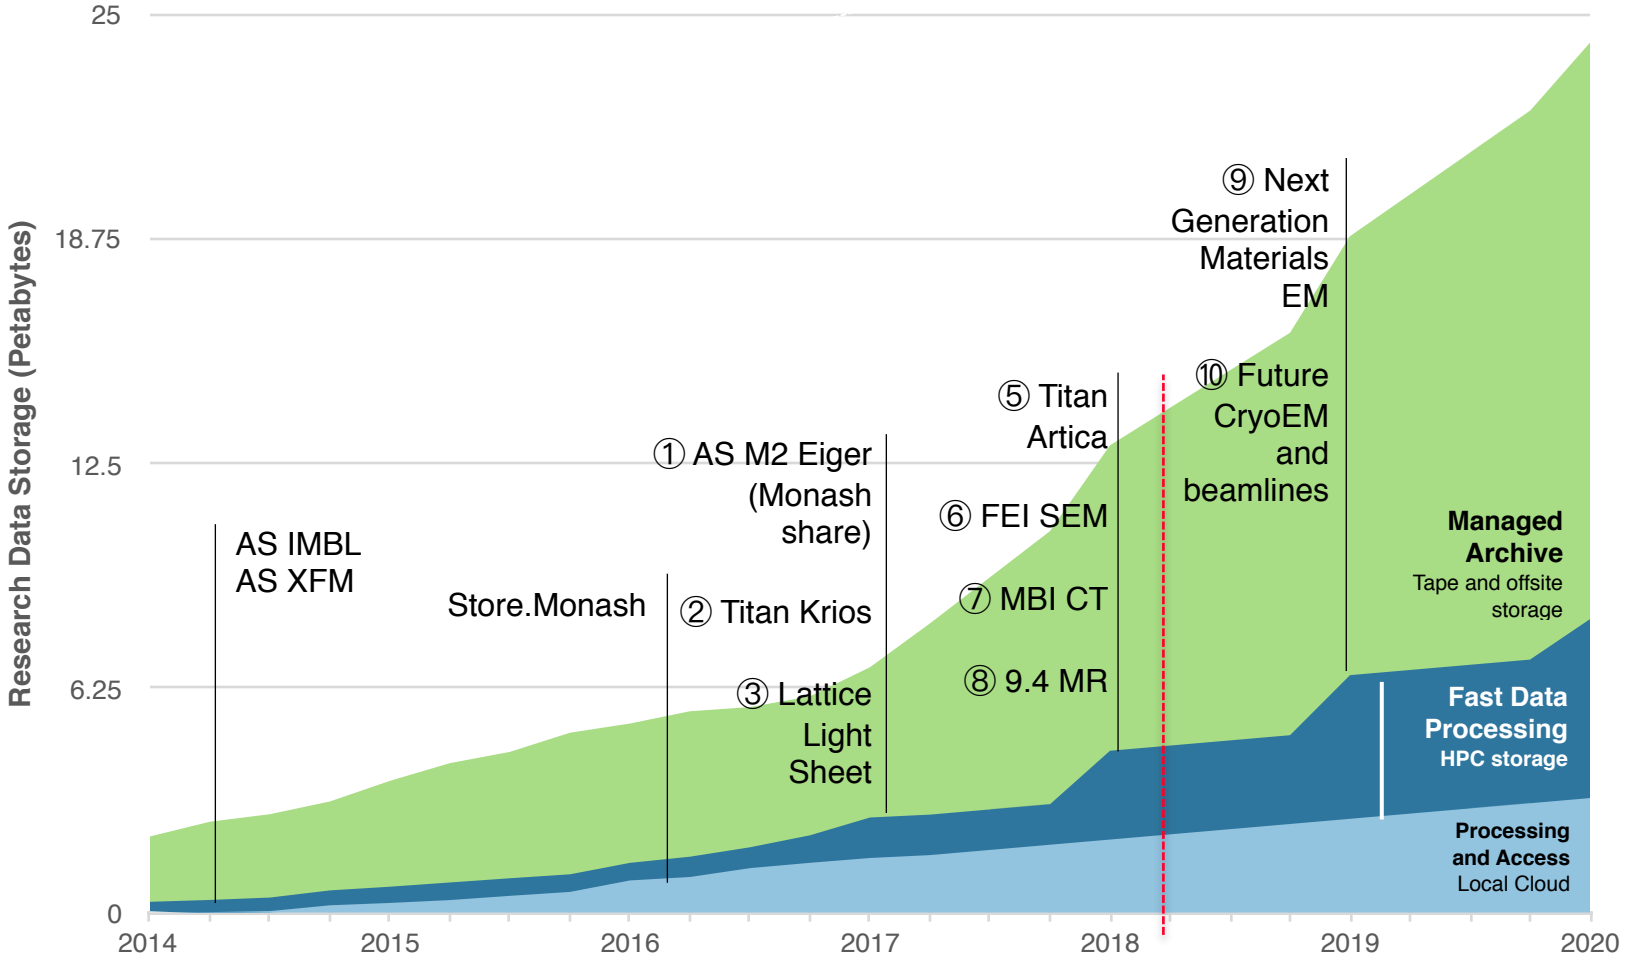

# Life sciences focus

## Breakdown of usage - FOR codes

MONASH University

| M1 & M2 (2011 onward) | |
|---|---|
| 09 ENGINEERING | 40.03% |
| 02 PHYSICAL SCIENCES | 13.06% |
| **11 MEDICAL AND HEALTH SCIENCES** | **10.64%** |
| 08 INFORMATION AND COMPUTING SCIENCES | 9.45% |
| **06 BIOLOGICAL SCIENCES** | **8.07%** |
| **17 PSYCHOLOGY AND COGNITIVE SCIENCES** | **7.23%** |
| Oher | 11.53% |

| M3 (2017 onward) | |
|---|---|
| **11 MEDICAL AND HEALTH SCIENCES** | **42.35%** |
| **06 BIOLOGICAL SCIENCES** | **30.14%** |
| **17 PSYCHOLOGY AND COGNITIVE SCIENCES** | **10.09%** |
| 08 INFORMATION AND COMPUTING SCIENCES | 5.17% |
| 09 ENGINEERING | 4.47% |
| 02 PHYSICAL SCIENCES | 4.27% |
| Oher | 3.49% |

# MonARCH

## Community

### Campus Cluster

**Provide Monash researchers with a local capability that focuses on engagement, education and community.**

### Investment

A co-investor model
Examples include Computational Chemistry, Astro and Fluid Dynamics
1/3rd of MonARCH is co-purchased

### Integrated into undergraduate study

**CHM3911 Advanced Physical Chemistry**

80 students across 3 practical sessions
Gaussian and GaussView for calculations

Students taught how to use a HPC system to perform their calculations

# Monash High Performance Computing

MONASH University

MeRC (Monash eResearch)

Monash HPC Infrastructure
- NCI and NCMAS
- MASSIVE
- MonARCH
- Characterisation VL and Instruments

MASSIVE-3

# Remote Desktops

## Workbenches

## Deployed on the research cloud and alongside HPC

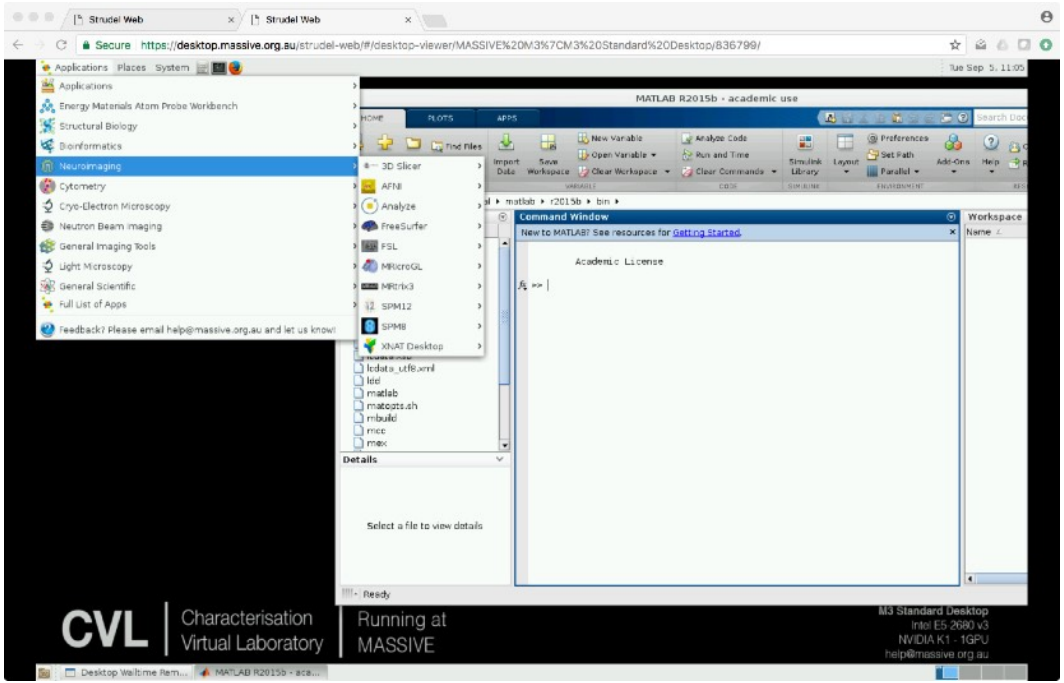


Atom Probe, Structural Biology, Bioinformatics, Cytometry, Cryo-Electron Microscopy, Neutron Beam Imaging, General Imaging Tool, Light Microscopy, General Scientific, X-ray

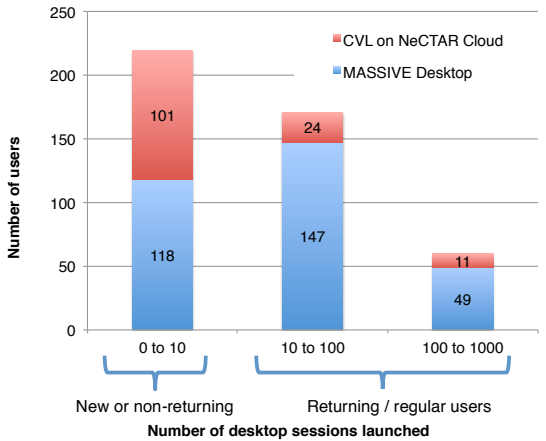

# Demo

http://desktop.massive.org.au

# Monash Research Cloud

## R@CMon

- A fabric of software defined infrastructure
  - networking: Cumulus + Mellanox RoCE
  - compute orchestration: OpenStack
  - disk: Ceph
  - tiering: watch this space
- CIFS, NFS, VDI, jupiter-VL, NSP, HPC, MyTardis, Figshare, safe havens, … (growth here)
- Network fabric spans 2x data centres and the “Clayton precinct”
- Network is heterogeneous but adaptable between: Ethernet & RoCE, and 10-100gb

| | |
|---|---|
| **Nodes** | 209 |
| **CPUs** | 6372 |
| **Threads** | 10136 |
| **RAM** | 46624 GBs |
| **GPUs** | 145 |
| **Persistent disk** | 8792 TBs |
| **Persistent tiered** | ~8000 TBs |
| **Users** | ~5000 |

R@CMon MONASH University

# Monash High Performance Computing

MONASH University

MeRC (Monash eResearch)

Monash HPC Infrastructure
- NCI and NCMAS
- MASSIVE
- MonARCH
- Characterisation VL and Instruments

MASSIVE-3

# Monash High Performance Computing

MONASH University

- The third generation MASSIVE supercomputer
- It's a bit different from traditional HPC
- HPC on the cloud
- Pass-through Mellanox CX-3 and CX-4 HCA
- Pass-through GPUs K1, K80, P100, V100
- High demand for Vis jobs
- Segregation and security
- RDMA over Ethernet (RoCE)
- Running software in singularity container
- Running Slurm v17 - fixes kmem cgroup constraint

**M3 at Monash University**

# Monash High Performance Computing

MONASH University

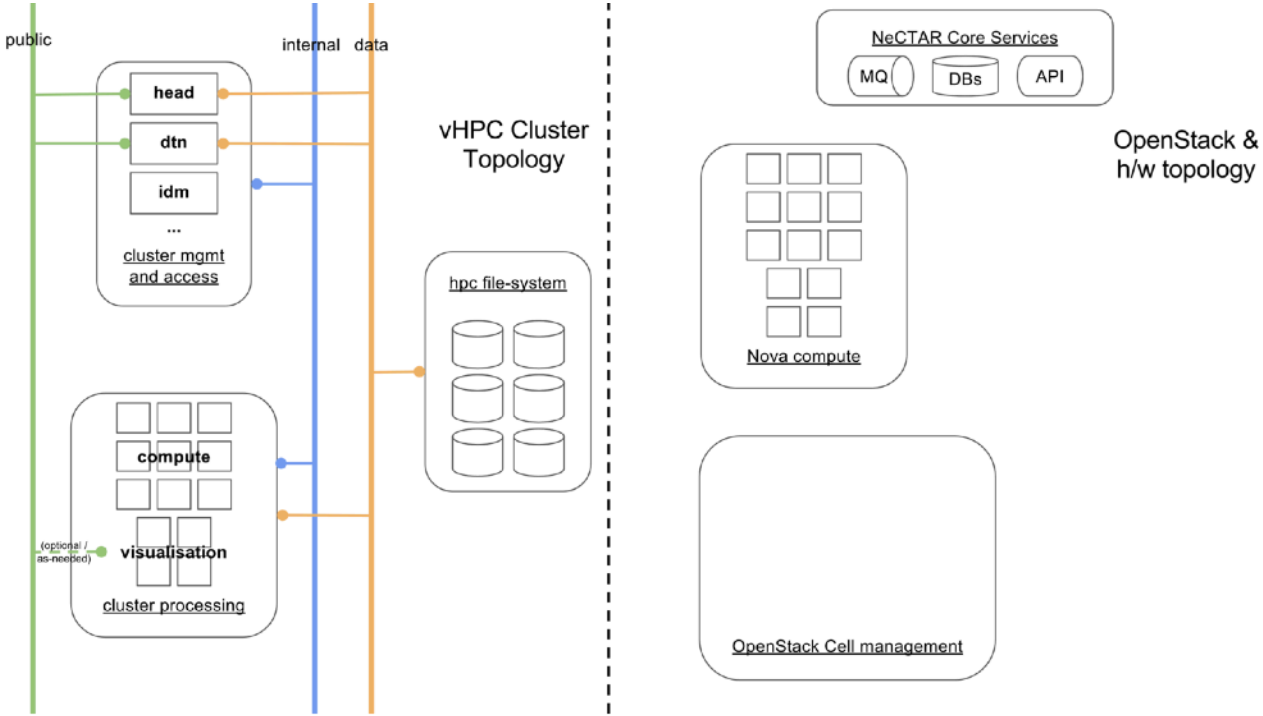


vHPC

# Monash High Performance Computing

MONASH University

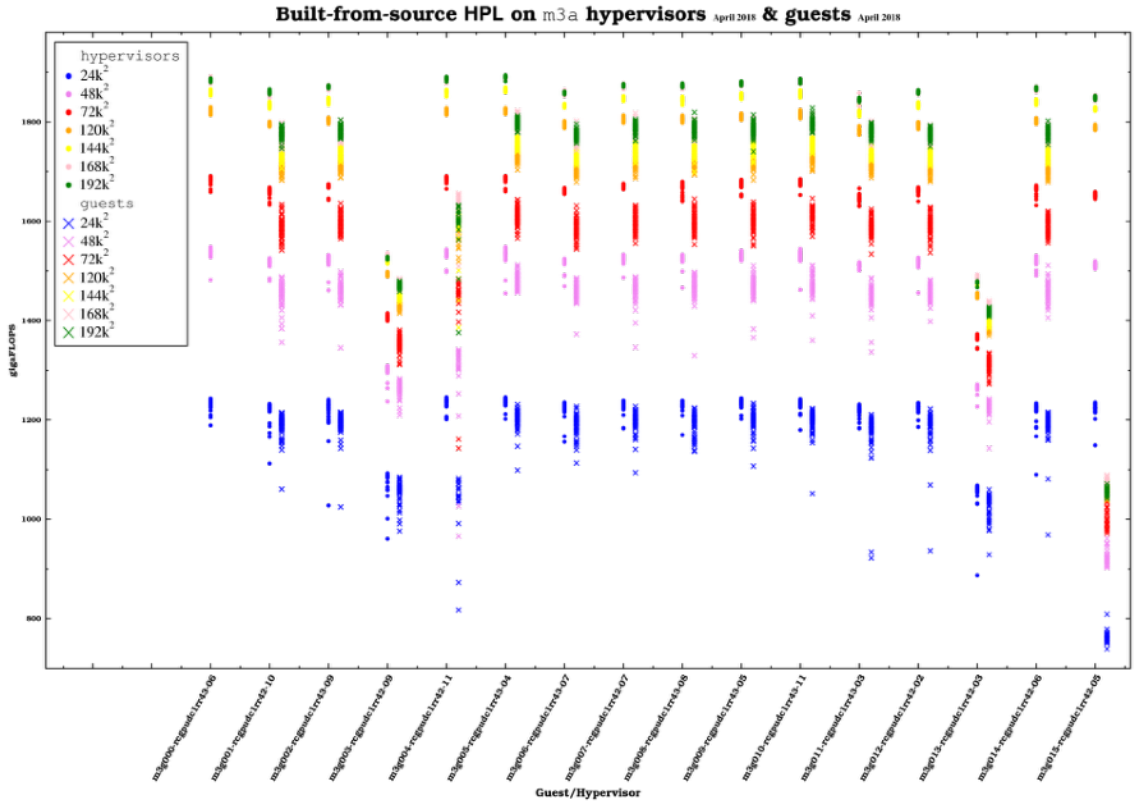


- numa topology
- CPU tuning
- THP
- around 40 runs
- 1.6 TFlops
- OpenMPI v2.1.3
- GCC v5.4.0
- Intel MKL 2017u4
- Singularity 2.4.5

HPL

# Monash High Performance Computing

MONASH University

## M3 at Monash University
### (2018 upgrade)

1,600 Intel Haswell CPU-cores
**2,520 Intel Skylake CPU-cores**

NVIDIA GPU coprocessors for data processing and visualisation:
- 48 NVIDIA Tesla K80
- 40 NVIDIA Pascal P100
- **60 NVIDIA Volta V100**
- **2 NVIDIA DGX1-V**

- 8 NVIDIA Grid K1 GPUs for medium and low end visualisation

A 1.1 petabyte Lustre parallel file system
**A 3 petabyte Lustre parallel file system usable after upgrade**

100 Gb/s Ethernet Mellanox Spectrum

# Monash High Performance Computing

MONASH University

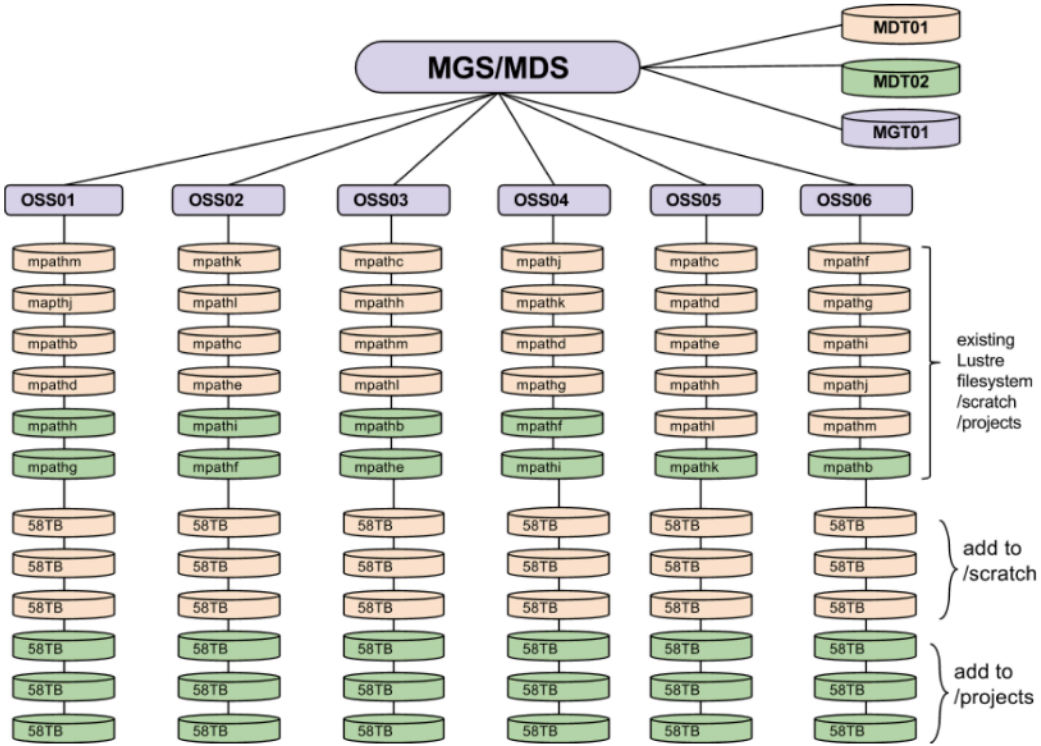


Lustre Storage upgrade to 3PB

# Monash High Performance Computing

MONASH University

## What don't we have?

- ZFS/NFS
- glusterFS
- IBM GPFS
- Ceph FS
- Lustre FS

## Why not GPFS:

- Proprietary
- Communities for the workflow
- General HPC e.g. simulation
- Maintenance cost & effort

## Why Lustre:

- **Clusters are sitting on Openstack research cloud**
- **Cinder driver for Openstack**
- **Community support**
- **Recently upgraded from 2.5 to 2.10.3**
- **Taking advantage of sub-dir mount in the cloud**
- **Progressive file layout for mixed use environment**

# Monash High Performance Computing

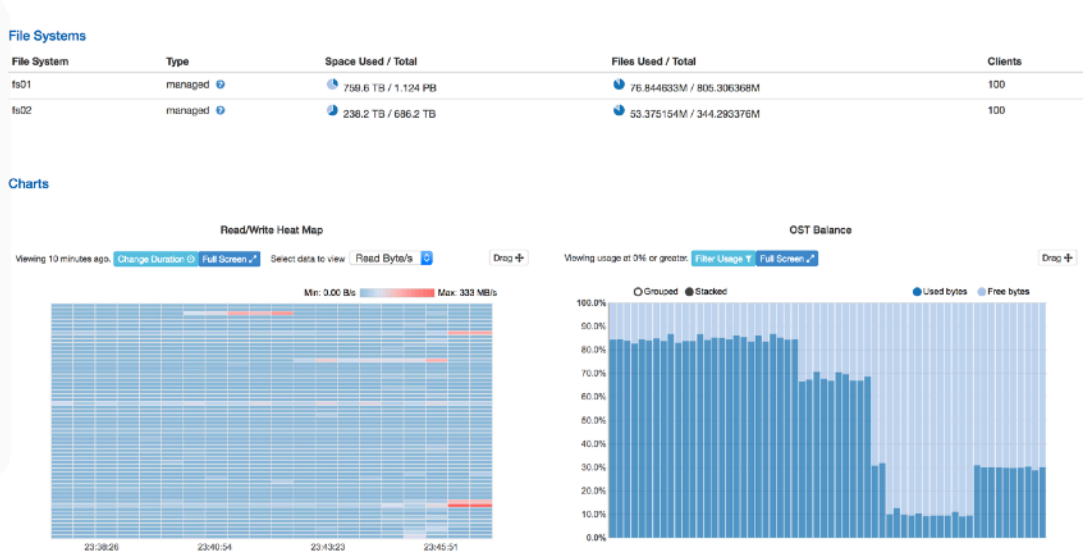


- Intel Manager for Lustre
- Managed mode
- Admin training
- Provide job stats
- And monitoring

IML

# Monash High Performance Computing

MONASH University

Thank You

email: gin.tan@monash.edu
url: www.massive.org.au

Questions?