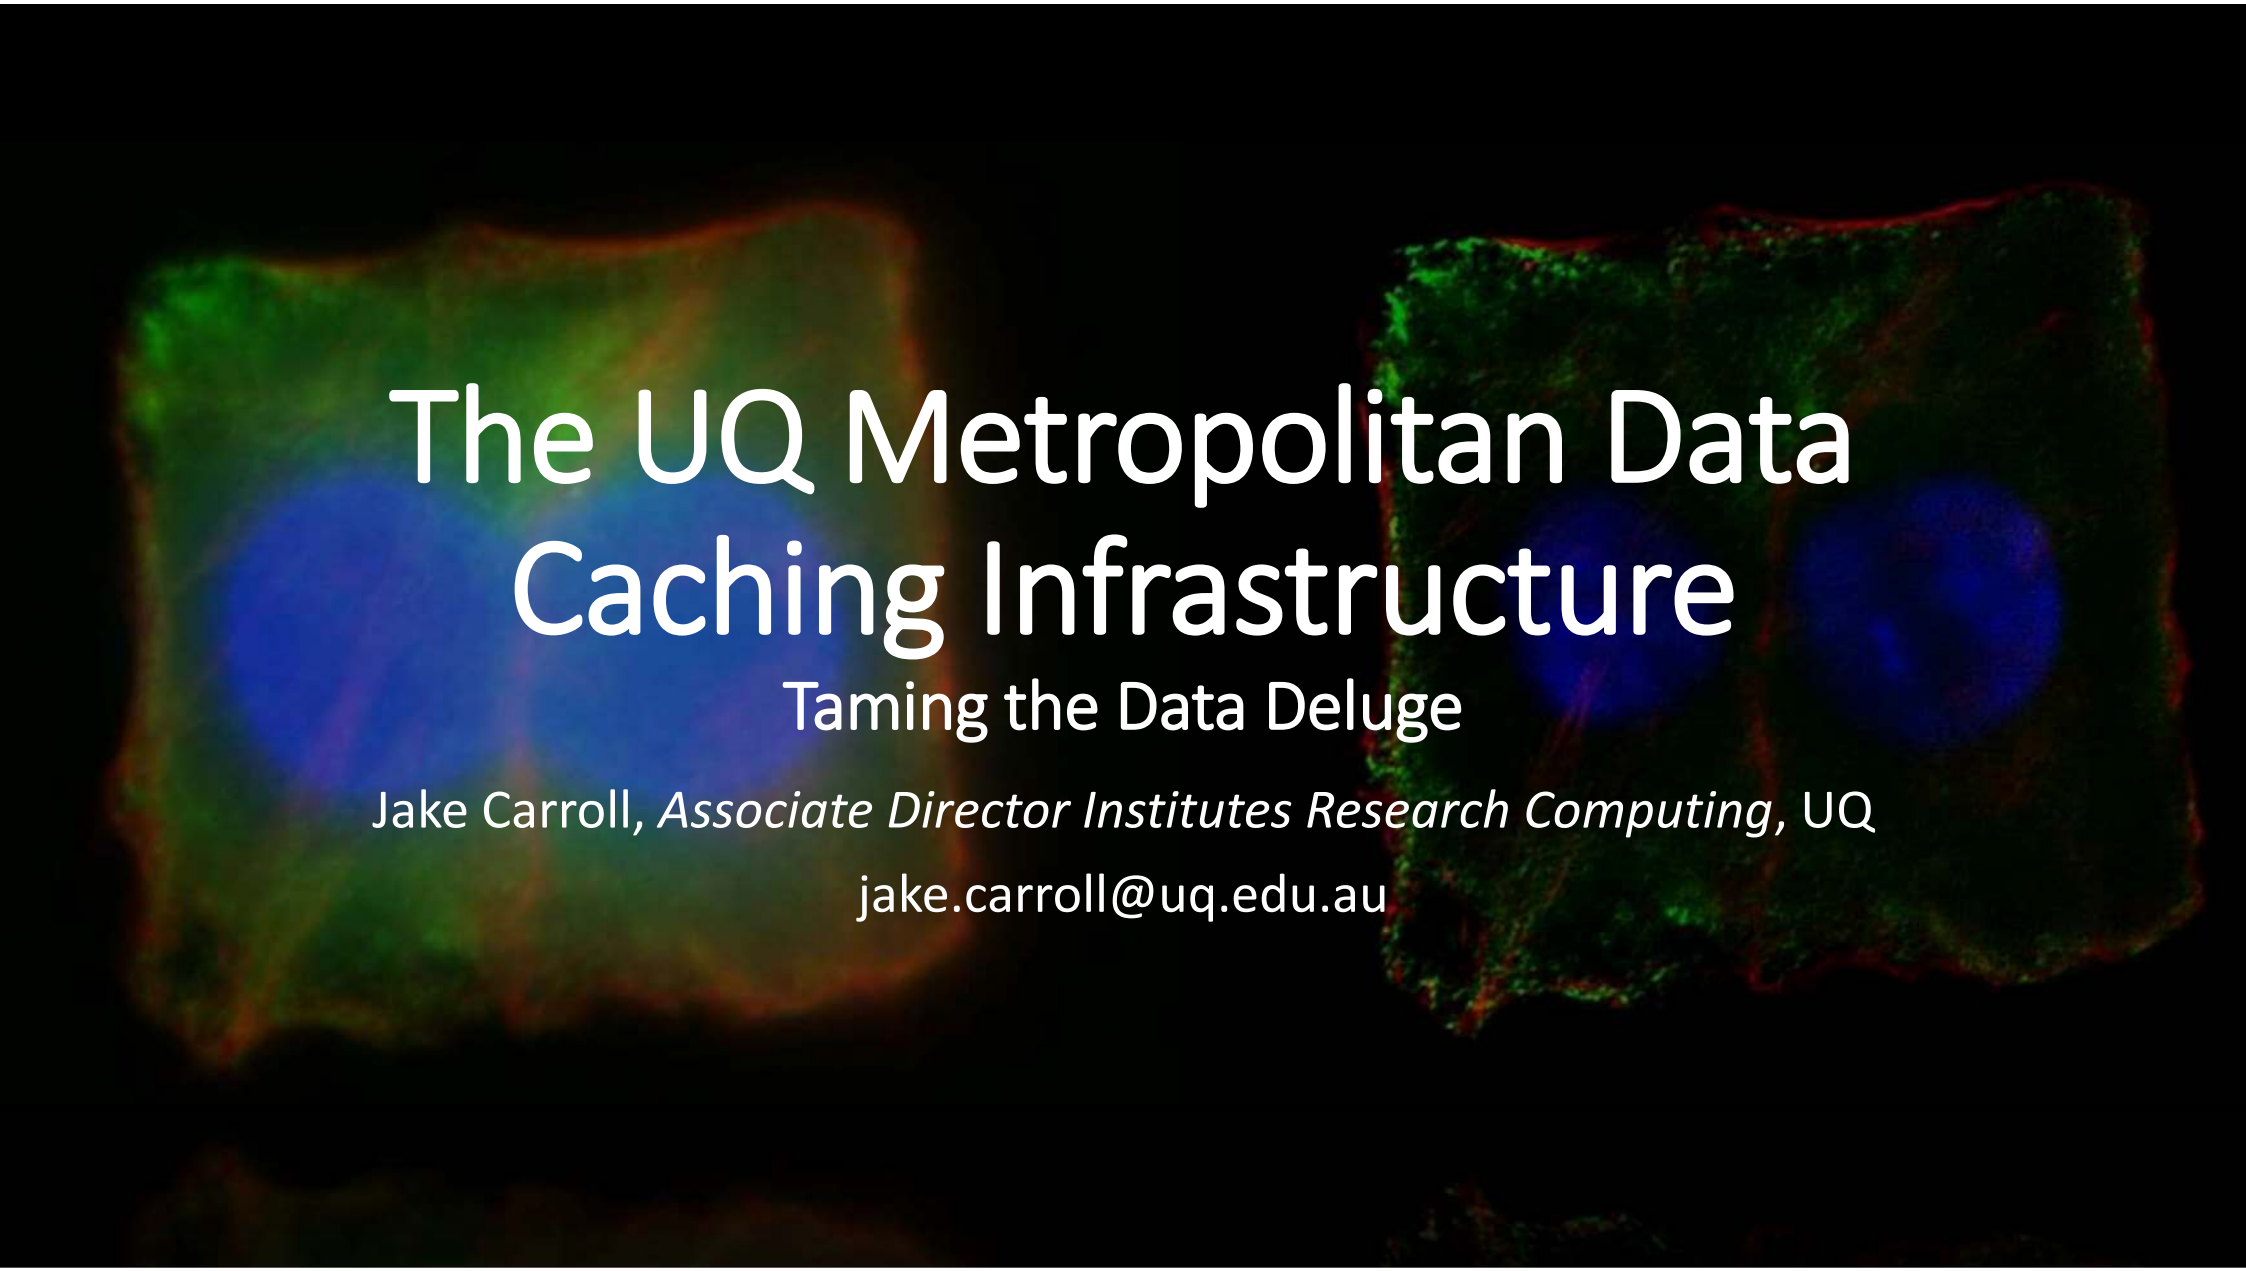

The UQ Metropolitan Data Caching Infrastructure

Taming the Data Deluge

Jake Carroll, *Associate Director Institutes Research Computing*, UQ

jake.carroll@uq.edu.au

This is a hard talk to give...

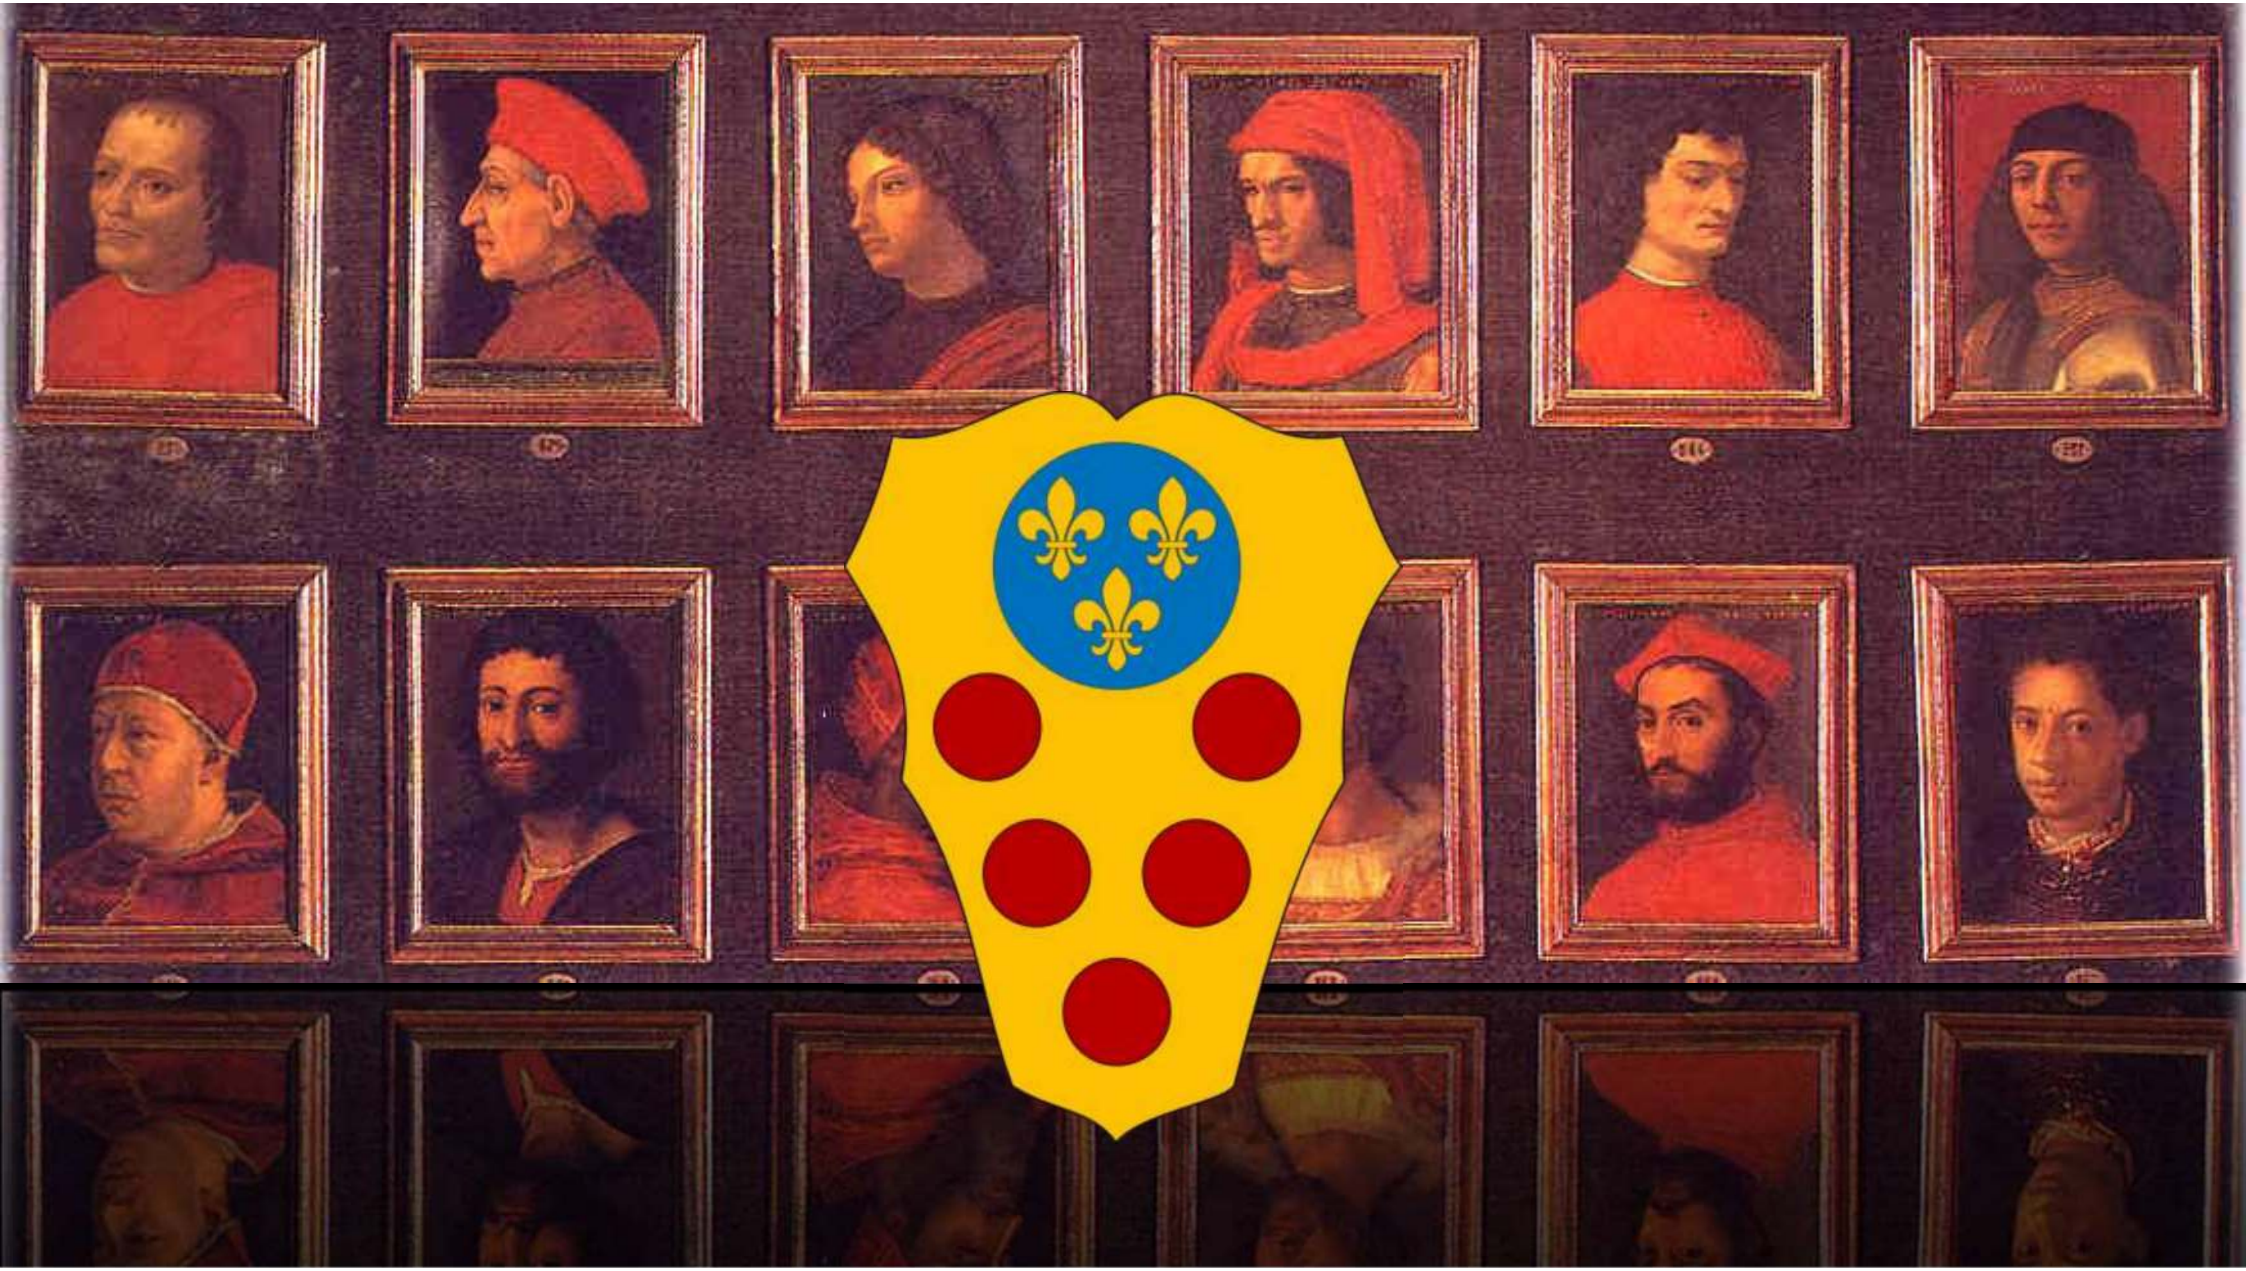


It was, legitimately, a dark and stormy night…

D.O.B.

11/03/2015 @ 22:18

Some fun stats about MeDiCI

- A combined on-disk capacity of 8.73PB of cache.
- A total write-workload since going into production of ~33PB.
- > 2000 collections of data hosted.
- Infrastructure powered two different ARC Laureate 2018 award winners with over $8.2m AUD in grant successes from data running atop this platform.
- 2.5 billion files under management.
- Total throughput capability of ~52GB/sec to front end caches.


I LIKE
IT.
WHAT
IS IT?

Meet the family.

CAI

QBI

* Imaging data f… super res, LLSM)

IMB

…

AIBN

Scientific infrastructure of immense scale

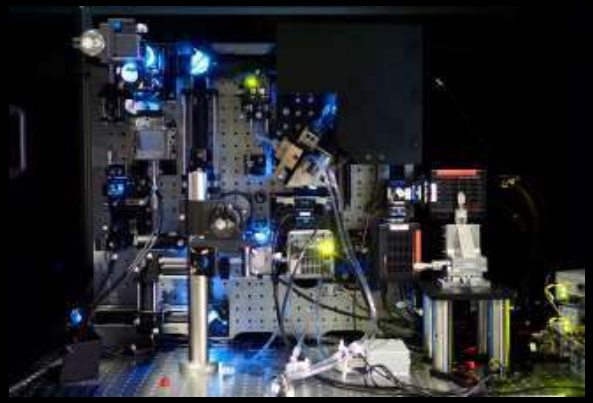

RESEARCH COMPUTING CENTRE

UQ's supercomputing strategy - Right supercomputer for the right task. "Best fit"

Tinaroo - 7000 cores of Intel Broadwell. Tight MPI, massively parallel, Infiniband FDR connected "Traditional" HPC.

FlashLite - 1632 cores of Intel Haswell. High memory footprint, virtual SMP (ScaleMP), high throughput. SSD /tmp in each node.

Awoonga - 1032 cores of Intel Broadwell. Loosely coupled, embarrassingly parallel, high latency tolerant workloads. *Ethernet* connected HPC.

The era of *accelerator based* supercomputing

To cope with 100’s of terabytes per day of imaging, genomics and sensor data, UQ turned to GPU accelerated supercomputing to solve its significant and complex problems.

First NVIDIA Volta based production system in Asia Pacific, 3rd in the world

MeDiCI

"Born" connected natively to UQ RCC's parallel filesystem data fabric; 2.75µs RDMA to instrument data from around UQ.

WIENER

~4.3 petaFLOPS of accelerated Compute capability; UQ's fastest HPC facility

189,440 CUDA cores, 23,680 dedicated hardware tensor cores

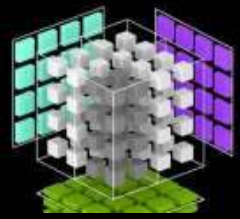

openHPC

OpenHPC Reference Site

BeeGFS

A dual parallel filesystem approach
BeeGFS + GPFS, delivering 180GB/sec and 25m IOPS of sustained performance in RDMA connected nVME flash

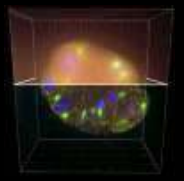

Deconvolution

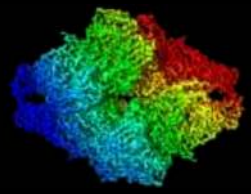

Accelerated EM

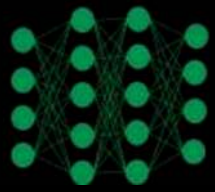

Deep Learning & AI

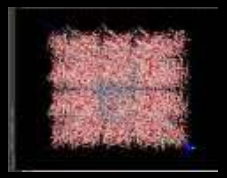

Amber protein structure seek acceleration

100Gbit/sec EDR non-blocking fabric; 1.9Tbit of combined signaling capability

The most powerful dedicated GPU volta-class system in the southern hemisphere, currently.

A pervasive set of problems...

- The movement of data.
- People not _wanting_ to move their data to the experiment.
- People not _wanting_ to micro-manage where the data is.
- People assuming it is “safe” where it is.
- People assuming it is “next to” the compute resource.
- People misunderstanding preservation + archiving.
- Retention.
- Locality to the instrument itself

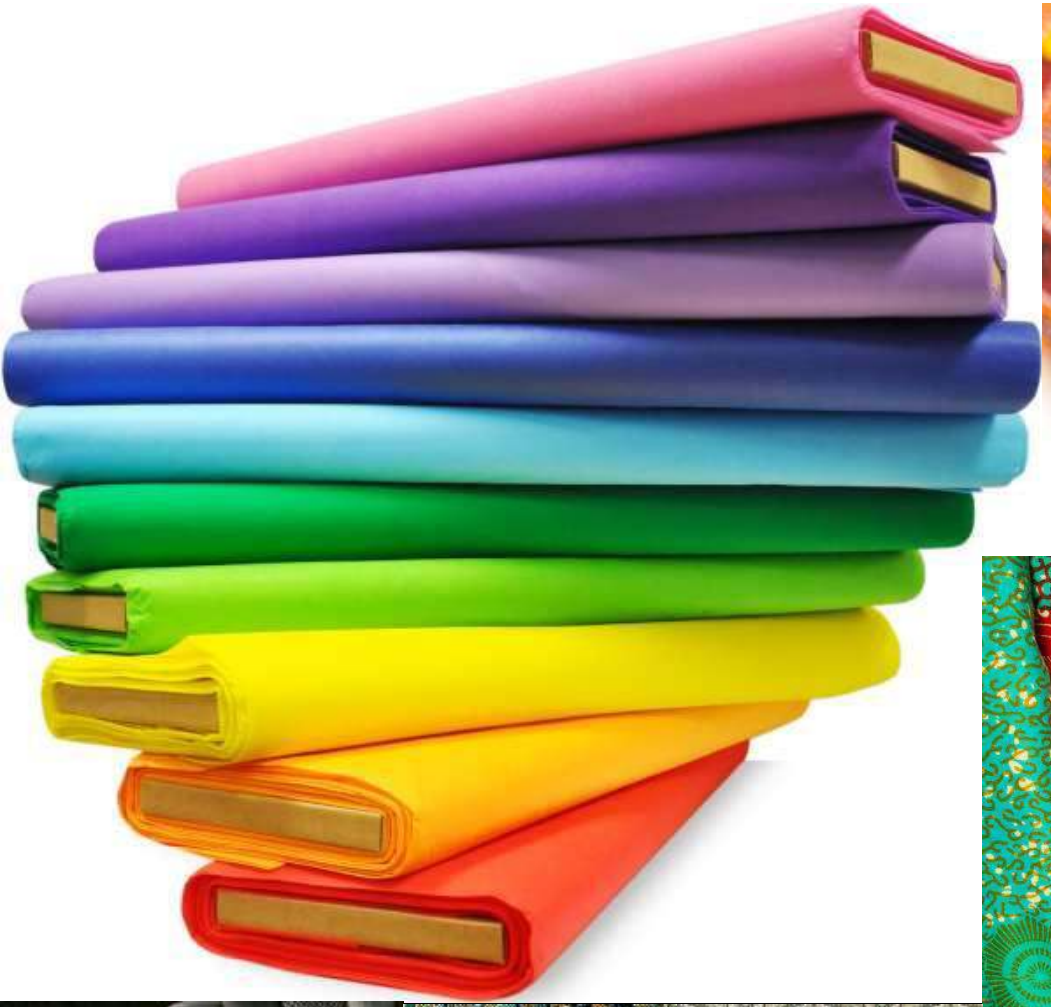

Thus, the problem (or question) definition:

"How do we provide parallel access to scientific data, through a multitude of protocols and give the illusion that the data is 'next to' the applications, on a budget, keeping the right data near the right type of computational infrastructure?

MeDiCI

UQ'S DATA FABRIC OF CHOICE

The University of Queensland's Metropolitan Data Caching Infrastructure (MeDiCI) spans urban Brisbane and provides seamless access to data regardless of where they are created, manipulated and archived.

MeDiCI holds copies of data on campus until they are not required for some time. Data are moved between on and off-campus storage on demand without user involvement. MeDiCI is underpinned by SGI DMF, Oracle StorageTek, DDN GridScaler storage, and IBM Spectrum Scale technologies.

CAMPUS COMPUTE & INSTRUMENTS

LOCAL CACHE STORAGE

LONG-TERM STORAGE

sgi ORACLE STORAGETEK IBM DDN STORAGE THE UNIVERSITY OF QUEENSLAND RCC RESEARCH COMPUTING CENTRE

Home, @ Polaris,
Springfield

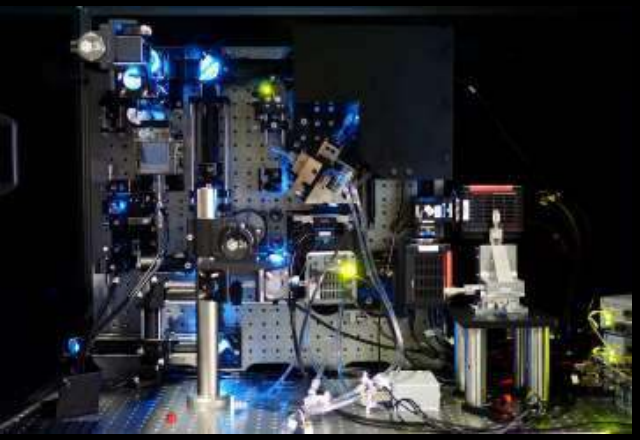

S3, Swift, Ceph,NFS, NSD,
gFTP, SMB...

Cache, @ St Lucia, BNE

SpectrumScale AFM (*cache*)

{Parallel IO via NSD protocol}

SpectrumScale AFM (*home*)

The family, with cache!

CAI

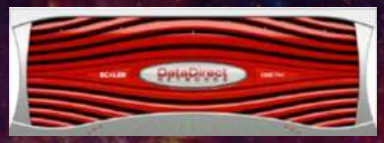

QBI

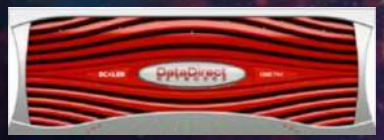

UQ ITS

IMB

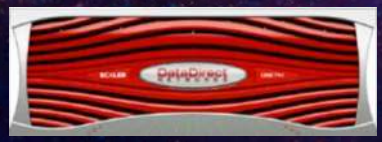

AIBN


BE
AUTHENTIC

When things go wrong…

Things that went *wrong*…

- The namespace ID mapping was inconsistent between cache and home
- The code didn’t exist in AFM to map between sites.
- We didn’t have a network big enough to transport that much IO around the campus.
- Our users started using the technology in *unusual ways.*
- Shared responsibilities for a big and complex beast.

Scale cluster “A”
using UQ creds

Scale cluster “B”
using *other* creds

Back at UQ

mmname2uuid
mmuuid2Name

Out at Polaris

uqjcarr1

someOtherName

Turns out, all that code was *missing* from SpectrumScale.

- We had, at best, 10GbE between our buildings and around the campus.
- Not made for the parallel IO aggression of SpectrumScale AFM over the NSD protocol.
- Needed to spawn an entire mini-project to upgrade campus networks for big storage IO to 40G around the “ring” of nodes.
- And even that was not enough.
- Rob and his team invoked a project and funding for the campus to put 40G/100G. That is now happening, thanks to our wonderful network team. We are (slowly) collapsing the initial 40G mini-fabric into a fresh 100G fabric.

Things we assumed users would do as per our mental model.

User puts data in cache from instruments to send to a supercomputer, at remote site

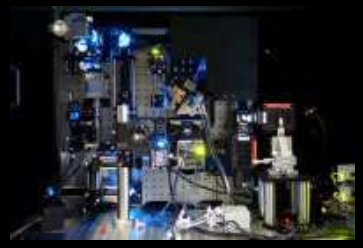

User processes data out at remote site on said supercomputer

Things people actually did, breaking our mental model.

User puts data in cache from instruments. They start processing on a supercomputer *locally*.

Simultaneously, they start using the storage fabric to process other "bits" of the outputs of the run on the other supercomputer for an additive workflow.
[culminating in the fabric becoming a means for both supercomputers to work on the same tasks at the same time]

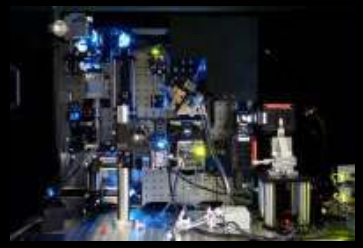

Same data namespace
ended up everywhere.

That much, was *intentional*.

IMB

QBI

RCC

As a result, user could leverage
every bit of the compute everywhere
simultaneously, if their workflow
is smart enough…

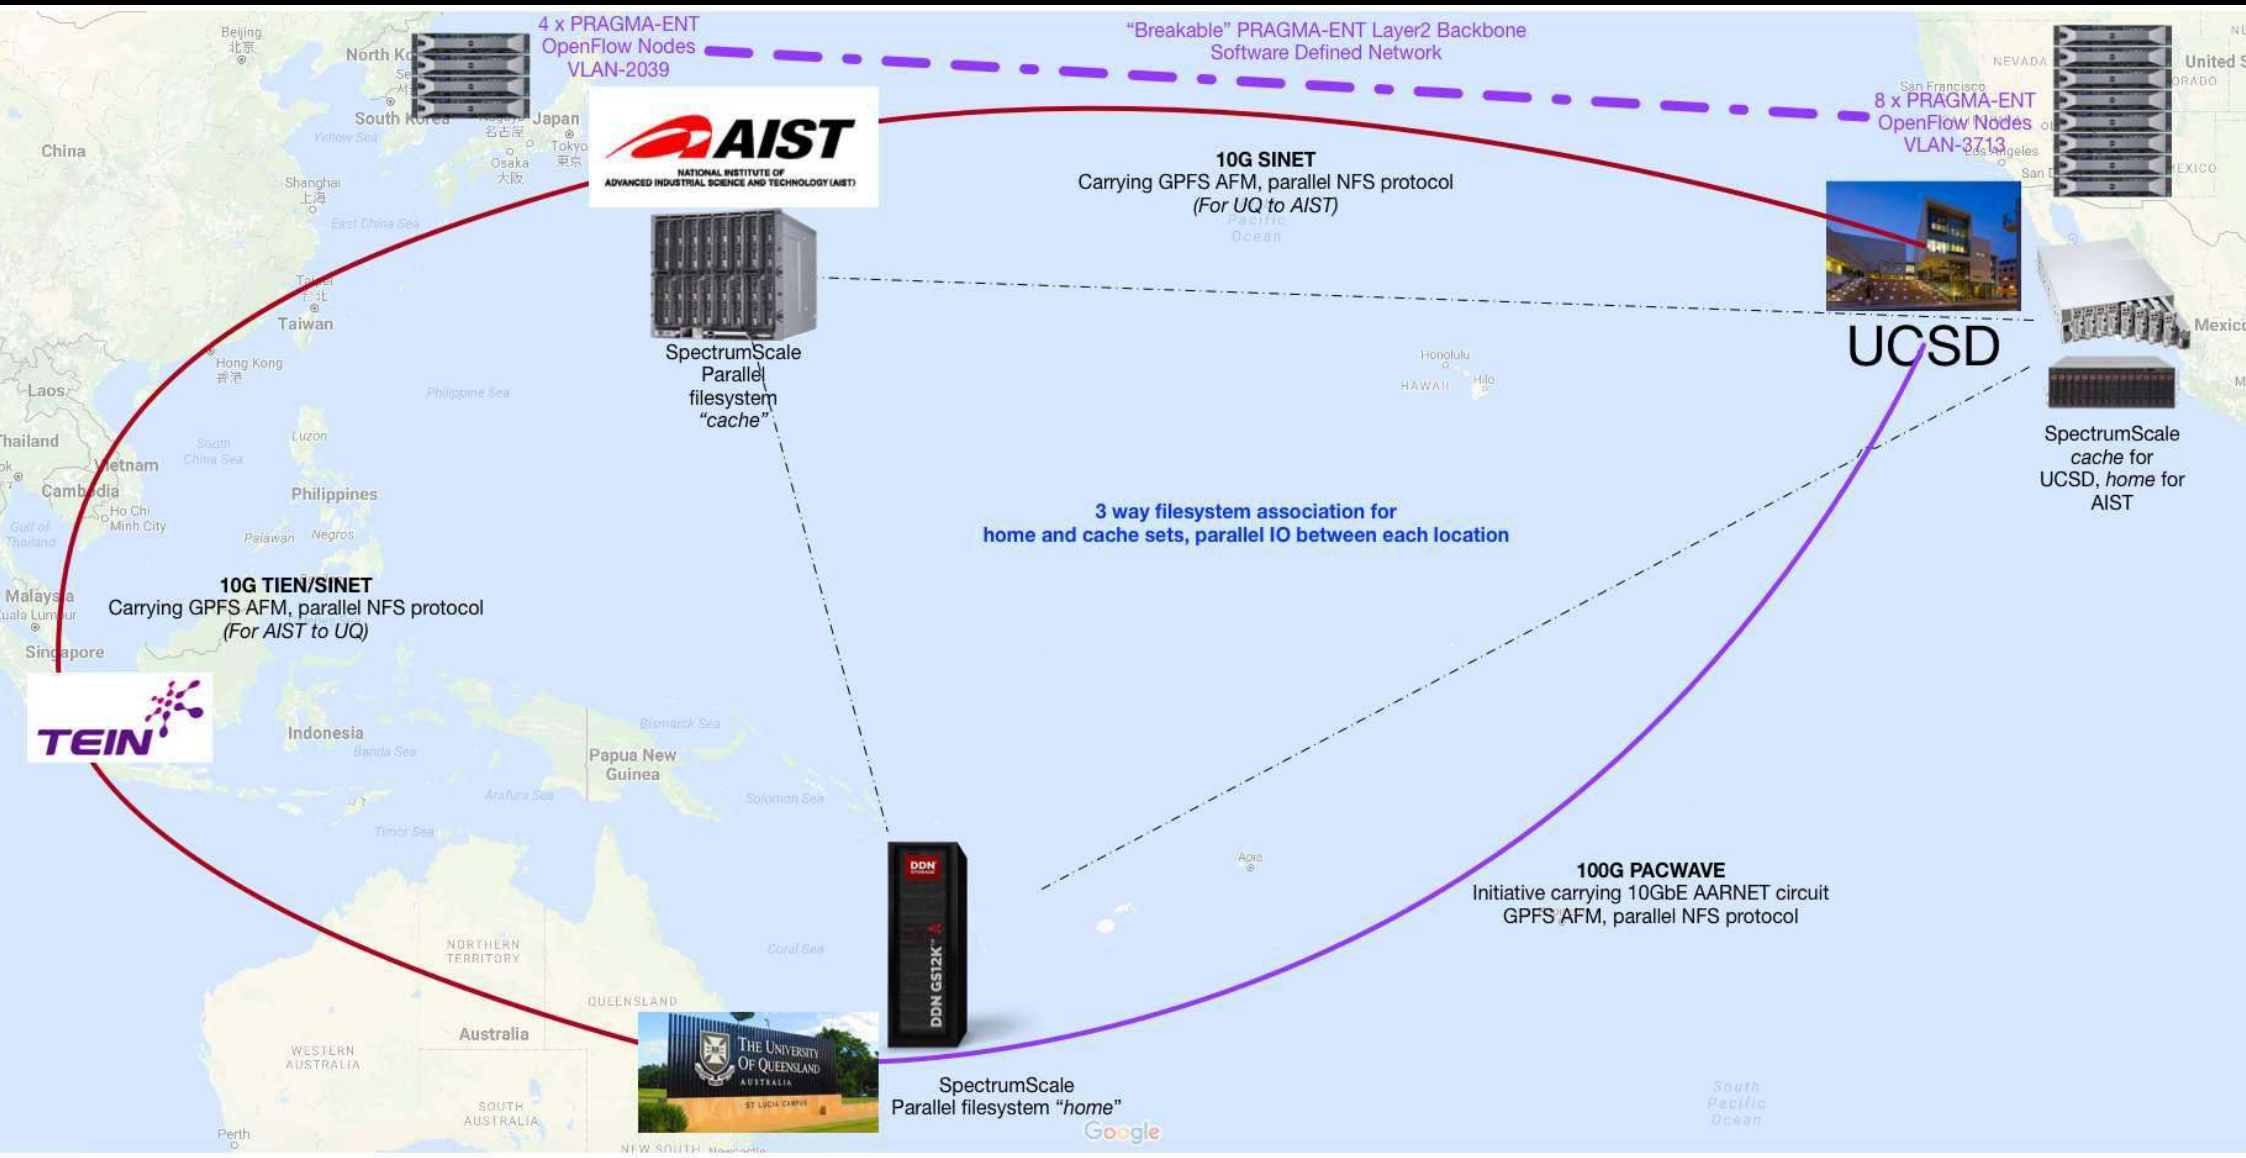
4 x PRAGMA-ENT
OpenFlow Nodes
VLAN-2039
"Breakable" PRAGMA-ENT Layer2 Backbone
Software Defined Network
8 x PRAGMA-ENT
OpenFlow Nodes
VLAN-3713
AIST
NATIONAL INSTITUTE OF ADVANCED INDUSTRIAL SCIENCE AND TECHNOLOGY (AIST)
10G SINET
Carrying GPFS AFM, parallel NFS protocol
(For UQ to AIST)
SpectrumScale
Parallel
filesystem
"cache"
UCSD
SpectrumScale
cache for
UCSD, home for
AIST
3 way filesystem association for
home and cache sets, parallel IO between each location
10G TIEN/SINET
Carrying GPFS AFM, parallel NFS protocol
(For AIST to UQ)
TEIN
100G PACWAVE
Initiative carrying 10GbE AARNET circuit
GPFS AFM, parallel NFS protocol
DDN G512K
THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
ST LUCIA CAMPUS
SpectrumScale
Parallel filesystem "home"
Google

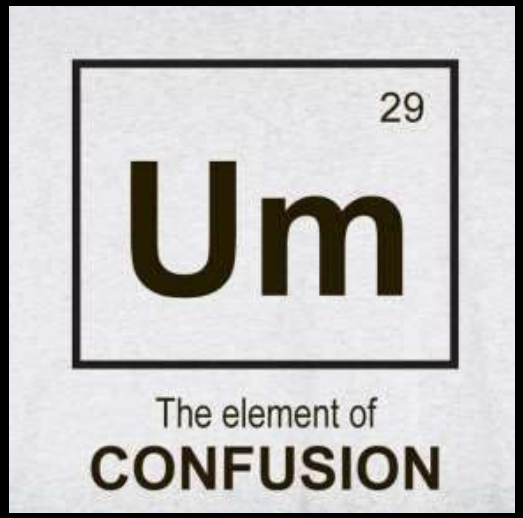

Who runs what bit? What are the rules? Who controls the inode count?

People noticed...

iAWARDS

itnews
BENCHMARK
AWARDS
2016


A preview of things to come.
This is the first time I've talked about this bit, publicly...

What comes next for MeDiCI?

- MeDiCI goes to the Amazon.
- MeDiCI goes to Azure.
- MeDiCI goes *south…..(to NCI).*

Why is that important?

- Our users want to run their very big datasets down to NCI, without the hassle of rsync, cp, scp or data locality drama.
- We want to "spring up" cache/home relationships close to every corner of the world, so that our collaborators can transit instrument data close to _them_, then have it transit back to us in real time – without them having to worry about the movement bit.
- A little bit more elegant than plain old *ascp*
- Easier to stand up "IO Pods" in these public clouds, then tear them down, when done – instead of building infrastructure that is only used for X, Y, or Z experimental data movement.

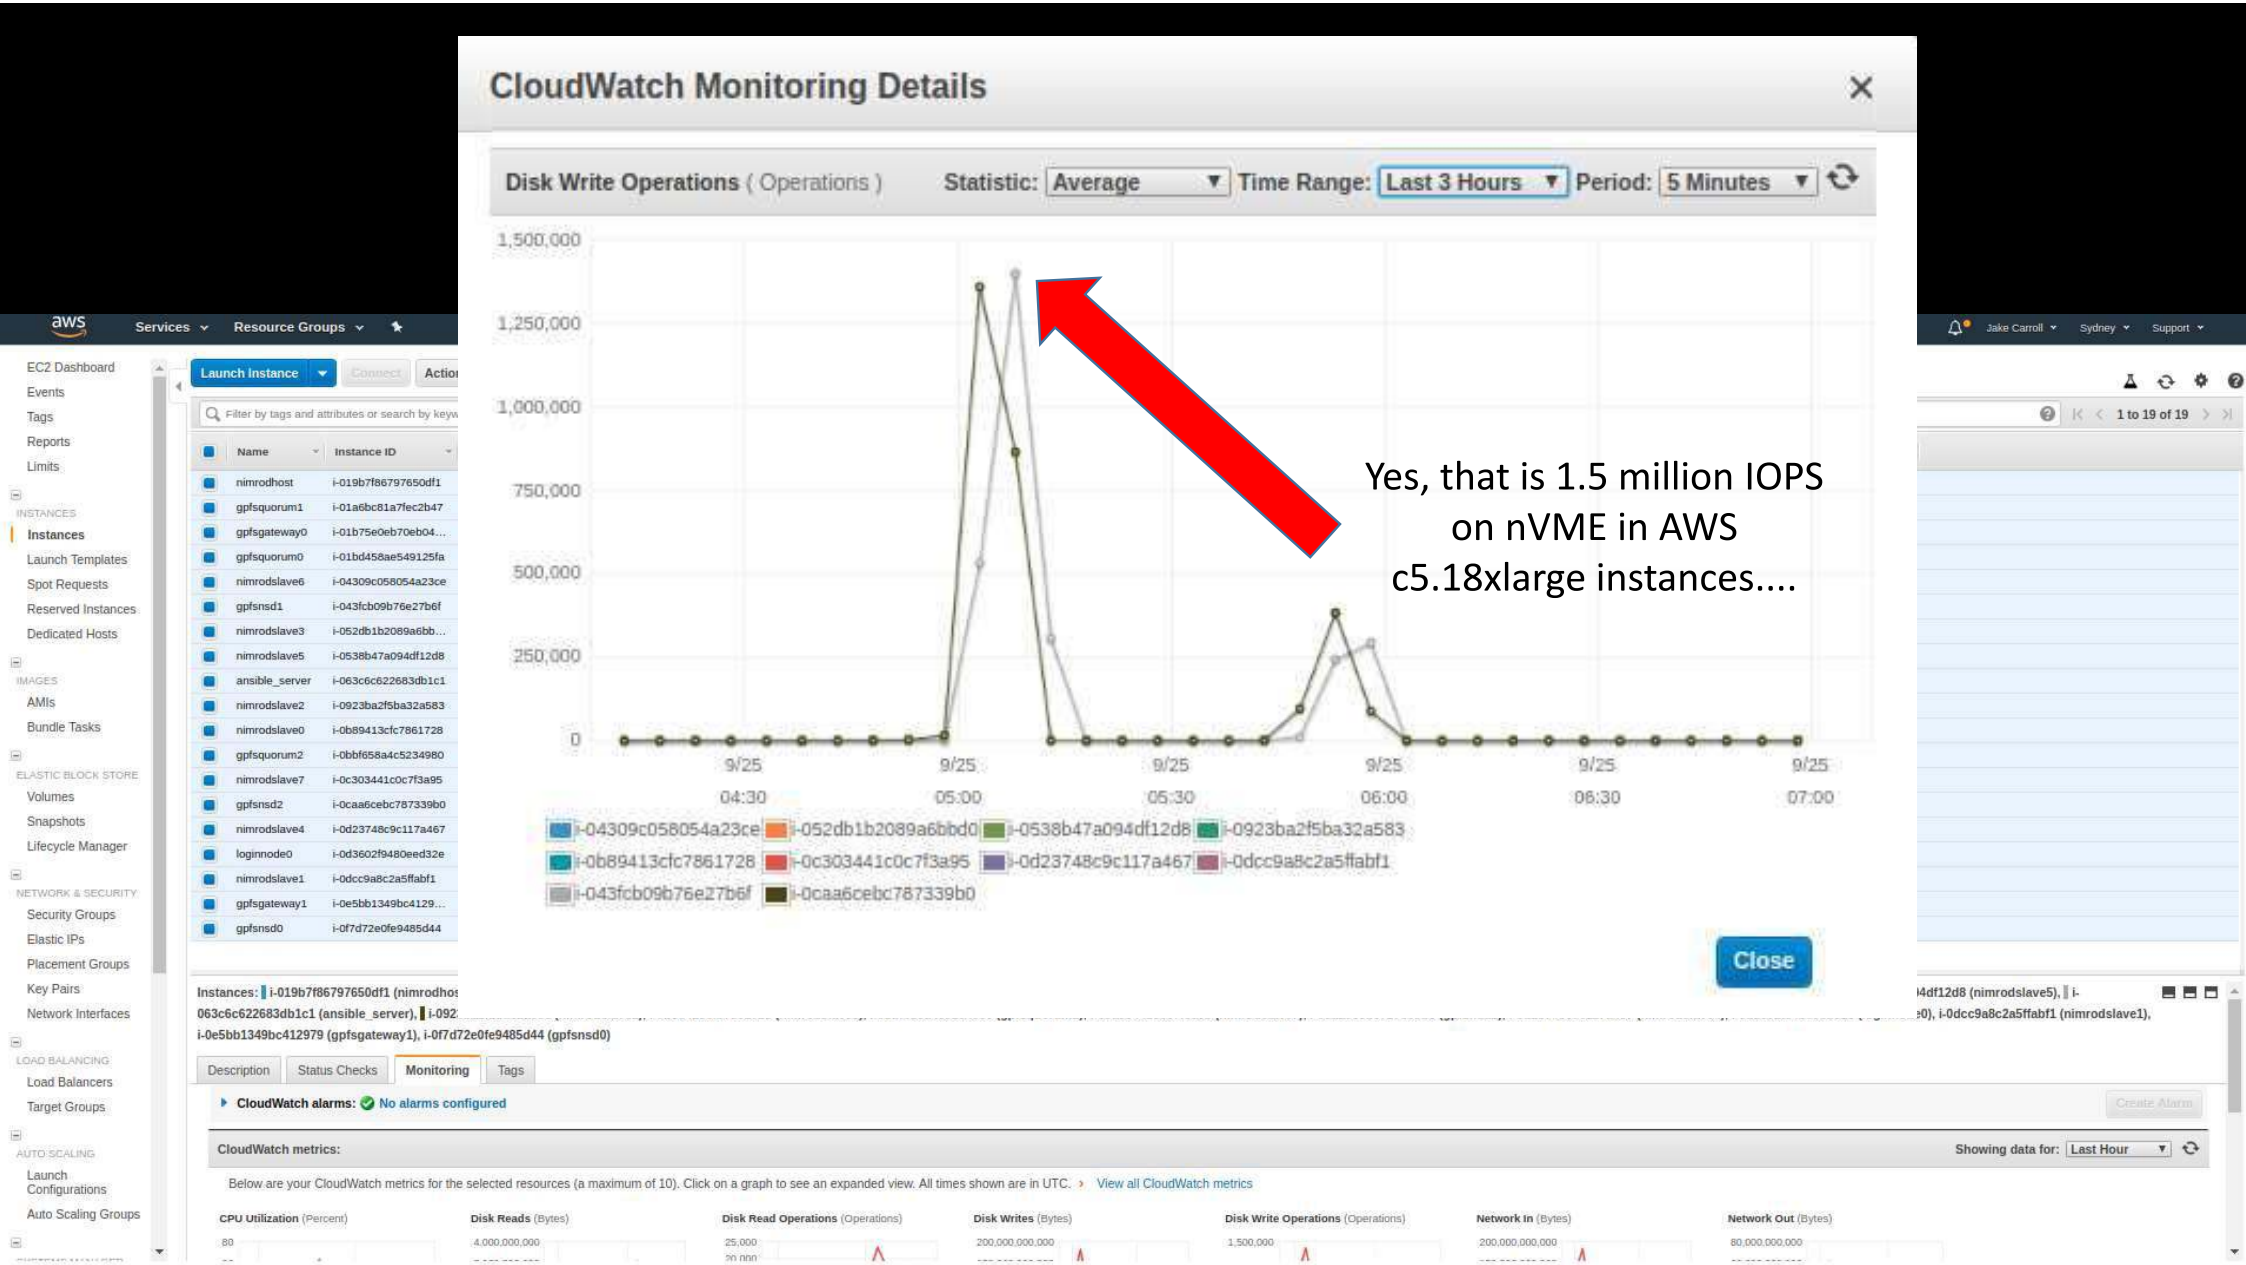
CloudWatch Monitoring Details
Disk Write Operations (Operations)
Statistic: Average
Time Range: Last 3 Hours
Period: 5 Minutes
1,500,000
1,250,000
1,000,000
750,000
500,000
250,000
0
9/25 04:30
9/25 05:00
9/25 05:30
9/25 06:00
9/25 06:30
9/25 07:00
i-04309c058054a23ce
i-052db1b2089a6bbd0
i-0538b47a094df12d8
i-0923ba2f5ba32a583
i-0b89413cfc7861728
i-0c303441c0c7f3a95
i-0d23748c9c117a467
i-0dcc9a8c2a5ffabf1
i-043fcb09b76e27b6f
i-0caa6cebc787339b0
Yes, that is 1.5 million IOPS on nVME in AWS c5.18xlarge instances....
Close
Launch Instance
Description
Status Checks
Monitoring
Tags
CloudWatch alarms: No alarms configured
CloudWatch metrics:
Showing data for: Last Hour
Below are your CloudWatch metrics for the selected resources (a maximum of 10). Click on a graph to see an expanded view. All times shown are in UTC.
View all CloudWatch metrics
CPU Utilization (Percent)
Disk Reads (Bytes)
Disk Read Operations (Operations)
Disk Writes (Bytes)
Disk Write Operations (Operations)
Network In (Bytes)
Network Out (Bytes)

Stuff getting busy, in H7C, AWS...1500 Skylake cores in AWS c5.18xlarge against MeDiCI filesystem, linked back in real time to the Polaris data centre over AFM...

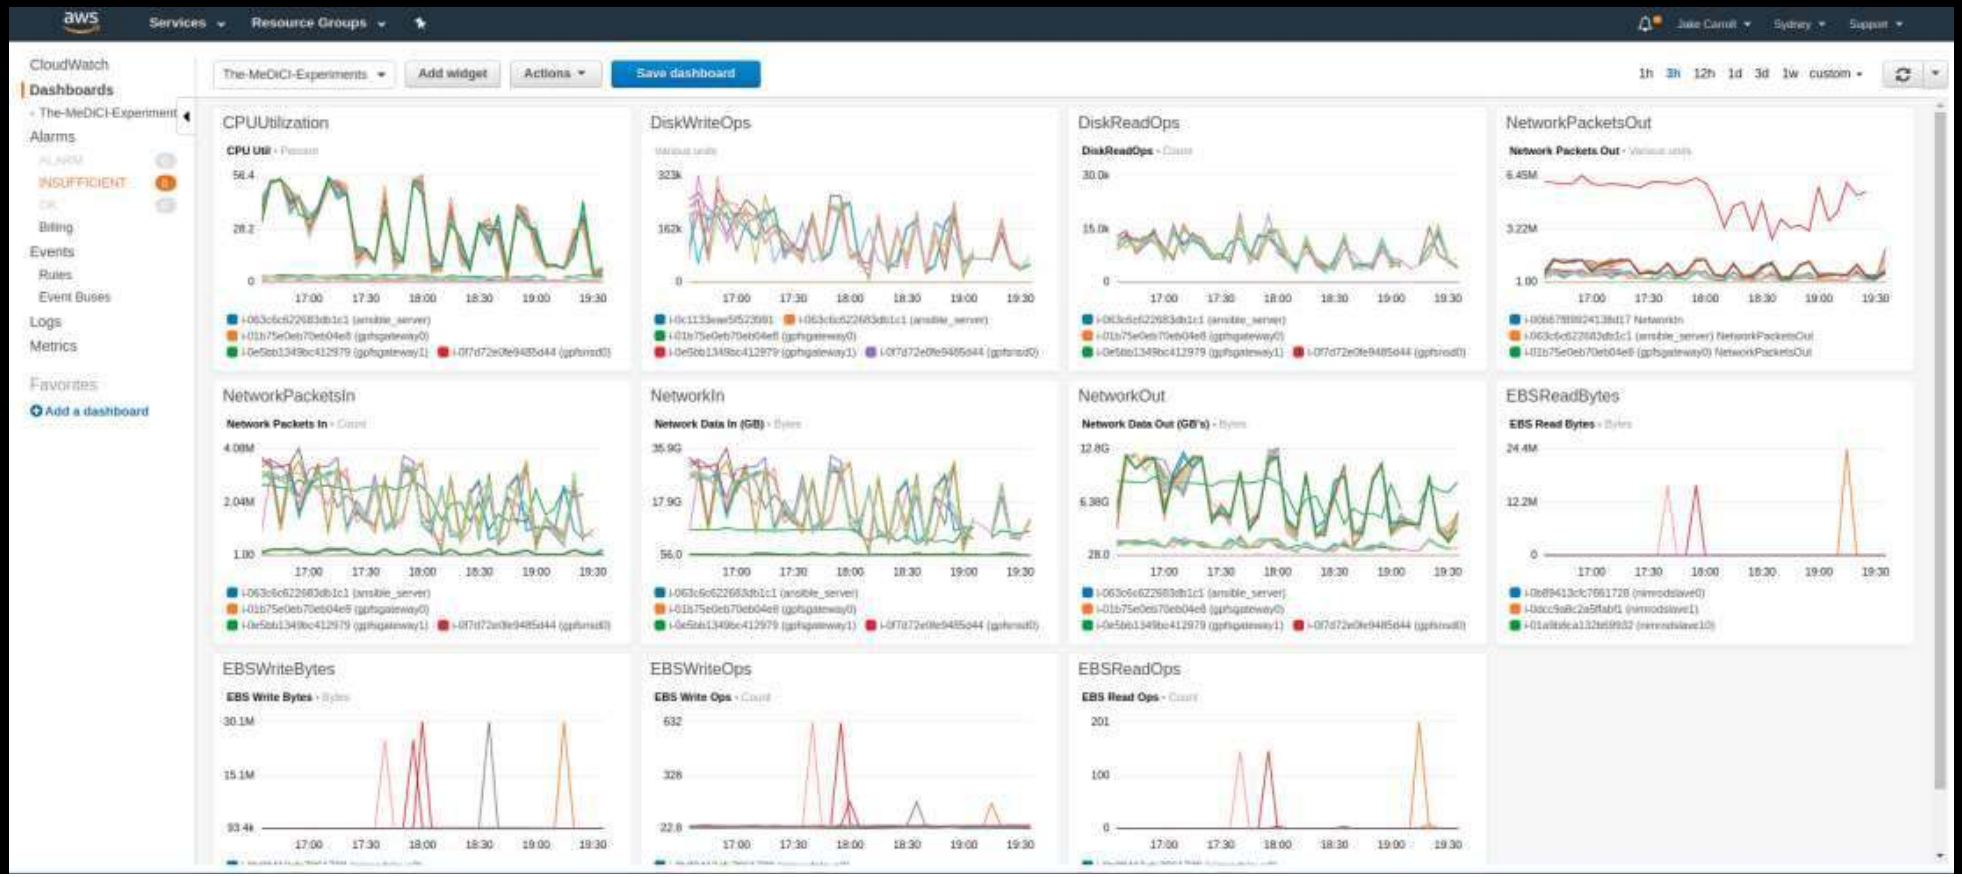

Dynamic IO pods, all over the world, near the workloads, when it counts...to get back to our supercomputers at UQ

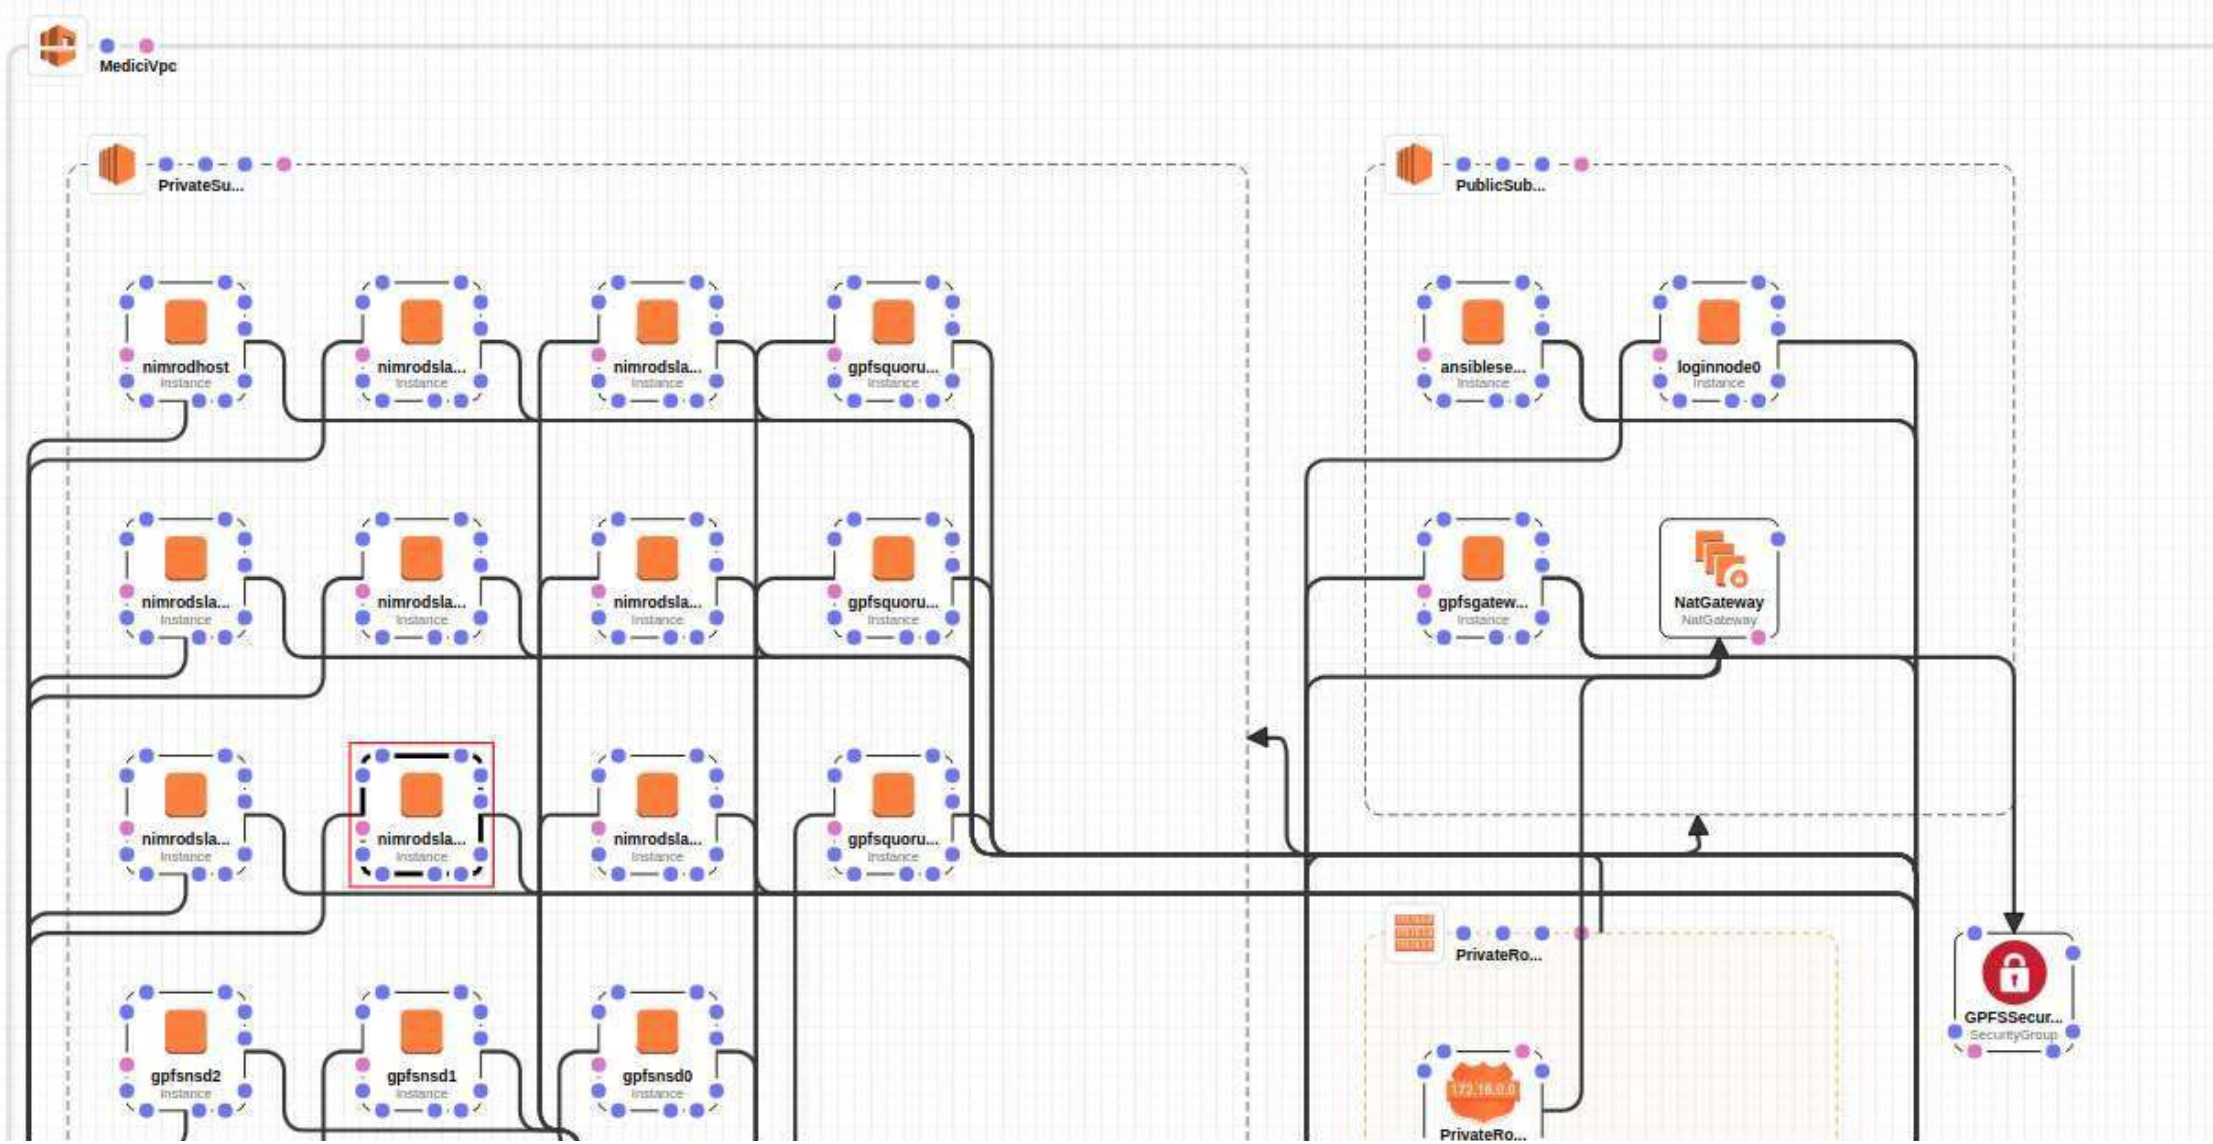

A cast of…quite a few people…

- David Abramson, RCC
- Michael Mallon, RCC/QCIF
- Leslie Elliot, ITS
- Chao Jin, RCC
- Hoang Nguygen, RCC
- Minh Dinh, RCC
- Zane Van Iperen, RCC
- The network team at ITS:
  - Alan Ewer, Felix Li, Jooil Lee, Pete Keeffe, Mike Rawle, Scott Phillips
- Stephen Bird, QCIF
- Irek Porebski, QBI
- Doug Stetner, IMB
- Matthew (Beanie) Bryant, IMB
- Chris Myers, AARNET
- Larry Smar, UCSD
- Venkatesawara Puuvada, IBM Spectrum Scale Engineering, India
- Andrew Beattie, IBM Spectrum Scale team, Australia

Thanks...

- David Abramson – for **adventure**, crazy, courage and vision.
- Rob Moffatt – for **trust**, for support, both financial and symbolic.
- The three institute directors – Pankaj, Brandon and Alan...
- All of the teams involved, all across our organisation.
- Our users and community – that pushed us into this crazy place, that changed the data-movement world.