



THE UNIVERSITY  
OF QUEENSLAND  
AUSTRALIA

CREATE CHANGE

# Caches all the way down: Building a wide area data fabric with Spectrum Scale.

David Abramson

Director, Research Computing Centre

Professor of Computer Science

University of Queensland

[david.abramson@uq.edu.au](mailto:david.abramson@uq.edu.au)

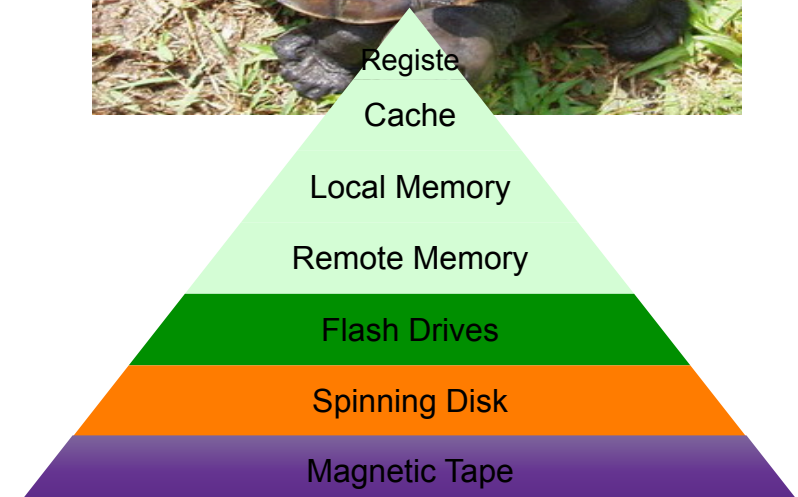
## Turtles Caches all the way down

“a jocular expression of the infinite regress problem in cosmology posed by the "unmoved mover" paradox.

The metaphor in the anecdote represents a popular notion of the theory that Earth is actually flat and is supported on the back of a World Turtle, which itself is propped up by a chain of larger and larger turtles.

Questioning what the final turtle might be standing on, the anecdote humorously concludes that it is turtles all the way down””

[https://en.m.wikipedia.org/wiki/Turtles\\_all\\_the\\_way\\_down](https://en.m.wikipedia.org/wiki/Turtles_all_the_way_down)





THE UNIVERSITY  
OF QUEENSLAND  
AUSTRALIA

CREATE CHANGE

# A Little Context

# ***RCC: Transforming research through the effective use of data and advanced computing***

The background is a vibrant blue with a digital aesthetic. In the center-left, a globe is covered in glowing binary code (0s and 1s). To the right, a complex network diagram with various nodes and connecting lines is visible. In the bottom right corner, a computer monitor is partially shown, displaying a similar digital interface. The overall scene is illuminated with a soft, ethereal light, suggesting a high-tech, data-driven environment.

New techniques and technologies allow us to both ask, and answer, bold new questions.

This is a significant shift in the way research is conducted

Researchers need to embrace the new ways of working

## RCC Goals

Lead innovation in e-Research Technologies

Partner in key University research priorities

Increase research uptake of leading edge equipment and software

Promote e-Research

Play leading role in national e-Research policy

Developing a 21st Century Workforce







THE UNIVERSITY  
OF QUEENSLAND  
AUSTRALIA

CREATE CHANGE

# Data Intensive Computing

## Data-Intensive Computing

Very large data-sets or very large input-output requirements  
Two data-intensive application classes are important and growing



## Infrastructure for Data Intensive Computing

### Computation

- Large amounts of main memory
- Parallel processors
- Smooth out memory pyramid


### Storage

- Significant long term storage
- Smooth out the memory pyramid
- Many views of same data
  - Parallel File System
  - Local access (POSIX)
  - Remote collaboration and sharing (Object store)
  - Sync-and-share
  - Web
  - Cloud

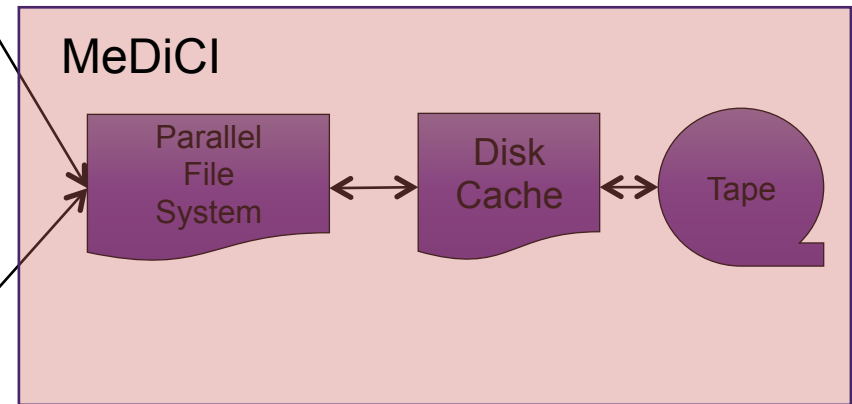



# Reference Architecture

## Cluster B



## FlashLite



Shared Memory Programming

Hierarchical File System

## MeDiCI

Centralising research data storage and computation

Distributed data is further from both the instruments that generate it, some of the computers that process it, and the researchers that interpret it.

Existing mechanisms manually move data


MeDiCI solves this by

- Augmenting the existing infrastructure,
- Implementing on campus caching
- Automatic data movement

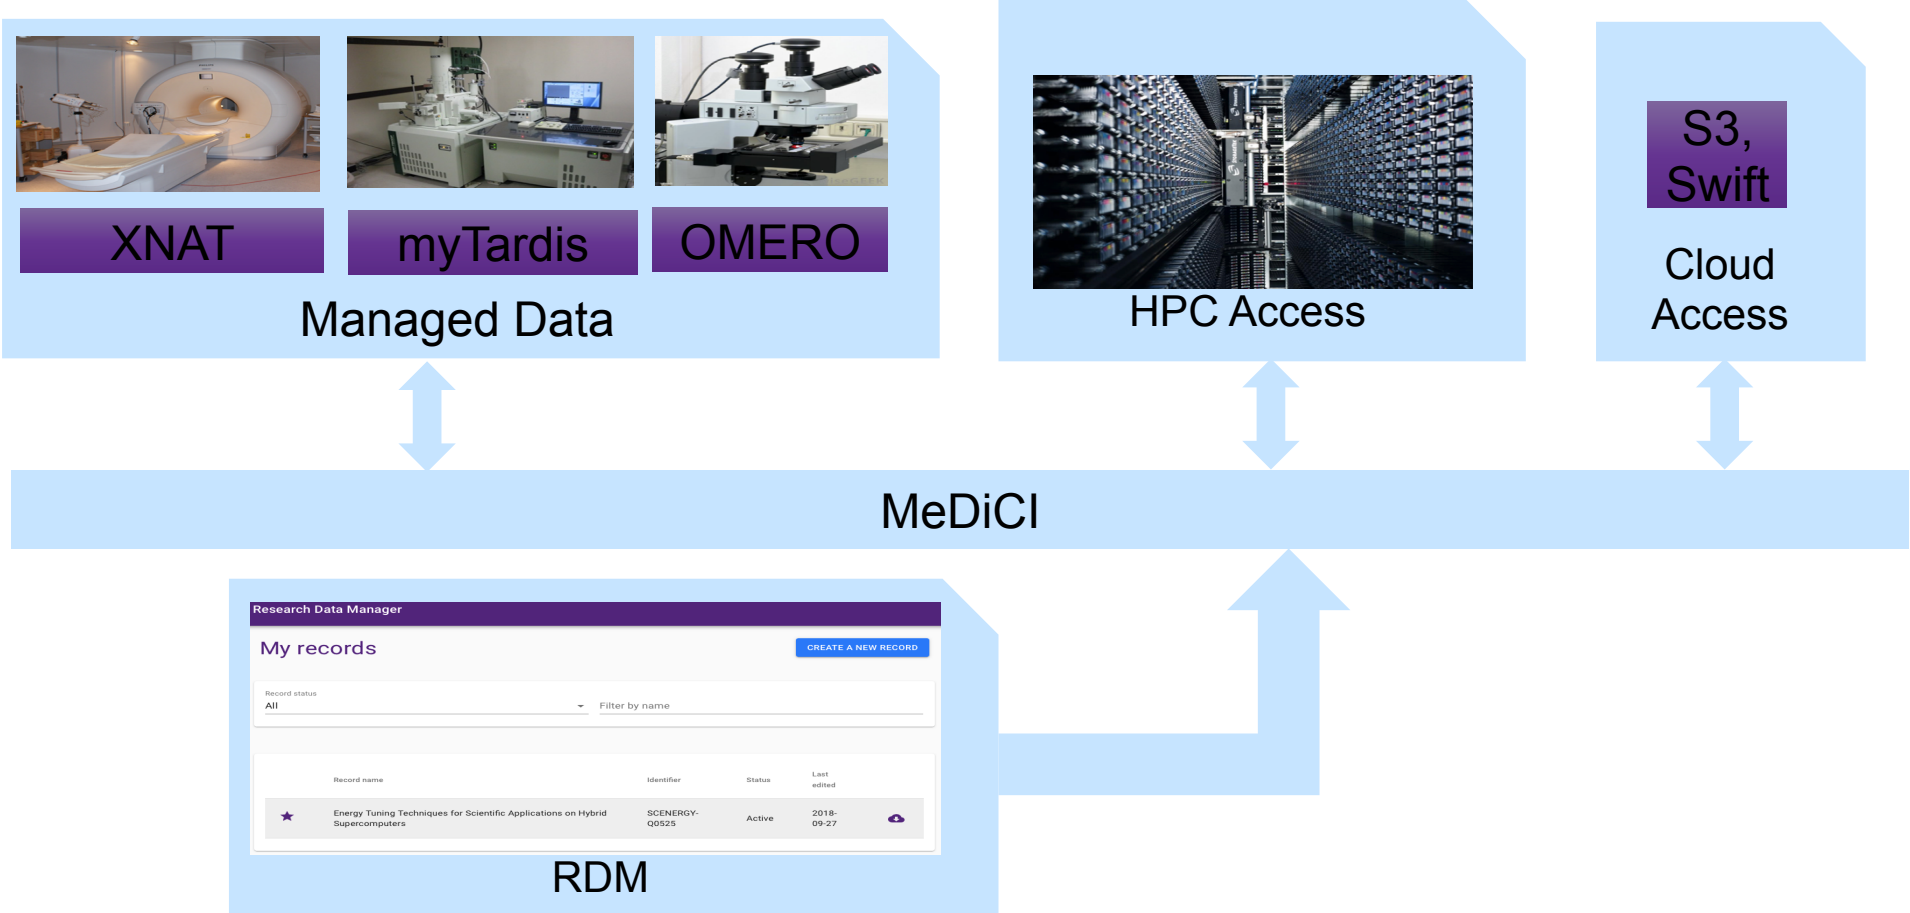
Current implementation based on IBM Spectrum Scale (GPFS) and HPE DMF




# MeDiCI unifies data access




# CAMERA: Focusing on Data

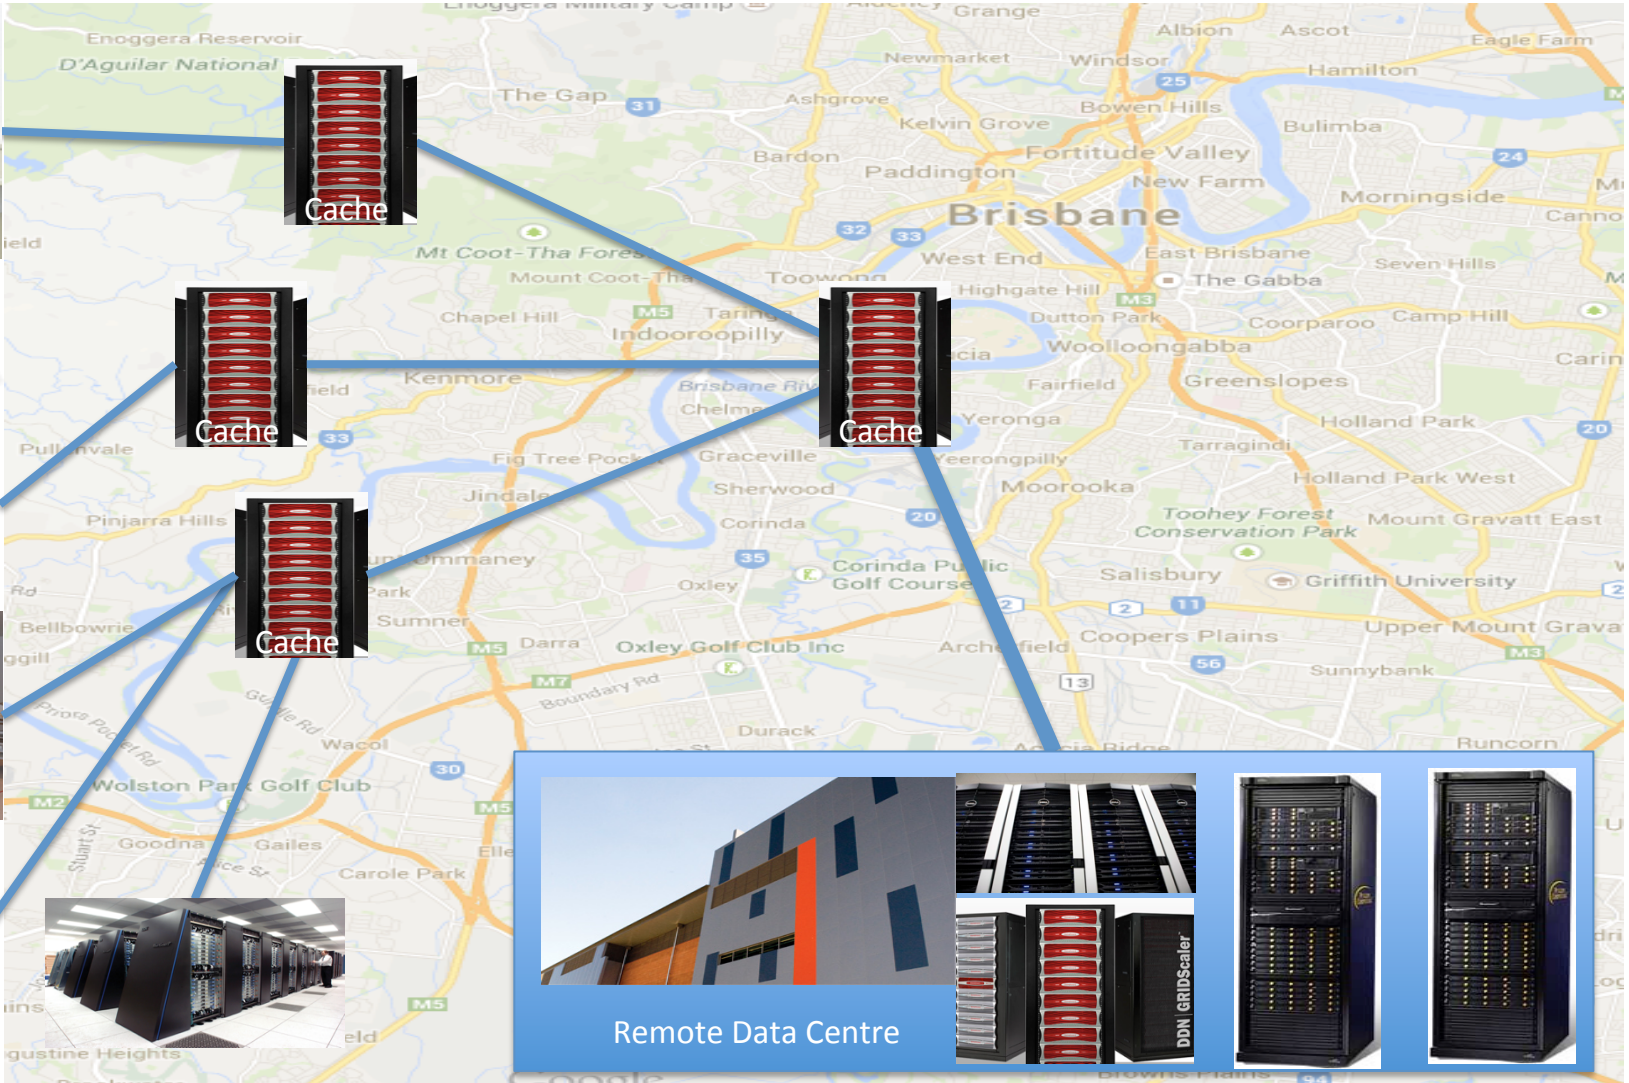


# UQ Landscape: An inconvenient truth



30 kms





Remote Data Centre

## MeDiCI as a parallel file system




FlashLite

DDN SFA12KXE

Parallel file system

## Accessing long term collections



FlashLite

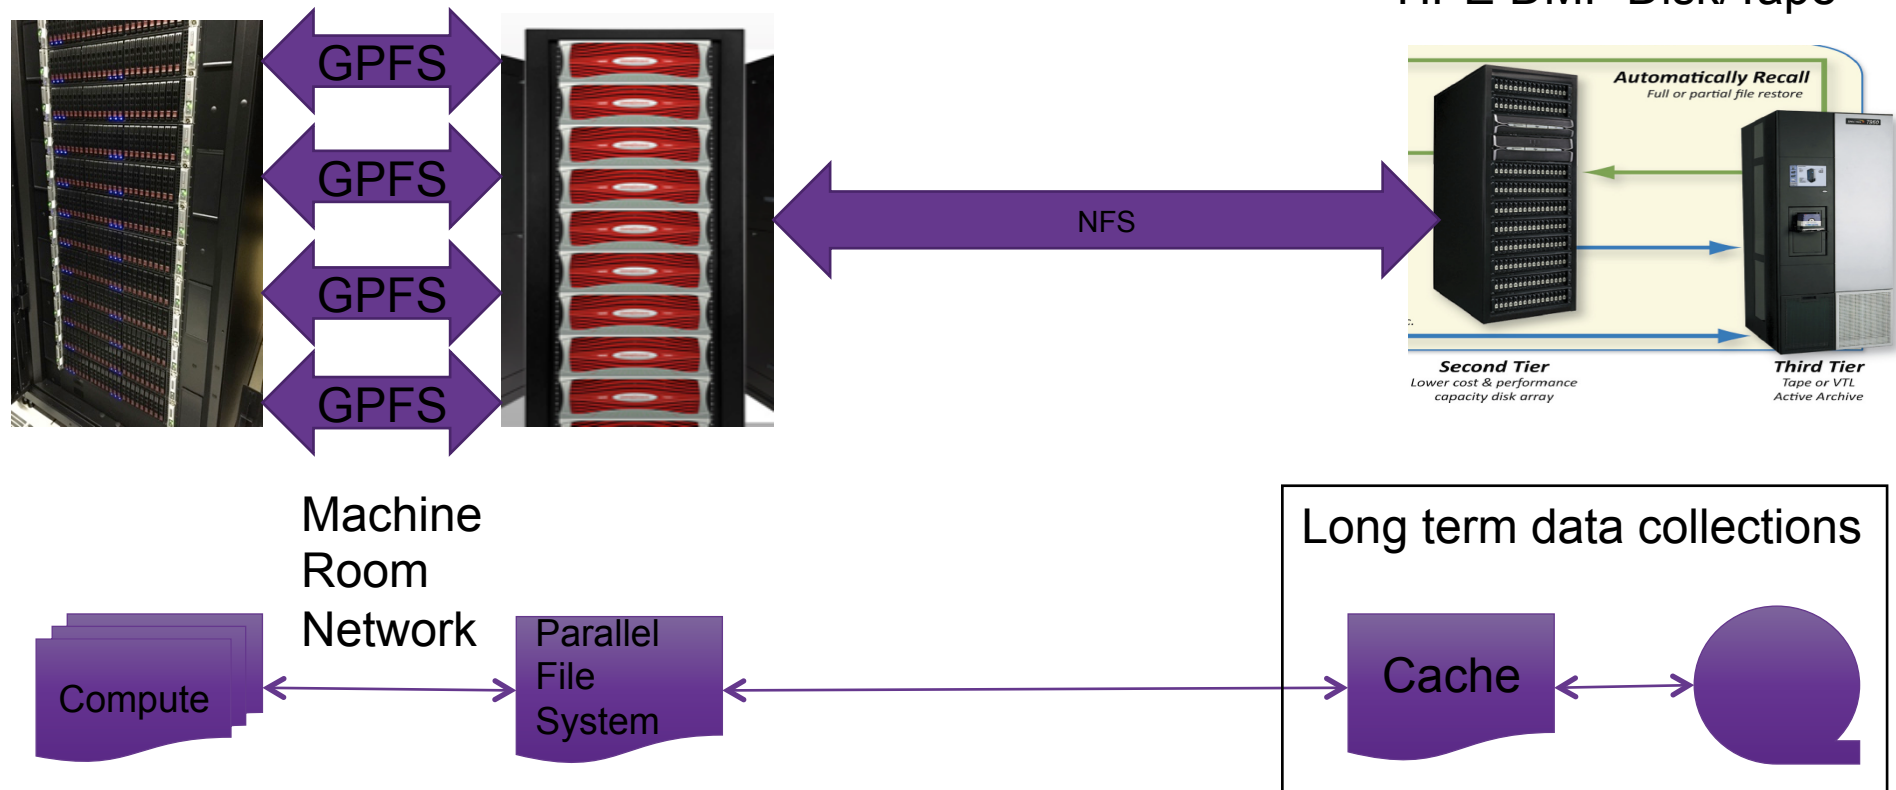
DDN SFA12KXE

HPE DMF Disk/Tape

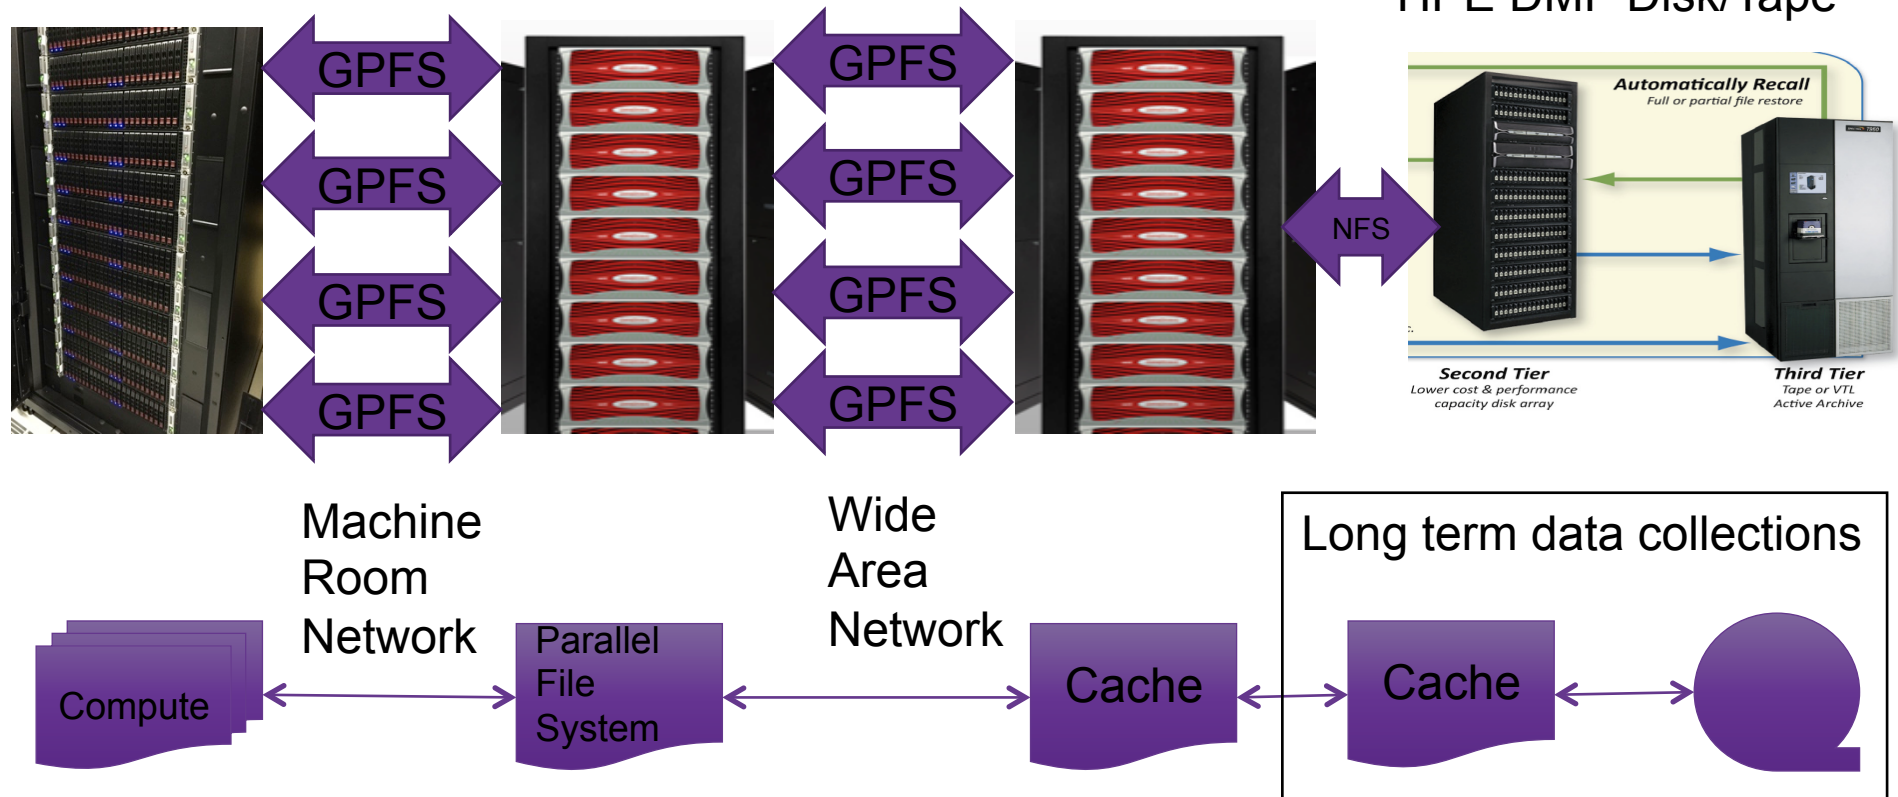
Parallel file system

Long term data collections



# MeDiCI Wide Area Architecture



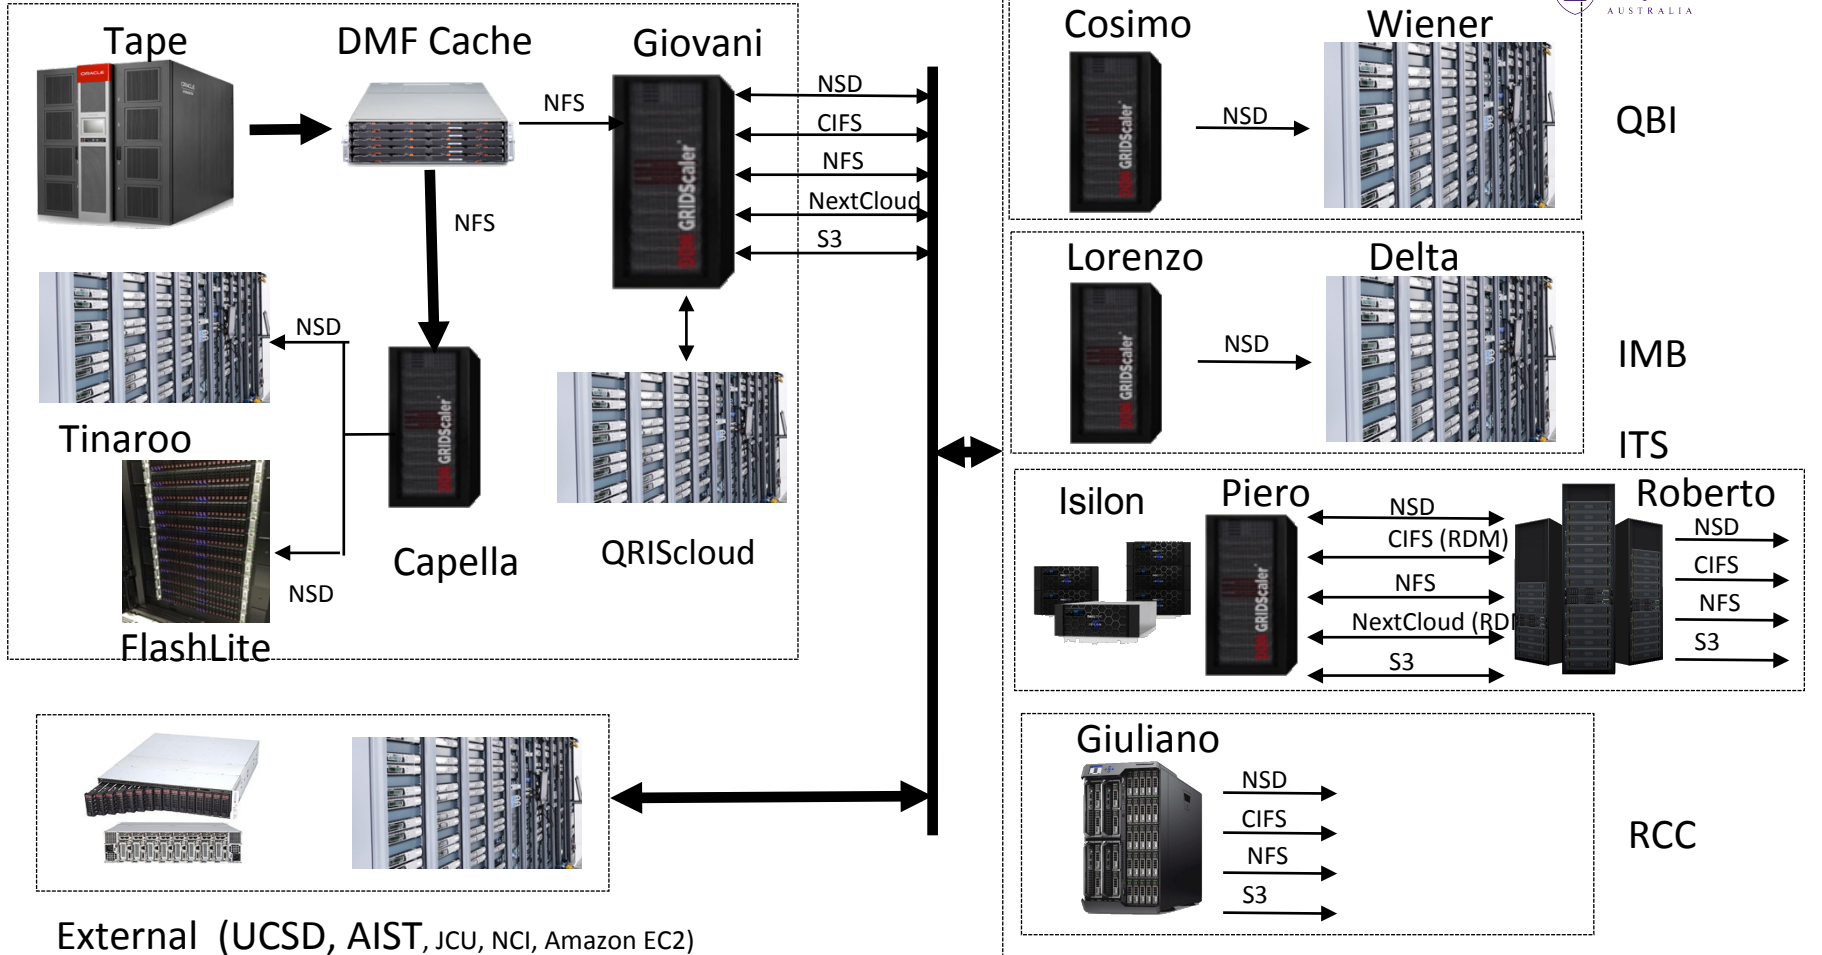
# MeDiCI Wide Area Architecture



# MeDiCI Wide Area Architecture



# Polaris Springfield



External (UCSD, AIST, JCU, NCI, Amazon EC2)

UQ St Lucia

## Building on basic architecture

A Declarative Machine Room

Leveraging Cloud Storage

Very Very Wide Area File Systems

Supporting repository stacks

Orchestrating Workflows

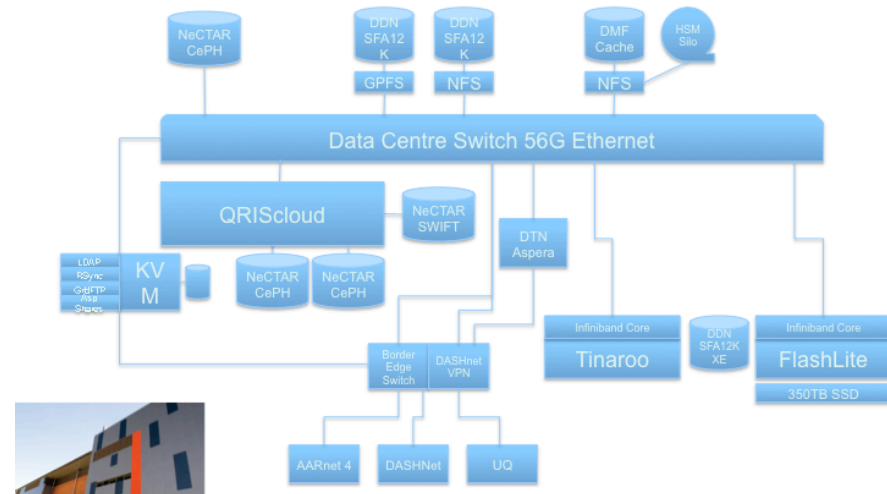
## A Declarative Machine Room?

Static allocation of disk  
and tape

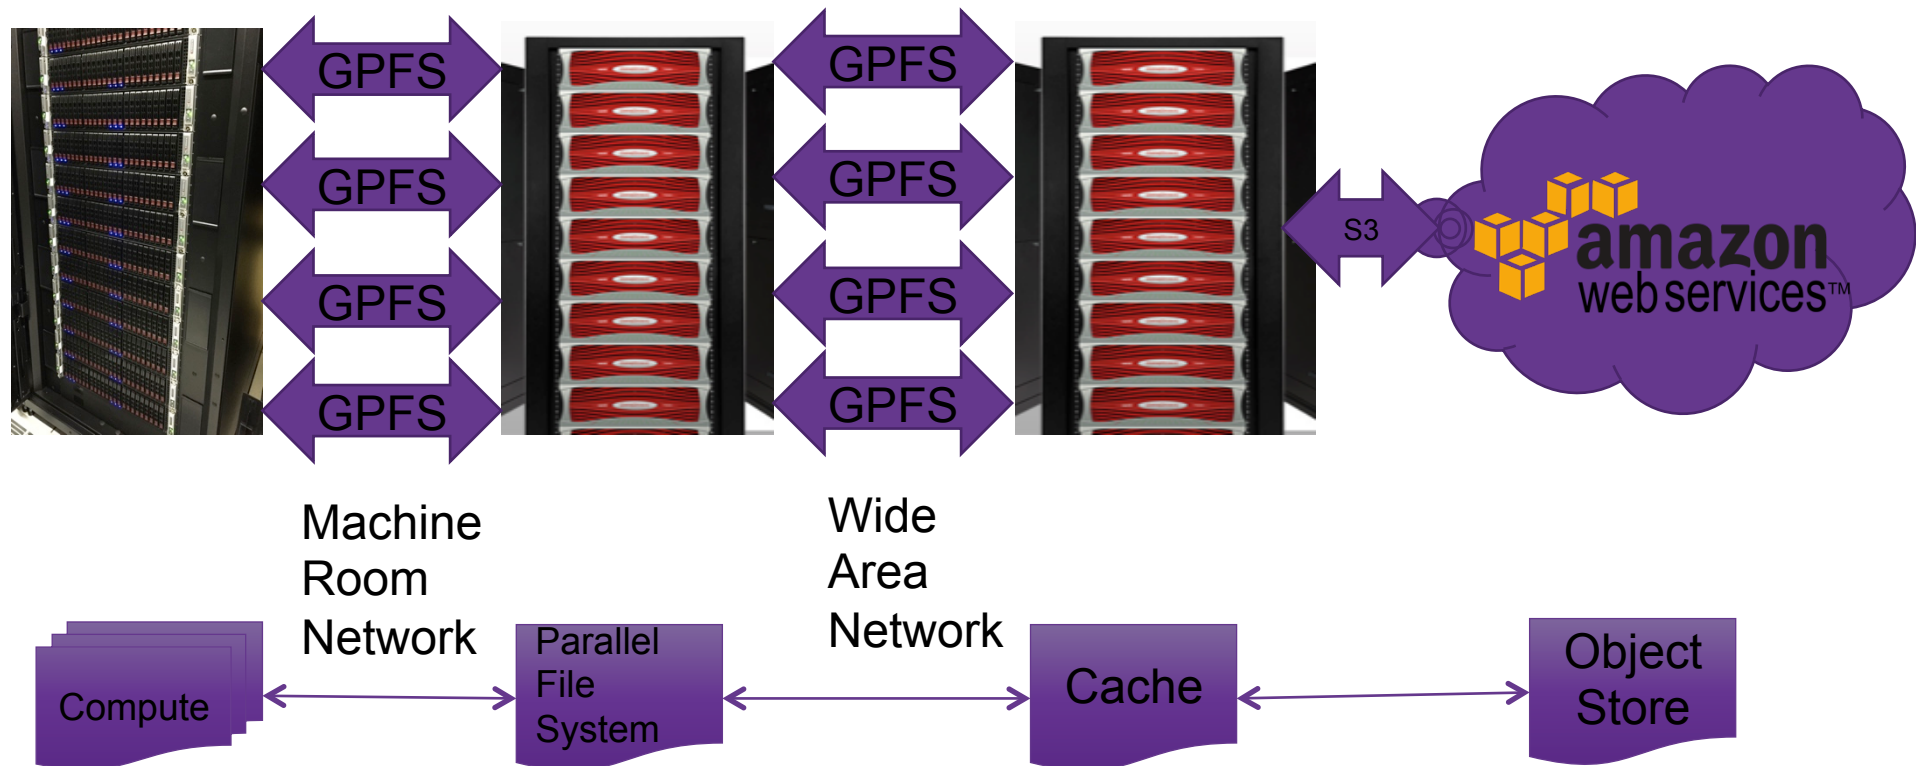
Policy driven allocation

```
RULE 'prefetch-list'
LIST 'toevict'
```

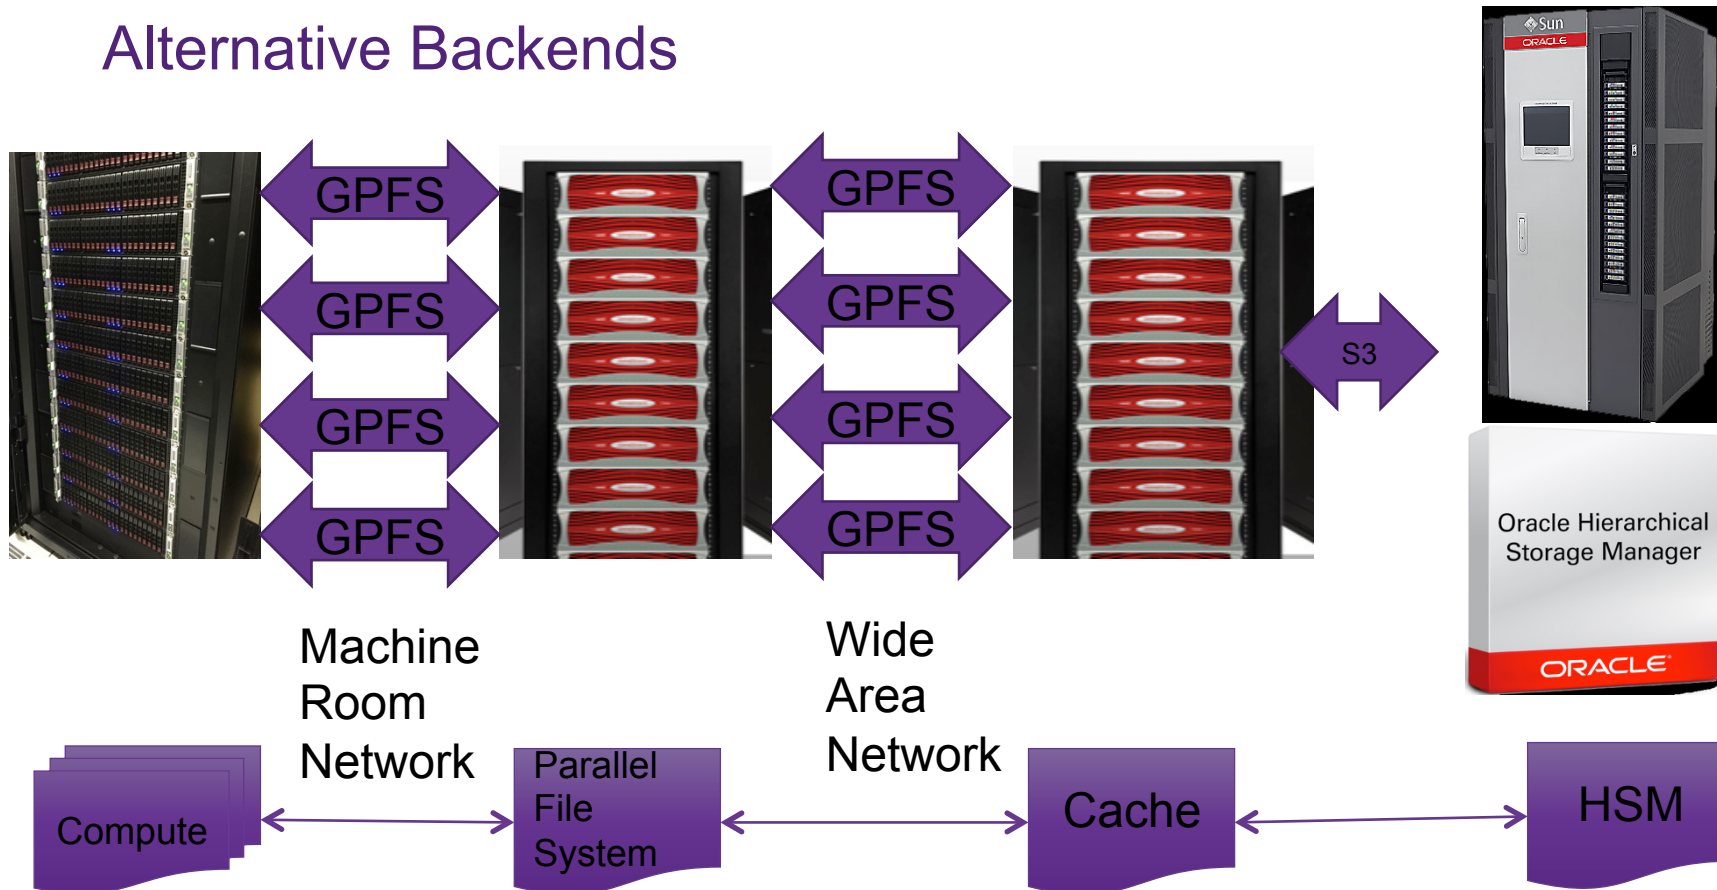
```
WHERE CURRENT_TIMESTAMP - ACCESS_TIME >
INTERVAL '7' DAYS
AND REGEX(misc_attributes, '[P]') /* only list AFM managed files
*/
```



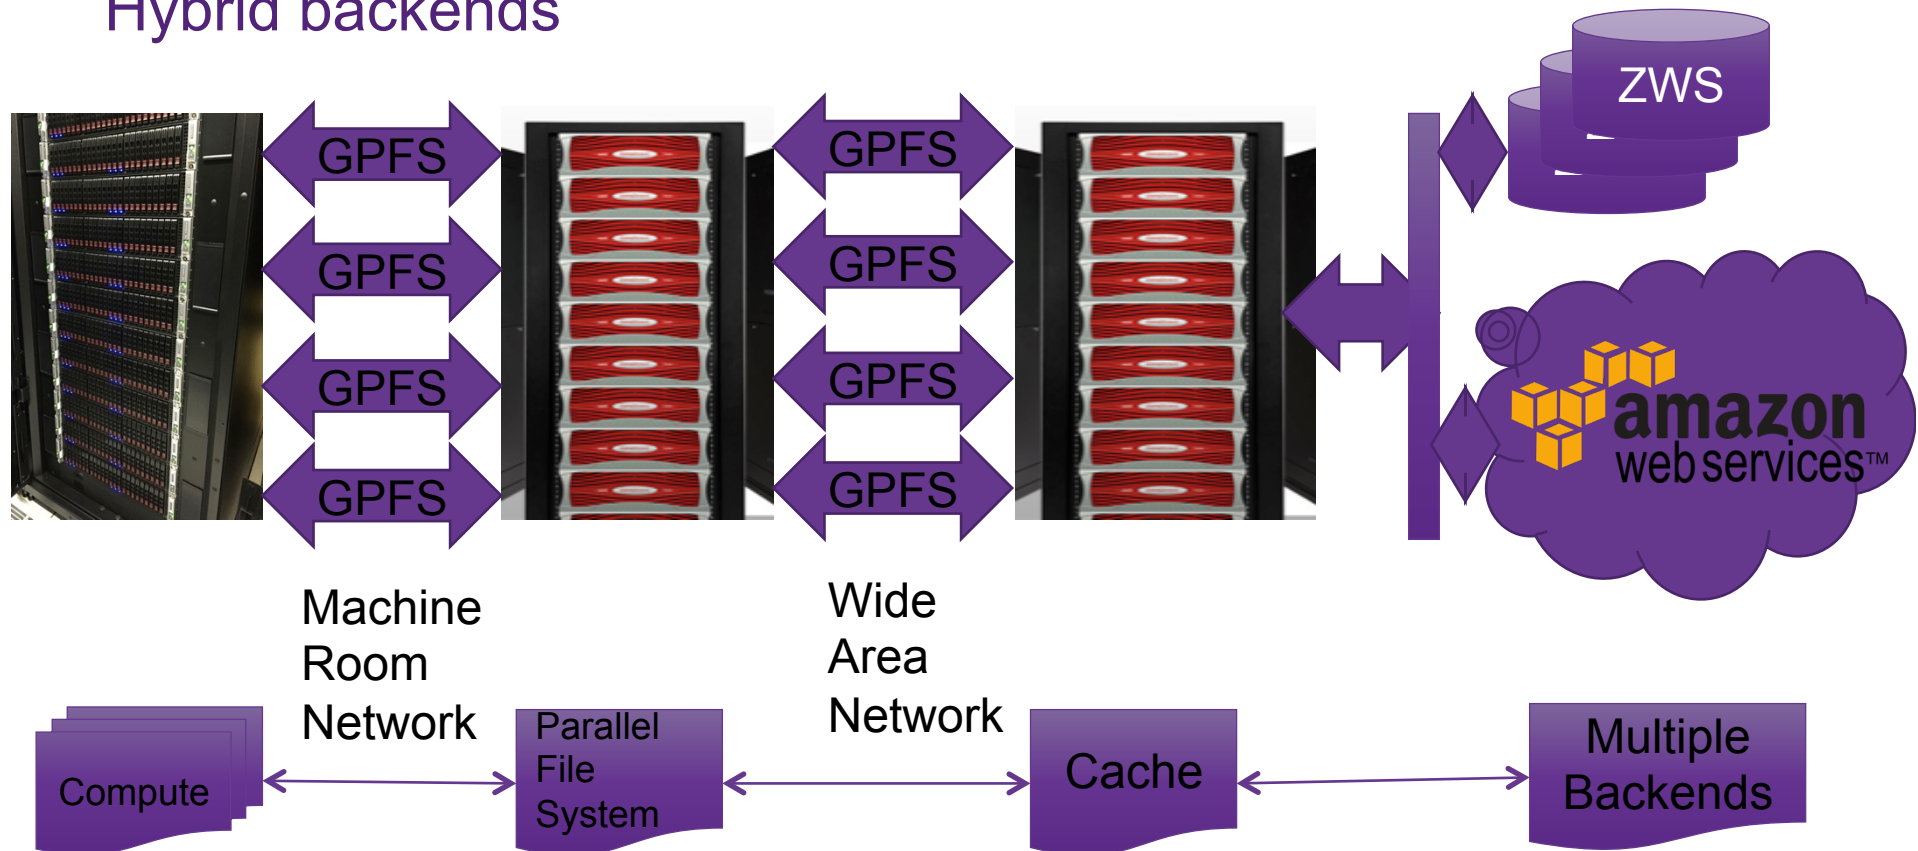
## MeDiCI Very Wide Area Architecture



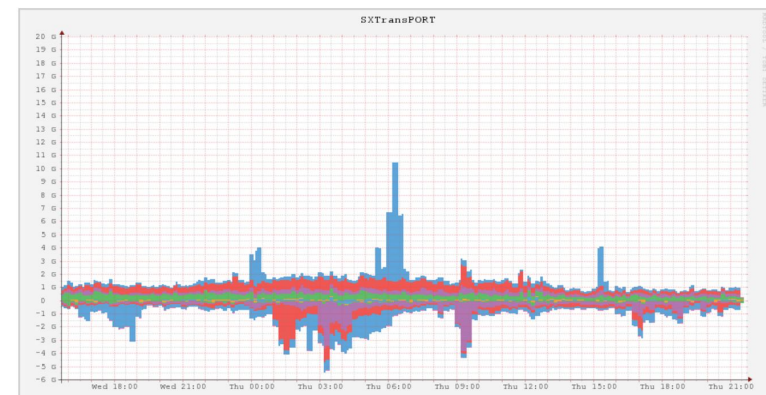
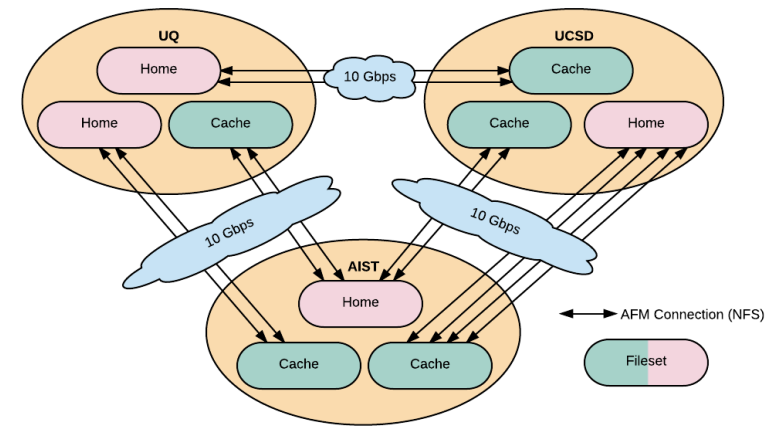
# Alternative Backends



## Hybrid backends



# MeDiCI goes East



## MeDiCI goes South and West



Remote Data Centre

AARnet X

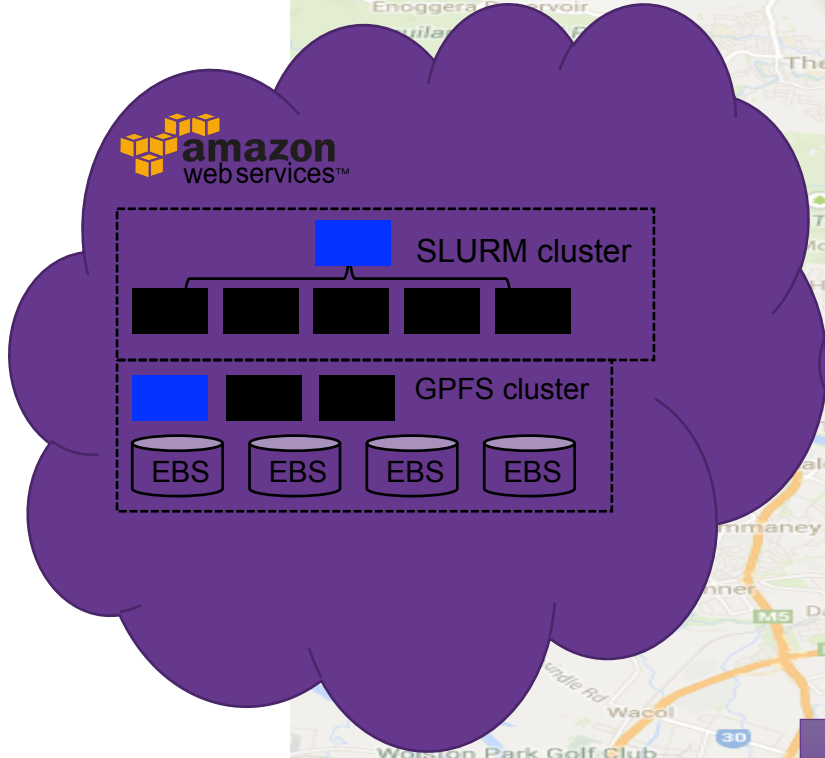


Magnus, Pawsey



Raijin, NCI

# MeDiCI goes to the Amazon



Remote Data Centre

## Conclusions

### MeDiCI

- Caches all the way down
- IBM Spectrum Scale & HPE DMF





THE UNIVERSITY  
OF QUEENSLAND  
AUSTRALIA

CREATE CHANGE

# Thank you

David Abramson | Professor and Director  
Research Computing Centre  
David.Abramson@uq.edu.au  
07 0000 000



[facebook.com/uniofqld](https://facebook.com/uniofqld)



[Instagram.com/uniofqld](https://Instagram.com/uniofqld)



[twitter.com/RCCUQ](https://twitter.com/RCCUQ)



[facebook.com/rccuq](https://facebook.com/rccuq)

[rcc.uq.edu.au](https://rcc.uq.edu.au)