

# Getting Started with Setonix

Setonix Phase 1 Release  
20 June 2022, Version 1.01



# Focus for this training

## Learning outcomes:

- Compare & contrast the existing and new supercomputer architectures.
- Determine at a high level how the Setonix architecture impacts your project.
- Discuss migration pathways to Setonix and the migration timeframe.
- Log into the new Setonix supercomputer!

## Core Migration Training Modules:



### 1. Getting Started with Setonix

2. Supercomputing Filesystems
3. Using Modules and Containers
4. Installing and Maintaining Your Software
5. Submitting and Monitoring Your Job
6. Using Data Throughout the Project Lifecycle



Section 1

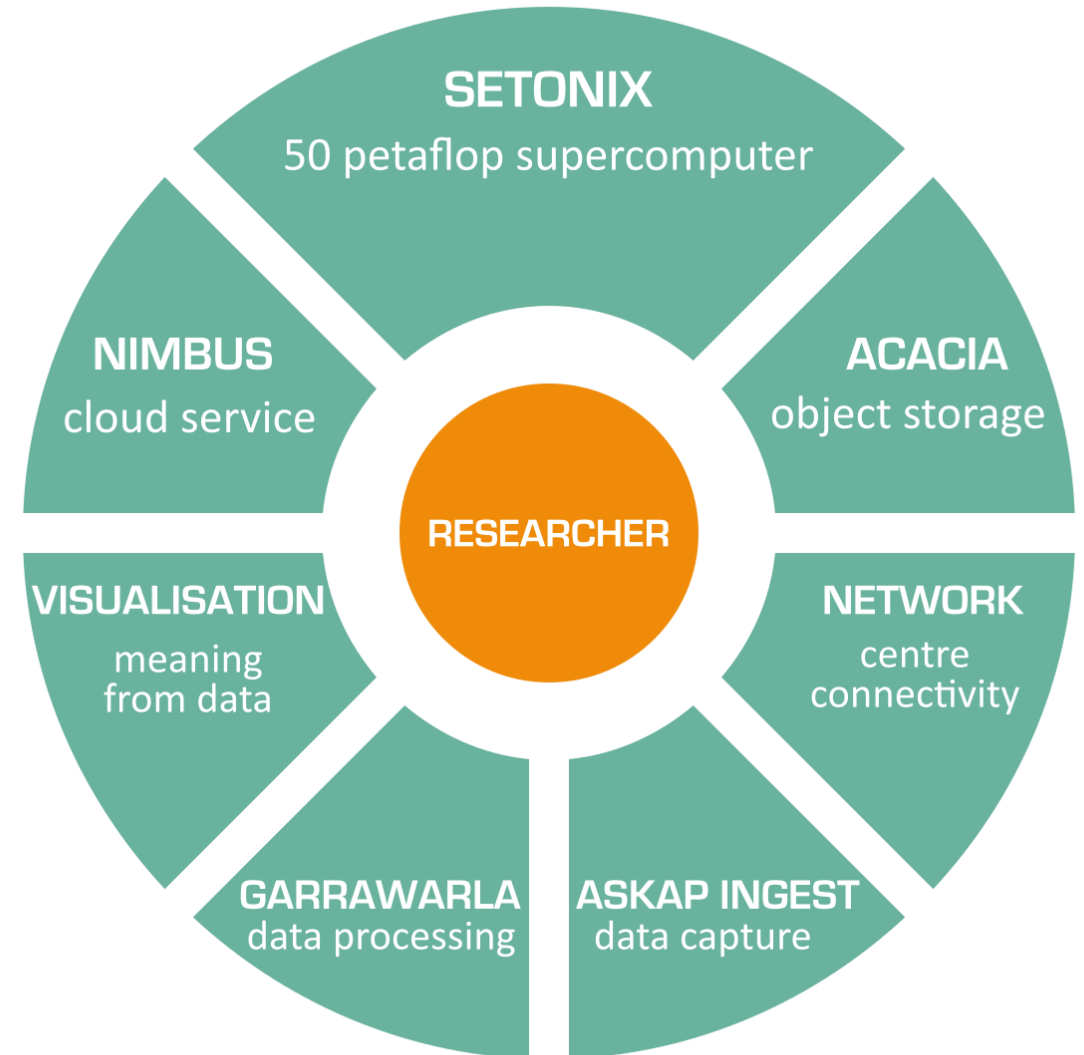
# What is Setonix?



**pawsey**

# Capital Refresh Overview

- Several initial procurements:
  - Garrawarla compute cluster
  - Astronomy high speed storage
  - Cloud high throughput computing
- Acacia object storage system
- Setonix:
  - Phase 1: CPU-based
  - Phase 2: Additional CPUs, GPUs
- More details @
  - [Changes in Supercomputing Services for 2022](#)



# Computational Capability

## What is the increase in computational capacity?

- Significant increase in computational capacity for the available schemes on Setonix.
- Increase in double-precision floating point operations available are:

Magnus  
(up to 2021)




1.1 petaflops

Setonix Phase 1  
(2022)



2.7 petaflops

Setonix Phase 2  
(2022+)



50 petaflops



# Setonix Hardware Overview

## Phase 1 provides:

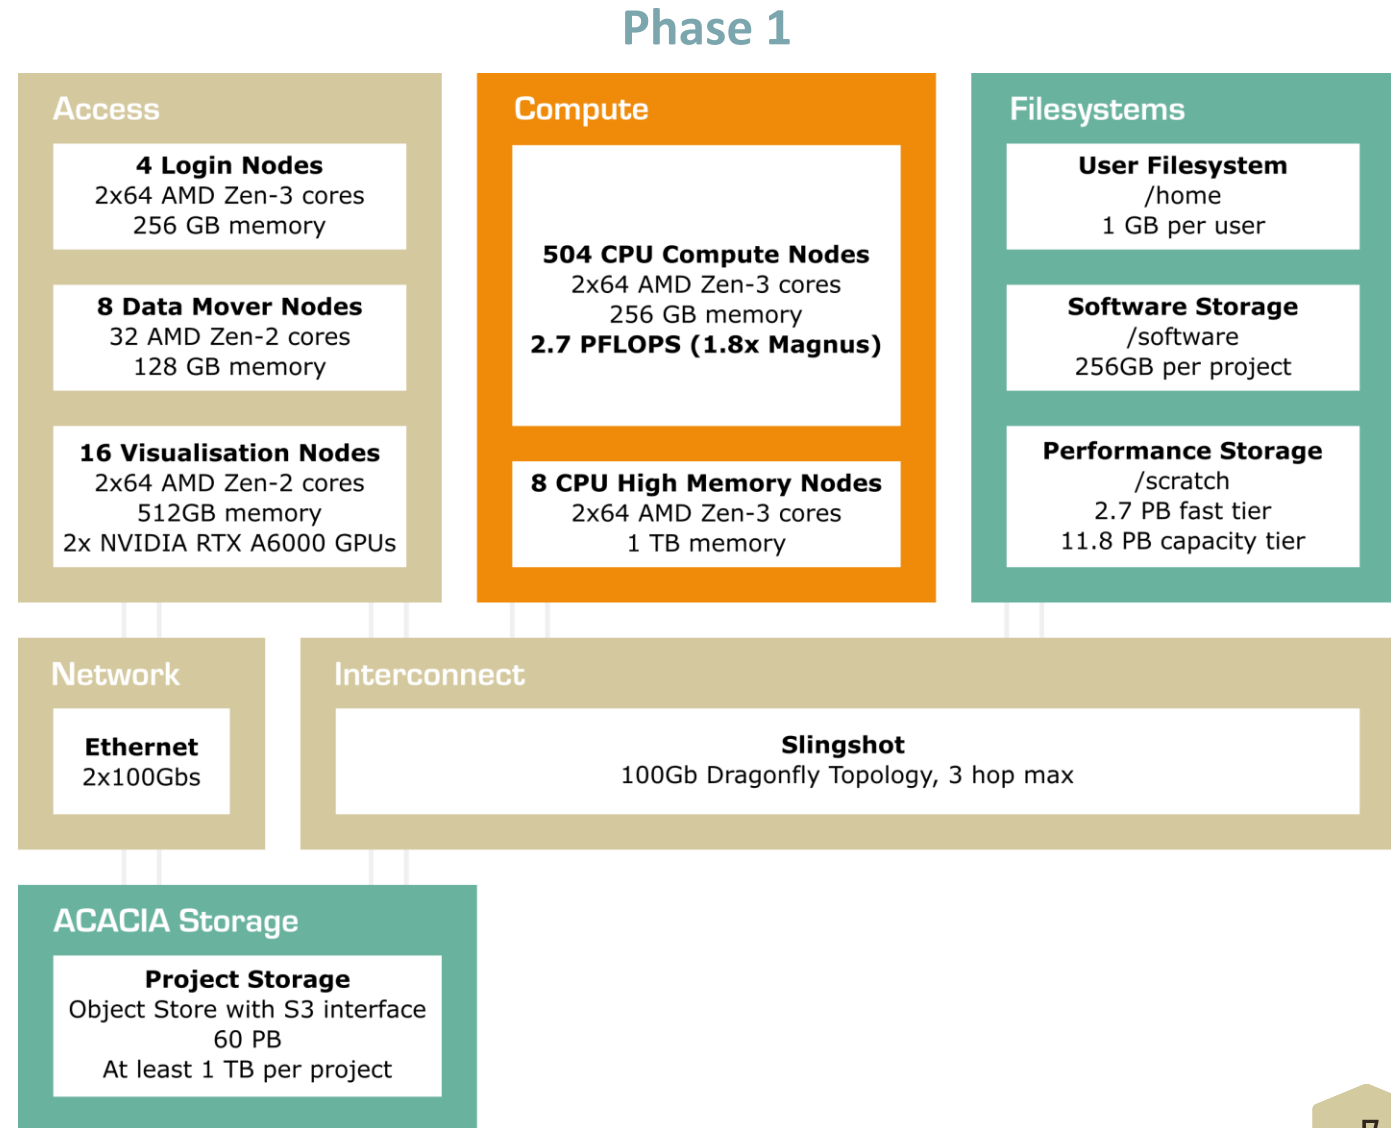
- CPU compute
- Fast interconnect
- High Memory and Visualisation Nodes
- High-performance filesystems (LustreFS)

## Phase 2 will add:

- Additional CPU compute
- Production-level GPU compute
- Slingshot upgrade to +200Gbs

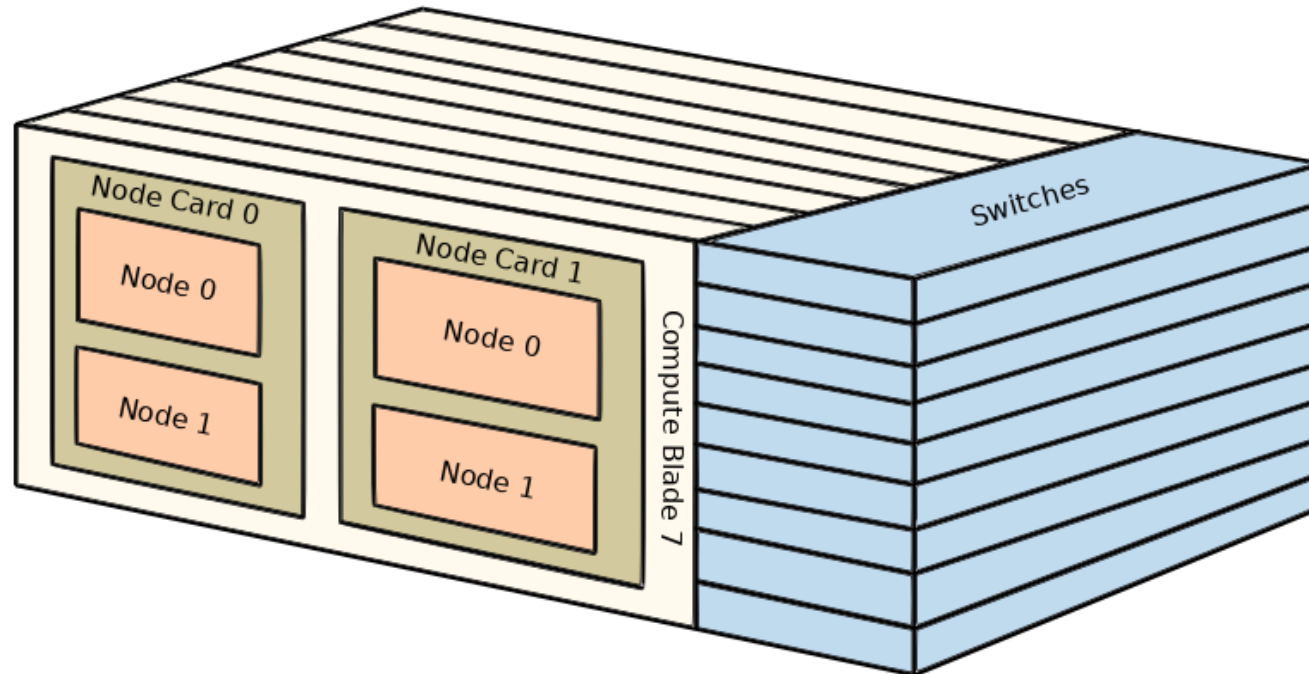
## Acacia system:

- Large-volume storage (Object Store, Ceph, S3)





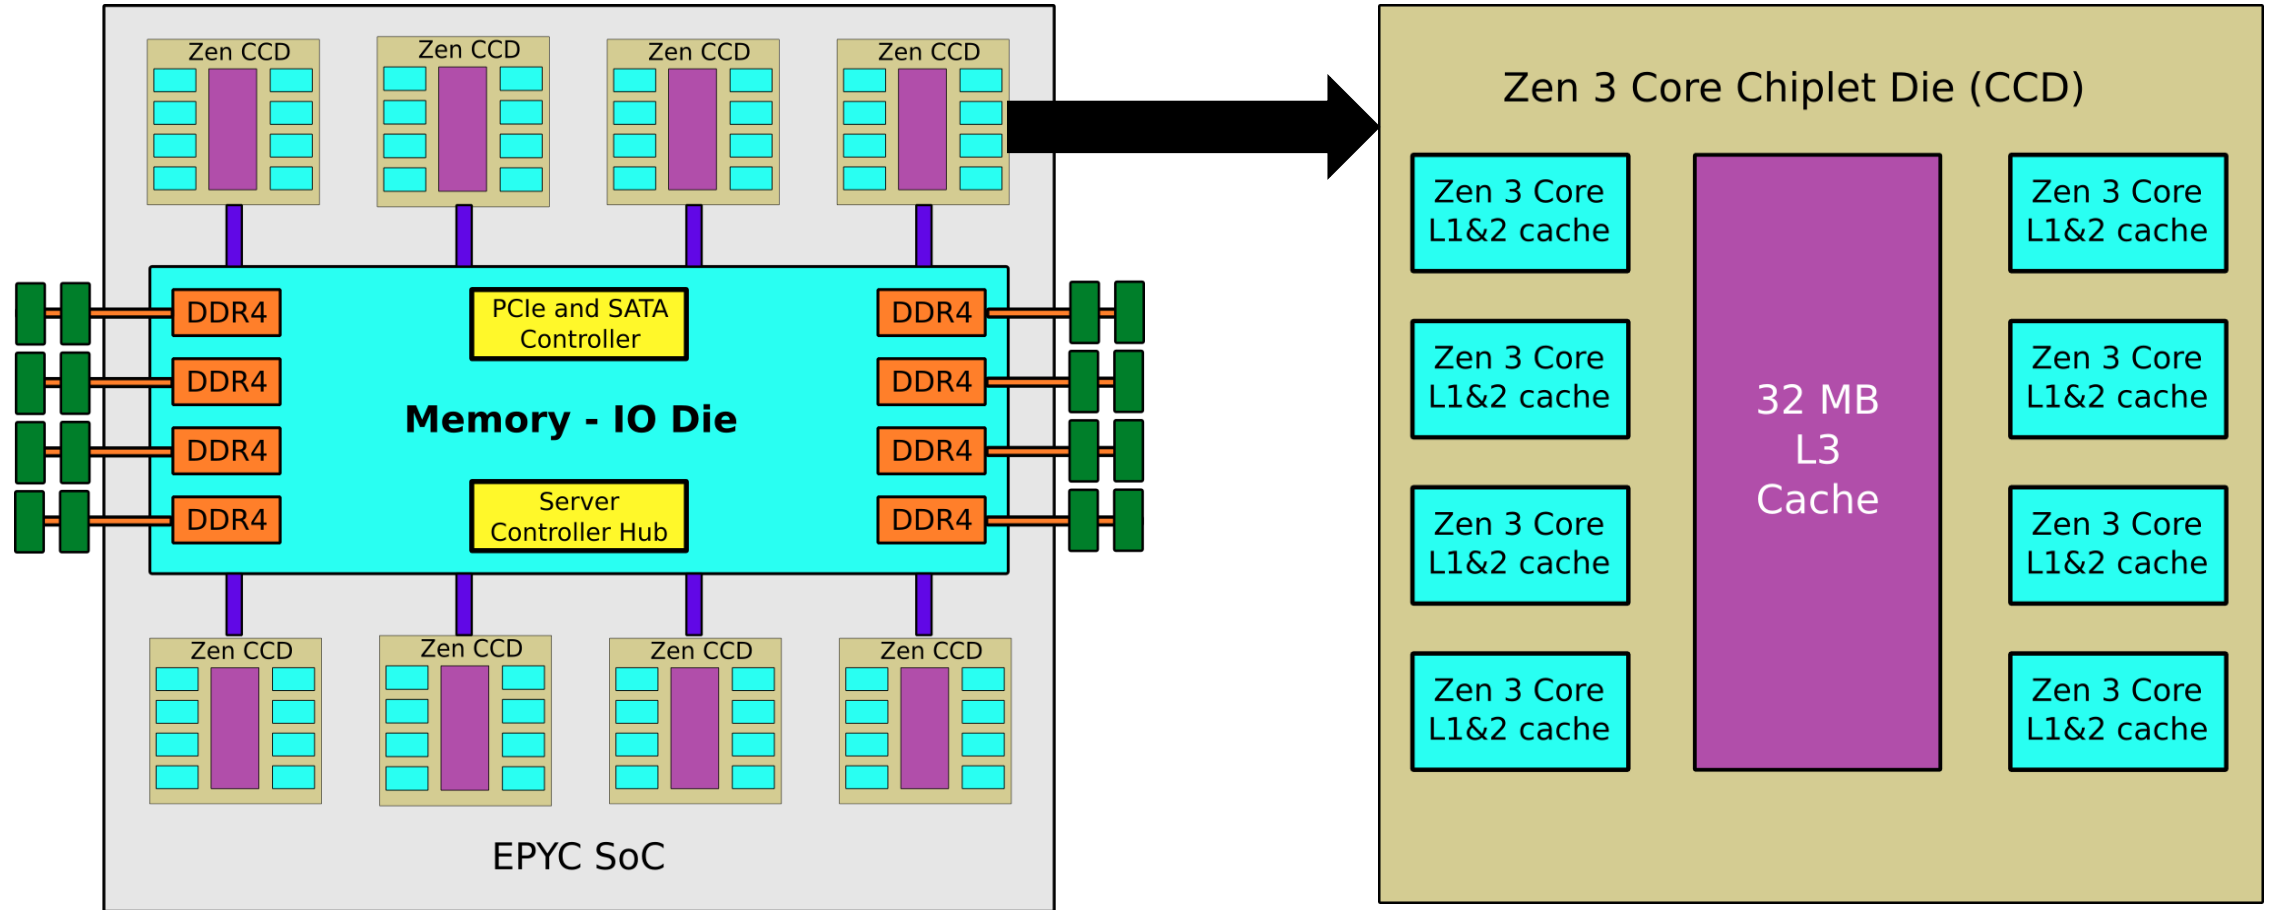
# Setonix Compute Node Architecture



- Each cabinet of Setonix contains 8 chassis.
- Each chassis contains 8 compute blades connected orthogonally to 8 interconnect switches.
- Each compute blade contains 4 nodes, each with 2 processors with a total of 128 cores.
- Note that adjacent nodes are not connected to the same switch.




# Setonix Compute Node Processor



- Further Zen-3 architecture information is available online:  
<https://www.amd.com/en/technologies/zen-core-3>

# Interconnect Layout

- The Setonix interconnect uses a dragonfly topology to connect nodes.
- Magnus has a similar topology, but Setonix lessens the number of hops, which reduces latency:
  - Setonix data transfers can take up to 3 hops.
  - While Magnus can take up to 5 hops between nodes.



# Key Hardware Changes

- Moving from 24 core Intel nodes to 128 core AMD nodes.
- Changing from 64 GB to 256 GB (more memory per node).
- Changing from 2.5 GB per core to 2 GB per core (slightly less memory per core).
- Exclusive node access to shared node access.
- Project storage on /group will move to the Acacia object store.
- Software installations on /group will move to the /software filesystem.
- New /scratch and /home filesystems.



# Storage Overview

## Supercomputing filesystems

### /home

- Like current Pawsey systems, minimal storage (NFS)

### /software

- LustreFS storage used for software, replaces some functionality of /group

### /scratch



- Fast LustreFS workflow storage. Data should be moved in or out of Object Store or offsite.

## Acacia Object Store

- Large-volume project storage, uses S3 interface

### ⚠ IMPORTANT:

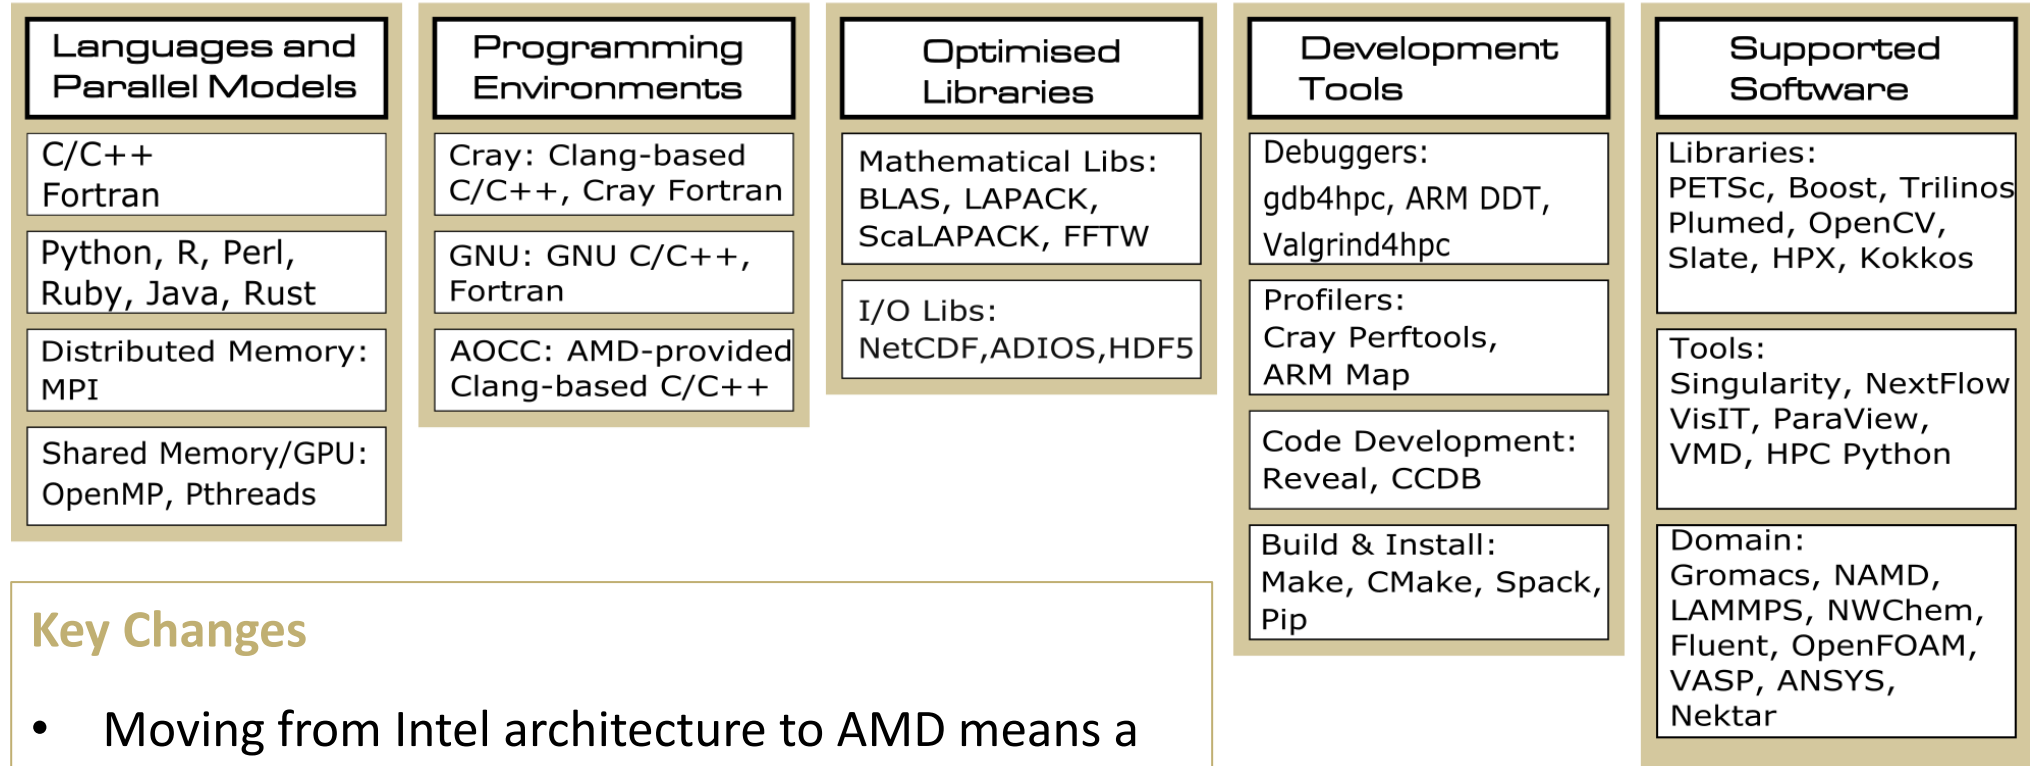
Acacia replaces the /group filesystem for project storage.



# Software Overview and Changes

## Overview

- HPE provides optimised compilers, libraries and tools.
- Pawsey supports software used in many scientific domains.





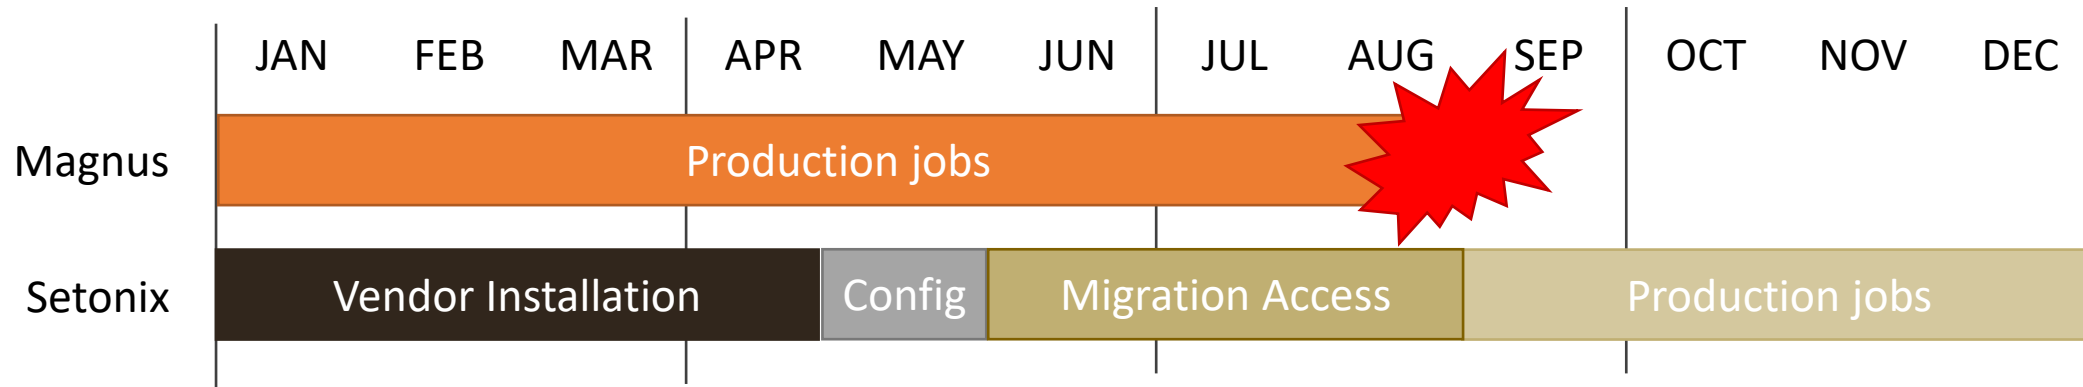
**pawsey**

Section 2

# Migration Overview



# Migration Schedule



- Projects will migrate from Magnus to Setonix, with access commencing across the first month of the migration.
- Following the migration, Magnus will be decommissioned to make way for expansions to Setonix.

 **IMPORTANT:** Magnus will be physically removed in August to make way for Setonix.



# Migration Support

- Support will be available throughout the migration to Setonix, including:
  - Regular migration module [training sessions](#), with [recordings](#) made available
  - Weekly [Q&A sessions](#)
  - [Supercomputing Documentation](#), including the [Setonix Migration Guide](#)
  - Advice from Pawsey staff through [User Support Portal](#) tickets



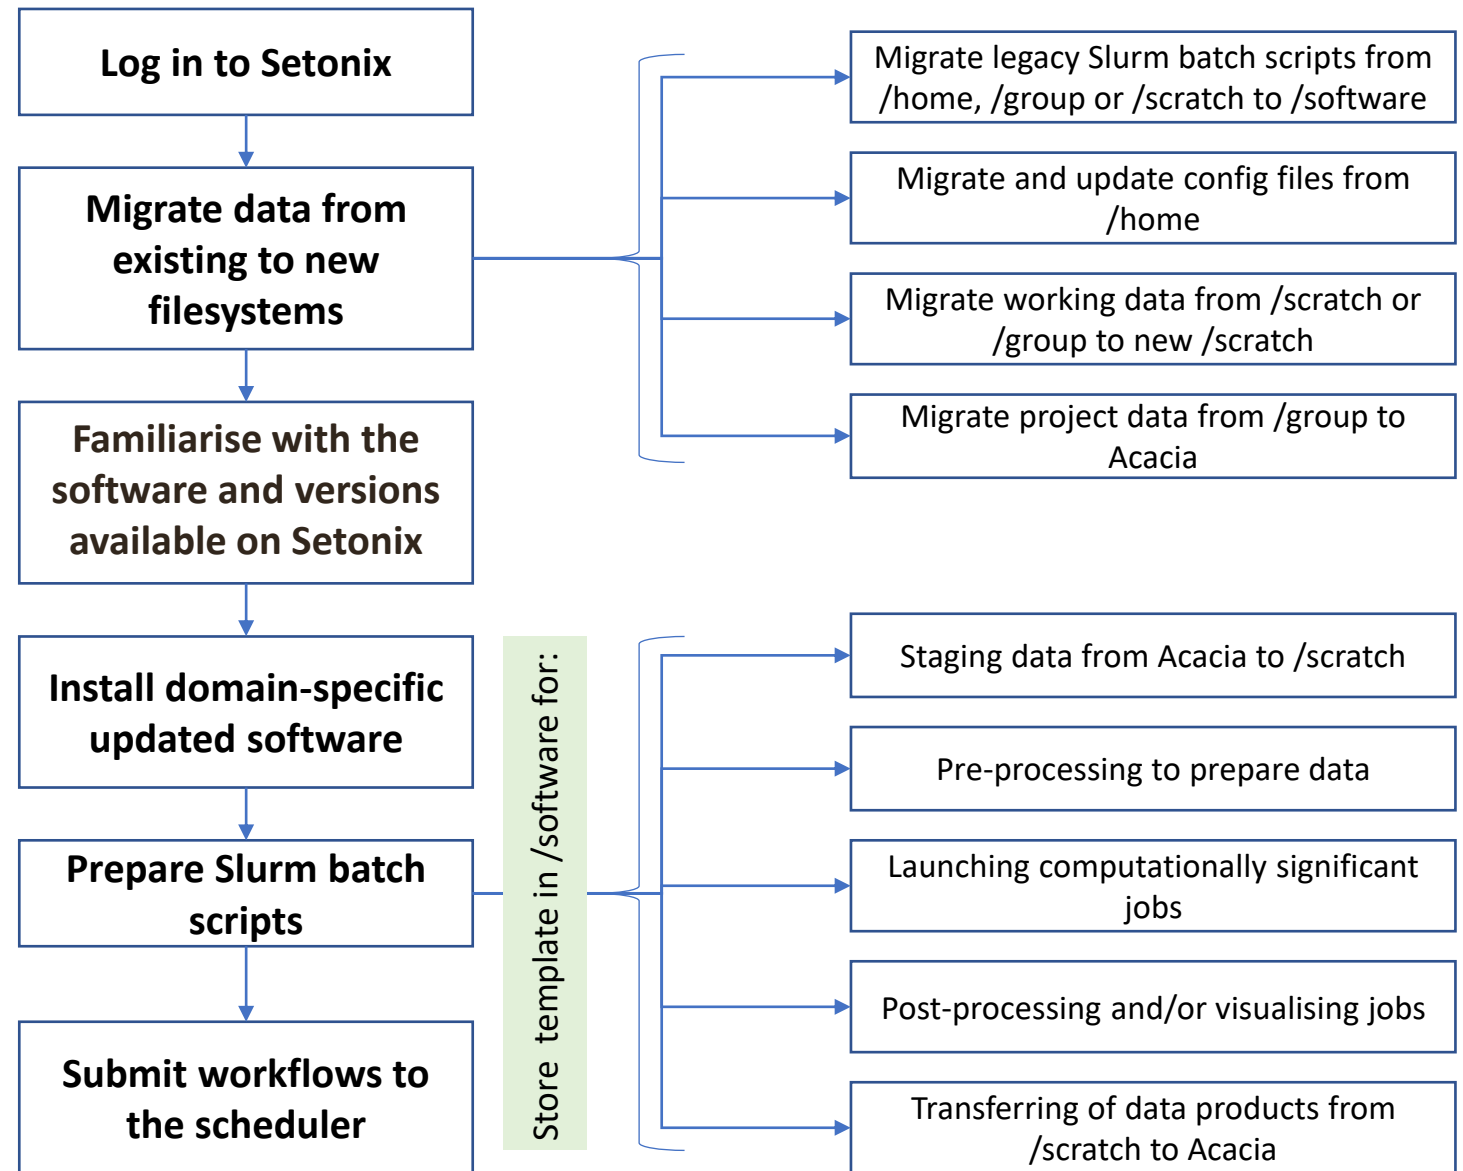
# Migration Pathway

## Overview

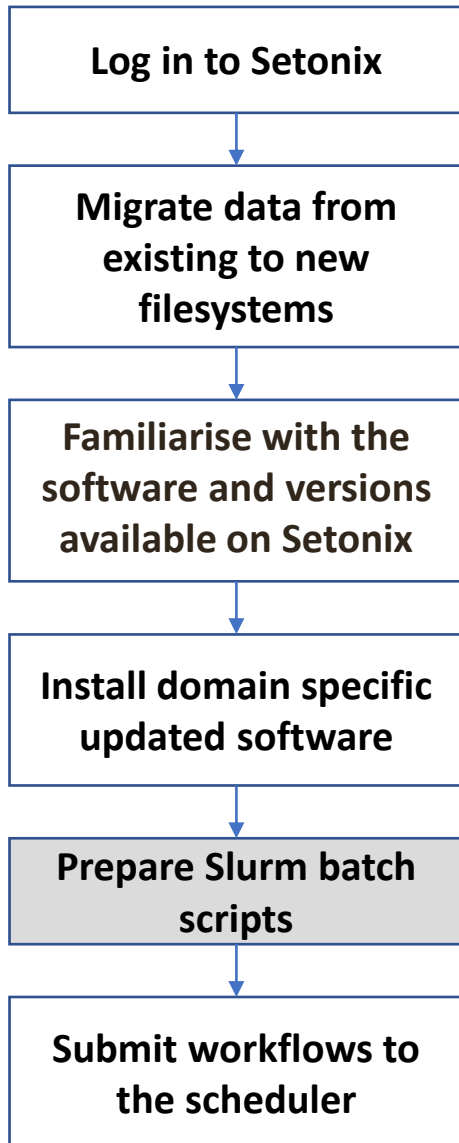
- This flowchart shows the key steps to migrating to Setonix from previous Pawsey systems.

## Best Practices

- Move important data, delete data that is no longer needed.
- Do not transfer software. Reinstall or use new versions.
- Set up environments in Slurm batch scripts, not login scripts.



# Use Case: Migrating Large-Scale MPI Jobs From Magnus



## Overview

- Same overall migration flow

## Key Changes

- There are now 128 cores and 256 GB of memory per node
- Benchmark to determine the best configuration for your code
- The Slurm directives shown are examples; use values that provide the best performance
- More details in Module 5: Submitting and Monitoring your Job ([Materials](#), [Recordings](#))

## Example 1024 core Slurm batch script

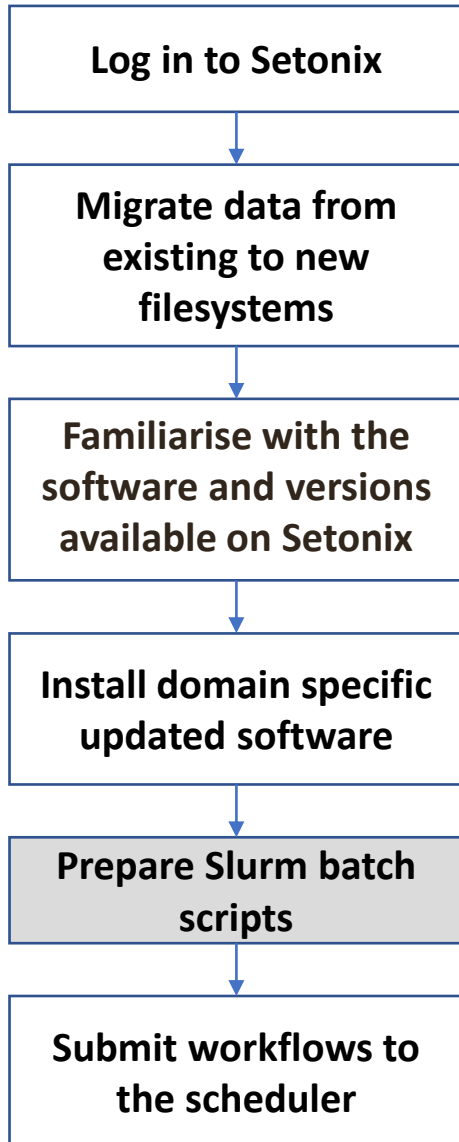
```
#!/bin/bash --login

#SBATCH --account=project
#SBATCH --partition=work
#SBATCH --ntasks=1024
#SBATCH --ntasks-per-node=128
#SBATCH --cpus-per-task=1
#SBATCH --time=24:00:00
#SBATCH --export=NONE
#SBATCH --exclusive

module load module/version

srun ./program
```

# Use Case: Migrating Throughput Workflow from Zeus



## Overview

- Same overall migration flow

## Key Changes

- There are now 128 cores and 256 GB of memory per node
- Memory usage should be specified if using part of a node
- Be aware of how shared node job accounting works
- More details in Migration Training Module 5: Submitting and Monitoring your Job ([Materials](#), [Recordings](#))

## Example single core job

```
#!/bin/bash --login

#SBATCH --account=project
#SBATCH --partition=work
#SBATCH --ntasks=1
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=1
#SBATCH --time=24:00:00
#SBATCH --export=NONE

module load module/version

srun ./program
```

# Use Case: Migrating GPU Workflow from Topaz

- Setonix Phase 1 does not contain GPU compute nodes; these will arrive with Phase 2.
- Continue using Topaz for GPU workflows.
- Setonix Phase 2 is expected later in 2022, and a separate migration will occur for Topaz.



To learn more about:

- AMD Instinct™ MI200 GPUs, refer to available [AMD Instinct videos](#).

You can also attend one of our upcoming *Introduction to AMD GPUs* training sessions. For information on upcoming training, see the [Pawsey Events calendar](#).





Section 4

# Logging in to Setonix



**pawsey**

# Setonix Login Nodes

- Log in using the same ssh client as before.
- There is a new hostname:  
**setonix.pawsey.org.au**
- Use the new hostname in place of old hostnames:  
magnus.pawsey.org.au  
galaxy.pawsey.org.au  
zeus.pawsey.org.au
- Login node provides access to:
  - Slurm scheduler
  - Cray Programming Environment
  - /home, /software, /scratch and Acacia Object Store
- Use the copy partition for significant data movement.

## Best practices for login nodes

Login nodes are shared. Use Slurm for:

- Compiling software (debug partition)
- Debugging software (debug partition)
- Benchmarking and profiling (debug partition)
- Transferring data (copy partition)
- Running programs (work partition)



# Log into Setonix Using a Terminal

```
$ ssh username@setonix.pawsey.org.au
#####
#                                                                    #
#          NOTICE TO USERS                                          #
#                                                                    #
# This computer system is operated by the Pawsey Supercomputing Centre #
# for authorised clients only. By using Pawsey facilities you agree to #
# the Conditions of Use available at                                 #
# https://support.pawsey.org.au/documentation/display/US/Conditions+of+Use #
#                                                                    #
# Your actions and activity on this system may be monitored and recorded. #
# Unauthorised or improper use of this system may result in administrative #
# disciplinary action and/or civil or criminal penalties.          #
#                                                                    #
# By continuing to use this system you indicate your awareness of and #
# consent to these terms and conditions of use. LOG OFF IMMEDIATELY if #
# you do not agree to the conditions stated in this warning.      #
#                                                                    #
#####
charris@setonix.pawsey.org.au's password:
```




## Login Options

We're logging into the system using SSH on the command line.

For details on how to log in using an alternative GUI program, see [How to log into Setonix](#) in the Setonix User Guide.

# How do I get help?



## Migration Documentation & Migration Guides

- [Setonix Migration Guide](#)
- [Setonix User Guide](#)
- [Supercomputing Documentation](#)

## Migration Training Materials & Video Recordings

- [Upcoming Migration Training](#)
- Recordings: [Pawsey YouTube Setonix Migration Phase 1 Playlist](#)
- Materials: [Setonix Migration Training Materials \(PDFs\)](#)

## Help Desk

- [Help Desk](#)
- Email: [help@pawsey.org.au](mailto:help@pawsey.org.au)



**Thank you for attending!**

Please complete this short survey:

<https://www.surveymonkey.com/r/Y3YFYHQ>



**pawsey**