The Bureau of Meteorology

HPC update for Australian Bureau of Meteorology

Adam Smith

Manager of Model Systems and Data
Research to Operations, Science and Innovation

October 2023

Australian Bureau of Meteorology's Purpose

The Bureau's purpose, as defined by its mission is:

To provide trusted, reliable and responsive weather, water, climate, ocean and space weather services for Australia – all day, every day.

These services benefit all Australians, contributing to a safe, prosperous, secure and healthy Australia.

The Bureau operates under the authority of the *Meteorology Act 1955* and the *Water Act 2007*. The Meteorology Act specifies that the Bureau must perform its functions in the public interest generally; for the purposes of the Defence Force, navigation and shipping and civil aviation; and to assist those engaged in primary production, industry, trade and commerce. The Bureau must also fulfil Australia's international obligations under the Convention of the World Meteorological Organization (WMO) and related international meteorological treaties and agreements.

Australian Bureau of Meteorology's Office Locations

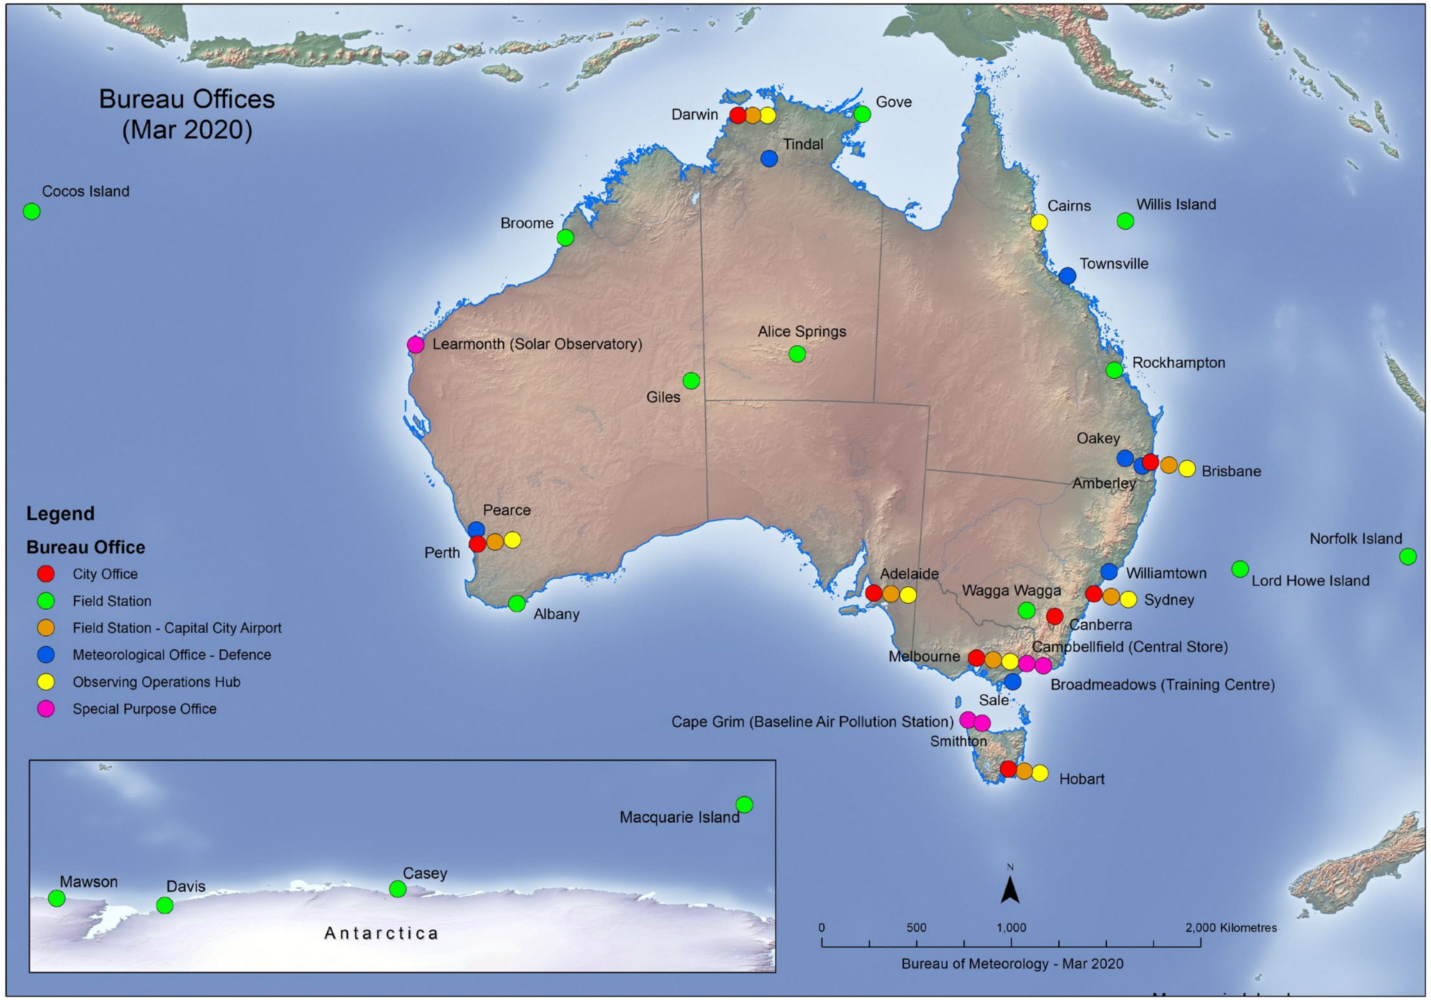

Australian Bureau of Meteorology's Region of Concern

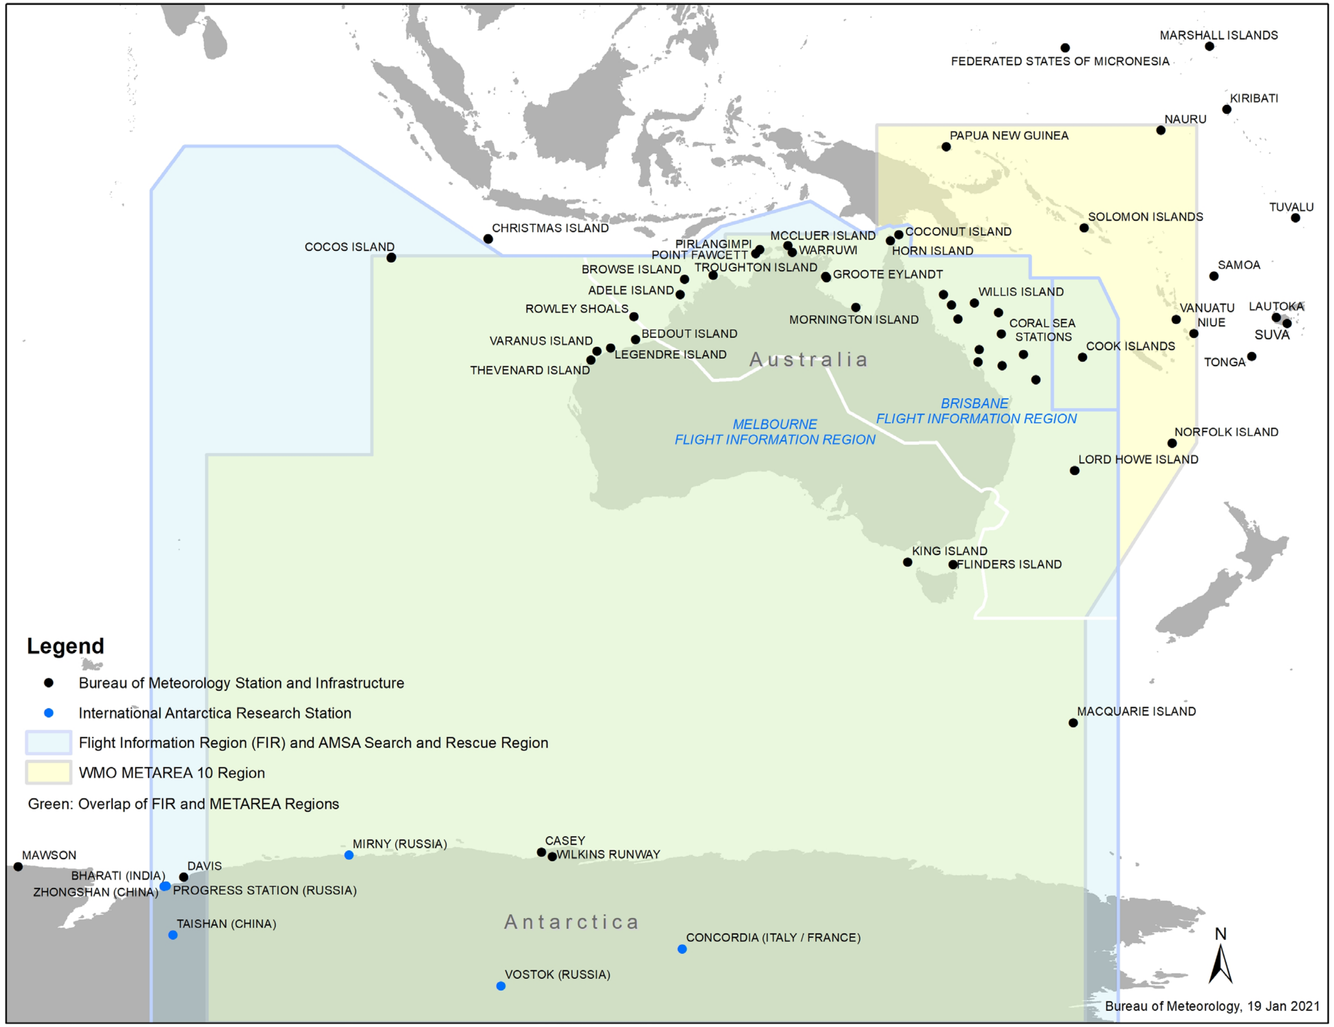

This area of responsibility covers the Australian continent, large areas of the Indian, Pacific and Southern Oceans, and the Australian Antarctic Territories, from the South Pole to the low equatorial latitudes.

At nearly 53 million square kilometres, this area covers one tenth of the earth's surface (~11%) and borders the regions of 10 other countries.

Supercomputing at the Bureau Today

Bureau's current production supercomputing platforms

Tim Pugh, Malcolm Cowe and HPC Services team

Overview of Australis

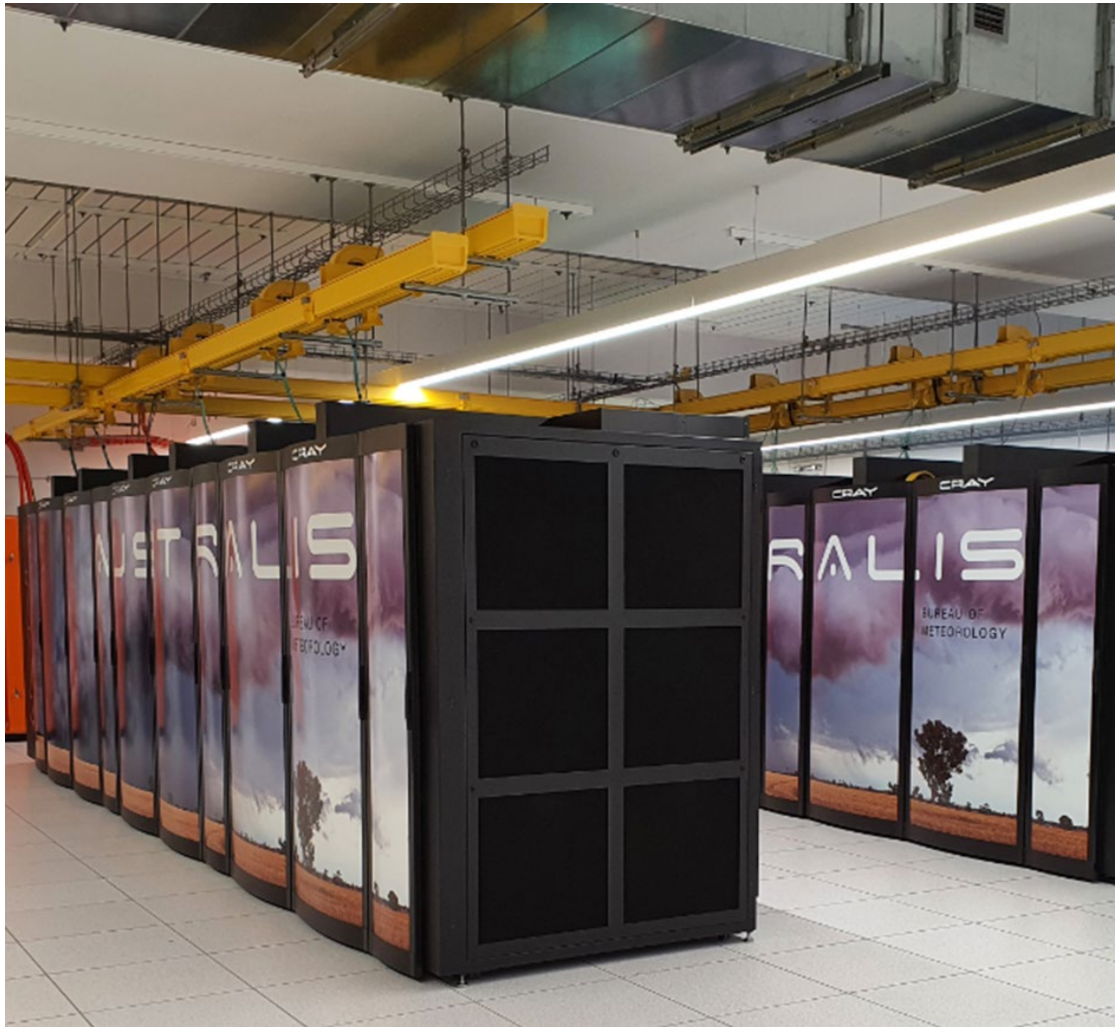

Australis I Supercomputer (2016)

- Production: July 2016
- Decommission: September 2024

Cray XC40 (simulation)

- 1.76 petaflops HPL
- 2152 nodes, 51,648 Haswell cores

Cray CS400 (data processing)

- 50 teraflops HPL
- 40 nodes, 1440 Broadwell cores

Cray Sonexion 2000 w/ Lustre

- 4.1 PB usable, 4 filesystems
- 135 GBps bandwidth

DDN GS14KXE w/ GPFS

- 5.8 PB usable – 3 filesystems
- 70 GBps bandwidth

NetApp FAS8040 home directory

FDR Infiniband network, Fat Tree

Australis II Supercomputer (2020)

- Development: January 2023
- Production: August 2024
- Decommission: December 2026

Cray XC50 (simulation)

- 3.51 petaflops HPL
- 1968 nodes, 94,464 Skylake cores

Cray CS500 (data processing)

- 97 teraflops HPL
- 48 nodes, 2304 Skylake cores

Cray ClusterStor L300N w/ Lustre

- 9.774 PB usable, 2 filesystems
- 162 GBps bandwidth
- 8.994 PB usable, 2 filesystems
- 108 GBps bandwidth

NetApp FAS8200 home directory

EDR Infiniband network, Fat Tree

Bureau's Australis Supercomputer Complex, 2023

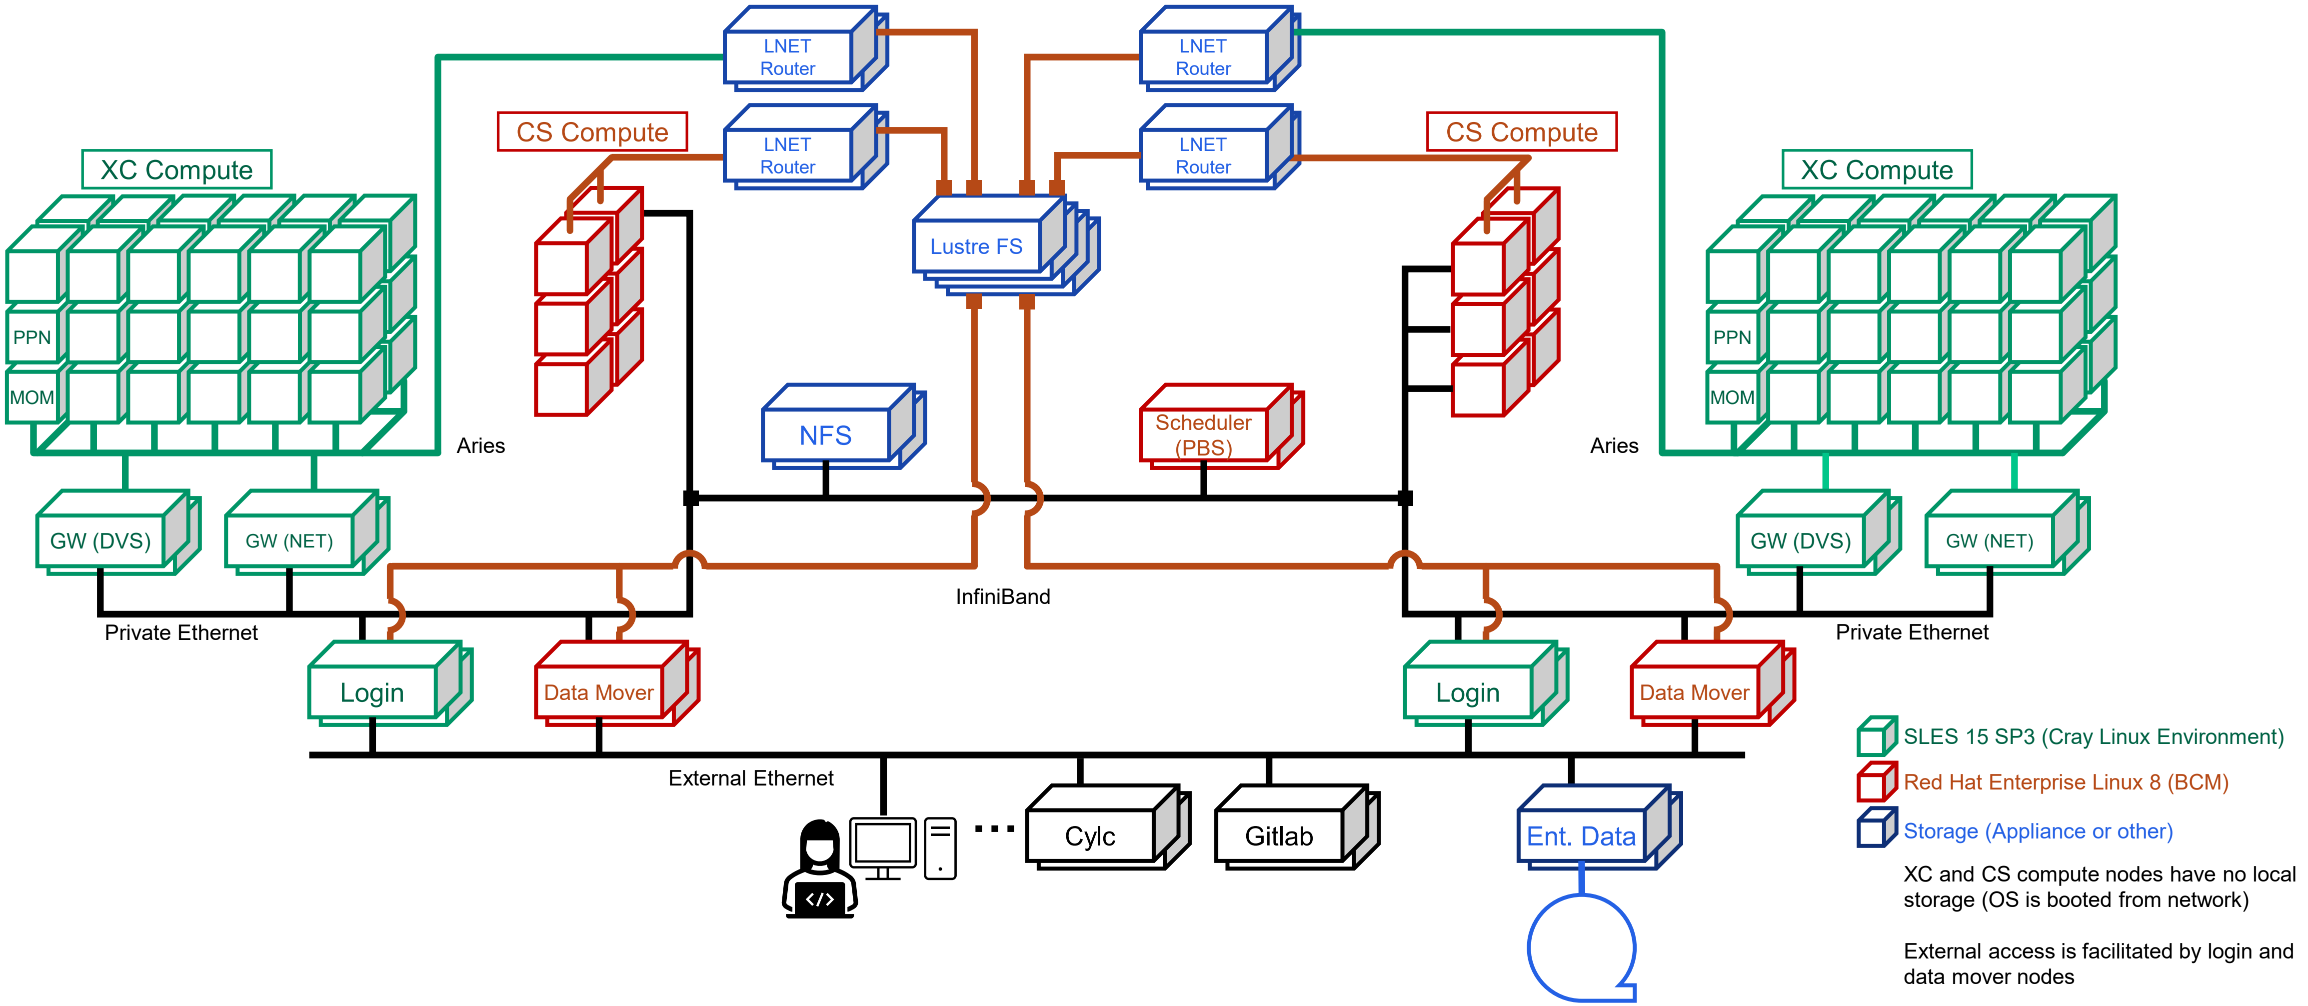

Networking information has been simplified for this diagram (some networks and routers omitted)
System management infrastructure has been excluded

Bureau HPC Resilience
Australis Supercomputer

The Australis supercomputer design is based on an enterprise high-availability paradigm

- Resilience is achieved through hardware redundancy
- Two individual supercomputers comprise the complex
- At any given time, one computer is the nominated production platform
- Second computer is used for pre-production development and testing
- If the production computer fails, the second computer takes over as the production platform
- The computers share a common pool of high-performance storage
- Data is written to primary storage and replicated to secondary storage
- The secondary storage provides redundancy for the primary storage in the event of a failure

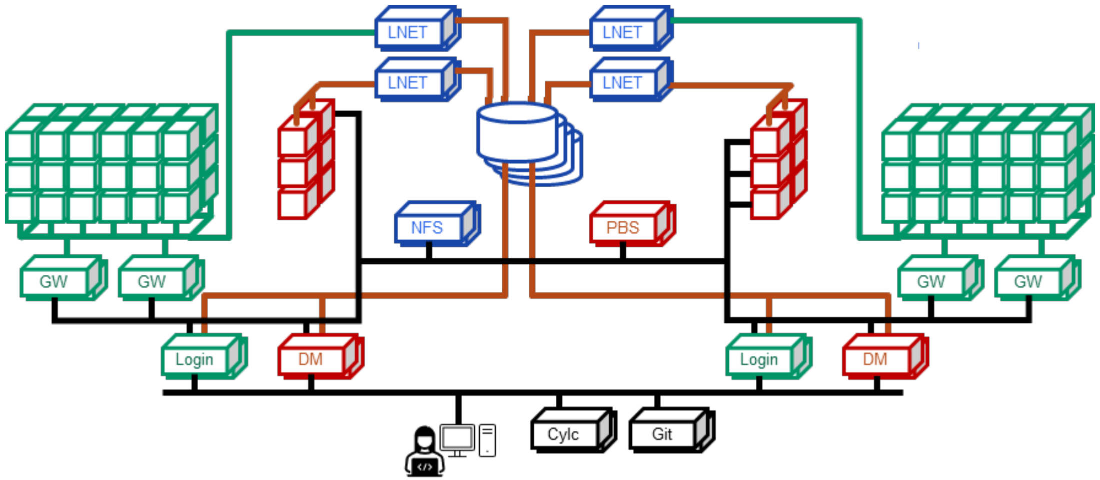

High Availability (active-passive failover of production workloads)

Limitations

- Machines in the complex are co-located (no DR)
- Trade-off between operational continuity and utilisation efficiency
  - Cannot assign >50% of the complex to production workflows (the size of the failure zone) without compromising availability and serviceability (i.e., maintenance)
- Secondary (non-production) zone has constraints that limit scope of research and other projects

Bureau HPC Resilience
Disaster Recovery for Australis

Secondary mirror of the Australis complex

- Functionally equivalent copy of the main supercomputer infrastructure installed in a geographically separate datacentre
- Each datacentre contains an independent supercomputer infrastructure
- Production runs at the primary site
- The supercomputer complex at the secondary site hosts workloads that are not critical to Bureau product delivery (e.g., R&D)
- The complexes at each site are resilient to failure using HA failover availability model
- Production data generated at the primary site is replicated to the secondary data centre
- In the event of a disaster, all workloads on the secondary site are halted and production workloads are transitioned to run at the secondary data centre
- Checkpoint recovery model for movement of production to secondary site
- Production data is sync'd back to primary site when the site is restored to service

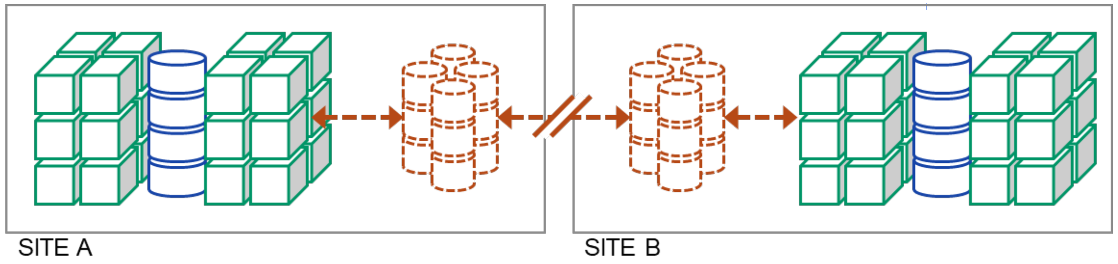

Basic dual-site Disaster Recovery
Design assumes presence of intermediate data store and transfer function
Note that HA model persists in each site to preserve existing resilience

Limitations

- Transition of production workloads to secondary site is by manual intervention, significant effort to recover operations
- Each site continues to trade-off operational continuity and utilisation efficiency (proportionally smaller capacity is available for production)
- Data must be synchronised between sites
  - Checkpointing requires explicit recovery to restore production
  - Risk of data corruption
- In the event of DR, production workloads displace the jobs that were running on the secondary site, creates uncertainty for users of the secondary site computer
- Loss of productivity for affected organisations for duration of secondary production running

Table of Bureau of Meteorology systems

Existing HPC Systems & new HPC Systems	Processor Technology	Aggregate Compute Nodes	Aggregate Compute Cores	Theoretical FLOPS	Top500 Rmax (HPL)	SSP benchmarks	2017 SPECfp base (kSFRP)
2016 to 2024 **Australis I (Cray XC40, CS400)** **2 HPC systems – East, West**	Intel Haswell 2690v3 12-core, 2.6 GHz with 4 ch. DDR4-2166	2,192 (2 x 1076, 2 x 20)	53,088 (51,648, 1440)	2.156 PF	1.765 PF (2 x 882 TF)	243.94	N/A
Terra (Cray XC40)		144	3,456	110 TF	90 TF		
2022 to 2026 **Australis II (Cray XC50, CS500)** **2 HPC systems – Casey, Davis**	Intel Skylake 8160v5 24-core, 2.1 GHz with 6 ch. DDR4-2666	2,016 (2 x 984, 2 x 24)	96,768 (94,464, 2304)	4.335 PF	3.514 PF (2 x 1757 TF)	409.50	423.36
2024 to 2029 **DR HPC (HPE Cray EX 4000)** **2 HPC systems**	Intel Sapphire Rapids 8470Q, 52-core, 2.1 GHz with 8 ch. DDR5-4800	896 (2 x 448)	93,184	6.262 PF	To be Confirmed	To be Confirmed	To be Confirmed
Test-Diagnostic System (TDS)		64	6,656	447 TF	To be Confirmed	To be Confirmed	To be Confirmed

Table of ANU NCI's systems (10% share for Bureau research)

Existing HPC System	Processor Technology	Aggregate Compute Nodes	Aggregate Compute Cores	Theoretical FLOPS	Top500 Rmax (HPL)	SSP benchmarks	2017 SPECfp base (kSFRP)
Gadi (Fujitsu/Lenovo) **Lenovo ThinkSystem SD650, Fujitsu PRIMERGY CX2570 M5**	Intel Cascade Lake 8274v6 24-core, 3.1GHz with 6 ch DDR4-2933	3,074	147,552 cores	15.142 PF	9.264 PF	N/A	N/A
	Intel Sapphire Rapids 8470Q, 52-core, 2.1 GHz with 8 ch. DDR5-4800	720	74,880	5.03 PF	4.398 PF	N/A	N/A
	Intel Cascade Lake + 4 x Nvidia V100 GPUs	160	7,680 cores 640 GPUs				

Compute Capacity Evolution Across System Generations

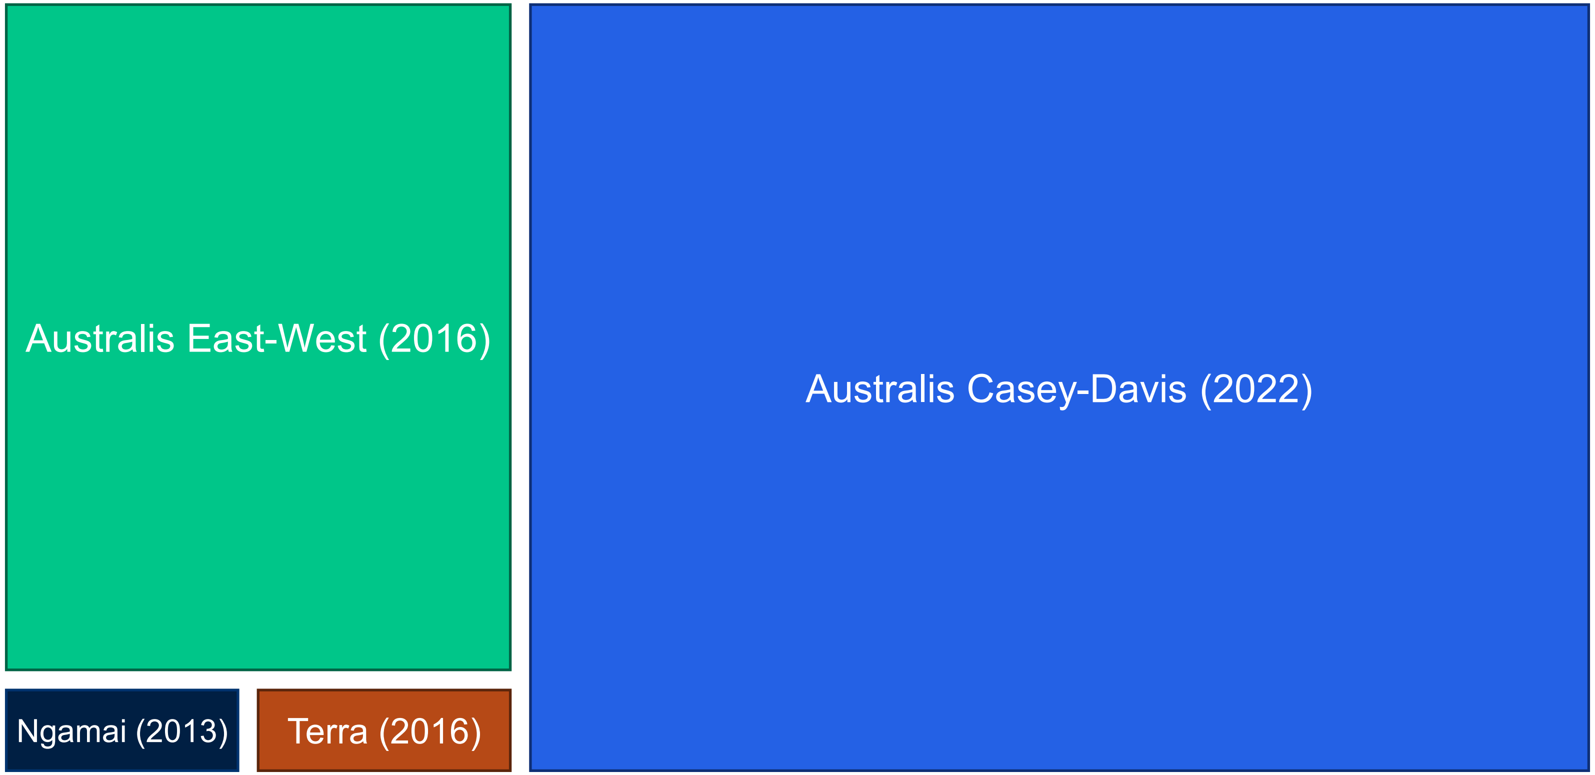

Compute capacity has increased significantly generation on generation

Australis MTU has over 2.5 times the capability of the existing East-West complex

Terra R&D platform has compute performance equivalent to the previous generation production supercomputer, Ngamai

2009 Solar: ~50TF (https://top500.org/system/176900/)

2013 Ngamai: ~104TF

2016 Australis: ~2PF (estimate, East + West combined)

2016 Terra: ~110TF

2022 Australis MTU: ~5PF (estimate, Casey + Davis combined)

High Performance Computing Systems at the Australian Bureau of Meteorology

HPC System	HPC model and Storage System	Processor Type	# Nodes: # Cores	Top500 Rmax (HPL)	Operation Dates
HPC for DR <to be named>	HPE Cray EX4000	Intel Xeon Sapphire Rapids 8470Q, 52-core, 2.1 GHz with 8 ch. DDR5-4800	896n: 93,184c	To be Confirmed	2024 - 2029
Australis II	HPE Cray XC50 HPE Cray CS500	Intel Xeon Skylake 8160v5, 24-core, 2.1 GHz with 6 ch. DDR4-2666	2,016n: 96,768c	3,514 TF	2022 - 2026
Australis I Aurora	Cray XC40 Cray CS400	Intel Xeon Haswell 2690v3,12-core, 2.6 GHz with 4 ch. DDR4-2166 Intel Xeon Broadwell E5-2695v4 18-core, 2.1GHz with 4 ch. DDR4-2400	2,200n: 53,280c	1,765 TF	2016 - 2024
Terra	Cray XC40	Intel Xeon Haswell 2690v3,12-core, 2.6 GHz with 4 ch. DDR4-2166	144n: 3,456c	110 TF	2016 - 2023
Ngamai	Oracle Blade 6000	Intel Xeon Sandy Bridge 2640v1, 6-core, 2.5 GHz with 4 ch. DDR3-1800	576n: 6,912c	104 TF	2013 - 2016
Solar	Sun Constellation	Intel Xeon Nehalem, 4-core	576n: 4,608c	50 TF	2010 - 2013
NEC SX-6	NEC SX-6/M224	NEC SX-6 node with 8 Vector Processors	28n: 224c	1.79 TF	2004 - 2010

* Estimated HPL value

Bureau's three generations of HPC systems

Milestone Schedule for the Bureau's HPC systems

- Australis II (Cray XC50) to run the Bureau's production suites by August 2024
  - Australis I (Cray XC40) to be decommissioned by September 2024
- New Secondary HPC (HPE Cray EX4000) going through implementation and testing
  - System delivery in April 2023, to be commissioned by mid-2024
  - Establish disaster recovery capability in late 2024

Existing HPC Systems & new HPC Systems	Processor Technology	Aggregate Compute Nodes	Aggregate Compute Cores	Theoretical FLOPS	Top500 Rmax (HPL)	SSP benchmarks	2017 SPECfp base (kSFRP)
2016 to 2023 Australis I (Cray XC40, CS400) 2 HPC systems – East, West	Intel Haswell 2690v3 12-core, 2.6 GHz with 4 ch. DDR4-2166	2,192 (2 x 1076, 2 x 20)	53,088 (51,648, 1440)	2.156 PF	1.765 PF (2 x 882 TF)	243.94	N/A
Terra (Cray XC40)		144	3,456	110 TF	90 TF		
2020 to 2026 Australis II (Cray XC50, CS500) 2 HPC systems – Casey, Davis	Intel Skylake 8160v5 24-core, 2.1 GHz with 6 ch. DDR4-2666	2,016 (2 x 984, 2 x 24)	96,768 (94,464, 2304)	4.335 PF	3.514 PF (2 x 1757 TF)	409.50	423.36
2023 to 2029 2nd HPC (HPE Cray EX 4000) 2 HPC systems – TBA	Intel Sapphire Rapids 8470Q, 52-core, 2.1 GHz with 8 ch. DDR5-4800	896 (2 x 448)	93,184	6.262 PF	To be confirmed	To be confirmed	To be confirmed
Test-Development System (TDS)		64	6,656	447 TF	To be confirmed	To be confirmed	To be confirmed

2023 Secondary HPC Supercomputer

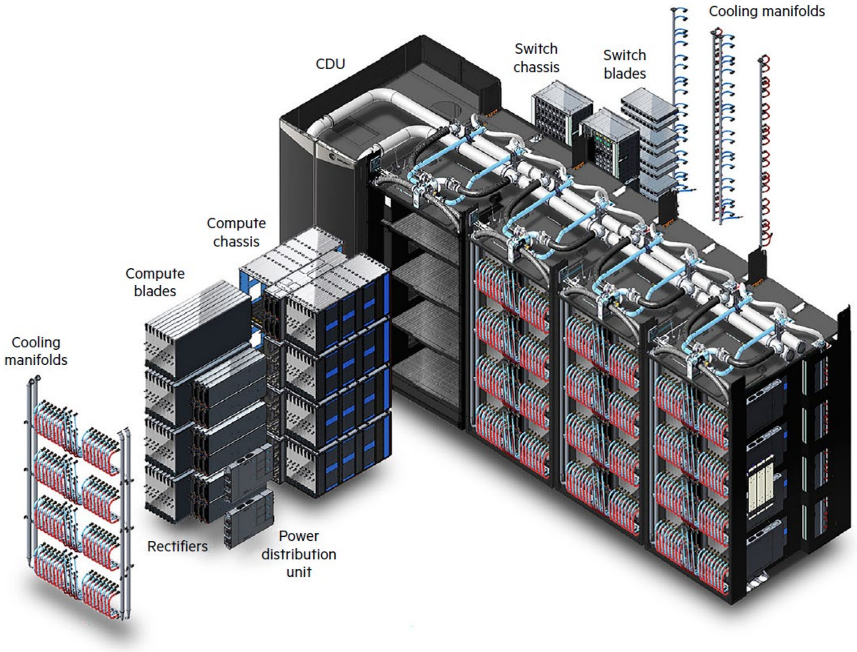

Enterprise Services	The Bureau Enterprise Common Servers
Developer Workflows	The Bureau Build Pipeline: GitLAB, Artifactory and Runners
Third Party Toolchains	Intel oneAPI Base and HPC Toolkits; Linaro Forge Profiler and Debugger
Toolchains	HPE Cray Programming Environment; Cray MPI
Data and Control Management	Altair PBS Professional; CYLC (Altair Weather Solution)
Operating Systems	HPE Cray OS; SLES; RHEL
System Management	HPE Performance ClusterMnager; HPE Slingshot Fabric Manager
Firmware	Redfish
System Networking	HPE Slingshot; HPE Aruba Ethernet
Compute and Storage Systems	HPE Cray EX; HPE Parallel File System Storage; HPE ClusterStor E1000

HPE EX Supercomputer (2023)

- Delivered: April 2023
- Commission: August 2024
- Decommission: 2029

Cray EX4000 (simulation)

- 5.0 petaflops HPL
- 832 nodes, 86,528 SPR cores

Cray EX4000 (data processing)

- 392 teraflops HPL
- 64 nodes, 6656 SPR cores

Cray Tiered ClusterStor 1000 w/ Lustre 2.12 and GridRAID

- 2 x 811 TB SSD usable
- 528 GBps bandwidth
- 2 x 5 PB HDD usable
- 232 GBps bandwidth

HPE PFSS shared/home directory

HDR Ethernet Slingshot network

Scientific Stack, Architectural Patterns & Tooling Fleet

Luke Garde and DevOps team

DevSecOps principles

Vulnerability Detection:
Detect and fix security issues in code and configurations early in the development process.

Faster Time to Market:
Streamline deployment processes for quicker software delivery.

Improved Compliance:
Ensure adherence to regulatory requirements and reduce compliance risks.

Enhanced Security Posture:
Proactively monitor and address security threats for a stronger security stance.

Reduced Security Risks and Costs:
Minimize the risk of breaches and associated costs through automation and proactive security practices.

Automation:
CI/CD processes speed up development by reducing manual tasks and enabling rapid code integration and testing.

Why this Scientific Stack?

1. Our software stack *needs to meet* security challenges and governance. This presents the opportunity for a standardised approach to tooling, libraries and applications to support our projects.

2. Our toolchain must be trusted and managed. By this we mean;
   - must be CVE compliant,
   - be sourced from known and approved locations
   - be security scanned before entering the network envelop
   - kept up to date with the latest patch or version levels.

3. Our toolchain must be Technical Reference Manual compliant. By this we mean, stated as part of the organisation’s Architectural roadmap, which means that, the organisation also has the onwards Operational skills and financial planning to support the product over its operational lifecycle.

4. The Industry standard for an organisation to manage such requirements is through a standardised software stack. However, developing a standard toolset is not a simple action for a scientific/IT agency. We are now of our third iteration of the concept. We call this version "*Kit*".

Scientific Stack Considerations

Move from custom-made tooling solutions to freeware or commercial off the shelf solutions (COTS)

Address Organizational governance, Whole of Australian Government (WoAG) guidelines, Cybersecurity mandates and DevSecOps Industry best practice

Identifying and addressing CVE patching and software compatibility across complex Scientific workflows

Acknowledging that Scientific Coder mentality is different to Software Engineer methodology

Approach to modernizing our Scientific Stack

BoM_Modules 1

One stack for all

- Completely manual.
- Each update = 8+ weeks work.
- Managing cross-software compatibility was like a “wack-a-mole” game.
- Stack deployment schedule could never keep up with release requests.
- Difficult to respond quickly to CVE issues.
- Impossible to deprecate.

Apps 2.0 2

One stack, some stacks

- Trialled with projects providing additional/updated packages individually.
- Added allow-listing to get new/updated packages from external sources.
- No tracking of additional packages in use made CVE response harder.

Kit 3

Everyone has stacks

- Packages delivered as a service.
- Mix of internally-built and externally sourced packages.
- On-demand for users/projects.
- Package manager (Conda) ensures package version compatibility per project.
- Packages in use tracked by service.
- CVE response and version updates are handled by the service.

Moving from BoM_modules to Kit

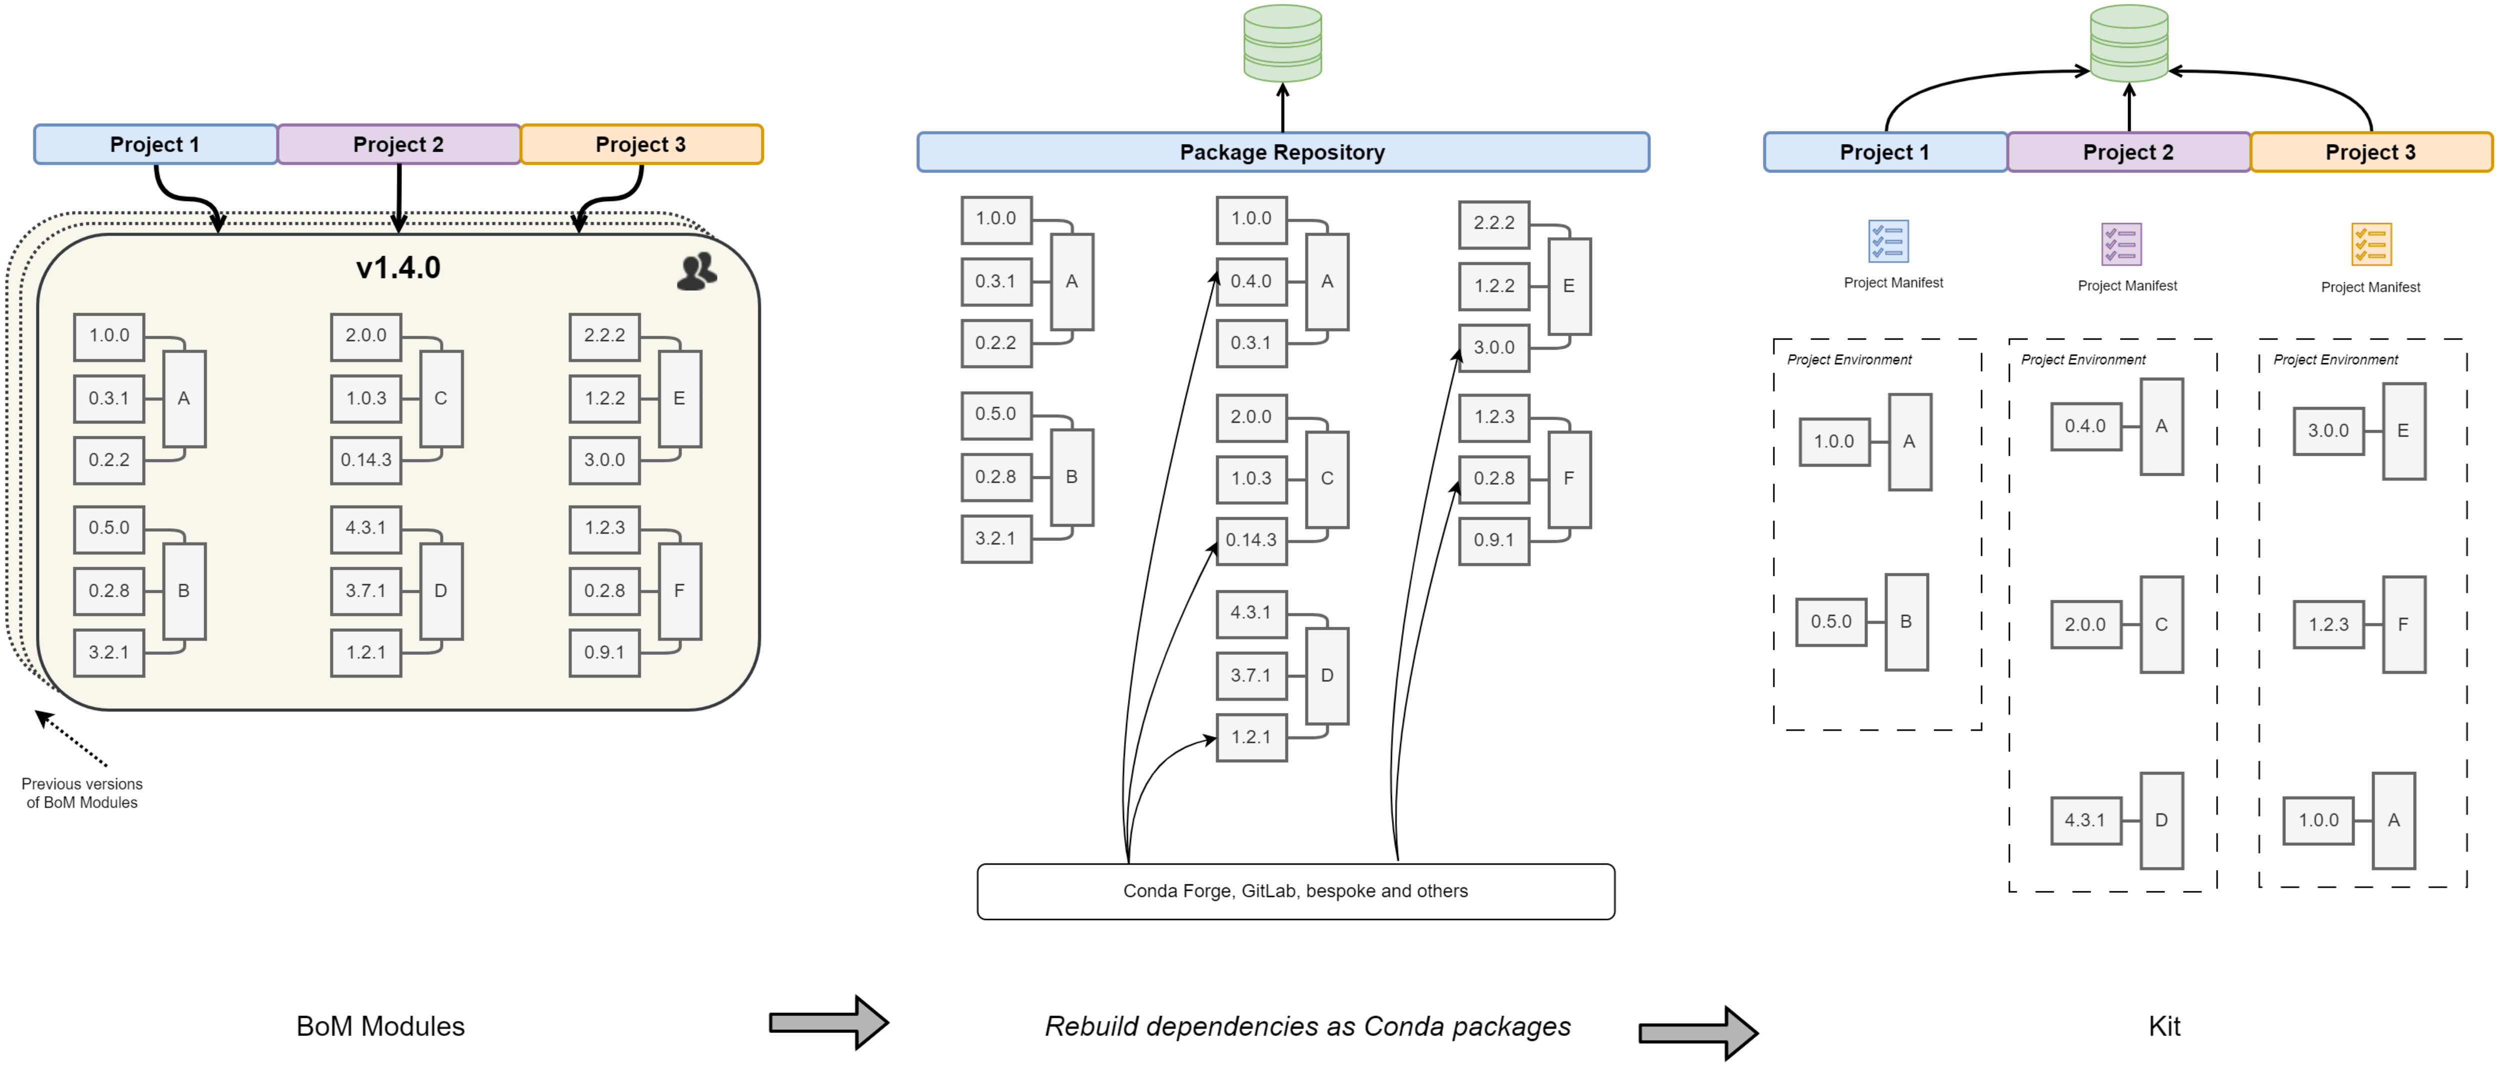

Scientific Stack Benefits

The benefit of the Kit approach

1. Centralized store of Bureau Security and Architecture approved tools, libraries and applications for use across multiple projects.
2. Improved access to upstream code sources, updates and security fixes.
3. Software version dependency independence between projects.

Efficiencies of the Kit Approach

1. Add new libraries and tools when each project needs them.
2. The most up to date, secure and compatible builds every time.
3. Enables multiple environments deploying variable builds.
4. Release windows measured in minutes, not weeks.
5. Built with Conda means software can be recreated without fuss using manifests.
6. Designed with automation as a foundation.

Why a patterned approach?

What is an Architectural Pattern?

- A "pattern" can be defined as: *"an idea that has been useful in one practical context and will probably be useful in others"*
- Patterns are a way of applying a previously agreed solution to a current problem.
- Architectural Patterns deliver repeatable, secure, robust and scalable solutions into an Enterprise environment.

Why do we have Architectural Patterns?

- Every time a new technical solution is developed it is reviewed by architecture teams ensure it confirms to the Bureau's agreed Technical Roadmap, Enterprise security standards, and Architectural and Competency standards.
- This ensures that all the Bureau's technical building blocks can fit together, be budgeted for and supported across their entire operational lifecycle.

Considerations for a patterned approach?

Must be documented and simple to follow and easy to use.

Must support all governance considerations.

Must deliver a secure, reliable and repeatable solution.

Must be able to respond to new technology usages.

Example of a Patterned approached - Design

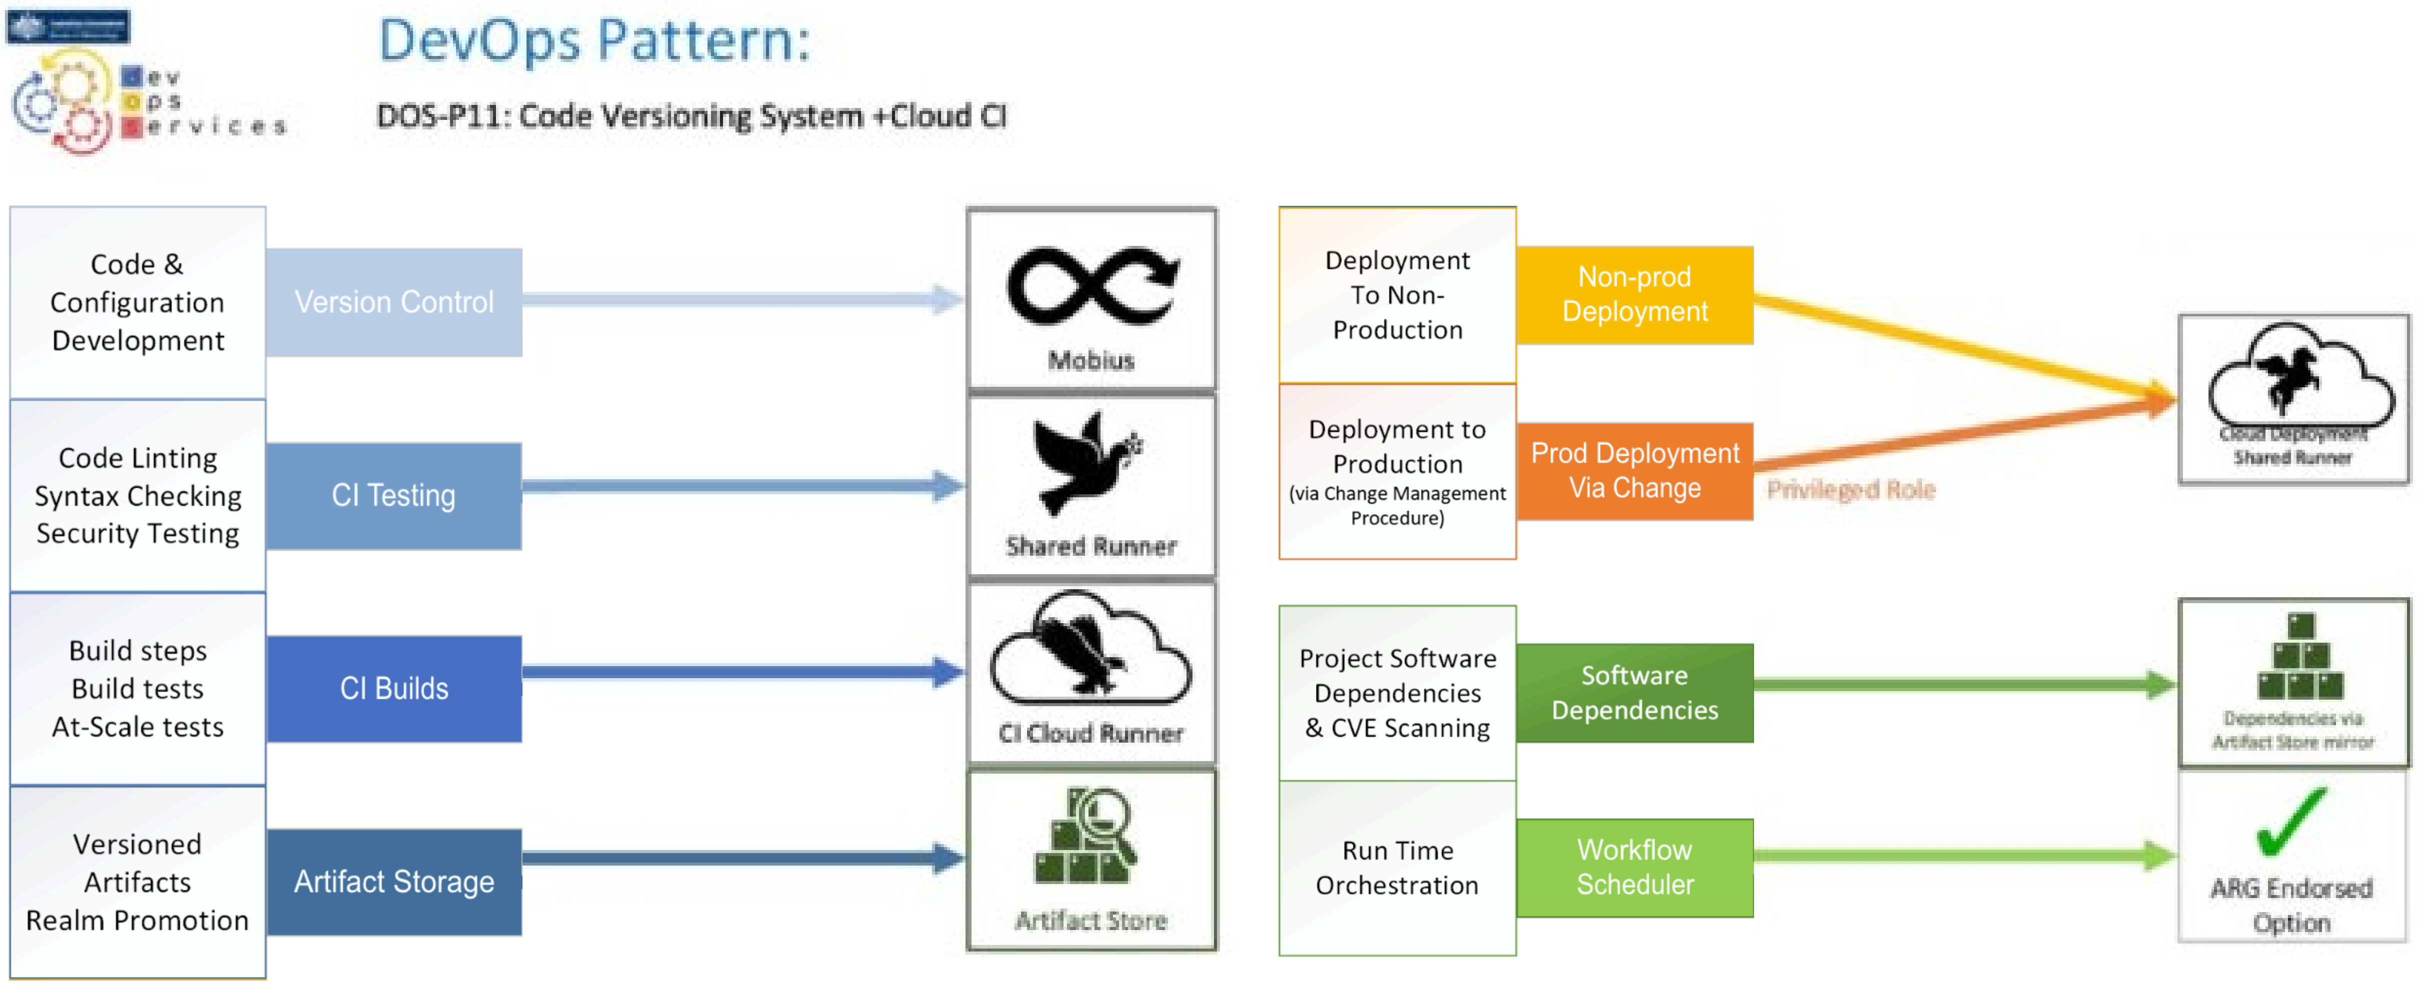

Example of our Patterned approach - Details

- Each DevOps pattern has a high-level design diagram and a detailed document which explains how it is used and itegrated.

- DevOps runs a consultation service, so that project/operational teams can engage to understand how they can be aligned and compliant in direction.

ID	Pattern details	What this means for your project
1	DOS-P11 allows projects with code in Bureau legacy Mobius (internal hyperlink) to securely deploy to cloud environments (currently AWS) from project artifacts and software dependencies held in Bureau legacy Artifact Store (internal hyperlink)	*DevOps infrastructure is ready to go, so your project can start working right away*
2	DOS-P11 permits connection via a proxy server to AWS public API endpoints to allow automated infrastructure deployment to AWS using approved tooling such as Terraform	*Activate automated deployments from Bureau to the Cloud*
3	DOS-P11 permits Mobius pipelines to connect to projects AWS assets via project SecureConnect (see Project Actions 2. below) to allow application maintenance/configuration, integration testing etc.	*Seamlessly update, maintain and test your Cloud deployments using DevOps tooling*
4	DOS-P11 uses a Secrets Management service to hold project secrets that are required during deployment – i.e. AWS Credentials, TLS certificates, SSH keys etc	*Secure handling of project credentials in the DevOps secrets management service*
5	Security testing & appraisal has been conducted on the DOSP11 solution infrastructure and can be referred to but does not need to be included in project security testing.	*The security testing/approval work has already been done so you can get to deployment quicker*

Patterned approach benefits

The benefit of the Pattern approach

- Projects are not reinventing the wheel each time, rather drawing from the work of previous projects, which saves the organization time.
- Governance compliance is easier to mange and plan for.
- With more known items the chance of unknown issues is lessoned.

Efficiencies of the Pattern approach

- Architectural approval for a project is easier to obtain and the Competency Team has done that work.
- Security profile has already been agreed to and understood.
- Unit and regression testing is already configured.
- Operational lifecycle costs are available from similar projects which helps with costing conversations.

Where we have been – People-led Pipelines

Science or Business groups

Deliver Scientific code & Operational Code
Deliver scientific & project outcomes

IT groups

Meet local and Whole of the Government governance
Be the run-time owner operator

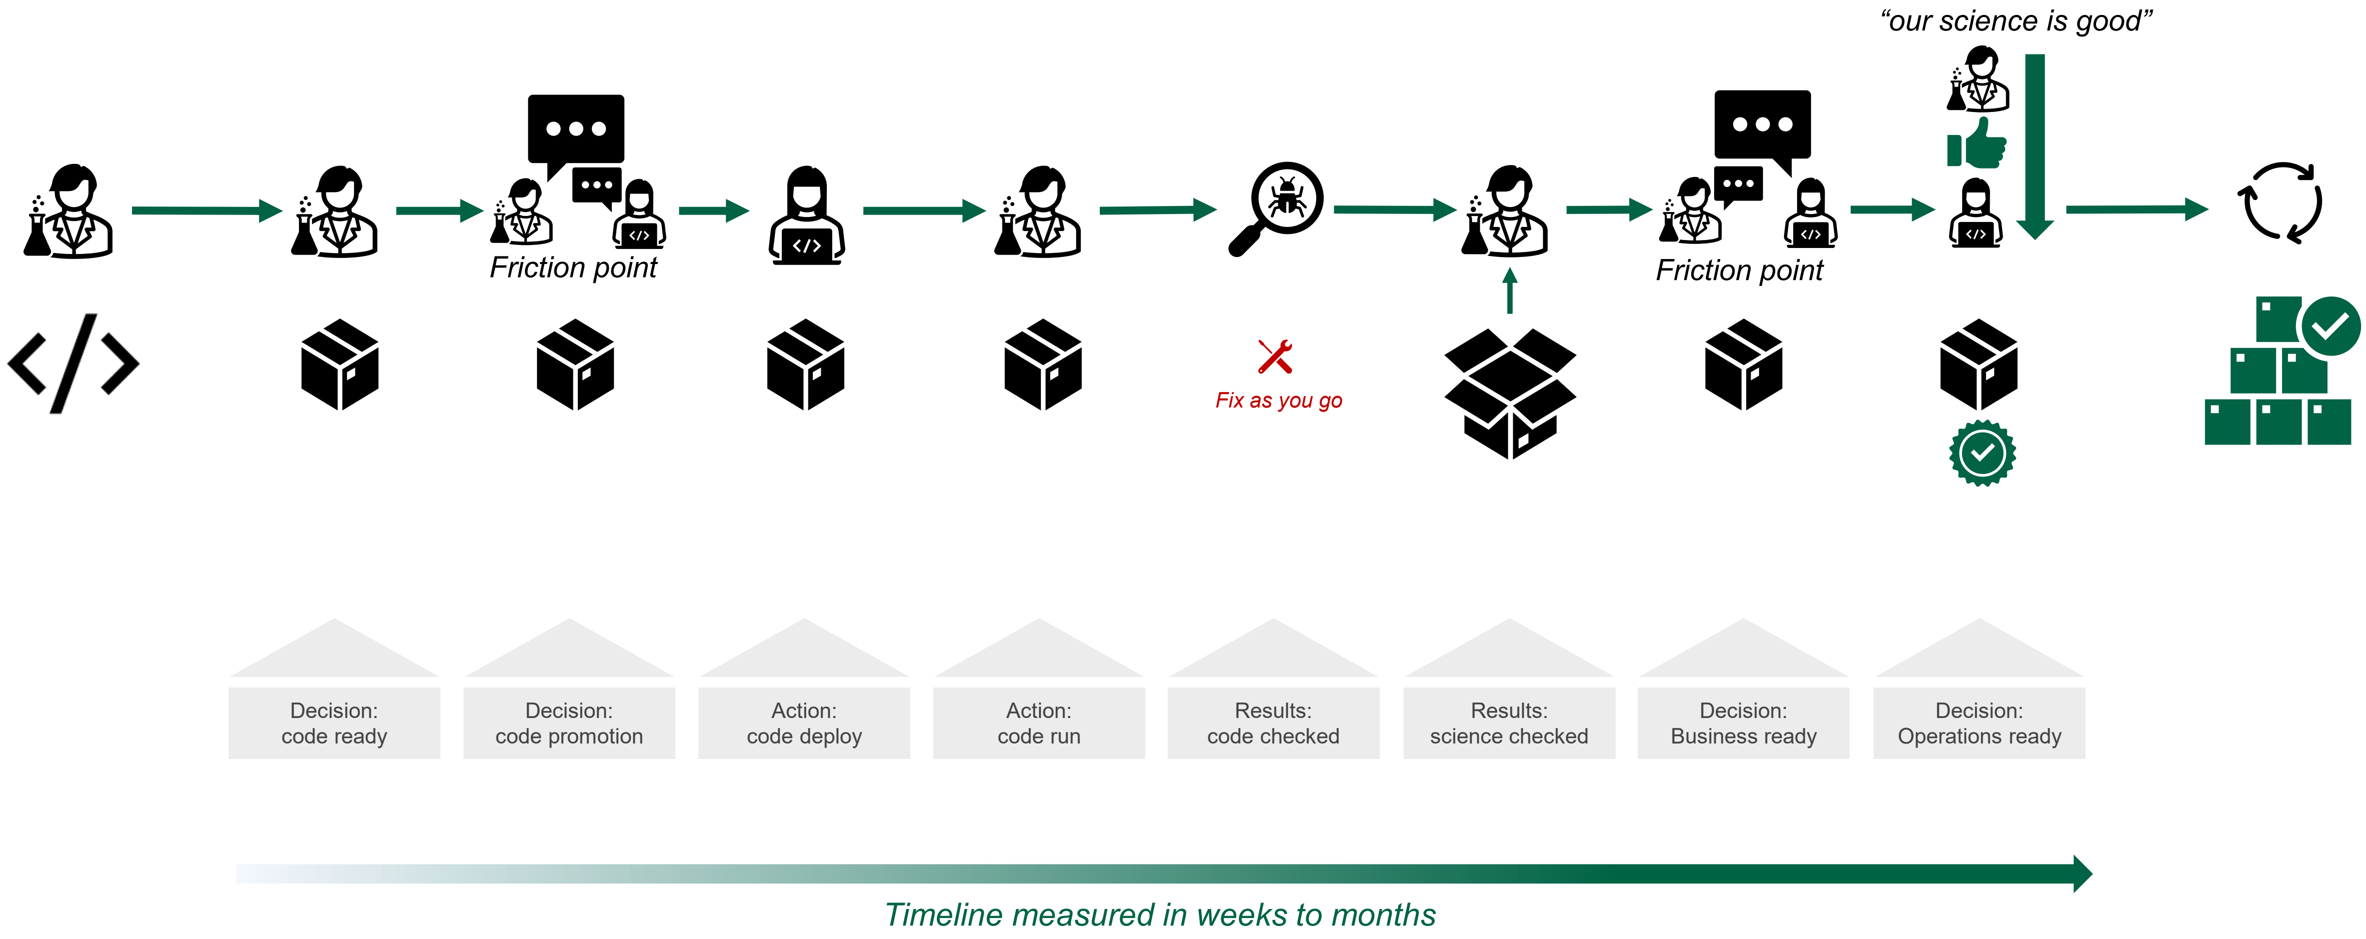

Where we are going – People-less Pipelines

Science or Business groups	*IT groups*	*Automated Build Pipeline*
Deliver Scientific code & Operational Code Deliver scientific & project outcomes	Meet local and Whole of the Government governance Be the run-time owner operator	Pre-set, pre-validated and standardized objective testing Only cares for code

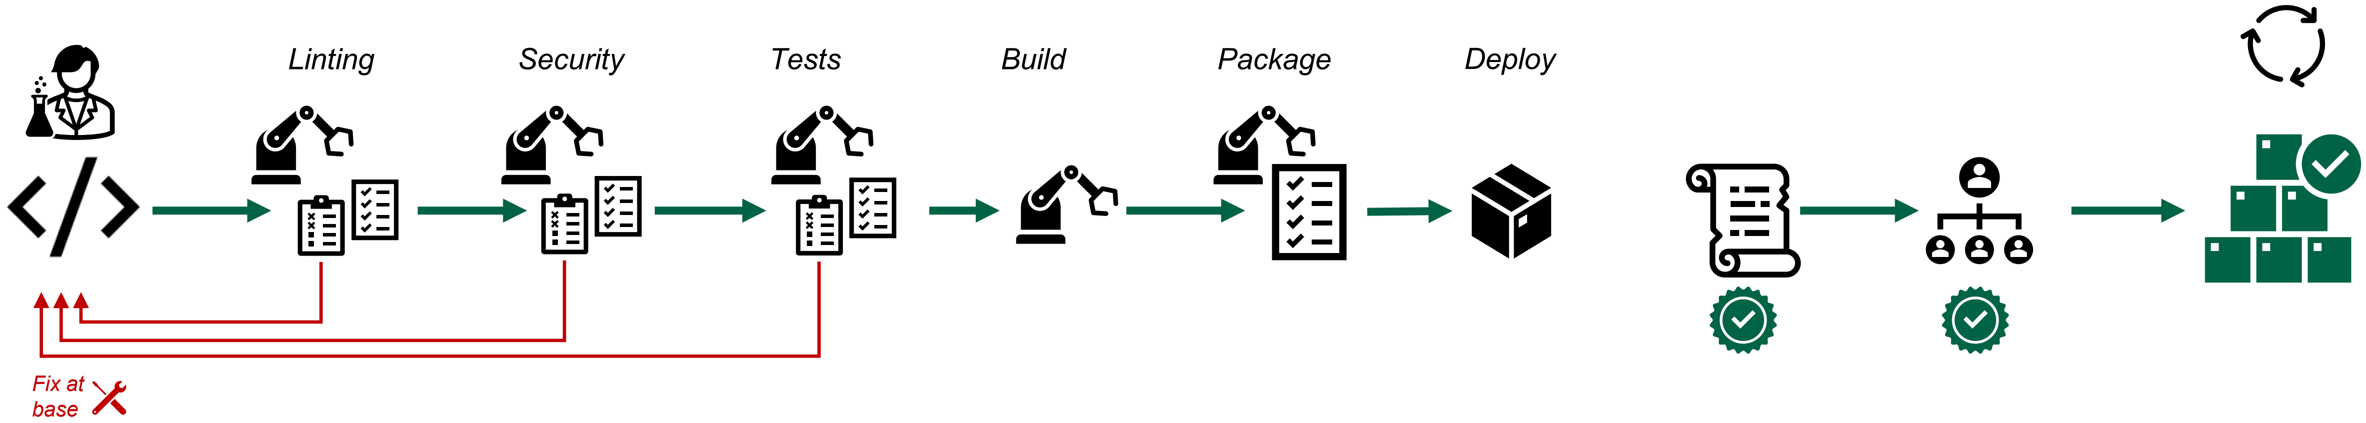

Automation: code quality	Automation: code secure	Automation: code ready	Automation: code builds	Automation: code packaged	Automation: ready for deploy	Decision: Business ready	Decision: Operations ready

Timeline measured in minutes to days

What DevOps is delivering for the Bureau

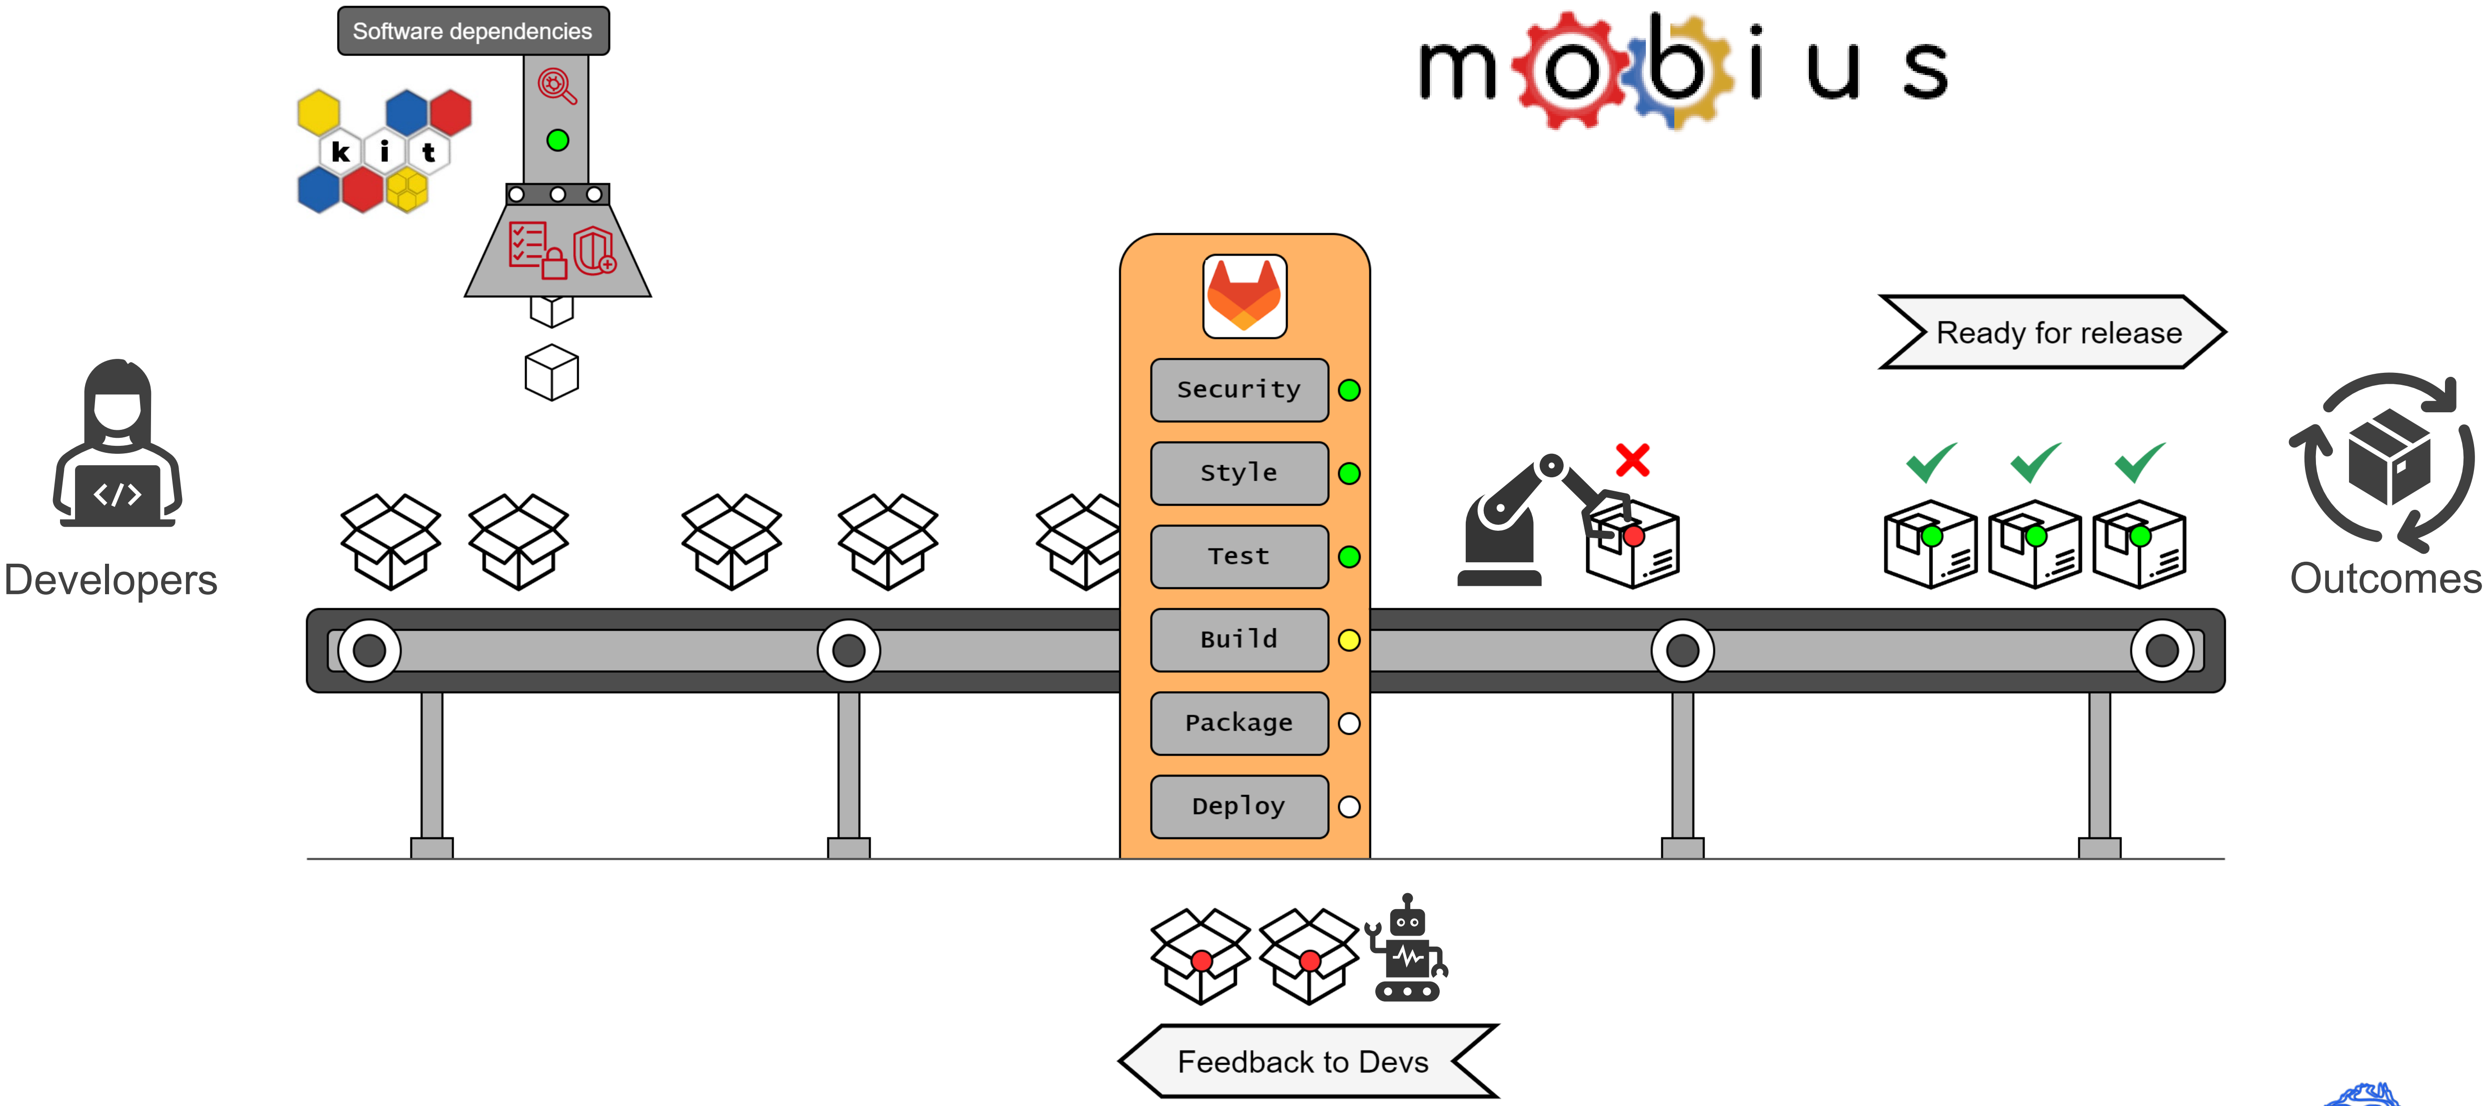

Modelling

Oscar Alves, Charmaine Franklin and Earth System Modelling team

Deterministic Models

Prediction System	APS1	APS2	APS3	APS4	APS5
HPC System	**Ngamai**	**Ngamai**	**Australis**	**MTU**	**MTU**
Weather Prediction	**2014**	**2015-16**	**2019-22**	**2024-25**	**2026**
ACCESS-G Global weather (deterministic) Global domain	40km spatial	25km spatial	12k.5m spatial	12.5km spatial	12.5km spatial
	70 vert levels	70 vert levels	70 vert levels	70 vert levels	70 vert levels
	2 x 10-day fc 2 x 3-day fc	2 x 10-day fc 2 x 3-day fc	2 x 10-day fc 2 x 3-day fc 4 x G4 analysis	2 x 10-day fc 2 x 3-day fc with hourly G4 analysis	2 x 10-day fc 2 x 3-day fc with hourly G4 analysis
ACCESS-R / ACCESS-A Regional weather (deterministic) National domain	12km spatial	12km spatial			1.5km spatial
	70 vert levels	70 vert levels			90 vert levels
	4 x 3-day fc	4 x 3-day fc			4 x 3-day fc with hourly analysis
ACCESS-C City/State weather (deterministic) State domains: (VicTas, NSW, Qld, N.Qld, NT, WA, SA)	4km spatial	2.2km spatial	1.5km spatial	1.5km spatial	1.5km spatial
	70 vert levels	70 vert levels	80 vert levels	80 vert levels	80 vert levels
	2 x 36-hour fc	2 x 36-hour fc	1 x 42-hour fc 3 x 36-hour fc	4 x 72 hour fc	4 x 72 hour fc
			hourly C3 analysis and 3-hr fc run	hourly C3 analysis and 3-hr fc run	hourly C3 analysis and 3-hr fc run
	5 domains	6 domains	7 domains	7 domains	7 domains

Ensemble Models

Prediction System	APS1	APS2	APS3	APS4	APS5
HPC System	**Ngamai**	**Ngamai**	**Australis**	**MTU**	**MTU**
Weather Prediction	**2014**	**2015-16**	**2019-22**	**2024-25**	**2026**
ACCESS-GE Global weather (probabilistic) Global domain		60km spatial	33km spatial	33km spatial	33km spatial
		70 vert levels	70 vert levels	70 vert levels	70 vert levels
		24 members	18 members	18 members	18 members
		2 x 10-day fc	4 x 10-day fc	4 x 10-day fc	4 x 10-day fc
ACCESS-AE Regional weather (probabilistic) National domain					2.2km spatial
					90 vert levels
					12 members
					4 x 3-day fc
ACCESS-CE City/State weather (probabilistic) State domains: (VicTas, NSW, Qld, N.Qld, NT, WA, SA)			2.2km spatial	2.2km spatial	2.2km spatial
			80 vert levels	80 vert levels	80 vert levels
			12 members	12 members	12 members
			4 x 36-hour fc	4 x 72-hour fc	4 x 72-hour fc
			7 domains	7 domains	7 domains

Forecast System Improvement Process

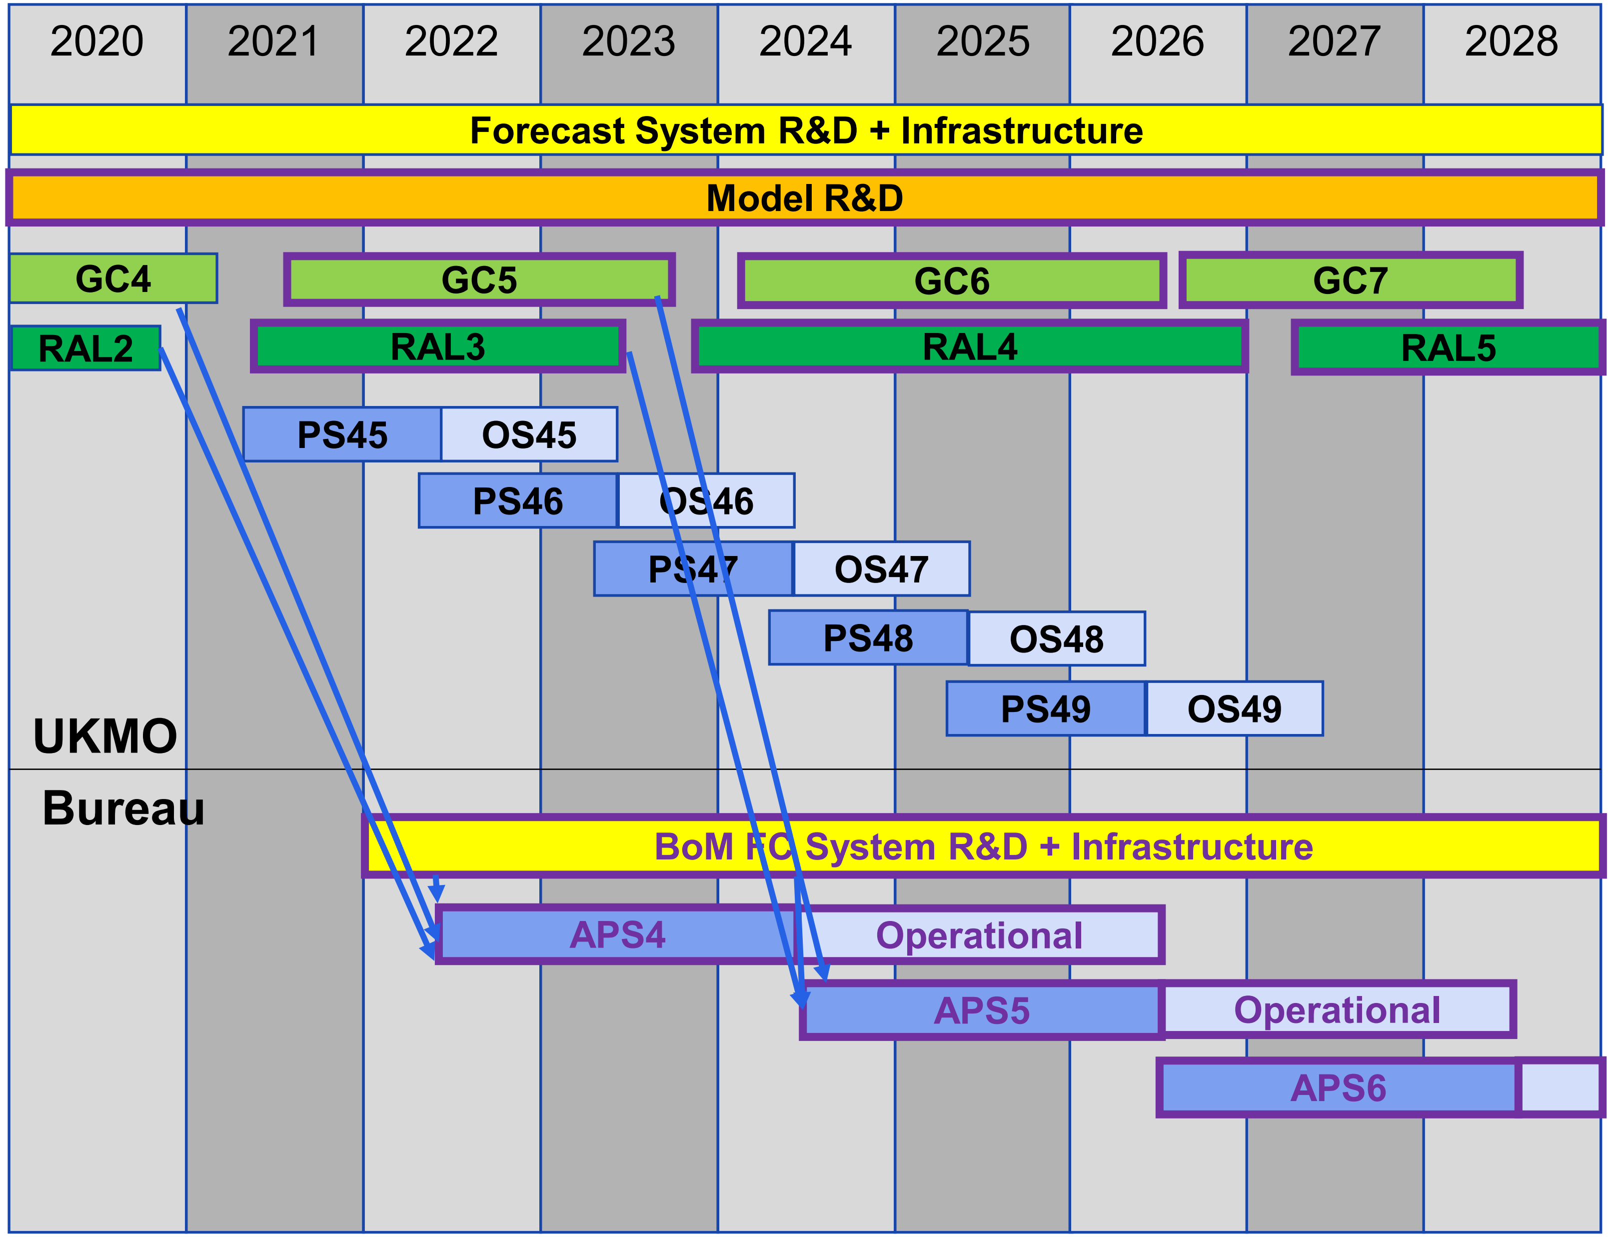

Proposed Bureau Approach

R&D upstream with UKMO
Join UKMO Config Phase (GC & RAL)

Adopt UKMO Parallel Suite Concept

- Every ~2 years
- Upgrade all in one go (dependencies etc)
- **Ready or Skip approach (R&S)**
- R&S also applies to infrastructure
- Parallel run ~1 year (downstream apps)
- Coordination of R&D and Infrastructure
- Cross-project/institute governance
- Decommissioning
- Good Long term planning
- Discipline

Adopt UKMO ‘Trailblaser’ Prototype approach

ACCESS-A

Motivation:

- 7 x 1.5 km operational domains covering most populated areas
- Desire to have high-resolution forecasts nationally
- Australia spans tropics to mid-latitudes – unified configuration

ACCESS-A:

- RAL3.2 configuration, including MORUSES urban scheme
- 1.5 km horizontal grid spacing, 90 vertical levels
- Variable resolution grid to deal with spin up effects at boundaries
- Full domain: 3690 x 3045 x 90 > billion points
- 48-hour lead time

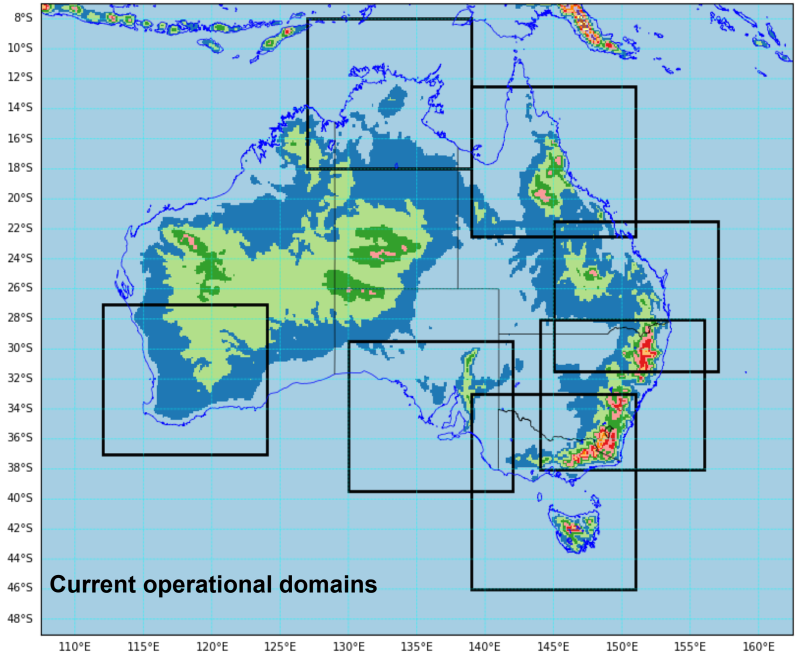

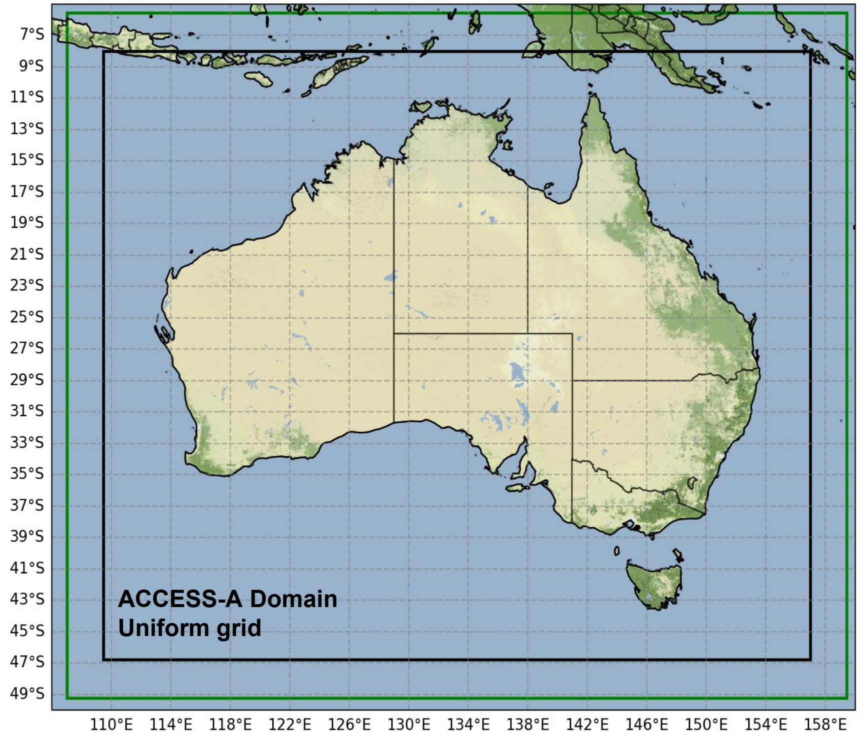

Initial work demonstrated urban-scale modelling benefits for convection

- Convection is under-resolved in km-scale models
- Urban-scale models can give a better representation of the structure of storms
- 100m model firmly in turbulence grey zone

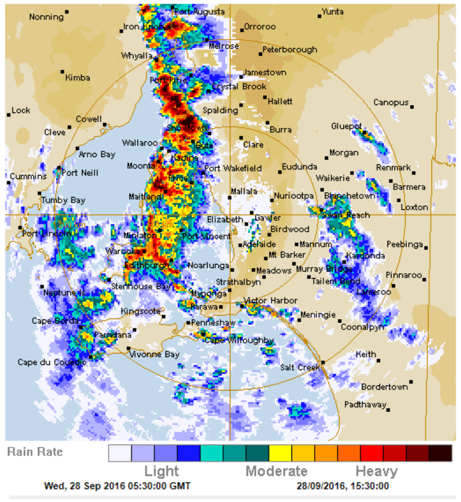

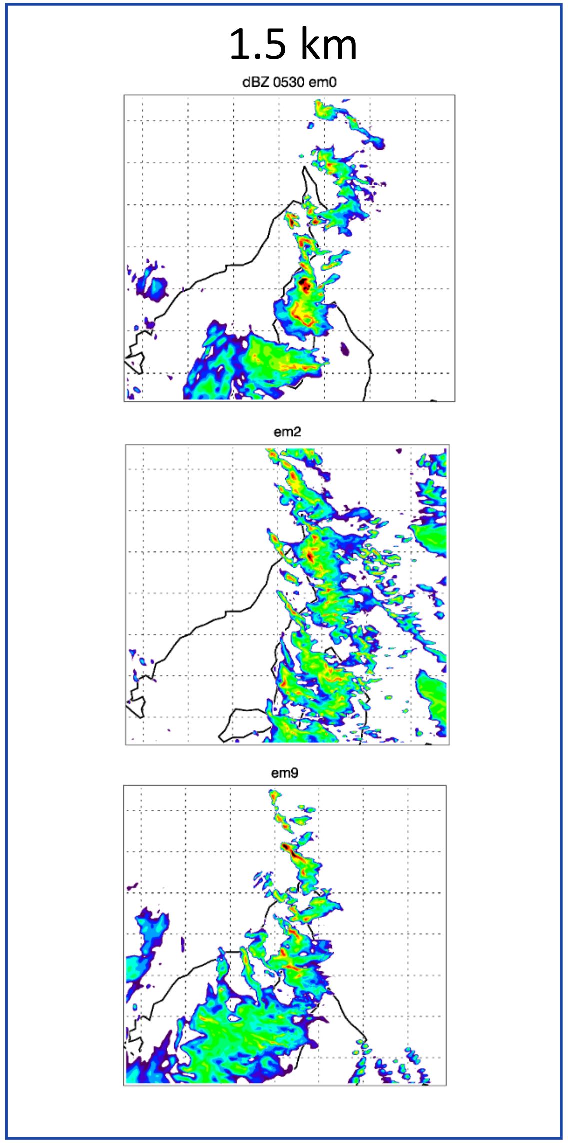

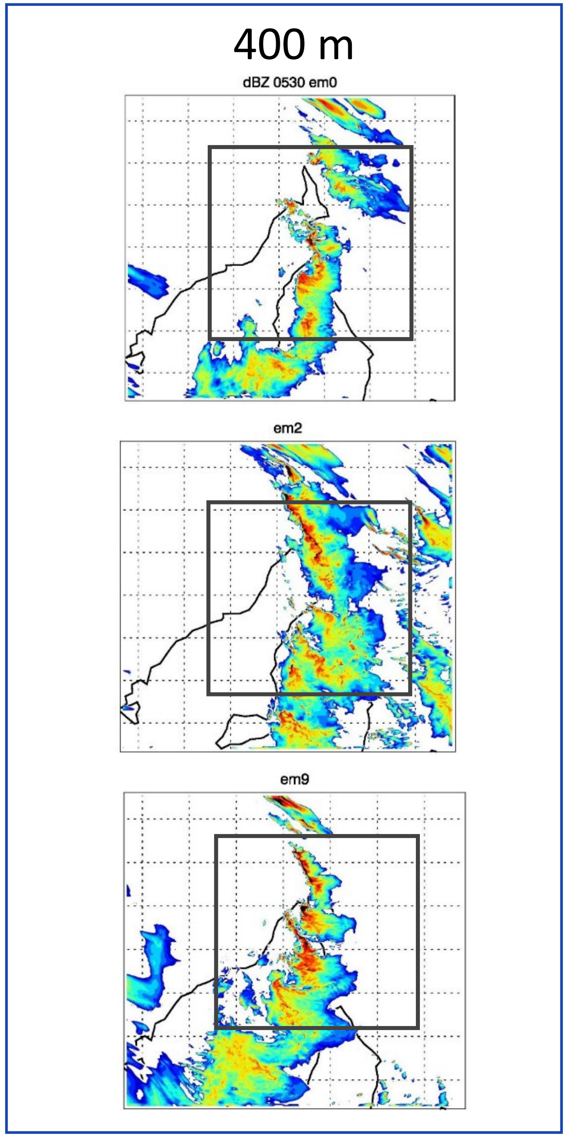

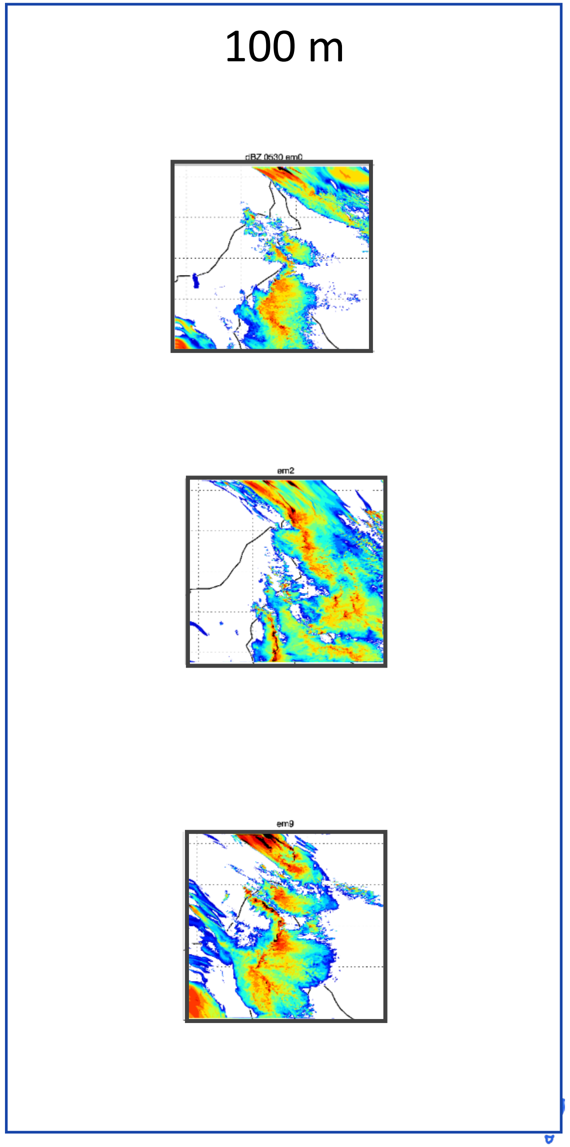

The Bureau of Meteorology

Thank you

Questions?