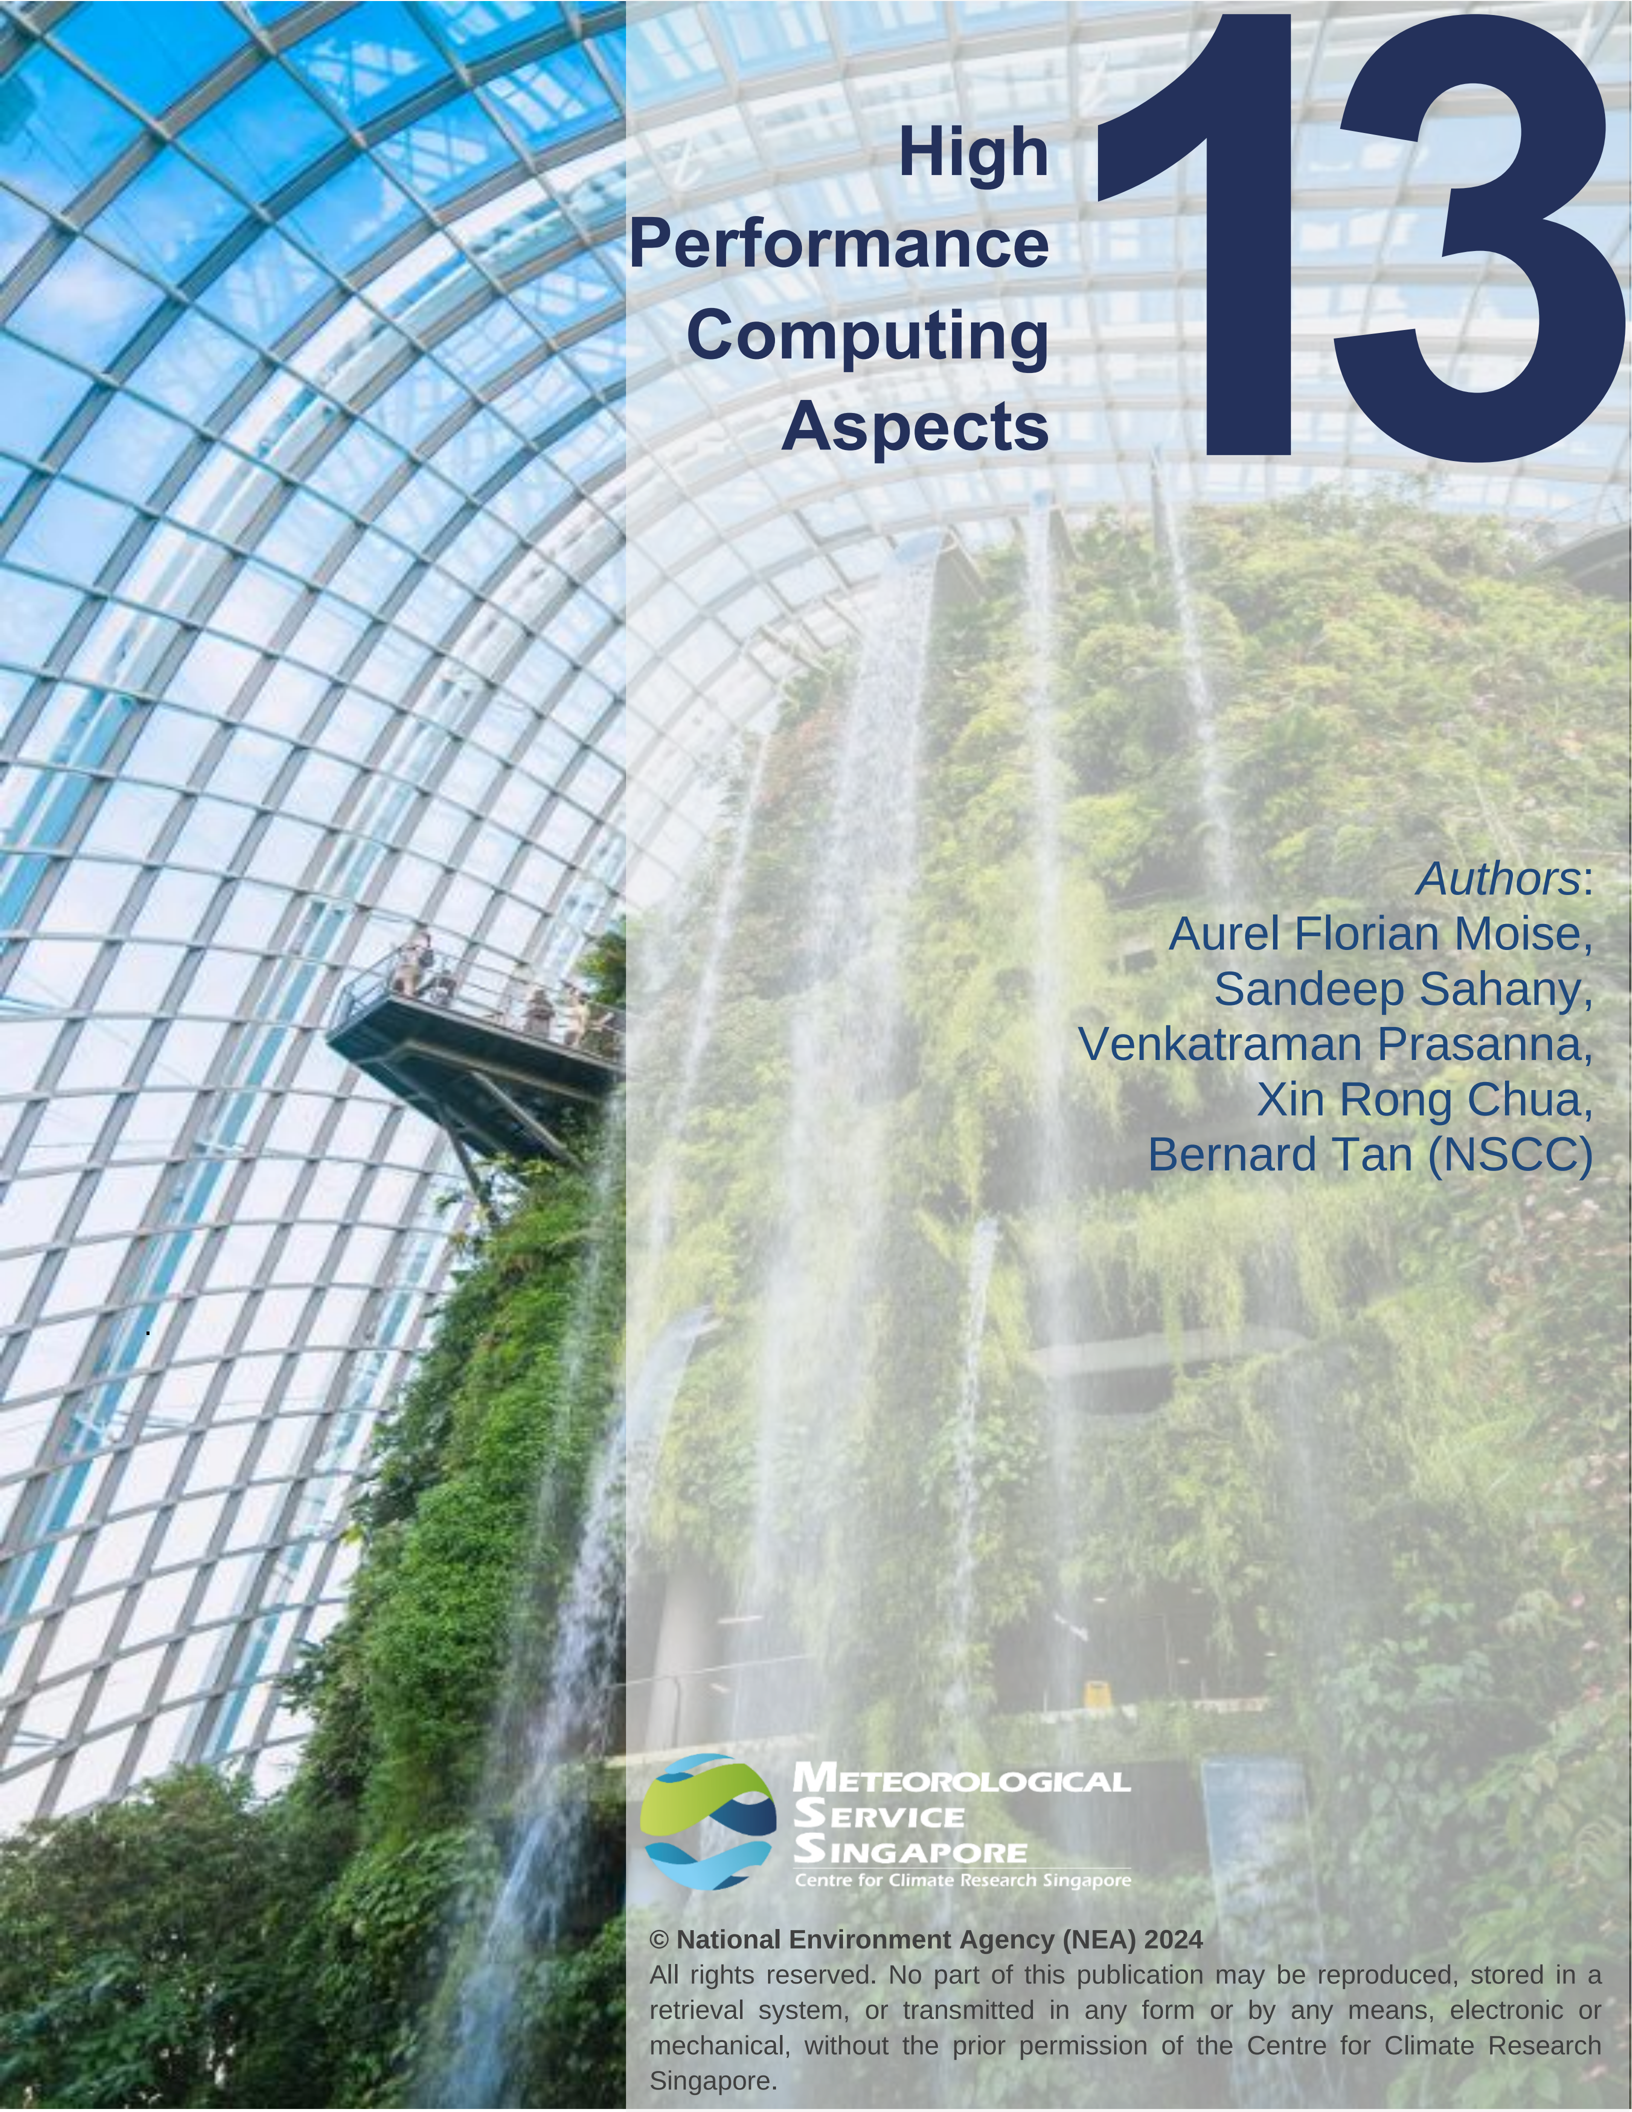

13 High Performance Computing Aspects

Authors:
Aurel Florian Moise,
Sandeep Sahany,
Venkatraman Prasanna,
Xin Rong Chua,
Bernard Tan (NSCC)

Meteorological Service Singapore
Centre for Climate Research Singapore

13.1 Introduction

The dynamical downscaling of CMIP6 GCMs from coarse resolution of ~75 km - 200 km to 8 km resolution over the SEA domain and further downscaling to 2 km resolution over the western Maritime Continent domain was the most time and resource consuming aspect of the entire V3 study. The high-performance computing (HPC) dimension of V3 documented in this chapter will not only serve as a benchmark for the next set of national climate change projections for Singapore (V4) but will also provide useful information for other similar efforts worldwide and to the HPC community in Singapore and elsewhere.

In this chapter we present numerical details of V3 dynamical downscaling simulations, such as details of the regional model used for downscaling, number of grid points in the 8 km and 2 km domains, and the time step size used for each of the resolutions in section 13.2. In the subsequent section (section 13.3) we present details of the 3 HPC systems (Koppen, and ASPIRE 2A at Singapore's National Supercomputing Centre NSCC, and Gadi at Australia's National Computational Infrastructure NCI) used for carrying out the dynamical downscaling simulations. In section 13.4 we present some details about the computing and storage requirements of the V3 study. Scalability tests that were carried out to design workflow of simulations such as length of each simulation chunk and number of parallel chunks are presented in section 13.5.

13.2 Numerical Details of V3 downscaling

The Singapore Variable Resolution (SINGV) model is the numerical weather prediction (NWP) model of the Meteorological Service Singapore that was developed in 2020 (Dipankar et al. 2020). This NWP version of the SINGV model domain covers Singapore, the Malay Peninsula and the Indonesian island of Sumatra, with a grid-resolution of 1.5 km, having 1,092 points in longitude and 1,026 points in latitude. Since the SINGV model was designed to run in an NWP mode, it could not be used as-is in climate mode to run long-term climate simulations. Hence, the SINGV model was tailored to run as a climate model and the climate version is called the SINGV Regional Climate Model (SINGV-RCM). Notably, the diurnal cycle of the sea surface temperatures (SSTs) was implemented in the SINGV-RCM for which the SST fields are interpolated from the driving model grid resolution to the SINGV-RCM grid resolution and updated every 6 hours. Further details of the evolution from SINGV to SINGV-RCM can be found in Chapter 5 of this report.

All numerical models of the atmosphere have a dynamical core, which is responsible for solving the governing equations of atmospheric motion. The dynamical cores used by all operational configurations of the UM prior to July 2014 are called "New-Dynamics" (Davies et al., 2005). Following the implementation of New Dynamics, the Met Office initiated the development of "ENDGame" (Even Newer Dynamics for General atmospheric modelling of the environment) (Wood et al., 2014). ENDGame is an evolution of New Dynamics designed to maintain its benefits, whilst improving its accuracy, stability and scalability. ENDGame uses a semi-Lagrangian advection scheme and a semi-implicit scheme for the temporal discretization of the non-hydrostatic, deep-atmosphere equations of motion (Wood et al., 2014). The V3 8km domain has 1120 points in longitude and 560 points in latitude, and a timestep size of 240s, whereas, the 2km domain has 960 points in longitude and 960 points in latitude, and a timestep size of 120s. The 8km and 2km downscaling domains are shown in Figure 13.1 below.

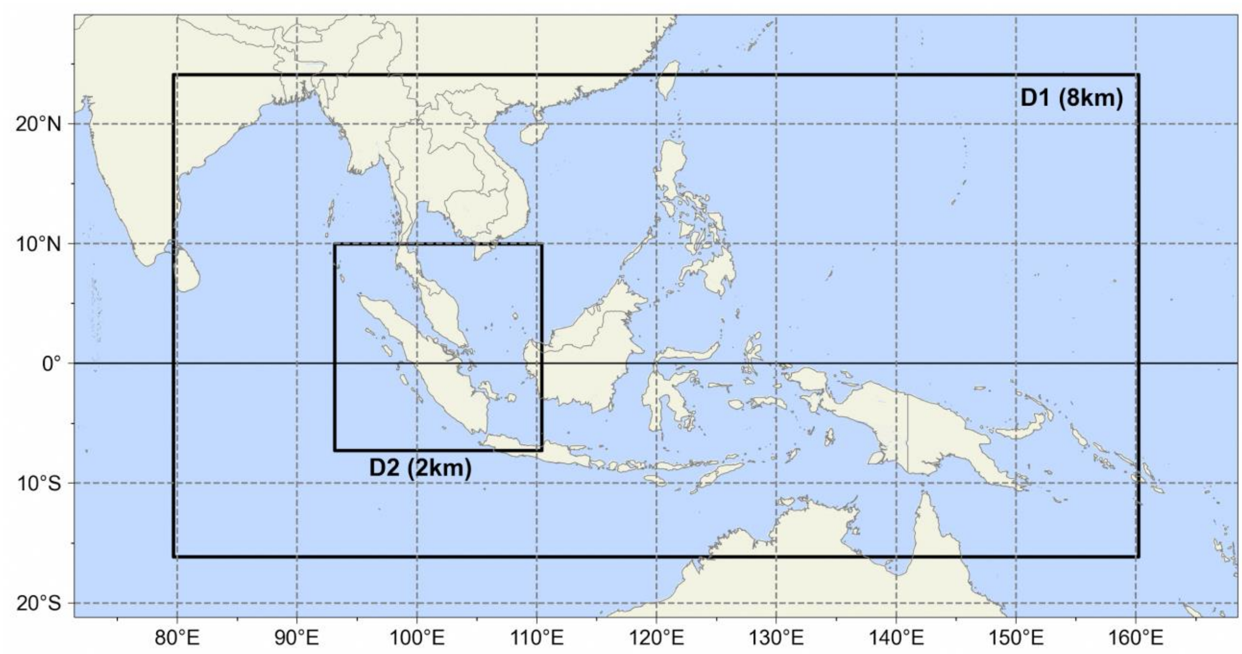

Figure 13.1: V3 downscaling domain

After an extensive period of testing the new SINGV-RCM and conducting various sensitivity experiments (see Chapter 6), the final version and configuration was decided in late 2019 and the initial simulations for V3 projections commenced in early 2020. Since then, V3 simulations were coordinated across 3 HPC systems and 2 continents for a period of 3 years. Final simulations stopped in March 2023. Using effectively tuned HPC resources, it took 1 month to conduct 10 years of 8km resolution simulations and 2 months for 10 years of 2km resolution. Because of the scaling properties of the UM model, simulations also had to be run in chunks of 10 years (for 2km simulations) and 30 years (for 8km simulations) only. Running all of this sequentially, this would have taken over 29 years to complete all V3 simulations. Because of running many simulations simultaneously in parallel, the time was reduced by a factor of 10. However, this took an enormous effort on every staff contributing to V3 to monitor all the parallel simulations, including restarting, trouble-shooting, house-keeping and post-processing.

13.3 HPC Systems Used

The V3 production runs were carried out on 3 HPC systems across 2 continents, two at NSCC (Koppen and Aspire 2A [A2A]), and one at NCI (Gadi) in Australia (refer Figure 13.2). NCI is made available through NSCC's network and arrangements.

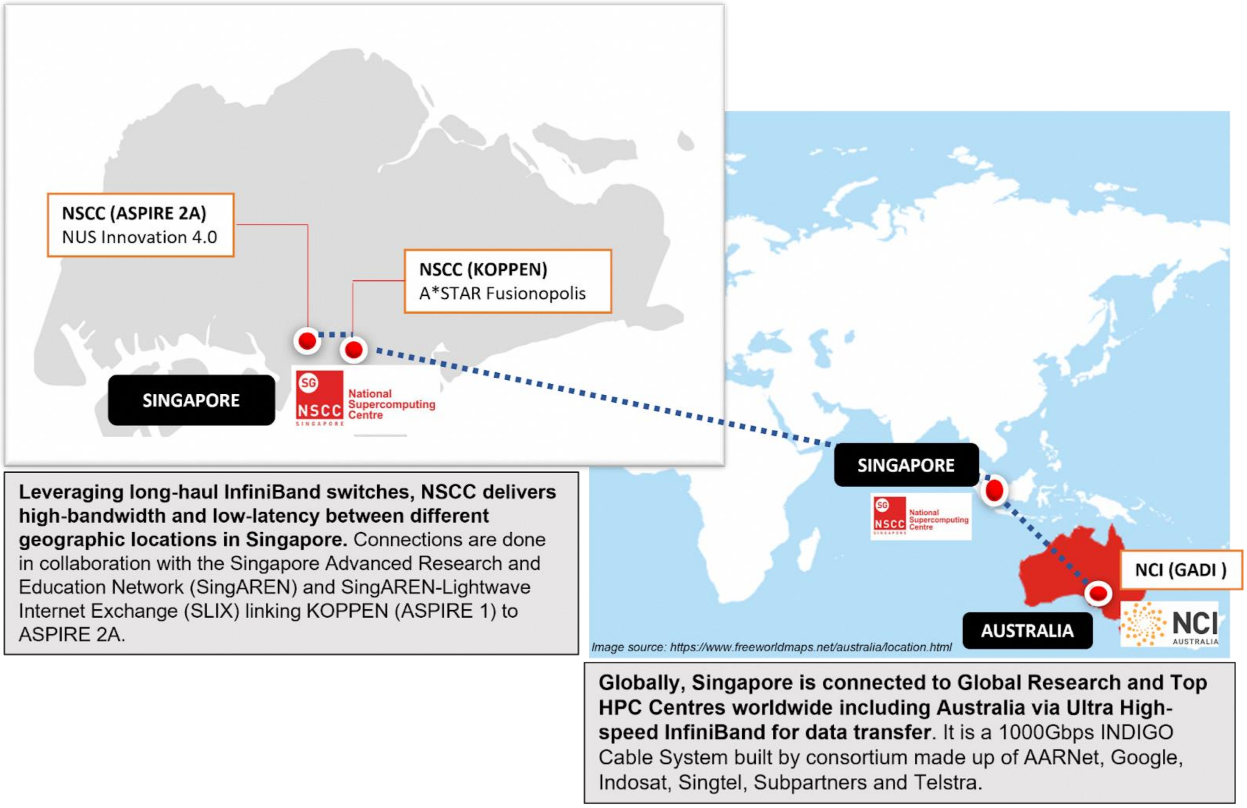

Figure 13.2: Geographical Distribution across 3 HPC Systems

The initial plan was to carry out the initial set of simulations on Koppen in 2020 and then later conduct the majority of the runs on A2A starting February 2021. However, due to COVID-19 related supply chain disruptions in the availability of semiconductor chips during 2020 and beyond, the A2A delivery was delayed by almost 15 months.

To mitigate the potential delay, NSCC proactively employed a multi-pronged strategy to address this matter. At the request of CCRS, NSCC purchased an additional 1 PB of storage equipment (valued at SGD400K) to upgrade Koppen storage from the initial 1 PB to 2 PB to support the V3 project on Koppen.

NSCC also engaged their international partner in Australia, National Computational Infrastructure (NCI), to secure 15 million CPU core-hours/month and 5 PB of storage/month, on their Gadi supercomputer (15 PFlops) for a period of 12 months, starting from July 2021. As a result, NSCC brought the V3 project up to speed within 5 months avoiding a potential 15 months delay due to A2A's late delivery.

Once A2A hardware is available, NSCC made special provisions by providing dedicated access to half of the entire A2A design capacity, and allocated 35 million CPU core-hours/month and 10 PB/month of storage, starting from June 2022 till Dec 2022, to accelerate the V3 research work.

In total, over the 3 supercomputers, NSCC has provided 306M CPU core-hours and 144 PB of storage from July 2021 till Dec 2022. This is about 12 M CPU core-hours and 108 PB more than what CCRS is entitled, based on the original arrangement which is 295 M CPU core-hours and 36 PB of storage over a 12 months period. The additional HPC resources, worth about SGD 3.06M, is provided by NSCC at a goodwill basis at no additional cost to CCRS.

The distribution of V3 simulations across the 3 HPC systems is shown in Table 13.1. As can be seen, almost two-thirds of the simulations were carried out in Singapore, and the remaining one-third in Australia.

Table 13.1: Distribution of Simulations across 3 HPC Systems

V3 Simulations	NSCC Koppen	NCI Gadi	NSCC A2A
ERA5_8km, ERA5_2km	X		
MIROC6_8km	X		
MPI_8km, MPI_2km		X	
NorESM_8km, NorESM_2km		X	
UKESM_8km		X	
UKESM_2km			X
ACCESS_8km, ACCESS_2km			X
EC-Earth_8km, EC-Earth_2km			X

Following is a description of the three HPC systems used for producing all V3 simulations.

Koppen: NSCC Koppen is specifically designed to support HPC research activities in climate and environment research in areas such as advanced modelling and simulation and weather pattern analysis. It is a Cray XC50 supercomputer system with 160 TFLOPS computing capacity, consisting of 52 nodes and 1.2 PB of lustre storage.

The special features of the Cray XC50 supercomputer include:

- the industry-leading Aries network interconnect, which is designed specifically to meet the performance requirements seen in today's emerging class of data center GPU accelerated applications, where high node-to-node communication performance is critical;
- a Dragonfly network topology tightly integrated with Aries that reduces communication latency for scale-out applications that rely heavily on the Message Passing Interface;
- optional SSD-enabled DataWarp I/O accelerator technology, enabling software-defined provisioning of application data for improved performance;
- innovative cooling systems to lower customers' total cost of ownership;
- the next-generation of the high performance and tightly integrated Cray Linux Environment that supports a wide range of applications;
- image-based systems management for easy upgrades, less downtime, and field-tested large-scale system deployment;
- enhancements to Cray's HPC optimized programming environment for improved performance and programmability of GPU environments;
- support for next-generation Intel Xeon and Intel Xeon Phi processors.

The detailed Koppen system information can be found in the Table below:

Table 13.2: Koppen HPC Systems details

Cray XC50-AC		
System peak performance		166.4 TFLOPS
Total system memory (compute node only)		12 TB
Compute nodes	Number of compute blades/nodes	13/(48*+4*=52)
	Processor type	Intel Xeon Cascade Lake 6248 20-core 2.5GHz
	Number of cores per socket/node	20/40
	Peak performance per socket/node	1,600 / 3,200 GFLOPS
	Memory type	*48 nodes with 16 GB DDR4-2933 DIMM *4 nodes with 64 GB DDR4-2933 DIMM
	Memory per socket/node	*48 nodes with 96/192 GB *4 nodes with 384/768 GB
	Memory bandwidth per node (peak)	281.6 GB/s
	Memory capacity per core	48 nodes with 4.8 GB per core (4.65 GB usable) 4 nodes with 19.2 GB per core
Inter-connection	Technology	Cray Aries
	Topology	Dragonfly
	Injection bandwidth	16 GB/s
	Node bandwidth	157.5 Gb/s (per direction, single cabinet of XC50-AC)
	Global bandwidth	1,260 GB/s
	Bi-sectional bandwidth	672 GB/s
Login Node	2 nodes of elogin	elogin1, elogin2

Storage System	
Model	Cray ClusterStor L300N with NXD flash accelerator
Parallel File System	Lustre
Number of System Management Unit	2
Raw Capacity	1,640 TB
Usable Capacity	Approximately 1,200 TB
Sustained I/O Bandwidth	10 GB/s
Connectivity	InfiniBand EDR

A simplified layout of the Koppen system is shown in Figure 13.3 below.

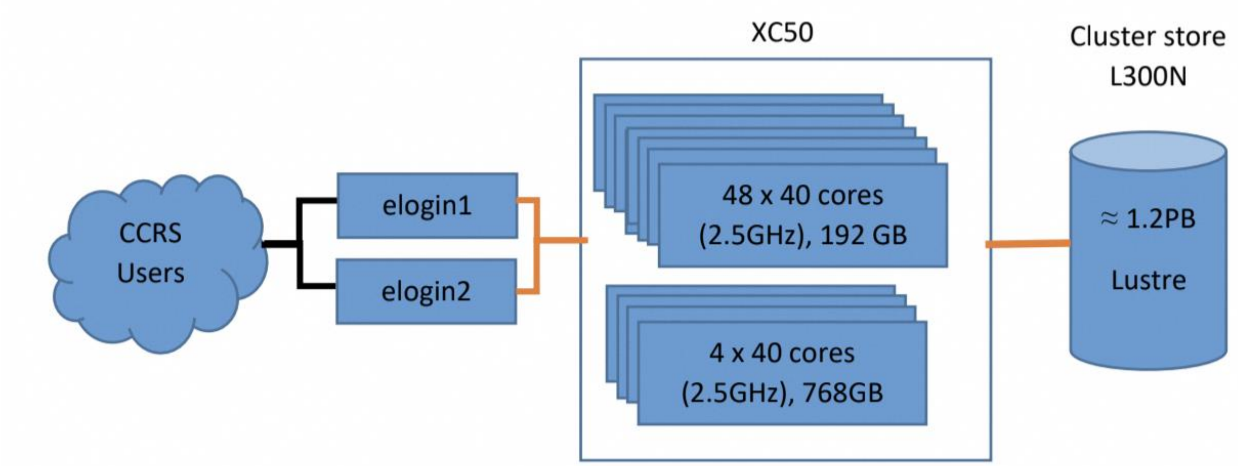

Figure 13.3: Simplified layout of Koppen (credit: NSCC)

ASPIRE 2A: NSCC ASPIRE 2A's (Figure 13.4) core computing capabilities deliver a level of performance and flexibility needed to support a multifaceted array of HPC applications. The computational components are balanced with high-speed storage subsystems and a low latency high speed interconnect that ensures to deliver the highest levels of performance across a broad spectrum of applications. It is an AMD-Based Cray EX supercomputer with 8 PB of GPFS FS and 10 PB of Lustre FS storage and Slingshot interconnect.

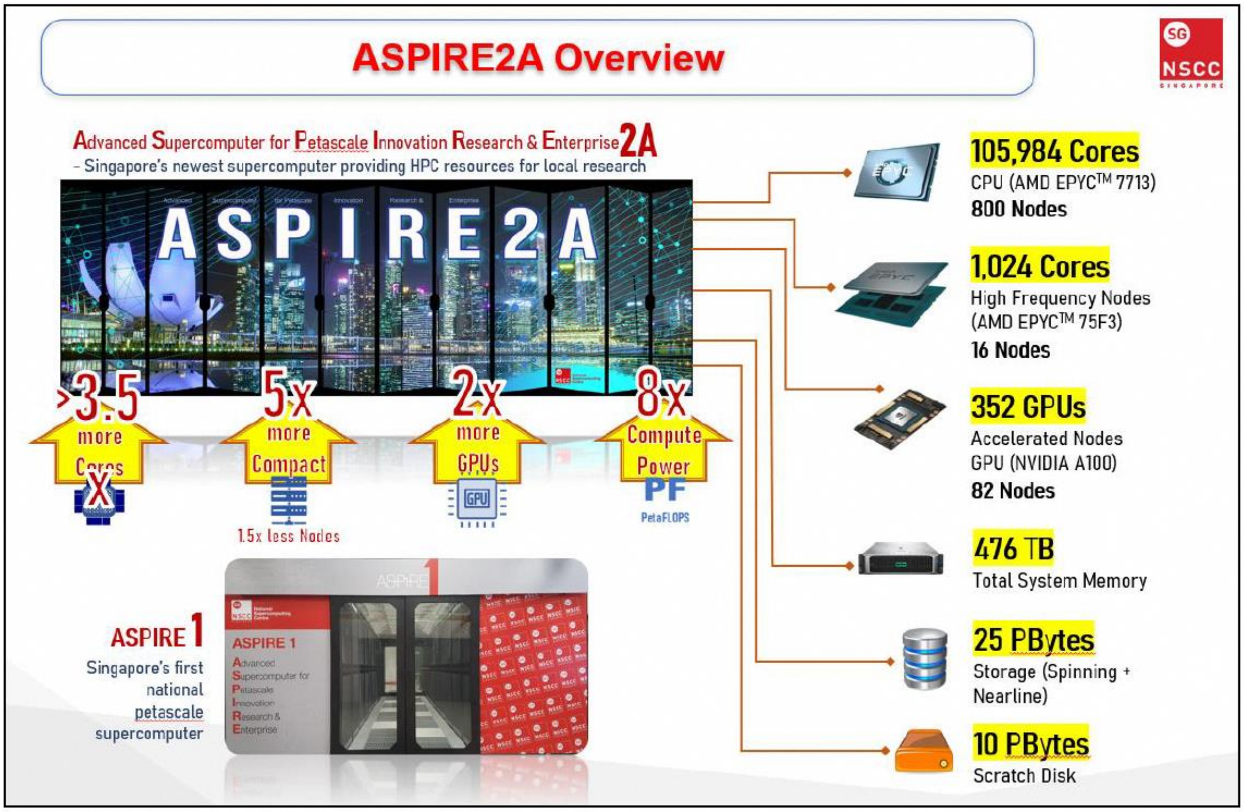

Figure 13.4: ASPIRE 2A specifications (credit: NSCC).

The building block of the HPE Cray EX supercomputer is the liquid cooled cabinet, a sealed unit that uses closed-loop cooling technology. (Figure 13.5) Each EX cabinet holds eight compute chassis and a total of 64 blades with eight Slingshot injection ports per blade. Each blade supports four dual-CPU nodes for a total of 512 processors per cabinet.

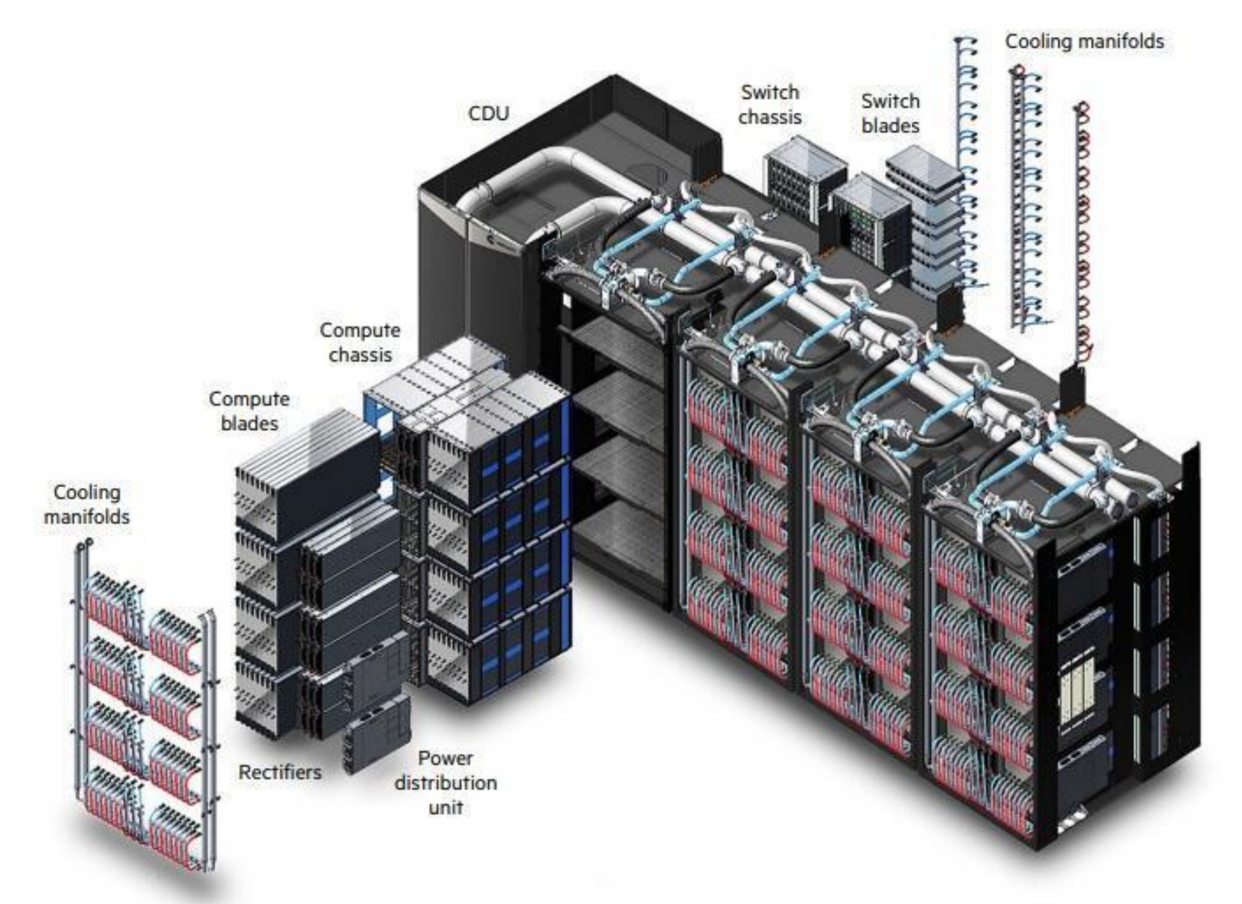

Figure 13.5: Building Block of HPE Cray EX supercomputer

This NSCC HPC **ASPIRE2A** system comprises of the following:

1. HPE Cray EX 2x AMD EPYC Millan 7713 providing total compute capacity of up to 10 PFlops, 512 GB memory and 128 cores per node. This includes a GPU compute capability with 4 x NVIDIA A100-40G SXM per node.
2. AI System: Total of 18 AI GPU nodes. This is divided into 12 nodes with 4x Nvidia A100 40GB and 12 TB nvme local storage; and further 6 nodes with 8x Nvidia A100 40GB and 14TB nvme local storage. The access to the AI systems is via ASPIRE2A "ai" queue.
3. High Frequency: 16 DL385 High Frequency Nodes. These are supported by a Dual-CPU AMD 75F3 (32 cores/CPU + 32 cores/CPU = 64 cores in a node). Additional components are a 100G High speed network, and 512GB DDR4 ECC RAM (User accessible RAM = 500 GB). The operating system is running Red Hat Enterprise Linux-8.
4. High Speed Network Interconnect (HPE Slingshot): All nodes are connected with HPE Slingshot Interconnect (Dragonfly Topology). HPE Slingshot provides a modern, high-performance interconnect for HPC and AI clusters that delivers high-bandwidth and low-latency for HPC, ML, and analytics applications.
5. Additional Features: Remote extended network connections to the A*STAR, NUS, NTU, SUTD and NEA sites; a parallel file system (Lustre & PFSS); Liquid cooled high-density Cray EX cabinets; Air cooled racks (specialized AI, Large memory, storage, login nodes) and an Altair Workload Manager.

NCI Gadi: Gadi (see Figure 13.6 below) is NCI's peak supercomputer and Australia's most powerful CPU-based research supercomputer. It is a 4,962 node supercomputer comprising Intel Sapphire Rapids, Cascade Lake, Skylake and Broadwell CPUs and NVIDIA V100 and DGX A100 GPUs, Gadi supports diverse workloads with well over 10 petaflops of peak performance.

Figure 13.6: NCI Gadi

Gadi contains more than 250,000 CPU cores, 930 Terabytes of memory and 640 GPUs. The technical specifications of Gadi are:

- **3,074 nodes** each containing **two 24-core Intel Xeon Scalable 'Cascade Lake' processors** and *192 Gigabytes of memory*
  - This includes 50 nodes each offering 1.5 Terabytes of Intel Optane DC Persistent memory
- **720 nodes** each with **two 52-core Intel Xeon Scalable 'Sapphire Rapids' processors** and *512 Gigabytes of memory*
- **804 nodes** each with **two 14-core Intel 'Broadwell' processors**
- **192 nodes** each with **two 16-core Intel 'Skylake' processors**
- **160 nodes** each containing **four Nvidia V100 GPUs** and two 24-core Intel Xeon Scalable 'Cascade Lake' processors.
- **10 nodes** each with **two 14-core Intel 'Broadwell' processors** and *512 Gigabytes of memory*

.

- **2 nodes** of the **NVIDIA DGX A100** system, with 8 A100 GPUs per node.
- Linking the storage systems and Gadi is Mellanox Technologies' latest generation HDR InfiniBand technology in a Dragonfly+ topology, capable of transferring data at up to 200 Gb/s.

13.4 Computing and Storage Requirements

Given that V3 involved running very high-resolution long-term simulations (a total of ~2000 model years at 8 km resolution and ~750 model years at 2 km resolution) for multiple GCMs (6 for 8 km runs and 5 for 2 km runs) and multiple scenarios (historical plus 3 SSP's), this required a large amount of computing and storage.

Considering the needs and available resources, both storage and compute was regularly monitored and adjusted throughout the project in close collaboration with NSCC (see Table 13.2). Monthly meetings assured that any issues were promptly dealt with and simulations were not significantly affected.

Although the required amount of computing and storage varied across the time span of the production runs which lasted for more than 3 years, the peak computing usage was around 35 million core-hours per month on A2A and around 15 million core-hours per month on Gadi.

The peak storage was around 12 PB on A2A and around 5PB on Gadi. The peak computing usage on Koppen was around 1.2 million core-hours per month, and for storage it was around 2 PB.

After the completion of the simulations on Gadi in Australia, around 4 PB of data was transferred back to Singapore (A2A) with the help of NSCC and the Singapore Advanced Research and Education Network (SingAREN). After employing parallel transfers, we obtained peak throughputs on the order of 1 Gbps

	Jun-21	Jul-21	Aug-21	Sep-21	Oct-21	Nov-21	Dec-21	Jan-22	Feb-22	Mar-22	Apr-22	May-22	Jun-22	Jul-22	Aug-22	Sep-22	Oct-22	Nov-22	Dec-22
NCI-Australia																			
Million Core-Hours	Testing **	Testing **	12.50	12.50	12.50	**15**	**15**	**15**	**15**	**15**	**15**	**15**	**15**	**15**					
PFLOPS	0.00	0.00	1.43	1.43	1.43	1.72	1.72	1.72	1.72	1.72	1.72	1.72	1.72	1.72					
Storage (Project) in PB	0.5	0.5	2.0	2.0	4.5	4.5	4.5	5.0	7.0	7.5	8.0	8.5	9.0	9.5	5.0	5.0			
Storage (Scratch) in PB	0.2	0.2	1.0	1.0	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5			
Aspire2A-NSCC																			
Million Core-Hours												Testing **	35	35	35	35	35	35	35
PFLOPS												0.00	4	4	4	4	4	4	4
Storage (Project) in PB												6	10	10	10	10	10	10	10
Storage (Scratch) in PB												1	2.0	2.0	2.0	2.0	2.0	2.0	2.0
Koppen-NSCC																			
Million Core-Hours	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40	1.40		
PFLOPS	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16		
Storage (Total) in PB	1.5	1.5	1.5	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0		

	Jan-23	Feb-23	Mar-23	Apr-23	May-23	Jun-23	Jul-23	Aug-23	Sep-23
Aspire2A-NSCC									
Million Core-Hours	35	35	**Postprocessing and analysis**						
PFLOPS	4	4							
Storage (Project) in PB	10	10	10	10	10	10			
Storage (Scratch) in PB	2.0	2.0	2.0						

Table 13.2: Screenshot of a resource table projecting the V3 computing and storage estimates on the 3 HPC systems from one of the monthly MSS-NSCC monthly meetings held on 08th March 2022.

After the processing of all the level-0 data to levels 1, 2, and 3, and final housekeeping, around 5-6 PB of storage would still be needed, with a large fraction of this data stored in tapes and the rest as readily accessible live storage for the purpose of data sharing with stakeholders and for further analysis by CCRS.

13.5 Scalability

Scalability is the process optimisation for running compute-heavy simulations on HPC. This optimisation could significantly affect both timing and costs of resources.

Carrying out scalability tests by changing the number of CPUs and parallelization options (e.g. Message Passing Interface [MPI], threading such as OpenMP) to design the workflow of model simulations is standard practice when producing long-term climate change projections in order to optimise computational resource utilisation. As a part of V3 workflow design, systematic scalability tests were carried out for both the 8km and 2km model configurations.

Two metrics were used to measure scalability - minutes per simulated day (MPSD) and the scalability factor (SF). While MPSD is defined as the wall-clock (real) time required to complete 1 day of model simulation (model time), SF is defined as the ratio of theoretical MPSD based on linear scaling with the number of CPUs to actual MPSD. The results from the performed scalability test for the 8 km and 2 km resolutions are shown in Tables 13.3 and 13.4, respectively. The optimisation process tries to minimise MPSD while making sure the SF is sufficiently high. Typically, SF should be at least 0.9 while aiming for a low MPSD.

The tables show the total number of nodes used, the number of CPUs along the x direction of the domain (zonal direction) denoted by Xprocs, the number of CPUs in the y direction of the domain (meridional direction) denoted by Yprocs, number of threads used by each process, MPSD and SF.

Table 13.3: 8 km Scalability Test

Number of compute nodes used for simulation	Number of CPUs used for longitude computes	Number of CPUs used for latitude computes	Minutes per simulated day	Scalability factor
6	16	48	15.7	1
9	30	38	11.3	0.92
10	28	42	11.3	0.83
10	30	42	10.3	0.91
15	30	62	50.7	0.12

It is to be noted that although we tried the multithreading option, and the model showed improved scalability performance with this option, the results failed bit-reproducibility tests considered an important criteria to carry out long-term simulations, and hence we had to use the single thread option for our simulations. This behaviour originates from the actual core model used (UM) and couldn't be adjusted.

Based on the results of the scalability tests shown in Tables 13.3 and 13.4, it was decided to use 10 nodes (1260 CPUs) for each chunk of the 8km simulation and 14 nodes (1764 CPUs) for each chunk of the 2km simulation.

Table 13.4: 2 km Scalability Test

Number of compute nodes used for simulation	Number of CPUs used for longitude computes	Number of CPUs used for latitude computes	Minutes per simulated day	Scalability factor

8	32	32	33	1
10	40	32	29.3	0.9
12	48	32	23	0.96
14	42	42	21	0.9
18	48	48	17.7	0.83

13.6 Summary

In summary, a total of ~2000 model years of dynamical downscaling simulations covering almost entire SEA at 8 km resolution and ~750 model years of dynamical downscaling simulations covering the western Maritime Continent at 2 km resolution using the SINGV-RCM were carried out for the current and future climates as a part of V3.

Based on the availability of computing resources and V3 timelines, the simulations had to be carried out on 3 different HPC systems spanning 2 continents, 2 systems in Singapore (NSCC Koppen and A2A) and 1 in Australia (NCI Gadi). Although the computing and storage requirements kept varying throughout the simulation period, the peak computing exceeded ~35 million core-hours per month and peak storage went as high as ~12 PB.

In addition, around 4PB of data was transferred over from Australia to Singapore at peak transfer speeds of ~1 Gbps with the help of NSCC and SingAREN.

Overall, the HPC journey to accomplish the V3 simulations was challenging and adventurous, comprising significant resources from CCRS staff and dedicated time for managing the resources and simulations. With dedicated teamwork within MSS/CCRS and support from NSCC, NCI and SingAREN, along with the hard work by the V3 team at CCRS, the challenges were successfully maneuvered, and the desired outcome was achieved in time to be able to meet the V3 delivery timelines.