



Technology to aid the analysis of large-volume multi-institute climate model output at a central analysis facility (PRIMAVERA Data Management Tool V2.10)

Jon Seddon¹, Ag Stephens², Matthew S. Mizielinski¹, Pier Luigi Vidale³, and Malcolm J. Roberts¹

¹Met Office Hadley Centre, FitzRoy Road, Exeter, EX1 3PB, UK

²Centre for Environmental Data Analysis, RAL Space, STFC Rutherford Appleton Laboratory, Harwell, Oxford, Didcot, OX11 0QX, UK

³NCAS-Climate, Department of Meteorology, University of Reading, Reading, RG6 6AH, UK

Correspondence: Jon Seddon (jon.seddon@metoffice.gov.uk)

Received: 5 March 2023 – Discussion started: 4 May 2023

Revised: 4 September 2023 – Accepted: 18 September 2023 – Published: 20 November 2023

Abstract. The PRIMAVERA project aimed to develop a new generation of advanced and well-evaluated high-resolution global climate models. As part of PRIMAVERA, seven different climate models were run in both standard and higher-resolution configurations, with common initial conditions and forcings to form a multi-model ensemble. The ensemble simulations were run on high-performance computers across Europe and generated approximately 1.6 PiB (pebibytes) of output. To allow the data from all models to be analysed at this scale, PRIMAVERA scientists were encouraged to bring their analysis to the data. All data were transferred to a central analysis facility (CAF), in this case the JASMIN super-data-cluster, where it was catalogued and details made available to users using the web interface of the PRIMAVERA Data Management Tool (DMT). Users from across the project were able to query the available data using the DMT and then access it at the CAF. Here we describe how the PRIMAVERA project used the CAF's facilities to enable users to analyse this multi-model dataset. We believe that PRIMAVERA's experience using a CAF demonstrates how similar, multi-institute, big-data projects can efficiently share, organise and analyse large volumes of data.

Creative Commons Attribution 4.0 License and the OGL are interoperable and do not conflict with, reduce or limit each other.

© Crown copyright 2023

1 Introduction

The PRIMAVERA project ran from 2015 until 2020 and aimed to develop a new generation of advanced and well-evaluated high-resolution global climate models. High-resolution simulations have been shown to better represent many different aspects of the climate, such as tropical cyclones (Roberts et al., 2020) and atmospheric blocking (Schiemann et al., 2020). Two “streams” of simulations were performed within PRIMAVERA, each consisting of seven different climate models (AWI-CM-1-1, CMCC-CM2, CNRM-CM6-1, EC-Earth3P, ECMWF-IFS, HadGEM3-GC31 and MPI-ESM1-2) that were run at their standard nominal resolution (Taylor et al., 2018) (typically 250 km in the atmosphere and 100 km in the ocean) and at a higher resolution (25 km atmosphere and 8–25 km ocean). All models were run with common initial conditions and forcings according to the HighResMIP (High Resolution Model Intercomparison Project) protocol (Haarsma et al., 2016). The simulations were run on high-performance computers (HPCs) across Europe, and the more than 100 scientists who analysed the data were based at 20 different institutes across Europe with assistance from other global scientists. Perhaps the most challenging aspect of working with

Copyright statement. The works published in this journal are distributed under the Creative Commons Attribution 4.0 License. This license does not affect the Crown copyright work, which is re-usable under the Open Government Licence (OGL). The

global, high-resolution simulations is the volume of data produced, which in this case was a total of 1.6 PiB (pebibytes) of data¹.

A data management plan (DMP) (Mizielinski et al., 2016) was developed to allow these data to be stored, analysed and archived at a central analysis facility (CAF), which in PRIMAVERA's case was the JASMIN super-data-cluster. The Data Management Tool (DMT) software (Seddon and Stephens, 2020) was written to implement the validate and analyse steps of the data workflow in the DMP and to assist with the ingest and disseminate steps. The DMT provided the ability to catalogue and then search the project's data. It tracked and controlled the movement of individual files as they were moved between tape and disk and allowed the data to be published to the global community for sharing at the end of the project.

In this paper we introduce the CAF and the PRIMAVERA project. We then describe the DMP and DMT and explain how they allowed the project's data to be managed. We provide a summary of how the data were transferred to the CAF, discuss the data that were accessed and users' views on the technology, and describe some of the hardware and software opportunities that may be suitable for managing similar projects in the future.

2 The central analysis facility

The CAF used in the PRIMAVERA project is JASMIN, a super-data-cluster that was installed in 2012 (Lawrence et al., 2013). It is funded by the Natural Environment Research Council (NERC) and the United Kingdom (UK) Space Agency and operated by the Science and Technology Facilities Council (STFC). It is located at the STFC Rutherford Appleton Lab (RAL), Harwell, UK. JASMIN's Phase 4 update in September 2018 added 38.5 PiB of new storage co-located with the existing 4000 cores of compute. The storage and compute are tied together with a low-latency network. RAL has a fast connection to Janet, the UK's academic network, which in turn is connected to the GÉANT European network, allowing the fast transfer of data from the HPCs used to run the PRIMAVERA simulations to JASMIN. In addition, JASMIN is connected to an on-site tape library for the offline storage of data. Similar CAFs exist at Centro Euro-Mediterraneo sui Cambiamenti Climatici (CMCC) (CMCC, 2021), Deutsches Klimarechenzentrum (DKRZ, German Climate Computing Centre) (DKRZ, 2021), Lawrence Berkeley National Laboratory on the Cori supercomputer (LBNL, 2021) and other sites.

The compute at JASMIN is split into interactive data analysis servers, the LOTUS batch processing system and some additional private cloud servers. During PRIMAVERA, Intel Xeon processors were used. The compute is primarily de-

signed for batch and parallel processing for scientific analysis, but it also has limited support for multi-node MPI jobs. Part of the storage is dedicated for the Centre for Environmental Data Analysis (CEDA) archive, which contains over 13 PiB of atmospheric and Earth observation data and is connected to the UK's Earth System Grid Federation (ESGF) (Cinquini et al., 2014; Petrie et al., 2021) node. The co-location of the storage, datasets in the archive and compute makes JASMIN a powerful facility to analyse climate simulations with.

All compute hosts at JASMIN run the Linux operating system and have a suite of modern software tools and programming languages installed for the analysis and manipulation of common Earth science data formats. The storage for individual projects is split into allocations called group workspaces (GWSs), each of which is up to 100 TiB (tebibytes).

3 PRIMAVERA

The Horizon 2020-funded PRIMAVERA project was made up of 11 work packages (WPs) (PRIMAVERA and the European Commission, 2015) covering a range of scientific, technical, management, communication and user engagement topics. The stream 1 and 2 simulations described here were run and managed by two of the work packages. Several of the other WPs analysed these simulations and compared them with existing simulations, observations and reanalyses. The remaining WPs carried out model development work or ran small simulations that did not need to be shared with other WPs. These WPs only required a small volume of storage on one of the GWSs.

The stream 1 and 2 simulations follow the HighResMIP protocol and have been submitted to the sixth phase of the Coupled Model Intercomparison Project (CMIP6) (Eyring et al., 2016). The stream 1 simulations consist of a single ensemble member from each model at a standard and a high resolution for six different experiments. The follow-on stream 2 simulations contain additional ensemble members for the experiments performed in stream 1 but with a reduced data output to minimise the volume of data generated and some additional simulations to exploit the new physics modules (for example, the OSMOSIS boundary layer scheme; Gurvan et al., 2019) that were developed in PRIMAVERA.

Due to the scientific complexity, data volumes and the geographical distribution of the project participants, it was recognised that developing and implementing a data management plan would require a significant amount of time and effort. 44 person months of leadership and developer resource was included in the 4-year project proposal to allow this work to take place.

¹Data volumes and rates in this paper use binary prefixes, and so 1 PiB is 1024^5 bytes (IEEE, 2009).

4 Data Management Plan

The PRIMAVERA Data Management Plan can be summarised as “taking the analysis to the data”. Figure 1 shows the workflow developed in the DMP. All data files from the PRIMAVERA simulations were uploaded to the CAF and made accessible to project members, who were able to undertake their analysis of them. It has been common practise in many climate science projects to download the required data and perform analyses locally at users’ home institutes. Early analysis of the data requirements of the PRIMAVERA project, along with experience from prior projects, indicated that the downloading and local analysis of data would lead to significant technical challenges for each institute. As such, central analysis facilities such as JASMIN that provide data storage and compute are key to the exploitation of big-data projects such as PRIMAVERA and CMIP6.

The first step in the DMP workflow is the generation of data by the modelling centres; simulations are run on a variety of HPCs across Europe, and each model typically has its own proprietary output file format. As PRIMAVERA made up the majority of the European contribution to HighResMIP, it was necessary to conform to the CMIP6 data standards (Balaji et al., 2018). These standards require data to be provided in CMOR3 (Climate Model Output Rewriter) (Nadeau et al., 2019)-compliant netCDF files following the CMIP6 conventions. Where data had been generated in a proprietary format, they were post-processed to comply with these standards.

The second step in the workflow is the transfer of the data from the HPCs or post-processing systems to the CAF. A discussion on the transfer techniques used and the rates achieved is given in Sect. 6. The data are uploaded to the PRIMAVERA storage volume at the CAF.

The datasets provided are passed through a quality control process in step 3, which includes checking that the data and metadata standards have been complied with. The metadata are then extracted to store in the DMT’s database (Seddon, 2020a). After completion of the validation process, files are archived to tape and removed from the GWS to create space for other uploads. After the upload and validation of the data, users internal to the project analyse and work with them in step 4. The movement of files between tape and disk is fully automated by the DMT so that the consistency of file paths in the database is maintained.

The final steps in the workflow (steps 5 and 6 in Fig. 1) are managed by CEDA rather than the PRIMAVERA project as these relate to the use of shared facilities. Uploaded data are ingested into the CEDA archives (CEDA, 2020a) and are then available for dissemination to the global community using the CEDA Earth System Grid Federation node (CEDA, 2020b). Steps 5 and 6 do not have to be run immediately after the data have been validated. A delay before dissemination to the global community allows the project’s users to have a pe-

riod of sole access to the data and the opportunity to generate the first publications from the simulations.

4.1 CAF resources

After reviewing the DMP, PRIMAVERA was allocated 440 TiB of storage at the CAF split across five volumes. A virtual machine in the internal cloud was provided, which was given a domain name and HTTPS access allowed to it from the Internet, and the DMT software and database were installed on this server.

The 440 TiB of storage at the CAF that was allocated to the project was 2.4 % of the almost 18 PiB of project storage allocated to all 242 group workspaces at JASMIN in March 2021 (Townsend et al., 2021). In total, 100 PRIMAVERA users were 4.1 % of the total number of the CAF’s users at the same time.

It was originally estimated that around 2.4 PiB of data would be generated by the project, and therefore only a subset of data could be held on disk at once. The remaining data would have to be held on tape and moved to the GWS as required. The DMP describes how the DMT would allow data to be efficiently and reliably moved between tape and GWS and their location tracked.

4.2 Typical analysis workflow

Figure 2 shows the workflow that users working with the PRIMAVERA data at the CAF were required to follow. Users began by using the DMT’s web interface to identify what data had been uploaded. If the data they required were not available on disk, then they used the DMT’s web interface to request that a subset of their data was restored from tape to disk. The DMT sent them an email when this requested sample dataset was available on disk, and they could then work on the CAF’s interactive servers to develop and test their analysis code. After testing and validating their analysis code, they could then use the DMT to request that all of the data they required were restored from tape to disk. Once the full dataset was available on the GWS, then the analysis was run on the LOTUS batch processing cluster. When their analysis was complete, users then marked the data as finished in the DMT’s web interface, and the DMT would then delete the data from disk to create space for other data. This workflow allowed for the efficient use of disk space; PRIMAVERA generated over 1.6 PiB of data, but the upload of data, analysis and some storage of additional observations and reanalyses was able to fit into the 440 TiB of disk space that was available to the project.

5 Data Management Tool

The DMT was developed to track and control the flow of PRIMAVERA data around the CAF and to allow users to query the available data and their location. It was built upon

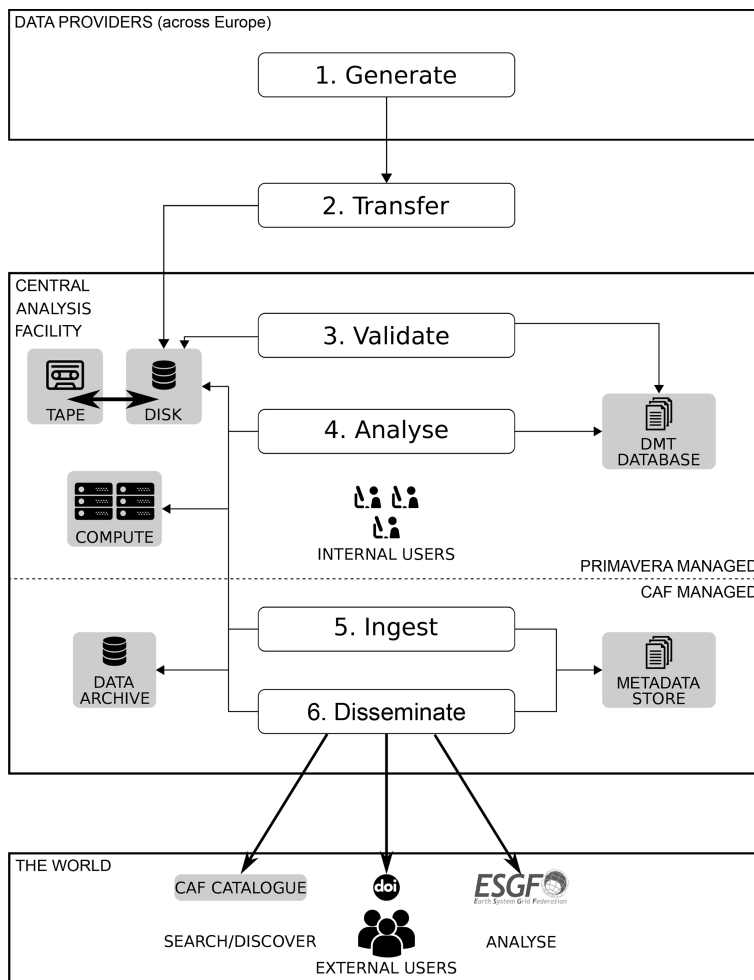


Figure 1. The workflow developed for PRIMAVERA in the Data Management Plan.




Figure 2. The workflow applied by users working with the PRIMAVERA data at the CAF.

a PostgreSQL database and custom software written in the Python programming language, using the Django web framework (Django, 2019). The database was installed on a dedicated server, along with web server software and the DMT application, and was accessible to users across all hosts at the CAF. Access to the database from the compute cluster allowed the batch submission of validation processes to the compute cluster, allowing significant parallelisation of the work.

A typical screenshot from a user querying the DMT’s web interface is shown in Fig. 3. In this example the variable

“rsut”, the top-of-atmosphere outgoing shortwave radiation, from the “highres-future” coupled future experiment is being queried. A value of “offline” in the “Online Status” column shows that the files for this simulation are currently only available on tape; a value of “partial” shows that some files are on disk, but others are only available on tape; and “online” shows that all files are available on disk. The “Request Retrieval?” column allows users to indicate that they want to work with this variable. If the variable’s files need to be restored from tape to disk, then they are queued for retrieval, and the user is emailed when the data become available. The




Figure 3. A screenshot from the DMT’s web interface showing the data available for one variable from the coupled future experiment.

DMT contains a similar page to allow users to view the data that they have requested and to mark them as used, so that they can be deleted from disk to release storage.

The DMT software is distributed under an open-source license. Development of the DMT began in April 2016, and the first data were made available via the DMT in May 2017, requiring the full time work of one developer. Development work continued as required to improve the flow of data through the system and to facilitate the publication of data to the ESGF.

5.1 DMT internal structure

Figure 4 shows a simplified UML (Unified Modelling Language) class diagram illustrating how data are represented as Django models, implemented as Python classes in the DMT. Internally, the DMT’s primary data object is a DataRequest. A DataRequest object consists of a VariableRequest object along with details of the institute, climate model, experiment and variant label. The VariableRequest object specifies a CMIP6 variable and MIP table along with additional metadata. VariableRequest objects are created programmatically from the CMIP6 HighResMIP data request (Juckes et al., 2020). During the development of the Data Management Plan, a shared spreadsheet was created from the High-ResMIP data request, and each institute edited the spreadsheet to indicate the variables that they would generate. The DataRequest objects were created programmatically from this spreadsheet. A facility to allow data providers to update

the data requests that they would provide was developed to allow for changes to providers’ plans.

A DataFile object was created for each file uploaded to the CAF during the validation process, and a checksum was calculated for each file submitted to the DMT. The checksums were used throughout the workflow as the DMT moved files from tape to disk to ensure their integrity. Each DataFile is related to a DataRequest, and all views in the DMT’s web interface are generated from these DataRequest objects. All objects are mapped to database tables using Django’s object-relational mapper.

5.2 Design considerations

An agile (Beck et al., 2001) methodology was used for the development of the DMT. Releases were regularly made available to users, and their feedback was incorporated into additional functionality and performance improvements in future releases. The initial releases of the DMT provided the functionality for the validate and analyse steps in the workflow shown in Fig. 1. Further development provided the functionality for the ingestion and dissemination steps later in the project. The choice of an agile approach was partly necessary because of the limited amount of developer resource available and so the need to concentrate the available resource on releasing functionality to the users as soon as possible.

The DMT’s users were not formally involved in the initial design of the tool. The developers used their experience of working with previous climate modelling projects to in-

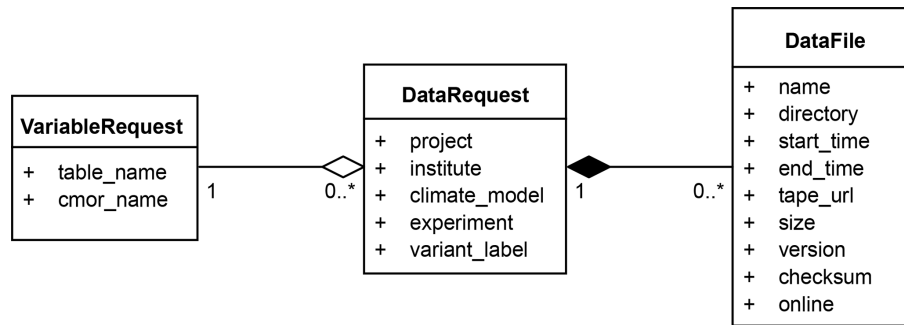


Figure 4. A simplified UML class diagram showing how data requests and data files are implemented as Django models in the DMT.

clude the search and browsing functionality that they thought would be the most useful to the users. This functionality was then refined in future releases based on informal feedback from users.

The underlying data structures used in the DMT were designed at the start of the project during the development of the DMP and were documented in there. The data structures were designed to hold enough information to uniquely identify each file and data request, and their design was based on experience gained during previous projects. The structures evolved as experience was gained during the initial DMT development. The Django framework allows the underlying data structures to be changed as a project evolves using “migrations”; this flexibility worked well with the agile methodology.

The data structures used in the DMT were based on the metadata attributes in the netCDF files specified by the CMIP6 data standards. The existing DMT software could be reused by other projects using the same data standards. The current version of the DMT is not immediately generalisable to other data standards as the search facilities rely on these CMIP6 metadata attributes. However, the DMT can be generalised to a wider range of file standards by either reducing the number of metadata attributes extracted from data files or by making the attributes extracted optional. Reduced metadata extraction would limit the usefulness of the current search fields, which may need to be removed from the search interface. Such changes would be relatively trivial. Some of the functionality to restore data from tape to disk is specific to the structure of the current CAF and would need to be generalised for use at other CAFs.

Users had to use the DMT’s web interface as no alternative was available. Several solutions are available to add a web application programming interface (API) to applications written in the Django framework. A web API should be added to future versions of the DMT to provide users with alternative methods to access the DMT. One user used web-scraping technology to extract information from the DMT to automate their own workflows, which could have been avoided had an API been available. Additionally, an API would allow command line software to interface with the

DMT, which some existing users may prefer to the web-based interface.

The DMT worked well, although some manual effort was required to maintain the set of expected data requests. The data request objects in the database that were planned to be produced by each modelling centre were created programmatically at the start of the project. As each modelling centre gained more experience with their model, and as scientists requested additional data, modelling centres generated additional output variables. Additional data request objects then had to be created manually by the system administrator. Programmatically creating data request objects when each file is validated and added to the DMT would have been more efficient than creating them at the start of the project and then having to manually handle such changes.

The upload of data, analysis and storage of additional observations was all able to fit into the 440 TiB of CAF disk storage allocated to the project. As discussed in Sect. 4.2 users marked their data as finished so that they could be deleted from disk if no other users had requested those data. All data requests were given equal priority for retrieval and deletion. Occasionally, the amount of free disk space became low, and so targeted emails were sent to high-volume users reminding them to mark any completed data as finished. If targeted emails had not worked, then the next step would have been to develop a utility that deleted the lowest-priority data requests. However, demand never exceeded the available resource, and so no prioritisation algorithm had to be developed.

6 Data transfer

Table 1 shows the sustained rates achieved when transferring data from the data providers to the CAF in mebibytes per second (MiB s^{-1}). JASMIN contains several servers in a data transfer zone, whose connection to the Internet and to the rest of the CAF has been designed for maximum data throughput. At the start of the project, it had been assumed that parallel transfer protocols such as BSCP (Hanushevsky, 2015) or GridFTP (Foster, 2011; Globus, 2019) would pro-

vide the best transfer rates, and data providers were encouraged to use these. In reality, most data providers used the techniques that they were most familiar with as long as they gave sufficiently fast transfer rates. The rates achieved depend on many factors including the file systems, the load on the servers and the network between the servers. The fastest transfer rate from a site was over 200 MiB s^{-1} , equivalent to 16.5 TiB d^{-1} . Typical rates were 30 MiB s^{-1} (2.5 TiB d^{-1}). These rates were sufficient to complete the project on time. It is hoped that with further optimisation work, a sustained minimum transfer rate of 5.0 TiB d^{-1} could be achieved.

The collection of all of PRIMAVERA's data at the CAF was only possible because of the magnitude of the transfer rates that were achieved. If data transfer rates between the HPC centres and the CAF had been insufficient, then the scientific analysis performed by PRIMAVERA would have been significantly impacted.

7 Data access and analysis

Because the DMT was used to restore data from tape to disk, it was possible to analyse the fraction of data produced by the project that was analysed. As of 20 July 2020, 31 048 data requests had been uploaded to JASMIN containing 1.5 PiB of data. A total of 6070 data requests had been restored from tape to disk at least once, which is 20 % of the data requests by number. However, these restored data requests contained 426 TiB of data, which is 27 % of the volume of data that had been uploaded (Seddon, 2020b). One reason that more data by volume than by number of data requests were requested could be because users were analysing the larger-volume variables, which tend to be variables with a higher temporal frequency.

The number of downloads of each variable was published (Seddon et al., 2020) as this information may be useful when designing the data requests in future projects. This information was used when developing the DMP for a successor project to PRIMAVERA (Savage and Wachsmann, 2023).

All of the data from the PRIMAVERA simulations were added to the CEDA archives and published to the ESGF. Because of the need to restore the data from tape, publication took over 1 year to complete.

In PRIMAVERA, users accessed the data using a variety of tools such as CDO (Schulzweida, 2022), Python, Iris (Met Office, 2010–2022), Jupyter Notebooks (Kluyver et al., 2016) and ESMValTool (Righi et al., 2020). PRIMAVERA provided some resource for the development of ESMValTool; therefore much scientific analysis was completed before ESMValTool was available. ESMValTool may allow for the more efficient analysis of multi-model datasets in future projects.

8 User feedback

A survey of all PRIMAVERA users of the CAF (approximately 100 users) was conducted, and responses were received from 24 users (Seddon et al., 2020). Users of the data were asked to rate five aspects of the process, with 1 being very difficult and 7 being very easy, and the results are shown in Table 2. The response to all of the questions was midway between neutral and very easy to use. Four users did indicate that they had various levels of difficulty accessing the data, and two users indicated some levels of difficulty analysing the data at the CAF; no additional feedback was provided explaining the difficulties that these users had encountered. The remaining 17 and 14 users respectively answered these two questions with neutral to very positive answers. These largely positive responses indicate that users had few problems performing their analyses at the CAF.

9 Future opportunities

9.1 Platforms

9.1.1 Central analysis facility

Access to the CAF has been essential to the success of PRIMAVERA. JASMIN has a finite capacity and so may not be able to support all further projects that would benefit from having access to it. CAFs similar to JASMIN providing additional capacity would be incredibly useful for collaborative projects like PRIMAVERA as they would reduce the need for all institutes participating in a project to have their own large data storage and data analysis facilities. Such facilities also reduce the volume of data that needs to be transferred, which is important as climate simulations increase in resolution and complexity and therefore increase in size.

Access to a CAF also allows scientists from some countries that do not currently have access to high-performance data analysis facilities locally to analyse large cutting-edge datasets. Some of these countries are on the frontline of climate change, and access to a CAF would allow them to further contribute to climate science. For example, the CP4-Africa (Stratton et al., 2018) convection-permitting regional climate simulations over Africa are now available at JASMIN (Senior, 2019) and have been analysed at JASMIN by scientists working from Africa. Tools to access CAFs must be robust to the potentially slow and unreliable Internet connections that may be encountered in some countries. For example, Secure Shell (Ylonen, 1996) (SSH) sessions do not retain state when they disconnect from unreliable connections (Senior et al., 2020). Further work is required to choose the most appropriate tools to allow all users to reliably access CAFs.

Table 1. The rates achieved while transferring data from the data providers to the CAF.

Transfer from	Rate achieved	Protocol used
Toulouse, France	13 MiB s ⁻¹	Four BSCP jobs in parallel, with four streams each
Hamburg, Germany	20 to 55 MiB s ⁻¹	globus-url-copy
Lecce, Italy	34 MiB s ⁻¹ average, 69 MiB s ⁻¹ peak	gridftp, with four concurrent FTP connections and eight processes in parallel
Bologna (Cineca), Italy	200 to 300 MiB s ⁻¹	Five parallel <code>rsync -av -e "ssh -c arcfour"</code>
Barcelona, Spain	13 MiB s ⁻¹	<code>rsync</code>
Exeter, UK	30 MiB s ⁻¹	Five groups of three FTP connections over dedicated link
Reading, UK	85 MiB s ⁻¹	Four parallel <code>rsync -rvz -rsh="ssh -c arcfour"</code>

Table 2. Number of responses for each score from the users of the PRIMAVERA facilities at the CAF, with 1 being very difficult and 7 being very easy.

	Querying the data request to see which variables were scheduled to be produced.	Use of the DMT to discover and query the PRIMAVERA data available at JASMIN.	Retrieving data from tape to group workspace/disk.	Accessing the data on the JASMIN group workspaces.	Analysing the data at JASMIN?
Score	1	0	1	1	1
	2	1	1	0	1
	3	0	0	0	2
	4	0	0	1	0
	5	3	1	1	0
	6	7	11	10	7
	7	4	7	6	10
No. responses		15	21	19	21
Mean		5.8	5.9	5.9	5.8

9.1.2 Public cloud

Public cloud computing technologies could provide an alternative to a facility such as JASMIN. As users have been using local Unix-based computers to analyse data for many years, the change from working locally to working remotely at JASMIN was therefore only a small change for users. The change was assisted by the CAF's comprehensive documentation (CEDA, 2020c) and the development of additional documentation and demonstration videos by the PRIMAVERA project. The CAF's mix of interactive servers and the batch processing cluster allowed for the easy development of analysis software and its running on the full multi-model datasets. Moving the analyses to the cloud would be a larger change for users, and significant development and support would be required. However, the use of public cloud technology would allow the compute to be easily scaled to the amount required at any one time and potentially provide large volumes of data storage and remote access to the data for all users.

STFC was a PRIMAVERA project partner and received GBP 141 000 of funding for the provision of the CAF and support services over the course of the project (Townsend and Bennett, 2020). STFC received an annual grant from NERC to cover the remainder of their running costs. If a project de-

ecided to use the public cloud rather than an existing facility such as JASMIN, then funding for the provision of the public cloud resources would need to be included in the project proposal. Compute in the public cloud is billed per second per processor core, and storage is billed per unit of storage; there can be a charge for moving data into and out of the storage. A typical cost for storage using Amazon Web Service's S3 storage service in September 2022 was USD 0.022 per gibibyte, resulting in a charge of USD 9912 per month for 440 TiB of storage or USD 475 791 for disk storage over the course of the 4-year project (AWS, 2022). Some discounts are available for the advanced and bulk purchase of cloud resource. Additional costs on top of this would be required for data upload and download and tape storage and compute.

Determining the level of funding required for public cloud computing in a proposal will be difficult but important to correctly determine so that the project does not overspend or run out of compute or storage before the end of the project. No record of the compute time used at the CAF was made for the PRIMAVERA project. However, data sharing outside of the project could be made easier in a public-cloud-based solution as external users could fund their own processing costs, whereas in PRIMAVERA they had to be invited onto the CAF.

Significant architectural design and software development is required before projects like PRIMAVERA can move their data storage and analysis to the public cloud.

9.2 Software

The DMT was a significant component in the success of the PRIMAVERA project and appears to be a novel tool in projects of this scale. The current implementation of the DMT for PRIMAVERA makes some assumptions about the layout of the storage allocated to PRIMAVERA at JASMIN. Applying the DMT to other projects and using it at other CAFs may require code modifications to remove these assumptions, but this implementation has demonstrated the value of such a tool. The current PRIMAVERA DMT implementation relies on the presence of CMIP6 metadata attributes in data files and may not work with other data standards. Development of a more generic version has begun in an alternative repository (Seddon, 2023), and the community has been invited to use and contribute to the new version.

In PRIMAVERA all analysis was performed at a single CAF. Multiple CAFs (perhaps called distributed analysis facilities) coupled together by improved software tools could reduce the need for data transfers in future projects. Pangeo (Pangeo, 2020) is an example of such a tool and is a collection of software packages to enable Earth science research in cloud and HPC environments. It allows data to be distributed across different storage areas, and it schedules processing on the compute attached to each storage area. Model output data would be uploaded to a storage location close to the data provider. Just the metadata would be uploaded to a catalogue such as the DMT. These tools would then run users' analysis software on the compute where each dataset is stored, and only the small volume of analysis results would need to be transferred back to the user. Such technologies would be a further extension to the DMP's aim of "taking the analysis to the data".

10 Conclusions

PRIMAVERA's Data Management Tool and access to the central analysis facility have been essential to the success of the PRIMAVERA project. The DMT and CAF have allowed over 100 researchers to collaboratively analyse a multi-model high-resolution set of climate simulations. Over 1.6 PiB of data was collected in a single location where users were able to analyse the data via both interactive analysis servers and a batch processing cluster. The Data Management Tool minimised the volume of expensive disk storage in use at any one time by allowing data to be seamlessly moved between tape and disk under the control of data users.

If a CAF had not been available for PRIMAVERA to use, then it would have been necessary for each modelling centre to make their own data available, typically on an ESGF

node or local File Transfer Protocol (FTP) server. Anyone wanting to analyse data would have needed to identify which variables were available and then download them from each modelling centre to their home institute. There would have been multiple copies of common datasets, and each institute would have required significant volumes of storage, transfer bandwidth and compute resources.

The use of the CAF and the DMT has laid down a marker for future projects, and their adoption is recommended. The DMT allowed the project to manage its storage resources to the best of its ability. The DMT's ability to allow users to move data between tape and disk minimised the use of expensive disk storage. A user survey indicated that project members found this solution straightforward to integrate into their working practises.

Access to the CAF has been essential for analysing data in the PRIMAVERA project, but continued expansion of the storage and compute of CAFs such as at JASMIN, CMCC, DKRZ and Cori is necessary to allow them to continue supporting such projects and to support additional projects in the future. In each CMIP era the data volume has increased by an order of magnitude (Balaji et al., 2018). The Earth sciences community would benefit from the development of additional CAFs, or the expansion of existing CAFs, to allow even more projects and a wider range of users to take advantage of such facilities and to handle the inevitable increase in future data volumes.

The tools developed by the PRIMAVERA project have successfully demonstrated the feasibility of such techniques. The tools have been made freely available but require some development before they can be seamlessly adopted by other projects. The inclusion in the project proposal of dedicated time resource for data management and development of the DMT contributed to the success.

Code and data availability. The PRIMAVERA Data Management Tool's source code can be found in a Zenodo repository at <https://doi.org/10.5281/zenodo.4011770> distributed under a BSD 3-Clause license (Seddon and Stephens, 2020). Additional code is similarly available in the validation tool at <https://doi.org/10.5281/zenodo.3596772> (Seddon, 2020a) and MIP tables at <https://doi.org/10.5281/zenodo.1245673> (Nadeau et al., 2018) with similar licenses. Development of a more generic version of the Data Management Tool has begun under the same license at <https://doi.org/10.5281/zenodo.8273457> (Seddon, 2023), and the community is invited to contribute there. The PRIMAVERA datasets are listed at <https://doi.org/10.5281/zenodo.10118150> (Seddon, 2020c).

Author contributions. MSM and AS developed the PRIMAVERA Data Management Plan. JS then developed and implemented the DMT, with input from AS, and managed the data upload and data availability. JS prepared the manuscript with contributions from all co-authors.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

Acknowledgements. The PRIMAVERA project is funded by the European Union's Horizon 2020 programme, grant agreement no. 641727. Matthew S. Mizielinski received additional funding from the EU Horizon 2020 CRESCENDO project, grant agreement no. 641816 and by the Newton Fund through the Met Office Climate Science for Service Partnership Brazil (CSSP Brazil). This work used JASMIN, the UK's collaborative data analysis environment (<https://jasmin.ac.uk>, last access: 26 October 2023). The IS-ENES3 project has provided funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 824084 to support continued access for the PRIMAVERA project to JASMIN. The authors would like to thank all of the STFC staff who design and maintain JASMIN, without whom this work would not have been possible. The data providers from the PRIMAVERA modelling centres should also be thanked for their hard work complying with the metadata standards and their enthusiasm to use the workflows developed in the project, along with all of PRIMAVERA's users.

Financial support. This research has been supported by the Horizon 2020 (grant nos. 641727, 641816, and 824084) and the Newton Fund (CSSP Brazil).

Review statement. This paper was edited by Xiaomeng Huang and reviewed by three anonymous referees.

References

- AWS: AWS Pricing Calculator, <https://calculator.aws/>, last access: 15 September 2022.
- Balaji, V., Taylor, K. E., Juckes, M., Lawrence, B. N., Durack, P. J., Lautenschlager, M., Blanton, C., Cinquini, L., Denvil, S., Elkington, M., Guglielmo, F., Guilyardi, E., Hassell, D., Kharin, S., Kindermann, S., Nikonov, S., Radhakrishnan, A., Stockhause, M., Weigel, T., and Williams, D.: Requirements for a global data infrastructure in support of CMIP6, *Geosci. Model Dev.*, 11, 3659–3680, <https://doi.org/10.5194/gmd-11-3659-2018>, 2018.
- Beck, K., Beedle, M., Van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., Grenning, J., Highsmith, J., Hunt, A., Jeffries, R., Kern, J., Marick, B., Martin, R. C., Mellor, S., Schwaber, K., Sutherland, J., and Thomas, D.: The agile manifesto, <https://agilemanifesto.org/> (last access: 23 August 2023), 2001.
- CEDA: Centre for Environmental Data Analysis Archive Catalogue, <https://catalogue.ceda.ac.uk/>, last access: 2 June 2020a.
- CEDA: CMIP6 CEDA ESGF Node, <https://esgf-index1.ceda.ac.uk/search/cmip6-ceda/>, last access: 2 June 2020b.
- CEDA: JASMIN help documentation, <https://help.jasmin.ac.uk/>, last access: 4 June 2020c.
- Cinquini, L., Crichton, D., Mattmann, C., Harney, J., Shipman, G., Wang, F., Ananthakrishnan, R., Miller, N., Denvil, S., Morgan, M., Pobre, Z., Bell, G. M., Doutriaux, C., Drach, R., Williams, D., Kershaw, P., Pascoe, S., Gonzalez, E., Fiore, S., and Schweitzer, R.: The Earth System Grid Federation: An open infrastructure for access to distributed geospatial data, *Future Gener. Comp. Sy.*, 36, 400–417, <https://doi.org/10.1016/j.future.2013.07.002>, 2014.
- CMCC: ECASLab, <https://ophidialab.cmcc.it/>, last access: 4 February 2021.
- Django: Django, <https://djangoproject.com/> (last access: 12 February 2020), 2019.
- DKRZ: Deutsches Klimarechenzentrum, <https://www.dkrz.de/>, last access: 4 February 2021.
- Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E.: Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization, *Geosci. Model Dev.*, 9, 1937–1958, <https://doi.org/10.5194/gmd-9-1937-2016>, 2016.
- Foster, I.: Globus Online: Accelerating and Democratizing Science through Cloud-Based Services, *IEEE Internet Comput.*, 15, 70–73, <https://doi.org/10.1109/MIC.2011.64>, 2011.
- Globus: globus-toolkit, GitHub [code], <https://github.com/globus/globus-toolkit> (last access: 4 June 2020), 2019.
- Gurvan, M., Bourdallé-Badie, R., Chanut, J., Clementi, E., Coward, A., Ethé, C., Iovino, D., Lea, D., Lévy, C., Lovato, T., Martin, N., Masson, S., Mocavero, S., Rousset, C., Storkey, D., Vancoppenolle, M., Müller, S., Nurser, G., Bell, M., and Samson, G.: NEMO ocean engine, Zenodo [code], <https://doi.org/10.5281/zenodo.3878122>, 2019.
- Haarsma, R. J., Roberts, M. J., Vidale, P. L., Senior, C. A., Bellucci, A., Bao, Q., Chang, P., Corti, S., Fučkar, N. S., Guemas, V., von Hardenberg, J., Hazeleger, W., Kodama, C., Koenigk, T., Leung, L. R., Lu, J., Luo, J.-J., Mao, J., Mizielinski, M. S., Mizuta, R., Nobre, P., Satoh, M., Scoccimarro, E., Semmler, T., Small, J., and von Storch, J.-S.: High Resolution Model Intercomparison Project (HighResMIP v1.0) for CMIP6, *Geosci. Model Dev.*, 9, 4185–4208, <https://doi.org/10.5194/gmd-9-4185-2016>, 2016.
- Hanushevsky, A.: bbcp, <https://www.slac.stanford.edu/~abh/bbcp/> (last access: 4 June 2020), 2015.
- IEEE: IEEE Standard for Prefixes for Binary Multiples, *IEEE Std 1541-2002 (R2008)*, c1–4, <https://doi.org/10.1109/IEEESTD.2009.5254933>, 2009.
- Juckes, M., Taylor, K. E., Durack, P. J., Lawrence, B., Mizielinski, M. S., Pamment, A., Peterschmitt, J.-Y., Rixen, M., and Sénési, S.: The CMIP6 Data Request (DREQ, version 01.00.31), *Geosci. Model Dev.*, 13, 201–224, <https://doi.org/10.5194/gmd-13-201-2020>, 2020.
- Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C.: Jupyter Notebooks – a publishing format for reproducible computational workflows, in: *Positioning and Power in Academic Publish-*

- ing: Players, Agents and Agendas, edited by: Loizides, F. and Schmidt, B., IOS Press, 87–90, <https://doi.org/10.3233/978-1-61499-649-1-87>, 2016.
- Lawrence, B. N., Bennett, V. L., Churchill, J., Juckes, M., Ker-shaw, P., Pascoe, S., Pepler, S., Pritchard, M., and Stephens, A.: Storing and manipulating environmental big data with JASMIN, in: 2013 IEEE International Conference on Big Data, Silicon Valley, CA, USA, 6–9 October 2013, 68–75, <https://doi.org/10.1109/BigData.2013.6691556>, 2013.
- LBNL: Lawrence Berkeley National Laboratory Cori supercom-puter, <https://www.nersc.gov/systems/cori/>, last access: 5 February 2021.
- Met Office: Iris: A powerful, format-agnostic, and community driven Python package for analysing and visualising Earth science data, Zenodo [code], <https://doi.org/10.5281/zenodo.595182>, 2010–2022.
- Mizielinski, M. S., Stephens, A., van der Linden, P., Bretonnière, P.-A., Fiore, S., von Hardenberg, J., Kolax, M., Lohmann, K., Moine, M.-P., Le Sager, P., Semmler, T., and Senan, R.: PRIMAVERA Deliverable D9.1 Data Management Plan, Zenodo, <https://doi.org/10.5281/zenodo.3598390>, 2016.
- Nadeau, D., Seddon, J., Vegas-Regidor, J., Kettleborough, J., and Hogan, E.: PRIMAVERA-H2020/cmip6-cmor-tables: Version 01.00.23, Zenodo [code], <https://doi.org/10.5281/zenodo.1245673>, 2018.
- Nadeau, D., Doutriaux, C., Mauzey1, Hogan, E., Kettleborough, J., Kjoti, TobiasWeigel, Durack, P. J., Nicholls, Z., Jmrgonza, Wach-sylon, Taylor13, Seddon, J., and Betts, E.: PCMDI/cmor: 3.5.0, Zenodo [code], <https://doi.org/10.5281/zenodo.3355583>, 2019.
- Pangeo: Pangeo project web page, <https://pangeo.io/>, last access: 2 June 2020.
- Petrie, R., Denvil, S., Ames, S., Levavasseur, G., Fiore, S., Allen, C., Antonio, F., Berger, K., Bretonnière, P.-A., Cinquini, L., Dart, E., Dwarakanath, P., Druken, K., Evans, B., Franchistéguy, L., Gardoll, S., Gerbier, E., Greenslade, M., Hassell, D., Iwi, A., Juckes, M., Kindermann, S., Lacinski, L., Mirto, M., Nasser, A. B., Nassisi, P., Nienhouse, E., Nikonov, S., Nuzzo, A., Richards, C., Ridzwan, S., Rixen, M., Serradell, K., Snow, K., Stephens, A., Stockhause, M., Vahlenkamp, H., and Wagner, R.: Coordinating an operational data distribution network for CMIP6 data, *Geosci. Model Dev.*, 14, 629–644, <https://doi.org/10.5194/gmd-14-629-2021>, 2021.
- PRIMAVERA and the European Commission: Grant Agree-ment number: 641727 – PProcess-based climate sIMulation: AdVances in high resolution modelling and Euro-pean climate Risk Assessment (PRIMAVERA), Zenodo, <https://doi.org/10.5281/zenodo.3874429>, 2015.
- Righi, M., Andela, B., Eyring, V., Lauer, A., Predoi, V., Schlund, M., Vegas-Regidor, J., Bock, L., Brötz, B., de Mora, L., Di-blen, F., Dreyer, L., Drost, N., Earnshaw, P., Hassler, B., Koldunov, N., Little, B., Loosveldt Tomas, S., and Zimmer-mann, K.: Earth System Model Evaluation Tool (ESMValTool) v2.0 – technical overview, *Geosci. Model Dev.*, 13, 1179–1199, <https://doi.org/10.5194/gmd-13-1179-2020>, 2020.
- Roberts, M. J., Camp, J., Seddon, J., Vidale, P. L., Hodges, K., Van-niere, B., Mecking, J., Haarsma, R., Bellucci, A., Scoccimarro, E., Caron, L. P., Chauvin, F., Terray, L., Valcke, S., Moine, M. P., Putrasahan, D., Roberts, C., Senan, R., Zarzycki, C., and Ull-rich, P.: Impact of model resolution on tropical cyclone simu-lation using the HighResMIP-PRIMAVERA multimodel ensem-ble, *J. Climate*, 33, 2557–2583, <https://doi.org/10.1175/JCLI-D-19-0639.1>, 2020.
- Savage, N. and Wachsmann, F.: EERIE Data Management Plan, Zenodo, <https://doi.org/10.5281/zenodo.8304509>, 2023.
- Schiemann, R., Athanasiadis, P., Barriopedro, D., Doblus-Reyes, F., Lohmann, K., Roberts, M. J., Sein, D. V., Roberts, C. D., Terray, L., and Vidale, P. L.: Northern Hemisphere blocking simulation in current climate models: evaluating progress from the Climate Model Intercomparison Project Phase 5 to 6 and sensitivity to resolution, *Weather Clim. Dynam.*, 1, 277–292, <https://doi.org/10.5194/wcd-1-277-2020>, 2020.
- Schulzweida, U.: CDO User Guide, Zenodo, <https://doi.org/10.5281/zenodo.7112925>, 2022.
- Seddon, J.: PRIMAVERA-H2020/primavera-val: Initial release, Zenodo [code], <https://doi.org/10.5281/zenodo.3596772>, 2020a.
- Seddon, J.: PRIMAVERA-H2020/stream2-planning: Data Request Summary, Zenodo [data set], <https://doi.org/10.5281/zenodo.3921887>, 2020b.
- Seddon, J.: Deliverable D9.5 Publication of the PRIMAVERA Stream 2 Data Set, Zenodo [data set], <https://doi.org/10.5281/zenodo.10118150>, 2020c.
- Seddon, J.: MetOffice/primavera-dmt, Zenodo [code], <https://doi.org/10.5281/zenodo.8273457>, 2023.
- Seddon, J. and Stephens, A.: PRIMAVERA Data Management Tool V2.10, Zenodo [code], <https://doi.org/10.5281/zenodo.4011770>, 2020.
- Seddon, J., Mizielinski, M. S., Roberts, M., Stephens, A., Hege-wald, J., Semmler, T., Bretonnière, P.-A., Caron, L.-P., D’Anca, A., Fiore, S., Moine, M.-P., Roberts, C., Senan, R., and Fladrich, U.: Deliverable D9.6 Review of DMP and lessons learnt for fu-ture projects, Zenodo, <https://doi.org/10.5281/zenodo.3961932>, 2020.
- Senior, C., Finney, D., Owiti, Z., Rowell, D., Marsham, J., Jack-son, L., Berthou, S., Kendon, E., and Misiani, H.: Techni-cal guidelines for using CP4-Africa simulation data, Zenodo, <https://doi.org/10.5281/zenodo.4316467>, 2020.
- Senior, C. A.: P25-Present: Present-day 25 km regional pan-Africa data, CEDA Archive [data set], <https://catalogue.ceda.ac.uk/uuid/4e362effa16146abbe45c2c58f1e54ed> (last access: 4 June 2020), 2019.
- Stratton, R. A., Senior, C. A., Vosper, S. B., Folwell, S. S., Boutle, I. A., Earnshaw, P. D., Kendon, E., Lock, A. P., Mal-colm, A., Manners, J., Morcrette, C. J., Short, C., Stirling, A. J., Taylor, C. M., Tucker, S., Webster, S., and Wilkinson, J. M.: A Pan-African Convection-Permitting Regional Climate Simulation with the Met Office Unified Model: CP4-Africa, *J. Climate*, 31, 3485–3508, <https://doi.org/10.1175/JCLI-D-17-0503.1>, 2018.
- Taylor, K. E., Juckes, M., Balaji, V., Cinquini, L., Denvil, S., Du-rack, P. J., Elkington, M., Guilyardi, E., Kharin, S., Lauten-schlager, M., Lawrence, B., Nadeau, D., and Stockhause, M.: CMIP6 Global Attributes, DRS, Filenames, Directory Structure, and CV’s, <http://goo.gl/v1drZI> (last access: 26 October 2023), 2018.
- Townsend, P. and Bennett, V.: CEDA Annual Report 2019–2020, <http://cedadocs.ceda.ac.uk/1489/> (last access: 15 Septem-ber 2022), 2020.

Townsend, P., Bennett, V., Jukes, M., Parton, G., Petrie, R., Stephens, A., Kershaw, P., Pepler, S., Jones, M., Pascoe, C., Gray, H., Tucker, W., Harwood, A., Williamson, E., Donegan, S., Garland, W., and Smith, R.: CEDA Annual Report 2020–2021, Zenodo, <https://doi.org/10.5281/zenodo.5592466>, 2021.

Ylonen, T.: SSH – Secure Login Connections Over the Internet, in: 6th USENIX Security Symposium (USENIX Security 96), USENIX Association, San Jose, CA, 22–25 July 1996, <https://www.usenix.org/conference/6th-usenix-security-symposium/ssh-secure-login-connections-over-internet> (last access: 23 August 2023), 1996.