Enabling high-performance cloud computing for Earth science modeling on over a thousand cores: application to the GEOS-Chem atmospheric chemistry model

Jiawei Zhuang[1], Daniel J. Jacob[1], Haipeng Lin[1], Elizabeth W. Lundgren[1], Robert M. Yantosca[1], Judit Flo Gaya[1], Melissa P. Sulprizio[1], Sebastian D. Eastham[2], Kevin Jorissen[3]

[1] School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA.

[2] Laboratory for Aviation and the Environment Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[3] Amazon Web Services, Seattle, WA 98109, USA

Corresponding author: Jiawei Zhuang (jiaweizhuang@g.harvard.edu)

Key Points:

- Cloud computing platforms provide immediate access to Earth science models and large datasets

- Recent advances in cloud network performance enable efficient model simulations over a thousand cores

- Performance and cost of cloud are now comparable to local supercomputing cluster

Abstract

Cloud computing platforms can facilitate the use of Earth science models by providing immediate access to fully configured software, massive computing power, and large input datasets. However, slow inter-node communication performance has previously discouraged the use of cloud platforms for massively parallel simulations. Here we show that recent advances in the network performance on the Amazon Web Services (AWS) cloud enable efficient model simulations with over a thousand cores. The choices of Message Passing Interface (MPI) library configuration and inter-node communication protocol are critical to this success. Application to the GEOS-Chem model of atmospheric chemistry at global 50 km horizontal resolution shows efficient scaling up to at least 1152 cores, with performance and cost comparable to the NASA Pleiades supercomputing cluster.

Plain Language Summary

Earth science model simulations are computationally expensive, typically requiring the use of high-end supercomputing clusters that are managed by universities or national laboratories. Commercial cloud computing offers an alternative. However, past work found that cloud computing platforms were not efficient for large-scale simulations over 100 CPU cores, because the network communication performance on the cloud was slow compared to local clusters. We show that recent advances in the cloud network performance enable efficient model simulations with over a thousand cores, and cloud platforms can now serve as a viable alternative to local clusters for simulations at large scale. Computing on the cloud has extensive advantages, such as providing immediate access to fully-configured model code and large datasets for any users, allowing full reproducibility of model simulation results, offering quick access to novel hardware that might not be available on local clusters, and being able to scale to virtually unlimited amounts of compute and storage resources. Those benefits will help advance Earth science modeling research.

1 Introduction

Cloud computing platforms can provide scientists immediate access to complex Earth science models and large datasets, thus greatly facilitating scientific research and collaboration (Zhuang et al., 2019). Past work has ported a number of Earth science models to the cloud (Vance et al., 2016), including the Weather Research and Forecast Model (WRF) (Molthan et al., 2015), the Community Earth System Model (CESM) (Chen et al., 2017), the Model for Prediction Across Scales - Ocean (MPAS-O) (Coffrin et al., 2019), and the GEOS-Chem model of atmospheric chemistry (Zhuang et al., 2019). However, applications have been mainly limited to a single compute node or a small cluster with less than 100 cores, because high-performance computing (HPC) applications with intensive inter-node communication did not scale well to a large cluster on the cloud (Chang et al., 2018; Jackson et al., 2010; Mehrotra et al., 2016; Walker, 2008; Yelick et al., 2011; Zhai et al., 2011). This performance weakness has severely limited the usage of cloud computing in Earth science modeling, considering that modern Earth science models are of much greater complexity and resolution (Balaji, 2015) and routinely use thousands of cores on local HPC clusters for efficient execution. Here we show that new technology makes it possible to conduct computationally- and cost-efficient simulations with over a thousand cores on cloud platforms. We apply this new capability to the GEOS-Chem model of atmospheric chemistry (Eastham et al., 2018), and present an easy-to-follow research workflow for GEOS-Chem users

on the Amazon Web Services (AWS) cloud that can also serve as template for other Earth science models.

Past studies attributed the inefficiency of cloud platforms to two factors (Yelick et al., 2011): (1) the slow inter-node networking on the cloud, as opposed to the dedicated high-speed networking on local HPC clusters; and (2) the performance overhead of the "virtualization" technology used by the cloud (which runs multiple virtual machines on the same physical hardware for better resource utilization), as opposed to directly running programs on the physical, "bare-metal" servers. However, the performance of cloud platforms has been improving rapidly over the past several years, and early conclusions about inefficient HPC applications on the cloud may no longer hold. For example, the Microsoft Azure cloud now offers dedicated HPC networking and achieves similar scalability as traditional supercomputers (Mohammadi & Bazhirov, 2018). The AWS cloud has recently introduced (1) the "Nitro Hypervisor" for virtual machine management with extremely low overhead and thus "near-bare-metal" performance (Gregg, 2017); (2) a new instance type "C5n" with 100 Gb/s bandwidth, 4× higher than the previous "C5" type (Amazon, 2018e); and (3) a low-latency network interface called Elastic Fabric Adapter (EFA) to improve the inter-node communication performance for Message Passing Interface (MPI) libraries (Amazon, 2019e). Due to these and other improvements, it is time to revisit the suitability of cloud computing platforms for HPC applications.

Here we demonstrate the high performance and scalability of GEOS-Chem on the AWS cloud by conducting global simulations at 50-km resolution using up to 1152 CPU cores on 32 compute nodes. GEOS-Chem (http://www.geos-chem.org) simulates the global 3-D evolution of over 200 coupled chemical species using NASA Goddard Earth Observing System (GEOS) assimilated meteorological data as input (Bey et al., 2001). It is used by a large international community for a wide range of atmospheric chemistry applications. We have previously developed a user-ready capability for running single-node GEOS-Chem simulations on the AWS cloud (Zhuang et al., 2019), including immediate access to 30 Terabytes (TB) of GEOS meteorological data and 4 TB of emission data (Keller et al., 2014) through the AWS cloud storage. Here we push this development further to high-performance, multi-node applications. This is to our knowledge the first demonstration that massively parallel Earth science model simulations on the cloud can achieve similar computational and cost efficiencies as local HPC clusters.

Much of this paper is devoted to the general principles for efficient HPC on cloud platforms, applicable to any compute-intensive Earth science models. Earth science models tend to share similar workflows, software requirements, and parallel computation patterns. We present a complete, easy-to-follow research workflow in an HPC cluster environment on the AWS cloud, and also show how to identify and resolve performance bottlenecks using micro-benchmarks and MPI profiling. Detailed online instructions are available for reproducing this work (See "Acknowledgments, Samples, and Data").

2 Benefits of cloud computing for Earth science research

Basic benefits of cloud computing for Earth science modeling are reviewed by Zhuang et al, (2019), including (1) immediate access to complex models for novice and occasional users; (2) sharing of model configurations and datasets for collaborative work; (3) offering a cost-effective

alternative to the acquisition of a local cluster. Here we review additional benefits that have fueled increasing recent interest in cloud computing.

Earth science big data analytics. Large volumes of Earth data are now shared through cloud platforms, allowing users to access the data immediately inside the cloud instead of enduring slow data transfer to local computers (Ansari et al., 2017; Sudmanns et al., 2019; Yang et al., 2017). The data can often be accessed freely through "open data programs" offered by commercial cloud vendors, such as the AWS Open Data Program (https://aws.amazon.com/opendata/), the Azure Open Datasets (https://azure.microsoft.com/en-us/services/open-datasets/), and the Google Cloud Public Datasets (https://cloud.google.com/public-datasets/). For example, 70 TB of the Community Earth System Model (CESM) Large Ensemble dataset (Kay et al., 2015) has recently been made available via the AWS Open Data Program (de La Beaujardière et al., 2019). 100 TB of the Model Intercomparison Project Phase 6 (CMIP6) data are now available in the Google Public Dataset (Google, 2019). The NASA Earth Observing System Data and Information System (EOSDIS) plans to move PetaBytes (PBs) of Earth observation data to the AWS cloud storage (Behnke et al., 2019; Lynnes et al., 2017; Ramachandran et al., 2018). The cloud further offers massive computing power to analyze big data. For example, the Pangeo big data framework (Robinson et al., 2019) provides a user-friendly way to process massive geoscientific data (in NetCDF or other formats) on cloud platforms or local HPC clusters, leveraging open-source Python libraries like Xarray (Hoyer & Hamman, 2017), Dask (Rocklin, 2015), and Jupyter notebooks (Perkel, 2018).

Enabling reproducible research. Access to code and data is increasingly required as part of a scientific publication to foster reproducible research (Irving, 2016; NAS, 2019). However, it may be impractical to reproduce a research project even if its source code and data are published online, due to the difficulty to ensure the exact same software environment on different machines, and the generally-slow data transfer rate from one institution to another. Public cloud platforms can guarantee a consistent environment and the same data source for different users, leading to a high degree of reproducibility (Howe, 2012; de Oliveira et al., 2017; Perkel, 2019). An example is MyBinder (Jupyter et al., 2018, https://mybinder.org), a free, online, cloud-based service for reproducible execution of Jupyter notebooks.

Continuous Integration for scientific code. Online cloud-based Continuous Integration (CI) services such as TravisCI (https://travis-ci.com/) and Azure pipelines (Microsoft, 2019) automatically build the software source code and run the unit tests at every code commit, and report any potential errors during the build or test stage. CI is a standard practice in software engineering (Duvall et al., 2007), and is also useful for scientific data analysis (Wessel et al., 2019) and HPC application development (Sampedro et al., 2018). Without CI, software developers would need to spend long time manually finding and tracking software issues. CI services can also run in parallel with multiple software environments using the "build matrix" feature (e.g. https://docs.travis-ci.com/user/build-matrix/); this is particularly useful for ensuring that a complicated HPC codebase works with a wide variety of library versions.

Providing special hardware like GPUs (graphical processing units). Next-generation weather and climate models are likely to be based on GPUs or other special hardware instead of general-purpose CPUs (Lawrence et al., 2018). However, configuring a local GPU environment can be a

challenge for scientific users. The cloud offers high-end, HPC-oriented GPU types with pre-configured GPU libraries (Amazon, 2017a), allowing users to prototype and test GPU code without upfront investment. Other special hardware types such as FPGAs (field programmable gate arrays) (Düben et al., 2015) and TPUs (tensor processing units) (Jouppi et al., 2017) are also available via cloud services (Amazon, 2017b; Google, 2018a).

Accelerating "embarrassingly parallel" computations. Certain problems like sensitivity study and ensemble forecasting require a large number of independent model simulations that can run on individual compute nodes ("embarrassingly parallel"). With the large resource pool on public cloud platforms, independent jobs can be executed simultaneously and finish much faster (Monajemi et al., 2019). For example, the AWS cloud has provided 40,000 compute nodes for one industrial HPC use case (Amazon, 2019g). The NSF IceCube Experiment utilizes 51,500 GPUs across three cloud vendors (Amazon, 2019b) via the HTCondor software tool (Thain et al., 2005). As long as the parallelization efficiency is near 100% (no overhead from inter-node communication), users can launch any numbers of compute nodes without increasing the total CPU time and total cost.

Machine learning enhanced modeling. A recent research trend is to incorporate big data and machine learning into Earth science models (Reichstein et al., 2019; Schneider et al., 2017), by having coarse-resolution global models continuously learn from observation data or high-resolution local simulations. Machine learning workflows can require massive resources for training data generation, data preprocessing, and machine learning model training. Cloud platforms provide abundant data storage and compute resources for such workflows, including the readily-available Earth science datasets as potential training data, and GPUs and TPUs for fast training of neural networks.

"Cloud bursting" for local clusters. Local clusters can be connected to cloud platforms to handle temporary surges in user demand or to provide extra hardware types like GPUs. Such usage is called "bursting into the cloud" and is supported by popular job schedulers like Slurm (Amazon, 2019d). For example, the NASA Pleiades supercomputer is now interfaced with the AWS cloud, allowing Pleiades users to submit jobs to AWS compute services and upload files to AWS cloud storage (NASA HECC, 2019). Such cloud bursting capability can also simplify the funding and cost management issues for commercial cloud resources, as the cloud charges can be managed by the local cluster administrators and potentially be absorbed by research grant allocations.

3 High-performance computing workflow on the AWS cloud

Here we describe the research flow in an HPC cluster environment on the AWS cloud. Similar concepts and services apply to other cloud platforms such as Microsoft Azure and Google Cloud (Google, 2018b; Microsoft, 2018). We refer the reader to a step-by-step online tutorial (Zhuang, 2019a) to supplement the conceptual overview presented here.

(a) Single-node workflow

(b) Multi-node cluster architecture and workflow
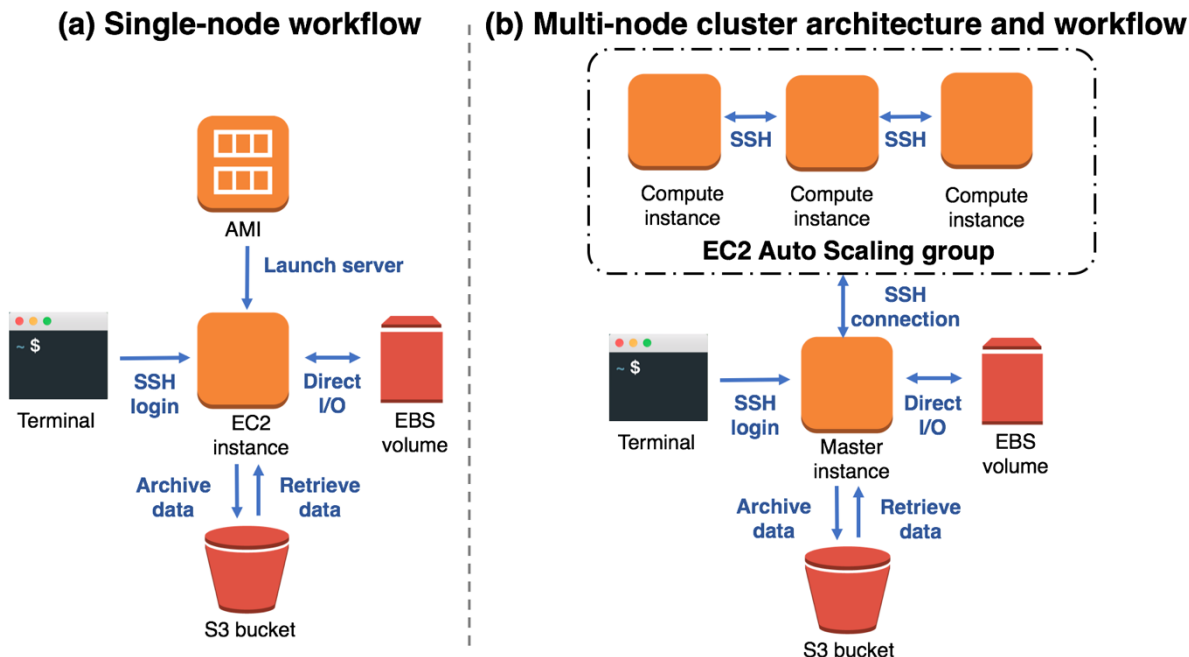
Figures 1. Single-node and multi-node workflows on the AWS cloud. The single-node workflow (panel a) is a simplified version of Fig. 2 in Zhuang et al. (2019), where the Amazon Machine Image (AMI) contains the software environment for the virtual machine (EC2 instance). The multi-node workflow (panel b) adds the auto-scaling group for compute nodes. The software environment for the master node and all compute nodes is still created from an AMI, but the creation of nodes is generally handled by a high-level framework, not manually by the user, so the AMI is not shown in the workflow (see main text in Section 3.2). The master node runs a Network File System (NFS) server to allow all compute nodes to access the EBS volume. If a Lustre parallel file system is further enabled (not shown in the figure), it will give direct I/O access to all compute nodes and avoid any I/O bottlenecks from the NFS.

3.1 Single-node workflow

For novice users, it is important to first get familiar with the single-node workflow (Zhuang et al. 2019), which is often sufficient for early-stage model testing and data analysis tasks. Fig. 1(a) shows the single-node workflow. It involves two core services on AWS: Elastic Compute Cloud (EC2) for computation and Simple Storage Service (S3) for data storage. The user launches a virtual machine (an "EC2 instance" in AWS terminology) to perform any computation tasks (e.g., perform a GEOS-Chem simulation, or run a data analysis script), and terminates the instance after finishing the computation. The EC2 instance behaves like a normal Linux server that users can log-into via Secure Shell (SSH) on their terminal. The hardware aspects (CPUs, memory, network speed) of the instance are defined by the "Instance Type" (Amazon, 2018a), and the software aspects (operating system, pre-installed software) are defined by an Amazon Machine Image (AMI). The EC2 instance uses Elastic Block Store (EBS) volumes as temporary disks to perform direct I/O during computation. S3 is used for persistent data storage.

3.2 Multi-node workflow

A multi-node cluster on the AWS cloud contains an ensemble of EC2 instances (each is a compute node), connected via SSH, allowing an MPI program to run across multiple nodes. A special master node handles administration tasks. It is also useful to install a job scheduler such as Slurm (Yoo et al., 2003) that is commonly used on local HPC clusters. The scheduler serves three purposes: (1) correctly map MPI processes (also called MPI ranks) to different compute nodes and cores; (2) manage the initialization and finalization of each simulation job; and (3) potentially integrate with the Auto Scaling capability (Amazon, 2018c) to request or terminate compute nodes as needed and thus avoid the cost of idle compute resources. Unlike job schedulers on local shared HPC clusters, the scheduler here is generally not used for resource sharing between users, because any user can create a new cluster for exclusive use.

Besides compute facilities, one needs a shared file system to host data files, software libraries, and model executables that can be accessed from all nodes. A simple choice is the Network File System (NFS) that runs on the master node and serves data to all compute nodes. The NFS architecture may cause an I/O bottleneck, and I/O-intensive models can benefit from a high-performance parallel file system such as Lustre (Schwan, 2003). The Amazon FSx for Lustre service (Amazon, 2018b) provides a fully-managed Lustre file system that can be mounted to multiple compute nodes.

Fig. 1(b) illustrates the basic form of a multi-node cluster, using a single EBS volume as a shared disk for all nodes, without a parallel filesystem. The software environment for the master node and all compute nodes is still created from an AMI, but the launch of nodes is generally handled by a high-level framework (see Section 3.3), not manually by the user. The user logs into the master node via SSH, and submits jobs to compute nodes via a job scheduler. The compute nodes form an Auto Scaling group (Amazon, 2018c) that automatically adjusts the number of nodes based on the jobs in the scheduler queue. After finishing the computation, the user archives important data to S3 storage and then terminates the entire cluster.

Building a multi-node HPC cluster environment on the clouds used to be difficult, prompting developments of "HPC-as-a-Service" (Baun et al., 2011; Wong and Goscinski, 2013; Huang, 2014; Church et al., 2015; Section 2.3 of Netto et al., 2017), where scientists can access a pre-configured HPC environment, often through a graphical web portal (Calegari et al., 2019), with no need to understand the underlying infrastructure. However, such black-box service has several drawbacks for research computing: (1) the available software libraries and applications are determined by the HPC service provider, and are not easy to extend to custom research code; (2) continuous maintenance by the service provider is required to keep software and hardware up-to-date. Recent developments in software tools have made it much easier for scientists to create their own HPC environment on the cloud for custom research needs, without requiring specific support from system administrators. On the AWS cloud, the user can now quickly launch a multi-node HPC cluster using the AWS ParallelCluster framework (https://github.com/aws/aws-parallelcluster), and then easily install the desired software libraries (e.g. specific versions of compiler, MPI, and NetCDF) using the Spack HPC package manager (Gamblin et al., 2015, https://github.com/spack/spack). Appendix A reviews other approaches and tools for deploying HPC cluster infrastructure on the cloud. Appendix B shows the advantages of Spack compared to other software installation methods. If desired, it is then possible to deploy HPC-as-a-Service for

other users on top of such a custom HPC environment, using auxiliary AWS functionalities (Amazon, 2019f, 2019a, 2019c) and multi-user serving frameworks like JupyterHub (Glick & MacHe, 2019; Milligan, 2018; Prout et al., 2017; Sarajlic et al., 2018); this can be useful for classes and workshops where the customization of software environment is not important.

4 Benchmarking inter-node communication performance

It is common practice to conduct micro-benchmarks to test individual hardware components of a cluster architecture, such as peak CPU performance via Linpack or DGEMM (J. J. Dongarra, 2014; J. Dongarra & Luszczek, 2013), memory bandwidth via STREAM benchmark (McCalpin, 1995), disk I/O rate via LLNL IOR benchmark (Shan and Shalf, 2010), and network communication performance via OSU Micro-Benchmarks (OMB) (Liu et al. 2003) or Intel MPI benchmarks (https://github.com/intel/mpi-benchmarks). Compared to actual model benchmarks, micro-benchmarks provide an application-agnostic assessment of system performance, help attribute actual model performance issues to specific hardware components (e.g. very slow networks), and catch potential hardware/system problems at an early stage. For example, we were able to catch a critical performance issue regarding OpenMPI and EFA on the AWS cloud using micro-benchmarks (reported to https://github.com/aws/aws-parallelcluster/issues/1143). If we had skipped the micro-benchmarks and directly tested the actual model, such performance problem would have been much harder to identify.

Early micro-benchmark results for cloud platforms have largely become obsolete (e.g. Evangelinos and Hill, 2008; Jackson et al., 2010; Iosup et al., 2011; Sadooghi et al., 2015), due to rapid hardware innovations. Breuer et al. (2019) conducted an up-to-date, systematic study of the CPU, memory, and network performance of the latest AWS EC2 "c5n.18xlarge" instance, and concluded that the CPU performance is "within 95% of the efficiency of the bare-metal system", and the memory performance also "behaves similar to the bare-metal machine". This is in part due to the extremely low overhead of the AWS Nitro Hypervisor (Gregg, 2017). However, Breuer et al. (2019) did not test AWS EFA which should further improve network performance. Here we do not repeat the micro-benchmarks on CPU and memory, and instead only focus on inter-node communication, which is known to be the most important difference between cloud platforms and local HPC clusters.

We compare the performance of the AWS HPC cluster to that of the NASA Pleiades supercomputer, here with OMB as a generic inter-node communication benchmark and in the next section (Section 5) for application to GEOS-Chem. On AWS, we use the biggest "C5n.18xlarge" EC2 instance type (36 physical cores). We use Spack to install the Intel compiler, OpenMPI (Gabriel et al., 2004), and NetCDF-Fortran on top of the AWS ParallelCluster environment. The cluster also contains a pre-installed Intel-MPI library optimized for AWS EFA. On Pleiades, we use the latest available software modules. Table 1 summarizes the hardware and software.

Table 1. Hardware and software used for performance benchmarks

	Node hardware	Inter-node connection	Software environment

AWS EC2 C5n.18xlarge[a]	Intel Xeon Platinum 8124M, 3.0 GHz, 36 physical cores	100 Gigabit Ethernet (100 Gb/s bandwidth)	Intel compiler 2019 OpenMPI 3.1.4 or Intel MPI 2019.4 NetCDF-Fortran 4.4.4
Pleiades Haswell nodes[b]	Intel Xeon E5-2680v3, 2.5 GHz, 24 physical cores	InfiniBand FDR (two independent fabrics of 56 Gb/s bandwidth)	Intel compiler 2018 SGI-MPT 2.0 NetCDF-Fortran 4.4.0

a. Detailed at https://aws.amazon.com/blogs/aws/new-c5n-instances-with-100-gbps-networking/
b. Detailed at https://www.nas.nasa.gov/hecc/resources/pleiades.html

4.1 Network performance factors

We measure the inter-node network performance with OMB version 5.6.1, originally described by Liu et al. (2003) and available for download at http://mvapich.cse.ohio-state.edu/benchmarks/. Network performance includes two factors: bandwidth and latency (Chapter 4.5.1 of Hager and Wellein, 2010). The total time T for passing a message can be computed as:

$$T = T_l + \frac{N}{B}$$ (1)

where T_l is latency [μs], B is bandwidth [MB/s], and N is the message size [bytes]. Latency indicates the initial connection overhead and bandwidth measures the sustained data transfer rate after the connection is established. Transferring small messages is latency-limited ($T_l \gg N/B$), while transferring large messages is bandwidth-limited ($T_l \ll N/B$). The actual size of messages is application-dependent and can be measured by MPI profiling tools, detailed in Section 5.3. In some literature as well as in the OMB code, "latency" may also refer to "the total time to send messages" (T) instead of just the initial overhead (T_l). Here we strictly use the term "latency" for T_l and "time to send messages" for T.

The network latency and bandwidth are determined by both network hardware (physical fabric and network adaptor) and software (network driver and communication protocol). On the hardware side, high-end supercomputing clusters typically use InfiniBand (Grun, 2010), while low-end clusters and most cloud platforms use Ethernet (Chapter 3 of Gavrilovska, 2009). On the software side, InfiniBand-clusters can communicate via Remote Direct Memory Access (RDMA), an HPC communication technology with ultra-low overhead and sub-microsecond latency; while Ethernet clusters generally rely on the Transmission Control Protocol (TCP), a general-purpose network protocol with much higher overhead and longer latency. The technical differences between TCP and RDMA are further explained in Appendix C. TCP is known to be inefficient for HPC applications, and various non-TCP mechanisms have been developed to improve the communication performance even on Ethernet clusters (Gavrilovska, 2009), such as RDMA over Ethernet (known as RoCE, Guo et al., 2016). Although AWS EC2 does not yet support RDMA, the newly-introduced AWS EFA provides a high-performance, non-TCP communication mechanism, supporting some RDMA-like features (Appendix C). EFA uses the open-source Libfabric library (Grun et al., 2015) for interfacing between the low-level AWS

network protocols and the high-level MPI libraries. Because the change is made in the MPI library level, there is no need to modify the application code (for example, the GEOS-Chem model code in this work) in order the utilize the EFA feature. When conducting this work, we found that only the Intel-MPI library can efficiently utilize EFA; the OpenMPI library still had to use TCP for inter-node communication.
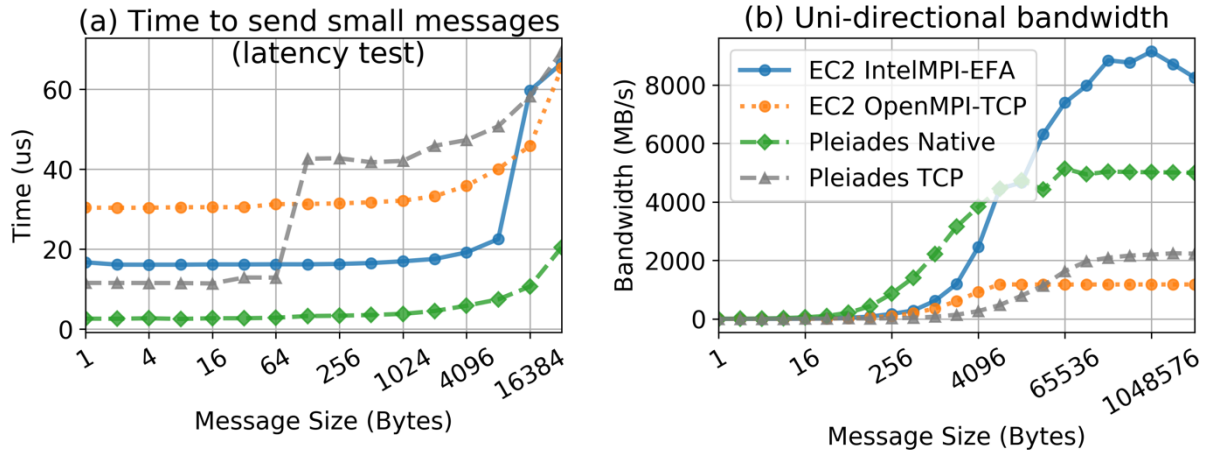
4.2 Network benchmark results

Figure 2. Performance comparison for message passing between two compute nodes. The figure shows OSU one-pair point-to-point latency and bandwidth micro-benchmarks, on the AWS EC2 c5n.18xlarge cluster and the NASA Pleiades supercomputing clusters. Latency is the asymptote for very small messages, i.e., the left limit of panel (a). The maximum bandwidth is achieved by relatively large messages, i.e., the right limit of panel (b).

Fig. 2 shows the OSU point-to-point, one-pair latency and bandwidth benchmark results (the "osu_latency" and "osu_bw" programs in OMB code). The benchmarks involve two compute nodes passing messages to each other ("point-to-point"), with only one MPI process running on each node ("one-pair"). The time or bandwidth is shown as a function of message size, following the standard practice for visualizing network performance (Liu et al. 2003). Latency is the asymptote for small messages in Fig. 2a. Between two EC2 c5n.18xlarge instances, the latency is about 30 µs with TCP and 16 µs with EFA. In contrast, the NASA Pleiades cluster with InfiniBand network has a latency of 2 µs with the native RDMA mechanism. To demonstrate the important effect of communication protocol (in addition to the network hardware), we also force TCP over InfiniBand (disabling the native RDMA capability following Section 9.1.4 of Yelick et al. 2011), and find a much higher latency of 11 µs.

For bandwidth, we focus on the maximum bandwidth achieved by relatively large messages (right limit in Fig. 2b). The EC2 c5n.18xlarge bandwidth with EFA is about 8000 MB/s, 7× higher than with TCP (1200 MB/s), and exceeds the bandwidth of NASA Pleiades InfiniBand (5000 MB/s). The bandwidth measured here is defined as "uni-directional bandwidth" as messages flow in only one direction. We also measure the "bi-directional bandwidth" (the "osu_bibw" programs in OMB code). which allows two MPI processes to simultaneously send

messages to each other (not plotted). The bi-directional bandwidth on EC2 is the same as the uni-directional value, while the bi-directional bandwidth on Pleiades doubles its uni-directional bandwidth, as expected for InfiniBand interconnect (Fig. 4 in Lockwood et al. 2014).

One would expect the 100 Gigabit Ethernet on EC2 C5n.18xlarge to deliver 12.5 GB/s bandwidth (8 Gigabit (Gb) = 1 GigaByte (GB)). However, TCP only delivers one tenth of theoretical bandwidth. This is because the data transfer with TCP is limited by the CPU processing power instead of network speed (Appendix C). The full bandwidth can be obtained by a multi-pair benchmark (the "osu_mbw_mr" program in OMB code.), which uses many MPI processes to simultaneously pass messages between two compute nodes and make use of all CPU cores (detailed instructions in Zhuang, 2019c). EFA and RDMA have much smaller CPU overhead and can achieve near-full bandwidth with only one pair of MPI processes.

Overall, the current network bandwidth and latency on AWS EC2 have improved by orders of magnitude over the old EC2 instance types with only 80 MB/s bandwidth and 80 μs latency (Evangelinos & Hill, 2008; Hill & Humphrey, 2009). Such improvement comes from the continuous updates on the AWS networking technology over the past decade, such as the support of Placement Group (Amazon, 2010), Enhanced Networking (also known as SR-IOV, Lockwood et al., 2014; Chapter 13.3.2 of Foster and Gannon, 2017), and most recently EFA.
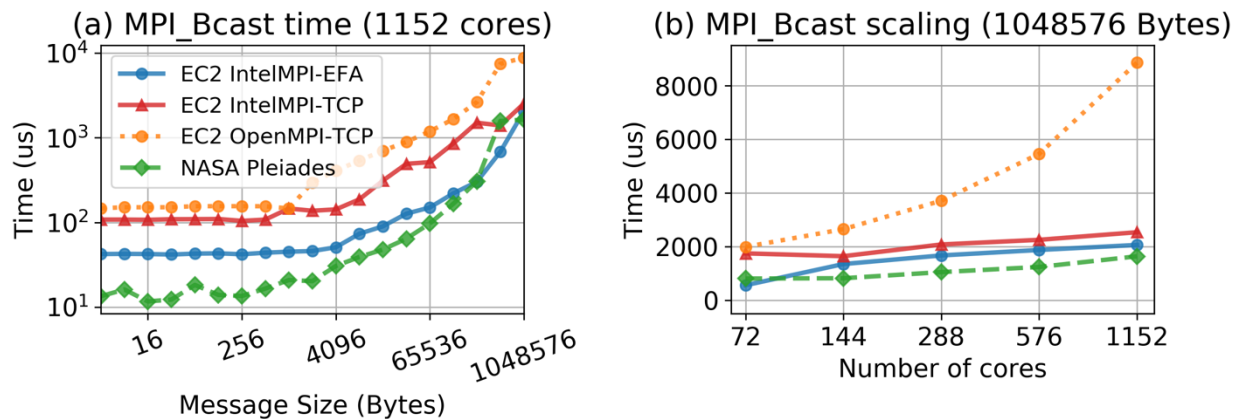

Figure 3. Performance comparison for message broadcasting across many compute nodes, and the scaling for large message sizes. The figure shows the OSU broadcast micro-benchmark on the AWS EC2 c5n.18xlarge cluster and the NASA Pleiades supercomputing cluster.

We also benchmark collective MPI functions (often called "collectives") that perform complex message transfer across many compute nodes. We find that the choice of MPI libraries and communication protocols has a major impact on the performance of MPI collectives, especially at large core counts. Fig. 3 show the performance of the "MPI_Bcast", one of the most frequently used collective functions (Table 3 of Parker et al., 2018). For example, GEOS-Chem uses "MPI_Bcast" to send input data from the master MPI process to all MPI processes (shown in Section 5.3). For small messages (left limit of Fig 3a), NASA Pleiades InfiniBand is 6× faster than EC2 with TCP (either for Intel-MPI or OpenMPI), and 3× faster than EC2 with EFA. This performance difference between EC2 and Pleiades is mainly due to the higher latency on EC2, which mostly affects small messages (Fig. 2a). For large messages (right limit of Fig 3a), EC2 with Intel-MPI (either for TCP or EFA) is as fast as Pleiades InfiniBand, while EC2 with

OpenMPI is 4× slower than Pleiades. Broadcasting with OpenMPI does not scale well with the number of cores, compared to the excellent scalability of Intel-MPI (Fig 3b). We also compare the "MPI_Allreduce" function, with a similar finding that OpenMPI is 5× slower than Intel-MPI at 1152 cores. This performance difference between OpenMPI and Intel-MPI is mainly due to their underlying collective algorithms (Chan & Heimlich, 2007; Patarasuk & Yuan, 2009; Pješivac-Grbović et al., 2007) (reported and discussed in detail at https://github.com/aws/aws-parallelcluster/issues/1436) The performance gap in MPI collectives can largely explain the performance difference of the GEOS-Chem model when built with different MPI libraries (Section 5.2).

5 Application to massively parallel GEOS-Chem simulations on the cloud

5.1 The GEOS-Chem model

We use GEOS-Chem version 12.3.2 in its high-performance implementation (GCHP, Eastham et al. 2018). GEOS-Chem simulates tropospheric-stratospheric chemistry by solving the 3-D chemical continuity equations in an Eulerian framework for over 200 chemical species. The equations include transport, chemistry, emission, and deposition terms. From a structural perspective, the model contains three main components: (1) the "column operators" (Long et al., 2015) for local or column-based computations such as chemical kinetics and convection; (2) the advection component using the GFDL-FV3 model (Putman & Lin, 2007); and (3) the I/O component using GMAO MAPL software (Suarez et al., 2007). The model parallelization is based on horizontal grid decomposition with MPI. The column operators do not involve MPI communication so have near-ideal scalability regardless of network performance. The GFDL-FV3 model is designed to achieve high scalability using the cubed-sphere grid (Chapter 5 of Putman, 2007). The MAPL I/O component in GEOS-Chem version 12.3.2 reads most data in serial, rather than parallelized across cores, which may become a bottleneck for a large number of cores (Fig. 4 in Eastham et al. 2018). More recent GEOS-Chem versions (12.5.0 and beyond) allow parallel I/O via the "GMAO_pFIO" module (ESMF, 2018).

5.2 Performance and scalability

We benchmark a 7-day global simulation at C180 cubed-sphere horizontal resolution (≈50 km) and 72 vertical layers. The model reads 150 GB of input data during the simulation and writes the global 3-D daily concentration fields of 163 chemical species resulting in 60 GB of output data (8.6 GB per day × 7 days). All data are hosted on a single throughput-optimized hard disk drive (HDD, the "st1" EBS volume type) with a theoretical throughput of 500 MB/s, mounted to all compute nodes via NFS. We do not enable the AWS FSx for Lustre service for parallel I/O, because the model version used here mostly relies on serial I/O.
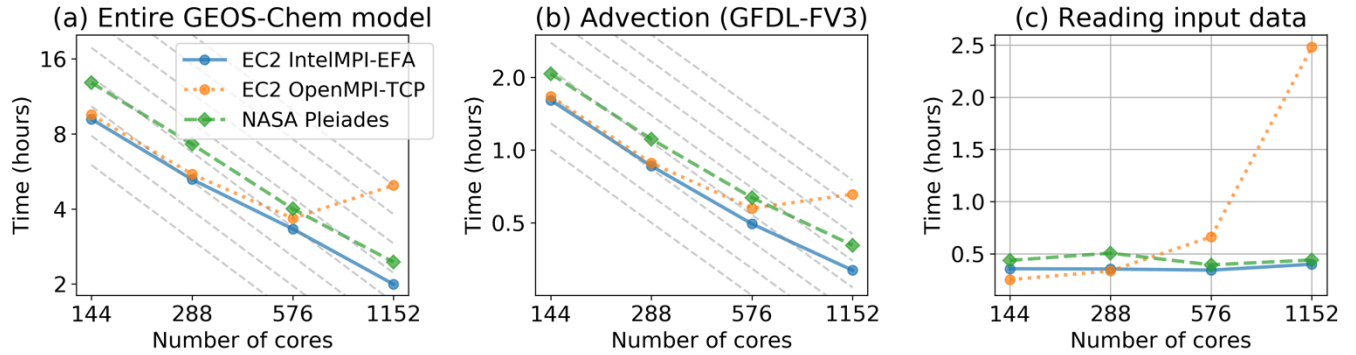
Figure 4. Performance and scalability of the GEOS-Chem model on the Amazon Web Services (AWS) cloud with up to 1152 cores. Results are for a 7-day global simulation at cubed-sphere C180 (≈50 km) horizontal resolution, and are compared to results for the same simulation on the NASA Pleiades supercomputing cluster (see hardware description in Table 1). The AWS cluster uses EC2 c5n.18xlarge instances with either Intel-MPI and EFA or OpenMPI and TCP for inter-node communication. Gray dashed lines indicate perfect scaling. Note difference in ordinate scales between panels.

Fig. 4 shows the benchmark results from 144 to 1152 cores. We achieve good performance and scalability on AWS EC2 with Intel-MPI and EFA, consistently faster than on the NASA Pleiades cluster at the same number of cores (Fig. 4a). This is in part to due to the newer processor generation (Intel Skylake) available on EC2, as opposed to the older generation on Pleiades (Intel Haswell). With OpenMPI and TCP, however, the model cannot scale beyond 576 cores, due to a major slowdown in reading input data (Fig. 4c) and a minor slowdown in advection (Fig. 4b). When the model reads input data, a single master MPI process handles most of the disk read operation and broadcasts the data to the rest of MPI processes. While the disk operation time roughly stays constant, the broadcasting takes longer with more cores, as each core runs an MPI process that has to receive a copy of data. The exact same I/O problem was also found in early results with the Community Atmosphere Model (CAM) on the AWS cloud (Section 9.1.4 of Yelick et al. 2011). Writing output data is fast compared to input, as each core can write its own domain without data exchange. The slow performance of "MPI_Bcast" with OpenMPI (Section 4.3, Fig. 3) is the major cause of the I/O bottleneck here, as further verified in the next section (Section 5.3). Importantly, we find that serial I/O does not limit performance with at least up to 1152 cores when using Intel-MPI and EFA (Fig. 4).

While our benchmark focused on the model simulation time, the job queuing time is also relevant for a real research workflow, especially for early-stage experiments that require frequent debugging and re-submission of jobs. On AWS EC2, we were able to request a thousand-core cluster in less than 5 minutes; on the NASA Pleiades cluster, the queuing time to start a thousand-core job ranged from an hour to a day, depending on the cluster utilization.

5.3 MPI Profiling

The weaker performance of OpenMPI compared to Intel-MPI is worth investigating further, because OpenMPI has otherwise the advantage of being popular and open-source (Intel-MPI is free but closed-source). To identify the exact performance bottleneck, we conduct MPI profiling using the open-source Integrated Performance Monitoring (IPM) software tool (Skinner, 2005;

Wright and Pfeiffer, 2009; https://github.com/nerscadmin/IPM). At the top level, IPM measures the communication time (i.e. all MPI function calls) versus computation time (non-MPI part of code). It further breaks down the MPI timing into different MPI function calls and message sizes. Jackson et al. (2010) used IPM to analyze the performance of several HPC applications on the AWS cloud. The IPM library can be linked at run-time and does not require recompiling the model source code. This allows us to apply IPM on GEOS-Chem simulations without modifying the model compiling procedure.

It is sufficient to use the lightweight IPM profiler here, instead of more powerful HPC profilers such as TAU (Shende & Malony, 2006) and HPCToolkit (Adhianto et al., 2010) which can capture all function calls. Here we are only interested in MPI function calls (a tiny subset of all function calls), because any potential network performance issues on AWS EC2 will slow down MPI.
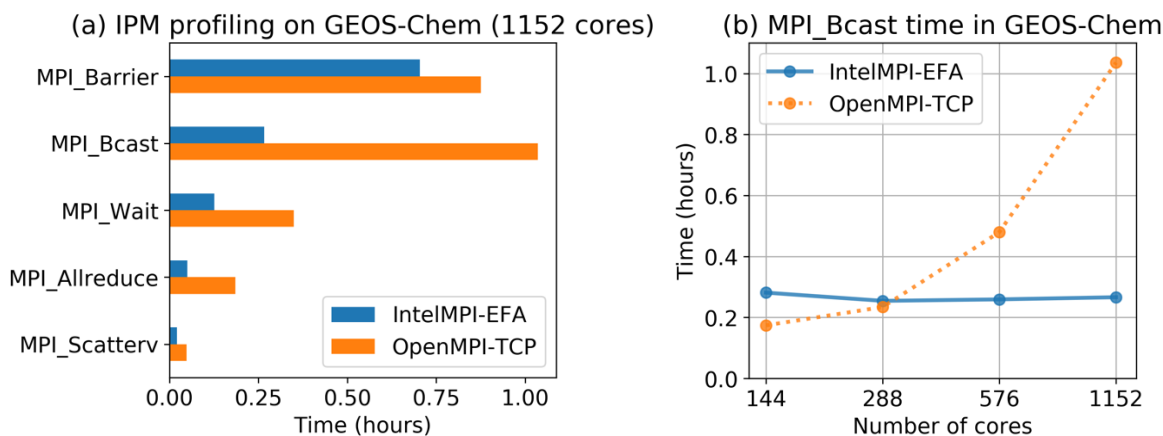

Figure 5. MPI profiling for GEOS-Chem on AWS cloud to identify performance bottlenecks. Results are from the Integrated Performance Monitoring (IPM) tool applied to a 7-day GEOS-Chem simulation presented in Section 5.2 and Fig. 4.

Fig. 5 shows the results of IPM profiling applied to the GEOS-Chem simulations presented in Section 5.2. The most time-consuming functions are "MPI_Barrier" (all MPI processes need to wait for the slowest process) and "MPI_Bcast" (after reading data, the master process broadcasts data to all processes), followed by "MPI_Wait" (waiting for asynchronous functions like "MPI_Isend" to finish) and "MPI_Allreduce" (computing the global masses of chemical species, to check mass conservation). With more cores, the time spent on "MPI_Bcast" grows rapidly when using OpenMPI, but stays constant when using Intel-MPI and EFA (Fig. 5b). This is consistent with the "MPI_Bcast" micro-benchmarks in Section 4.3 and Fig. 3. The idling time in "MPI_Barrier" is relatively insensitive to the MPI configuration and is not specific to the cloud. It reflects a load imbalance in parallelization, and further investigation shows that it is caused by slow integration of the chemical kinetics near sunrise/sunset (Appendix D).

We found MPI profiling to be a critical step in the early stage of this work for identifying and resolving the model performance bottleneck. Initially, we only built the model with OpenMPI 3.1.4, which is the standard library configuration used for the official GEOS-Chem benchmarks conducted on Harvard's local cluster (GEOS-Chem, 2019). However, this software configuration

does not work efficiently in the AWS cloud environment. By conducting profiling with IPM, we were able to narrow down the performance issue to a single "MPI_Bcast" function, and then focus on improving "MPI_Bcast" performance using OSU micro-benchmarks with different MPI configurations, leading to success when using Intel-MPI with EFA. Without such further investigation, we might have erroneously concluded that the AWS cloud is inefficient for HPC, following on the body of literature cited in the Introduction.

5.4 Cost analysis

AWS EC2 charges per computing second and per unit storage. For our application, the cost of storage is cheap compared to the cost of computing (Zhuang et al., 2019). There are several pricing options for computing, including the standard "on-demand" pricing and the cheaper "spot" pricing (about a factor of 3 cheaper). The difference between the two is that spot instances may be interrupted but this happens rarely (Amazon, 2018d). Here we used spot pricing and did not experience any interruptions. The spot price varies depending on time and region. When conducting this work in the AWS "us-east-1" region, one EC2 "c5n.18xlarge" node (36 cores) cost 1.17 USD per hour. The cost would have been 0.72 USD per hour if we were in the "us-east-2" region. To compare these costs to the NASA Pleiades cluster, we use the Standard Billing Unit (SBU) cost model to estimate true hourly cost, following the calculations in Chang et al. (2018) and Zhuang et al. (2019). One Pleiades Haswell node (24 cores) costs 0.53 USD per hour in this cost model.
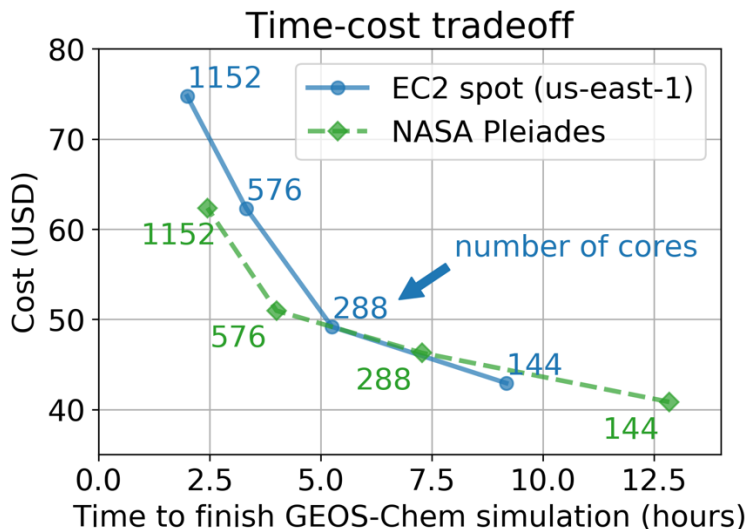

Figure 6. Cost of a 7-day global GEOS-Chem simulation of tropospheric-stratospheric chemistry at cubed-sphere C180 (≈50 km) resolution. The time required to finish the simulation (x-axis) varies with the number of cores (indicated as text next to the data point), which in turn affects the total cost because scalability is less than 100%,. Cost of the AWS EC2 cluster is compared to that of the NASA Pleiades cluster. The AWS EC2 time and cost are based on Intel-MPI and EFA performance and assume "us-east-1" spot pricing. The NASA Pleiades cost is based on the Standard Billing Unit (SBU) model.

Fig. 6 shows the cost versus the completion time of the 7-day global GEOS-Chem simulation at cubed-sphere C180 (≈50 km) resolution, for the AWS EC2 and the NASA Pleiades clusters. The

AWS EC2 time and cost are based on Intel-MPI with EFA and assume "us-east-1" spot pricing. There is a trade-off between time and cost -- using more CPU cores reduces the time to finish simulation, but increases the total CPU hours (thus the cost). Even though Fig. 4 shows a relatively successful parallelization of GEOS-Chem on both AWS EC2 and Pleiades in going from 144 to 1152 cores, there is still a 42% departure from perfect scaling with AWS EC2 and 35% with Pleiades, as indicated by the increase in total cost (perfect scaling means zero increase in total CPU hours and cost). If the turnaround time required for simulation completion can be longer than 5 hours, then AWS EC2 is cheaper than Pleiades; if it must be less than 5 hours, then Pleiades is cheaper than AWS EC2. If the amount of money available for the simulation is less than 50 USD, then AWS EC2 is the faster choice; if it is more than 50 USD, then Pleiades is the faster choice. Obviously, the trade-offs will depend on the model configuration and resolution. Overall, the cost difference between AWS EC2 and Pleiades is within 20% for the same turnaround time. Our results contrast with previous studies concluding that the AWS cloud was not cost-efficient compared to local HPC clusters (e.g. Chang et al., 2018), as these studies did not take advantage of recent advances in cloud networking.

6 Conclusions

We have shown that the Amazon Web Services (AWS) cloud can be used for massively parallel Earth science model simulations with performance comparable to local supercomputing clusters. This is enabled by recent improvements in the cloud inter-node network performance. We present a complete research workflow in a custom high-performance computing (HPC) environment on the AWS cloud, using the AWS ParallelCluster framework and the Spack HPC package manager. The workflow can be easily applied to any Earth science model. Online instructions (see Acknowledgments, Samples, and Data) and tutorials (Zhuang, 2019a) are available for reproducing the workflow. Making models readily available in an HPC cloud environment can save researchers considerable time that would otherwise be spent in configuring complicated software and downloading input data.

We demonstrated the HPC-on-cloud capability with the GEOS-Chem global 3-D model of atmospheric chemistry, and show that we can make it scale efficiently up to at least 1152 cores on the AWS cloud, with comparable performance and cost to the NASA Pleiades supercomputer. This is in sharp contrast to previous studies finding that cloud platforms were much slower and more costly than local HPC clusters. The AWS cloud has recently become much more suited for large-scale HPC applications, due to various hardware and software improvements including the new "c5n.18xlarge" EC2 instance type with much higher network bandwidth, the new Elastic Fabric Adapter (EFA) for more efficient inter-node communication, and the fast MPI collective algorithms in the latest Intel-MPI library. We demonstrate these improvements in inter-node communication performance using micro-benchmarks applicable to general HPC applications, and with MPI profiling applied to the GEOS-Chem model.

An open question is how to fund researchers for utilizing commercial cloud resources. Researchers are used to having their computing costs subsidized through their institution or covered through research grants. For this work, we were able to obtain free credits from the AWS Cloud Credit for Research Program. In the long term, however, there must be coordination between funding agencies, cloud vendors, and universities to build a more cloud-friendly funding procedure. For example, the NSF Cloud Access Program (NSF, 2019) aims to support research

and education on cloud computing platforms. For Earth science specifically, "Jupyter meets the Earth" (Pérez et al., 2019) is a recent NSF-funded project that focuses on geoscientific software development for cloud platforms and HPC centers.

Appendix A. Approaches to deploy HPC cluster infrastructure on the cloud

To deploy the cluster architecture described in Section 3.2, there are two main approaches: bottom-up and top-down. In the bottom-up approach, the user launches individual resources (EC2 instances, EBS volumes) and manually glues them together. For connecting EC2 instances, this just means copying the SSH key to each instance so they access each other. A step-by-step tutorial is given by Lockwood (2013). This is easy to do even for a novice AWS user. Other tasks, such as deploying an auto-scaling-enabled job scheduler and configuring a shared file system, involve more complicated steps. This bottom-up deployment can be done by manually clicking on the AWS web console (Appendix A of Jung et al., 2017) or be partly automated by custom scripts (Chapter 9 of Yelick et al., 2011; Chang et al., 2018) based on the AWS Command Line Interface (AWSCLI).

A better approach is top-down, or "Infrastructure as Code" (IaC) in cloud computing terminology (Morris, 2016; Chapter 1 of Brikman, 2018; Chapter 4 of Wittig and Wittig, 2018). Unlike the bottom-up approach that specifies the exact operations to build a complicated cluster infrastructure (a "procedural" approach), here the user simply describes the desired architecture in a text file (a "declarative" approach). The text file, often written in JSON or YAML format, contains the complete details of the cluster including the EC2 Instance Type for master and compute nodes, the number of compute nodes, the size and type of EBS volumes, how the volumes are mounted to the instances, the network configurations between the instances, etc. This description file is then digested by a cloud management service (CloudFormation in case of AWS; Deployment Manager in case of Azure and Google Cloud) to produce the full architecture without human intervention. A key difference from the bottom-up approach is that IaC handles inter-component dependencies much more robustly. In the bottom-up approach, one AWSCLI command typically only handles one component ("launching an EC2 instance", "creating an EBS volume", or "mounting the volume to the instance", etc.), but there can be complicated dependencies between each instruction. For example, before mounting an EBS volume to an EC2 instance, one needs to validate that both the instance and the volume are running correctly. However, executing a single EC2 launch command does not guarantee a ready-to-use instance, due to the warm-up time of the instance and even potential errors during instance creation. IaC frameworks ensure that all prerequisites are met before executing the next stage. When deleting the cluster after the model computation is done, IaC also helps resolve the dependency problem in the reverse order.

Compared to the bottom-up approach of manually assembling a cluster, the top-down IaC approach has important advantages: the cluster architecture and deployment steps are self-documented by the description file, making it easy to reproduce the same architecture and share it with other people. The cluster architecture as a text file can even be version-controlled by standard version control software like Git. The manual approach, in contrast, is vulnerable to human errors and unclear documentation. For frequent users, the manual deployment task will become repetitive and time-consuming, while the IaC approach only involves executing a single

command on the existing configuration file. The IaC approach also scales better. For example, suppose that one wishes to increase the number of compute nodes from 2 to 16. The top-down approach only involves modifying a single number in the configuration file and using a single command to update the cluster; the bottom-up approach requires manually launching 14 more instances and connecting them to the existing cluster, a complicated and error-prone process. Due to those advantages, IaC has become the standard practice in modern IT management (Morris, 2016). Thus we advise scientific users to adopt the top-down approach for production deployment. The bottom-up approach is still a useful technical exercise for novice AWS users to better understand the cluster infrastructure (Lockwood, 2013).

Although it is possible to create an HPC cluster on AWS using the standard AWS CloudFormation IaC service, scientific users should use higher-level IaC tools that are designed specifically for managing HPC clusters. Those cloud-HPC tools are built on top of general IaC frameworks like CloudFormation, but have a simpler interface and a smoother learning curve. An example is AWS ParallelCluster, an open-source HPC cluster management tool developed and supported by AWS. It is the successor of the CfnCluster tool used in previous literature (Freniere et al., 2016; Madhyastha et al., 2017). There are other third-party tools with different strengths that might be of interest to scientists. AlcesFlight (http://docs.alces-flight.com) provides a comprehensive, scientist-friendly documentation, and supports both AWS and Azure cloud; Slurm-GCP (https://github.com/SchedMD/slurm-gcp) provides an easy way to launch Slurm-enabled clusters on Google Cloud; ElasticCluster (https://github.com/elasticluster/elasticluster) is designed to support multiple cloud vendors; Ronin (https://ronin.cloud/) provides a user-friendly web portal on top of command-line-based tools like AWS ParallelCluster. StarCluster (http://star.mit.edu/cluster) was once a popular cloud-HPC tool and still used in recent literature (Chen et al., 2017), but we do not recommend it due to the stalled development and the lack of new feature support (e.g. new instance types and networking capabilities). We recommend AWS ParallelCluster for the AWS cloud environment, because it supports the latest AWS HPC features (e.g. FSx for Lustre, EFA) that are not available in other tools.

In practice, the cluster involves additional AWS concepts and services that are not shown in Fig. 1(b), including Virtual Private Cloud (VPC) and Subnet for network address configuration; EC2 Security Groups for network permission; EC2 Placement Groups for network topology configuration; Identity and Access Management (IAM) for cross-service permission; and messaging services (Amazon SNS, SQS) to monitor the cluster state. When using a high-level framework like AWS ParallelCluster, these additional components can be handled automatically.

Appendix B. Approaches to install HPC software libraries on the cloud

After the cluster infrastructure is deployed, the next step is to install software libraries such as MPI and NetCDF. Building scientific software libraries is notoriously difficult, due to complicated inter-library dependencies (Dubais et al., 2003). Traditionally, users need to choose between two extreme ways of software installation: using default package managers or building from source code. Most operating systems contain a default package manager (*apt* for Debian/Ubuntu and *yum* for RedHat/CentOS) that can install pre-built binary files ("binaries") in a few seconds and ensure correct inter-library dependencies. However, the available pre-built

binaries are limited to very few versions and compile settings, while HPC codebases with legacy software modules often require specific library versions to compile correctly (for example, the case with the GEOS model in Chang et al., 2018). Furthermore, pre-compiled binaries are built for general environments and are not optimized for a particular computer system and compiler. Therefore, HPC model developers generally have to build each library from its source code and manually handle inter-library dependencies, a slow, laborious, and error-prone process.

To solve this problem, we use the Spack supercomputing package manager (Gamblin et al., 2015) to install the required libraries on the AWS cloud. Spack was originally developed at the Lawrence Livermore National Laboratory (LLNL) and is now used in many HPC centers. To install a particular library, Spack first builds a Directed Acyclic Graph (DAG) representing the inter-library dependencies, and solves the DAG bottom-up to ensure that all dependencies are met. This DAG approach is similar to how default package managers (*apt/yum*) resolve dependencies. Different from default package managers, Spack builds libraries from source code instead of downloading pre-compiled binaries. This allows users to compile arbitrary software versions and even insert custom compile flags to optimize software performance. Compiling from source code takes relatively long time, but the build process for the entire DAG is automated and does not require human intervention. Other HPC-oriented package managers like EasyBuild (Geimer et al., 2014) and OpenHPC (Schulz et al., 2016) serve a similar purpose as Spack.

Spack provides advanced software management utilities (McLay et al., 2011) beyond library installation. It allows multiple environments to coexist and provides an easy way to switch between them (via the "spack env" command). Such utility is particularly valuable for model development that often requires testing many combinations of libraries. Spack has been chosen by US DOE's Exascale Computing Project (ECP) to deliver the entire ECP software stack (Heroux et al., 2019), which contains many cutting-edge HPC development tools that can be useful for Earth science model developers.

After the software libraries and models are correctly built on the cloud, other model users can simply reuse the environment and do not need to repeat the installation. On a single EC2 instance, the environment can be shared through an AMI (Zhuang et al., 2019). For a multi-node cloud cluster, the most straightforward way is to archive the entire Spack software directory in a persistent storage such as AWS S3, and let users pull the directory to their clusters' shared disks to replicate the same software environment. An alternative way to share pre-built models is modifying the compute node's AMI to include custom software, but this is discouraged by the AWS ParallelCluster official documentation (https://docs.aws.amazon.com/parallelcluster/latest/ug/tutorials_02_ami_customization.html), because "after you build your own AMI, you no longer receive updates or bug fixes with future releases of AWS ParallelCluster".

Containers like Docker, Singularity (Kurtzer et al., 2017) and CharlieCloud (Priedhorsky et al., 2017) are additional ways to simplify software deployment. Spack and containers serve different purposes and may be used together. Containers help the software delivery process once the software has been correctly installed, while Spack help the installation process itself. To create a container image, one can use any of the approaches reviewed above: pull pre-built binaries, build

from source manually, or use Spack. The advantage of containers is the guarantee of the exact same software environment on different systems. On the native non-container environment, in contrast, even if the user installs exactly the same software versions with Spack, there is still chance for operating-system-specific errors to occur, especially for legacy software modules that are not extensively tested on multiple platforms. Despite some pioneering examples of using containers within multi-node MPI environments (Younge et al., 2017; Zhang et al., 2017), such usage can still be challenging in practice, as stated in the Singularity container documentation (https://sylabs.io/guides/3.3/user-guide/mpi.html) that: "the MPI in the container must be compatible with the version of MPI available on the host" and "the configuration of the MPI implementation in the container must be configured for optimal use of the hardware if performance is critical". In this work, we use the native system without containers.

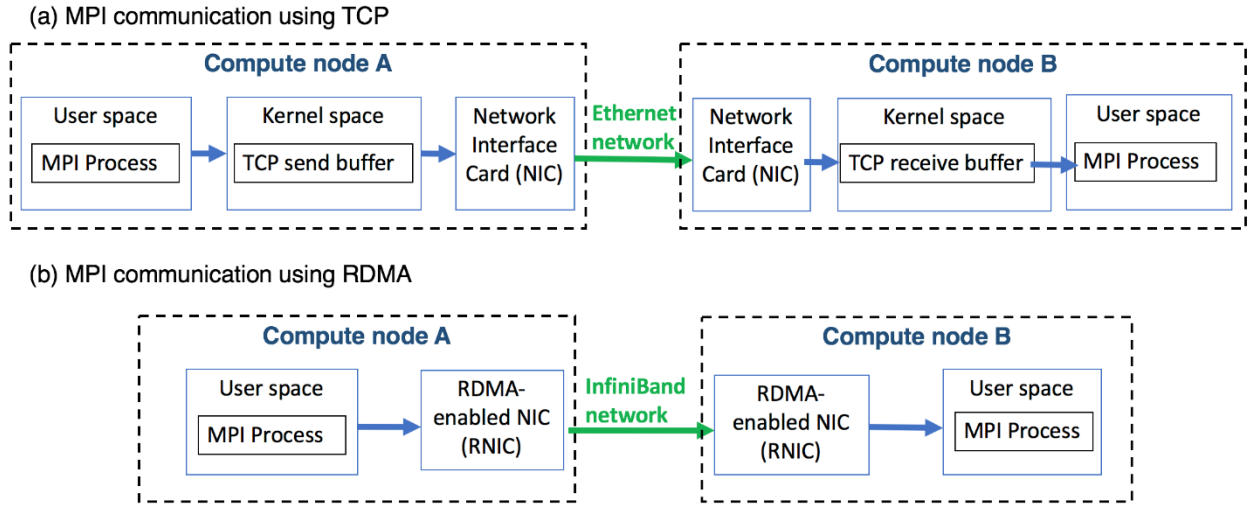## Appendix C. Details of network protocols for inter-node communication

Figure C1. Comparison between Transmission Control Protocol (TCP) and Remote Direct Memory Access (RDMA) for MPI communications.

TCP's large overhead mainly comes from excessive data copying (Chapter 7.1 of Donahoo and Calvert 2009). One would expect that sending data from one compute node to another would require a single copying operation from source memory to destination memory. Although RDMA does offer such direct-copy capability, TCP requires multiple copying operations under the hood, as illustrated in Fig. C1. First, the data residing in the user-application's memory are copied to the TCP Send Buffer residing in the kernel space. Next, the kernel sends the data over the network via the Network Interface Card (NIC). The destination node then receives the data from the network and put them into the TCP Receive Buffer. Finally, the data are copied from the Receive Buffer to the actual program's memory. In fact, even more copying operations can happen inside the network protocol stack (Mellanox Technologies, 2014). Those copying operations degrade the performance of MPI programs in three different ways (Chapter 2 and 3 of Gavrilovska 2009):
(1) Copying from one memory location to another is limited by memory bandwidth. In early days, the memory bandwidth was much faster than network bandwidth, so such extra memory

copying did not incur much performance penalty. However, modern networks have comparable bandwidth as memory, so extra memory copying can be a communication bottleneck.

(2) The CPU is responsible for the copying operation and protocol processing (e.g. "checksum" to verify message integrity), and during the same time the CPU cannot perform useful numerical computations. The CPU processing speed can even become the bottleneck for very fast message transfer. This explains why a single pair of MPI processes (which only utilize a single pair of CPU cores) cannot drive the full 12.5 GB/s bandwidth on C5n.18xlarge, and a much higher bandwidth is achieved by adding more MPI pairs (which utilize more CPU cores), as shown in Section 4.2.

(3) Copying from user space to kernel space requires a Linux "system call" which further adds latency and disturbs the user program. A system call is for switching a Linux program between "user mode" and "kernel mode". A program typically runs in the user mode with limited permissions; accessing external devices (e.g. disk I/O or network communication) requires the kernel mode with higher privileges.

On the contrary, InfiniBand RDMA offers low-overhead communication features. Below are three main features, each addresses one of the above TCP limitations (Chapter 8 of Gavrilovska 2009):

(1) *Zero-copying.* The data in the user program are directly sent over the network, and get directly placed in the receiver program's memory. This removes the memory copy overhead as in TCP.

(2) *Protocol offloading.* The protocol processing workload is offloaded from the CPU to the network device (RNIC in Fig 2b). This allows the CPU to do more useful numerical computations (Squyres, 2009).

(3) *Kernel-bypassing.* The communication does not invoke the operating system (OS) kernel. The program can keep running in user mode without being disturbed. It is also called OS-bypassing or user-space networking (Squyres, 2015).

AWS EFA supports some RDMA-like features such as kernel-bypassing, to reduce CPU overhead and increase network performance.

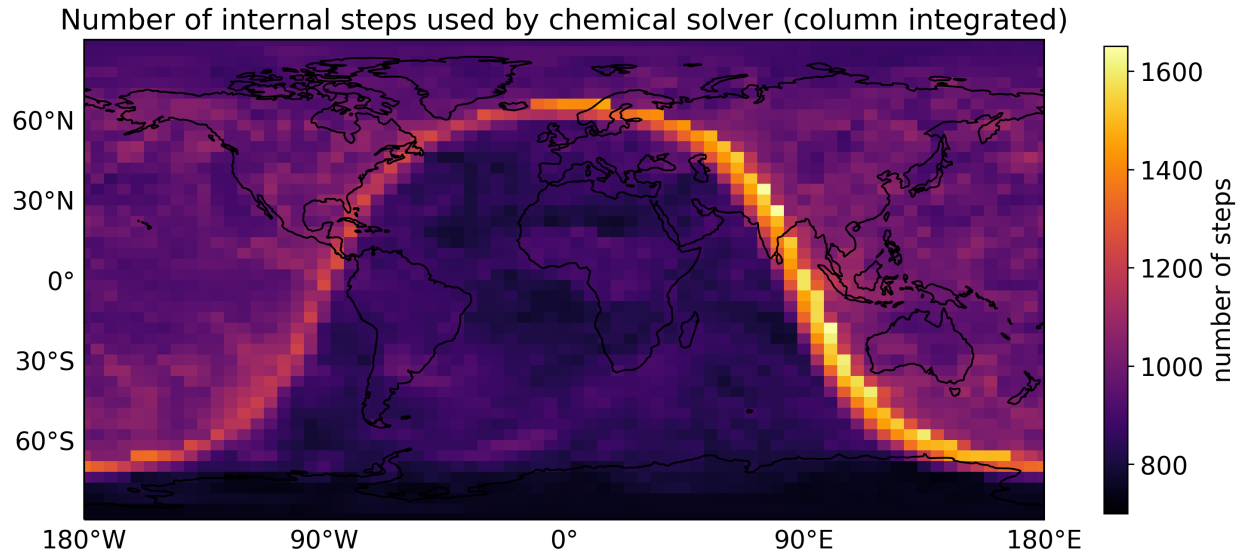**Appendix D. MPI Barrier caused by load imbalance in the chemical integrator**

Number of internal steps used by chemical solver (column integrated)

Figure D1. Number of internal steps taken by the Rosenbrock solver for integrating chemical kinetics in GEOS-Chem over a 20-minute external timestep, sampled at 2016-07-01 00:00 UTC for a global simulation at 4°×5° resolution. Numbers are summed over the 72 vertical layers for each column of the model grid.

The MPI profiling in Section 5.3 reveals a GEOS-Chem performance bottleneck with the "MPI_Barrier" function (Fig. 5). This bottleneck is not caused by inter-node communication or any cloud-platform-specific limitation. It comes instead from the load-imbalance of the chemical kinetics solver (the Rosenbrock algorithm of Sandu and Sander, 2006) when applied globally to the GEOS-Chem domain. Fig. D1 shows the number of internal time steps required for integrating the chemical kinetics over one 20-minute external time step, in a GEOS-Chem simulation at 4°×5° resolution using the standard tropospheric-stratospheric chemistry. The internal time steps are selected adaptively by the solver to achieve a certain error tolerance. The numbers are summed over the 72 vertical layers for each model grid column. The solver takes much more internal steps around sunrise/sunset when photolysis frequencies change rapidly. We find here that 93% of model columns use 700~1100 internal steps, while the sunrise/sunset columns use near 1650 internal steps. Thus different MPI processes may require very different time to complete the chemical integration depending on their number of sunrise/sunset columns. We verified that this was the problem by rerunning the model with chemistry turned off and observing a much shorter "MPI_Barrier" time. A similar pattern was shown by the Fig. 1 of Alvanos and Christoudias (2017) in their global atmospheric chemistry model. The code and further discussions for this problem are at https://github.com/geoschem/gchp/issues/44 and https://github.com/geoschem/geos-chem/issues/77.

One might attempt to correct this load imbalance by parallelizing the global domain by latitude bands, so that each compute node has a comparable number of sunrise/sunset grid boxes. But the advection component of the model would then require more inter-node communication and thus scale less efficiently. Alternatively, one might seek to speed up the chemical integration at sunrise/sunset. There has been interest within the GEOS-Chem community in developing adaptive chemical solvers to speed up the chemical integration (Santillana et al., 2010; Shen et

al., 2019). Our results indicate that particular attention should be placed on sunrise/sunset in order to accelerate massively parallel computation.

Acknowledgments, Samples, and Data

Scripts, notebooks and instructions to reproduce the GEOS-Chem performance benchmarks are available at https://github.com/JiaweiZhuang/cloud-gchp-paper (DOI: 10.5281/zenodo.3601521). Micro-benchmark scripts for AWS cloud and NASA Pleiades cluster are at https://github.com/JiaweiZhuang/aws-mpi-benchmark (DOI: 10.5281/zenodo.3601523) and https://github.com/JiaweiZhuang/pleiades-mpi-benchmark (DOI: 10.5281/zenodo.3601525) respectively. A Python library to process IPM output files is available at https://github.com/JiaweiZhuang/ipm_util (DOI: 10.5281/zenodo.3601527). Official GEOS-Chem repositories are all hosted under https://github.com/geoschem (DOI: 10.5281/zenodo.1343546). The specific branch used for this work is at https://github.com/JiaweiZhuang/gchp/tree/12.3.2-intelmpi-cloud-gchp-paper (DOI: 10.5281/zenodo.3601529).

References

Adhianto, L., Banerjee, S., Fagan, M., Krentel, M., Marin, G., Mellor-Crummey, J., & Tallent, N. R. (2010). HPCTOOLKIT: Tools for performance analysis of optimized parallel programs. *Concurrency Computation Practice and Experience*, *22*(6), 685–701. https://doi.org/10.1002/cpe

Alvanos, M., & Christoudias, T. (2017). GPU-accelerated atmospheric chemical kinetics in the ECHAM/MESSy (EMAC) Earth system model (version 2.52). *Geoscientific Model Development*, *10*(10), 3679–3693. https://doi.org/10.5194/gmd-10-3679-2017

Amazon. (2010). New Amazon EC2 Instance Type – The Cluster Compute Instance. Retrieved from https://aws.amazon.com/blogs/aws/the-new-amazon-ec2-instance-type-the-cluster-compute-instance/

Amazon. (2017a). Amazon EC2 Instances with Up to 8 NVIDIA Tesla V100 GPUs (P3). Retrieved November 10, 2019, from https://aws.amazon.com/blogs/aws/new-amazon-ec2-instances-with-up-to-8-nvidia-tesla-v100-gpus-p3/

Amazon. (2017b). EC2 F1 Instances with FPGAs – Now Generally Available. Retrieved November 10, 2019, from https://aws.amazon.com/blogs/aws/ec2-f1-instances-with-fpgas-now-generally-available/

Amazon. (2018a). Amazon EC2 Instance Types. Retrieved December 20, 2018, from https://aws.amazon.com/ec2/instance-types/

Amazon. (2018b). Amazon FSx for Lustre. Retrieved December 20, 2018, from

https://aws.amazon.com/fsx/lustre/

Amazon. (2018c). AWS Auto Scaling. Retrieved December 20, 2018, from https://aws.amazon.com/autoscaling/

Amazon. (2018d). New Amazon EC2 Spot pricing model: Simplified purchasing without bidding and fewer interruptions. Retrieved December 20, 2018, from https://aws.amazon.com/blogs/compute/new-amazon-ec2-spot-pricing/

Amazon. (2018e). New C5n Instances with 100 Gbps Networking. Retrieved December 20, 2018, from https://aws.amazon.com/blogs/aws/new-c5n-instances-with-100-gbps-networking/

Amazon. (2019a). AWS API Gateway for HPC job submission. Retrieved November 10, 2019, from https://aws.amazon.com/blogs/opensource/aws-api-gateway-hpc-job-submission/

Amazon. (2019b). *AWS helps researchers study "messages" from the universe*. Retrieved from https://aws.amazon.com/blogs/publicsector/aws-helps-researchers-study-messages-from-the-universe/

Amazon. (2019c). AWS ParallelCluster with AWS Directory Services Authentication. Retrieved November 10, 2019, from https://aws.amazon.com/blogs/opensource/aws-parallelcluster-aws-directory-services-authentication/

Amazon. (2019d). Deploying a Burstable and Event-driven HPC Cluster on AWS Using SLURM. Retrieved November 10, 2019, from https://aws.amazon.com/blogs/compute/deploying-a-burstable-and-event-driven-hpc-cluster-on-aws-using-slurm-part-1/

Amazon. (2019e). Elastic Fabric Adapter (EFA) for Tightly-Coupled HPC Workloads. Retrieved April 30, 2019, from https://aws.amazon.com/blogs/aws/now-available-elastic-fabric-adapter-efa-for-tightly-coupled-hpc-workloads/

Amazon. (2019f). How to run AWS ParallelCluster from AppStream 2.0 and share S3 data. Retrieved November 10, 2019, from https://aws.amazon.com/blogs/opensource/run-parallelcluster-from-appstream-2-0-share-s3-data/

Amazon. (2019g). Western Digital HDD Simulation at Cloud Scale – 2.5 Million HPC Tasks, 40K EC2 Spot Instances. Retrieved from https://aws.amazon.com/blogs/aws/western-digital-hdd-simulation-at-cloud-scale-2-5-million-hpc-tasks-40k-ec2-spot-instances/

Ansari, S., Del Greco, S., Kearns, E., Brown, O., Wilkins, S., Ramamurthy, M., et al. (2017). Unlocking the potential of NEXRAD data through NOAA's Big Data Partnership. *Bulletin of the American Meteorological Society*, BAMS-D-16-0021.1. https://doi.org/10.1175/BAMS-D-16-0021.1

Balaji, V. (2015). Climate Computing: The State of Play. *Computing in Science and Engineering*, *17*(6), 9–13. https://doi.org/10.1109/MCSE.2015.109

Baun, C., Kunze, M., Kurze, T., & Mauch, V. (2011). High performance computing as a service. *Advances in Parallel Computing*, *20*, 109–123. https://doi.org/10.3233/978-1-60750-803-8-109

Behnke, J., Mitchell, A., & Ramapriyan, H. (2019). NASA's earth observing data and

information system - Near-term challenges. *Data Science Journal, 18*(1), 1–9. https://doi.org/10.5334/dsj-2019-040

Bey, I., Jacob, D. J., Yantosca, R. M., Logan, J. A., Field, B. D., Fiore, A. M., et al. (2001). Global modeling of tropospheric chemistry with assimilated meteorology: Model description and evaluation. *Journal of Geophysical Research: Atmospheres, 106*(D19), 23073–23095.

Breuer, A., Cui, Y., & Heinecke, A. (2019). Petaflop seismic simulations in the public cloud. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 11501 LNCS*, 167–185. https://doi.org/10.1007/978-3-030-20656-7_9

Brikman, Y. (2018). *Terraform: Up & Running, 2nd Edition, Writing Infrastructure as Code*.

Calegari, P., Levrier, M., & Balczyński, P. (2019). Web portals for high-performance computing: A survey. *ACM Transactions on the Web, 13*(1), 1–36. https://doi.org/10.1145/3197385

Chan, E., & Heimlich, M. (2007). Collective Communication: Theory , Practice , and Experience, 1749–1783.

Chang, S., Hood, R., Jin, H., Heistand, S., Chang, J., Cheung, S., et al. (2018). Evaluating the Suitability of Commercial Clouds for NASA ' s High Performance Computing Applications : A Trade Study, (May). Retrieved from https://www.nas.nasa.gov/assets/pdf/papers/NAS_Technical_Report_NAS-2018-01.pdf

Chen, X., Huang, X., Jiao, C., Flanner, M. G., Raeker, T., & Palen, B. (2017). Running climate model on a commercial cloud computing environment: A case study using Community Earth System Model (CESM) on Amazon AWS. *Computers and Geosciences, 98*, 21–25. https://doi.org/10.1016/j.cageo.2016.09.014

Church, P., Goscinski, A., & Lefèvre, C. (2015). Exposing HPC and sequential applications as services through the development and deployment of a SaaS cloud. *Future Generation Computer Systems, 43–44*, 24–37. https://doi.org/10.1016/j.future.2014.10.001

Coffrin, C., Arnold, J., Eidenbenz, S., Aberle, D., Ambrosiano, J., Baker, Z., et al. (2019). The ISTI Rapid Response on Exploring Cloud Computing 2018, (August), 1–72. Retrieved from http://arxiv.org/abs/1901.01331

Donahoo, M. J., & Calvert, K. L. (2009). *TCP/IP Sockets in C, Second Edition: Practical Guide for Programmers*. *TCP/IP Sockets in C: Practical Guide for Programmers*.

Dongarra, J., & Luszczek, P. (2013). HPC Challenge: Design, History, and Implementation Highlights. In *Contemporary High Performance Computing*.

Dongarra, J. J. (2014). Performance of various computers using standard linear equations software, technical report CS-89-85. Retrieved from http://www.netlib.org/benchmark/performance.ps

Dubais, P. F., Epperly, T., & Kumfert, G. (2003). Why Johnny can't build. *Computing in Science and Engineering, 5*(5), 83–88. https://doi.org/10.1109/MCISE.2003.1225867

Düben, P. D., Russell, F. P., Niu, X., Luk, W., & Palmer, T. N. (2015). On the use of programmable hardware and reduced numerical precision in earth-system modeling.

Journal of Advances in Modeling Earth Systems, *7*(3), 1393–1408.
https://doi.org/10.1002/2015MS000494

Duvall, P. M., Matyas, S., & Glover, A. (2007). *Continuous integration: improving software quality and reducing risk*. Pearson Education.

Eastham, S. D., Long, M. S., Keller, C. A., Lundgren, E., Yantosca, R. M., Zhuang, J., et al. (2018). GEOS-Chem high performance (GCHP v11-02c): A next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications. *Geoscientific Model Development*, *11*(7), 2941–2953. https://doi.org/10.5194/gmd-11-2941-2018

ESMF. (2018). MAPL Refactoring for ESMF Regrid. Retrieved November 10, 2019, from https://www.earthsystemcog.org/projects/cupid/mapl_refactor

Evangelinos, C., & Hill, C. N. (2008). Cloud Computing for parallel Scientific HPC Applications: Feasibility of Running Coupled Atmosphere-Ocean Climate Models on Amazon's EC2. *Cca'08*, *2*(2.40), 2–34.

Foster, I., & Gannon, D. B. (2017). *Cloud Computing for Science and Engineering*. Retrieved from https://cloud4scieng.org/

Freniere, C., Pathak, A., Raessi, M., & Khanna, G. (2016). The feasibility of amazon's cloud computing platform for parallel, GPU-accelerated, multiphase-flow simulations. *Computing in Science and Engineering*, *18*(5), 68–77. https://doi.org/10.1109/MCSE.2016.94

Gabriel, E., Fagg, G. E., Bosilca, G., Angskun, T., Dongarra, J. J., Squyres, J. M., et al. (2004). Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation, 97–104. https://doi.org/10.1007/978-3-540-30218-6_19

Gamblin, T., LeGendre, M., Collette, M. R., Lee, G. L., Moody, A., de Supinski, B. R., & Futral, S. (2015). The Spack Package Manager: Bringing Order to HPC Software Chaos. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis on - SC '15*. https://doi.org/10.1145/2807591.2807623

Gavrilovska, A. (2009). *Attaining high performance communications: a vertical approach*. Chapman and Hall/CRC.

Geimer, M., Hoste, K., & McLay, R. (2014). Modern scientific software management using easybuild and lmod. In *Proceedings of HUST 2014: 1st International Workshop on HPC User Support Tools - Held in Conjunction with SC 2014: The International Conference for High Performance Computing, Networking, Storage and Analysis* (pp. 41–51). https://doi.org/10.1109/HUST.2014.8

GEOS-Chem. (2019). GEOS-Chem Benchmarks. Retrieved November 10, 2019, from http://wiki.seas.harvard.edu/geos-chem/index.php/GEOS-Chem_Benchmarks

Glick, B., & MacHe, J. (2019). Jupyter notebooks and user-friendly HPC access. *Proceedings of EduHPC 2018: Workshop on Education for High-Performance Computing, Held in Conjunction with SC 2018: The International Conference for High Performance Computing, Networking, Storage and Analysis*, 11–20. https://doi.org/10.1109/EduHPC.2018.00005

Google. (2018a). Cloud TPU machine learning accelerators now available in beta. Retrieved

November 10, 2019, from https://cloud.google.com/blog/products/gcp/cloud-tpu-machine-learning-accelerators-now-available-in-beta

Google. (2018b). Google Cloud Platform for AWS Professionals. Retrieved December 20, 2018, from https://cloud.google.com/docs/compare/aws/

Google. (2019). New climate model data now in Google Public Datasets. Retrieved December 20, 2019, from https://cloud.google.com/blog/products/data-analytics/new-climate-model-data-now-google-public-datasets

Gregg, B. (2017). AWS EC2 Virtualization 2017: Introducing Nitro. Retrieved February 1, 2019, from http://www.brendangregg.com/blog/2017-11-29/aws-ec2-virtualization-2017.html

Grun, P. (2010). Introduction to InfiniBand ™ for End Users. *White Paper, InfiniBand Trade Association*. Retrieved from http://www.mellanox.com/pdf/whitepapers/Intro_to_IB_for_End_Users.pdf

Grun, P., Hefty, S., Sur, S., Goodell, D., Russell, R. D., Pritchard, H., & Squyres, J. M. (2015). A Brief Introduction to the OpenFabrics Interfaces - A New Network API for Maximizing High Performance Application Efficiency. *Proceedings - 2015 IEEE 23rd Annual Symposium on High-Performance Interconnects, HOTI 2015*, 34–39. https://doi.org/10.1109/HOTI.2015.19

Guo, C., Wu, H., Deng, Z., Soni, G., Ye, J., Padhye, J., & Lipshteyn, M. (2016). RDMA over Commodity Ethernet at Scale. *Proceedings of the 2016 Conference on ACM SIGCOMM 2016 Conference - SIGCOMM '16*. https://doi.org/10.1145/2934872.2934908

Hager, G., & Wellein, G. (2010). *Introduction to high performance computing for scientists and engineers. Introduction to High Performance Computing for Scientists and Engineers*. https://doi.org/10.1201/EBK1439811924

Heroux, M. A., Carter, J., Thakur, R., Vetter, J., McInnes, L. C., Ahrens, J., & Neely, J. R. (2019). *Second release of ECP software technology capability assessment report*. Retrieved from https://www.exascaleproject.org/ecp-software-technology-capability-assessment-report-second-release/

Hill, Z., & Humphrey, M. (2009). A quantitative analysis of high performance computing with Amazon's EC2 infrastructure: The death of the local cluster? *Proceedings - IEEE/ACM International Workshop on Grid Computing*, 26–33. https://doi.org/10.1109/GRID.2009.5353067

Howe, B. (2012). Virtual appliances, cloud computing, and reproducible research. *Computing in Science and Engineering*, *14*(4), 36–41. https://doi.org/10.1109/MCSE.2012.62

Hoyer, S., & Hamman, J. J. (2017). xarray: N-D labeled Arrays and Datasets in Python. *Journal of Open Research Software*, *5*, 1–6. https://doi.org/10.5334/jors.148

Huang, Q. (2014). Development of a SaaS application probe to the physical properties of the Earth's interior: An attempt at moving HPC to the cloud. *Computers and Geosciences*, *70*, 147–153. https://doi.org/10.1016/j.cageo.2014.06.002

Iosup, A., Ostermann, S., Yigitbasi, M. N., Prodan, R., Fahringer, T., & Epema, D. H. J. (2011). Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. *IEEE Transactions on Parallel and Distributed Systems*, *22*(6), 931–945.

https://doi.org/10.1109/TPDS.2011.66

Irving, D. (2016). A minimum standard for publishing computational results in the weather and climate sciences. *Bulletin of the American Meteorological Society*, *97*(7), 1149–1158. https://doi.org/10.1175/BAMS-D-15-00010.1

Jackson, K. R., Ramakrishnan, L., Muriki, K., Canon, S., Cholia, S., Shalf, J., et al. (2010). Performance Analysis of High Performance Computing Applications on the Amazon Web Services Cloud. *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, 159–168. https://doi.org/10.1109/CloudCom.2010.69

Jouppi, N. P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., et al. (2017). In-Datacenter Performance Analysis of a Tensor Processing Unit. *Nature*, *264*(5581), 79–81. Retrieved from http://arxiv.org/abs/1704.04760

Jung, K., Cho, Y.-K., & Tak, Y.-J. (2017). Performance evaluation of ROMS v3.6 on a commercial cloud system. *Geoscientific Model Development Discussions*, 1–43. https://doi.org/10.5194/gmd-2017-270

Jupyter, Bussonnier, M., Forde, J., Freeman, J., Granger, B., Head, T., et al. (2018). Binder 2.0 - Reproducible, interactive, sharable environments for science at scale. *Proceedings of the 17th Python in Science Conference*, (Scipy), 113–120. https://doi.org/10.25080/majora-4af1f417-011

Kay, J. E., Deser, C., Phillips, A., Mai, A., Hannay, C., Strand, G., et al. (2015). The community earth system model (CESM) large ensemble project : A community resource for studying climate change in the presence of internal climate variability. *Bulletin of the American Meteorological Society*, *96*(8), 1333–1349. https://doi.org/10.1175/BAMS-D-13-00255.1

Keller, C. A., Long, M. S., Yantosca, R. M., Da Silva, A. M., Pawson, S., & Jacob, D. J. (2014). HEMCO v1.0: A versatile, ESMF-compliant component for calculating emissions in atmospheric models. *Geoscientific Model Development*, *7*(4), 1409–1417. https://doi.org/10.5194/gmd-7-1409-2014

Kurtzer, G. M., Sochat, V., & Bauer, M. W. (2017). Singularity: Scientific containers for mobility of compute. *PLoS ONE*, *12*(5), e0177459. https://doi.org/10.1371/journal.pone.0177459

de La Beaujardière, J., Banihirwe, A., Shih, C.-F., & Paul, K.Hamman, J. (2019). NCAR CESM LENS Cloud-Optimized Subset. https://doi.org/10.26024/wt24-5j82

Lawrence, B. N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., et al. (2018). Crossing the chasm: How to develop weather and climate models for next generation computers? *Geoscientific Model Development*, *11*(5), 1799–1821. https://doi.org/10.5194/gmd-11-1799-2018

Liu, J., Chandrasekaran, B., Wu, J., Jiang, W., Kini, S., Yu, W., et al. (2003). Performance Comparison of MPI Implementations over InfiniBand, Myrinet and Quadrics. *Proceedings of the 2003 ACM/IEEE Conference on Supercomputing - SC '03*, 58. https://doi.org/10.1145/1048935.1050208

Lockwood, G. K. (2013). Quick MPI Cluster Setup on Amazon EC2. Retrieved April 2, 2019, from https://glennklockwood.blogspot.com/2013/04/quick-mpi-cluster-setup-on-amazon-

ec2.html

Lockwood, G. K., Tatineni, M., & Wagner, R. (2014). SR-IOV: Performance Benefits for Virtualized Interconnects. *Proceedings of the 2014 Annual Conference on Extreme Science and Engineering Discovery Environment - XSEDE '14*, 1–7. https://doi.org/10.1145/2616498.2616537

Long, M. S., Yantosca, R., Nielsen, J. E., Keller, C. A., Da Silva, A., Sulprizio, M. P., et al. (2015). Development of a grid-independent GEOS-Chem chemical transport model (v9-02) as an atmospheric chemistry module for Earth system models. *Geoscientific Model Development*, 8(3), 595–602. https://doi.org/10.5194/gmd-8-595-2015

Lynnes, C., Baynes, K., & McInerney., M. A. (2017). Archive Management of NASA Earth Observation Data to Support Cloud Analysis, 0–3. Retrieved from https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20170011455.pdf

Madhyastha, T. M., Koh, N., Day, T. K. M., Hernández-Fernández, M., Kelley, A., Peterson, D. J., et al. (2017). Running Neuroimaging Applications on Amazon Web Services: How, When, and at What Cost? *Frontiers in Neuroinformatics*, 11(November), 1–15. https://doi.org/10.3389/fninf.2017.00063

McCalpin, J. D. (1995). Memory Bandwidth and Machine Balance in Current High Performance Computers. *IEEE Computer Society Technical Committee on Computer Architecture Newsletter*, 19–25.

McLay, R., Schulz, K. W., Barth, W. L., & Minyard, T. (2011). Best practices for the deployment and management of production HPC clusters. *2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 1–11. https://doi.org/10.1145/2063348.2063360

Mehrotra, P., Djomehri, J., Heistand, S., Hood, R., Jin, H., Lazanoff, A., et al. (2016). Performance evaluation of Amazon Elastic Compute Cloud for NASA high-performance computing applications. *Concurrency Computation* , 28(4), 1041–1055. https://doi.org/10.1002/cpe.3029

Mellanox Technologies. (2014). Introduction to InfiniBand. Retrieved from http://www.mellanox.com/blog/2014/09/introduction-to-infiniband/

Microsoft. (2018). Azure for AWS Professionals. Retrieved December 20, 2018, from https://docs.microsoft.com/en-us/azure/architecture/aws-professional/

Microsoft. (2019). Azure Pipelines. Retrieved November 10, 2019, from https://azure.microsoft.com/en-us/services/devops/pipelines/

Milligan, M. B. (2018). Jupyter as common technology platform for interactive HPC services. *ACM International Conference Proceeding Series*. https://doi.org/10.1145/3219104.3219162

Mohammadi, M., & Bazhirov, T. (2018). Comparative benchmarking of cloud computing vendors with High Performance Linpack. *ACM International Conference Proceeding Series*, (60), 1–5. https://doi.org/10.1145/3195612.3195613

Molthan, A., Case, J., & Venner, J. (2015). Clouds in the cloud: weather forecasts and applications within cloud computing environments. *Bulletin of The*, 96(8), 1369–1379.

Retrieved from http://journals.ametsoc.org/doi/abs/10.1175/BAMS-D-14-00013.1

Monajemi, H., Murri, R., Jonas, E., Liang, P., Stodden, V., & Donoho, D. (2019). Ambitious Data Science Can Be Painless. *Harvard Data Science Review*. https://doi.org/10.1162/99608f92.02ffc552

Morris, K. (2016). *Infrastructure As Code: Managing Servers in the Cloud*. Retrieved from http://shop.oreilly.com/product/0636920039297.do

NAS. (2019). *Reproducibility and Replicability in Science*. National Academies Press. https://doi.org/10.17226/25303

NASA HECC. (2019). HECC AWS Cloud: Overview. Retrieved November 10, 2019, from https://www.nas.nasa.gov/hecc/support/kb/hecc-aws-cloud-overview_581.html

Netto, M. A. S., Calheiros, R. N., Rodrigues, E. R., Cunha, R. L. F., & Buyya, R. (2017). HPC Cloud for Scientific and Business Applications: Taxonomy, Vision, and Research Challenges, *51*(1), 1–29. https://doi.org/10.1145/3150224

NSF. (2019). Enabling Access to Cloud Computing Resources for CISE Research and Education (Cloud Access). Retrieved November 10, 2019, from https://www.nsf.gov/pubs/2019/nsf19510/nsf19510.htm

de Oliveira, A. H. M., de Oliveira, D., & Mattoso, M. (2017). Clouds and Reproducibility: A Way to Go to Scientific Experiments? In *Cloud Computing* (pp. 127–151). Springer International Publishing. https://doi.org/10.1007/978-3-319-54645-2_5

Parker, S., Chunduri, S., Harms, K., & Kandalla, K. (2018). Performance Evaluation of MPI on Cray XC40 Xeon Phi Systems. *Cray User Group Proceedings*. Retrieved from https://cug.org/proceedings/cug2018_proceedings/includes/files/pap131s2-file1.pdf

Patarasuk, P., & Yuan, X. (2009). Bandwidth optimal all-reduce algorithms for clusters of workstations. *Journal of Parallel and Distributed Computing*, *69*(2), 117–124. https://doi.org/10.1016/j.jpdc.2008.09.002

Pérez, F., Hamman, J., Larsen, L., Paul, K., Heagy, L., Holdgraf, C., & Panda, Y. (2019). Jupyter meets the Earth: Enabling discovery in geoscience through interactive computing at scale. *Zenodo*. https://doi.org/http://doi.org/10.5281/zenodo.3369939

Perkel, J. M. (2018). Why Jupyter is data scientists' computational notebook of choice. *Nature*, *563*(7729), 145–146. https://doi.org/10.1038/d41586-018-07196-1

Perkel, J. M. (2019). Make code accessible with these cloud services. *Nature*, *575*(7781), 247–248. https://doi.org/10.1038/d41586-019-03366-x

Pješivac-Grbović, J., Angskun, T., Bosilca, G., Fagg, G. E., & Dongarra, E. G. J. J. (2007). Performance Analysis of MPI Collective Operations. *Cluster Computing*, *10*(2), 127–143. https://doi.org/10.1007/s10586-007-0012-0

Priedhorsky, R., Randles, T. C., & Randles, T. (2017). Charliecloud: Unprivileged containers for user-defined software stacks in HPC Charliecloud: Unprivileged containers for user-defined software stacks in HPC. *SC17: International Conference for High Performance Computing, Networking, Storage and Analysis*, *17*, p1-10. https://doi.org/10.1145/3126908.3126925

Prout, A., Arcand, W., Bestor, D., Bergeron, B., Byun, C., Gadepally, V., et al. (2017). MIT

SuperCloud portal workspace: Enabling HPC web application deployment. *2017 IEEE High Performance Extreme Computing Conference, HPEC 2017*, 1–6. https://doi.org/10.1109/HPEC.2017.8091097

Putman, W. M. (2007). *Development of the Finite-Volume Dynamical Core on the Cubed-Sphere* (Vol. D). The Florida State University. Retrieved from https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=3&ved=2ahUKEwi bkt_at87hAhUKvlkKHRyYDtkQFjACegQIABAC&url=https%3A%2F%2Fdiginole.lib.fsu .edu%2Fislandora%2Fobject%2Ffsu%3A168667%2Fdatastream%2FPDF%2Fdownload%2 Fcitation.pdf&usg=AOvVaw0ERsMpAV53e

Putman, W. M., & Lin, S.-J. (2007). Finite-volume transport on various cubed-sphere grids. *Journal of Computational Physics*, *227*(1), 55–78. https://doi.org/10.1016/j.jcp.2007.07.022

Ramachandran, R., Lynnes, C., Baynes, K., Murphy, K., Baker, J., Kinney, J., et al. (2018). Recommendations to Improve Downloads of Large Earth Observation Data. *Data Science Journal*, *17*(0), 1–6. https://doi.org/10.5334/dsj-2018-002

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., & Prabhat, &. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*. https://doi.org/10.1038/s41586-019-0912-1

Robinson, N. H., Hamman, J., & Abernathey, R. (2019). Science needs to rethink how it interacts with big data: Five principles for effective scientific big data systems, 1–8. Retrieved from http://arxiv.org/abs/1908.03356

Rocklin, M. (2015). Dask : Parallel Computation with Blocked algorithms and Task Scheduling. *Proceedings of the 14th Python in Science Conference*, (Scipy), 130–136.

Sadooghi, I., Hernandez Martin, J., Li, T., Brandstatter, K., Zhao, Y., Maheshwari, K., et al. (2015). Understanding the Performance and Potential of Cloud Computing for Scientific Applications. *IEEE Transactions on Cloud Computing*, *PP*(99), 1–1. https://doi.org/10.1109/TCC.2015.2404821

Sampedro, Z., Holt, A., & Hauser, T. (2018). Continuous integration and delivery for HPC: Using Singularity and Jenkins. *ACM International Conference Proceeding Series*, 1–6. https://doi.org/10.1145/3219104.3219147

Sandu, A., & Sander, R. (2006). Technical note: Simulating chemical systems in Fortran90 and Matlab with the Kinetic PreProcessor KPP-2.1. *Atmospheric Chemistry and Physics*, *6*(1), 187–195.

Santillana, M., Le Sager, P., Jacob, D. J., & Brenner, M. P. (2010). An adaptive reduction algorithm for efficient chemical calculations in global atmospheric chemistry models. *Atmospheric Environment*, *44*(35), 4426–4431. https://doi.org/10.1016/j.atmosenv.2010.07.044

Sarajlic, S., Chastang, J., Marru, S., Fischer, J., & Lowe, M. (2018). Scaling JupyterHub using Kubernetes on Jetstream cloud: Platform as a service for research and educational initiatives in the atmospheric sciences. *ACM International Conference Proceeding Series*. https://doi.org/10.1145/3219104.3229249

Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth System Modeling 2.0: A Blueprint

for Models That Learn From Observations and Targeted High-Resolution Simulations, 1–22. https://doi.org/10.1002/2017GL076101

Schulz, K. W., Baird, C. R., Brayford, D., Georgiou, Y., Kurtzer, G. M., Simmel, D., et al. (2016). Cluster computing with OpenHPC.

Schwan, P. (2003). Lustre: Building a File System for 1,000-node Clusters. *Proceedings of the Linux Symposium*, 401–409. https://doi.org/10.1.1.2.456

Shan, H., & Shalf, J. (2010). Using IOR to analyze the I/O performance for HPC. Retrieved from https://escholarship.org/uc/item/9111c60j

Shen, L., Jacob, D. J., Santillana, M., Wang, X., & Chen, W. (2019). An adaptive method for speeding up the numerical integration of chemical mechanisms in atmospheric chemistry models: application to GEOS-Chem version 12.0.0 (in review). *Geosci. Model Dev. Discuss.*, (October), 1–20. https://doi.org/https://doi.org/10.5194/gmd-2019-279

Shende, S. S., & Malony, A. D. (2006). The TAU parallel performance system. *International Journal of High Performance Computing Applications*, *20*(2), 287–311. https://doi.org/10.1177/1094342006064482

Skinner, D. (2005). Performance monitoring of parallel scientific applications. https://doi.org/10.2172/881368

Squyres, J. (2009). Network hardware offload. Retrieved from https://blogs.cisco.com/performance/network_hardware_offload

Squyres, J. (2015). MPI newbie: What is "operating system bypass"? Retrieved from https://blogs.cisco.com/performance/mpi-newbie-what-is-operating-system-bypass

Suarez, M., Trayanov, A., Hill, C., Schopf, P., & Vikhliaev, Y. (2007). MAPL: a high-level programming paradigm to support more rapid and robust encoding of hierarchical trees of interacting high-performance components. In *Proceedings of the 2007 symposium on Component and framework technology in high-performance and scientific computing - CompFrame '07* (p. 11). ACM. https://doi.org/10.1145/1297385.1297388

Sudmanns, M., Tiede, D., Lang, S., Bergstedt, H., Trost, G., Augustin, H., et al. (2019). Big Earth data: disruptive changes in Earth observation data management and analysis? *International Journal of Digital Earth*, *0*(0), 1–19. https://doi.org/10.1080/17538947.2019.1585976

Thain, D., Tannenbaum, T., & Livny, M. (2005). Distributed computing in practice: the Condor experience: Research Articles. *Concurrency and Computation: Practice & Experience*, *17*(2–4), 323–356. https://doi.org/10.1002/cpe.v17:2/4

Vance, T. C., Merati, N., Yang, C. P., & Yuan, M. (2016). *Cloud Computing in Ocean and Atmospheric Sciences*. Elsevier. https://doi.org/10.1016/C2014-0-04015-4

Walker, E. (2008). Benchmarking Amazon EC2 for high-performance scientific computing. *Login:: The Magazine of USENIX & SAGE*, *33*(5), 18–23.

Wessel, J. R., Gorgolewski, K. J., & Bellec, P. (2019). Switching Software in Science: Motivations, Challenges, and Solutions. *Trends in Cognitive Sciences*, *xx*, 1–3. https://doi.org/10.1016/j.tics.2019.01.004

Wittig, M., & Wittig, A. (2018). *Amazon Web Services in Action, Second Edition*. Manning.

Wong, A. K. L., & Goscinski, A. M. (2013). A unified framework for the deployment, exposure and access of HPC applications as services in clouds. *Future Generation Computer Systems*, *29*(6), 1333–1344. https://doi.org/10.1016/j.future.2013.01.014

Wright, N., & Pfeiffer, W. (2009). Characterizing parallel scaling of scientific applications using IPM. *The 10th LCI International Conference on High-Performance Clustered Computing, March 10-12*, 1–21. Retrieved from http://users.sdsc.edu/~nwright/IPM_LCI09_Final.pdf

Yang, C., Huang, Q., Li, Z., Liu, K., & Hu, F. (2017). Big Data and cloud computing: innovation opportunities and challenges. *International Journal of Digital Earth*, *10*(1), 13–53. https://doi.org/10.1080/17538947.2016.1239771

Yelick, K., Coghlan, S., Draney, B., Ramakrishnan, L., Scovel, A., Sakrejda, I., et al. (2011). The Magellan Report on Cloud Computing for Science, 170.

Yoo, A. B., Jette, M. A., & Grondona, M. (2003). SLURM: Simple Linux Utility for Resource Management, 44–60. https://doi.org/10.1007/10968987_3

Younge, A. J., Pedretti, K., Grant, R. E., & Brightwell, R. (2017). A Tale of Two Systems: Using Containers to Deploy HPC Applications on Supercomputers and Clouds. *Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom*, *2017–Decem*, 74–81. https://doi.org/10.1109/CloudCom.2017.40

Zhai, Y., Liu, M., Zhai, J., Ma, X., & Chen, W. (2011). Cloud Versus In-house Cluster: Evaluating Amazon Cluster Compute Instances for Running MPI Applications. *State of the Practice Reports on - SC '11*, 1. https://doi.org/10.1145/2063348.2063363

Zhang, J., Lu, X., & Panda, D. K. (2017). Is Singularity-based Container Technology Ready for Running MPI Applications on HPC Clouds? *Proceedings of The10th International Conference on Utility and Cloud Computing - UCC '17*, 151–160. https://doi.org/10.1145/3147213.3147231

Zhuang, J. (2019a). A Scientist's Guide to Cloud-HPC: Example with AWS ParallelCluster, Slurm, Spack, and WRF. Retrieved April 20, 2019, from https://jiaweizhuang.github.io/blog/aws-hpc-guide/

Zhuang, J. (2019b). MPI over Multiple TCP Connections on EC2 C5n Instances. Retrieved November 10, 2019, from https://jiaweizhuang.github.io/blog/mpi-tcp-ec2/

Zhuang, J., Jacob, D. J., Gaya, J. F., Yantosca, R. M., Lundgren, E. W., Sulprizio, M. P., & Eastham, S. D. (2019). Enabling immediate access to earth science models through cloud computing. *Bulletin of the American Meteorological Society*, *100*(10), 1943–1960. https://doi.org/10.1175/BAMS-D-18-0243.1