Deploying RMACC Summit: An HPC Resource for the Rocky Mountain Region

Jonathon Anderson
Research Computing Group
University of Colorado Boulder
jonathon.anderson@colorado.edu

Patrick J. Burns
Vice President for IT
Colorado State University
patrick.burns@colostate.edu

Daniel Milroy
Research Computing Group
University of Colorado Boulder
daniel.milroy@colorado.edu

Peter Ruprecht
Research Computing Group
University of Colorado Boulder
peter.ruprecht@colorado.edu

Thomas Hauser
Research Computing Group
University of Colorado Boulder
thomas.hauser@colorado.edu

Howard Jay Siegel
Electrical and Computer Engineering
Colorado State University
HJ@colostate.edu

ABSTRACT

RMACC Summit is a heterogeneous supercomputer cluster with an aggregate floating point performance of 379 TFLOPS (Rmax, as currently configured) that provides about 85 million core-hours/year to researchers from institutions participating in the Rocky Mountain Advanced Computing Consortium (RMACC). The development of Summit was a collaborative effort toward specifying a system that meets the needs of researchers at multiple universities, and included implementation and testing of several new technologies. We discuss our experiences in creating and maintaining a successful ongoing collaboration between the two universities that are RMACC Summit's primary operators, and consider both the technical and support challenges of extending that collaboration to other regional users.

CCS CONCEPTS

• **Computer systems organization** → **Grid computing**; *Heterogeneous (hybrid) systems*; • **Information systems** → Information storage systems; • **Social and professional topics** → Management of computing and information systems;

KEYWORDS

High-performance computing; supercomputers; clusters

ACM Reference format:

Jonathon Anderson, Patrick J. Burns, Daniel Milroy, Peter Ruprecht, Thomas Hauser, and Howard Jay Siegel. 2017. Deploying RMACC Summit: An HPC Resource for the Rocky Mountain Region. In *Proceedings of PEARC17, New Orleans, LA, USA, July 09-13, 2017*, 7 pages.
DOI: 10.1145/3093338.3093379

1 INTRODUCTION

1.1 Overview

The Rocky Mountain Advanced Computing Consortium (RMACC), led by the University of Colorado Boulder (CU Boulder) and Colorado State University (CSU), has designed and procured a supercomputer that will be used by several hundred researchers in the Rocky Mountain region for computational science and engineering research. The supercomputer allows researchers to explore the impact of new technology and high core-count on application performance, including the Intel Omni-Path interconnect and Intel's Xeon "Haswell" and Xeon Phi "Knights Landing" many-core processor nodes.

Furthermore, this joint activity encourages and facilitates collaborative research across multiple institutions and a wide variety of disciplines. CU and CSU were founding members of the Front Range Consortium for Research Computing (FRCRC), which was instantiated with CSU, CU Boulder, the Colorado School of Mines, and the University of Wyoming, along with three national labs: the National Center for Atmospheric Research, the National Oceanic and Atmospheric Administration, and the National Renewable Energy Lab. This activity was so successful in Colorado that it was recently expanded to multiple states, as the Rocky Mountain Advanced Computing Consortium (see <http://www.rmacc.org>), encompassing research universities and national labs in Colorado, Idaho, Utah, Wyoming, New Mexico, and Montana.

This paper explains the motivation and use case for a regional HPC resource; narrates the collaborative process that supported its design, proposal, procurement, and testing; presents issues and solutions resulting from its distributed user community; and presents experiences working with the cutting-edge technology present in the system's architecture. This information should help future collaborative HPC site developers by providing a model from our experience while signaling potential pitfalls that might otherwise unexpectedly complicate such a deployment.

1.2 RMACC goals and activities

The RMACC mission is to facilitate widespread effective use of high performance computing throughout the Rocky Mountain region by:

This work is licensed under a Creative Commons Attribution International 4.0 License.

PEARC17, New Orleans, LA, USA

© 2017 Copyright held by the owner/author(s). 978-1-4503-5272-7/17/07...\$15.00
DOI: 10.1145/3093338.3093379

- (1) Educating graduate and undergraduate students, faculty, researchers, and industry partners on the use of computational science and high performance computing.
- (2) Coordinating multi-institutional efforts to advance research, practice, and education in computational science and to address important regional problems.
- (3) Bringing together a broad range of researchers, faculty, and industry partners with a depth of experience and expertise not available at any single institution and facilitate their collaboration in multi-disciplinary and multi-institutional teams.

Since 2011, FRCRC/RMACC has held an annual HPC symposium, sponsored thrice-yearly HPC system administrators meetings, and provided training for faculty, students, and staff in new and emergent areas of HPC. Five RMACC partners also participated in NSF grant ACI-1440642 “CC*III Region: Rocky Mountain Cyberinfrastructure Mentoring and Outreach Alliance (RMCMAO),” extending our activities to smaller institutions in the region. Indeed, that grant provided motivation for the regional focus of the Summit supercomputer, ten percent of which is made available to the other RMACC members, with emphasis on serving sites without their own significant HPC resources.

1.3 Motivation for a regional HPC resource

A supercomputing cluster requires significant supporting infrastructure, including a data center with sufficient power and cooling, dedicated full time system administration staff, and high-performance local and wide-area networking. Spending resources on duplicated infrastructure across multiple regional institutions can substantially reduce the funds available for computing capability. Thus, in 2014 when we – the Colorado State University Information Science and Technology Center (CSU ISTeC) and the University of Colorado Boulder Research Computing group (CU Boulder RC) – began planning to replace our respective aging supercomputing systems, a collaborative effort made obvious sense. Furthermore, it was clear that a new HPC cluster could benefit other RMACC institutions in the region: smaller colleges and universities that could not individually support a large cluster as well as other flagship universities and national labs who would be interested in testing cutting-edge HPC technology.

CU Boulder and CSU submitted a joint proposal to the National Science Foundation’s Major Research Instrumentation (Advanced Cyberinfrastructure) grant program in 2015, which resulted in awards ACI-1532235 and ACI-1532236. The resulting HPC cluster, named RMACC Summit, is operated by RC and jointly supported by RC and ISTeC for the benefit of researchers at CU Boulder, CSU, and members of RMACC.

The CSU/CU Boulder collaboration has benefitted from differing strengths and emphases in the two universities. ISTeC was established in the late 1990s to enhance activities in IT fields. More recently, CU Boulder leadership enhanced support for HPC as a strategic imperative, and provided substantial funding for a new department in the Office of Information Technology (OIT) to support the cyberinfrastructure needs of researchers across campus.

2 COLLABORATION PLANNING AND INITIAL SYSTEM DESIGN

2.1 Multi-site collaboration model

The division of labor that emerged after discussion between the two institutions was that CU Boulder would house, power, cool, monitor, and perform system administration for the cluster, as it is axiomatic that system administration must be performed by a single, responsible entity, rather than being a shared activity. Moreover, CU Boulder had the expertise and available staff in this area, having operated its previous supercomputing cluster, named Janus [19], for over six years. Additionally, the extra effort to support the system for CSU users in addition to CU Boulder users was mostly required up front, during implementation. Each institution would provide consulting support to its own users, and for the most part would provide its own applications within its user and project storage spaces. However, applications that are commonly used may be installed, after discussion, in a shared space for all users. User support is handled via a single ticket system (ServiceNow) provided by the CU Boulder OIT. OIT also provided overall project management for the acquisition and implementation to ensure timely deployment of the system, while CSU has provided project management for aspects of the deployment that require multi-site coordination, including account and authorization federation, user-facing communications, network access and data transfer, and shared software support and installation.

Remote user access to RMACC Summit from CSU is facilitated because both sites are connected via high-bandwidth networks through the Front Range GigaPoP, our Regional Optical Network. Both institutions have recently received NSF Campus Cyberinfrastructure grants to upgrade their networks to provide better support for research, and each had implemented a local high-performance Science DMZ [13], increasing their WAN bandwidth commensurately. This provides sufficient access to RMACC Summit for CSU users, even though Summit is physically housed at CU Boulder.

2.2 Usage-driven design considerations

About six months prior to submitting the NSF MRI grant proposal, the RC staff began work to determine what features of a new cluster would best meet existing and expected user needs. We analyzed patterns of usage of Janus and other RC-managed clusters to characterize job width, duration, and memory usage. From this data, it was clear that, even though Janus had a fully non-blocking QDR InfiniBand interconnect network, the large majority of jobs used fewer than 30 nodes. Even in terms of overall core-hours consumed, most of the usage of Janus was by these smaller jobs. However, the number of jobs requiring thousands of cores and substantial MPI communication was big enough to imply that the future cluster would need to support MPI jobs of at least 256 nodes.

In addition, RC reached out to a representative set of users about their expected computational needs, which yielded four major themes. First, many dozens of different applications would need to be supported on a new cluster, including commercial binary distributions, well-supported “community” codes, domain-specific applications, and home-built codes used by individual research groups. Next, the majority of users identified themselves primarily

as domain researchers without substantial experience in parallel programming; their preference was for ease in getting their applications working well without requiring a substantial porting effort. Also, increased RAM per core (as compared to Janus) would be needed for anticipated growth in data-intensive applications, especially those used by CSU life sciences researchers. Finally, small-file-intensive workflows are increasingly prevalent, particularly in research domains that are relatively new to traditional HPC.

2.3 Physical installation considerations

Though a small set of potential sites were considered for deployment of a new HPC resource on the CU Boulder campus, we eventually selected the existing High-Performance Computing Facility (HPCF) [19] for its ability to support dense power and cooling. The HPCF provides rack space in two separate hot-aisle containment “pods,” both of which were occupied by Janus. While Janus was meant to be replaced by RMACC Summit, a period of overlapping operation was desired to support a gradual transition. As such, we made the decision to decommission Janus hardware from one pod to provide space for Summit. This freed ten 48RU racks to house Summit, with each rack providing up to 35 kW of power at 208 V.

The HPCF hot-aisle containment system also necessitated the use of air cooling at the server. Though chilled water provides the heat exchange for the in-row air-handling system, the water supply layout, coupled with the existing narrow hot-aisle, was not suitable for incorporation of more direct water cooling.

2.4 Technical and performance requirements for RFP

We envisioned installing RMACC Summit in two phases: a large initial deployment of CPU- and GPU-based compute nodes, followed by a smaller installation of Xeon Phi-based compute nodes. Thus, when developing a Request for Proposal (RFP) in mid-2015 we specified Intel’s Omni-Path Architecture (OPA) for the cluster’s high-performance interconnect because we believed that it would provide the most advanced interconnect available for the conventional CPU and GPU nodes, and also integrate seamlessly with the “Knights Landing” generation of Phi nodes. Because published performance specs for OPA were not readily available at the time of the RFP, we focused the fabric acceptance criteria mainly on meeting straightforward results from synthetic tests (e.g., Ohio State University MPI Micro-Benchmarks [14]) and on providing stable run-to-run consistency for several well-characterized MPI-intensive user applications.

The choice of parallel filesystem between IBM SpectrumScale (previously known as GPFS) and Intel Lustre was left to the proposing vendors, with a general consensus that Lustre would not likely meet our requirements for small-file metadata performance. Acceptance test criteria for the scratch storage were straightforward: >21 GB/s parallel bandwidth as measured with IOR [7], and >10K file creations per second and >6K file deletions per second in a directory containing at least 8M files, as measured with mdtest [10].

Specifying not-yet-available technologies like OPA and “Knights Landing” Phi added an element of risk to both the expected performance of Summit and its installation timeline. Despite repeated assurances from prospective vendors regarding availability dates

for OPA, delivery slippage did add about four months to the planned deployment date. In addition, while the OPA fabric’s overall performance has been excellent, successfully integrating it into Summit’s scratch storage system (along with appropriately configuring and testing the resulting parallel filesystem performance) required an additional two months of work.

3 ARCHITECTURE OF THE RMACC SUMMIT SUPERCOMPUTER

3.1 Compute nodes

When initially installed, RMACC Summit consisted of three compute node partitions. The majority of its floating-point performance comes from 380 Dell C6320 “general compute” nodes, each with two Intel Xeon E5-2680v3 CPUs, 128 GiB DDR4 RAM, a 200 GB SSD, and an OPA Series 100 HFI. The second partition consists of ten Dell C4130 “GPU” nodes. Each node is equipped with the same CPU, RAM, SSD, and OPA configuration as the “general compute” nodes, but each also includes two NVIDIA K80 accelerators. A third partition includes five Dell R930 “HiMem” nodes, each with four Intel Xeon E7-4830v3 CPUs, 2 TiB (2048 GiB) RAM, a local 10 x 1 TB SATA RAID-0 scratch storage array, and an OPA Series 100 HFI. Together these three compute partitions delivered approximately 323 TFLOPS (Rmax [18]).

A fourth partition composed of 20 “Phi” compute nodes, each equipped with an Intel Xeon Phi 7250 CPU, 64 GiB RAM, one 200 GB SSD, and an OPA series 100 HFI, is anticipated to be in production by May 2017. This partition was originally designed to be based on the Intel Xeon Phi 7250F CPU with integrated dual-port OPA connectivity.

These four compute-node architectures were chosen to meet the widest range of anticipated user requirements without making the cluster too heterogeneous. While it would have been possible to achieve much higher aggregate High-performance Linpack (HPL) [15] results by buying a larger fraction of many-core GPU or Phi nodes, the primary design consideration was to enable a large constellation of diverse applications—some of which are not well-parallelized—to port straightforwardly to Summit; today conventional CPUs still best meet that need. Similarly, specifying more nodes with less RAM per node would also boost a potential TOP500 score but would hinder real-world workflows that are increasingly data- and memory-intensive.

Some of our most memory-intensive HPC applications come from the life sciences, especially in the fields of genetics and genomics. Thus, an initial specification of 1 TiB of RAM in the high-memory nodes was doubled in the final configuration to meet anticipated demands from CSU’s existing research programs in biological and agricultural sciences.

The incorporation of a “Knights Landing” Xeon Phi design represents perhaps the most exploratory component of Summit’s architecture, being both new and not specifically supported by a broad range of existing applications. However, this many-core partition, especially when considered alongside the now-traditional GPGPU partition, provides a viable testbed for comparing programming and processing models across the two competing architectures. It is also likely to provide a simpler learning environment for users making the jump to many-core programming.

In early 2017 the original deployment was augmented with 72 additional “general compute” nodes and one additional “GPU” node via “condo” buy-ins by CU Boulder and CSU researchers.

3.2 High-performance interconnect

Summit’s 100 Gb/s Omni-Path interconnect consists of 22 discrete 48-port OPA switches, with 8 switches making up the network core and 14 providing edge ports for up to 32 compute nodes. (In 2017, the edge grew to 17 switches to accommodate condo expansion.) Each edge switch is fed by 16 core uplinks for a maximum 2:1 oversubscription between otherwise nonblocking groups of up to 32 nodes.

Summit’s predecessor at CU Boulder, Janus, was equipped with a fully-nonblocking QDR InfiniBand interconnect; but experience with Janus indicated that most jobs in our environment were not wide enough to justify the additional networking cost and complexity required to provide a fully-nonblocking interconnect. As such, it was decided that 32-node nonblocking “islands” within a single edge switch would be sufficient for the large majority of production jobs.

Two Dell R430s serve as redundant subnet managers, using the default “shortestpath” routing algorithm.

3.3 Scratch storage

Janus’s scratch storage system was based on Lustre 2.1.6, which we observed to exhibit instability under metadata intensive I/O. Summit needs to support a diverse set of applications that have varied I/O access patterns and that may not be designed to use parallel filesystems. Furthermore, the scratch filesystem must tolerate resource contention from applications running simultaneously with different I/O requirements. While Lustre has since incorporated significant advancements to rebalance its performance profile toward metadata-heavy workloads, we selected the DDN SFA14K-based GRIDScaler appliance largely for its use of the GPFS filesystem, which performs particularly well under many-file workloads due to its distributed token management system.

The GRIDScaler appliance was also specifically attractive for its expected native OPA support, but this support was not ready in time for Summit deployment. Thus, four discrete NSD servers provide OPA connectivity from the compute nodes to the InfiniBand-connected storage appliance.

Summit scratch storage provides approximately 1.2 PB of filesystem storage capacity, including 2.9 TB of SSD for metadata and storage of files smaller than 3.9 KiB.

4 DEPLOYMENT AND TESTING

4.1 Overview

RMACC Summit was installed in the HPCF in July 2016. Testing of the compute nodes and OPA fabric was performed in August, and the compute/fabric part of the system passed the necessary acceptance tests by early September. Bandwidth and metadata performance testing and tuning of the scratch storage system continued through October with the entire cluster being accepted on November 1, 2016.

4.2 Performance Testing

The CPU and memory performance of individual compute nodes were validated using a combination of HPL and STREAM [9]. The “general compute” nodes averaged 762 GFLOPS with a maximum run-to-run variation of 1.7% on any given node. Those nodes’ STREAM (Triad) score averaged 116,507 with a maximum run-to-run variation of 1.8%. We also ran HPL across all nodes of each type: The 380 “general compute” nodes scored 279 TFLOPS, the 10 “GPU” nodes scored 37 TFLOPS, and the 5 “HiMem” nodes scored 7 TFLOPS, for an aggregate of 323 TFLOPS.

Summit’s OPA interconnect was validated using the OSU MicroBenchmarks [14], particularly `osu_bibw` and `osu_alltoall`. Bidirectional bandwidth between pairs of nodes on different OPA edge switches averaged 175 Gbps. Latency as measured with 2-node all-to-all at a message size of 2 averaged 1.5 μ s between “general compute” nodes on different edge switches.

The IOR bandwidth acceptance tests were performed with a GPFS filesystem block size of 8 MiB. However, such a large block size potentially wastes space if most user files are much smaller than that size. To select a production filesystem configuration, we performed a study of filesystem bandwidth as a function of GPFS block size. We used `fio` [1] (manually dispatched on multiple filesystem client nodes simultaneously) to measure data volume read or written over a fixed period of time. Results are shown in Table 1. The dropoff in I/O bandwidth between the 8 MiB and 4 MiB block sizes led us to accept potential space inefficiencies for files between 4 KiB and 8 MiB in favor of higher parallel bandwidth.

To test the scratch filesystem bandwidth available to a single, nearly ideally behaved application, we used a version of IOR [7] updated by Los Alamos National Laboratory. We performed the test with 350 Summit general compute nodes, using 8 MPI tasks per node for the IOR MPIIO API (by specifying the “-a MPIIO” flag). The access pattern was strided, with 24 segments (-s 24) at sequential offsets to a single share file. Both the transfer size and block size were 1GiB, chosen to be a multiple of the GPFS filesystem block size of 8 MiB. This translated to an aggregate file size of 67,200 GiB (a factor of 1.5 greater than client memory to avoid caching effects [6]). The file written in the write test (-w) was read for the read test (-r). We present the results for three consecutive tests in Table 2.

Initial tests of bandwidth to the scratch storage had showed performance that was well below the acceptance criteria. Analysis by our vendor partners suggested that the issue was due to a mismatch between GPFS, which relies on verbs for communication, and OPA, which natively uses Performance Scaled Messaging instead. In the interest of passing the acceptance tests in a timely way, the raw performance of the SFA14K appliance was improved through extensive tuning and also by adding enough additional NL-SAS drives to make up for the slower-than-expected node-to-storage speed. Future enhancements to the OPA software stack are expected to enhance verbs performance and thus we expect commensurate improvements in GPFS bandwidth at that time.

To evaluate the scratch filesystem metadata performance, we chose `mdtest` [10] from Lawrence Livermore National Laboratory. We configured the GPFS filesystem with a 4KiB inode size, permitting files less than about 3.9 KiB to be stored in the inode. Files

Table 1: Storage aggregate bandwidth with different GPFS block sizes

Block size (MiB)	# Clients	Read (MiB/s)	Write (MiB/s)
4	16	13,754	13,816
4	32	13,718	13,918
4	64	14,470	14,122
4	128	17,494	14,150
8	16	21,306	21,254
8	32	21,212	21,370
8	64	21,631	21,581
8	128	22,357	21,677

Table 2: Results of three IOR tests with 2800 MPI tasks across 350 nodes, using 1GiB block and transfer size, with 24 segments

IOR Operation	Max (MiB/s)	Min (MiB/s)	Mean (MiB/s)
Write	21,069	20,996	21,039
Read	21,894	21,864	21,875

smaller than this threshold are written to SSD and larger files to NSDs comprised of RAID pools of spinning NL-SAS disks. Thus, we were interested in testing three file sizes (1 KiB, 4 KiB, and 16 KiB) for an adequate description of metadata performance. Given observed user propensity to write large numbers of files to a single directory and our previous filesystem’s struggle to remain responsive during this adverse workload, we specified that mdtest write all files to a single directory (unique_dir_per_task False and -z 0). For each file size, we instructed mdtest to read the entire written file (e.g., -e 1024 for 1KiB file size) and used a stride of two between neighbor tasks (-N 2) for file statting. We chose four MPI tasks per node with a variable number of files to create, stat, and remove per task to maintain 2^{23} (8,388,608) files in the directory (e.g., -n 131072 for 64 MPI tasks on 16 nodes). We present test results for 16, 32, and 64 nodes in Table 3.

Overall cluster performance was further validated by running several well-characterized user applications to demonstrate consistency of performance between individual runs and on different segments of the cluster, though not to meet any pre-specified application performance targets. Tests by early-adopters showed MILC [16], an MPI-intensive lattice quantum chromodynamics code, scaling well to 256 nodes. Dedalus [3], a spectral partial differential equation solver, matched the performance of a NASA supercomputer using the same number and type of processors as Summit but InfiniBand interconnect.

Several weeks after acceptance testing was completed Summit experienced a serious performance regression in user application performance that was not reflected in the synthetic CPU or MPI tests used for acceptance. We eventually traced the problem to the mmsysmon daemon, a part of GPFS, which had recently been enabled automatically as part of a configuration change. mmsysmon adds just a few percent to the system CPU load, but this is apparently enough to disrupt MPI performance across an OPA fabric when

all cores are also in heavy use by a user application. We disabled mmsysmon, immediately restoring application performance; but the experience has led us to modify our future acceptance and performance validation plans to incorporate tests that fully utilize a node’s CPU and interconnect resources simultaneously, which HPL, osu_bibw, and osu_alltoall do not do in their default configuration.

4.3 Statistical Consistency Testing

New supercomputers like Summit are composed of many innovative software and hardware technologies. As a natural consequence, their results may not be bit-for-bit (BFB) identical to results generated on another machine. Baker et al. created the Community Earth System Model Ensemble Consistency Test (CESM-ECT) to provide an objective measurement of statistical consistency between new CESM simulation outputs and an accepted ensemble of CESM outputs [2]. While CESM-ECT has been used to characterize the variability of compilers, code modifications, and machine ports [11], we used it as an experimental acceptance test of Summit’s software (e.g., compiler and MPI) and hardware. Obtaining a CESM-ECT result indicating consistency would increase our confidence that the new machine was functioning properly. However, demonstrating inconsistency could help identify a problem with the machine early in its production life. Suggestive evidence for this capability can be seen in [11], where statistically inconsistent results on Argonne National Laboratory (ANL) Mira led to the identification of Fused Multiply-Add (FMA) operations as the root cause. In fact, FMA instructions have generated statistically inconsistent results on other CESM supported machines (see Table 4).

The annual results for Summit in Table 4 were produced by running sets of 30 yearly simulations of CESM with the same configuration as in [2]. We also ran 30 “ultra-fast” simulations (nine model time steps) on Summit with the same configuration as in [12]. Both the annual and ultra-fast runs were compiled with the Intel 17 and Intel MPI 17 compiler suites. The accepted ensembles were generated on the National Center for Atmospheric Research (NCAR) Yellowstone machine, and were of size 300 for the yearly simulations [11] and 350 for the ultra-fast runs [12]. The ninth time step UF-CAM-ECT (Ultra Fast Community Atmospheric Model Ensemble Consistency Test) and annual CAM-ECT (Community Atmospheric Model Ensemble Consistency Test) failure percents reported in Table 4 are computed using the Ensemble Exhaustive Test [11] [12]. While a CESM-ECT pass or fail determination can be made by as few as three experimental runs (as in [2]), we generated between 24 and 31 simulation outputs for results reported here. The failure percents are thus the number of individual test failures divided by the total number of tests (i.e., the binomial coefficient: the number of test runs choose 3) [11] [12]. This exhaustive testing yields a failure spectrum that provides more detailed information about the statistical distinguishability of the test and ensemble distributions. From Table 4, we see that Summit is well within the range of failure percents of other CESM-supported machines with FMA disabled.

Table 3: Results of mdtests for varying node counts and file sizes. An asterisk (*) indicates the results are the average of two tests

File size (KiB)	# clients	Creation (ops/s)	Stat (ops/s)	Read (ops/s)	Removal (ops/s)
1	16	24,361*	28,420*	26,185	7,626*
1	32	21,057*	60,628*	51,154*	12,208*
1	64	20,386	117,969	109,863	17,981
4	16	19,824	30,027	5,146	8,325
4	32	19,436	60,765	9,760	12,567
16	16	19,416*	30,118*	5,056*	8,118*
16	32	19,024*	60,911*	9,641*	12,701*
16	64	18,348	115,261	16,438	18,079

Table 4: Results of UF-CAM-ECT and CAM-ECT testing for CESM supported machines and Summit

Machine	FMA	UF-CAM-ECT failure (%)	CAM-ECT failure (%)
NERSC Edison	N/A	0.1	0.5
ANL Mira	enabled		47.2
ANL Mira	disabled		4.3
NERSC Cori	enabled	95.7	
NERSC Cori	disabled	9.7	
Summit	enabled	77.2	49.9
Summit	disabled	0.0	3.4

5 MULTI-SITE CHALLENGES AND SUCCESSES

5.1 Accounts, authentication, and allocations

Because Summit is primarily a joint effort between CU Boulder and CSU we wanted to maintain continuity with existing CU Boulder Research Computing accounts and authentication while incorporating accounts and authentication to support incoming CSU users. CSU was already in the processes of deploying Duo [4] for two-factor authentication, so CU Boulder installed a local instance of the Duo authentication proxy [5] to serve as a local authentication front-end for CSU accounts stored in the CSU Active Directory.

Authorized CSU accounts are realized in the local RC LDAP instance upon request and successful authentication. Multi-domain support in sssd [17] is used to provide simultaneous access to the CU Boulder and CSU accounts that exist in the RC LDAP. To prevent conflicts between the CU Boulder and CSU username namespaces, CSU users access the system using a domain-qualified username of the form `eid@colostate.edu` which, despite its unconventional form, has produced virtually no negative effects in practice.

In addition to the shared access between CU Boulder and CSU, up to 10% of Summit compute resources are available for use by other RMACC partner institutions. The possibility of supporting several dozen partner sites, each with potentially only a handful of Summit users, makes individual direct federation with all sites impractical at present. Instead, we are working to federate with the Extreme Science and Engineering Discovery Environment (XSEDE) [20] as part

of an in-development use case of the XSEDE Community Infrastructure [21]. If successful, this should allow RMACC users to access Summit using an XSEDE portal identity, which will correspond to an RC user identity of the form `portalid@xsede.org`.

To equitably allocate compute Service Units (SUs; roughly equivalent to core-hours) between the various Summit partners, we use the “Fair Tree” fairshare algorithm in the Slurm workload manager [8]. Each partner can prioritize usage by groups within their branch of the tree without affecting the proportional distribution of resources between partners.

5.2 Outreach and training

We undertook substantial outreach to other existing and potential RMACC sites to encourage them to take advantage of Summit and to inform them about its more novel features. This outreach has taken the form of presentations and announcements at the RMACC HPC symposium, advertisement on RMACC mailing lists and conference calls, and in-person visits. In-person visits have been especially effective in encouraging new and increased RMACC participation among smaller institutions in the region.

Effective use of a modern cluster like Summit requires a solid understanding of multi-node and many-core applications and workflows and thus we have provided a variety of training opportunities centered on these topics. All together, we have presented or provided access to nearly 20 hours of pedagogic sessions including symposium talks on parallelization and vectorization, vendor- and XSEDE-led tutorials on GPU programming and Deep Learning, and Meetups covering job scheduling and Xeon Phi programming.

Most of these sessions have been streamed online for the benefit of remote users.

6 CONCLUSIONS

Although at the time of writing Summit has been in full operation for just over a month, we see many promising signs for its long-term success. During its three-month early-user testing phase, nearly 60 unique users ran over 14,000 jobs and consumed more than 16M SUs. In the subsequent five weeks, participation has accelerated with over 80 users submitting more than 25,000 jobs and using 6.7M SUs. Interest in a “condo” expansion of Summit has been greater than anticipated, with 16 research groups or departments at CSU and CU Boulder contributing 73 compute nodes by early 2017.

The success of RMACC relies on a good match among members of the consortium. Ideally our strengths and capabilities will complement and supplement each other, displacing the weaknesses, idiosyncrasies, or limitations of any one site. Summit not only is a manifestation of solid existing relationships among the partner institutions, but is also strengthening the foundation of ongoing teamwork and providing impetus for future collaborative work. For example, since the award for Summit, together with the University of Utah (UU) and under the auspices of the RMACC, we have received NSF Award ACI-1659425 “CC* Cyber Team: Creating a Community of Regional Data and Workflow Cyberinfrastructure Facilitators.” This proposal is to add one Cyberteam expert at each of the three lead institutions to assist all RMACC institutions: CU Boulder - “cradle to grave” next generation HPC workflows; CSU - data management and preservation; UU - IT security, privacy, and access controls. We also have engendered numerous collaborations among our HPC researchers via Birds of a Feather activities, and have taught HPC courses across multiple institutions. The RMACC is indeed thriving and growing. Bolstered by the acquisition and deployment of Summit, it has significantly enhanced our research environments, and we firmly believe that it will continue to do so far into the future.

ACKNOWLEDGMENTS

RMACC Summit is supported in part by the National Science Foundation under awards ACI-1532235 and ACI-1532236. We thank the staff of CSU ISTeC and CU Boulder RC for the implementation of many cyberinfrastructure features that help to support RMACC Summit. We are grateful to Anna Hasenfratz, Ethan Neil, and Ben Brown for application benchmarking. Bruce Loftis and Becky Yeager coordinated and performed much of the Summit-related outreach to existing and potential regional partners. Special thanks go to the many staff at Intel, Dell EMC, and DataDirect Networks who provided very helpful collaboration during the design, installation, and testing of RMACC Summit.

REFERENCES

- [1] J. Axboe. 2017. Flexible IO Tester. (2017). Retrieved 2017-03-13 from <http://git.kernel.dk/?p=fio.git>
- [2] A. H. Baker, D. M. Hammerling, M. N. Levy, H. Xu, J. M. Dennis, B. E. Eaton, J. Edwards, C. Hannay, S. A. Mickelson, R. B. Neale, D. Nychka, J. Shollenberger, J. Tribbia, M. Vertenstein, and D. Williamson. 2015. A New Ensemble-Based Consistency Test for the Community Earth System Model (pyCECT v1.0). *Geoscientific Model Development* 8, 9 (2015), 2829–2840. DOI:<https://doi.org/10.5194/gmd-8-2829-2015>
- [3] Dedalus Collaboration. 2017. A Flexible Framework for Spectrally Solving Differential Equations. (2017). Retrieved 2017-03-13 from <http://dedalus-project.org>
- [4] Inc. Duo Security. 2017. Duo Trusted Access. (2017). Retrieved 2017-03-13 from <https://duo.com/product>
- [5] Inc. Duo Security. 2017. Duo Trusted Access. (2017). Retrieved 2017-03-13 from https://duo.com/docs/authproxy_reference
- [6] IOR. 2016. IOR Benchmark. (2016). Retrieved 2016-10-09 from <http://www.nersc.gov/users/computational-systems/cori/nersc-8-procurement/trinity-nersc-8-rfp/nersc-8-trinity-benchmarks/ior/>
- [7] IOR-LANL. 2016. IOR-LANL/ior. (2016). Retrieved 2016-10-09 from <https://github.com/IO-IO-IO/ior>
- [8] M. Jette and M. Grondona. 2003. SLURM: Simple Linux Utility for Resource Management. (2003). Retrieved 2017-03-13 from https://slurm.schedmd.com/slurm_design.pdf
- [9] J. D. McCalpin. 2017. STREAM: Sustainable Memory Bandwidth in High Performance Computers. (2017). Retrieved 2017-03-13 from <https://www.cs.virginia.edu/stream/>
- [10] MDTEST-LANL. 2016. MDTEST-LANL/mdtest. (2016). Retrieved 2016-10-22 from <https://github.com/MDTEST-LANL/mdtest>
- [11] D. J. Milroy, A. H. Baker, D. M. Hammerling, J. M. Dennis, S. A. Mickelson, and E. R. Jessup. 2016. Towards Characterizing the Variability of Statistically Consistent Community Earth System Model Simulations. *Procedia Computer Science* 80 (2016), 1589 – 1600. DOI:<https://doi.org/10.1016/j.procs.2016.05.489>
- [12] D. J. Milroy, A. H. Baker, D. M. Hammerling, and E. R. Jessup. 2017. Nine Time Steps: Ultra-Fast Statistical Consistency Testing of the Community Earth System Model (pyCECT v3.0). *Geoscientific Model Development Discussions* 2017 (2017), 1–22. DOI:<https://doi.org/10.5194/gmd-2017-49>
- [13] Energy Sciences Network. 2013. Science DMZ Implemented at CU Boulder. (2013). Retrieved 2017-03-13 from <https://www.es.net/science-engagement/case-studies/science-dmz-case-studies/science-dmz-implemented-at-cu-boulder/>
- [14] OSU-Microbenchmarks. 2016. MVAPICH::Benchmarks. (2016). Retrieved 2016-10-30 from <http://mvapich.cse.ohio-state.edu/benchmarks/>
- [15] A. Petitet, R. C. Whaley, J. Dongarra, and A. Cleary. 2016. HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers. (2016). Retrieved 2016-02-24 from <http://www.netlib.org/benchmark/hpl/>
- [16] MILC Research Program. 2017. MIMD Lattice Computation (MILC) Collaboration. (2017). Retrieved 2017-03-13 from <http://physics.indiana.edu/~sg/milc.html>
- [17] Fedora Project. 2017. System Security Services Daemon. (2017). Retrieved 2017-03-13 from <https://fedorahosted.org/sssd/>
- [18] TOP500.org. 2017. TOP500 Description. (2017). Retrieved 2017-05-10 from https://www.top500.org/project/top500_description/
- [19] H. M. Tufo, M. K. Patterson, M. Oberg, M. Woitaszek, G. Cobb, R. Strong, and J. Gutowski. 2011. Janus: Co-designing HPC Systems and Facilities. In *State of the Practice Reports (SC '11)*. ACM, New York, NY, USA, Article 16, 9 pages. DOI:<https://doi.org/10.1145/2063348.2063370>
- [20] XSEDE. 2017. Extreme Science and Engineering Discovery Environment. (2017). Retrieved 2017-03-13 from <https://www.xsede.org/>
- [21] XSEDE. 2017. XSEDE Community Infrastructure Use Cases. (2017). Retrieved 2017-03-13 from <https://software.xsede.org/use-case-area/community-infrastructure>