The Sol Supercomputer at Arizona State University

DOUGLAS M. JENNEWEIN, JOHNATHAN LEE, CHRIS KURTZ, WILLIAM DIZON, IAN SHAEFFER, ALAN CHAPMAN, ALEJANDRO CHIQUETE, JOSH BURKS, AMBER CARLSON, NATALIE MASON, ARHAT KOBAWALA, THIRUGNANAM JAGADEESAN, PRAFUL BHARGAV BASANI, TOREY BATTELLE, REBECCA BELSHE, DEB MCCAFFREY, MARISA BRAZIL, CHAITANYA INUMELLA, KIRBY KUZNIA, JADE BUZINSKI, DHRUVIL SHAH, SEAN M. DUDLEY, GIL SPEYER, and JASON YALIM†, Arizona State University, USA

The Sol supercomputer provides ASU researchers access to a state-of-the-art system with an observed GPU-only HPL speed of 2.272 PetaFLOP/s. This short paper provides a motivation for the supercomputer as well as a technical reference of its many components, in the hopes that the implicit template will encourage peer institutions to also generate corresponding publications and foster discussion. The supplemental git repository [2] contains finer details on the system's configuration and has the ability to be updated as the system evolves.

CCS Concepts: • **Social and professional topics → User characteristics**; **Management of computing and information systems**; **Computing profession**; • **Applied computing**; • **Mathematics of computing**; • **Hardware**; • **Computer systems organization → Parallel architectures**; • **Networks → Network architectures**;

Additional Key Words and Phrases: cyberinfrastructure, supercomputer design, system design, supercomputing implementation, high-performance computing, top500 HPC

ACM Reference Format:
Douglas M. Jennewein, Johnathan Lee, Chris Kurtz, William Dizon, Ian Shaeffer, Alan Chapman, Alejandro Chiquete, Josh Burks, Amber Carlson, Natalie Mason, Arhat Kobawala, Thirugnanam Jagadeesan, Praful Bhargav Basani, Torey Battelle, Rebecca Belshe, Deb McCaffrey, Marisa Brazil, Chaitanya Inumella, Kirby Kuznia, Jade Buzinski, Dhruvil Shah, Sean M. Dudley, Gil Speyer, and Jason Yalim. 2023. The Sol Supercomputer at Arizona State University. In *Practice and Experience in Advanced Research Computing (PEARC '23), July 23–27, 2023, Portland, OR, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3569951.3597573

1 INTRODUCTION

Following a community practice established by computational teams that have built resources at the national scale, the intention of this report is to establish a process for centralized compute facilities at research institutions to describe the specifications, implementation, and benchmarking of these instruments [1, 16, 20, 21, 23, 25]. Such a report would benefit (1) the computational research community by informing about the successes and pitfalls of standing up a petaflop-scale cluster configuration, (2) the home institution's user community by providing a convenient means to cite the cluster in publications and (3) broader research investigations, by providing a template for new HPC systems with significant investment and simplifying, with an aggregation of these published records, subsequent analyses surveying these instruments' architectures and impact.

2 MOTIVATION

Arizona State University operates campus HPC systems funded through periodic university investments that establish a base amount of free computing capacity and sufficient headroom in the system's interconnect fabric, shared parallel filesystem, liquid cooling capacity, rack space, and power distribution to accommodate researcher-purchased "condo" nodes. So-called condo nodes are researcher-purchased compute nodes incorporated into a shared campus cluster, prioritizing their owner's compute jobs, with any unused cycles offered opportunistically to the rest of campus.

By 2020 Sol's predecessor had outgrown its network fabric and shared filesystem (both had been upgraded more than once during the system's lifetime), and most of the system's GPU capacity resided in researcher-purchased "condo" nodes. The opportunistic capacity provided by GPU-equipped condo nodes ultimately proved insufficient in meeting the demands of GPU-hungry ASU researchers. Thus, Sol was built to provide (1) sufficient general-purpose GPU capacity to address current computing needs and (2) sufficient expansion capacity to accommodate three to five years of condo nodes.

3 OVERVIEW

At its release, the Sol supercomputer provides ASU researchers with free access to 17 904 Zen3 cores, 224 80GiB A100 and 15 24GiB A30 Graphics Processing Units (GPUs), 104 TB of RAM, all over an InfiniBand network with access to 4 PiB of policy-enforced temporary scratch space and 100 GiB of persistent home storage.

4 INFRASTRUCTURE

Sol is housed in an off-campus Tier 3+ commercial datacenter operated by Iron Mountain (NYSE:IRM) and provides FISMA High security. The caged space has room for up to eighty-racks. Sol currently consists of twenty APC 42U-racks with CoolIT liquid-chilled doors and liquid cooling for compute and high-memory nodes. Cooling is provided by three CoolIT rack DCLC CHx750 chillers which have redundant pumps. One rotating chiller is always on standby. Extensive under-tile plumbing had to be completed before bringing any systems online, requiring careful coordination between the on-site team, datacenter, and subcontracted installers. All initial configuration, cabling, and labeling was performed by Dell. All

PEARC '23, July 23–27, 2023, Portland, OR, USA

ACM ISBN 978-1-4503-9985-2/23/07…$15.00
https://doi.org/10.1145/3569951.3597573

equipment is on a five-year warranty. NetBox is used for inventory management.

Virtualization is provided by the Citrix Xen Hypervisor. Critical systems and services are virtualized (e.g., the administration node) to allow for daily or arbitrary snapshotting and ease of restoration. Redundant Dell PowerStore block storage devices provide virtual machine (VM) disks and are replicated between Iron Mountain and the on-campus data center (four miles away in Tempe). Daily system and volume snapshots are kept up to seven and thirty days, respectively. XenOrchestra allows researchers and administrators to manage VMs over a web interface.

The datacenter provides up to one megawatt of redundant power, fully sourced from renewable wind or solar energy (via Iron Mountain's, "Green Power Pass"). As configured, Sol's server racks are supplied with either two (CPU-only racks) or four (GPU racks) 400 volt / 60 amp feeds, dependent on hardware power requirements.

Sol connects directly to the ASU research network, which routes to both Internet2 and a commodity network over redundant 100 Gib/s links. Internet2 is the primary path and is provided by SunCorridor's Research and Educational Network. Cox Communications provides the commodity service for ASU which serves as a backup connection for Sol.

5 STORAGE

User home directories are hosted on a 2 PiB Isilon array. Researchers are provided 100 GiB allocations that are snapshotted daily (rolling 30-day window).

Scratch storage is provided by a 4 PiB BeeGFS array of Dell ME4084 devices configured by System Fabric Works. Four object storage servers point to eight ME4084 storage enclosures. Two metadata servers and one management server all point to one all-flash SSD ME2024. The BeeGFS service is available over both the 200 Gib/s InfiniBand and 25 Gib/s Ethernet interfaces. BeeGFS data are subject to a Starfish-enforced ninety-day retention policy [3, 22].

Researchers also have access to an existing project-based storage offering that is housed within the on-campus Tempe datacenter four miles away from Sol, connected via a 100 Gib/s link over private, 2ms-latency fiber. This service is currently provided to researchers for $50/TiB/year. Project-based storage resides on two storage arrays: a 4 PiB Isilon and a 3 PiB Qumulo. Snapshots are taken daily and stored for up to thirty-days.

Finally, an archive service is provided to researchers for $10/TiB/year, provided by a 2 PiB Spectra Logic LTO-tape array with a BlackPearl object storage frontend. The array is housed in Iron Mountain and is accessible through Globus [6].

6 COMPUTE NODE SPECIFICATIONS

All Sol compute nodes are dual-socket with hyperthreading disabled, eight non-uniform memory access (NUMA) regions that map sequentially to the core indices, and dual-channel DDR4-3200 PC4-25600 ECC RAM manufactured by Hynix or Samsung. All compute nodes also have Mellanox MT27800 ConnectX-5 and MT28908 ConnectX-6 providing 25 Gib/s (Ethernet) and 200 Gib/s (4x HDR InfiniBand) respectively. All nodes have at least 16x 32GiB (540 GB total) of RAM, the exception being the *high-memory* nodes with 32x 64GiB (2.2 TB total). Sol may be subdivided by several homogeneous compute partitions: 112 *CPU-only* nodes with AMD EPYC 7713 Zen3 (base clock 2 GHz, boost 3.6 GHz, 128 physical cores total) [Dell PowerEdge C6525], 5 *high-memory* nodes with the same characteristics as the *CPU-only* nodes but four-times more memory [Dell PowerEdge R6525], 56 *A100* nodes with AMD EPYC 7413 Zen3 (base clock 2.65 GHz, boost 3.6 GHz, 48 physical cores total) and four NVIDIA 80GiB A100 SXM4 graphics processing units (GPUs) [Dell PowerEdge XE8545], and 5 *A30* nodes with AMD EPYC 7413 Zen3 (base clock 2.65 GHz, boost 3.6 GHz, 48 physical cores total) and three NVIDIA 24GiB A30 PCIe GPUs [Dell PowerEdge R7525].

In aggregate, Sol's compute nodes provide 17 904 Zen3 cores (~84% EPYC 7713), 224 NVIDIA 80GiB A100 SXM4 GPUs, 15 NVIDIA 24GiB A30 PCIe GPUs, and 104 TB of DDR4 RAM. At least one *A100* node is run with NVIDIA's Multi-Instance GPU (MIG) enabled, providing seven slices per GPU (slices of one streaming multiprocessor with 9.5 GiB). A development environment has four additional *CPU-only* nodes. The *A30* nodes and five of the *A100* nodes are part of AFORCE (NSF OAC-2126303): one of each will be dedicated to external use which includes the Open Science Grid and a tri-university initiative (with Northern Arizona University and the University of Arizona).

7 DATA TRANSFER NODES AND SERVICES

To enable large data transfers and host the Globus service, Sol has a dedicated data transfer node (DTN). The physical node has the same characteristics as the *CPU-only* compute nodes, but has 100 Gib/s Ethernet and no InfiniBand. The system is positioned behind a pair of 100 Gib/s firewalls, and connectivity is provided via static NAT. Traffic is monitored but stateful packet inspection is disabled to prevent performance loss. The DTN mounts the scratch filesystems on Sol and its predecessor to facilitate migrations between the two supercomputers. Researchers may access the DTN explicitly through SSH or implicitly via Globus. ASU's Globus endpoint includes the Spectra Logic plugin for accessing the long-term tape-storage system, and the Google Drive plugin.

8 NETWORKING

Sol is constructed with two-tier spine-leaf networks for the various fabrics. The HDR 200 Gib/s InfiniBand network consists of Mellanox 40-port QM8700 series switches and was designed with a 2.3:1 oversubscription rate. There are twelve spine and thirteen leaf switches. The twelve spines will allow for an additional twenty-seven leaf switches to enable future expected growth.

The frontend Ethernet network consists of two Dell Z9264 64-port spine-switches and thirteen Dell S5248F-ON 48-port leaf-switches (all running OS10). Each leaf switch provides a 25 Gib/s link to compute nodes and has a 100 Gib/s link back to each spine switch.

The backend management network has two Dell S5248F-ON spine-switches and thirteen Dell S3048-ON 48-port leaf-switches. Each leaf switch provides a 1 Gib/s Ethernet link to compute nodes and has a 10 Gib/s link back to each spine switch. The management network also supports out-of-band systems, i.e., IPMI, PDUs, chillers, and the cooled doors, with room for future sensors or systems.

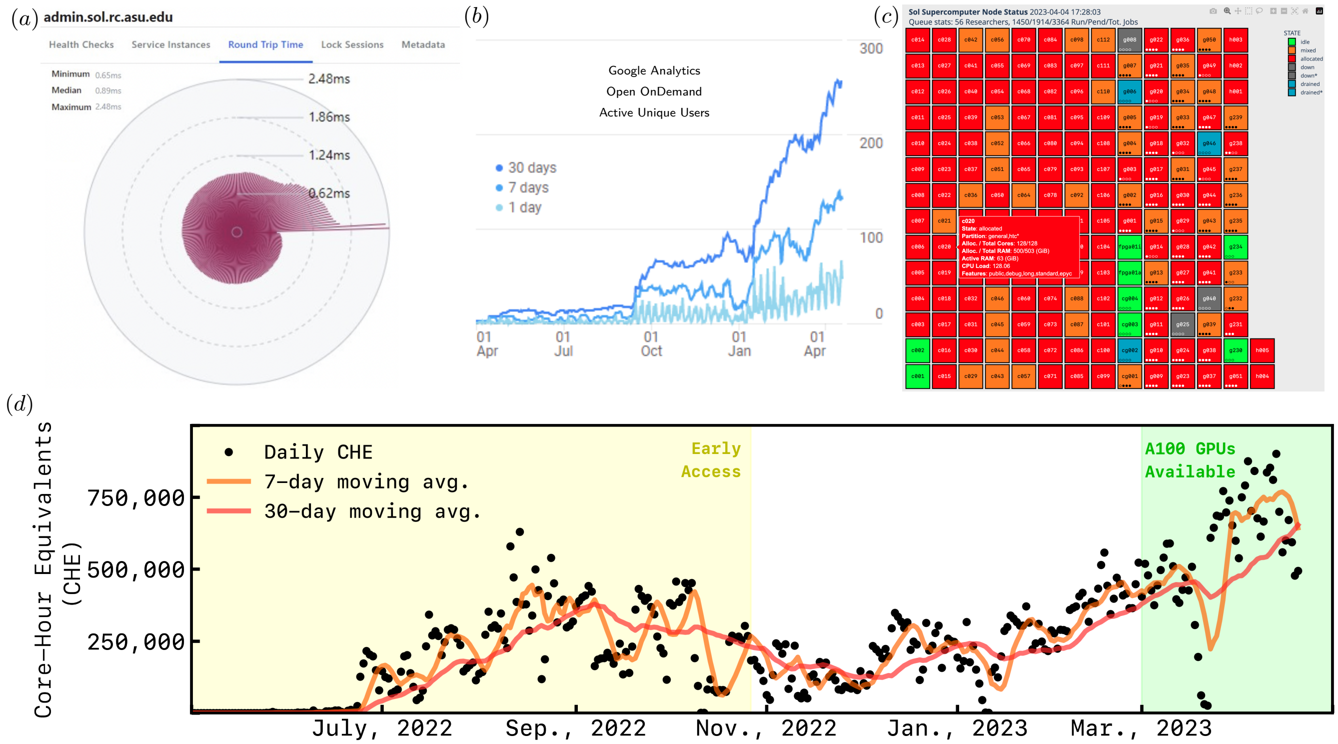

Fig. 1. (a). Example visualization from Consul illustrating the difference in latency between hosts on the two supercomputers housed in separate datacenters (4 miles apart). (b). Google Analytics traffic on Sol's Open OnDemand (OOD) nodes. The curves are unique users over a rolling time period. (c). An example of Sol's status page, as part of the OOD web portal. Each square is a compute node; color indicates SLURM status. Hovertext provides node details. Filled or empty circles indicate allocated or idle GPUs, respectively. (d). Sol compute-hour equivalent (CHE) usage as defined by equation (1).

A layer 3 boundary is provided by two 100 Gib/s Fortinet FortiGate-3401E firewalls. These also facilitate VLAN-to-VLAN communication and connectivity back to the on-campus datacenter (four miles away from Sol) via private virtual routing and forwarding (VRF).

9 SOFTWARE

Images are deployed via PXE by Cobbler [4]. SaltStack [19] is used to deploy and maintain system state. Salt is a client-server configuration management software, with clients on every node monitoring configuration changes. The main salt server runs on a virtualized administration node, and the salt states are version-controlled by git (via an independent but local gitlab instance).

Sol runs Rocky Linux 8.x (chosen over Alma and RedHat after discussions with the HPC community) with the Mellanox OFED 5.x drivers. Spack 0.18 [7] is used to provide the majority of researcher software lmod 8.x modules. Apptainer 1.x is maintained by salt outside of the module system, as is rclone 1.x. Python package management is done outside of spack with mamba 1.x [14] (chosen for speed and license), with a semi-annual update schedule. Mamba was patched to prevent messages about `conda activate` and ultimately `conda init`, which causes troublesome shell configuration injections. Outside of global python environments which are installed to system-mounted `/packages/envs` and Open OnDemand-powered Jupyter lab 3.x with ipywidget and voila support, researchers are expected to create their own environments.

Open OnDemand 2.x [9] is maintained by salt, and provides simplified access to salt or admin maintained VirtualGL enabled virtual desktops, Jupyter lab servers, MATLAB, and Rstudio Server (semi-annual updates) [12, 18].

10 SCHEDULING

Sol runs SLURM 23.x [26] in a headless configuration and compiled with PMIx support. The control daemon, `slurmctld`, has high-availability and failover enabled and runs on a dedicated virtual machine (VM). A separate VM is dedicated for the MariaDB backend. Sol is free for researchers and job priorities are enforced through SLURM's Fairshare, which does not have fairtree and is depth oblivious. One hour on a GPU is tracked as 20-25 core-hour equivalents (CHE) depending on the GPU model (A30-A100). One hour of a Multi-Instance GPU (MIG) slice of an A100 is charged as three CHE. Every hour of 4 GiB of memory allocation is tracked as one CHE of usage. CHE are then tracked as

$$CHE = (cores) + (4GiB\,slices) + 3\times(MIGs) + 20\times(A30s) + 25\times(A100s), \quad (1)$$

Figure 1(d) shows Sol's CHE usage since its inception, which is heavily discounted when considering the raw compute ability of the A100 relative to a CPU core (based on LINPACK, one A100 is ~500 times faster than a theoretical EPYC 7713 core). Additionally, MIG slices (provided only in the finest subdivision, roughly one-seventh of an A100) are further discounted. This discount helps motivate researchers to migrate workflows to the GPU.

The SLURM account hierarchy was motivated by XDMoD requirements (to provide faculty sponsors with an aggregate report on their students' utilization) and influenced Linux groups. Internal users are within their sponsoring faculty's group and thus account, which are under the university department, collected into "groups," and then registered as internal, that is, ASU/groups/department/lab account/user. Adjacent to "groups" is a class hierarchy for Academic course support and workshops. Adjacent to "ASU" is an external hierarchy for OSG support.

Sol's partitions are kept as generic as possible to match researcher use cases. The SLURM job submit plugin allows highly precise hardware allocations with minimal user job specification effort. Node features are used to describe node hardware and also further constrain jobs (e.g., to not land on a condo 40GiB A100 instead of a public 80GiB A100). The four partitions are general, HTC (High Throughput Computing), highmem, and lightwork. General and HTC contain all compute nodes sans *high-memory* and two *CPU-only* nodes oversubscribed and dedicated to the lightwork partition. Highmem contains the five *high-memory* nodes. General and highmem allow jobs with wall times for up to one week (by regular QOS access, up to two weeks with special "long" QOS access). Lightwork, which is oversubscribed and there to support classes and light interactive sessions, allows jobs with wall times for up to one day. HTC will only accept jobs that request four hours or less of wall time, but these jobs will run without interruption over the partition, which includes faculty purchased condo nodes. These condo nodes are also available in the general partition, but opportunely (jobs may be cancelled, via preemption, to make room for the purchasing lab's work).

Typical researchers will use one of four main Quality of Services (QOS): public, long, private, and debug. The default is public, which works as described in the previous paragraph. Vetted researchers may use the long QOS to increase wall times to up to two weeks. The private QOS enables up to one week of opportunistic access on faculty purchased condo nodes (that is, faculty and their labs may preemptively cancel a private-QOS job to allocate on their own purchased hardware, and these labs have a special QOS to access their nodes in this way). Finally, the debug QOS boosts priority on up to two jobs, for up to two nodes, for up to fifteen minutes of wall time.

Backfilling is a primary job allocation mechanism, and thus it was important to ensure that the backfilling window was at least as long as the longest possible job (fourteen days).

The `text/affinity` plugin is enabled, which allows for core binding at the job level. The select type parameter, `CR_CORE_DEFAULT_DIST_BLOCK`, was enabled to improve the default performance of non-I/O bound compute-work by up to a factor of two.

Additional details are provided in the supplementary materials [2], including the output of `scontrol show config` and the `job_submit.lua` source code.

11 USER MANAGEMENT

Sol has its own independent LDAP, which probes a university-maintained active directory service to obtain UIDs and GIDs. Linux groups associate a user to their sponsoring faculty and are tied to user SLURM accounts and upstream reporting tools like XDMoD. User access is reviewed annually. Class accounts are time-limited to a semester plus a short grace period.

12 MONITORING AND ALERTING

Grafana [8] is used to visualize and alert on a variety of system-wide data. The data are provided by different tools. Node metrics and status are collected by Prometheus [17] and `node_exporter`. Telegraf [10] is used to collect GPU and SNMP data (i.e., from the three CoolTerra chillers, the network devices, and PDUs) which is stored by InfluxDB. Grafana alerts are sent to Victorops and Slack for critical system failures.

Elasticsearch is used with Filebeat for log collection. Consul is used to monitor latency between systems. An example view is shown in figure $1(a)$. Finally, node health check (nhc) [13] performs regular system checks to take compute nodes offline when issues arise.

13 RESEARCHER INTERFACES

In addition to the third-party interface to Sol's filesystems via Globus, researchers typically access Sol by first connecting to a university-provided virtual private network, which has two-factor authentication with Duo, and then using their ASURITEs (university-provided credentials) may access the Sol supercomputer via `sol.asu.edu` through a shell (traditional `ssh`) or a web browser (Open OnDemand, [9]). The web portal is very popular with researchers and students. At the time of this writing, Sol had 349 users, and as figure $1(b)$ indicates, the web portal received up to 60, 140, and 260 unique users over a 1-, 7-, and 30-day period respectively. The web portal provides a custom status page, as illustrated in figure $1(c)$ and also links an XDMoD instance for rich usage statistics [15]. Popular interactive applications, e.g., Virtual Desktops, Jupyter, MATLAB, and RStudio are facilitated through the web portal.

14 BENCHMARKS

During a March maintenance, two benchmarks were conducted with NVIDIA's HPC-X LTS suite for MPI over the InfiniBand: high-performance linpack (HPL) [5] and IO500. The HPL benchmark is a long-standing metric for quantifying the capabilities of a supercomputer, and was ultimately run across 52 *A100* nodes to obtain a speed of 2.272×10^{15} 64-bit floating-point operations per second (FLOP/s), with a passed residual (value of 6.4×10^{-6}) and parameters WR03L2L2, $N = 1469440$, $NB = 224$, $P = 13$, and $Q = 16$. Care was taken to match the core and memory affinity of the GPUs as reported by `nvidia-smi topo -m`. This result has placed Sol at #388 on the Top500 June 2023 list [24]. For comparison, a single *A100*-node with all four SXM4 GPUs achieved a maximum speed of ~5.7×10^{13} FLOP/s, suggesting a 52-node theoretical speed of ~3×10^{15} FLOP/s and ~25% inefficiency likely due to network losses.

The IO500 benchmark was run on ten InfiniBand-adjacent *CPU-only* nodes with care to bind every eight-core Core Cache Complex

(CCX) to an independent MPI task (resulting in 160 tasks with 16-tasks per node). With a stonewall-time of 300, the overall score was 16.48, with 61.76 IOPS and a bandwidth of 4.4 GiB/s. As a result, Sol ranked 66$^{\text{th}}$ on the June 2023 ten-node IO500 list [11].

15 RESEARCHER ENGAGEMENT

When first announcing that a new supercomputer was purchased, a naming contest was conducted and the person that submitted the winning name, "Sol," won a new Dell laptop [27].

Beginner tutorials are provided on a monthly basis in either an in-person, hybrid, or online-only mode. In total, fifty trainings and one expo are conducted annually on introductory to advanced topics. A day-long "GPU Day" was conducted in April to on-board fifty-attending researchers to the newly stood-up Sol GPUs.

In May of 2023, Sol hosted the 2023 Rocky Mountain Advanced Computing Consortium HPC Symposium in Scottsdale, Arizona. In preparation for the May 16-18 event, a new temporary account system was created. Three transient open ondemand instances and five-hundred temporary accounts were created. Attendees were given access to ten "MIGified" *A100* nodes which supplied a dedicated pool of 280 GPU slices. The ability to generate temporary accounts will be carried over to new workshops, tutorials, and trainings.

16 CONCLUSION

Researchers are able to buy into the supercomputer by purchasing Research Computing (RC) approved blades that are racked in an RC-provided chassis. These compute nodes are available to the supercomputing community for high-throughput computing for no more than four hours or opportunely (pre-emptable) for up to seven days. This model was based on the success of the previous flagship supercomputer, Agave. In fact, many of the previous lessons learned from Agave were used to motivate the many decisions involved with standing up Sol but, prior to this report, remained esoteric among our team.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Grant Numbers OAC-2018886 and OAC-2126303. Thanks to Dell, Iron Mountain, System Fabric Works, CoolIT, and Comfort Systems for help with the delivery, install, and the setup of the Sol supercomputer.

REFERENCES

[1] Jonathon Anderson, Patrick J. Burns, Daniel Milroy, and et al. 2017. Deploying RMACC Summit: An HPC Resource for the Rocky Mountain Region *(PEARC17)*. Assoc. for Comp. Mach., New York, NY, USA, Article 8, 7 pages. https://doi.org/10.1145/3093338.3093379

[2] ASU. 2023. Sol. https://github.com/ASU-KE/Sol https://github.com/ASU-KE/Sol.

[3] BeeGFS. 2023. BeeGFS. https://git.beegfs.io/pub https://git.beegfs.io/pub.

[4] Cobbler. 2023. Cobbler. https://github.com/Cobbler/Cobbler https://github.com/Cobbler/Cobbler.

[5] Jack J. Dongarra, Piotr Luszczek, and Antoine Petitet. 2003. The LINPACK Benchmark: past, present and future. *Concurrency and Computation: Practice and Experience* 15, 9 (2003), 803–820. https://doi.org/10.1002/cpe.728

[6] I. Foster. 2011. Globus Online: Accelerating and Democratizing Science through Cloud-Based Services. , 70–73 pages.

[7] Todd Gamblin, Matthew LeGendre, Michael R. Collette, and et al. 2015. The Spack Package Manager: Bringing Order to HPC Software Chaos *(SC '15)*. Assoc. for Comp. Mach., Article 40, 12 pages. https://doi.org/10.1145/2807591.2807623

[8] Grafana. 2023. Grafana. https://github.com/grafana/grafana https://github.com/grafana/grafana.

[9] Dave Hudak, Doug Johnson, Alan Chalker, and et al. 2018. Open OnDemand: A web-based client portal for HPC centers. *J. of Open Source Software* 3, 25 (2018), 622. https://doi.org/10.21105/joss.00622

[10] Influxdata. 2023. Telegraf. https://github.com/influxdata/telegraf https://github.com/influxdata/telegraf.

[11] IO500. 2023. 10-Node ISC23 IO500 List. https://io500.org/list/isc23/ten

[12] Jupyter Lab. 2023. Jupyter Lab. https://github.com/jupyterlab/jupyterlab https://github.com/jupyterlab/jupyterlab.

[13] LLBL Node Health Check (nhc). 2023. LLBL Node Health Check (nhc). https://github.com/mej/nhc https://github.com/mej/nhc.

[14] Mamba Org. 2023. mamba. https://github.com/mamba-org/mamba https://github.com/mamba-org/mamba.

[15] Jeffrey T. Palmer, Steven M. Gallo, Thomas R. Furlani, and et al. 2015. Open XDMoD: A Tool for the Comprehensive Management of High-Performance Computing Resources. *Computing in Science & Engineering* 17, 4 (2015), 52–62. https://doi.org/10.1109/MCSE.2015.68

[16] Marcelo Ponce, Ramses van Zon, Scott Northrup, and et al. 2019. Deploying a Top-100 Supercomputer for Large Parallel Workloads: The Niagara Supercomputer *(PEARC '19)*. Assoc. for Comp. Mach., New York, NY, USA, Article 34, 8 pages. https://doi.org/10.1145/3332186.3332195

[17] Prometheus. 2023. Prometheus. https://github.com/prometheus/prometheus https://github.com/prometheus/prometheus.

[18] Rstudio. 2023. Rstudio. https://github.com/rstudio/rstudio https://github.com/rstudio/rstudio.

[19] Salt Stack. 2023. Salt. https://github.com/saltstack/salt https://github.com/saltstack/salt.

[20] Dan Stanzione, Bill Barth, Niall Gaffney, and et al. 2017. Stampede 2: The Evolution of an XSEDE Supercomputer *(PEARC17)*. Assoc. for Comp. Mach., New York, NY, USA, Article 15, 8 pages. https://doi.org/10.1145/3093338.3093385

[21] Dan Stanzione, John West, R. Todd Evans, and et al. 2020. Frontera: The Evolution of Leadership Computing at the National Science Foundation *(PEARC '20)*. Assoc. for Comp. Mach., New York, NY, USA, 106–111. https://doi.org/10.1145/3311790.3396656

[22] StarfishStorage. 2023. Starfish. https://github.com/StarfishStorage https://github.com/StarfishStorage.

[23] Shawn Strande, Haisong Cai, Mahidhar Tatineni, and et al. 2021. Expanse: Computing without Boundaries: Architecture, Deployment, and Early Operations Experiences of a Supercomputer Designed for the Rapid Evolution in Science and Engineering *(PEARC '21)*. Assoc. for Comp. Mach., New York, NY, USA, Article 47, 4 pages. https://doi.org/10.1145/3437359.3465588

[24] Top500. 2023. June List. https://www.top500.org/lists/top500/2023/06 Sol #388.

[25] Verónica G. Vergara Larrea, Wayne Joubert, Michael J. Brim, and et al. 2019. Scaling the Summit: Deploying the World's Fastest Supercomputer. In *ISC High Performance 19*. Springer International Publishing, 330–351.

[26] Andy B. Yoo, Morris A. Jette, and Mark Grondona. 2003. SLURM: Simple Linux Utility for Resource Management. In *Job Scheduling Strategies for Parallel Processing*, Dror Feitelson, Larry Rudolph, and Uwe Schwiegelshohn (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 44–60.

[27] Pete Zrioka. 2022. A new dawn for supercomputing at ASU. *ASU Knowledge Enterprise Research News* (23 May 2022). https://research.asu.edu/new-Sol-supercomputer-ASU