WELCOME TO

CUG
Cray User Group

CUG 2025 WELCOME

Dear Cray User Group Colleagues,

The Institute for Advanced Computational Science (IACS) at Stony Brook University (SBU) and the HPE Cray User Group (CUG) Board are thrilled to extend a warm invitation to all members of CUG to the 2025 CUG Conference, from May 4 – May 8, 2025, within minutes of New York City, the big apple – a city that is always open for adventure.

The theme for the 2025 event is Computing Horizons. With this theme, we want to capture the sense of the great opportunities arrayed in front of us, with the tight integration of traditional HPC workloads and powerful AI enabled by innovations in algorithms and special purpose hardware. The 2025 conference is about this bright future and the exploration of these limitless possibilities and their associated challenges.

CUG serves as an annual gathering for experts, researchers, and enthusiasts from the computing community to come together, exchange ideas, and explore the latest advancements in our field. The CUG2025 event promises to be an exciting and enriching experience featuring a diverse range of keynote presentations, panel discussions, workshops, and networking opportunities.

Since 2019, IACS at SBU in partnership with the Center for Computational Research (CCR) at the University of Buffalo, has been operating Ookami, a Cray Apollo 80 system. Ookami is the Japanese word for Wolf, and an homage to the mascot of Stony Brook University, Wolfie Seawulf, and the Japanese origin of the Fujitsu A64FX processor. This testbed is supported by the US National Science Foundation (NSF) and provides researchers worldwide with access to the A64FX processor which also powers the famous Fugaku Supercomputer.

SBU is honored to host this event, which is surely an annual highlight of all CUG members. We look forward to welcoming you to the NYC area!

Sincerely,

Dr. Robert J Harrison,

Director of the Institute for Advanced Computational Science (IACS)
at Stony Brook University

Stony Brook University

Invited Speaker
Keynote

Michael Zingale

Department of Physics and Astronomy Professor

Michael Zingale is a Professor of Physics and Astronomy at Stony Brook University. He earned a BS in Physics and Astronomy (1996) from the University of Rochester and a PhD in Astronomy and Astrophysics (2000) from the University of Chicago. He was part of the Flash Code development team that won a Gordon Bell Prize in 2000, and received a Presidential Early Career Award for Scientists and Engineers (PECASE) through DOE NNSA in 2005, and an Outstanding Junior Investigator award for the DOE Office of Nuclear Physics in 2006. Michael's research involves the development of new algorithms for efficiently modelling convection in stellar interiors and coupling nuclear reactions and hydrodynamics in explosive flows. He is a codeveloper of the Open Source astrophysical, multiscale/multiphysics simulation codes Castro and MAESTROeX and the nuclear astrophysics library pynucastro. He applies these codes to studies of early phases of Type Ia supernovae, novae, and X-ray bursts.

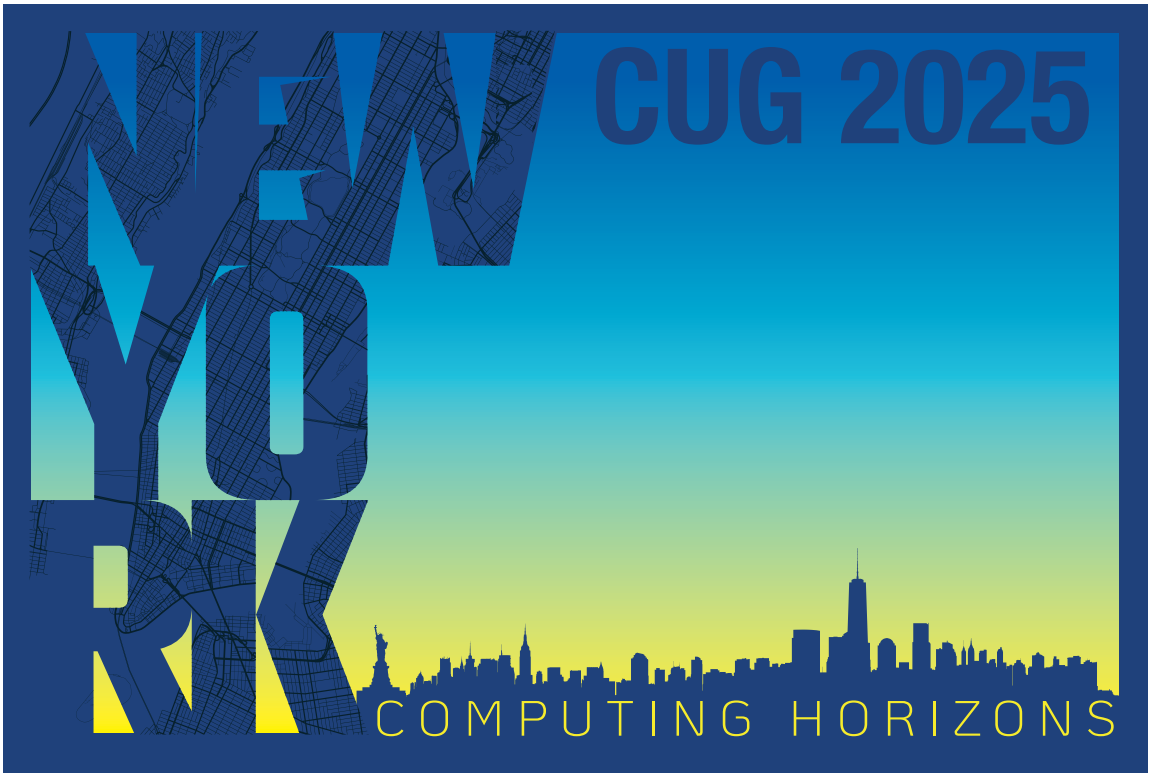

Sessions and Abstracts

Sunday, May 4th

8:30am-10:10am

XTreme (Approved NDA Members Only) **Newport III**

10:00am-10:30am

Coffee Break **Newport Foyer**

10:30am-12:00pm

XTreme (Approved NDA Members Only) **Newport III**

12:00pm-1:00pm

Lunch (open to PEAD and XTreme participants) **Newport Foyer**

1:00pm-2:30pm

Programming Environments, Applications, and Documentation (PEAD)

Newport I

Description

The PEAD (Programming Environments, Applications, and Documentation) is a CUG Special Interest Group that provides a forum for discussion and information exchange between CUG sites and Cray/HPE. The group focus includes system usability, performance of programming environments (including compilers, libraries, and tools), scientific applications running on Cray/HPE systems, user support, communication, and documentation.

The group host meetings at CUG each year to help foster discussions surrounding these topics between HPE and member sites. Following a successful event at last year's CUG, this year the PEAD SIG will meet Sunday, May 05, from 1:00 PM - 5:00 PM. We are planning topics surrounding the HPE PE roadmap, training collaborations, HPE documentation, as well as Fortran support. All topics will be interactive and discussion based. Registration for the event is required. Lunch will be available for everyone who registers for the meeting.

PEAD Introduction

1pm - 1:05pm

Chris Fuson (Oak Ridge National Laboratory)

CPE in a Container

1:05pm - 1:45pm

Kaylie Anderson (HPE), Ben Cumming (Swiss National Supercomputing Centre), Subil Abraham (Oak Ridge National Laboratory), and Panchapakesan Chitra Shyamshankar (Argonne National Laboratory)

Abstract

Containers provide a way to package applications along with their dependencies in a single unit. Containers can aid common HPC workflow processes such as building, testing, and execution. Many CUG centers are expanding use of containers to provide additional workflow functionality to their user community. However, integration with a center's HPC environment including HPE's CPE can be complex. In this BOF, representatives from multiple centers will discuss their center's effort and goals to integrate and utilize containers. Representatives from HPE will discuss the CPE integration options. Discussion is the focus of each PEAD BOF. To help promote discussion, each BOF will open with short presentations followed by open floor discussions.

Sessions and Abstracts

Python Management

1:45pm - 2:30pm

Chun Sun (HPE); Cristian Di Pietrantonio (Pawsey); Dave Carlson (Stony Brook University); and Juan Herrera (EPCC, The University of Edinburgh)

Abstract

Chun Sun (HPE), Cristian Di Pietrantonio (Pawsey), Dave Carlson (Stony Brook), Juan Herra (EPCC) Python continues to play a growing role in workflows used within the HPC community. Integrating Python environments with varying needs into HPE programming environments can be complex. For CUG centers, managing multiple Python environments and ensuring performance can be non-trivial. In this BOF, representatives from multiple centers will discuss their center's Python environment use cases, management, and lessons learned. Representatives from HPE will discuss the CPE provided Python. Discussion is the focus of each PEAD BOF. To help promote discussion, each BOF will open with short presentations followed by open floor discussions.

XTreme (Approved NDA Members Only) Newport III

2pm - 2:30pm

Nick Hagerty (Oak Ridge National Laboratory) and David Carlson (Stony Brook University)

Abstract

Programming Environment Management BoF with ORNL, Pawsey and Stony Brook representative.

XTreme (Approved NDA Members Only)

2:30pm-3:00pm

Coffee Break Newport Foyer

3:00pm-5:00pm

Programming Environments, Applications, and Documentation (PEAD) Newport I

CPE Update

3pm - 3:15pm

Barbara Chapman (HPE)

CPE Testing

3:15pm - 4:00pm

Barbara Chapman (HPE), Cristian Di Pietrantonio (Pawsey), Brian Vanderwende (NCAR), Brandon Cook (Lawrence Berkeley National Laboratory/National Energy Research Scientific Computing Center), and Cedric Jourdain (CINES)

Abstract

HPC programming environments can be very complex, containing libraries, compilers, and tools that must work together to provide an effective resource to a center's user community. Over a resource's lifespan, upgrades can impact not only the individual component, but the ability for multiple components to successfully work together. Testing at various stages of a resource's lifespan is crucial to ensure the numerous hardware and software components are in working order. The goal of this BOF is to provide a venue for CUG member sites to share techniques, best practices, and lessons learned for resource testing. During the BOF HPE representatives will discuss the environ-

ment, process, and tools used to test the CPE. Discussion is the focus of each PEAD BOF. To help promote discussion, each BOF will open with short presentations followed by open floor discussions.

Exploring the Challenges of the World-Class HPE Cray Programming Environment for Modern Software Development in Fortran

4:00pm - 4:30pm

Manuel Arenaz (Codee)

Abstract

Modernizing the software development workflow for Fortran developers is crucial to enhance productivity, code quality, and maintainability. Despite being around since 1950, Fortran remains widely used in Aerospace, Automotive, Climate & Weather, Defense, Energy & Utilities, High Performance Computing, Manufacturing, Oil and Gas, Scientific Research, and other industries. However, traditional Fortran development often lacks modern tools that do exist for C/C++, which empower these developers to implement modern DevOps best practices in the organization.

The HPE Cray Programming Environment (CPE) is recognized as the best-in-class ecosystem for Fotran/C/C++, providing the CCE compilers integrated with the tools CrayPat and Reveal, among others. Recent community feedback revealed the need for better Fortran code formatting tools and static code analyzers, highly customizable and seamlessly integrated with popular Integrated Development Environments (IDEs) like VS Code, version control systems like Git, and Continuous Integration/Continuous Deployment (CI/CD) frameworks like GitLab and GitHub. This begs an open question: what are the current challenges for the world-class HPE CPE to better support modern software development best practices in Fortran?

The key is to provide Fortran developers with new automated tools that speed up frequent coding tasks carried out daily. For example, by employing code formatting tools, developers can ensure consistent code style, which improves readability and simplifies collaboration. Static code analyzers play a vital role in detecting potential bugs, enforcing coding guidelines, and ensuring code correctness by identifying issues early in the development process. Integration with modern IDEs provides advanced features such as syntax highlighting, autocompletion, and debugging support, significantly improving the development experience. Utilizing version control systems enables efficient collaboration, change tracking, and rollback capabilities. Furthermore, embedding Fortran workflows into CI/CD pipelines ensures continuous testing, automated builds, and faster deployment cycles. So, what can we learn from the 20-year-old successful tooling ecosystem available for C/C++ embedded software development and adopt it for Fortran in HPE CPE?

In this BoF we aim to discuss what important capabilities and tools are desirable that would enhance today's HPE CPE. Brainstorming and networking with HPE/Cray users and experts through short presentations, discussions with panelists, and interactive Q&A is the main objective of this BoF.

Open Discussions Birds of a Feather

4:30pm - 5:00pm

Chris Fuson (Oak Ridge National Laboratory)

Sessions and Abstracts

Open Floor Discussions

XTreme (Approved NDA Members Only) Newport III

5:30pm-6:30pm

Welcome Reception Newport Foyer

All attendees and their guests are invited to attend the welcome reception.

6:30pm-8:30pm

Program Committee Dinner (invite only) Batello

Participants that helped with the reviews and program committee are invited to a private event. .

CUG Program Committee

Monday, May 5th

8:30am-10:00am

Tutorial 1A Newport III

Monitoring HPE Cray HPC systems Tutorial

Harold Longley, Sue Miller, Raghul Vasudevan, and Pete Guyan (HPE)

Abstract

This tutorial covers the current monitoring architecture for HPE Cray HPC systems running HPCM (HPE Performance Cluster Manager) and CSM (Cray System Management) and how to configure the monitoring components to ensure log and telemetry data flows into the proper location for analysis. The various available analysis tools are discussed including OpenSearch dashboards for log data, Grafana for telemetry such as compute node resources and fabric metrics, infrastructure monitoring and alerting whilst looking at the latest new data inspection tools. The Monitoring Pipeline Visualization Tool (MPVT) enhances system monitoring by collecting real-time metrics from monitoring stack components such as Kafka, OpenSearch, and VictoriaMetrics via Prometheus exporters, and storing the data in VictoriaMetrics. MPVT generates detailed graphs using Grafana's Node Graph panel to represent the health, performance, and data flow of the monitoring pipeline. Alerting configuration and customization are presented.

The tutorial summarizes commonality in addition to differences between HPCM and CSM monitoring tools and architecture and previous versions of those. Verifying system performance and looking for node outliers is also considered. An overview of AIOps (artificial intelligence for operations) plus its installation and base configuration is provided.

The hands-on portion of the tutorial provides access to a system to explore the analysis tools to use dashboards, customize queries, and analyze system alerts while being guided by HPE experts about the tooling.

Feedback will be sought from the community about their needs for further improvements for monitoring.

8:30am-10:00am

Tutorial 1B Newport I

Hands on with uenv and CPE in a container with Grace Hopper on Alps Tutorial

PBen Cumming and Tim Robinson (CSCS) and Jonathan Gallmeier (HPE)

Abstract

Management of the Cray Programming Environment (CPE) can pose challenges for CUG sites, including dependency issues within a monolithic stack, disruptive updates, and the need for administrative privileges. For these reasons, the CUG community has expressed considerable interest in exploring alternative approaches for deploying and running software on HPE Cray systems that do not rely on the CPE being included in the operating system image. This tutorial offers a hands-on experience with building and managing software on an HPE Cray system that does not have the CPE installed, specifically using uenv and a containerized CPE. Participants will get access to the cutting-edge NVIDIA Grace-Hopper (ARM64) and AMD EPYC (x86-64) nodes on Alps -- an HPE Cray Supercomputing EX system at CSCS.

8:30am-10:00am

Tutorial 1C **Newport III**

Best Practices For Operating and Maintaining Slingshot Fabrics-Tutorial

Ryan Tidwell and Forest Godfrey (HPE)

Abstract

Whether performing an initial bringup, applying software updates, or adding capacity, handling routine maintenance tasks is a fact of life. This tutorial will guide attendees through best practices for handling common scenarios in the lifecycle of Slingshot fabrics. Emulated Slingshot environments and sample fabric designs will be used to provide hands-on experiences working through important workflows in the lifecycle of a Slingshot deployment.

Attendees wishing to follow along with the tutorial hands-on should bring a laptop running Docker including Docker Compose (or a compatible container technology such as podman). USB drives with the tutorial containers and sample data will be provided.

Tutorial 1D **Newport IV**

Exploring High Performance Storage with DAOS Tutorial

Mohamad Chaarawi (HPE); Adrian Jackson (EPCC, The University of Edinburgh); and Jerome Soumagne (HPE)

Abstract

The diversity of applications utilizing HPC has been increasing beyond computational simulation approaches to a more varied mix including machine learning and data analytics. This introduces changes in I/O patterns and requirements on data storage. Traditional data storage technologies in HPC have long been optimized for bulk data movement, focused on high bandwidth with relatively low volumes of metadata operations. However, many applications now exhibit non-optimal I/O patterns for Parallel File Systems (PFS), with significant amounts of small I/O operations, non-contiguous data access, and increases in read as well as write I/O activities. PFS today are not optimized for all such patterns and are struggling with the diversified application workloads.

The Distributed Asynchronous Object Storage (DAOS) was designed to address those challenges. DAOS is an open-source software defined high performance data store designed

Sessions and Abstracts

from the ground up for massively distributed Non-Volatile Memory and Storage Class Memory. It presents a key-value interface and provides features such as transactional non-blocking I/O, a versioned data model, and global snapshots – all completely in user space. DAOS also provides "legacy" interface mechanism, such as file access, making porting applications to DAOS doable without code modifications.

This tutorial will provide hands-on experience using and programming against DAOS. The tutorial will detail the mechanisms for leveraging the DAOS tools, programming with the DAOS KV/KArray API and using the File System layer for DAOS. We aim to provide the knowledge and skills required to evaluate DAOS for applications and systems and leverage new access mechanisms where appropriate.

10:00am-10:30am

Coffee Break **Newport Foyer**

10:30am-12:00pm

Tutorial 1A *(continued)* **Newport I**

Monitoring HPE Cray HPC systems Tutorial

Tutorial 1B *(continued)* **Newport III**

Hands on with uenv and CPE in a container with Grace Hopper on Alps Tutorial

Tutorial 1C *(continued)* **Newport III**

Best Practices For Operating and Maintaining Slingshot Fabrics-Tutorial

Tutorial 1D *(continued)* **Newport IV**

Exploring High Performance Storage with DAOS Tutorial

12:00pm-1:00pm

Lunch (by Codee) **Newport Foyer**

1:00pm-2:30pm

BoF 1D **Newport IV**

Security BoF **Birds of a Feather**

Aaron Scantlin (National Energy Research Scientific Computing Center)

Abstract

This is the first BoF at CUG for the new HPC Security Special Interest Group. Its mission is to provide a forum for both exchanging and discussing strategies and ideas related to ensuring the secure configuration (of both hardware and software), operation (from the system administrator's perspective) and utilization (from the user's perspective) of Cray/HPE systems. Related topics in HPC security (e.g. awareness training for both users and admins) will also covered by the SIG.

Tutorial 1A *(continued)* **Newport III**

Monitoring HPE Cray HPC systems Tutorial

1:00pm-2:30pm

Tutorial 2B **Newport I**

Automated Inspection of Fortran/C/C++ Code Using Codee for Correctness, Modernization, Optimization, and Security on HPE/Cray Tutorial

Sessions and Abstracts

Manuel Arenaz (Codee)

Abstract

Codee is a software development tool that facilitates the development, maintenance, modernization, and optimization of Fortran/C/C++ codes by providing a systematic and predictable approach to finding and fixing issues related to correctness, modernization, optimization, and security vulnerabilities. Codee provides automated checkers for the rules documented in the Open Catalog of Code Guidelines for Correctness, Modernization, and Optimization. It also features AutoFix capabilities to enable semi-automatic source code rewriting, by modifying Fortran statements and inserting OpenMP or OpenACC directives. Codee integrates seamlessly with popular IDEs, Control Version systems and CI/CD frameworks. Overall, Codee helps developers uncover hidden bugs, avoid introducing new ones, and pinpoint suggestions for various code improvements. In this tutorial, participants will explore Codee and the Open Catalog through short demos and hands-on exercises, with step-by-step instructions for HPE/Cray systems such as Perlmutter. The session begins with simple, well-known kernels and quickly progresses to large HPC codes like WRF, enabling participants to effectively use Codee tools in real-world scenarios.

1:00pm-2:30pm

Tutorial 2C **Newport II**

Performance Analysis on AMD GPUs

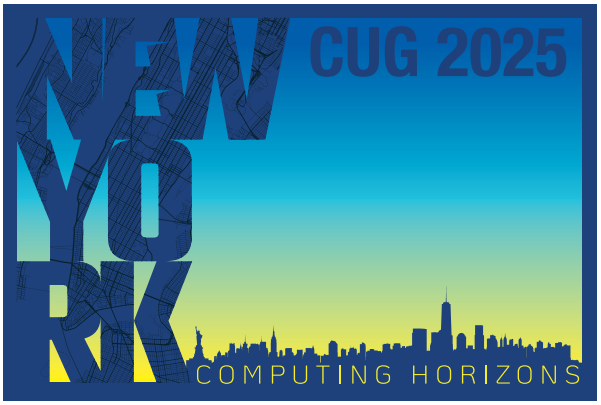

Georgios Markomanolis and Samuel Antao (AMD)

Abstract

Over the past few years, AMD has released several profiling tools for AMD GPUs. These tools, called rocprofv3, ROCm Systems Profiler (rocprof-sys, ex-Omnitrace) and ROCm Compute Profiler (rocprof-compute, ex-Omniperf), are now AMD products. In this tutorial, divided in 4 parts, in the first part we discuss the used applications and GPUs for the current tutorial, including MI300A. Afterwards, we will showcase the latest advancements and offer guidance on selecting the most suitable tool based on your specific needs. We will introduce rocprofv3, showcasing its new features, usage, and functionalities designed to enhance user experience, instrumenting MPI applications is straight forward compared to previous rocprof version. Notable additions include support for OpenMP offloading profiling for Fortran, which we will demonstrate among also other improvements, as also improved overhead efficiency. In the third part, we will explain the capabilities of the rocprof-sys tool, for timeline analysis of an application execution, visualize traces, and understanding the insights among also performance characteristics. Participants will learn how to use profiling tools with key applications and various programming models, including MPI, OpenMP, Python, and Kokkos. In the last part, the rocprof-compute will be used for the roofline analysis of kernel performance, presenting the new developments and improvements. Additionally, we will delve into identifying inefficient metrics affecting specific kernel performance and provide deeper insights into optimization strategies.

2:30pm-2:45pm

Coffee Break (by SchedMD) Newp. Foyer

Sessions and Abstracts

2:45pm-4:15pm

BoF 2D **Newport IV**

Kubernetes on HPE Supercomputers BoF

Sadaf Alam (University of Bristol), Dino Conciatore (Swiss National Supercomputing Centre), and Jesse L. Treger (HPE)

Abstract

Cloud native technologies including containers, orchestration engines like Kubernetes and virtualisation on the compute plane for HPE EX platforms will be discussed in this BOF with the CUG community. There are diverse use cases driving these somewhat non-native HPC requirements for user or custom defined platforms. These range from AI and ML workflows to the Trusted Research Environments to users-managed CI/CD pipelines. Essentially, we will explore how community needs are evolving for not only running containers on HPC but also MLOps style workflows, which require a level of user autonomy that is not feasible within a batch scheduling system alone. This interactive BOF will include brief presentations (in cases summaries of topics covered in detailed sessions and submitted papers) on the progress made in prototyping Kubernetes environment on the compute plane with the HPE Slingshot interconnect, use cases from sites running diverse supercomputing platforms including variants of CPUs and GPU technologies from different vendors, and a panel-led BoF discussion on future directions, priorities, and challenges.

2:45pm-4:15pm

Tutorial 1A *(continued)* **Newport III**

Monitoring HPE Cray HPC systems **Tutorial**

Tutorial 2B *(continued)* **Newport I**

Automated Inspection of Fortran/C/C++ Code Using Codee for Correctness, Modernization, Optimization, and Security on HPE/Cray **Tutorial**

2:45pm-4:15pm

Tutorial 2C *(continued)* **Newport II**

Performance Analysis on AMD GPUs **Tutorial**

4:20pm-5:30pm

BoF 1A **Newport III**

CSM updates, iSCSI boot content projection, and other CSM topics

THarold Longley, Dennis Walker, Ravi Bissa, Jason Coverston, Siri Vias Khalsa, Ashalatha A. M, and Ravikanth Nalla (HPE)

Abstract

This BOF has an overview of recent changes in the Cray System Management (CSM) software stack, a technical dive into iSCSI-based boot content projection, a method for workload manager-controlled reboots of compute nodes, and then an open discussion amongst attendees.

The overview highlights software stack improvements for operating a CSM system including upgrades, changes in boot content projection with iSCSI, multi-tenancy features, monitoring improvements and alert management, tuning compute nodes, containers on compute nodes and user access nodes (UANs),

and CSM diagnostics and hardware triage tools.

Boot content projection of "rootfs" and "programming environment" images for compute nodes and UANs has been done using Cray Data Virtualization Service (DVS), but this encounters various challenges related to reliability, availability, security, deployment speed, and ease of management. HPE has implemented an alternate boot content projection known as "Scalable Boot Content Projection Service (SBPS)". SBPS is a simple user space solution which is lightweight, modular, highly available, easy to deploy and manage, and based on Linux-standard iSCSI.

Upgrading OS and firmware has challenges like managing job interruptions, dealing with node dependencies, minimizing performance degradation, ensuring system availability, handling scheduling conflicts, ensuring data integrity, and preventing human errors. An architecture for addressing this is described that ensures zero job preemption and continuous service delivery, with no noticeable delays, errors or disruptions. This has been implemented on a multi-zoned system with CSM and HPCM using PBS Pro with high priority queues to decouple dependencies and dynamically compute number of down nodes to throttle reboot operations.

4:20pm-5:30pm

BoF 1B **Newport I**

CUG SIG System Monitoring Working Group BoF

Massimo Benini (CSCS - ETH Zurich), Lena Lopatina (Los Alamos National Laboratory), and Jeff Hanson and Pete Guyan (HPE)

Abstract

The System Monitoring Working Group (SMWG) is a CUG SIG (Special Interest Group) established to promote collaboration between HPE Cray and its customers in enhancing system monitoring capabilities. The group includes representatives from numerous HPE Cray member sites and meets regularly to discuss and address system monitoring-related topics. We meet virtually throughout the year, and the annual CUG meeting is an opportunity for participants from normally-incompatible time zones to swap ideas about their monitoring setup, needs, and findings. This year's main key topics are new monitoring capabilities provided by HPE, member sites creation of a standardized set of dashboards for monitoring HPC centers, exploring new approaches to application code profiling, energy profiling, and data centers metadata management. This BoF is a unique opportunity to engage collaboratively with HPE Cray and with other HPE Cray sites to shape future enhancements of observability and operational data analytics (ODA). This forum also provides a platform to share best practices, showcase common problems and their solutions, and offer HPE a platform to present the current status of their monitoring platform directly to relevant HPC staff

4:20pm-5:30pm

BoF 1C **Newport III**

Sharing is Caring: Tackling Node-Sharing Challenges at CUG Sites

Sessions and Abstracts

Tim Robinson (Swiss National Supercomputing Centre, ETH Zurich); Tim Wickberg (SchedMD LLC); Pengfei Ding (Lawrence Berkeley National Laboratory); and Cristian Di Pietrantonio (Pawsey Supercomputing Research Centre)

Abstract

HPC centres have traditionally allocated computational resources in entire nodes, a practice rooted in the architectural and operational simplicity of earlier systems, and the belief that "HPC" meant the ability to scale to hundreds or thousands of nodes. However, as technology has advanced, nodes have become increasingly powerful and expensive, incorporating hundreds of CPU cores, multiple GPUs, and large amounts of high-bandwidth memory. This evolution makes whole-node allocation inefficient for many modern workloads—high-throughput computing, data-intensive workflows and interactive computing, to name a few. Consequently, subdividing nodes to allocate resources at finer granularity—by socket, core, GPU, or memory—has emerged as a necessary alternative.

While node subdivision can improve resource utilization and lower costs, it introduces several challenges, including technical, operational, and business concerns. Fine-grained resource allocation complicates job scheduling and resource management, requiring enhancements to workload managers like Slurm to handle mixed-resource requests effectively. Subdividing nodes also increases the risk of resource contention, such as memory bandwidth bottlenecks, interconnect congestion, or power sloshing, which can all affect application performance. Finally, the heterogeneity of modern nodes, with CPUs, GPUs, and other accelerators, creates challenges for monitoring and accounting: how should researchers be charged for using fractions of these components?

In this BOF, we will explore topics such as resource contention, performance variability, security and isolation, scheduling complexity, accounting, and scalability of shared workloads. Our focus will extend beyond static resource allocations at the Slurm level to include finer, dynamic granularities, such as GPU sharing (e.g., NVIDIA's Multi-Instance GPU and the recently open-sourced Run:ai). Ahead of the conference, we will gather discussion points and questions from the wider CUG community. At the session, speakers from Pawsey, NERSC, and CSCS will share their current practices and future needs, while SchedMD will offer insights from the workload management perspective.

4:20pm-5:30pm

BoF 3D **Newport IV**

Rethinking Interactive HPC Resource Access: Enhancing Security and Flexibility

Maxime Martinasso (Swiss National Supercomputing Centre), Sadaf Alam (University of Bristol), and Isa Wazirzada and Alex King (HPE)

Abstract

The traditional approach to accessing HPC resources relies on login nodes and SSH connections authenticated through POSIX Identity and Access Management (IAM). While this method has served the community well, it presents significant challenges in today's land-

scape of cybersecurity threats and evolving user needs, such as maintaining a secure shared login node or managing identity life-cycle. This Birds of a Feather (BoF) session aims to explore innovative approaches to modernize interactive HPC resource access with CLI, addressing the dual goals of enhancing security and increasing service customization flexibility for users. Emerging practices, such as SSH signed keys, offer a promising alternative to traditional log-in names and passwords, mitigating risks associated with credential theft by enabling more advance authentication flow like multi-factor authentication. Virtualized login nodes, implemented as containerized environments, could allow user-defined environments with for instance advanced debugging capability, AI stacks or a higher integration with IDE while improving isolation, scalability of users, and individual session management. Additionally, the generation of temporary POSIX accounts from OpenID Connect (OIDC) tokens could seamlessly integrate modern federated and non-local identity providers, reducing administrative overhead and attack surfaces. The session will showcase existing solutions, discuss opportunities for innovation, challenge classic IAM HPC and login nodes workflow and highlight the potential benefits of these new approaches. Attendees will hear from practitioners actively exploring these paradigms, sparking discussions on how the community can collectively advance this shift and benefit for a common solution. We invite participants to contribute their ideas, share experiences, and help shape a future where interactive HPC resource access is not only more secure but also more adaptable to the diverse and continuously evolving needs of its users.

Organiser: CSCS has implemented MFA by introducing SSH services and SSH-signed keys as the sole method for accessing login nodes. BriCS has adopted a similar approach, enhancing user management by generating POSIX accounts on login to eliminate the need for manual identity management. Additionally, HPE has introduced the concept of User Access Node (UAN), which function as on-demand containers for login nodes.

Agenda: To kick off the session, we will have three brief 5-minute presentations introducing key concepts: the CSCS SSH service, BriCS identity management approach, and the HPE User Access Node (UAN) concept. Following these presentations, we will present a list of thought-provoking questions to ignite discussion and encourage debate. See examples in the attached document.

Tuesday, May 6th

8:30am-10:00am

Plenary: Welcome, Keynote Newport III

8:30am - 8:45am

Opening

Ashley Barker (Oak Ridge National Laboratory)

Welcome to from the CUG President

8:45am - 9:35am

What I've Learned About Supercomputing from Blowing Up Stars

Michael Zingale (Stony Brook University)

Abstract

Stars shine throughout their lives by convert-

ing light elements into heavier elements via nuclear burning. While there are different pathways that low and high mass stars take in their evolution as they exhaust their fuel, explosions of both groups (or their remnants) are possible, leading to a wide-range of stellar transients. Modeling these events requires capturing the interplay between hydrodynamics, nuclear reactions, gravity, radiation, rotation, and more physics. These models are also inherently multi-dimensional and span a vast range of timescales. Both algorithmic developments and leveraging of modern supercomputer architectures are key to performing accurate and efficient simulations of these explosions. In this talk, I will discuss some of the lessons I've learned from more than two decades of writing simulation codes for these problems. I will show examples of where new algorithms needed to be developed, instead of using general codes, and when complete rewrites of codes were needed to support new architectures. Finally, I will talk about how we will train our students to write the next generation of codes.

9:35am-9:45am

New CUG Sites

New Member Site: Introducing LRZ

9:45am - 10:00am

Canidates Statements

10:00am-10:30am

Coffee Break (by Pier Group) Newp. Foyer

10:30am-12:00pm

Plenary: CUG site, HPE update Newport III

10:30am - 10:45am

Welcome by Stony Brook

Robert Harrison (Stony Brook University)

Abstract

Overview of Institute for Advanced Computational Science at Stony Brook University.

10:45am-10:50am

AI/ML Intelligent Scheduling for HPC with Altair®

Bill Nitzberg (Altair Engineering, Inc.)

Abstract

What's possible when you start with a deep understanding of HPC usage patterns (via Altair InsightPro™) and use that data to build predictive AI models (via Altair RapidMiner®) to augment traditional HPC (via Altair HPCWorks®)? Better utilization (a lot better), faster turnaround (a lot faster), and more throughput (a lot more).

10:50am-11:00am

NVIDIA HPC Software - Expanding HPC with Python & AI

Becca Zandstein (NVIDIA)

Abstract

NVIDIA's HPC Software enables developers to build applications that take advantage of every aspect of the hardware available to them: CPU, GPU, and interconnect. In this presentation you will learn the latest updates on NVIDIA's HPC software that is being used in HPC centers, including the latest AI for Science and

Sessions and Abstracts

Python products. This presentation will provide an overview of the NVIDIA HPC Software stack, taking you from traditional HPC compilers, CPU optimized libraries, Python HPC tooling, and AI for Science software that you can start using today.

11:00am-12:00pm

Gerald Kleyn (HPE)

Abstract

HPE corporate update by Gerald Kleyn

12:00pm-1:00pm

Lunch (by Nvidia) **Newport Foyer**

1:00pm-2:30pm

Technical Session 1A: Multitenancy **Newport III**

1pm - 1:30pm

Infrastructure as a Service with Strong Tenant Separation on a Supercomputer

Riccardo Di Maria, Chris Gamboni, Manuel Sopena Ballesteros, Hussein Harake, Mark Klein, Marco Passerini, Miguel Gila, Maxime Martinasso, and Thomas C. Schulthess (Swiss National Supercomputing Centre) and Alun Ashton, Derek Feichtinger, Marc Caubet, Elsa Germann, Hans-Nikolai Viessmann, Achim Gsell, and Krisztian Pozsa (Paul Scherrer Institute)

Abstract

This paper explores the innovative implementation of Infrastructure-as-a-Service (IaaS) on a HPE Cray Shasta EX supercomputer. In cloud environments, IaaS offers scalable, on-demand access to virtualized resources. However, applying IaaS principles to high-performance computing (HPC) systems without relying on virtualization technologies poses some challenges, since they typically have a tightly coupled software stack. We address these challenges in a co-design partnership between an HPC provider, CSCS, and an end-user institution, PSI, by developing a suite of technologies for the HPE Cray Shasta EX system architecture that supports resource isolation and granular control. This approach not only provides the IaaS model on supercomputing environments but also enables dynamic resource management. Our contributions include a detailed exploration of the technological advancements necessary for integrating IaaS into HPC, together with the lessons learned from our collaborative efforts. By extending IaaS capabilities to supercomputers, we aim to provide scientific institutions with unprecedented flexibility and control over their computational resources.

1:30pm - 2pm

Dynamic Network Perimeterization: Isolating Tenant Workloads With VLANs, VNIs, & ACLs

Nikhil Mukundan, Dennis Walker, Stephen Han, Atif Ali, Siri Vias Khalsa, Amit Jain, Vishal Bhatia, and Vinay Karanth (HPE)

Abstract

There is a growing trend in the high-performance computing (HPC) community where separate user groups share HPC infrastructure with varying security clearances (tenants). In such cases, tenants require robust security boundaries to ensure data privacy, results integrity, and intellectual property secrecy.

Sessions and Abstracts

Additionally, sensitive transactions within a tenant may need to be further insulated from lower-clearance workloads. Join us as we show how product-agnostic, version-controlled configuration data can be used to dynamically isolate infrastructure resources supporting workloads, including compute node groups, data-at-rest (storage), and data-in-motion within high-speed and management networks. On the high-speed network (HSN), we'll examine how switch port VLAN filters and VNIs (traffic labels) isolate TCP/IP and RDMA traffic per tenant. On the management network, we'll demonstrate how to segment compute node groups via switch ACLs, VLANs, and iptables. Complete dynamic network segmentation will be applied at various levels of infrastructure from chassis, nodes, and within the OS. Finally, we'll review architecture features in Slingshot, CSM, and other products that enable elastic tenant reallocation. We'll compare and contrast the differences in the number of configuration options and security posture when applying segmentation at the switch vs the node.

2pm - 2:30pm

CSCS' journey towards complete platform automation in a multi-tenant environment

Miguel Gila, Ivano Bonesana, and Alejandro Dabin (Swiss National Supercomputing Centre, CSCS)

Abstract

The Swiss National Supercomputing Centre operates a complex ecosystem of high-performance computing resources. With a focus on scalability and efficiency, we have implemented a multi-tenancy approach to serve diverse scientific communities. This layered architecture encompasses infrastructure, platform, and application layers, each with unique automation challenges.

To drive process efficiencies and reduce manual intervention, we are embarking on a journey towards complete platform automation. Currently, we are making progress in several areas: integrated automated node health analysis and recovery with job schedulers (Slurm); Terraform pipelines for streamlined infrastructure provisioning and platform re-configuration through rolling updates; and complete API-based operations for simplified workflows. Whereas in a multi-tenant environment automation is crucial, traditional single-tenant HPC systems can also benefit from the added consistency, efficiency and versioning provided. Future work will focus on integrating these efforts into a cohesive automation framework leveraging modern cloud technologies.

This presentation will detail our successes and challenges, outlining the path forward towards a more autonomous CSCS platform that better serves our users' needs.

1:00pm-2:30pm

Technical Session 1B: Workload Manager — Newport I

1pm - 1:30pm

Slinky: The Missing Link Between Slurm and Kubernetes

Tim Wickberg (SchedMD LLC)

Abstract

Slinky is SchedMD's collection of projects to integration the Slurm Workload Manager with the Kubernetes Orchestrator.

This presentation will cover the Slurm Operator - designed to manage ephemeral Slurm clusters, and which was released under the Apache 2.0 license in November 2024 - as well as the upcoming Slurm Bridge - which will provide for converged scheduling between Slurm and Kubernetes compute workloads.

Time permitting, a brief review of the forthcoming Slurm 25.05 release will be provided as well.

1:30pm - 2:00pm

How Best to Leverage Cloud for (Big) HPC Sites

Bill Nitzberg and Ian Littlewood (Altair Engineering, Inc.)

Abstract

Cloud (finally) works for HPC, but the devil is still in the details. HPC in the Cloud has transitioned from "proof-of-concept engagements" to "hmm, but what about security" to "maybe, but what about the data" to "OK, but only if we carefully manage expenses". Today, significant big sites have voted with their dollars that on-premise HPC is here to stay, adding Cloud judiciously, with an eye towards resilience, and only where the end-to-end ROI makes sense.

HPC sites tend to run at 85-99% utilization. At these levels, Cloud, especially Cloud tuned for HPC, remains 2-10x more expensive than on-premise HPC. For this reason, the best use cases for Cloud HPC are hybrid, with most work executing on-premise, while using Cloud for transient activities: to meet project deadlines, to gain access to hard-to-get resources, to use for testing, and to recover from disasters. This hybrid model -- marrying on-prem and Cloud -- requires a flexible approach, including access to multiple Clouds, DevOps-ready automation, ability to embed site-specific policy, openness to schedulers, bulletproof accountability, and AI readiness. We present a proven set of use cases, with examples from real sites over the past several years, for hybrid HPC-Cloud, that deliver real ROI.

2:00pm - 2:30pm

Divide and Rule: Automated Workload Distribution for Efficient User Support Services

Luca Marsella (Swiss National Supercomputing Centre)

Abstract

User support services for High-Performance Computing (HPC) systems help users conduct simulations and optimize resources. The complexity of HPC platforms has increased, making user support more challenging. Site Reliability Engineering (SRE) best practices suggest that technical staff should focus on project work and automation to reduce repetitive tasks. Artificial intelligence and machine learning can help create effective knowledge bases from internal reports and user tickets, easing the burden on support staff.

Work management tools can further reduce toil by organizing user communications and managing status pages. The Swiss National Supercomputing Center (CSCS) uses a functional structure with Working Structures (WS) to group staff by expertise. The Service Reliability Management (SRM) WS aims to automate tasks and improve efficiency by implementing SRE practices and a service-oriented architecture (SOA).

WS-SRM focuses on eliminating toil, improving velocity, and monitoring production services for reliability. User support services are measured by response times and support request traffic. CSCS's distributed user support system relies on WS for end-to-end service delivery. Future challenges include automated task review and workload management. The presentation aims to share best practices and gather feedback to enhance user support at CSCS.

1:00pm-2:30pm

Technical Session 1C Newport II

1pm - 1:30pm

Deploying and Tracking Software with NCCS Software Provisioning

Asa Rentschler, Nicholas Hagerty, Elijah Maccarthy, and Edwin F. Posada Correa (Oak Ridge National Laboratory)

Abstract

The National Center for Computational Sciences (NCCS) at Oak Ridge National Laboratory has developed NCCS Software Provisioning (NSP), a unified framework for deploying software stacks and monitoring their usage on HPC systems. NSP leverages Ansible to automate the deployment of Spack environments using templates and to manage installation procedures for non-Spack software through custom roles. Additionally, NSP enhances vendor-managed LMOD installations using hooks, enabling dynamic and responsive software layouts that adapt seamlessly to changes in the programming environment. For software usage tracking, NSP configures LMOD to track module usage data, which is then stored in a centralized database, and can be visualized with Grafana dashboards. This paper discusses the design and motivation behind NSP, its strategies for managing software complexity under vendor-provided programming environments, and its implementation on Frontier, the world's fastest supercomputer for open science.

1:30pm - 2:00pm

Modern Software Deployment on a Multi-Tenant Cray-EX System

Ben Cumming, Simon Pintarelli, Andreas Fink, Mikael Simberg, and Jonathan Coles (CSCS)

Abstract

User-facing software -- libraries, tools, applications and programming environments tuned for the node and network architecture -- is a key part of HPC centers' service offering. Teams that maintain and support this software face challenges: providing a stable software platform for users with long running projects while also providing the latest versions of software for developers; giving full responsibility to build, modify and deploy the whole software stack to staff who do not have root acccess; and reproducible deployment based on GitOps practices. CSCS addresses these challenges on Alps by using small independent software environments called uenv, which deploy from text-file recipes without requiring installation of the Cray Programming Environment. This paper discusses installing communication libraries from HPE and NVIDIA with SlingShot support; the CI/CD pipeline that builds uenv and deploys them in a container registry; and the command line tools and SLURM plugin that

interface users with the software environments. We demonstrate diverse use cases such as JupyterHub, summarize the user and support team experience, and document how to build and deploy CPE containers.

2:00pm-2:30pm

Employing a Software-Driven Approach to Scalable HPC System Management

Aaron Barlow (Oak Ridge National Laboratory)

Abstract

Managing Frontier and other HPE Cray and Apollo clusters at Oak Ridge National Laboratory involves thousands of users, projects, and security policies across multiple HPC systems. With diverse research needs, varying security enclaves, and massive resource allocations, manual processes don't scale, and administrative burden increases as HPE systems grow. To manage HPC systems at this scale, we developed RATS (Resource Allocation Tracking System), a software platform that centralizes operations.

By providing RESTful API endpoints, RATS defines and distributes policy updates while system-specific scripts enforce configurations on HPE systems. This enables user and group management, node allocations, scheduler configurations, directory management, sudo rule configuration, file system provisioning, purge exemptions, and security compliance across all HPC systems. In addition, RATS integrates with myOLCF, providing users with a self-service portal for managing projects, tracking allocations, and accessing system analytics.

This session will provide an in-depth look at RATS, covering its core features, policy enforcement mechanisms, and automation of HPC operations at scale. We'll explore the software behind RATS—its integration with job schedulers, identity systems, and myOLCF—along with how it operates and enforces policies on HPE systems via API-driven automation and script execution..

2:30pm-3:00pm

Coffee Break (by Linaro) Newport Foyer

3:00pm-5:00pm

Technical Session 2A: Slingshot Newport III

3:00pm - 3:30pm

The HPE Slingshot 400 Expedition

Houfar Azgomi, Duncan Roweth, Gregory Faanes, and Jesse Treger (HPE)

Abstract

HPE Slingshot 400 is a high-performance interconnect for classic and AI supercomputing clusters. As the successor of HPE Slingshot, it comprises a PCIe Gen5 NIC (Cassini-2) and a 64-port switch (Rosetta-2), linking over standard 400 Gbps Ethernet physical interfaces, and enabling dragonfly and fat-tree networks with up to 260,000 endpoints. HPE Slingshot is currently deployed in 7 of the 10 largest supercomputers worldwide and dominates the top 3 list as the interconnect solution for El Capitan, Frontier, and Aurora machines. With such success, the Slingshot Transport protocol has become the cornerstone for HPC-optimized Ethernet networking standardization efforts led by the Ultra Ethernet Consortium (UEC). Growing around its foundational adaptive routing and congestion management feature set, the

HPE Slingshot 400 interconnect doubles its predecessor's bandwidth with significant enhancements: exact match forwarding increasing route visibility across the cluster; dedicated ACL tables for security and cloud isolation; feature hardening flexibility with P4-programmability; and improved quality of service with 50% more traffic classes. It is supported across HPE's portfolio of rack- and chassis-based supercomputing platforms including HPE Cray XD, HPE Cray EX, and the latest HPE Cray GX. This paper presents the key features and some early performance results of Rosetta-2 and Cassini-2 devices.

3:30pm - 4:00pm

Introduction To HPE Slingshot NIC Libfabric Environment Variables

Jesse Treger and Ian Ziemba (HPE)

Abstract

Libfabric, a high-performance fabric software library, provides a rich API for applications to communicate efficiently over various networking technologies. The Libfabric provider for the specific networking interface type translates the API requested communications into optimal protocols that strive to optimize use of hardware. The HPE Slingshot NIC in particular has extensive hardware offloads to improve performance and reduce memory overhead. While Libfabric aims for seamless integration, achieving optimal performance may require users to configure environment variables to fine-tune the software for specific workloads and hardware setups. This presentation will demystify the role of environment variables in Libfabric, explaining trade-offs and why they matter to performance and stability under various conditions. We will begin with an overview of how the HPE Slingshot NIC uses Libfabric to optimize performance with various messaging requirements. Next, we will explore the most common environment variables users may need to adjust from default values, using examples learned on different applications, MPI middleware, processors, or job scale. Finally, we will briefly touch on some best-known methods to troubleshoot application failures that can be addressed with environment settings.

4pm - 4:30pm

Math in Your Network: Slingshot Hardware Accelerated Reductions

orest Godfrey and Duncan Roweth (HPE)

Abstract

In high performance computing applications, the use of collective operations such as reductions and barriers is commonplace. The performance of collectives is critical to overall performance in many applications, especially those where collectives are increasingly large part of the runtime as jobs scale. Collective operations are typically performed by software, requiring packets carrying contributions to the collective to go all the way to endpoint memory and be acted upon only for the result to have to transit back out to the network. This occurs at each level of a collective tree. By performing collective operations inside the network switch hardware itself the round trips to memory are

removed and significant improvements in latency can be achieved. The Slingshot Rosetta switch fabric supports the hardware acceleration of many collective operation such as barriers and 64-bit IEEE floating point reductions. Upcoming Slingshot software will enable this functionality and present it to the end user transparently through the industry standard libfabric network communication library. This presentation will cover the details of this upcoming feature and how it can be used to accelerate applications.The implementation inside libfrabric, interaction with the job scheduler and fabric manager, as well as initial benchmarking results will be discussed.

4:30pm - 5:00pm

Slingshot Host Software Ethernet Tuning

Ravi Bissa, Ian Ziemba, Duncan Roweth, and Forest Godfrey (HPE)

Abstract

High-performing Ethernet is the cornerstone of Exascale supercomputers, enabling seamless communication, minimizing latency, and supporting massive scalability. Without robust Ethernet infrastructure, these systems cannot achieve their goal of solving the world's most complex computational problems efficiently and within a reasonable time and energy budgets.

Tuning Slingshot host software Ethernet on supercomputers is challenging due to system's massive scale, diverse workloads, and performance requirements. Non-Uniform Memory Access (NUMA) architecture presents unique challenges due to system's memory and CPU layout. NUMA architectures are designed for scalability, but their distributed memory model introduces complexities in achieving optimal network performance. It demands a holistic approach, including hardware optimization, software tuning, advanced monitoring, and workload-specific adjustments.

This presentation has two sections: key challenges and potential solutions & mitigations. We will start by highlighting networking concepts that can significantly degrade performance. Next, we will cover how to mitigate problems by fine tuning these parameters. These recommendations are derived from numerous tests conducted on large systems including ORNL Frontier, UK Met Office, and LANL Venado. We will also present a comparative study with before & after results. Finally, we will conclude with a step-by-step guide for different NUMA Architectures that users can follow to attain optimal performance.

3:00pm-5:00pm

Technical Session 2B: Newport I
Security & Configuration Management

3pm - 3:30pm

Pragmatic Security Audits: Fortifying HPC Environments at a Consumable Pace

Alden Stradling (Los Alamos National Laboratory) and Monica Dessouky and Dennis Walker (HPE)

Abstract

Do you know your security posture? Are you overwhelmed by the latest reports? Don't let a sea of security findings parallelize your progress. By establishing a frequent, recurring cadence for audits and remediation, organi-

zations ensure continuous protection against emerging threats. This paper represents the practical, scalable approach to HPC security implemented at a recent customer to secure its many new environments.

Start by scanning each node type—e.g., compute, login, storage, and management—for Common Vulnerabilities and Exposures (CVEs), security benchmark adherence, and open port threats. Aggregate and stack-rank these results according to a site threat model, factoring in user adjacency and exploit severity.

Next, build your backlog using the threat model and level of effort. Structure work in consumable batches. This will empower your organization to maximize its security posture, remediating vulnerabilities within the constraints of human resourcing and tolerance to potential operational disruption.

Finally, automate these scans to increase efficiency and reduce human error. This will enable security teams to focus on remediation rather than data collection, ensuring the environments remain resilient against increasingly sophisticated threats and creating a dynamic defense strategy that adapts to evolving security landscapes.

3:30pm - 4pm

Experimenting with Security Compliance Checking using ReFrame

Victor Holanda Rusu, Matteo Basso, Chris Gamboni, Fabio Zambrino, Gianna Marano, and Massimo Benini (Swiss National Supercomputing Centre)

Abstract

Security is a critical aspect of High-Performance Computing (HPC) systems, where the implementation of security compliance checks and hardened configurations is essential to safeguard resources and data. Continuous security checking is fundamental, especially for detecting indications of compromise, but its implementation must balance effectiveness and efficiency to avoid unnecessary strain. Open Source and freely available security-focused tools, such as OSCAP, are less known and accessible to engineers from other disciplines, who may not be familiar with their functionality or utility. This creates a barrier to collaborative efforts in improving system-wide configurations and promoting shift-left security in HPC centers. We leveraged ReFrame to perform robust security compliance testing to address these limitations. ReFrame enables the creation of customizable tests to evaluate system configuration, generic exploits, and execute tests in parallel, optimizing testing workflows without significant performance penalties. We will present the latest developments at CSCS, showcasing how we plan to use ReFrame to enhance security compliance testing in HPC environments using three different standards STIG DoD, ANSSI BP-28-enhanced, and CSCS'. We aim to create a community to develop, maintain, and benefit from a shared set of security checks tailored to HPC systems based on customer-specific, industry-specific, or government-mandated requirements.

4pm - 4:30pm

From Weeks to Hours: Harnessing Configuration Management and Deployment Pipelines

Dennis Walker and Siri Vias Khalsa (HPE) and

Alex Lovell-Troy (Los Alamos National Laboratory)

Abstract

Ensuring peak reliability, current functionality, and up-to-date security requires cultivating the capability to continuously update and integrate a complex array of dependencies, spanning hardware, firmware, system management software, network configurations, API services, OS distros, job schedulers, AI libraries, and analytics tools. This paper presents a simple yet contemporary DevOps methodology designed to automate, validate, and replicate changes effectively to one or many production environments.

Modernized HPC configurations increasingly span multiple sites and zones to guarantee cloud-like availability. In such cases, configuration management, infrastructure-as-code, and deployment automation are paramount for reproducing the configuration details necessary to provide a coherent user interface across the topologies.

By creating these capabilities in-house, an organization empowers its vendors to deliver more quickly. This results in providing current functionality and a better guarantee of operational reliability to stakeholders. Vendors become empowered to provide software earlier in cases where the urgency warrants bypassing potentially lengthy, deep, and complete integration test suites. Stakeholders can experiment with a variety of software stacks, leading to innovation.

Employ this simple DevOps framework to streamline the update process. It will provide system security, current functionality, environment reproducibility, and more reliable supercomputing operations, lowering deployment times from weeks to hours.

4:30pm - 5pm

Rev Up Compute Node Reboots: 2x to 5x Faster

Dennis Walker (HPE) and Paul Selwood (Met Office, UK / NERC CMS)

Abstract

Join us as we race to the bottom, showcasing the innovations developed to speed up reboot times by 300% for UK Met Office's latest, multi-zoned, CSM-based HPE Cray EX systems. With increasing complexity in software and site-specific customizations, node reboot times ballooned to over 35 minutes by November 2023, far exceeding operational requirements. In response, HPE developed automation to download logs, parse metrics, and graph boot stages by duration to understand better what was happening. Guided by data, the following changes were implemented, sorted by impact: - CFS/Ansible run-time plays were moved into local Systemd execution, running earlier in the boot cycle and without blocking other nodes. - Software installation and select configuration activity were moved into image build, streamlining deployment. - CSM boot settings were tuned for optimal performance. - Node Health Checks (NHC) were moved into Systemd, to run preemptively before the job scheduler agent to ensure nodes are consistently job-ready as early as possible.

As a result of the above changes, reboot times for system integration acceptance testing were reduced to 10-12 minutes as of July 2024; a

Sessions and Abstracts

300+% improvement.

3:00pm-5:00pm

Technical Session 2C: Climate applications **Newport III**

3pm - 3:30pm

Bit reproducibility in UK Met Office weather and climate applications

David Acreman (HPE)

Abstract

Weather and climate applications solve partial differential equations which are highly sensitive to small perturbations in model variables. Changes in even the least significant bit of a variable can have an observable impact on scientific results ("Butterfly effect"). The nature of floating point arithmetic means that subtle changes to code or the order of a summation can change results at the bit level due to different round-off errors. Consequently achieving bit reproducible results is challenging.

The UK Met Office Unified Model has been designed to give reproducible results even with different MPI and OpenMP parallelisations. However even when the same executable is used additional steps are required to achieve reproducible results. Differences in dynamically linked libraries have the potential to change results, which has important implications for updates to the Programming Environment and Operating System on operational weather and climate systems. An additional consideration is that the Met Office HPC systems comprise two phases with different generations of AMD EPYC processors. Achieving bit identical results on the two processor types also requires care.

We will describe the work undertaken to achieve bit reproducibility in this challenging scenario.

3:30pm - 4pm

Enabling km-scale coupled climate simulations with ICON on AMD GPUs

Jussi Enkovaara (CSC - IT Center for Science Ltd.)

Abstract

The Icosahedral Nonhydrostatic (ICON) weather and climate model is a modelling framework for numerical weather prediction and climate simulations. ICON is implemented mostly in Fortran 2008 with the GPU version based mainly on OpenACC. ICON is used on a large variety of hardware, ranging from classical CPU cluster to vector architecture and different GPU systems. In coupled simulations ICON can utilize heterogeneous architectures, i.e. ocean runs on CPUs while a atmosphere runs concurrently on GPUs.

An ICON model configuration developed for km-scale climate simulations is used as a scientific prototype for the digital twin of the Earth within the Destination Earth program of the European Comission. Here we focus on our effort to run these coupled ICON configurations at km-scale on LUMI, a HPE Cray system with a GPU partition based on AMD MI250x's.

Even though OpenACC is a hardware and compiler agnostic standard, code is not always directly portable due to differences in implementations. In this presentation we discuss the

various challenges and their resolutions encountered in porting ICON to LUMI. We discuss also profiling and optimization work, and present performance results, including strong and weak scaling of fully coupled climate simulations up to 2048 nodes on LUMI.

4pm - 4:30pm

MARBLChapel: Fortran-Chapel Interoperability in an Ocean Simulation

Brandon Neth and Ben Harshbarger (HPE); Scott Bachman ([C]Worthy); and Michelle Mills Strout (HPE, University of Arizona)

Abstract

As the climate crisis continues to have widespread effects on the biosphere, scientists increasingly turn to computer modeling to understand the impacts of different interventions. Modeling one such intervention, ocean carbon dioxide removal, requires incorporating multiple sources of interaction (air-sea gas exchange, biogeochemical processes, etc.) and high spatial and temporal resolutions. To address the need for scalable and high-resolution simulations, scientists at [C]Worthy have written the core of an ocean modeling code using Chapel, a parallel programming language for writing high-performance, distributed programs. Although Chapel has enabled rapid development, an important library for modeling biogeochemical processes, MARBL, is written in Fortran. MARBL is a robust, standalone library and is used in several state-of-the-art models including MOM6, MPAS, POP, and ROMS. Rather than rewriting the MARBL library in Chapel, we use Chapel and Fortran's C interoperability to integrate MARBL into the distributed Chapel simulation. This allows us to re-use reliable scientific code while using Chapel to orchestrate parallelism. In this talk, we demonstrate how the distributed Chapel simulation sets up data structures needed by MARBL, calls out to the Fortran library, and brings results back to update the simulation. We show performance results on Perlmutter, an HPE Cray EX system.

4:30pm - 5pm

Redefining Weather Forecasting Systems: The Transition to ICON and Alps

Matthias Kraushaar, Roberto Aielli, and Mauro Bianco (ETH Zurich); Oliver Fuhrer (Federal Office of Meteorology and Climatology MeteoSwiss); and Thomas Schulthess (ETH Zurich)

Abstract

The transition of MeteoSwiss operational weather forecasting from COSMO to the ICON model represents a major modernization in meteorological services, integrating software-defined infrastructure to improve flexibility, scalability, and resilience. The migration also involved significant hardware upgrades, from fixed systems with K80 GPUs to flexible architectures using V100 and later A100 GPUs, supported by the Cray-Shasta EX "Alps" infrastructure developed by CSCS.

The evolution included consolidating previously separate systems into a shared, cloud-like infrastructure with high-speed networking and advanced resource management using Cray System management (CSM), and Versatile Software Defined Cluster (vCluster) and User Environments (UENV) technologies. This approach allows the deployment of isolated and customizable configurations for various services on the infrastructure by minimizing

the overall complexity of the system. Geo-distributed infrastructure, by means of the AlpsE facility at EPFL, further improved resiliency, enabling updates without service disruptions and supporting critical applications like weather forecasting, disaster response, and research.

The transition not only upgraded technical capabilities but also established a foundation for scalable, future-proof meteorological services. This evolution ensures MeteoSwiss can meet the growing demands of modern weather and climate science.

6:00pm-8:00pm

HPE Networking Event

Lokal Eatery & Bar

HPE will host their annual CUG community networking reception from 6:00 to 8:00 pm ET at the Lokal Eatery & Bar. Lokal is located at 2 2nd St, Jersey City, NJ 07302, along Jersey City's waterfront, allowing CUG guests to enjoy expansive views of the Manhattan Skyline. Co-presented by AMD, all registered CUG attendees and their guests are invited to attend for a reception with light hors d'oeuvres and drinks. The Lokal is about a 10 min walk from the CUG hotel, but transportation will be provided if inclement weather.

Wednesday, May 7th

8:30am-10:00am

Plenary: CUG Updates (Open) and Best papers Newport III

10:00pm-10:30pm

Coffee Break (by VAST) Newport Foyer

10:30am - 10:35am

Unlocking Exascale Debugging and Performance Engineering with Linaro Forge

Rudy Shand (Linaro Ltd)

Abstract

Dive into the future of code development and see how Linaro Forge is reshaping what's possible in the world of parallel computing. Linaro Forge unveils the latest advancements: with Linaro DDT, MAP and Performance Reports, we're setting new standards in scalability and ease of use. Discover how these tools have become the go-to solution for developers seeking to push the boundaries of code optimization and performance engineering.

10:35am - 10:45am

Codee: A Tool to Enhance Correctness, Modernization, Security, Portability and Optimization in Fortran and C/C++ Software Applications

Manuel Arenaz (Codee)

Abstract

Fortran/C/C++ developers are under constant pressure to deliver increasingly complex simulation software that is correct, secure and fast. It is critical to empower development teams with tools to automate code reviews, enforce compliance with industry standards, and prioritize reducing the risk of security vulnerabilities. Codee features unique capabilities for Deep Analysis of Fortran/C/C++ code, helping to catch bugs, enforce coding guidelines, modernize legacy code, ensure code portability, address security vulnerabilities, and optimize code efficiency. Codee provides automated checkers for

the rules documented in the Open Catalog as well as AutoFix capabilities for semi-automatic source code rewriting, including modification of source code statements and insertion of OpenMP or OpenACC directives. Codee integrates seamlessly with popular editors, IDEs, Control Version systems and CI/CD frameworks, making it easy to incorporate into existing development workflows. Overall, Developers who are actively writing, modifying, testing and benchmarking Fortran code will increase their productivity by using Codee. Developers, team leaders and managers will benefit from DevOps and DevSecOps best practices, mitigating risks, boosting productivity, and reducing costs. In this presentation we will also talk about how to use Codee in conjunction with the Cray tools, including compilers (CCE) and performance tools (e.g. CrayPat, Reveal).

10:45am - 11:15am

The Unreasonable Effectiveness of FP64 Precision Arithmetic

Nicholas Malaya (AMD)

Abstract

Double precision datatypes, also known as FP64, has been a mainstay of high performance computing (HPC) for decades. Recent advances in AI have extensively leveraged reduced precision, such as FP16, or more recently, FP8 for Deepseek. Many HPC teams are now exploring mixed and reduced precision to see if significant speed-ups are possible in traditional scientific applications, including methods such as the Ozaki scheme for emulating FP64 matrix multiplications with INT8 datatypes. In this talk, we will discuss the opportunities, and significant challenges, in migrating from double precision to reduced precision. Ultimately, AMD believes a spectrum of precisions are necessary to support the full range of computational motifs in HPC, and that native FP64 remains necessary in the near future.

11:15am - 12:00pm

HPE 1 on 100 (HPE Customers only: no HPE partners or CUG sponsors)

Trish Damkroger (HPE)

Abstract

Session closed to only CUG sites, HPE representatives and sponsors not allowed.

12:00pm-1:00pm

Lunch (by Codee) Newport Foyer

1:00pm-2:30pm

Technical Session 3A: Data Centers Newport III

1:00pm-1:30pm

Causality inference for Digital Twins in GPU Data Centers and Smart Grids.

Rolando Pablo Hong Enriquez, Pavana Prakash, Ebad Taheri, and Aditya Dhakal (HPE); Matthias Maiterth and Wesley Brewer (Oak Ridge National Laboratory); and Dejan Milojicic (HPE)

Abstract

To the benefit of both technologies, data centers and smart grids will likely get evermore integrated in the near future. The downside is that effectively managing those systems will rapidly become burdensome if we neglect to prepare accordingly. Digital twins that can potentially wrap the benefits of advance analytics

and visualizations to manage such complex environments. Yet even today's AI systems lack the proper causal understanding of the data. Here we embark on a journey to collect proper causal data for validating causal inference methods based on three fundamentally different theoretical foundations: causal calculus, information theory, and dynamical system theory. Subsequently, we apply such methods to two target datasets from a smart grid and a GPU data center. We finally analyze the success and failures of applying these methodologies and the indications they offered to create more insightful and energy-efficient prediction strategies for digital twins in support of smart grids and GPU data centers.

1:30pm - 2:00pm

Co-design, deployment and operation of a Modular Data Centre (MDC) with air and direct-liquid cooled supercomputers

Sadaf Alam (University of Bristol); Emma Akinyemi, Martin Podstata, and Jan Over (HPE); and Simon McIntosh-Smith, Ross Barnes, Naomi Harris, and Dave Moore (University of Bristol)

Abstract

The Bristol Centre of Supercomputing (BriCS) has deployed its first HPE modular data centre (MDC), also known as a Performance Optimised Data Centre (POD), in March 2024. This has been a collaborative, co-design project between HPE and the University of Bristol. The MDC has enabled the rapid commencement of operations for the research community for the direct liquid cooled (DLC) Isambard-AI phase 1 (HPE Cray EX2500) and the air-cooled Isambard 3 (HPE Cray XD224), with NVIDIA Grace-Hopper and Grace-Grace superchips, respectively. The second set of MDCs have been deployed for Isambard-AI phase 2 containing 5,280 NVIDIA Grace-Hopper superchips in HPE Cray EX4000 DLC cabinets, together with the management and storage ecosystems. This manuscript outlines key features of the HPE POD MDCs for sustainability, efficiency, flexibility and observability in the era where data centre cooling and power needs are changing with growing demands for AI and HPC. We leverage the community efforts, specifically, the Energy Efficient High Performance Computing Working Group (EE HPC WG) that aims to sustainably support science through committed community action by encouraging the implementation of energy conservation measures and energy efficient design in HPC [1]. We outline notable advantages of the MDC approach for constraints and requirements that are unique for the Isambard-AI project that led to a co-design approach. We conclude by highlighting the key lessons drawn from this work.

2pm - 2:30pm

AlpsB – a Geographically Distributed Infrastructure to Facilitate Large-Scale Training of Weather and Climate AI Models

Alex Upton, Jerome Tissieres, and Maxime Martinasso (Swiss National Supercomputing Centre)

Abstract

AI-based models are transforming weather forecasting; these models are high quality and inexpensive to run compared to traditional physics-based models, and are already outperforming existing forecasting systems for many standard scores. The size of training datasets, however, remains a challenge. The widely-used ERA5, for example, is over 5PB, and these datasets are not typically located close to the large-scale compute power required for training AI models. As such, new solutions are required.

To address this, we present AlpsB, an HPE/Cray storage-focussed system, part of the Alps research infrastructure operated by CSCS. AlpsB is installed at the ECMWF data centre in Bologna, Italy, 250kms away from the CSCS data centre in Lugano, Switzerland. AlpsB is connected to CSCS via two dedicated, geo-redundant 100Gbps Ethernet links, and to ECMWF via four 100Gbps Ethernet links. AlpsB functions as a data hypercube, offering CSCS users high performance access to ECMWF data, including operational weather prediction and archive data. Additionally, an API-based query system deployed on AlpsB allows retrieval of specific dataset chunks, supporting different weather and climate use cases. Finally, we will present key use cases that have contributed to the co-design of AlpsB, and discuss the challenges of deploying a geographically distributed system.

1:00pm-1:30pm

Technical Session 3B:HPCM

Newport 1

A Brief Summary of the HPCM (HPE Performance Cluster Manager) Evolution Over Recent Releases

Sue Miller, Lee Morecroft, and Peter Guyan (HPE)

Abstract

This presentation will present the enhancements to the HPE Performance cluster manager (HPCM) over the current releases. It will cover 1.11 to 1.13 and patches to those versions as the release mechanisms for these enhancements.

There will be a brief summary of OS and hardware support followed by more detailed discussion of changes to the releases.

The enhancements discussed, amongst others, will include iSCSI as a root filesystem, a summary of changes to the monitoring infrastructure, repo priorities, the introduction of local post-install script variants, image activation whilst nodes are running and changes to the slot copy process. Various gluster improvements have been implemented including a move to using Ganesha as the NFS server, fixes for I/O errors when using an additional gluster export from the admin node, and changes to strengthen security. Improvements to the cm command line will be discussed as well as the new cpwrcli and mpwrcli.

The presentation will touch on forthcoming features where possible as 1.13 will have been recently released.

As there have been changes to the format of the virtual machine image format, there will be a discussion on backing up that as well as general cluster information backups.

1:30pm - 2:00pm

System Visualization Using Rackmap

Troy Dey and Peter Guyan (HPE)

Abstract

Understanding the health of an HPC system across various dimensions such as power, environment, fabric, and job performance is a challenging task, and continues to increase in difficulty as these systems become larger and more complex. To address this problem, the HPE Performance Cluster Manager (HPCM) now provides Rackmap an extensible CLI tool capable of rendering telemetry data as a dense 2D representation of the physical layout of components within an HPC system. This dense display of information within the CLI allows a system administrator to view, for example, the power status of thousands of nodes instantly, on a single screen, without having to context switch to a separate application. In addition, system administrators can easily create their own maps to display information of interest to them such as whether nodes have passed acceptance tests. This presentation will provide an overview of the Rackmap tool, describe the maps currently being shipped with HPCM, and go over how system administrators can create their own maps.

2:00pm - 2:30pm

Harvesting, Storing and Processing Data from our HPCM Systems

Ben Lenard, Eric Pershey, Brian Toonen, Peter Upton, Doug Waldron, Lisa Childers, Micheal Zhang, and Bryan Brickman (Argonne National Laboratory)

Abstract

With the Argonne Leadership Computing Facility (ALCF) acquiring more HPE Supercomputers, each equipped with its own HPCM stack, alongside other operational programs, we devised a strategy to centralize monitoring data from these systems. This centralized system aggregates data from various sources and securely distributes it to different consumers, including various teams and platforms within the ALCF.

Initially developed for the deployment of the Polaris supercomputer, this strategy was expanded to include the Aurora supercomputer. We now feed data into a Data Lake, enabling multi-stage processing and the application of code that utilizes bitmaps to quickly identify issues within the Aurora system.

The contributions made by this paper are as follows: *) Discuss the centralization of data from HPCM's Kafka bus to a secure and centralized location for multi-consumer access and the harvesting of log data from ephemeral nodes *) Highlight and show the value of a data warehouse using row and columar based storage *) Highlight the value of a data (Delta) lake for data processing as well as demonstrate the use of bitmaps and Apache Spark to efficiently search data generated by a supercomputer at Aurora's scale

While some of the concepts discussed are specific to HPCM systems, the principles of centralized data brokering using Kafka, and the use of data lakes and bitmaps, may be generalized to other platforms.

1:00pm-2:30pm

Technical Session 3C: Future Technology
Newport II

1pm - 1:30pm

Evolving Sarus to augment Podman for HPC on Cray EX

Alberto Madonna, Gwangmu Lee, and Felipe Cruz (Swiss National Supercomputing Centre)

Abstract

Podman provides modern and flexible containerization solution but lacks the specialized features required for high-performance computing. The evolution of the Sarus project aims to integrate Podman into a modular, open-source HPC container suite, bridging mainstream container technologies and supercomputing. This presentation highlights how Sarus deploys and optimizes Podman for HPC on CSCS's Alps infrastructure, a Cray EX system, focusing on the following areas: HPC-Optimized Podman: secure and scalable rootless containers for supercomputing environments with HPC-specific configuration templates. Workload Management Integration: seamless job orchestration of containerized workloads via SLURM-compatible SPANK component. Transparent HPC Resource Access: Open Container Initiative (OCI) hooks and Container Device Interface (CDI) provide pluggable access to compute, network, and storage resources on Cray EX systems. Parallel Filesystems Support: Squashfs-based image store for efficient usage of HPC storage systems. Secure Multi-Tenancy: rootless subid synchronization for Podman on shared distributed systems. This presentation will include test results on Alps, demonstrating how Sarus enables Podman to handle containerized job submissions efficiently and seamlessly. By augmenting community container tools like Podman to meet HPC needs, Sarus delivers a modern and flexible container stack optimized for CSCS's vClusters architecture on Cray EX systems.

1:30pm - 2:00pm

What is RISC-V and why should we care?

Nick Brown (EPCC)

Abstract

RISC-V is an open Instruction Set Architecture (ISA) standard which enables the open development of CPUs and a shared common software ecosystem. With billions of RISC-V cores already produced, and this is accelerating rapidly, we are seeing a revolution driven by open hardware. Nonetheless, for all the successes that RISC-V has enjoyed, it is yet to become mainstream in HPC. This comes at a time when HPC is facing new challenges especially around performance and sustainability of operations, and recent advances in RISC-V such as data centre RISC-V hardware make this technology a more realistic proposition with potential to address these. In this survey paper we explore the current state of art of RISC-V for HPC, identifying areas where RISC-V can benefit the HPC community, the level of maturity of the hardware and software ecosystem for HPC , and identify areas where the HPC community can contribute. The outcome is a set of recommendations around where the HPC and RISC-V communities can can come together and focus on high priority action points to help increase adoption.

2:00pm - 2:30pm

A Full Stack Framework for High Performance Quantum-Classical Comput-

ing

Xin Zhan, K. Grace Johnson, and Soumitra Chatterjee (HPE); Barbara Chapman (HPE, Stony Brook University); and Masoud Mohseni, Kirk Bresniker, and Ray Beausoleil (HPE)

Abstract

To address the growing needs for scalable distributed High Performance Computing (HPC) and Quantum Computing (QC) integration, we present our HPC-QC full stack framework and its hybrid workload development capability with modular hardware/device-agnostic software integration approach. The latest development in extensible interfaces for quantum programming, dispatching, and compilation within existing mature HPC programming environment are demonstrated. Our HPC-QC full stack enables high-level, portable invocation of quantum kernels from commercial quantum SDKs within HPC meta-program in compiled languages (C/C++ and Fortran) as well as Python through a quantum programming interface library extension. An adaptive circuit knitting hypervisor is being developed to partition large quantum circuits into sub-circuits that fit on smaller noisy quantum devices and classical simulators. At the lower-level, we leverage Cray LLVM-based compilation framework to transform and consume LLVM IR and Quantum IR (QIR) from commercial quantum software front-ends in a retargetable fashion to different hardware architectures. Several hybrid HPC-QC multi-node multi-CPU and GPU workloads (including solving linear system of equations, quantum optimization, and simulating quantum phase transitions) have been demonstrated on HPE EX supercomputers to illustrate functionality and execution viability for all three components developed so far. This work provides the framework for a unified quantum-classical programming environment built upon classical HPC software stack (compilers, libraries, parallel runtime and process scheduling).

2:30pm-3:00pm

Coffee Break (by Altair) **Newport Foyer**

3:00pm-5:00pm

Technical Session 4A: New Deployment **Newport III**

3:00pm - 3:30pm

A journey to provide GH200

Mark Klein, Thomas Schulthess, Jonathan Coles, and Miguel Gila (Swiss National Supercomputing Centre, ETH Zurich)

Abstract

Bringing a new hardware architecture to production involves a multi-phase process that ensures optimal performance, stability, and integration with existing infrastructure. The process begins with the physical installation of hardware and networking components. Once the hardware is in place, system engineers configure the operating system and necessary software stacks for performance monitoring and fault detection. This is followed by rigorous testing, including stress tests and benchmark runs, to verify the system's capabilities and identify any hardware or software anomalies.

The Swiss National Supercomputing Centre (CSCS) has recently expanded their Alps infrastructure with an additional 24 cabinets of Grace Hopper (GH200) blades. This pro-

cess has taken many months, and uncovered challenges related to the new hardware, along with existing software and infrastructure. This presentation will cover the discovery of these issues, temporary work-arounds, proper fixes, and system tunings required to bring the Grace Hopper blade architecture to production at CSCS.

Evaluation of the Nvidia Grace Superchip in the HPE/Cray XD Isambard 3 supercomputer

3:30pm - 4:00pm

Thomas Green (University of Bristol)

Abstract

The Bristol Centre for Supercomputing (BriCS) has recently deployed 55,296 Arm Neoverse V2 CPU cores in a supercomputing platform, via 384-nodes of NVIDIA Grace CPU Superchips with LPDDR5 memory as part of the Isambard 3 HPC service for the UK HPC research community. Isambard 3 is an HPE/Cray XD series system using the Slingshot 11 interconnect. As one of the first systems of this kind, this manuscript overviews details of the hardware and software configuration and presents early performance evaluation and benchmarking results using a representative subset of scientific applications. The focus is to evaluate Isambard 3 as a "plug-and-play" environment for researchers, especially who are familiar with Cray software environment. We include microbenchmark results to provide insights into the performance behaviour of this unique architecture. We present a small scale scaling comparison between the NVIDIA Grace CPU Superchip with other mainstream CPUs, including Intel Sapphire Rapids and AMD Genoa and Bergamo. We attempted to compare performance across several major software toolchains available for Arm, such as the HPE Cray Compiler Environment (CCE), the Arm Compiler for Linux, the NVIDIA Compiler, and GNU. Our findings include key opportunities for improvements that were discovered during our benchmarking, evaluation and regression testing on the system as we transitioned the service into operations from January 2025.

Alps, a versatile research infrastructure

4pm - 4:30pm

Maxime Martinasso (Swiss National Supercomputing Centre, ETH Zurich) and Mark Klein and Thomas Schulthess (Swiss National Supercomputing Centre)

Abstract

The Swiss National Supercomputing Centre (CSCS) has a long-standing tradition of delivering top-tier high-performance computing systems, exemplified by the Piz Daint supercomputer. However, the increasing diversity of scientific needs has exposed limitations in traditional vertically integrated HPC architectures, which often lack flexibility and composability. To address these challenges, CSCS developed Alps, a next-generation HPC infrastructure designed with a transformative principle: resources operate as independent endpoints within a high-speed network. This architecture enables the creation of independent tenant-specific and platform-specific services, tailored to diverse scientific requirements.

Alps incorporates heterogeneous hardware, including CPUs and GPUs, interconnected by a high-performance Slingshot network, and

offers a modular storage system. A key innovation is the versatile software-defined cluster (vCluster) technology, which bridges cloud and HPC paradigms. By abstracting infrastructure, service management, and user environments into distinct layers, vClusters allow for customized platforms that support diverse workloads. Current platforms on Alps serve various scientific domains, including numerical weather prediction, and AI research.

Evaluating AMD's MI300A APU: Performance Insights on LLM Training via Knowledge Distillation

4:30pm - 5:00pm

Dennis Dickmann (Seedbox); Philipp Offenhäuser (HPE); Rishabh Saxena (HLRS, University of Stuttgart); George Markomanolis (AMD); Alessandro Rigazzi (HPE HPC/AI EMEA Research Lab); Patrick Keller (HPE); and Kerem Kayabay and Dennis Hoppe (HLRS, University of Stuttgart)

Abstract

AMD (Advanced Micro Devices) has recently launched the MI300A Accelerated Processing Unit (APU), which integrates Central Processing unit (CPU) and Graphical Processing Unit (GPU) compute in a single chip with unified High Bandwidth Memory (HBM). This study assesses the AMD Instinct™ MI300A capabilities and examines how this new architecture handles real-world generative Artificial Intelligence (AI) workloads. While performance data for Large Language Model (LLM) use cases exists for the MI250X and MI300X, to the best of our knowledge, such assessments are absent for the MI300A. We apply a Knowledge Distillation (KD) use case to distil the knowledge of the Mistral-Small-24B-Base-2501 teacher model to a student model that is 50% sparse using a 2:4 sparsity pattern. The results show quasi-linear scaling of raw performances on up to 256 APUs. Lastly, we discuss the challenges, research and practical implications, and outlook.

3:00pm - 5:00pm

Technical Session 4B: GPU Energy Efficiency — Newport I

Optimizing GPU Frequency for Sustainable HPC: Lessons Learned from a Year of Production on Adastra, an AMD GPU Supercomputer

3pm - 3:30pm

Gabriel Hautreux, Naïma Alaoui, and Etienne Malaboeuf (CINES)

Abstract

Power consumption is a critical concern for GPU-based high-performance computing (HPC) systems as rising energy costs and environmental challenges push for energy-efficient solutions. Modern GPUs, such as the AMD MI250X used in the Adastra supercomputer, offer features like frequency scaling to manage power consumption dynamically. However, optimizing frequency configurations for diverse HPC and AI workloads is complex due to their varied computational demands. CINES, operating the Adastra system in Montpellier, France, conducted a study analyzing the impact of reducing the GPU frequency from 1.7 GHz to 1.5 GHz. Adastra, ranked #3 on the Green500 list of energy-efficient supercomputers, supports French researchers across scientific domains.

Sessions and Abstracts

Previous findings presented at SuperComputing 23 showed that frequency downscaling improved energy efficiency while slightly impacting performance, prompting CINES to adopt the lower frequency in July 2024. A year-long analysis revealed a 15% reduction in energy consumption per node, aligning with sustainability goals without requiring hardware modifications. The study also assessed application performance, user satisfaction, hardware reliability, and differences between HPC and AI workloads. The results provide actionable insights for HPC centers aiming to enhance energy efficiency while avoiding the complexity and overhead of dynamic strategies.

Fine-Grained Application Energy and Power Measurements on the Frontier Exascale System

3:30pm - 4pm

Oscar Hernandez and Wael Elwasif (Oak Ridge National Laboratory)

Abstract

The increasing complexity and power/energy demands of heterogeneous exascale systems, such as the Frontier supercomputers, present significant challenges for measuring and optimizing power consumption in applications. Current tools either lack the resolution to capture fine-grained power and energy measurements or fail to integrate this information with application performance events. This paper introduces a novel open-source performance toolkit that integrates extended PAPI components with Score-P to enable fine-grained millisecond-level power and energy measurements for AMD MI-250x GPUs and CPUs.

Our toolkit combines coarse-grained measurements from Cray PM counters with high-resolution metrics from rocm_smi converting instantaneous accumulated energy to power to capture both transient and steady-state power behavior. By mapping these metrics to specific application regions, developers can identify energy hotspots, address inefficiencies in GPU kernel execution, and optimize resource usage. We demonstrate the effectiveness of this approach through case studies using benchmarks like rocblas_sgemm tests and applications, highlighting GPU variability and the impact of transient power spikes on kernel-level efficiency.

EVeREST: An Effective and Versatile Runtime Energy Saving Tool for GPUs

4pm - 4:30pm

Anna Yue, Torsten Wilde, Sanyam Mehta, and Barbara Chapman (HPE)

Abstract

The widespread adoption of GPUs combined with the significant power consumption of GPU applications prepares a strong case for an effective power/energy saving tool for GPUs. Interestingly, however, GPUs present unique challenges (that are traditionally not seen in CPUs) towards this goal, such as very few available low-overhead performance counters and fewer optimization opportunities. We propose Everest, a proof-of-concept tool that dynamically characterizes applications to find novel and effective opportunities for power and energy savings while providing desired performance guarantees. Specifically, Everest finds two unique avenues for saving energy using

DVFS (Dynamic Voltage Frequency Scaling) in GPUs in addition to the traditional method of lowering core clock for memory bound phases. Everest does not require any application modification or apriori profiling and has very low overhead. Everest relies on a single chosen performance event that is available across both AMD and NVIDIA GPUs that we show to be sufficient and effective in application characterization, which also makes Everest portable across GPU vendors. Experimental results of our PoC across 8 HPC and AI workloads demonstrate up to 25% energy savings while maintaining 90% performance relative to the maximum application performance, outperforming existing solutions on the latest NVIDIA and AMD GPUs.

HPE Cray EX225a (MI300a) Blade Power Capping and HBM Page Retirement

4:30pm - 5pm

Steven Martin, Randy Law, Leo Flores, Ron Urwin, and Larry Kaplan (HPE)

Abstract

HPE Cray Supercomputing EX255a is a density-optimized, accelerated blade featuring eight AMD MI300a Accelerated Processing Units (APUs). To deploy the HPE Cray EX255a blades at maximum density in the HPE Cray EX4000 cabinet, the nodes need to be power capped to limit total cabinet power to the 400KVA maximum cabinet power constraint. Managing node-level power to enforce the cabinet-level constraint while maximizing node-, cabinet-, and system-level performance drove the engineering team to a new power-capping design that will be described in this presentation. This new power capping design is configured out-of-band via Redfish and is complementary to in-band capping that can be configured via rocm-smi. This presentation will show power and performance data collected on a large customer system and from a smaller system internal to HPE. This design is expected to be leveraged for future HPE Cray EX blades.

AMD MI300a APUs support High Bandwidth Memory (HBM) page retirement, but do not provide an on-socket data store for the required data to persist across reboots, or power cycles. The presentation will cover the motivation and implementation details for the HPE Cray EX225a blade, including data from a large customer deployment.

Technical Session 4C: Monitoring — Newport II

LLM Serving With Efficient KV-Cache Management Using Triggered Operations

3pm - 3:30pm

Nikolay A. Simakov, Joseph P. White, and Matthew D. Jones (SUNY University at Buffalo) and Eva Siegmann, David Carlson, and Robert J. Harrison (Stony Brook University)

Abstract

High-Performance Computing (HPC) resources are essential for scientific and engineering computations but come with substantial initial and operational costs. Therefore, ensuring their optimal utilization throughout their lifecycle is crucial. Monitoring utilization and performance helps maintain efficiency and proactively address user needs and performance issues. This is particularly important for technological

testbed systems, where frequent software updates can mask localized performance degradations with improvements elsewhere.

This presentation shares our experience monitoring Ookami, an HPE Apollo 80 system, using XDMoD. Ookami features the ARM-based Fujitsu A64FX CPU, the first to support SVE-512. After more than three years of monitoring, we found that while the total number of jobs remained relatively constant, total CPU-hours consumed by users steadily increased. This suggests that shorter test runs are being replaced by longer test runs and, in some cases, production runs. Another find is that application and numerical library performance saw the most significant improvements with the initial release of new technology support, followed by smaller incremental gains. Additionally, we observed that vendors with significant investment optimize numerical libraries more quickly than community-driven efforts.

HPE Slingshot Monitoring Software: Actionable Insights for HPC and AI Systems

3:30pm - 4pm

Sahil Patel (HPE)

Abstract

Modern HPC and AI systems produce vast telemetry data, making performance monitoring and root cause analysis increasingly challenging. Traditional troubleshooting methods often lead to inefficiencies, lengthy resolution times, and costly downtime, unable to meet the demands of today's high-performance computing environments.

HPE Slingshot Monitoring Software (SMS) revolutionizes this process with cutting-edge automation and real-time insights. By analyzing data comprehensively and providing intuitive visualizations, it offers clear insights into system performance and fabric health, along with actionable recommendations. This empowers system administrators to rapidly identify and resolve issues, minimizing manual intervention while improving diagnostic accuracy and speed.

This presentation will explore how SMS optimizes fabric utilization and enhances system reliability. We'll begin with its role in translating complex telemetry into actionable insights, then examine its ability to pinpoint bottlenecks and recommend corrective actions. Finally, we'll discuss best practices for preventing downtime and streamlining troubleshooting.

Paired with the System Diagnostic Utility (SDU), HPE Slingshot Monitoring Software delivers a robust solution. While SMS provides real-time insights and visualization, SDU automates capturing and transmitting critical configuration and performance data, enabling proactive identification of unhealthy components and bugs. Together, they set a new standard for efficiency, reliability, and innovation in HPC and AI systems.

LDMS New Features for Deployment in Advanced Environments and Feedback for Operations

4pm - 4:30pm

Jim Brandt, Ben Schwaller, Jennifer Green, Ben Allan, Cory Lueninghoener, Evan Donato, Vanessa Surjadidjaja, Sara Walton, and Ann Gentile (Sandia National Laboratories)

Abstract

Sessions and Abstracts

The Lightweight Distributed Metric Service (LDMS) monitoring, transport, and analysis framework has been deployed on large-scale Cray and HPE systems for over a decade. Over that time its capabilities have improved dramatically. In this talk we provide updates on capabilities including deployment and management methods in bare metal, containerized, and cloud (including hybrid on+off prem) environments. We describe how LDMS is being used to collect application data concurrent with system data and how the low-latency availability of this data for analysis can be used for real-time data analysis and feedback in order to support efficient, resilient, and reliable system operations. Finally, we will describe current related research areas including 1) use of machine learning for modeling application and system behavioral characteristics and 2) use of new features in the bi-directional communication capability of LDMS to provide low-latency communication and feedback from a distributed analysis system to user, system, and application processes on disparate clusters and to inform data center orchestration decisions.

On the product development, management, and distribution fronts we will present our planned improvements over the next year, release cadence, and package distribution methods including how we plan to stay in sync on HPE's CSM and HPCM releases.

Proactive Health Monitoring and Maintenance of High-Speed Slingshot Fabrics in HPC Environments

4:30pm - 5pm

Michael Cush, Jeff Kabel, Michael Schmit, Michael Accola, and Forest Godfrey (HPE)

Abstract

This whitepaper addresses the critical need for maintaining the health of high-speed Slingshot fabrics in high-performance computing (HPC) environments. Identifying and resolving known issues swiftly is essential for optimizing HPC workload performance, yet pinpointing common and emerging problems can be highly challenging. We propose a proactive solution that leverages automated capture of key configuration and performance metrics, coupled with sophisticated event logic, to detect unhealthy components/known bugs within the fabric. This is achieved through the System Diagnostic Utility (SDU), integrated with HPCM and CSM software, which automates data capture and securely transmits it to HPE using HPE Remote Device Access (RDA). This solution is complimentary to other system monitoring solutions such as SMS and AIOps. In fact, this solution can also capture data from those tools to consolidate and enhance the level of data captured for analysis.

The impact of this approach is significant: it enables faster identification and resolution of issues, thereby enhancing fabric performance and overall HPC job efficiency. Furthermore, the ability to analyze both historical and current data allows for real-time rule additions to address newly discovered bugs. The comprehensive visibility into issues across customer environments also aids HPE R&D in developing timely fixes and enhancements.

5:05pm-5:45pm

BoF 2A **Newport III**

CPE Futures

Barbara Chapman (HPE, Stony Brook Uni-

versity) and Kaylie Anderson and Chun Sun (HPE)

Abstract

The HPE Cray Programming Environment (CPE) provides a suite of integrated programming tools for application development on a diverse range of HPC systems delivered by HPE, including those with integrated node architectures or whose nodes are configured with AMD or NVIDIA GPUs. Its compilers, math libraries, communications libraries, debuggers, and performance tools enable the creation, enhancement and optimization of application codes written using mainstream HPC programming languages and the most widely used parallel programming models.

Supercomputing workloads are rapidly evolving, with an emphasis on workflows, as well as the rapid integration of AI into – and deployment alongside - traditional simulations, and more focus on large-scale data analytics. CPE is undergoing several changes that will enable it to better meet the needs of these new workloads, support emerging hardware platforms, and accommodate the requirements of a more diverse user base. In this interactive session, we will describe our plans for a future CPE and seek feedback to guide us in this transition.

BoF 2B **Newport I**

Managing System Reliability: From system acceptance through production

Pete Guyan and Sue Miller (HPE)

Abstract

Managing System Reliability: From system acceptance through production can be a messy situation. This BoF will start with a synopsis of what HPE plans to do in the short, mid and long term to make a cohesive strategy for reproducibility. We will discuss the current tools in HPE Performance Cluster Manager [HPCM] and how we weave these tools into the strategy, how they interact with monitoring and reporting tools.

We'd like to start conversations with customers on their own strategies and to address these questions: What makes them successful? How do they track issues? How do they track success? What can HPE do to help to make acceptances simpler and give the user community confidence in the systems? Which tools do they use and are confident in? Are there tools missing that need to be created?

While HPE will be leading this BoF, we will be asking a couple of customers we know are attending to contribute to the presentation to help start the conversation.

BoF 2C **Newport II**

HPE Slingshot Birds-of-a-Feather

Jesse Treger (HPE)

Abstract

This birds-of-a-feather session will provide an opportunity for users to ask questions and share advice on managing and using HPE Slingshot systems, as well as to hear and provide input into HPE Slingshot's software roadmap. HPE Slingshot software scope covers capabilities both for the administrators who operate and manage the system and the fabric, and for HPC and AI application writers/users of the HPE Slingshot NIC's Libfabric provider. Users will be encouraged to share desired use cases, learnings, and best-known methods.

Sessions and Abstracts

6:00pm-10:00pm

CUG Night Out **Hudson House**

We invite all registered attendees and guests with a paid CUG night out ticket to join us for an unforgettable evening at Hudson House. Situated at the end of Port Liberty in Jersey City, NJ, this structure is an arms' length away from the Hudson River and boasts a panoramic view of the Statue of Liberty, Brooklyn, Manhattan, and Verrazano Bridges, and of course the NYC Skyline. Coaches will depart outside the Westin Jersey City Hotel at 6:15pm to arrive at Hudson House for a drink's reception before seating for dinner at approximately 7:15pm. If you are making your own way to the venue, please use the full address as Google Maps takes you to a different address! Hudson House, 2 Chapel Ave is approx. a 15 – 20-minute drive. Our first bus will return to the hotel at approximately 9:00pm with shuttles then running until approximately 10:30pm.

Thursday, May 8th

8:30am-10:00am

Plenary: CUG 2026, Panel **Newport III**

8:30am-8:40am

HPC Activities at Austrian Weather Service (GeoSphere)

Martin Shivraj Saini (Geosphere)

8.40-8.50

New Member Site: Introducing Cyfronet

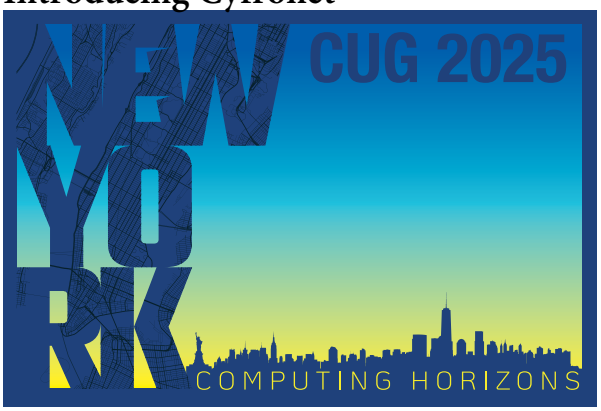

Patryk Lasoń (Academic Computer Centre Cyfronet AGH)

8:50am-8:55am

VAST Data Platform

Jan Heichler (VAST)

Abstract

The VAST Data Platform leverages its Disaggregated Shared Everything (DASE) architecture to seamlessly unify HPC and AI data management, offering a data platform that prioritizes high speed, scalability, and simplicity. By eliminating downtime during system upgrades and streamlining parallel I/O, VAST ensures continuous and efficient operation at any scale. Its support for multiple data protocols facilitates effortless integration into existing infrastructures, requiring no custom tuning or complex configurations. This session will explore how the VAST Data Platform addresses the evolving needs of modern technical environments, enabling enhanced performance and operational efficiency.

8:55am - 9:05am

CUG 2026 site presentation

TBD

9:00am-10:00am

Panel on : The Future of Precision in HPC, which FP is the Right One?

Ashley Barker (Oak Ridge National Laboratory)

Abstract

This panel will explore the evolving role of floating-point precision in high-performance computing (HPC) and AI workloads, analyzing the trade-offs between FP64 and emerging alternatives such as mixed-precision techniques, lower-precision formats, and emulation methods. The panelists will share a nuanced discussion on balancing accuracy, performance, and

hardware constraints in modern HPC and AI workloads.

Moderator: Robert Wisniewski, HPE Fellow and Chief Architect in the HPC and AI Solutions Organization at HPE Panelists: Membership Computing Facility; Nic Malaya, Fellow, HPC, AMD Dan Ernst, Director of Super-computing System Technology at Nvidia; Piotr Luszczek,
Bronson Messer, Director of Science, Oak Ridge Leadership Computing Facility Nic Malaya, Fellow, HPC, AMD Dan Ernst, Director of Supercomputing System Technology at Nvidia Piotr Luszczek, Technical Staff at MIT Lincoln Lab and Research Associate Professor, University of Tennessee.

10:00am - 10:30am

Coffee Break Newport Foyer

10:30am - 12:00pm

Technical Session 5A: Slingshot & MPI Tuning Newport III

10:30am-11:00am

MPI implementation optimization for Slingshot network

Rahulkumar Gayatri, Adam Lavely, Neil Mehta, and Brandon Cook (Lawrence Berkeley National Laboratory)

Abstract

Optimizing MPI performance is one of the keys to improving performance of HPC applications. While algorithmic improvements such as overlap of communication and computation are used to improve MPI parallelism within a workflow, the choice of MPI implementation for a given application can also affect overall performance. We have characterized how OpenMPI, MPICH, and Cray MPICH perform on Perlmutter for small to moderate problem sizes using the OSU micro-benchmarks and have shown that the vendor-tuned Cray MPICH typically outperforms the other MPI implementations. We plan to expand this work for additional microbenchmarks, identify ways to improve the MPI implementations, and show how MPI performance differences impact the performance of full HPC applications for the final paper.

11:00am - 11:30am

Using Different MPI Implementations on the LUMI Supercomputer for Native and Containerized Applications Execution

Maciej Szpindler (Academic Computer Centre CYFRONET), Marcin Krotkiewski (University of Oslo), Alfio Lazzaro (HPE), and Maciej Pawlik (Academic Computer Centre CYFRONET)

Abstract

Message Passing Interface (MPI) implementations have to be tailored for specific system architectures to maximize application performance. This applies to optimizations for network transport and provision of efficient data movement between CPU and GPU memories. The default MPI on HPE Cray Ex systems such as LUMI is a proprietary and optimized implementation based on the open source MPICH Ch4 implementation. There are situations where having an alternative such as OpenMPI benefit users who have applications or con-

tainers that are targeting OpenMPI and where some effort would be needed to change. Having alternative MPI implementations allows for performance comparisons, investigating bugs, and checking new MPI functionalities. In this paper, we report on our experience with installing containerized and native OpenMPI environments on LUMI, showing how users can build and run containers and get the expected performance. We show a performance comparison with respect to HPE Cray MPI executions using OSU benchmarks and an example real-world application for the solution of Dirac equations using GPUs. Although we only refer to LUMI, similar concepts can be applied to the case of other supercomputers.

11:30am - 12:00pm

Scaling MPI Applications on Aurora

Nilakantan Mahadevan (HPE); Premanand Sakarda (Intel Corporation); and Scott Parker, Servesh Muralidharan, Longfei Gao, Kris Rowe, Vitali Morozov, and Victor Anisimov (Argonne National Laboratory)

Abstract

Aurora supercomputer at Argonne National Laboratory is one of the three Exascale machines in the world. Aurora is s the fastest AI system in the world dedicated to AI and HPC simulations for open science. Aurora integrates world-class compute capabilities with powerful memory, storage, and networking to advance the frontiers of science. Aurora machine consists of 10,624 Compute Blades, 21,248 Intel Xeon CPU Max Series, 63,744 Intel Data Center GPU Max Series connected using Slingshot high-performance fabric interconnect. Aurora fabric is the largest fabric with 5600 switches connected in a dragonfly topology. This session will cover the various aspects, challenges, and learnings in utilizing the Aurora network to scale MPI applications and include industry standard network benchmarks on the Aurora system. We will analyze and discuss results of different MPI applications and benchmarks at various scales. System level optimizations are critical to deliver performance at scale for supercomputers. On Aurora, as node count increases, the performance for throughput, latency and bandwidth across compute, memory, and network subsystems is critically important for scaling HPC and AI applications.

Technical Session 5B: Regression Testing — Newport I

10:30am-11:00am

Hardware Triage Tool: Enhancements and Extensions

Isa Muhammad Wazirzada, Abhishek Mehta, Vinanti Phadke, and Bhuvan Meda Rajesh (HPE)

Abstract

In 2023, HPE released the Hardware Triage Tool (HTT) with the mission to provide high-fidelity diagnoses and minimize time to repair of hardware faults across HPE Cray EX compute and accelerator blades irrespective of the system manager being used.

The past 12 months have seen an increase in the adoption of HTT, which has not only led to efficient system deployments but also helped customers maintain healthy node counts.

Our presentation focuses on the enhancements made to HTT since its introduction. This includes enhancements in the areas of

diagnosability, usability, logging, and customer requested additions as well. A major highlight includes the ability to analyze Common Platform Error Record (CPER) files that are generated at the time of a hardware fault. CPERs are a part of the Unified Extensible Firmware Interface (UEFI) standard and give hardware vendors the ability to serialize hardware error information through a standard format. In future releases HTT will be able to decode CPER files for AMD MI300 based EX255a accelerator blades and provide accurate repair actions by cross-referencing the AMD service action guide.

We will also unveil a set of new features that extend HTT to support non HPE Cray EX hardware.

11:00am-11:30am

Detecting operating system noise with detect-detour

Nagaraju KN, Clark Snyder, Dean Roe, and Larry Kaplan (HPE)

Abstract

HPC applications, especially those that frequently perform global synchronization operations, can be negatively affected by background operating system (OS) activity. The background actions of interest are processing hardware interrupts, software interrupts, and process context switches. While these actions are necessary to the operation of the OS, from the application's point of view, they are viewed as "OS noise" that affects performance, and the system should be tuned to minimize them. Identifying sources of OS noise is crucial for application performance but can be difficult. Few options exist to identify sources of OS noise without getting into the intricacies of the underlying kernel internals. The detect-detour tool makes use of the Linux kernel enhanced Berkeley Packet Filter (eBPF) [1] feature to help system administrators identify sources of OS noise without requiring them to be kernel experts.

11:30am-11:00pm

Analyzing a Lifetime of Failures on a Cray XC40 Supercomputer

Kevin Brown and Tanwi Mallick (Argonne National Laboratory), Zhiling Lan (University of Illinois Chicago), Robert Ross (Argonne National Laboratory), and Christopher Carothers (Rensselaer Polytechnic Institute)

Abstract

We analyze hardware errors over the seven-year lifetime of the Theta supercomputer, a large-scale Cray XC40 system at the Argonne Leadership Computing Facility. To ensure accurate interpretation of the logs, we leverage expert knowledge to clean the dataset and remove redundant information. Temporal and spatial analysis techniques are then used to expose how failures and errors trend over time and across components in the system. Additionally, we correlate hardware error logs to system downtime logs to capture the relationship between critical errors and outages over the lifetime of the system. The results in this work represent a state-of-the-practice report highlighting how severe error types vary over time and across different component types, such as on-node and off-node (network) components. We also demonstrate the effectiveness of our technique in simplifying log analysis by using a unified error classification across

components from different vendors, providing valuable insights into normal and anomalous system behaviors.

Technical Session 5C: Filesystems & I/O **Newport II**

10:30am-11:00am

E2000 Performance From Microbenchmarks to Applications

William Loewe, Michael Moore, Sakib Samar, and Chris Walker (HPE)

Abstract

With the advance of the Exascale Age and its continued gains in FLOPS performance, the associated I/O demands of performance increase commensurately. To address this, the HPE Cray Supercomputing Storage Systems E2000 is the next generation of the HPE Cray Supercomputing Storage product line with a focus on performance. This paper discusses the architecture changes in the E2000 and provides node and file system microbenchmarks measuring bandwidth, IOPS, and metadata performance. The improved PCIe and NVMe drive speeds in addition to the higher density enclosure in the E2000-F allow for more than twice the throughput and IOPS performance compared to the previous generation with nearly all of the performance achievable by optimal application workloads. System configuration choices, such as number of storage targets and BIOS settings, which influence system level performance will be compared with an aim to optimize the gains and determine ideal client/server tunings. Finally, performance of application-relevant workloads including random access, shared file, and AI/ML storage workloads will be presented along with discussion of application and job changes to utilize the E2000 performance improvements.

11am-11:30am

Towards Empirical Roofline Modeling of Distributed Data Services: Mapping the Boundaries of RPC Throughput

Philip Carns, Matthieu Dorier, Rob Latham, Shane Snyder, and Amal Gueroudji (Argonne National Laboratory); Seth Ockerman (University of Wisconsin-Madison); Jerome Soumagne (HPE); Dong Dai (University of Delaware); and Robert Ross (Argonne National Laboratory)

Abstract

The scientific computing community relies on distributed data services to augment file systems and decouple data management functionality from applications. These services may be native to HPC or adapted from cloud environments, and they encompass diverse use cases such as domain-specific indexing, in situ analytics, AI data orchestration, and special-purpose file systems. They unlock new levels of performance and productivity, but also introduce new tuning challenges. In particular, how do practitioners assess performance, select deployment footprints, and ensure that services reach their full potential? Roofline models could address these challenges by setting practical performance expectations and providing guidance to achieve them.

This paper outlines initial steps towards establishing an empirical roofline modeling methodology for distributed data services, focusing exclusively on network characteristics as a proof of concept. We first explore how to max-

imize performance on modern platforms. We then propose an adaptation of the classic roofline model and a methodology for collecting model parameters using the Mochi data service framework. Finally, we survey four large-scale high performance computing systems and construct a roofline model for each of them. We evaluate the models and identify next steps towards a comprehensive roofline modeling framework for distributed data services.

11:30am-12:00pm

HPC workload characterization using eBPF

Shubh Pachchigar and Brandon Cook (Lawrence Berkeley National Laboratory/National Energy Research Scientific Computing Center)

Abstract

Efficient interactions with filesystems are essential for scientific workflows operating at scale on HPC systems. In order to design new filesystems, tune system configurations, effective I/O characterization is needed. Darshan is a widely used tool for I/O characterization that relies on injecting code into application binaries but has some limitations in providing low-level insights. In this work, we propose leveraging eBPF that enables the execution of user-defined programs within the kernel, to develop a new I/O characterization tool. Our approach aims to complement the capabilities of Darshan by using eBPF to gain deeper insight into application interactions with the underlying filesystems. This is achieved by deploying dynamic instrumentation techniques below the application layer to extract detailed I/O metrics. In this work, we demonstrate the collection of read/write operations, and their associated latencies with various available filesystems. The metrics are periodically sampled with a custom eBPF-based LDMS sampler to enable the collection of data at scale. Finally, to demonstrate its feasibility in production HPC environments, we establish that the overhead of the tool is low. This work demonstrates the potential of eBPF to enhance I/O characterization in HPC environments, providing valuable insights that can lead to improved performance and resource utilization.

12:00pm-1:00pm

Lunch (by Nvidia) Newport Foyer

1:00pm-2:30pm

Technical Session 6A: DAOS Newport III

1pm - 1:30pm

Enhancing RPC on Slingshot for Aurora's DAOS Storage System

Jerome Soumagne, Alexander Oganezov, Ian Ziemba, and Steve Welch (HPE); Philip Carns and Kevin Harms (Argonne National Laboratory); and John Carrier, Johann Lombardi, Mohamad Chaarawi, Zhen Liang, and Scott Peirce (HPE)

Abstract

DAOS, an open-source software-defined high-performance storage solution designed for massively distributed NVMe SSDs and Non

Volatile Memory (NVM), is a key component of the Aurora Exascale system that aims to deliver high storage throughput and low latency to application users. Utilizing the Slingshot interconnect, DAOS leverages Remote Procedure Call (RPC) to communicate between compute and storage nodes. While the preexisting RPC mechanism used by DAOS was already designed for High-Performance Computing (HPC) fabrics, it required a number of scalability, performance, and security enhancements in order to be successfully deployed on Aurora.

We present and discuss in this paper the improvements that were made to address this set of challenges in each of the components that DAOS relies upon: the Collective and RPC Transport (CaRT) layer, the Mercury RPC library and the libfabric Slingshot (cxi) provider. While some of these enhancements were tailored to the DAOS storage system specifically, they also serve as a broader reference for implementing scalable high-performance data services over RDMA fabrics. This paper focuses on the design and implementation of each functionality that was improved, enabling DAOS to provide both high throughput and low latency to Aurora users.

1:30pm - 2pm

Global Distributed Client-side Cache for DAOS

SClarete R. Crasta, John L Byrne, Abhishek Dwaraki, David Emberson, Harumi Kuno, Sekwon Lee, Ramya Ahobala Rao, Shreyas Vinayaka Basri K S, Amitha C, Chinmay Ghosh, Rishi Kesh Kumar Rajak, Sriram Ravishankar, Porno Shome, and Lance Evans (HPE)

Abstract

HPC/AI workloads process large amounts of data and perform complex operations on the data at exascale rates, for time-critical insights/results. Distributed workloads are often bottlenecked by communication when storage systems are used to co-ordinate and share results. Storage solutions supporting effective, scalable parallel access from compute clusters are critical to HPC architectures. Caching data on storage servers and/or clients are known techniques used by storage systems to ameliorate the communication costs. Current server-side caching methodologies are constrained by amount of memory and network bandwidth on the fixed and finite server nodes. Furthermore, most client-side caches are node-local, meaning the cached data is accessible solely by the node on which the data is stored. DAOS is a promising exascale storage stack recently acquired by HPE. Global client-side caching for DAOS is an attractive proposition due to higher aggregate client-side resources (e.g., DRAM and network bandwidth) that can scale independent of the number of server nodes. In addition to providing faster data access, a client-side cache should also be efficient as it consumes expensive resources and requires an efficient caching framework with its associated policies. In this paper, we cover the details of realizing efficient shared client-side caching for DAOS.

2pm - 2:30pm

DAOS - New Horizons for High Performance Storage

Michael Hennecke and Jerome Soumagne (HPE)

Abstract

DAOS is an open-source scale-out storage sys-

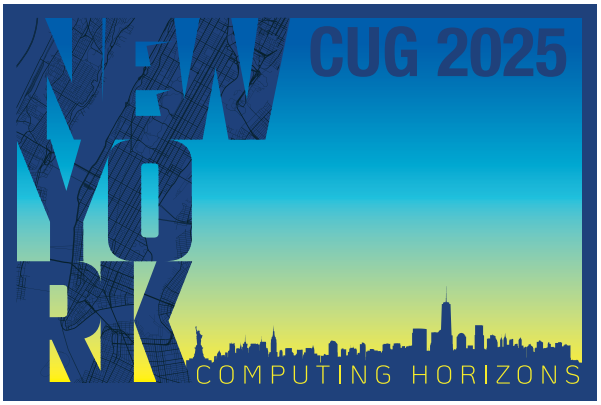

tem that has redefined performance for a wide spectrum of HPC and AI workloads (https://daos.io/). It is an all-flash solution that can be deployed as a stand-alone storage system, or it can be a high-performance storage tier used in combination with traditional Lustre, GPFS, or (cloud) object storage environments.

DAOS is the primary storage system for the Aurora supercomputer. With the creation of the DAOS Foundation in 2023 and the support of PMem-less DAOS servers, DAOS has seen increasing community contributions and growing adoption in both on-prem and cloud environments. The DAOS software engineering team has recently transitioned from Intel to HPE, complementing the existing HPE Cray Supercomputing Storage Systems engineering team and product portfolio.

This presentation will give an overview of the DAOS community roadmap, product plans for future DAOS-based HPE Cray Supercomputing Storage Systems, and other ongoing research and development activities. We will also discuss the APIs available to the application programmer to efficiently use DAOS. This includes legacy POSIX support, the DAOS Filesystem (DFS) API, DAOS backends for MPI-IO, HDF5, Tensorflow-IO and other middleware, as well as the native DAOS key-value/key-array API and Python bindings.

1:00pm-2:30pm

Technical Session 6B: Framework for HPC-AI workflows Newport I

Framework for tracking metadata, lineage and model provenance in hybrid simulation-AI HPC exascale workflows

1pm - 1:30pm

Martin Foltin, Andrew Shao, Rishabh Sharma, Shreyas Kulkarni, Annmary Justine Koomthanam, Aalap Tripathy, and Cong Xu (HPE); Wenqian Dong (Oregon State University); Suparna Bhattacharya (HPE); Brian Sammuli (General Atomics); and Paolo Faraboschi (HPE)

Abstract

Integration of AI in HPC workflows can have a profound impact on HPC scale and usability, for example, by accelerating simulations with surrogate models or intelligently steering simulations based on previous results. New workflows are explored in which AI models are iteratively improved by continual learning to better reflect input data distributions and avoid outliers and drifts. Tracking of model provenance in these workflows is important to understand how new data affect model performance, allow unwinding to previous iterations, and provide a better understanding of conditions where AI models perform well for future reuse. This is more challenging in hybrid HPC-AI workflows because the lineage and provenance must be tracked across multiple software components at different levels of scale. In this work, we extend HPE Common Metadata Framework to hybrid simulation – AI workflows. We demonstrate benefits of CMF tracking across simulation, AI training and inference along HPE SmartSim system on a simple computational fluid dynamics problem with Eddy Kinetic Energy parameterized by AI. We track out-of-distribution data for continuous learning and employ adaptive switching between different models to improve the quality of results. We are working with fusion energy and materials science communities enhancing their workflows in a similar fashion.

Sessions and Abstracts

Search and Query Framework for Workflows with HPC and AI Models

1:30pm - 2pm

Christopher Rickett, Sreenivas Sukumar, and Karlon West (HPE)

Abstract

Modern computational science workflows increasingly involve complex, interactive, and iterative search through data from simulations of physics-based equations coupled with analytic, predictive, generative, and agentive tasks. Unfortunately, there are no query engines that empower scientists to search through scientific data with AI, analytic and physics-based models similar to searches with SQL query engines on structured data or keyword-search/prompt engines for textual data.

This paper presents an intelligent data search framework that empowers scientists to query scientific data. The framework functionality currently consists of: (i) a query engine and index for both keyword, set-theoretic and linear-algebraic search (ii) a repository of domain-specific models or pre-trained AI models, and (iii) a query planner for mixed HPC simulations and AI inferences. We demonstrate this framework on a life-sciences use-case in collaboration with the National Center for Natural Products. We leverage AI models such as AlphaFold, MolGAN, DTBA as AI surrogates for structure prediction, molecular generation, and binding-affinity prediction along with theoretical HPC codes such as Autodock Vina and domain-specific functions for protein similarity estimation to search and hypothesize healthy alternatives to existing compounds such as caffeine or zero-calorie sugar. We show how our search framework answers the "what-is," "what-else," "what-if," and "what-could-be" questions.

FirecREST v2: Lessons Learned from Redesigning an API for Scalable HPC Resource Access

2pm - 2:30pm

Elia Palme and Juan Pablo Dorsch (CSCS - ETH Zurich); Ali Khosravi and Giovanni Pizzi (PSI Center for Scientific Computing, Theory, and Data); and Francesco Pagnamenta, Andrea Ceriani, Eirini Koutsaniti, Rafael Sarmiento, Ivano Bonesana, and Alejandro Dabin (CSCS - ETH Zurich)

Abstract

Introducing FirecREST v2, the next generation of our open-source RESTful API for programmatic access to HPC resources. FirecREST v2 delivers a ~100x performance improvement over its predecessor. This paper explores the lessons learned from redesigning FirecREST from the ground up, with a focus on integrating enhanced security and high throughput as core requirements. We provide a detailed account of our systematic performance testing methodology, highlighting common bottlenecks in proxy-based APIs with intensive I/O operations. Key design and architectural changes that enabled these performance gains are presented. Finally, we demonstrate the impact of these improvements, supported by independent peer validation, and discuss opportunities for further improvements.

1:00pm-2:30pm

Technical Session 6C: Programming Models — Newport II

Designing GPU-aware OpenSHMEM

for HPE Cray EX and XD Systems

1pm - 1:30pm

Danielle Sikich, Naveen Namashivayam Ravichandrasekaran, Elliot Joseph Ronaghan, Nathan Wichmann, Robert Kierski, William Okuno, and Ted Peters (HPE)

Abstract

OpenSHMEM is a Partitioned Global Address Space (PGAS) based library interface specification. It is a culmination of a standardization effort among many implementers and users of the SHMEM programming model. The existing OpenSHMEM specification is not GPU-aware, the programming model does not enable managing the data movement operations involving a GPU-attached memory buffer. However, the OpenSHMEM users are exploring options to enable the execution of their data-driven workloads on heterogeneous system architectures.

HPE Cray OpenSHMEMX is a key implementation of the standard OpenSHMEM programming model that supports multiple vendor platforms. This work introduces the design for a vendor and network-agnostic GPU-aware OpenSHMEM library using the HPE Cray OpenSHMEMX implementation as an example across different system platforms supported by HPE. It includes performance characterization of the implementation across different heterogeneous architectures and further explores the need for supporting advanced GPU-centric communication schemes in introducing GPU autonomy in communication. This work is used to advance the OpenSHMEM programming model and allow standardizing vendor-agnostic support for GPU awareness in the programming model.

Quantifying Message Aggregation Optimisations for Energy Savings in PGAS Models

1:30pm - 2pm

Aaron Welch and Oscar Hernandez (Oak Ridge National Laboratory) and Stephen Poole and Wendy Poole (Los Alamos National Laboratory)

Abstract

Upon breaking past the exascale barrier, HPC systems are facing their greatest challenge yet - a power wall that must be addressed through new methods in both hardware and software. While energy costs are becoming a major issue at all levels, of particular concern is that of the network, as the relative cost of moving data is increasing faster than ever. The Partitioned Global Address Space (PGAS) model is critical within certain HPC domains, but is known to suffer from the small message problem, where irregular many-to-many access patterns result in congesting the network with excessive numbers of small messages. To address this, an extension to the OpenSHMEM PGAS library called was developed to defer individual messages and group them for subsequent bulk processing with nearly no code changes required. In this paper, we investigate its impact on energy use related to the network, with a focus on the Slingshot 11 interconnect. We will demonstrate that this strategy is not only highly performant, but also crucial to reducing energy footprints to remain within target power envelopes.

Sessions and Abstracts

Accelerating LArTPC Simulations: Enhancing larnd-sim with GPU Optimization Techniques

2pm - 2:30pm

Madan Timalsina (National Energy Research Scientific Computing Center/Lawrence Berkeley National Laboratory); Matt Kramer (Lawrence Berkeley National Laboratory); Pengfei Ding (National Energy Research Scientific Computing Center/Lawrence Berkeley National Laboratory); Ronan Doherty (Trinity College Dublin); Rishabh Dave (UC Berkeley); Nicholas Tyler, Urjoshi Sinha, and William Arndt (National Energy Research Scientific Computing Center/Lawrence Berkeley National Laboratory); and Callum Wilkinson (Lawrence Berkeley National Laboratory)

Abstract

Advancements in general-purpose computing on GPUs have enabled highly parallelized Monte Carlo simulations for particle physics experiments, including for the Deep Underground Neutrino Experiment (DUNE), which will use the world's most powerful neutrino beam to study the properties of these elusive particles. Here, we present our efforts on the optimization of larnd-sim, a microphysical simulation for liquid argon time projection chambers (LArTPCs) with light and pixelated charge readout, originally developed for the DUNE Near Detector (ND-LAr). Implemented in Python and utilizing GPU acceleration via Numba and CuPy, larnd-sim processes energy depositions from Geant4 to simulate physical phenomena (such as ionization electron drift) and the response of the detector electronics. By profiling with NVIDIA tools and optimizing memory transfers, adjusting register counts, tuning block and grid dimensions, altering floating-point precision, enabling the "fast-math" option for transcendental functions, converting arrays to a jagged format, and tuning CUDA kernels, we achieved an over-50% GPU-memory reduction, a ~30% wallclock speed improvement, and individual kernel speedups of 10–500%. In addition to these ongoing tests on the NERSC Perlmutter supercomputer, we are working with collaborators at ANL to run these simulations on the Polaris machine, further expanding larnd-sim's reach.

2:30pm-3:00pm

Coffee Break **Newport Foyer**

3:00pm-4:30pm

Technical Session 7A: AI/ML GPU Workloads **Newport III**

Evaluating the Performance of Containerized ML and LLM Applications on the Frontier and Odo Supercomputers

3pm - 3:30pm

Bishwo Dahal (University of Louisiana Monroe, Oak Ridge National Laboratory) and Elijah Maccarthy and Subil Abraham (Oak Ridge National Laboratory)

Abstract

Containers are transforming scientific computing by simplifying the packaging and distribution of applications. This enables researchers to create and deploy their applications in isolated environments with all necessary dependencies, enhancing portability and deployment flexibility. These advantages make containers especially

suitable for High Performance Computing (HPC) facilities like the Oak Ridge Leadership Computing Facility (OLCF), where complex scientific applications are being developed and deployed. In this work, we investigate the performance of containerized machine learning (ML) applications in comparison to bare-metal execution on the Frontier Exascale supercomputer. Specifically, we aim to determine whether ML models, when trained and tested within containers on Frontier using Apptainer, exhibit performance similar to that of bare-metal implementations. To achieve this, we use containers to package and run Convolutional Neural Network (CNN)-based ML applications on the OLCF Frontier and Odo supercomputers and assess their performance against bare-metal runs. After conducting scalability tests across up to 30 nodes with 1680 AMD EPYC CPU cores and 240 GPUs, we find that the performance of the containerized ML applications is at par with that of bare-metal runs. We apply the lessons learned from our containerized ML model to containerizing and evaluating performance of LLMs like AstroLLaMA, and CodeLLaMA on Frontier.

Evolving HPC services to enable ML workloads on HPE Cray EX

3:30pm - 4pm

Stefano Schuppli, Fawzi Mohamed, Henrique Mendonca, Nina Mujkanovic, Elia Palme, Dino Conciatore, Lukas Drescher, Miguel Gila, Pim Witlox, Joost VandeVondele, Maxime Martinasso, Torsten Hoefler, and Thomas Schulthess (Swiss National Supercomputing Centre)

Abstract

The Alps Research Infrastructure leverages GH200 technology at scale, featuring 10,752 GPUs. Accessing Alps provides a significant computational advantage for researchers in Artificial Intelligence (AI) and Machine Learning (ML). While Alps serves a broad range of scientific communities, traditional HPC services alone are not sufficient to meet the dynamic needs of the ML community. This paper presents an initial investigation into extending HPC service capabilities to better support ML workloads. We identify key challenges and gaps we have observed since the early-access phase (2023) of Alps by the Swiss AI community and propose several technological enhancements. These include a user environment designed to facilitate the adoption of HPC for ML workloads, balancing performance with flexibility; a utility for rapid performance screening of ML applications during development; observability capabilities and data products for inspecting ongoing large-scale ML workloads; a utility to simplify the vetting of allocated nodes for compute readiness; a service plane infrastructure to deploy various types of workloads, including support and inference services; and a storage infrastructure tailored to the specific needs of ML workloads. These enhancements aim to facilitate the execution of ML workloads on HPC systems, increase system usability and resilience, and better align with the needs of the ML community. We also discuss our current approach to security aspects. This paper concludes by placing these proposals in the broader context of changes in the communities served by HPC infrastructure like ours.

Transforming Hybrid Workflows: The Role of HPE Cray Supercomputing User Services Software in Bridging HPC and AI

4pm - 4:30pm

Tulsi Mishra, Dean Roe, and Larry Kaplan (HPE)

Abstract

As the convergence of HPC and AI reshapes computational workflows, the complexity of managing hybrid environments has become a significant challenge for organizations. HPE Cray Supercomputing User Services Software (USS) offers a transformative approach to simplify, scale, and optimize workflows across HPC and AI landscapes. In this session, we will explore how USS aims to bridge the gap between traditional HPC workloads and AI-driven innovations, providing a unified platform for containerized environments, hybrid deployment orchestration, and energy-efficient operations.

Key topics will include the challenges of managing hybrid HPC-AI workflows, the evolving capabilities of USS such as complex workload orchestration, container curation, deployment and management, and sustainability features, as well as real-world use cases in industries like life sciences, autonomous systems, and supply chain optimization. Participants will have the opportunity to engage in discussions around USS's open-source integrations, strategic initiatives, and its role in accelerating hybrid adoption. Whether you are an HPC administrator, AI researcher, or IT decision-maker, this session will offer insights into how USS can drive operational excellence in your organization.

Technical Session 7B: Access Nodes & Kubernetes Management Newport I

Porting Radio Astronomy Correlation to Setonix, a HPE Cray EX system powered by AMD GPUs

3pm - 3:23pm

Peter Upton, Ben Lenard, Ben Allen, and Cyrus Blackworth (Argonne National Laboratory)

Abstract

This paper presents Administrator Access Nodes (AANs) as an alternative to the traditional reliance on a single Admin Node for all aspects of system administration in an HPE Performance Cluster Manager (HPCM) managed supercomputer cluster. At the Argonne Leadership Computing Facility (ALCF), managing the Aurora supercomputer, a large HPE Cray EX system, requires a team of skilled developers and administrators. These professionals require access to many tools for tasks such as parsing log files, issuing power commands, and connecting to nodes via SSH. These tasks have typically been performed solely on the Admin Node. However, this centralization can lead to resource constraints due to simultaneous resource requirements by multiple administrators. To address these issues, the paper details the implementation and operation of AANs, including custom tools for interacting with HPCM, scripts to replicate some Admin Node functionality on AANs, and synchronization tools for configuration files. The introduction of AANs has alleviated resource constraints, streamlined workflows, and enhanced system manageability. Possible future work is also discussed, focusing on further integrating HPCM's APIs and improving usability, aiming to enhance AAN capabilities and administrative efficiency for Aurora's complex environments

HPE Slingshot in the Kubernetes Ecosystem

3:23pm - 3:46pm

Caio Davi and Jesse Treger (HPE)

Abstract

The convergence of traditional HPC systems with AI increases expectations for supercomputing sites to deliver new capabilities (beyond traditional batch scheduling, single tenancy, and bare-metal application deployment methodologies) for more dynamic provisioning. Convergence with enterprise cloud computing techniques such as containerized applications and Kubernetes have become a priority. But transitioning high-performance computing (HPC) environments and applications to Kubernetes is complex because of the critical requirement to maintain low-latency networking for high-performance. In this context, we have HPE Slingshot, a modern high-performance interconnect for HPC and AI clusters that delivers industry-leading performance, bandwidth, and low-latency for HPC, AI/ML, and data analytics applications through innovations in the fabric to overcome congestion and innovations in the NIC to significantly offload communications and message processing from the hosts. Because the HPE Slingshot NICs run native Ethernet alongside its optimized RDMA transport and connectionless protocol, ensuring that the RDMA transport is operating as intended is critical to delivering the high performance expected in HPC and AI. This requires careful configuration of Kubernetes because If not configured, the system can fall back to using standard TCP/IP over Ethernet instead of the expected HPC and AI performance. Our proposed solution is composed of a number of Kubernetes components such as device plugins, CNIs, Operator, and Admission Policies. These contributions represent a significant advancement in deploying and operating HPC applications within containerized environments and offering a robust framework for future developments in distributed computing, ensuring both high performance and ease of management for the continuing convergence of HPC/AI and cloud computing and the coming transition from siloed HPC interconnects to interoperable Ultra-Ethernet transport.

Building non-standard images for CSM systems

3:46pm - 4:08pm

Harold Longley, Isa Wazirzada, Dennis Walker, Andy Warner, and Davide Tacchella (HPE)

Abstract

HPC scientists increasingly must innovate using diverse toolchains crafted from various Linux distributions ensuring they meet individual and project-specific needs to tackle complex challenges.

Advancements in User Services Software (USS) 1.2 and Cray System Management (CSM) 1.6 now enable non-standard image utilization on managed nodes, moving beyond Cray Operating System (COS)/USS dependencies. This is facilitated by iSCSI-based OS image projection.

This paper explores creation of stock Linux images from RPMs--SUSE, OpenSUSE, RHEL, Rocky Linux--using CSM tools and recent CSM architectural enhancements. Images developed or externally sourced can be injected into the artifact storage or CSM tools could be used to create the non-standard OS image. CSM tools further enable Ansible-driven image customization, booting the image, and applying

post-boot Ansible configuration. Functionality is extended through enhancements with clients or agents for Slingshot Host Software, parallel filesystem, workload manager, node heartbeat, BOS, and CFS. Adhering to portable Ansible code best practices allows seamless integration for standard and non-standard OS images. SAT bootprep automates the creation of versioned boot and configuration artifacts. OS images can be transformed into OCI containers for deployment as UAIs on UANs or on compute nodes via Podman, Singularity, or Kubernetes.

Enabling NCCL on Slingshot 11 at NERSC

4:08pm - 4:30pm

JRiccardo Di Maria and Chris Gamboni (Swiss National Supercomputing Centre), Davide Tacchella and Isa Wazirzada (HPE), and Mark Klein (Swiss National Supercomputing Centre)

Abstract

The Alps research infrastructure consists of 32 HPE Cray EX cabinets which are managed by Cray System Management (CSM). A critical component of this system is the Slingshot fabric manager and it is responsible for managing the high-speed network fabric. Presently, the fabric manager is deployed as a Kubernetes Pod and runs amongst other services on the system management nodes. An ongoing effort aims at separating the fabric manager from Kubernetes, and deploying it on bare metal hardware. The architectural decision-making process is examined in detail, accompanied by a walkthrough of the newly proposed design. The discussion of the design is framed within the context of key quality attributes, including reliability, resiliency, availability, observability, and performance. Subsequently, the focus transitions from the "what" to the "how," providing a comprehensive overview of the execution of the migration of the fabric manager from a Kubernetes-based deployment to a bare-metal environment. Insights are presented regarding aspects that were successful, challenges encountered, and whether the overall outcome of this effort achieved the intended objectives.

Technical Session 7C: Application Performance — Newport II

Porting Radio Astronomy Correlation to Setonix, a HPE Cray EX system powered by AMD GPUs

3pm - 3:22pm

Cristian Di Pietrantonio (Pawsey Supercomputing Research Centre, Curtin Institute for Radio Astronomy); Marcin Sokolowski (Curtin Institute of Radio Astronomy); Christopher Harris (Pawsey Supercomputing Research Centre); and Daniel Price and Randal Wayth (SKAO)

Abstract

In low-frequency radio astronomy correlation of signals coming from hundreds of radio antennas is an early and fundamental step to create science-ready data products such as images of the sky at radio wavelength. Because of the high volume of data to process and the rate at which they are produced, correlation is most of the time performed in real time by dedicated hardware, a FPGA or GPU cluster, installed near the telescope. However, there are science cases when an astronomer would like to correlate data later with customised settings such as time and frequency averaging of signals. Setonix, Pawsey Supercomputing Centre's HPE Cray Ex supercomputer based on AMD CPUs and GPUs, provides radio astronomers with

enough computational power for such processing, but the only established GPU correlator works only on NVIDIA GPUs and proved hard to port. In this paper we discuss the process of providing Australian astronomers an implementation of the correlation algorithm that harnesses the computational power of Setonix.

Task-decomposed Overlapped Pressure Preconditioner for Sustained Strong Scalability on Accelerated Exascale Systems

3:22pm - 3:45pm

Niclas Jansson (KTH Royal Institute of Technology)

Abstract

Computational Fluid Dynamics is a natural driver for exascale computing with a virtually unbounded need for computational resources for accurate simulation of turbulent fluid flow, both for academic and engineering usage. However, with exascale computing capabilities on the horizon, we have seen a transition to more heterogeneous computer architectures with various accelerators. While offering high theoretical peak performance and high memory bandwidth, complex programming models and significant programming investments are necessary to efficiently exploit these systems. We detail our work on improving the performance and scalability of key numerical methods in the high-fidelity spectral element code Neko on accelerated exascale machines. Efficient preconditioners are essential in incompressible fluid dynamics; however, the most efficient method (with respect to convergence) might be challenging to implement with good performance on an accelerator. We present our development of a GPU-optimised preconditioner with task overlapping for the pressure-Poisson equation, improving the preconditioner's throughput (in TDoF/s) by more than 61%. The new preconditioner is explained in detail, together with detailed performance studies on Cray EX platforms, including strong scalability studies on Frontier, a performance comparison between AMD and NVIDIA accelerated nodes, and an assessment of the feasibility of mixing both node types in a single simulation.

Supernovae in HPC: Benchmarking FLASH Across Advanced Computing Clusters

3:45pm - 4:07pm

Joshua Martin, Eva Siegmann, and Alan Calder (Stony Brook University, Institute of Advanced Computational Science)

Abstract

Astrophysical simulations are highly demanding in terms of computation, memory, and energy, requiring new advancements in hardware. Stony Brook University recently expanded its "SeaWulf" computing cluster by adding 94 new nodes with Intel Sapphire Rapids Xeon Max series CPUs. This benchmarking study evaluates the performance and power efficiency of this new hardware using FLASH: a multi-scale, multi-physics software instrument that utilizes adaptive mesh refinement. Our study also compares the performance of Stony Brook's Ookami testbed which features ARM-based A64FX-700 processors as well as SeaWulf's existing AMD EPYC Milan and Intel Skylake nodes. The focus of our simulation is the evolution of a bright stellar explosion known as a thermonuclear (Type Ia) supernova—a complex 3D problem

that incorporates various operators for hydrodynamics, gravity, nuclear burning, and routines for the material equation of state. We perform strong-scaling tests on a 220 GB problem size and assess both single-node and multi-node performance. We analyze the performance of various MPI mappings and processor distributions across nodes. From our strong-scaling tests, we conclude the optimal configuration for balancing the minimization of runtime and energy consumption for our application.

Expanding Community Access to Real-world HPC Application I/O Characterization Data Using Darshan

4:07pm - 4:30pm

Shane Snyder, Philip Carns, Robert Ross, Robert Latham, and Kevin Harms (Argonne National Laboratory)

Abstract

HPC systems are deployed with massive, distributed storage subsystems to meet the demands of data-intensive applications. While these storage systems offer impressive peak performance, it is often only attainable in idealized scenarios not reflective of production workloads. In general, there continues to be a lack of community understanding of the I/O performance characteristics of real-world applications.

The Darshan tool can help characterize application I/O behavior, but unfortunately Darshan logs are generally kept private to protect sensitive user-identifying information. This drastically restricts the impact of Darshan data, precluding the ability to analyze I/O behavior in a broader context. Public exchange of Darshan data could increase community understanding and inspire novel systems research, but requires established processes for publishing anonymized, full-resolution Darshan logs.

To help address this problem, we developed an automated workflow for anonymizing and publishing Darshan data captured on Polaris, an HPE Apollo system at the ALCF. This continuously updated dataset includes over 600,000 logs and is growing daily. In this paper, we describe this new community resource, demonstrate its utility by highlighting platform-wide characteristics and outlier job behavior, and explore best practices for continuously capturing and publishing system telemetry data

4:30pm-4:40pm

2025 CUG Closing — Newport III

May 9th-10th

All Day

Expanding Horizons in AI with HPC Workshop

This workshop, located at Stony Brook University on Expanding Horizons in AI with HPC, aims to explore the dynamic intersection of AI and HPC, focusing on how advanced computing can accelerate AI research and applications. As AI models become more complex and data-intensive, traditional computing systems struggle to meet the demand for scalability, efficiency, and speed. HPC offers a solution by providing the necessary infrastructure for training large-scale models, enhancing AI algorithms, and enabling breakthroughs in fields such as deep learning, natural language processing, and autonomous systems.

Registration and more details are available here: https://cug.org/cug-2025-aiwithhpc-workshop-2/

30+ YEARS OF HPC AND CLOUD EXPERTISE

In today's fast-paced world, staying ahead means embracing innovation and precision. Achieving results faster and more efficiently is more important than ever. Your high-performance computing (HPC) environment needs to be agile, scalable, and fully optimized.

Altair® HPCWorks® covers every aspect of HPC, from seamless cloud scaling across any environment to flexible workload management, enhanced user experiences, in-depth monitoring, and more. Push boundaries, solve challenges, and change tomorrow.

Scan the QR code or go to altair.com/altair-hpcworks to find out more.

ALTAIR

Social Events

Sunday, 4th May 2025

5:30pm – 6:30pm CUG Welcome Reception - Newport Foyer

Please head to Newport Foyer from 3:00 pm onwards to register for CUG and obtain your badge.

Tuesday, 6th May 2025

6:00pm – 8:00pm HPE etworking Event at Lokal, 2 2nd Street, Jersey City, NJ 07302

HPE will host their annual CUG community networking reception from 6:00 to 8:00 pm ET at the Lokal Eatery & Bar. The Lokal is located at 2 2nd St, Jersey City, NJ 07302, along Jersey City's waterfront, allowing CUG guests to enjoy expansive views of the Manhattan Skyline. Co-presented by AMD, all registered CUG attendees and their guests are invited to attend for a reception with light hors d'oeuvres and drinks. The Lokal is about a 10 min walk from the CUG hotel, but transportation will be provided if inclement weather.

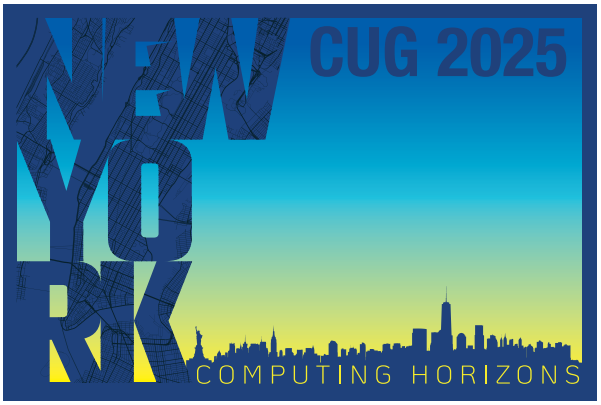

Social Events *(cont.)*

Wednesday, 7th May 2025

6:15pm – 9:30pm CUG Night out at Hudson House, 2 Chapel Ave, Jersey City, NJ

We invite all registered attendees and guests with a paid CUG night out ticket to join us for an unforgettable evening at Hudson House. Situated at the end of Port Liberte in Jersey City, NJ, this structure is an arms' length away from the Hudson River and boasts a panoramic view of the Statue of Liberty, Brooklyn, Manhattan, and Verrazano Bridges, and of course the NYC Skyline.

Coaches will depart outside the Westin Jersey City Hotel at 18:15 to arrive at Hudson House for a drink's reception before seating for dinner at approximately 19:15.

If you are making your own way to the venue, please use the full address as Google Maps takes you to a different address! Hudson House, 2 Chapel Ave is approx. a 15 – 20-minute drive.

Our first bus will return to the hotel at approximately 21:00 with shuttles then running until approximately 22:30.

Menu & Bar

This event will feature water, soda, beer and wine, and the restaurant will offer a cash bar for those wanting an alternative. We will be serving passed Hors D'Oeuvres and a cold station display during the welcome reception. Once seated for dinner, the main menu will include a fish, chicken, beef or vegetarian entree. Your selection would have been pre-selected at the time of CUG registration. If there are any changes to this, please let us know ASAP and we will see what the caterers can do.

Guest Tickets

CUG Night Out is included in the registration fee for attendees, but if you would like to bring a guest, (someone not registered for the conference), they will need a guest ticket. Guest tickets are available for purchase through May 4th, 2025, at a cost of $100 (Early Bird) or $125 (Normal). No guest tickets can be purchased on-site.

Local Arrangements

How to Contact Us

After the conference:

Cray User Group
Attn: Scott Michael
2709 E Tenth Street
Bloomington, IN, 47408
cug2025@cug.org

During the conference:

You can find us in the Executive Boardroom (aka CUG Office) on the 3rd floor or at the registration desk located in the Newport Foyer on the 3rd Floor.

Conference Registration

Cray User Group
Attn: Scott Michael
2709 E Tenth Street
Bloomington, IN, 47408
cug2025@cug.org

Attendance and Registration

Badges and registration materials will be available:
Sunday: 3:00pm to 6:30pm
Registration Desk – Newport Foyer – 3rd Floor
Monday: 07:30 to 17:30
Registration desk – Newport Foyer – 3rd Floor
Tuesday: 07:30 to 10:30
Registration desk – Newport Foyer – 3rd Floor
To register after Tuesday morning visit the CUG office (Executive Boardroom) on the 3rd Floor.
All attendees must wear badges during CUG Conference activities (including the Welcome Reception, HPE Networking Event and the CUG Night Out)

Smoking Policy

There is no smoking allowed at the Conference.

Special Assistance

Any requests for special assistance during the conference should be noted on the "Special Requirements" area of the registration form.

Conference Registration Fees

Your registration fee includes:

• Admission to all program sessions, meetings, and tutorials

• Arrival, morning and afternoon breaks, and lunch Monday through Thursday

Proceedings

Proceedings login details:
Username: CUG2025
Passwords: NYC

The Westin, Jersey City
Floor Plan

1. Newport Grand Ballroom
2. Executive Boardroom
3. Intrepid

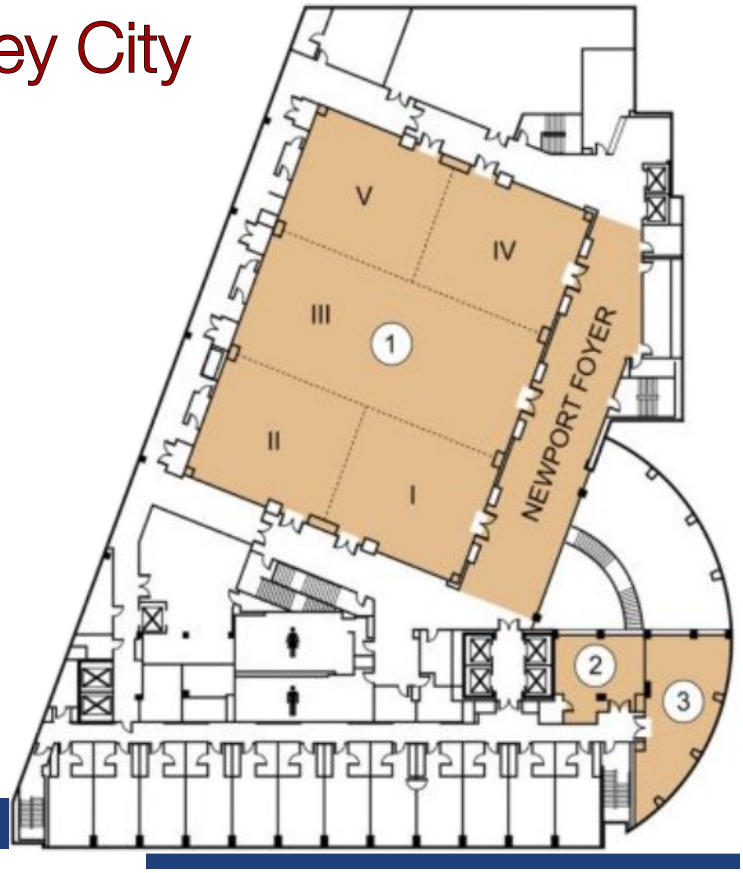

3RD FLOOR

Sponsors

DIAMOND SPONSOR

Hewlett Packard Enterprise

PLATINUM SPONSOR

AMD

GOLD SPONSORS

codee

NVIDIA

SILVER SPONSORS

linaro

ALTAIR

SCHEDMD
The Slurm Company

VAST

PIER Group
Partners In Education & Research

Contacts

CUG Board

President
Ashley Barker
Oak Ridge National Laboratory

Vice-President
Bilel Hadri
King Abdullah Science and Technology

Secretary
Jim Rogers
Oak Ridge National Laboratory

Treasurer
Scott Michael
Indiana University

Past President
Trey Breckenridge
Mississippi State University

Director-at-Large
Juan Rodriguez Herrera
EPCC

Director-at-Large
Jim Williams
Los Alamos National Laboratory

Director-at-Large
Maciej Cytowski

Pawsey Computing Centre

Board Liasions and Advisors

**HPE liaison to the CUG Board
Azita F. Sadri

**Note: This is not a CUG Board position.

The CUG Board of Directors (BoD) encourages anyone with a concern or question about the Cray User Group to contact the BoD directly, board@cug.org. The Program Lead Committee can be reached by emailing plc@cug.org and the CUG Office by emailing office@cug.org.

Special Interest Groups

Programming Environments, Applications and Documentation

Chair:
Chris Fuson (ORNL)

Deputy Chair:
Brandon Cook (NERSC)

Deputy Chair:
Eva Siegmann (Stony Brook)

Deputy Chair:
Ben Cumming (CSCS)

HPE SIG Liaisons:
Kaylie Anderson (HPE)

Systems Monitoring Working Group

Chair: Massimo Benini (CSCS)
Ann Gentile (SNL)
Steve Leak (LBL)
HPE Liaison: Jeff Hanson

XTreme Systems

Chair: Paul Peltz Jr. (ORNL)

Deputy Chair: Jared Baker (NCAR)

HPE SIG Liaison: Azita F. Sadri

HPC Security

Chair (Acting): Aaron Scantlin (LBNL)

Deputy Chair (Acting): Alan Stradling (LANL)

HPE SIG Liaison: Alex King

Thanks for attending

CUG
Cray User Group