What Time Taught Us: Monitoring a Computing Technology Testbed Across Multiple Years

OOKAMI

Eva Siegmann[1], **David Carlson**[1], Nikolay A. Simakov[2], Anthony Curtis[1], Alan Calder[1], and Robert J. Harrison[1]

6th ISC HPC International Workshop on MODA

[1]Stony Brook University, Institute for Advanced Computational Science, USA
[2]Center for Computational Research, SUNY University at Buffalo, NY 14203, USA

Stony Brook University

iACS INSTITUTE FOR ADVANCED COMPUTATIONAL SCIENCE

University at Buffalo
The State University of New York

Meet the Institute for Advanced Computational Science

Who we are: an interdisciplinary group of faculty and staff dedicated to research with a computational focus

Where we're located: Stony Brook University, NY, USA.

What we do: Operate and manage two HPC clusters

OOKAMI

iACS INSTITUTE FOR ADVANCED COMPUTATIONAL SCIENCE

Overview of the talk

1. The “What” and “Why” of the Ookami HPC cluster
2. Why would we want to monitor the system, anyway?
3. Introducing XDMoD
4. 4+ years of monitoring results

Fugaku
#1 Fastest computer in the world until 6/2022

First machine to be fastest in all 5 major benchmarks

- Green-500 benchmark
- Top-500 benchmark
- HPCG benchmark
- HPL-AI benchmark
- Graph-500 benchmark

(Currently 7, 104, & 2 on Top500, Green500 and HPCG)

- 432 racks
- 158,976 nodes
- 7,630,848 cores
- 440 PF/s dp (880 sp; 1,760 hp)
- 32 Gbyte memory per node
- 1 Tbyte/s memory bandwidth/node
- Tofu-2 interconnect

Ookami

1.5 mil node hours per year

Node	
Processor	A64FX 700
#Cores	48
Peak DP	2.76 TOP/s
Memory	32GB@**1TB/s**
System	
#Nodes	176
Peak DP	486 TOP/s
Peak INT8	3886 TOP/s
Memory	5.6 TB
Disk	0.8 PB Lustre
Comms	IB HDR-100

Additional CPUs

To facilitate users exploring current computer technologies and contrasting performance and programmability with the A64FX, Ookami also includes:

- 2 nodes with dual socket **Thunder X2** (64 cores)
- 2 nodes with dual socket **NVIDIA Grace CPU** superchips (144 cores)
- 1 node with dual socket **AMD Milan** (64 cores)
- 1 node with dual socket **Intel Skylake** (36 cores)
- 1 node with an **AmpereOne** CPU and a **QualcommAI** accelerator

Ookami - 狼

- Ookami is Japanese for wolf
  - Homage to the origin of the processor and the Stony Brook mascot
- A computer technology testbed supported by NSF
- Operated by Stony Brook in cooperation with the University at Buffalo

Stony Brook University

iACS INSTITUTE FOR ADVANCED COMPUTATIONAL SCIENCE

University at Buffalo The State University of New York

A64FX NUMA Node Architecture

- Arm V8-64bit
- Supports high calculation performance and low power consumption
- 32 (4x8) GB HBM @ 1TB/s
- Supports Scalable Vector Extensions (SVE) with 512-bit vector length
- 4 Core Memory Groups (CMGs)
  - 12 cores
  - 64KB L1$ per core - 256b cache line
  - 8MB L2$ shared between all cores - 256b cache line
  - Zero L3$

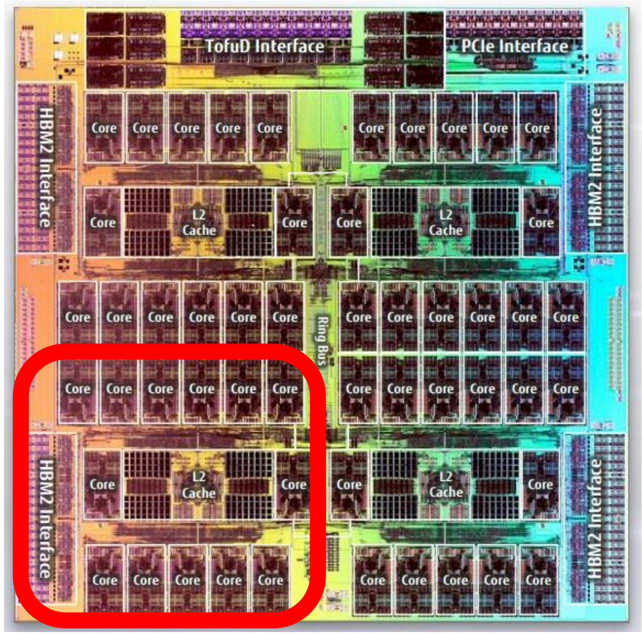

A64FX chip layout on Fugaku

Memory Statistics

These 86% of jobs using < 32GB memory correspond to 85% of the total XSEDE cpu-hour usage

N.A. Simakov, J.P. White, R.L. DeLeon, S.M. Gallo, M.D. Jones, J.T. Palmer, B. Plessinger, T.R. Furlani, "A Workload Analysis of NSF's Innovative HPC Resources Using XDMoD," arXiv preprint arXiv:1801.04306 (2018).

Figure 6. Memory used (GB) per node including OS for all XSEDE jobs on all resources in 2018. All queues, even large memory, are included. 86% of jobs ran on less than 32 GB/node (left of red line).

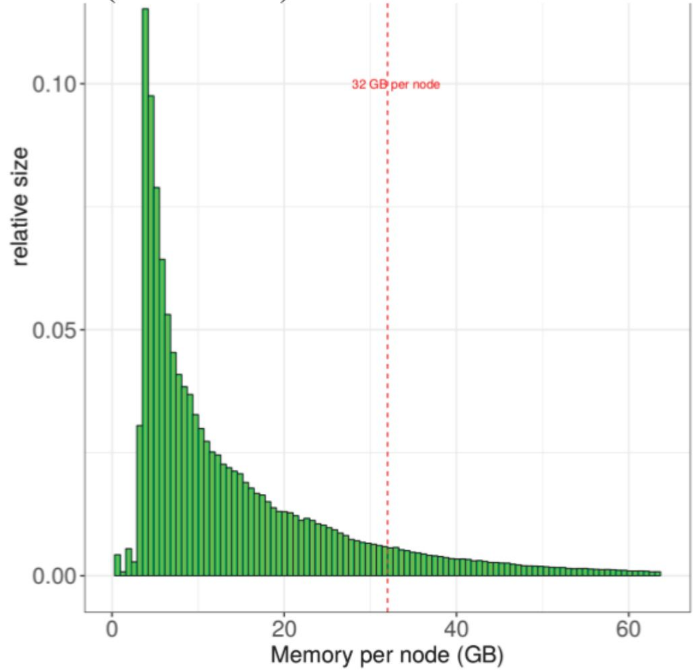

Ookami Environment

- Rocky Linux 8
- Lustre file system providing ~800TB
- Slurm workload manager
- Full bisection bandwidth HDR100
- Module environment
    - > 400 modules
- Various compilers, profilers & debuggers (GCC, Arm, CPE, forge, OSACA, etc.)
- Available via Open OnDemand

Ookami support mechanisms

There are several Ookami support mechanisms and a dedicated project team that is happy to support you!

- Online documentation & FAQ
- Slack channel
- Ondemand virtual office hours
- Ticketing system
- Regular webinars
- Email: ookami_computer@stonybrook.edu

Why Monitor Utilization and Performance?

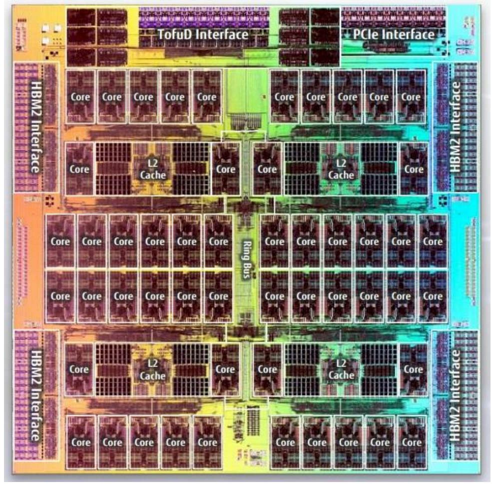

- **Are users on the system**?
  - It is a testbed that poses different types of challenges
  - Higher level of engagement with users
- **What are the users doing with the system?**
  - Applications, job sizes, walltime, etc.
  - Report the usage and analyse whether it matches the desired state of the project
- **How well is the system operating?**
  - Continuous performance monitoring to ensure optimal software and hardware state
  - Monitor the performance improvement with higher adoption of technologies
  - Benchmarking and comparison with other platforms

(Almost) five years of monitoring Ookami

What we collected:

- XDMoD - automated utilization, performance and job stats
- Manual collection of other metrics
  - project/user creation
  - events (webinars, etc.)
  - Google Scholar searches
  - User survey

XDMoD: A Comprehensive Tool for HPC System Management

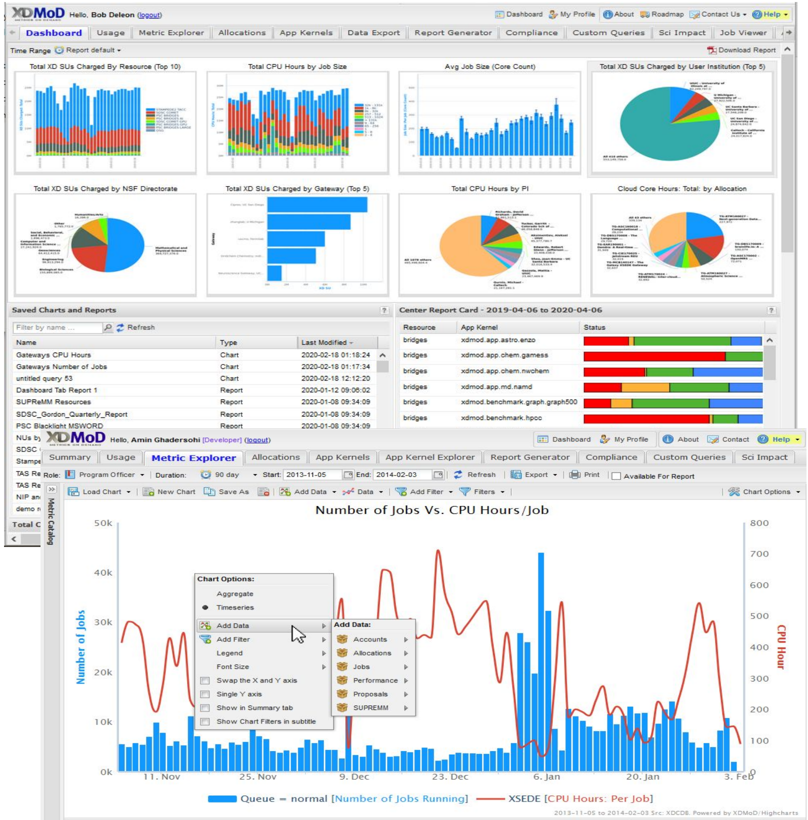

Goal: Optimize Resource Utilization and Performance

- Provide detailed information on utilization
- Continuous performance monitoring to proactively identify underperforming hardware and software
- Measure and improve job and system-level performance

NSF ACCESS Measurement and Metrics Service (MMS)

- Develop & deploy the **XDMoD** for monitoring ACCESS-CI

Open XDMoD: Open Source version for Data Centers

- Used to measure and optimize the performance of HPC centers

XDMoD history

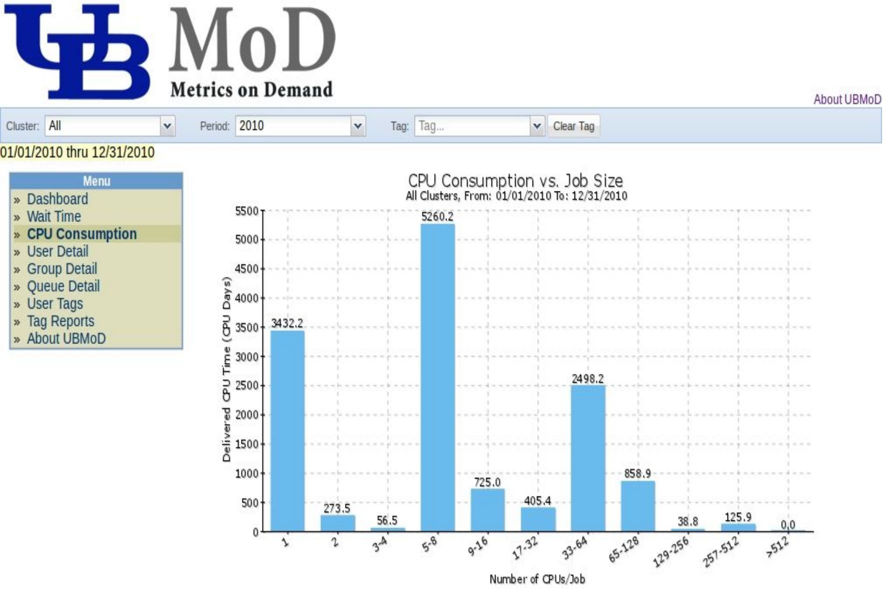

2007: UBMoD (UB Metrics on Demand) - tool for basic usage and accounting information visualization

XDMoD history

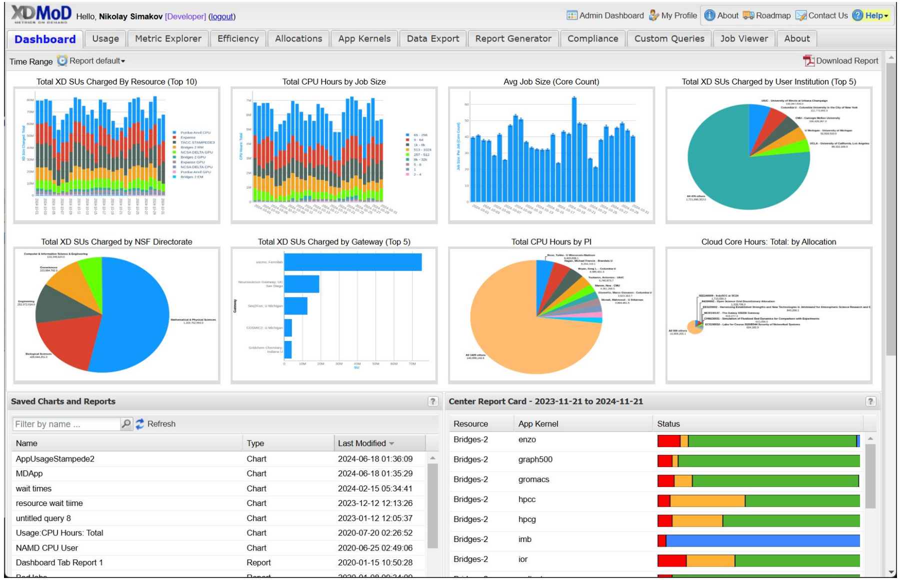

More info: Palmer, Jeffrey T., et al. "Open XDMoD: A tool for the comprehensive management of high-performance computing resources." *Computing in Science & Engineering* 17.4 (2015): 52-62.

open.xdmod.org

2010: XDMoD – NSF-funded tool for monitoring XSEDE cyberinfrastructure, a collection of NSF-funded HPC resources.

2014: Open XDMoD – open source version of XDMoD for HPC centers released. Used in 300+ academic & industrial systems worldwide

Ookami monitoring results

Steady increase in the number of users

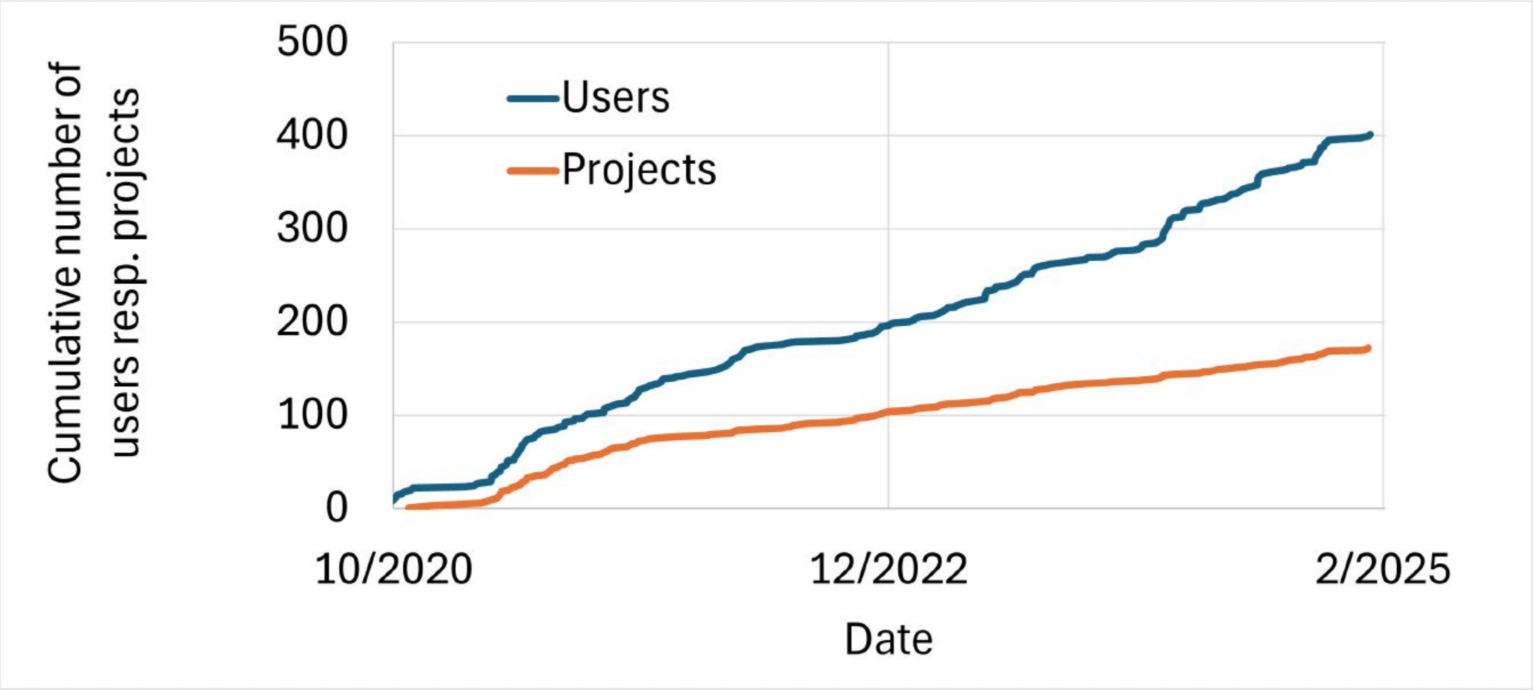

Allocations

- 97 projects via SBU allocation
- 80 projects via ACCESS allocation
  (allocation mechanism for researchers within the US)
- Projects from > 90 institutions
- 95.5% academia
- 4.48% industry (using so far only
  0.2% of all available node hours,
  no industry production projects)
- 94.19% USA, 5.81% Europe

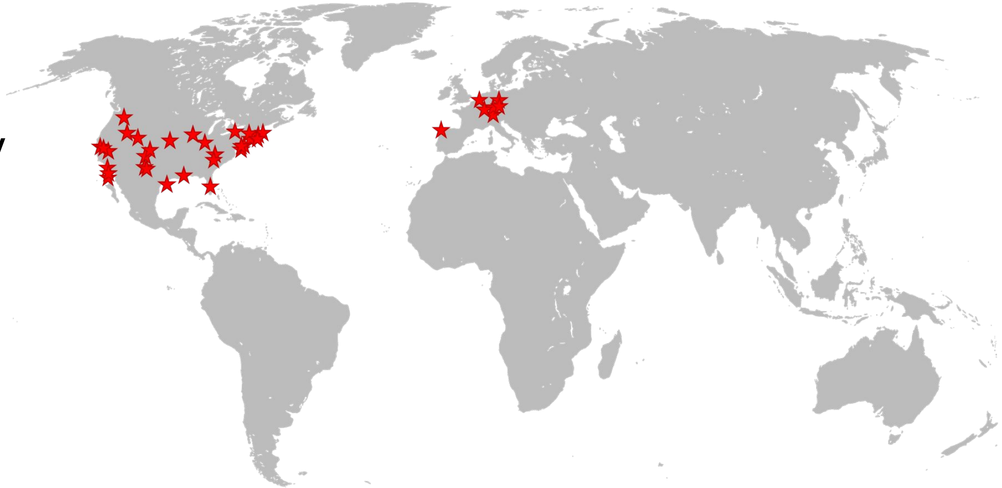

Active users remained consistent

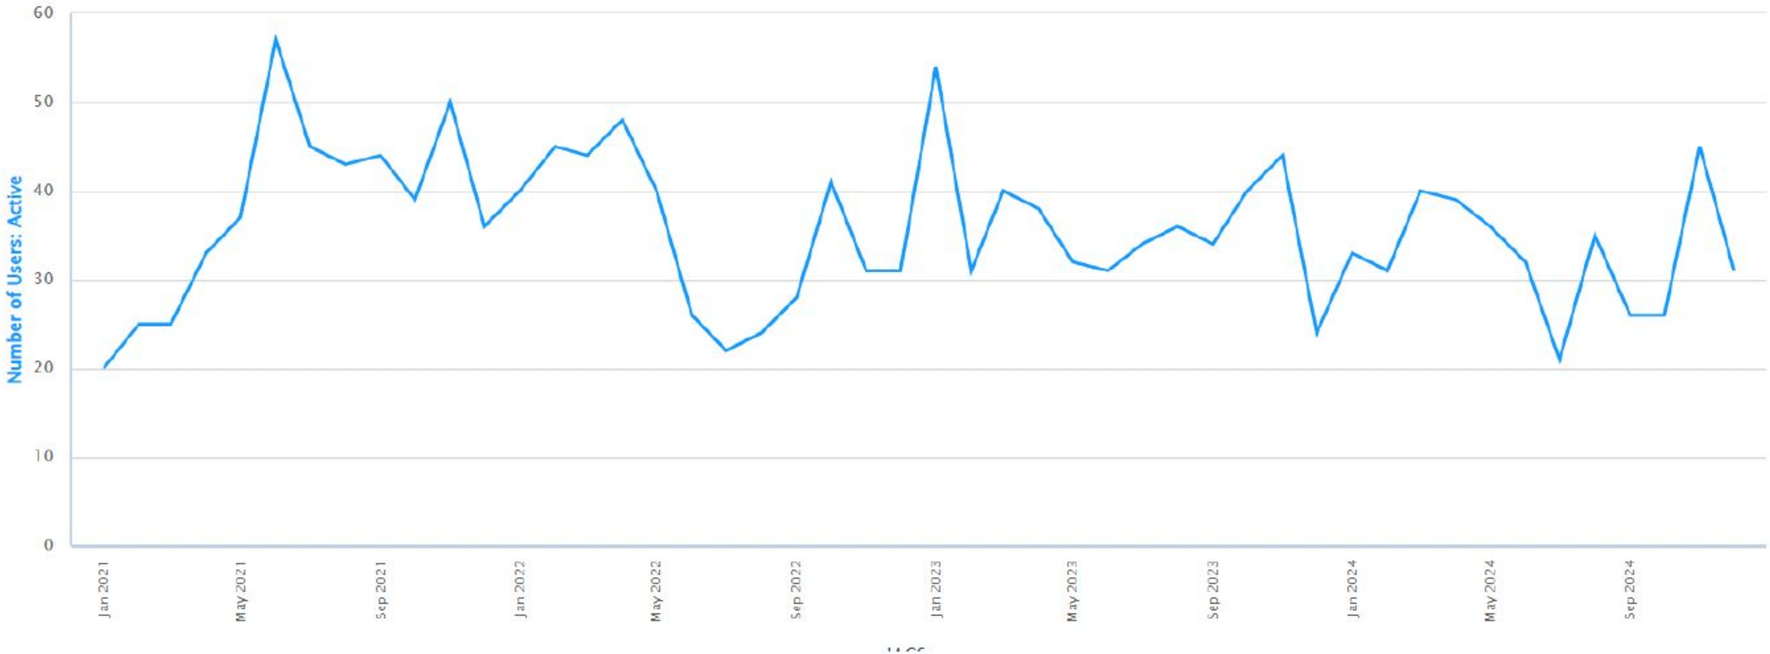

Overall utilization fluctuates (but grows!)

Blue: monthly averages

Orange: daily averages over recent months

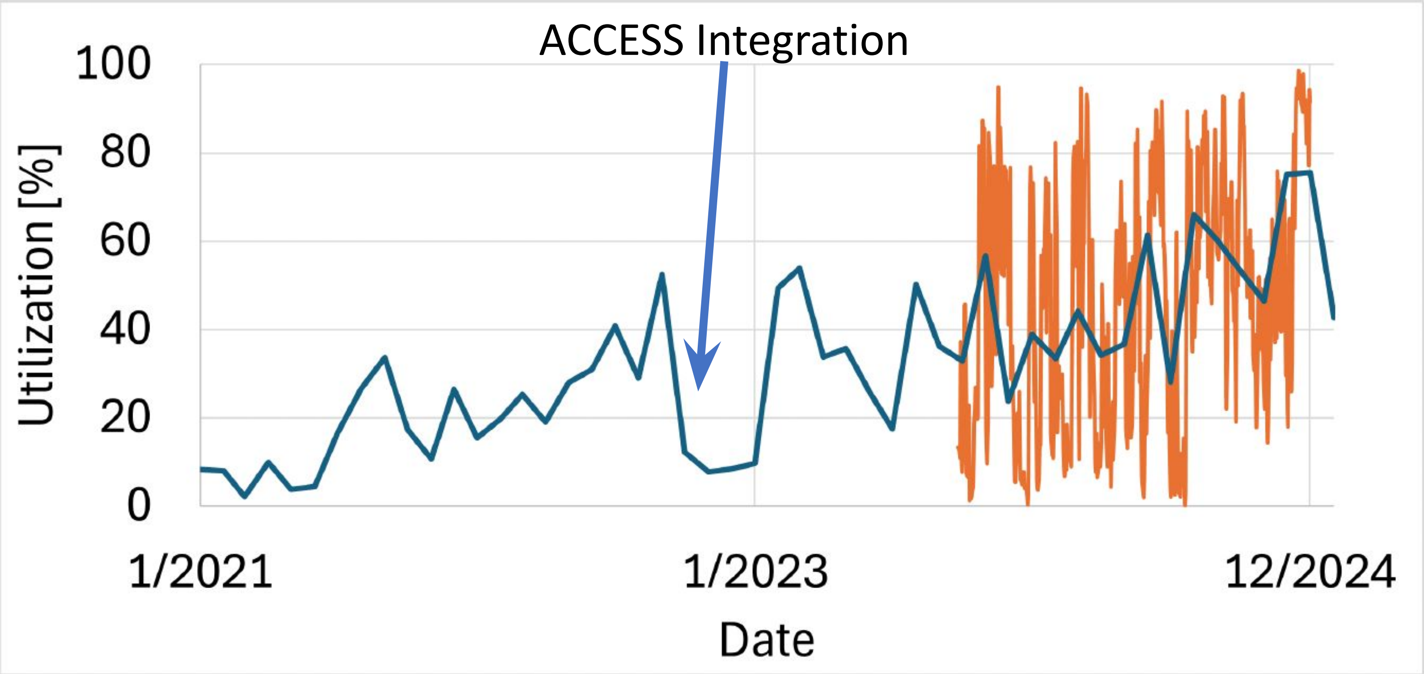

Patterns of utilization shift over time

Total core hours: **increasing**

Total jobs: **flat** (mostly)

Job size: **larger**

Movement from test jobs toward production

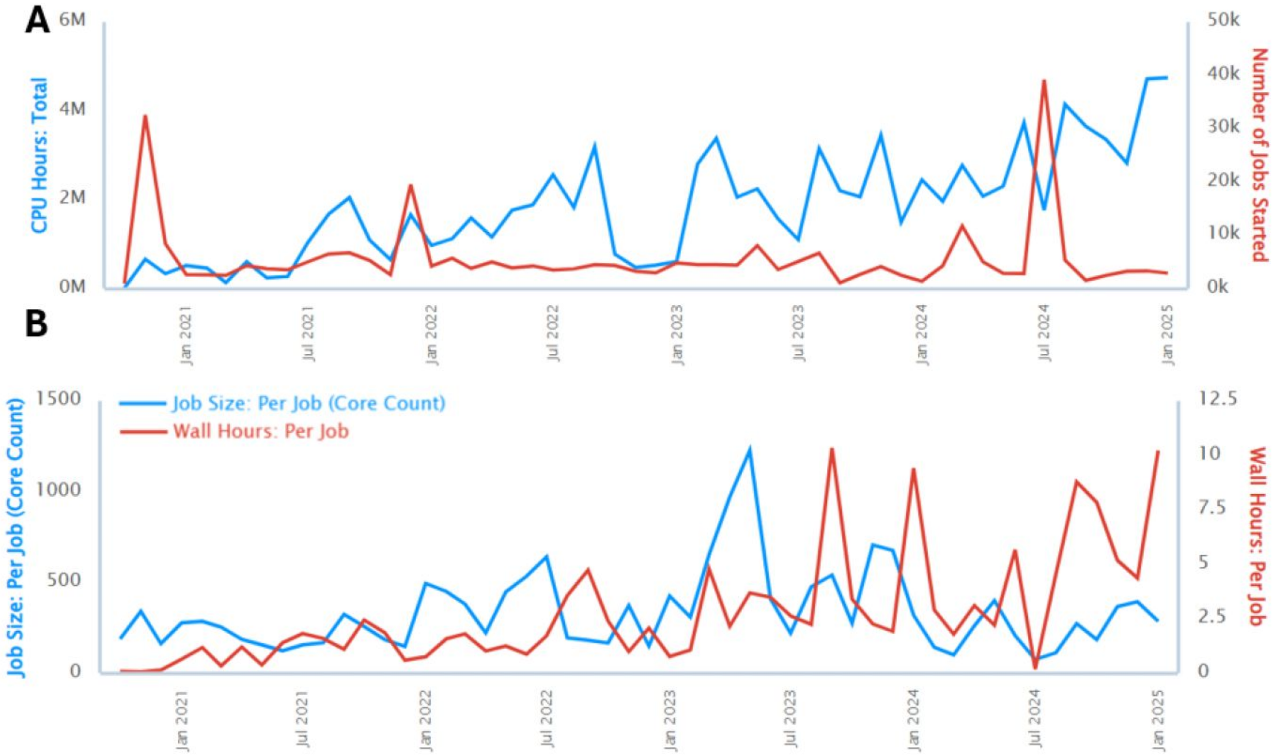

Ookami supported a diversity of different workloads

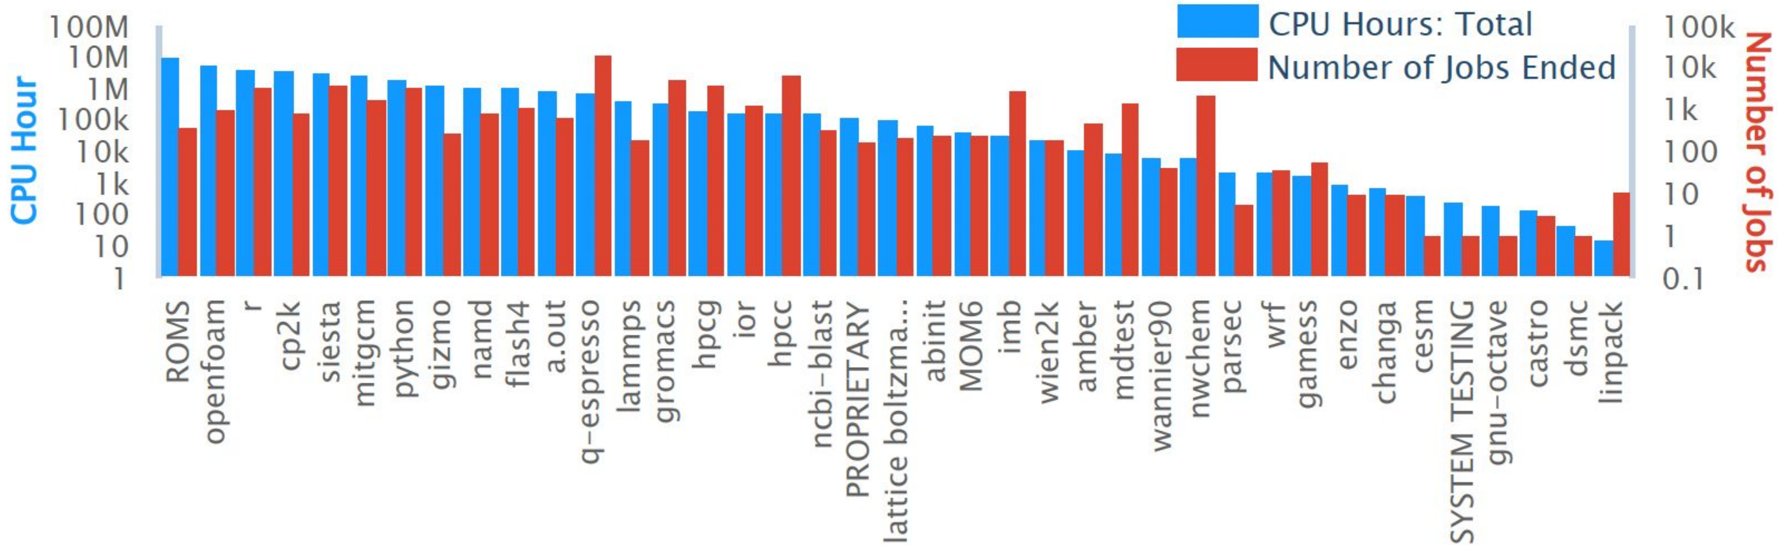

Support events held over time

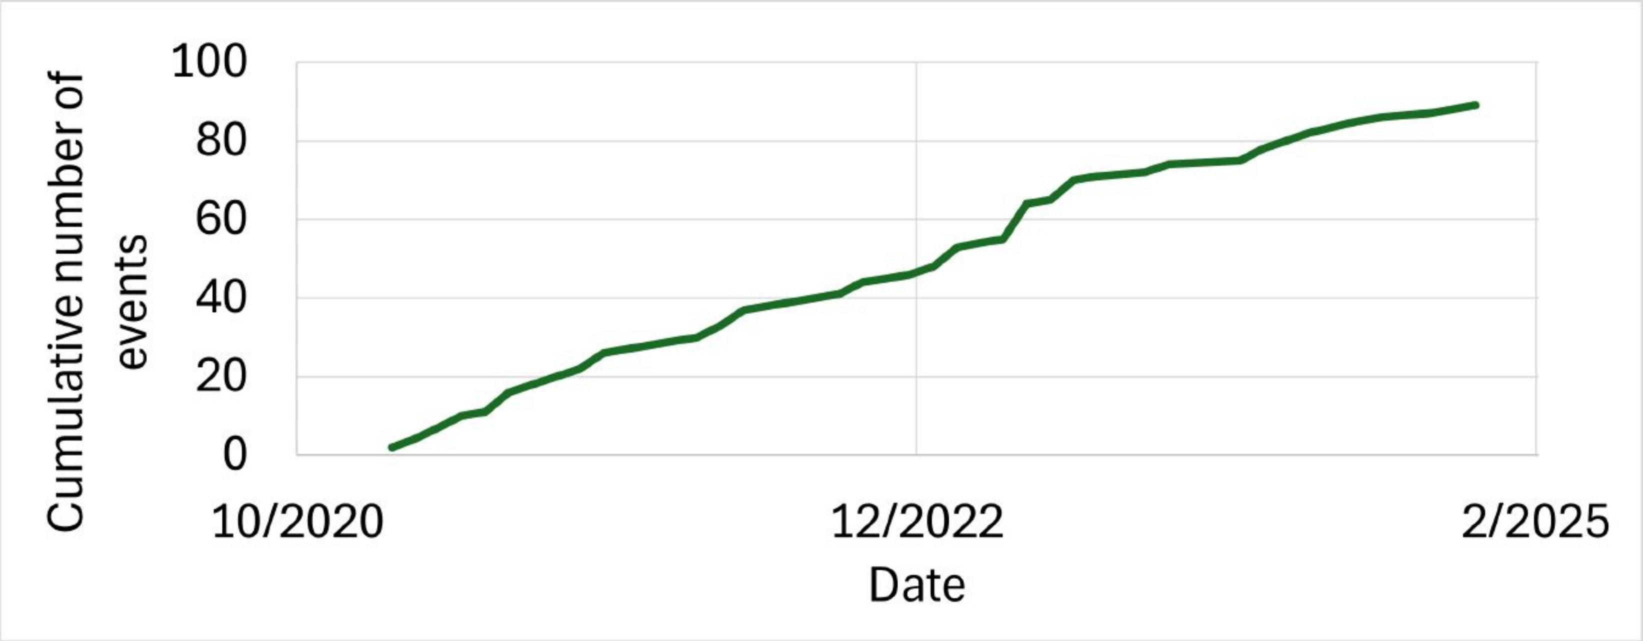

Steady grown in publications from Ookami users

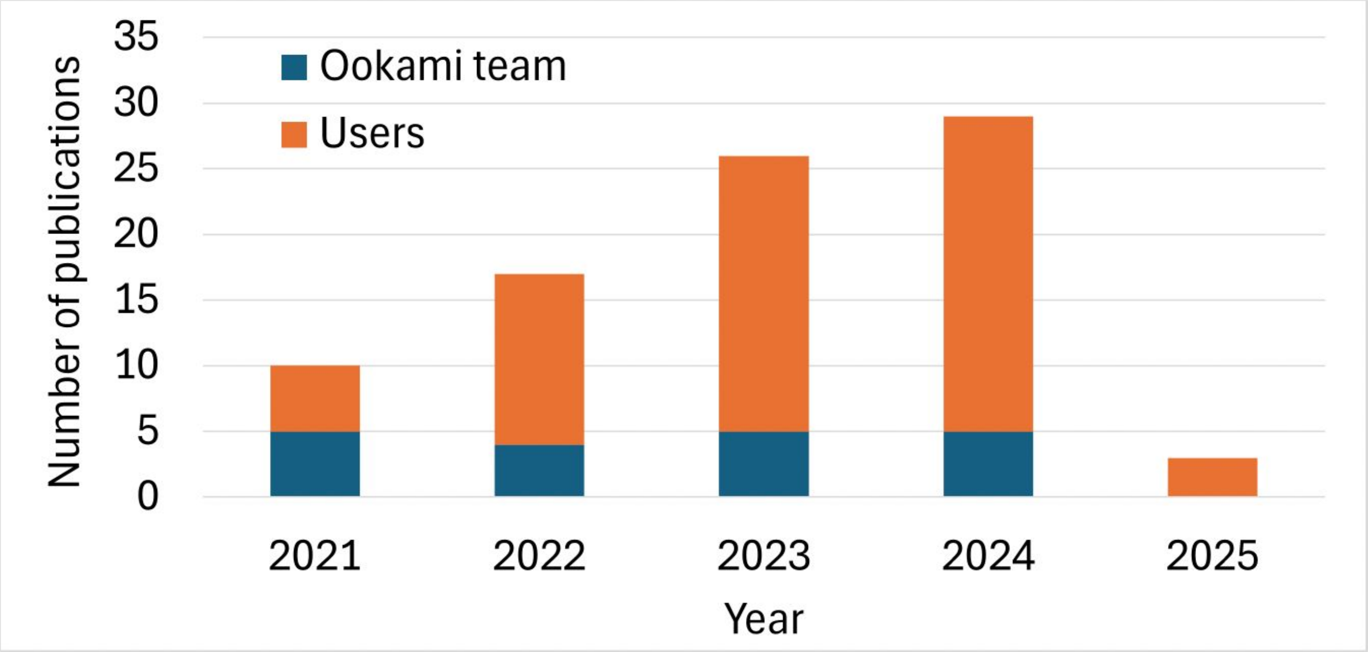

Compute node power draw

- Little difference in power draw between idle and load on A64FX CPU

- Low power spikes correspond to maintenance or system upgrades

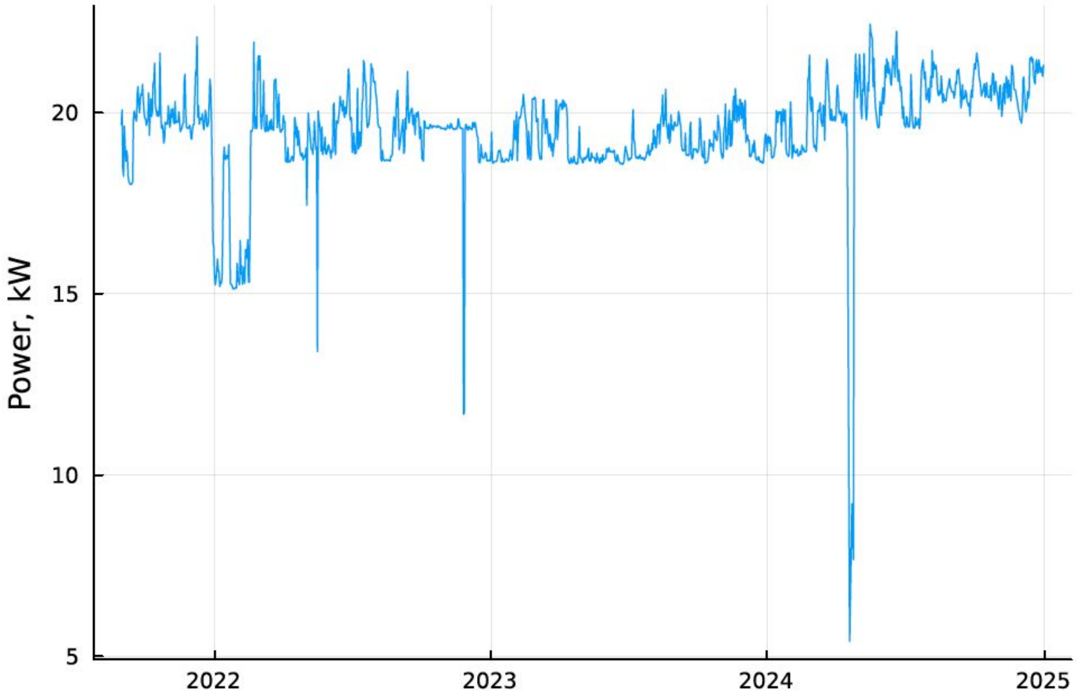

Science application performance

System	Cores	Simulation Speed, ns/day		Simulation Speed per Core, ns/day/core	Power, W	Energy Efficiency, ns/kWh	
CPU Only Calculation							
ARM Futjitsu A64FX, SVE 512bit (SBU-Ookami, Fujitsu)	48	22.8	± 0.3	0.48	105 ± 5	9.1	± 0.4
ARM Cavium ThunderX2 (SBU-Ookami)	64	28.8	± 4.2	0.45			
ARM Amazon Graviton 2, Neoverse N1 (AWS)	48	37.8	± 0.1	0.79			
ARM Amazon Graviton 3, Neoverse V1, SVE 256bit (AWS)	64	71.4	± 1.0	1.12			
ARM Ampere Altra, Neoverse N1 (Azure)	64	56.5	± 0.6	0.88			
ARM Ampere One A192-32X, Neoverse N1 (Ampere)	192	172.1	± 2.3	0.90	512 ± 5	14.0	± 0.1
ARM NVIDIA Grace, Neoverse V2, SVE 128bit (SBU)	144	235.2	± 0.4	1.63	709 ± 41	18.9	± 0.8
x86 AMD EPYC 7742 Zen2(Rome), AVX2 (PSC Bridges-2)	128	109.6	± 4.8	0.86			
x86 AMD EPYC 7763 Zen3(Milan), AVX2 (Purdue Anvil)	128	169.9	± 4.4	1.33			
x86 Intel Xeon Plat. 8160, Skylake-X, AVX512 (TACC-Stampede 2)	48	70.4	± 0.8	1.47			
x86 Intel Xeon Plat. 8380, Ice Lake, AVX512 (TACC-Stampede 2)	80	133.3	± 6.0	1.67			
x86 Intel Xeon Gold 6130, Skylake-X, AVX512 (UBHPC)	32	39.3	± 0.9	1.23	367 ± 35	4.5	± 0.5
x86 Intel Xeon Gold 6330, Ice Lake, AVX512 (UBHPC)	56	103.0	± 2.0	1.84	619 ± 17	6.9	± 0.2
x86 Intel Xeon Max 9468, Sapphire Rapids, AVX512 (SBU)	96	193.08	± 2.3	2.01	820 ± 7	9.8	± 0.1
CPU-GPU Calculations							
x86 Intel Xeon Gold 6130, NVIDIA V100x2 (UBHPC)	32	145.1	± 2.8		435 ± 7	13.9	± 0.3
x86 Intel Xeon Gold 6330, NVIDIA A100x2 (UBHPC)	56	236.5	± 10.8		707 ± 9	13.9	± 0.8
AMD Ryzen 9 7950X (16 Cores Used)/NVIDIA RTX 4090	16	284.82					
NVIDIA Grace Hopper Superchip ES	72	429					

Continuous Performance Monitoring - application performance

Continuous Performance Monitoring - system performance

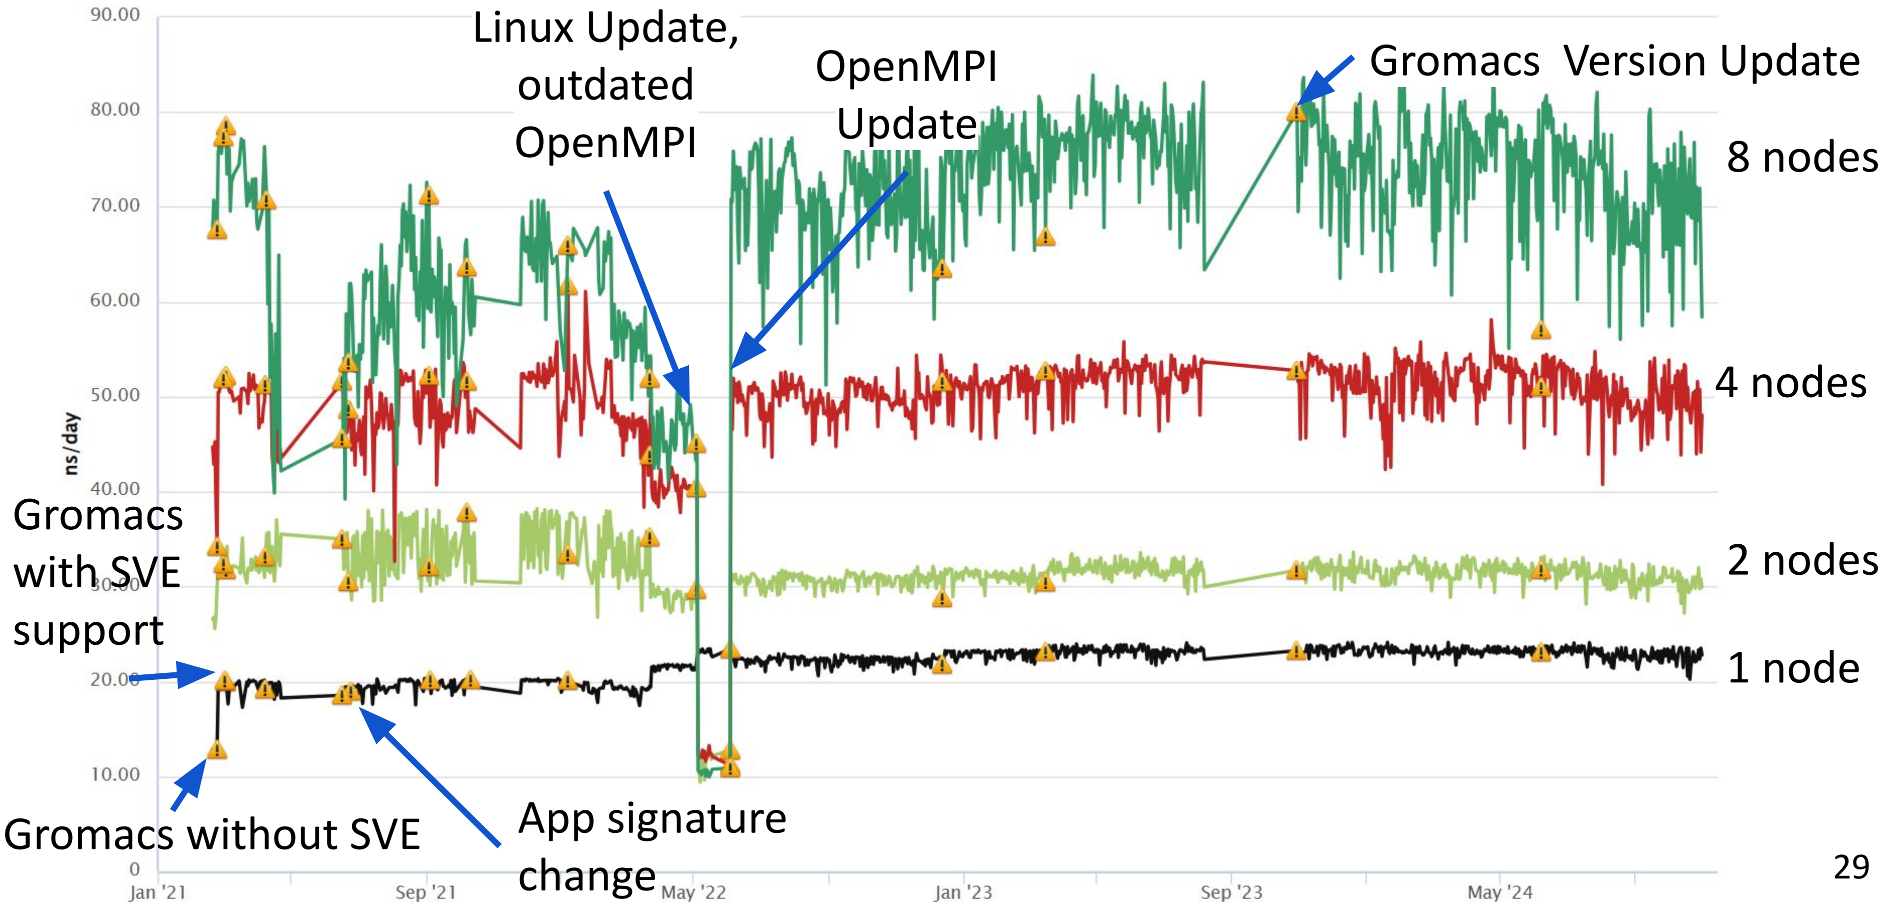

User survey results

- Survey given to users 1-2 months into their project
- ~90% reported satisfaction as “good” or “very good”
- GCC was most popular compiler
  - Nearly 80% usage compared with ~40% for Arm and Fujitsu

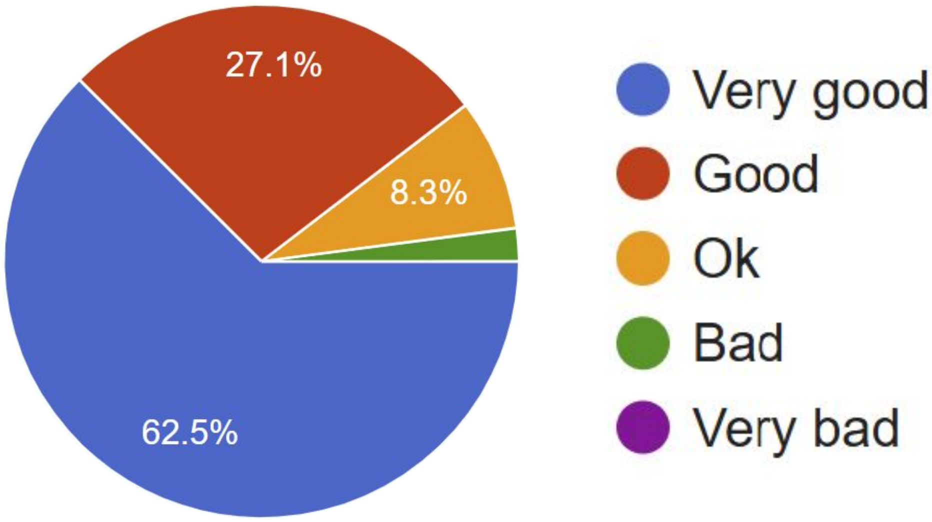

What did we learn?

- For a testbed, 100% utilization is not necessarily expected nor desired
  - Users must be able to experiment
- Consistent, multimodal engagement with users is key
  - Engagement helps to identify where gaps where more support is needed
- As a project shifts focus (e.g., toward production), the user base may shift too
  - constant reinforcement of best practices is necessary

What would we do differently?

- Better tracking of user participation in office hours and webinars
  - Link participation to job efficiency and volume data
  - Assess impact of training efforts
- Track degrees granted as result of Ookami use
  - Assess the testbed’s contribution to research and education
- More systematic logging and categorizing of compiler errors and bugs
- More automation of data collection and analysis

A big thanks to the entire team

University at Buffalo
Matt Jones (Co-PI)
Nikolay Simakov
Joseph White

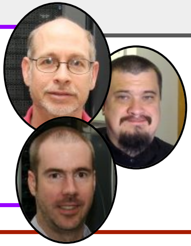

SBU Scientific SW support
Tony Curtis
Eva Siegmann
Dave Carlson

SBU graduate students
Joshua Martin (Astro.)
Rodrigo Ristow Hadlich (Engineering)
Gaurav Verma (CS)
Abdul-Wasay Butt* (HPC support)
Smeet Chheda (CS)
Yuzhang Wang (Phys. & Quant. Biology)
George Liang (HPC support)
Catherine Feldman (Astro.)
Kedarsh Kaushik (Physics)
Youwei Ma (Marine & Atm. Sciences)
Logan Swanson (Linguistics)
Chengpeng Sun (Appl. Math.& Statistics)

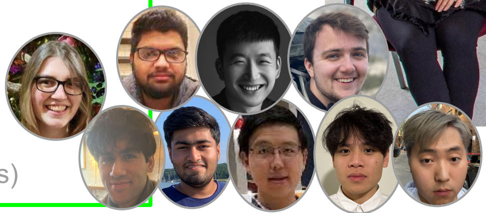

Chief Research Information Officer: David Cyrille

Assistant Director, Advanced Systems & Operations: Firat Coşkun
Admin: Daniel Wood
Admin: John Dey
Admin: Raul Gonzalez

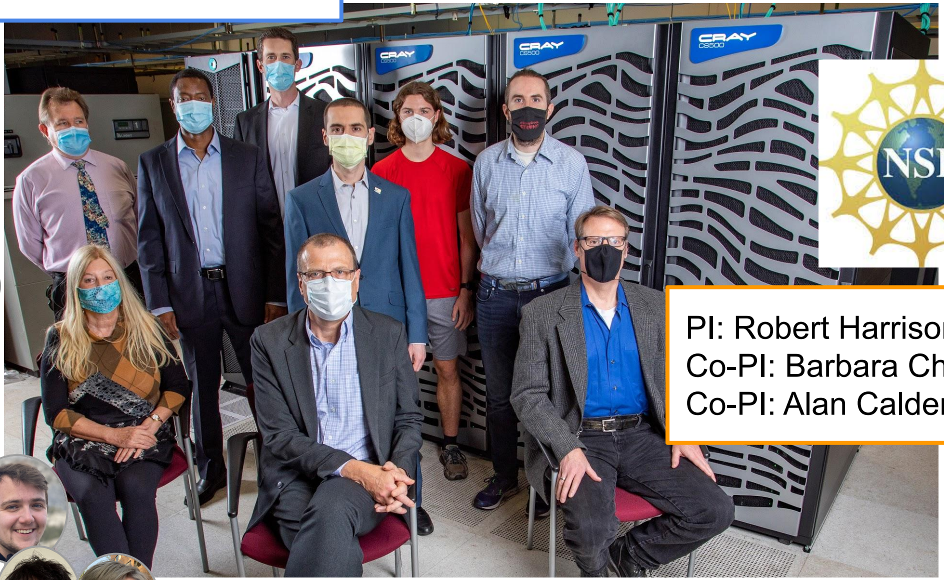

PI: Robert Harrison
Co-PI: Barbara Chapman
Co-PI: Alan Calder

Questions?

www.stonybrook.edu/ookami