



OOKAMI

XSEDE Service Provider Call 2/17/2022

Robert Harrison, Eva Siegmann

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
- Available for researchers worldwide
(excluding ITAR prohibited countries & restricted parties on the EAR entity list)
- Usage is free for non-commercial and limited commercial purposes




Stony Brook
University



University at Buffalo
The State University of New York

Fugaku #1

Fastest computer in the world



First machine to be fastest in
all 5 major benchmarks:

- Green-500
- Top-500 – 415 PFLOP/s in double precision – nearly 3x

Summit!

- HPCG
- HPL-AI
- Graph-500



- 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

<https://www.r-ccs.riken.jp/en/fugaku>

Ookami



Node	
Processor	A64FX
#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

What is Ookami




- 176 Fujitsu **A64FX** compute nodes each with 32GB of high-bandwidth memory and a 512 Gbyte SSD
 - Same as in currently fastest machine worldwide, Fugaku
 - First deployment outside Japan
 - HPE/Cray Apollo 80
- Ookami also includes:
 - 1 node with dual socket AMD Rome (128 cores) with 512 Gbyte memory and 2 NVIDIA V100 GPU
 - 1 node Intel Skylake Processors (32 cores) with 192 Gbyte memory
 - 2 nodes with dual socket Thunder X2 (64 cores) each with 256 Gbyte memory
- Delivers ~ 1.5M node hours per year

A64FX NUMA Node Architecture



- Arm V8-64bit
- Supports high calculation performance and low power consumption
- Supports Scalable Vector Extensions (SVE) with 512-bit vector length
- 4 Core Memory Groups (CMGs)
 - 12 cores (13 in the FX1000)
 - 64KB L1\$ per core - 256b cache line
 - 8MB L2\$ shared between all cores - 256b cache line
 - Zero L3\$
- 32 (4x8) GB HBM @ 1 TB/s
- PCIe 3 (+ Tofu-3) network

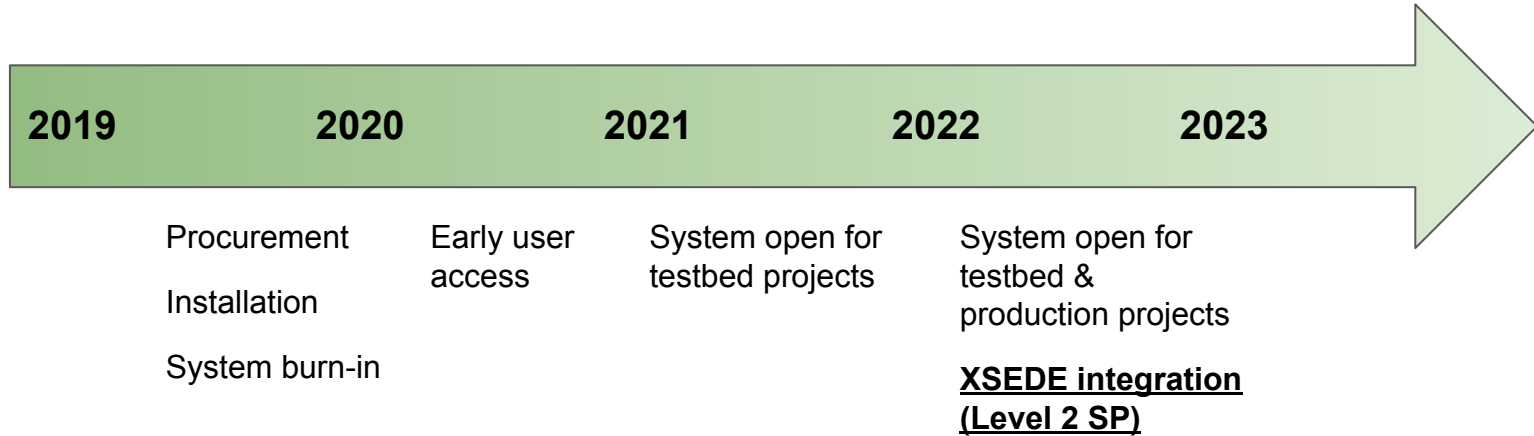


What else



- CentOS 8 operating system
- DUO Authentication
- High-performance Lustre file system (~800TB of storage)
- Slurm workload manager
- Compilers: GNU, Arm, Cray, Fujitsu, Intel, Nvidia
- Continuous growing stack of preinstalled software
 - MPI implementations
 - Math libraries
 - Performance analysis & debugging:
(Arm Forge, Cray, GNU, TAU, ..)

Project Timeline



Allocations



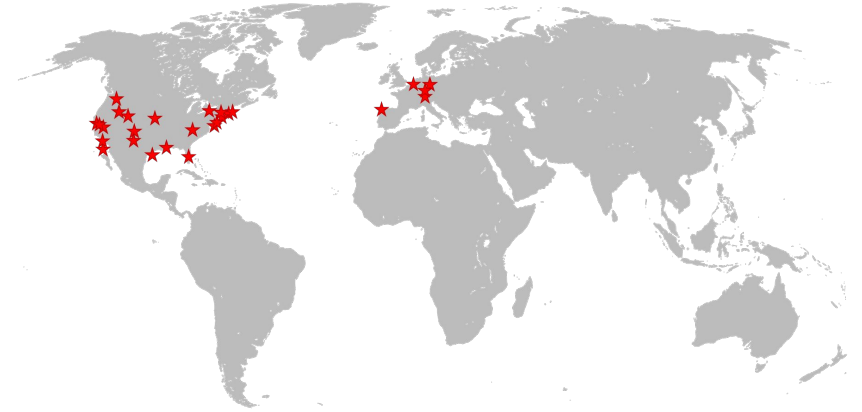
- Currently all allocations through Stony Brook
- Accounts have to be renewed every year
- Ookami is in the process of becoming an XSEDE level 2 service provider
- From October 2022 onwards 90% of allocations will be through XSEDE
- Current testbed projects will still have access though at reduced priority

Projects



- Total: 211 users & 71 projects
- 90% projects from within the US
- 10% from Europe
- 93% from academia
- Complete list of projects:

<https://www.stonybrook.edu/ookami/projects/>



Get in Contact



- <https://www.stonybrook.edu/ookami/>
- Ticketing system: <https://iacs.supportsystem.com/>
- Ookami_computer@stonybrook.edu
- Bi-weekly office hours (Tue 10am – noon, Thu 2 – 4pm EST)
- Slack Channel