

Systems

ISL maintains and provides access to a number of experimental systems:

[Main -> Systems](#)

| System Name | Description | Storage | Networking | Hostname | Status | Service dates | Access |
|--|--|--|---|------------|----------------|---------------|----------------------------|
| Clusters | | | | | | | |
| HAL | Cluster of 16 IBM POWER9 nodes with 4 NVIDIA V100 GPUs per node | 244 TB usable, NVME SSD-based storage by DDN Peak system bandwidth ~100GB/s | 10GbE external Dedicated 10GbE for data transfer dual-channel EDR IB internal | hal | in service | 2019 | open |
| Nano | Cluster of 8 SuperMicro servers with Intel Haswell /Broadwell CPUs and NVIDIA Tesla P100/V100 GPUs | Local 256GB SSD per node NFS-mounted 30TB /home (2x 6-drive RAID z2 with 4TB drives) GlusterFS w/ 2-node fault tolerance 45TB usable | 10GbE external 1GbE + IB QDR internal | nano | in service | 2016 | open |
| Kingfisher | Cluster of 5 Dell R730 servers with Intel Broadwell CPUs and NVIDIA Tesla K80 GPUs (Student Cluster Competition 2016-present) (Jump) | Local 400GB SSD per node 5x 800GB SSD in RAID 5 | 10GbE + IB EDR internal | kingfisher | in service | 2016 | restricted |
|  | | | | | | | |
| KNC | Xeon Phi 5-node co-processor cluster | | | knc | not in service | | |
| KNL | Xeon Phi 4-node processor cluster (access via KNC) | | | knl* | not in service | | |
| Clouds | | | | | | | |
| VLAD OpenStack cluster | OpenStack Cloud deployment | | | | in service | | |
| Servers | | | | | | | |
| DGX A100 | NVIDIA DGX A100 | 14TB RAID0 | 1GbE | hal-dgx | in service | 2021- | open |
| Ac33 | TYAN server with Intel Westmere-EP CPUs and NVIDIA Tesla C2050 GPUs | Local 300GB + 1TB HDD | | ac33 | not in service | 2014 | open |
| Hybrid | SuperMicro server with Intel Haswell-EP CPUs and NVIDIA Tesla P100/V100 GPUs | Local 75GB SSD + one (3x 4TB) RAID 0 arrays | 10GbE | hybrid | in service | 2016 | restricted |

Contact us

Request access to ISL resources: see individual systems

Contact ISL staff: [Email Address](#)

Visit: [NCSA](#), room 3050E

| | | | | | | | |
|----------------------------------|---|---|--------------------------|--------------|----------------|---------|-------------------------------|
| Vlad-hm | High memory server with 3TB of ECC RAM | Local 3TB SSD + two (6x 1TB) RAID 0 arrays CephFS 257TB (reported) | 10GbE | vlad-hm | in service | | open |
| | AMD ARM server | | | | not in service | | |
| Power9 AC922 server | IBM POWER9 4x NVIDIA Tesla V100 16GB | Local 2x 3.84TB SSD Local 1.4 TB NVME | 1GbE | p9 | in service | | restricted to XPACC project |
| NVIDIA ARM HPC DevKit | 1x Ampere Altra Q80-30 2x NVIDIA A100 GPU | | | | | | restricted to NCSA users only |
| FPGA dev | | | | | | | |
| | Altera DE5 FPGA board | | | | | | |
| pynq | XILINX PYNQ-Z1 | | | pynq | offline | 05/2017 | restricted |
| XilinxDevKit | Xilinx Kintex UltraScale FPGA KCU1500 Acceleration Development Kit | | | iridium | in service | | restricted |
| CAPI2.0 platform | Nallatech 250S+in Power 9 server (hal000) 4x NVIDIA Tesla V100 16GB | Local 2x 3.84TB SSD | 10GbE | hal000 | in service | 2018 | restricted to NSF MRI project |
| OpenCAPI platform (IBM P9 IC922) | | | | | | | |
| Xilinx u250 | x86 server with Alveo U250 Data Center Accelerator Card and Vitis development environment | | | hal-fpga-x86 | in service | 2020 | restricted to NSF MRI project |
| Misc | | | | | | | |
| | Ceph storage | | | | offline | | |
| | Raspberry Pi boards | | | | offline | | |
| Jetson | NVIDIA Jetson TK1 cluster | | | | offline | | restricted |
| Jetson TX2 | NVIDIA Jetson TX2 cluster | 1TB NFS-mounted /home | 1GbE internal & external | jetson-tx2-* | offline | | open |
| | | | | | | | |