

Nano cluster

System Description

[Main -> Systems](#)
-> Nano

<p>Host name: nano.ncsa.illinois.edu</p> <p>Hardware</p> <ul style="list-style-type: none"> 8x SuperMicro SYS-4028GR-TR <ul style="list-style-type: none"> X10DRG-O+-CPU motherboard 128 GB DDR4 (8x 16 GB Micron 2133 MHz 36ASF2G72PZ-2G1A2) 8 PCI-E 3.0 ports, switched Mellanox MT27500 Family [ConnectX-3] QDR IB 1x 256 GB Samsung SSD 850 NFS-mounted 30TB /home (2x 6-drive RAID z2 with 4TB drives) GlusterFS w/ 2-node fault tolerance - 45TB usable <p>Software</p> <ul style="list-style-type: none"> CentOS 7 CUDA 9.2/10.0 PGI 16.10 Intel ICC 16 gcc 4.8 gcc 5.3 via 'scl enable devtoolset-4 bash' 	<p>To request access please fill out this form. (Use the link on the confirmation page to sign up for a new account. The same link is also included in the confirmation email.)</p> <p>Instructions for running Jupyter Notebooks on compute nodes</p> <p>Usage notes:</p> <ul style="list-style-type: none"> nano (141.142.204.5) is the head node of the cluster, it should not be used for any computations! to connect to the cluster, ssh <code>username@nano.ncsa.illinois.edu</code> to get access to a particular node for interactive use, use <code>qsub</code>, e.g., <ul style="list-style-type: none"> to get one GPU and one CPU core on node 7 for 1 hour for interactive use: <pre>qsub -I -l nodes=nano7:ppn=1:gpus=1,walltime=3600</pre> to get entire node 1 for 1 hour for exclusive interactive use: <pre>qsub -I -l nodes=nano1:ppn=12,walltime=3600</pre> better yet, do not allocate nodes for interactive use, instead just submit batch jobs, see for example <i>Job Scripts</i> section at https://kb.iu.edu/d/avmy for details. This is a much better way to share computing resources. interactive jobs are limited to 12 hours maximum walltime per job. batch jobs are limited to 96 hours submit request to staff for longer batch jobs (up to 240 hours) to see what's running on the cluster, just run <code>qstat</code> this is a shared resource, please keep in mind that other users are using it as well; do not take over the system beyond what you really need. home directory is cross-mounted and accessible from all nodes Current System Status: https://nano.ncsa.illinois.edu:3000d/3QVrDIFmz/nano-status
--	--

Contact us

Request access to ISL resources: Application

Contact ISL staff: [Email Address](#)

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

DL frameworks

- TensorFlow 1.10

Node configuration (see login message for the exact configuration):

nano1	nano2	nano3	nano4
<ul style="list-style-type: none"> 2x Intel Xeon CPU E 5-2620 v3 @ 2.40 GHz 2x NVIDIA V100 GPUs <ul style="list-style-type: none"> 5120 cores 16 GB HBM2 CUDA 10.1 	<ul style="list-style-type: none"> 2x Intel Xeon CPU E 5-2680 v4 @ 2.40 GHz 2x NVIDIA V100 GPUs <ul style="list-style-type: none"> 5120 cores 16 GB HBM2 CUDA 11.6 	<ul style="list-style-type: none"> 2x Intel Xeon CPU E 5-2680 v4 @ 2.40 GHz 2x NVIDIA V100 GPUs <ul style="list-style-type: none"> 5120 cores 16 GB HBM2 CUDA 10.1 	<ul style="list-style-type: none"> 2x Intel Xeon CPU E 5-2620 v3 @ 2.40 GHz 2x NVIDIA V100 GPUs <ul style="list-style-type: none"> 5120 cores 16 GB HBM2 CUDA 10.1
nano5	nano6	nano7	nano8

<ul style="list-style-type: none">• 2x Intel Xeon CPU E 5-2620 v3 @ 2.40 GHz• 2x NVIDIA P100 GPUs<ul style="list-style-type: none">• 3584 cores• 16 GB HBM2• CUDA 11.6• UNSCHEDULABLE - reserved for project	<ul style="list-style-type: none">• 2x Intel Xeon CPU E 5-2620 v3 @ 2.40 GHz• 2x NVIDIA P100 GPUs<ul style="list-style-type: none">• 3584 cores• 16 GB HBM2• CUDA 11.6	<ul style="list-style-type: none">• 2x Intel Xeon CPU E 5-2620 v3 @ 2.40 GHz• 4x NVIDIA P100 GPUs<ul style="list-style-type: none">• 3584 cores• 16 GB HBM2• CUDA 10.1	<ul style="list-style-type: none">• 2x Intel Xeon CPU E 5-2620 v3 @ 2.40 GHz• 4x NVIDIA V100 GPUs<ul style="list-style-type: none">• 5120 cores• 32 GB HBM2• CUDA 11.6
--	---	---	---