

# Nano cluster

## System Description

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

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

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

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